IMPROVING JET ENGINE AERODYNAMIC DESIGN VIA NOVEL COMPONENT SHAPING AND ANALYSIS

Alistair John

Department of Mechanical Engineering
University of Sheffield

This dissertation is submitted for the degree of
Doctor of Philosophy

April 2018

I would like to dedicate this thesis to my wife, Laura. Thanks for always being there by my side.

Acknowledgements

I would like to thank my supervisor Professor Ning Qin for all of his guidance and advice, not only throughout my PhD but also as my Undergraduate and MEng project tutor. It is thanks to Professor Qin that I have had the opportunity to undertake my PhD and work with Rolls-Royce.

I would also like to sincerely thank Dr Shahrokh Shahpar for his support as my industrial mentor within Rolls-Royce. Dr Shahpar gave me the opportunity to work within his team on a Summer placement in 2013, and following this for my final year project. His industrial experience, advice and contacts within Rolls-Royce have enabled me to carry out state-of-the-art, interesting and industrially relevant research.

I would like to thank the University of Sheffield for funding my PhD via the Faculty of Engineering scholarship.

I would like to thank Rolls-Royce for their provision of computing resources to carry out the work, permission to publish my findings and support towards attending conferences.

I would also like to thank the Royal Aeronautical Society for their support towards my attendence at ASME Turbo Expo 2017, and the Centennial scholarship which has supported my research following the end of my PhD funding.

I would like to thank everyone in Shahrokh's team at Rolls-Royce for their support over the last few years, as well as the experts in the fan and installations departments at Rolls-Royce.

Finally, I would like to thank my peers in Sheffield for their support and friendship and my family, for everything they have done for me.

Abstract

Jet engines are comprised of many complex components, and overall engine efficiency and fuel consumption depend on the individual performance of these. This thesis investigates the use of novel analysis/design approaches that can lead to improvements in engine component aerodynamic performance, and therefore have the potential to reduce jet engine fuel consumption.

This work focuses on design improvements that can be used to benefit future advanced high bypass ratio engines. Several key engine components are studied: fan blades, compressor blades and engine intakes. The various investigations and outcomes are summarised here.

A novel compressor blade shaping approach that uses free-form parameterisation is first investigated. This is shown to offer the potential to increase blade efficiency over alternative parameterisations. Increased design flexibility and the ability to combine 3D blade shaping with aerofoil section modification allows the maximum benefit to be achieved. It is demonstrated that the off-design performance of the optimised blade is satisfactory, meeting the same choke mass flow (once skewed) and achieving reasonable performance at different rotational speeds.

Engine intake performance at high angles of incidence is limited by shock-induced separation that blocks the inlet and can stall the fan. Shock control bumps are demonstrated to be capable of reducing the shock strength and separation on engine inlets at high incidence. It is shown that shock control bumps can completely eliminate shock-induced separation near to the limiting angle of incidence and, crucially, extend the un-separated angle of incidence by at least three degrees.

The use of shock control bumps is also found to be beneficial on transonic fan and compressor blades. It is shown that shock control can delay and reduce separation for a highly-loaded compressor blade, increasing its efficiency without the need for 3D deformations of the blade. The benefit of shock control bumps in reducing shock-induced separation is also demonstrated for a modern, low-speed fan blade, providing an increase in stall margin.

Tip clearance effects for modern fan blades are also investigated. The influence of trenches that become worn in the abradable liner is analysed and a model developed that can predict fan blade efficiency for combinations of non-uniform tip clearances and liner trenches. This model allows designers to predict fan efficiency variation over their lifetime and set cold build clearances for optimal lifetime performance. It is also shown how tip leakage behaviour can influence flow behaviour and separation in other regions of the span, and the influence that the shock position has on the leakage aerodynamics.

Through the use of the approaches demonstrated and knowledge developed in this thesis, engineers have the potential to improve future engine designs in terms of efficiency, fuel consumption and also engine operability.

Publications and Awards

Journal articles

- John, A., Shahpar, S., & Qin, N. (2017). The impact of realistic casing geometries and clearances on fan blade tip aerodynamics. Journal of Turbomachinery, TURBO-17-1166 (accepted, available online).
- John, A., Shahpar, S., & Qin, N. (2017). Novel Compressor Blade Shaping Through a Free-Form Method. Journal of Turbomachinery, 139(8), 081002.

Patent applications

- John, A., Smith, A., Shahpar, S., & Qin, N. Novel engine intake shaping, (pending).

Conference papers

- John, A., Shahpar, S., & Qin, N. (2018, June (submitted)). Using shock control bumps to improve fan/compressor blade performance. In ASME Turbo Expo 2018: Turbomachinery Technical Conference and Exposition, Oslo, Norway. American Society of Mechanical Engineers.
- John, A., Shahpar, S., & Qin, N. (2017, June). Fan blade tip aerodynamics with realistic operational casing geometries and clearances. In ASME Turbo Expo 2017: Turbomachinery Technical Conference and Exposition, Charlotte, NC, USA. American Society of Mechanical Engineers.
- John, A., Shahpar, S., & Qin, N. (2017, April). Understanding fan blade tip aerodynamics. In 12th European Turbomachinery Conference, Stockholm, Sweden.
- John, A., Shahpar, S., & Qin, N. (2016, June). Alleviation of Shock-Wave Effects on a Highly Loaded Axial Compressor Through Novel Blade Shaping. In ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition, Seoul, South Korea. American Society of Mechanical Engineers.

Funding and awards

- University of Sheffield, Faculty of Engineering scholarship to undertake Ph.D
- SET for Britain poster competition invitation to present in Parliament
- Royal Aeronautical Society Centennial Scholarship
- Royal Aeronautical Society Aerospace Speakers Travel Grant
- 1st prize in University doctoral researcher poster competition
- 1st prize in Department of Mechanical Engineering PhD poster competition

Table of contents

List of figures

List of tables

Nomenclature

Roman Symbols

a	Acceleration
B	Set of boundary faces
c	SA closure constant
$\mathscr{C}$	Class function
C	Polynomial coefficient
C	Convection term
d	Distance to nearest wall
D	Destruction term
D	Dissipation term
E	Internal energy
E	Set of connected nodes
f	Blade efficiency in polynomial equation
F	Total flux
F^I	Inviscid flux
F^V	Viscous flux
F	Numerical flux
F	Force
I	Adjoint objective function
k	Thermal conductivity
L	Pseudo-Laplacian
L^{lp}	Linear-preserving form of the Laplacian
$\dot{m}$	Mass flow rate

m	Mass
n	Control surface area vector
n	Surface normal
P	Production term
p	Pressure
Pr	Prandtl number
Q	Vector of conserved variables
$\mathscr{R}_{ij}$	Reynolds stress term
R	Gas constant
R	Navier Stokes residual
$\mathscr{S}$	Shape function
S	Source term
T	Trip term
T	Temperature
$\bar{u}$	Mean x-direction velocity
U	Flow adjoint vector
u	X-direction velocity
u'	Fluctuating x-direction velocity
V	Measure of the volume
v	Y-direction velocity
w	Z-direction velocity
x	Unit vector in x, y and z directions
X	Clearance parameter
y^+	Non-dimensionised near-wall distance

Greek Symbols

β	Adjoint design variable
ε	Relative bump height
η	Isentropic efficiency
γ	Specific heat ratio

κ	Flux limiter
λ	Dynamic viscosity
ν_t	Turbulent eddy viscosity
$\tilde{\nu}$	Spalart variable
ϕ	Normalised distance between bump start/end points
ρ	Density

Subscripts

0	Total
s	Ideal
t	Turbulent
tot	Total

Other Symbols

$\triangle s$	Control surface area

Acronyms / Abbreviations

2D	Two Dimensional
SA	Spalart-Allmaras
3D	Three Dimensional
ADP	Aerodynamic Design Point
BFS	Backward Facing Step
CFL	Courant–Friedrichs–Lewy number
CR	Contraction Ratio
CO2	Carbon Dioxide
CST	Class Shape Transformation
DNS	Direct Numerical Simulation
DoE	Design of Experiments
FFD	Free Form Deformation
FFS	Forward Facing Step
JST	Jameson-Schmidt-Turkel

LE Leading Edge

LES Large Eddy Simulation

MAM Multi-point Approximation Method

MUSCL Monotone Upstream-centered Scheme for Conservation Laws

NOx Nitrogen Oxides

OTL Over Tip Leakage

PADRAM Parametric Design and Rapid Meshing

PR Total Pressure Ratio

RANS Reynolds Averaged Navier Stokes

RSM Response Surface Model

SFC Specific Fuel Consumption

SOFT Smart Optimisation for Turbomachinery

SOPHY SOFT-PADRAM-HYDRA

TE Trailing Edge

TR Total Temperature Ratio

Chapter 1

Background

1.1 Motivation

Civilian aircraft are a key part of the global economy, however their use has a significant impact on the environment. In 2012, commercial aircraft emitted approximately 700 million metric tons of carbon dioxide (CO2) worldwide [63], contributing significantly towards climate change. As a major contributor to global CO2 emissions, aviation would be ranked 7th behind Germany if it were a country. The amount of aviation CO2 emissions has been rising with an almost linear trend for the last 50 years, and is predicted to continue to do so (see Figure 1.1).

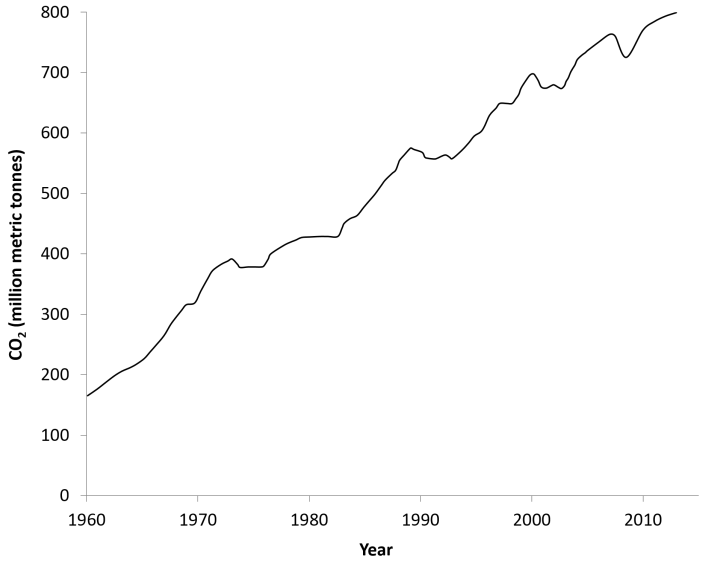

Fig. 1.1 CO2 emissions from aviation (including military and general aviation), 1960 to 2015 [69].

Between 1960 and 2006, aviation fuel use and CO2 emissions have quadrupled, and by 2050 are set to triple again. As can be seen in Figure 1.1, despite many factors, including the oil crises in the 70s, the gulf crisis in the 90s, the World Trade Centre attacks, and the 2008 financial crisis, the overall levels of fuel use and emissions have continued to steadily rise.

Fuel cost is a significant contributor to airline operating costs. Along with the limited supply of fossil fuels and concern over the environmental damage caused by civil aviation, there is a strong desire to reduce the fuel consumption of aircraft jet engines. Aircraft efficiencies have steadily improved over the years, but this has been far outstripped by demand, leading to an overall increase in fuel use and emissions. To try and improve the outlook for the future, the aviation industry and governments around the world have begun to set policies to reduce the impact of aviation on the environment.

Flightpath 2050: Europe's Vision for Aviation sets out the goal of a 75% reduction in individual aircraft CO2 emissions by 2050 (with respect to 2000 levels) [65]. In 2010, the International Civil Aviation Organization (ICAO) Assembly affirmed two main goals for aviation: to achieve an annual 2% average fuel efficiency improvement until 2020 and an aspirational 2% improvement per annum from 2021 to 2050 [46]. To help achieve these goals, much work is required of aerospace jet engine manufacturers to develop the technologies needed to produce more efficient and greener aerospace engines. This is the driver behind this research.

1.2 Aerospace jet engines

Jet engines are a type of internal combustion engine that generate thrust via jet propulsion. In a turbo-jet (the simplest form of jet engine) air enters the engine intake and is compressed by a rotating compressor. It is then mixed with fuel and burnt in the combustion chamber, before passing through a turbine and being expelled through a nozzle to produce thrust. The turbine is used to drive the compressor, and it is the leftover power in the exhaust gas that provides the thrust.

1.2.1 The history of the jet engine

The credit for the invention of the jet engine belongs to Great Britain's Frank Whittle [7], who registered a patent for his turbojet design in 1930. Hans von Ohain [10] of Germany however was actually the designer of the first operational jet engine. He helped to develop the world's first jet plane (the experimental Heinkel He 178) which flew in 1939, two years before the first flight test of Whittle's engine.

Jet engines are capable of producing far more thrust than their piston based predecessors, and it was throughout the second world war that the development of the jet engine first began to intensify, with jet aircraft produced by Britain, Germany and the US. By the end of the war it was clear that the future of aircraft propulsion lay with jet engines.

The early turbojet designs were extremely inefficient however and had very high fuel consumption. Since then, the primary concern of jet engine designers has been to improve fuel consumption whilst also producing large amounts of thrust. In 1948 the "dual spool" concept was introduced [27]. In this orientation the compressor and turbine are both split into two sections, with each part of the compressor attached to its own turbine. This leads to improvements in fuel consumption by allowing the different sections of the compressor and turbine to rotate at more favourable speeds.

In the early 1950s, variable guide vanes were first used [41]. These help to combat compressor stall and improve efficiency by keeping the flow entering the compressor well aligned for the rotor blades across the range of engine speeds.

To produce far greater power and also further increase efficiency, the next major development was the concept of the turbofan, introduced in the late 1950s [4]. In a turbofan, the low pressure turbine drives a fan which produces thrust via the bypass duct, with this bypass flow not passing through the engine core. This invention more than doubled the thrust of earlier engines, improved fuel economy and also helped to reduce noise.

Turbine blade cooling was first introduced around the same time [129]. This is where cold compressor air is used to cool the turbine blades, allowing the turbine entry temperature to increase above the melting point of the blade materials. The result is significantly higher engine thrust and efficiency.

Since this point, civil jet engines have maintained a similar architecture, with performance improvements coming from even higher turbine inlet temperatures, higher bypass ratios and more complex blade and sub-system design developed over many iterations.

1.2.2 Propulsive and thermal efficiencies

The overall efficiency of a jet engine can be split into two components, the thermal efficiency and the propulsive efficiency, namely

$$\eta_{overall} = \eta_{thermal} \times \eta_{propulsive} \tag{1.1}$$

The thermal efficiency relates to how efficiently the energy in the fuel is converted into kinetic energy. This depends on the efficiency of every component in the engine. The propulsive efficiency is how efficiently the engine uses this kinetic energy to produce thrust and drive the aircraft forward. Propulsive efficiency is increased by reducing the over speed of the engine exhaust, through increasing the bypass ratio to move a greater volume of air at a lower velocity. Increasing a turbofan engine bypass ratio increases the diameter of the entire engine and nacelle, and comes at the cost of increased weight.

1.2.3 The future of jet engines

Modern jet engines are highly efficient after undergoing years of design iterations and improvements. As can be seen from Figure 1.2 however, there is still room for improvement. Significant increases in both propulsive and thermal efficiency are possible, with a further 30% reduction in fuel consumption theoretically attainable (though this is limited by the need to limit NOx production). These improvements can be achieved through two methods of development, namely significant change in overall engine architecture and iterations on the current best practice.

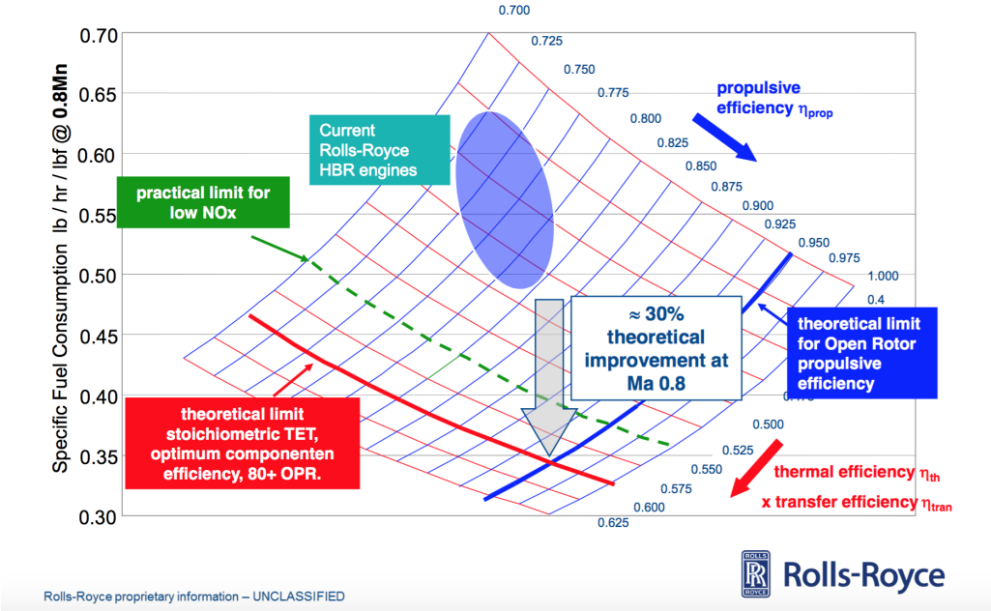

Fig. 1.2 Turbofan thermodynamic cycle efficiencies. Image credit Rolls-Royce Plc. [88].

Future jet engine architectures fall into two main types [63]. These are the advanced/geared turbofan and the open rotor. The principles of each of these are described here.

Open rotor

The open rotor concept provides a significant change in jet engine design and has the potential to result in a step improvement in aero engine performance. In this configuration, instead of a large fan inside a nacelle, two sets of counter-rotating fan blades are mounted outside of the cowling (see Figure 1.3). This design operates under similar principles to a turbofan, by moving a large volume air at lower speed to efficiently produce thrust. The benefits of the open rotor design are that it provides significant savings in engine weight and drag, allowing extremely high bypass ratios to be achieved for the same engine weight. However, there are remaining issues related to noise (the external rotors produce more noise than turbofan fan blades) and safety (the risk from losing a blade is far more severe without an engine casing to contain it). Interest in open rotor designs first began in the 1970s

after the fuel crisis (only to fall as prices did), but has recently been revived and looks promising for the future [97].

Fig. 1.3 Photograph of an open rotor concept. Image credit Aviation Week [126].

Advanced and geared turbofan

Developments in engine design that will be realised in the much nearer future are the advanced/geared turbofan concepts [99].

The term 'advanced turbofan' refers to the trend of moving to turbofans with very high bypass ratios and incorporating various other technologies to improve performance. A geared turbofan is a turbofan architecture where a gearbox is placed between the fan and the shaft that drives it, allowing the fan and low pressure turbine to operate at different (more beneficial) speeds. An additional benefit is that fewer stages are required for the low pressure turbine, reducing engine weight and cost. The geared turbofan also makes it possible to move to higher bypass ratios with a larger fan, without exceeding the limiting tip Mach number and producing unacceptable noise.

The Ultrafan is Rolls-Royce's geared engine architecture of the future [99] (see Figure 1.4). It is a geared turbofan design with an overall pressure ratio of 70:1 and a bypass ratio of 15:1. It has variable pitch carbon fibre / titanium alloy fan blades that save weight and eliminate the need for a thrust reverser, further saving significant weight. The Ultrafan aims to provide 25% lower specific fuel consumption than the Trent 700, or 8-9% lower than the Trent XWB (currently the world's most efficient large aero engine [98]). The focus of the research in this thesis is technologies that are relevant to and have the potential to improve jet engines of the Ultrafan type.

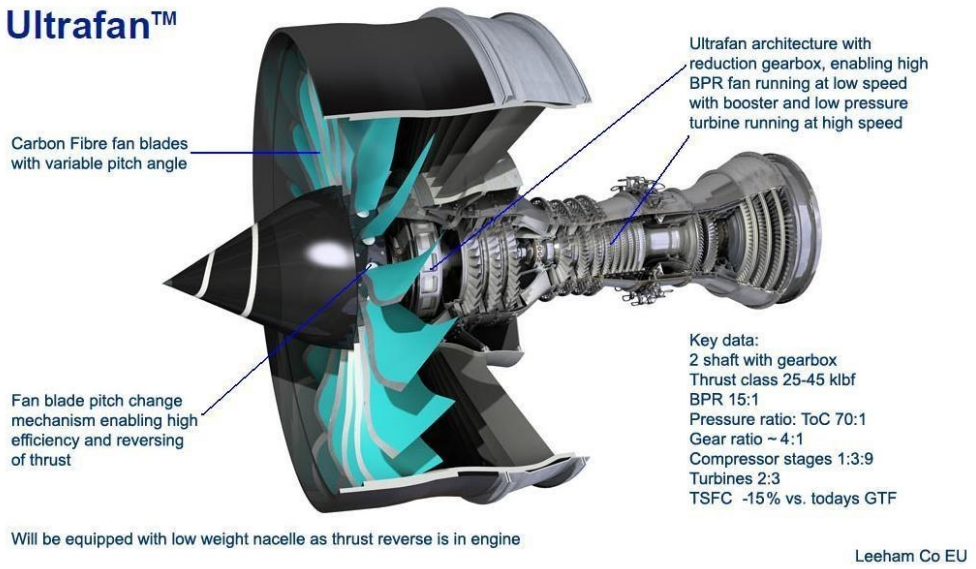

Fig. 1.4 Rolls-Royce Ultrafan. Image credit Rolls-Royce Plc. [99].

Potential areas of development

Alongside the step changes in engine performance from future engine architectures, continuous benefits in engine thermal efficiency are needed to help reduce fuel burn and emissions. These come from improving the individual aerodynamic performance of every component in the engine. The current design of these components is quite mature, and hence novel approaches are needed to find continued benefit.

The International Civil Aviation Organization report on medium and long term goals for aviation fuel burn reduction from technology [33] gives an interesting list of the potential development topics that need to be worked on to provide future reductions in jet engine fuel consumption. These are summarised as follows:

- Advanced turbofan
- Geared turbofan
- Open rotor
- Improved global design through advanced analytical tools
- Optimised components through advanced analytical tools
- Reduced nacelle weights which displace the optimum toward increased bypass-ratio
- Advanced composites / alloys
- Increased loading
- Manufacturing technology
- Combined feature functionality

- Integrated installations
- Zero hub fan
- Improved materials to allow higher temperature
- Improved compressor and turbine with 3D aerodynamics, blowing and aspiration
- Active tip clearance

This research comprises of work relevant to several topics listed above, namely optimising components through advanced analytical tools, improved 3D compressor aerodynamics, highly loaded blades, understanding tip clearance issues and alleviating issues associated with moving towards shorter (lighter) nacelles.

1.3 Aims of this work

This work focuses on the novel aerodynamic design and analysis of gas turbine engine components to improve their performance and reduce engine fuel consumption. This research aims to be relevant to the advanced turbofan and Ultrafan type engine designs of the near future. The three main topics covered in this thesis are - novel fan/compressor blade aerodynamic shaping, tip leakage aerodynamics and engine inlet shock control.

1.3.1 Objectives

A summary of the key objectives of this work are as follows:

- The redesign of fan/compressor blades using novel 3D shaping and optimisation - novel blade shaping with the help of optimisation has the potential to find improvement in blade performance by creating designs that would normally be missed by a designer.
- Engine intake shock control - at high incidence, engine intakes become prone to separation due to the presence of strong shock waves. The novel shaping of engine inlets could improve operability given the trend towards shorter (lighter) intakes (which exacerbate the problem of separation).
- Tip leakage - tip leakage is a significant source of loss for fan blades. With the move to composite fan blades, the tip clearance may need to be increased to prevent rubbing. Understanding fan blade tip clearance behaviour is necessary to be able to set optimum build clearances for best lifecycle cost and performance.

1.4 Organisation of thesis

The following chapters of this thesis are organised as follows:

- Computational methods (Chapter 2)
- Novel compressor blade shaping through a free-form method (Chapter 3)
- Using shock control bumps to improve engine intake operability and performance (Chapter 4)
- The novel shaping of fan/compressor blades using shock control bumps (Chapter 5)
- Fan blade tip aerodynamics with realistic casing geometries (Chapter 6)
- Conclusions (Chapter 7)

Due to the range of topics covered in this thesis, additional motivation, background literature and theory is given individually in the relevant chapters for each topic of work.

Chapter 2

Computational methodology

Computational Fluid Dynamics (CFD) is a critical part of the modern aerodynamic design process alongside physical experimentation. CFD cannot entirely replace experimental work, but it offers many advantages over experimental testing in terms of both cost and speed. The accuracy of modern CFD along with the availability of high performance computing means that it is a widely used, powerful tool in the design of aerodynamic components. CFD software is used in this work in the analysis and re-design of various aero-engine components, as described in Section 1.3.

The fundamentals of the computational methodology used for this research, with details of the geometry and meshing software, flow solver and optimisation tools, are described in this chapter.

2.1 The SOPHY system

In this work, the automated re-design and analysis of various aero-engine components is carried out. To enable this, the set of Rolls-Royce internal tools that make up the SOPHY system [108] is used. This combines the Rolls-Royce geometry parameterisation and meshing software PADRAM, CFD software Hydra [67] and optimisation toolbox SOFT (Smart Optimisation For Turbomachinery) [108]. Using the SOPHY (SOft-Padram-HYdra) system, geometries can be automatically redesigned, re-meshed, simulated and evaluated. A series of Python scripts manage the work-flow and submit jobs to the cluster. The simulations are carried out on the Rolls-Royce CFMS cluster. Figure 2.1 demonstrates the process.

2.2 Geometry and parameterisation

The first step in the design process is to define/modify the geometry. To modify geometry, a parameterisation approach is needed. Parameterisation is the method of describing a geometry by a series of variables (parameters) that are used to define its shape. The selection of the parameterisation approach is critical to the design/optimisation process. The cost of an optimisation/design study is

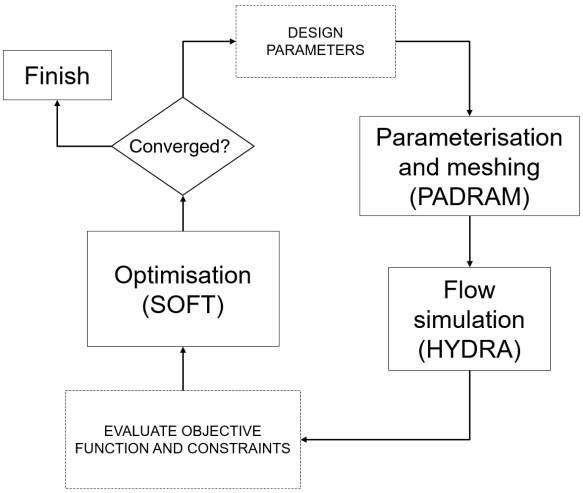

Fig. 2.1 Optimisation Flowchart.

usually proportional to the number of parameters, so allowing too much freedom can be prohibitively expensive. The approach used must also be capable (flexible enough) to generate the most beneficial designs. Sensible limits must also be set to restrict the study to feasible designs.

PADRAM (the Rolls-Royce meshing software [110]) is capable of modifying the geometry during the meshing process. Both the 3D geometry and mesh are constructed out of many 2D sections. In the case of blades, the sections are radially stacked and aligned with flow streamlines. The geometry can therefore be modified by perturbing each 2D section to change the shape of the overall 3D design.

In this work several forms of parameterisation are used. Some use the internal parametric design capability of PADRAM, while for others the geometry that is input to PADRAM is first perturbed using Python scripting. The specific parameterisation approaches for each design study are described in detail in the relevant chapters.

2.3 Meshing

To carry out a CFD analysis, a mesh is needed across which to discretise the governing equations, allowing them to be solved iteratively. PADRAM is used to generate the meshes used throughout this work. PADRAM is a tailor-made automatic multi-block meshing and geometry modification system capable of meshing complex gas turbine geometries, including blades, cavities, intakes and exhausts. It is capable of rapidly meshing two-dimensional, quasi-three-dimensional, and three-dimensional single passage and multi-passage turbomachinery blades. PADRAM produces high quality, hex based structured meshes. Structured meshes are of higher quality and produce more accurate results than unstructured meshes, so provide a benefit for the CFD solver.

Details and examples of the meshes used for each computational study are given in the relevant chapters.

2.4 Flow solver

The Rolls-Royce CFD software Hydra [67] is used throughout this work to simulate the flow around the various internal and external jet engine components studied. HYDRA is an unstructured solver that uses an edge-based data structure with the flow data stored at the cell vertices. A flux-differencing algorithm based on a Jameson-Schmidt-Turkel (JST) [58] with matrix dissipation scheme is used for the space discretisation. For the steady-state solution, a five-stage Runge-Kutta scheme is used. An element-collapsing multi-grid algorithm accelerates the convergence to steady state. A more detailed description of the governing equations, simulation approach and numerical scheme follows.

2.4.1 Governing equations

The governing equations of all CFD software are the mathematical statements of the three governing physical principles of fluid dynamics. These are as follows:

- The conservation of mass
- Newton's second law of motion (F = ma)
- The conservation of energy

The process of simulating fluid flow is to first define the fundamental physical principles (as above), apply these principles to a suitable model of the flow and finally construct the mathematical equations that embody these physical principles. Constructing a model of the flow can be done in two ways, by using a finite control volume approach (the method used here) or by considering an infinitesimal fluid element. Both approaches can either be fixed in space (conservation - used here) or move with the fluid through the flow (the nonconservation form). The fluid flow equations obtained by applying the fundamental physical principles to a finite control volume are in integral form, and are given here (using the conservation approach). All equations are taken from [77].

The three-dimensional Navier Stokes equations take the form:

$$\frac{\partial Q}{\partial t} + \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} + \frac{\partial F_z}{\partial z} = 0 \tag{2.1}$$

where

$$Q = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho E \end{pmatrix} \tag{2.2}$$

$$F_x = \begin{pmatrix} \rho u \\ \rho u^2 + p - \tau_{xx} \\ \rho uv - \tau_{yx} \\ \rho uw - \tau_{zx} \\ (\rho E + p)u - u\tau_{xx} - v\tau_{yx} - w\tau_{zx} + q_x \end{pmatrix} \tag{2.3}$$

$$F_y = \begin{pmatrix} \rho v \\ \rho uv - \tau_{xy} \\ \rho v^2 + p - \tau_{yy} \\ \rho vw - \tau_{zy} \\ (\rho E + p)v - u\tau_{xy} - v\tau_{yy} - w\tau_{zy} + q_y \end{pmatrix} \tag{2.4}$$

$$F_z = \begin{pmatrix} \rho w \\ \rho uw - \tau_{xz} \\ \rho vw - \tau_{yz} \\ \rho w^2 + p - \tau_{zz} \\ (\rho E + p)w - u\tau_{xz} - v\tau_{yz} - w\tau_{zz} + q_z \end{pmatrix} \tag{2.5}$$

Q are the quantities that vary over time in the control volume and F are the fluxes that enter or leave the control volume. ρ, u, v, w, p and E denote the density, the three Cartesian velocity components, the static pressure and the total internal energy, respectively. The equation of state for an ideal gas is required to complete the equations:

$$p = \rho \mathsf{R} T = (\gamma - 1)\rho \left(E - \frac{1}{2}(u^2 + v^2 + w^2) \right) \tag{2.6}$$

where R, T and γ are the gas constant, temperature and specific heat ratio, respectively. The equations defining the heat fluxes are given by:

$$q_x = -k\frac{\partial T}{\partial x}, \quad q_y = -k\frac{\partial T}{\partial y}, \quad q_z = -k\frac{\partial T}{\partial z},$$

with the coefficient of thermal conductivity $k = \frac{C_p \mu}{Pr}$ and Pr the Prandtl number (Pr = 0.72 for air). The shear stresses are given by:

$$\tau_{xx} = 2\mu\frac{\partial u}{\partial x} + \lambda\left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}\right) \tag{2.7}$$

$$\tau_{yy} = 2\mu\frac{\partial v}{\partial y} + \lambda\left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}\right) \tag{2.8}$$

$$\tau_{zz} = 2\mu\frac{\partial w}{\partial z} + \lambda\left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}\right) \tag{2.9}$$

$$\tau_{xy} = \tau_{yx} = \mu\left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}\right) \tag{2.10}$$

$$\tau_{xz} = \tau_{zx} = \mu\left(\frac{\partial u}{\partial z} + \frac{\partial w}{\partial x}\right) \tag{2.11}$$

$$\tau_{yz} = \tau_{zy} = \mu\left(\frac{\partial v}{\partial z} + \frac{\partial w}{\partial y}\right) \tag{2.12}$$

where the molecular viscosity is modelled by Sutherland's law,

$$\mu = \frac{1.461x10^{-6}T^{3/2}}{T + 110.3}, \tag{2.13}$$

and the bulk viscosity λ is defined via the Stokes' hypothesis [119]:

$$\lambda = -\frac{2}{3}\mu \tag{2.14}$$

2.4.2 Reynolds Averaged Navier Stokes (RANS)

Discretising the Navier-Stokes equations described above and directly solving them iteritevly in time and space is known as Direct Numerical Simulation (DNS) [85]. For high Reynolds number applications this is prohibitively expensive [85] and hence an alternative approach is needed here. For design studies, such as these investigations, Large Eddy Simulation (LES - where the turbulent vortices are only resolved above certain turbulence length scales) is also too computationally expensive to be of use. Hence, the time-averaged Reynolds Averaged Navier Stokes (RANS) approach is used. This far simplifies the simulation requirements while maintaining an accurate prediction of the mean flow.

The RANS equations are formed through the decomposition of the full Navier Stokes equations into the mean and fluctuating velocities through Reynolds decomposition [94]. Reynolds decomposition refers to separation of the flow variables (e.g. instantaneous velocity u) into the mean (time-averaged) component $\bar{u}$ and the fluctuating component u', i.e.

$$u(\mathbf{x},t) = \bar{u}(\mathbf{x}) + u'(\mathbf{x},t) \tag{2.15}$$

where $\mathbf{x} = (x,y,z)$.

Time averaging is then carried out on the resulting equations to form the RANS equations. Due to the non-linearity of the Navier Stokes equations, some fluctuating velocity terms remain. These extra terms that appear are known as Reynolds stresses [15], which have the form:

$$\mathscr{R}_{ij} = \rho\overline{u'_i u'_j} \tag{2.16}$$

Reynolds stresses pose a problem, (known as the closure problem), as there are not enough equations to solve for all of the variables. The effect of these extra terms on the flow is similar to that of a stress term. The Reynolds Stresses can therefore be represented by shear stresses (via Newton's law of viscosity) and then the Boussinesq hypothesis [15] used to represent this by the turbulent (eddy) viscosity. The averaged equations then take the same form as the Navier-Stokes equations, if the viscosity and conductivity terms are modified to incorporate molecular and turbulent contributions (see equations (2.17) and (2.18)). The total viscosities then become

$$\mu_{tot} = \mu + \mu_t, \quad \lambda_{tot} = -\frac{2}{3}\mu_{tot}, \tag{2.17}$$

and the total thermal conductivity is given by

$$k_{tot} = \frac{\gamma\mu}{Pr} + \frac{\gamma\mu_t}{Pr_t}, \tag{2.18}$$

where μ_t is the turbulent eddy viscosity and Pr_t is the turbulent Prandtl number (Pr_t=0.9 for air [77]). A turbulence model can be then be used to predict the turbulent viscosity and therefore simulate the turbulent aspects of the flow which cannot be resolved.

2.4.3 Turbulence modeling

The Spalart-Allmaras (SA) turbulence model is used throughout this work. The model has been used to good effect and validated for many turbomachinery test cases in the past. It was selected for use here due to the following reasons:

- When comparisons of the simulation results are made to the experiment (whenever data is available) the SA model compares well, and no benefit is seen in using an alternative model.
- Using SA allows adjoint analysis to be carried out (in the Hydra system, adjoint is currently only available using SA), allowing surface sensitivities to be analysed.
- Using the SA model generally provides the best convergence behaviour in Hydra.
- Unlike other models such as k-epsilon or k-omega, SA is a one-equation model, which is faster computationally.
- Using SA means this work is consistent with the current best practice of the fans, compressors and installation aerodynamic departments in Rolls-Royce, ensuring these results are comparable and relevant to industry.

Validation of the simulation results versus experiment, and in some cases other turbulence models, is given in the relevant chapters where data is available.

Model details

The Spalart-Allmaras turbulence model [117] is a one-equation model which solves a transport equation for a viscosity like variable, $\tilde{\nu}$,

$$\frac{\partial \tilde{\nu}}{\partial t} + u\frac{\partial \tilde{\nu}}{\partial x} + v\frac{\partial \tilde{\nu}}{\partial y} + w\frac{\partial \tilde{\nu}}{\partial z} = \frac{1}{\sigma}\left(\nabla.[(\nu+\tilde{\nu})\nabla\tilde{\nu}] + c_{b2}(\nabla\tilde{\nu})^2\right) + S \tag{2.19}$$

where ν is the molecular kinetic viscosity, $\tilde{\nu}$ is the Spalart variable and S is the source term.

The source terms are:

$$S = c_{b1}\tilde{S}\tilde{\nu} - \left(c_{w1}f_w - \frac{c_{b1}}{\kappa^2}f_{t2}\right)\left(\frac{\tilde{\nu}}{d}\right)^2 + f_{t1}\triangle u^2, \tag{2.20}$$

which may be divided into production (P), destruction (D) and trip (T) contributions:

$$\mathsf{S} \equiv \mathsf{P}(\tilde{\nu}) - \mathsf{D}(\tilde{\nu}) + \mathsf{T} \tag{2.21}$$

using the definitions

$$\mathsf{P}(\tilde{\nu}) = c_{b1}\tilde{S}\tilde{\nu}, \qquad \mathsf{D}(\tilde{\nu}) = \left(c_{w1}f_w - \frac{c_{b1}}{\kappa^2}f_{t2}\right)\left(\frac{\tilde{\nu}}{d}\right)^2, \qquad \mathsf{T} = f_{t1}\triangle u^2.$$

The trip term, T, provides a means for triggering transition at a specified location. However, in keeping with many implementations of the Spalart-Allmaras model, this is set to zero [67] as transition is not modelled in this work.

The turbulent eddy viscosity is defined by

$$\nu_t = \tilde{\nu} f_{v1}, \tag{2.22}$$

and the auxiliary relations used to construct the production and destruction terms are as follows:

$$f_{v1} = \frac{X^3}{X^3 + c_{v1}^3}, \qquad X = \frac{\tilde{\nu}}{\nu}, \qquad f_{v2} = 1 - \frac{X}{1 + X f_{v1}}, \tag{2.23}$$

$$f_w = g\left(\frac{1+c_{w3}^6}{g^6 + c_{w3}^6}\right)^{\frac{1}{6}}, \qquad g = r + c_{w2}(r^6 - r), \qquad r = \frac{\tilde{\nu}}{\tilde{S}\kappa^2 d^2}, \tag{2.24}$$

$$\tilde{S} = S + \frac{\tilde{\nu}}{\kappa^2 d^2} f_{v2}, \qquad S = \sqrt{\left(\frac{\partial w}{\partial y} - \frac{\partial v}{\partial z}\right)^2 + \left(\frac{\partial u}{\partial z} - \frac{\partial w}{\partial x}\right)^2 + \left(\frac{\partial u}{\partial y} - \frac{\partial v}{\partial z}\right)^2}. \tag{2.25}$$

Here, d is the distance to the nearest wall and the closure constants are given by:

$$c_{b1} = 0.1355, \quad c_{b2} = 0.622, \quad \sigma = \frac{2}{3}, \quad c_{v1} = 7.1, \tag{2.26}$$

$$c_{w1} = \frac{c_{b1}}{\kappa^2} + \frac{1 + c_{b2}}{\sigma}, \tag{2.27}$$

$$c_{w2} = 0.3, \quad c_{w3} = 2, \quad \kappa = 0.41. \tag{2.28}$$

In Hydra, the Spalart-Allmaras turbulence model can be used with or without wall functions. Wall functions can be used to model near-wall flow, if the grid resolution near the wall is not high enough. In this work, wall functions are not used and the near wall flow is resolved, unless the y^+ is detected to be too high, in which case the wall function is automatically switched on. All of the meshes employed in this work are therefore created with y^+ values of the order of one, in order to ensure the near wall flow can be accurately resolved.

2.4.4 Discretisation approach

The discretisation approach that allows the RANS equations to be solved is presented in this section. All equations are taken from [77]. The approach used in Hydra is the finite volume method.

The 3D compressible RANS equations are here expressed in a more concise way as follows:

$$\frac{\partial Q}{\partial t} + \nabla . \mathsf{F}(Q, \nabla Q) = S(Q, \nabla Q) \tag{2.29}$$

where $Q(\mathbf{x})$ is the vector of conserved variables, $(\rho, \rho u, \rho v, \rho w, \rho E)^T$. $\mathsf{F}(Q, \nabla Q)$ is the total flux and the source term S is of the form $(0,0,0,0,0)^T$, as it is zero for these cases.

Using the finite volume approach, equation (2.29) is integrated over the control volume. The divergence theorem (sometimes known as Gauss' theorem) [5] is then applied. The divergence theorem states that the outward flux of a vector field through a closed surface is equal to the volume integral of the divergence over the region inside the surface. Essentially, the overall change within the volume is equal to the summation of what passes through the volume's surfaces. The application of this theorem gives the expression:

$$R_j = \frac{1}{V_j} \left(\oint_{\partial V_j} \mathsf{F}(\mathbf{n}, Q, \nabla Q) ds - \int_{V_j} S(Q, \nabla Q) dv \right) = 0, \quad \forall j \tag{2.30}$$

where R_j is the residual (the imbalance of physical conservation within the control volume) and V_j is the measure of the control volume associated with index j. In Hydra, the unknowns are stored at the nodes of the computational mesh (the cell vertex approach), and the control volume is constructed

around these nodes. This control volume is the 'median-dual' [78] which is constructed from the centroids of the cells surrounding the node and the midpoints of the edges (see example in Figure 2.2).

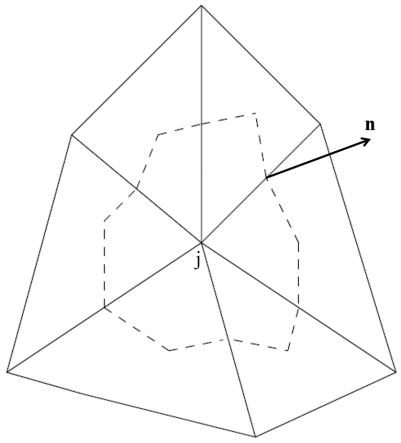

Fig. 2.2 The median dual (dashed line) formed around an internal node j [78].

The flux integration in equation (2.30) is approximated using the dual control volume surface area vectors (**n**) [59] [79]. These are the directional contributions to the overall control volume surface area of each control volume surface. These control surface area vectors are anti-symmetric; i.e. the contribution of the surface appears with opposite signs at opposite faces of the control volume, thus ensuring conservation by construction. The integration of the fluxes can therefore be approximated by the summation [77]:

$$\oint_{\partial V_j} \mathsf{F}(\mathbf{n}, Q, \nabla Q) ds \approx \sum_{i \in \mathsf{E}_j} \mathsf{F}(\mathbf{n}_{ij}, Q, \nabla Q)|_{x=\frac{1}{2}(x_i+x_j)} \triangle s_{ij} \tag{2.31}$$

where E_j is the set of all nodes connected to node j via an edge, $\mathbf{n}_{ij}$ is a unit vector and $\triangle s_{ij}$ an area in 3D (edge in 2D) associated with the edge connecting nodes i and j.

For nodes on a boundary, extra terms are needed. These extra terms from the boundary faces (such as the inlet, exit, far field, walls etc.) are added, so that the approximation of the flux may be written as follows:

$$\oint_{\partial V_j} \mathsf{F}(\mathbf{n},,) ds \approx \sum_{i \in E_j} \mathsf{F}(\mathbf{n}_{ij},,)|_{x=\frac{1}{2}(x_i+x_j)} \triangle s_{ij} + \sum_{k \in B_j} \mathsf{F}(\mathbf{n}_k,,)|_{x_j} \triangle s_k \tag{2.32}$$

where B_j is the set of boundary faces associated with node j, $\mathbf{n}_k$ is the corresponding normal and $\triangle s_k$ the edge area. The edge weight conservation implies that

$$\sum_{i \in E_j} \mathbf{n}_{ij} \triangle s_{ij} + \sum_{k \in B_j} \mathbf{n}_k \triangle s_k = 0 \tag{2.33}$$

Hence, the discrete equivalent to equation (2.30) becomes :

$$R_j = \frac{1}{V_j}\left(\sum_{i\in E_j} F_{ij}\triangle s_{ij} + \sum_{k\in B_j} F_k \triangle s_k - S_j V_j\right) \quad \forall j \tag{2.34}$$

where F_{ij} is the numerical flux in the direction $\mathbf{n}_{ij}$ associated with an antisymmetric edge (i, j), and F_k is the one associated with the boundary face k.

As shown in equation (2.34), to calculate the residual (that is needed to be able to update the nodal variable values at the next iteration) the fluxes at each edge are required. In Hydra these are split as follows:

$$\mathsf{F}(\mathbf{n}, Q, \nabla Q) = \mathsf{F}^I(\mathbf{n}, Q) + \mathsf{F}^V(\mathbf{n}, Q, \nabla Q) \tag{2.35}$$

For any unit normal $\mathbf{n}$, the discrete approximation F of F will have an inviscid and viscous part. As can be seen in equation (2.35), the inviscid part does not depend on the derivatives of Q, as there are no shear stress terms. Each of these are described separately in the next two sections.

Inviscid fluxes

The scheme described here is motivated by the well-known monotone upstream-centered scheme for conservation laws (MUSCL) [53] approach, in which a functional representation of Q is used within each control volume to arrive at a Riemann problem at the interface between control volumes. To avoid the cost of reconstructing Q^+ and Q^- at the edges of the control volumes however, an approximation is used where just the fluxes at the nodes are used to calculate the flux at the edge between the control volumes. Hence, the inviscid discrete flux at the edge between nodes i and j is given by [77]:

$$F^I_{ij} = \frac{1}{2}\left(\mathsf{F}^I_{ij}(Q_i) + \mathsf{F}^I_{ij}(Q_j) - |A_{ij}|(Q^+ - Q^-)\right) \tag{2.36}$$

where $\mathsf{F}_{ij} = \mathsf{F}(\mathbf{n}_{ij},,)$ and $A_{ij} = \partial \mathsf{F}^I / \partial Q$.

The flux terms are calculated from the nodal variables and by solving the discretised form of the Navier Stokes equations. The expression (2.36) can be interpreted as a central difference method with numerical smoothing. The term $|A_{ij}|(Q^+ - Q^-)$ is the numerical smoothing. This is used to smooth changes in the variables between cells and without it the scheme would be unstable [54]. The approach is based on that by Lohner [71] and constructs the values for Q_{j+} and Q_{i-} at equidistance spaced points along a straight line with Q_j and Q_i. A four-point scheme can then be used to evaluate the flux and the inviscid flux F^I_{ij} can therefore be written using the pseudo-Laplacian as follows:

$$F^I_{ij} = \frac{1}{2}\left(\mathsf{F}^I_{ij}(Q_i) + \mathsf{F}^I_{ij}(Q_j) - \frac{1}{2}(1-\kappa)|A_{ij}|(L_j(Q) - L_i Q)\right) \tag{2.37}$$

where κ acts as a flux limiter and L is an undivided pseudo-Laplacian with a negative unit central coefficient. A modification to the pseudo-Laplacian is required however as L will not preserve a linear solution. Therefore a 'linear preserving' form of the Laplacian is instead used (L_j^{lp}), such that the smoothing part of equation (2.36) is switched off if the flow behaviour is linear [78].

An additional treatment is required to aid the capturing of shocks. The artificial dissipation term is therefore adapted using a nonlinear blend of second and fourth-order differences, and a limiter is introduced so that the smoothing reverts to first-order characteristic-upwinding at shocks. The formulation is based on that used in [56] [31]:

$$F_{ij}^I = \frac{1}{2}\left(\mathsf{F}_{ij}^I(Q_i) + \mathsf{F}_{ij}^I(Q_j) - |A_{ij}|\left(-\frac{1}{3}(1-\psi)(\hat{L}_i^{lp}(Q) - \hat{L}_j^{lp}(Q)) + \psi(Q_i - Q_j)\right)\right) \tag{2.38}$$

where

$$\psi = min\left(\varepsilon^2 \left|\frac{p_j - p_i}{p_i + p_j}\right|^2, 1\right) \tag{2.39}$$

Here ε is 8 [77] and p is the static pressure at the corresponding node.

Viscous fluxes

The viscous flux is approximated half-way along each edge. The integration rule around each volume, i.e. equation (2.30) can then be used, giving a consistent finite volume treatment of the inviscid and viscous terms. Additionally to when calculating the inviscid fluxes, calculating the viscous fluxes also requires the approximation of ∇Q at the midpoint of each edge. The use of a smoothing term is not required as they are themselves dissipative. The gradients of the flow variables can be approximated at the nodes using the existing edge-weights. An approximation at the midpoint of the edge can then be obtained by a straightforward average, namely,

$$\overline{\nabla Q_{ij}} = \frac{1}{2}(\nabla Q_i + \nabla Q_j) \tag{2.40}$$

However, as this is the average of two central differences, it does not dampen high frequency modes. The inviscid flux includes numerical dissipation terms that will dampen these modes, but this is insufficient when the viscous terms dominate, such as in the boundary layer. To remedy this, the component of ∇Q in the direction along the edge is replaced by a simple difference along the edge as follows:

$$\nabla Q_{ij} = \overline{\nabla Q_{ij}} - \left(\nabla Q_{ij}.\delta\mathbf{s}_{ij} - \frac{(Q_i - Q_j)}{|\mathbf{x}_i - \mathbf{x}_j|}\right)\delta\mathbf{s}_{ij} \tag{2.41}$$

where

$$\delta \mathbf{s}_{ij} = \frac{\mathbf{x}_i - \mathbf{x}_j}{|\mathbf{x}_i - \mathbf{x}_j|} \tag{2.42}$$

In the boundary layer, it is the simple differences along the shortest edges which contribute to the dominant viscous flux terms, and so this formulation dampens the high-frequency error modes.

2.4.5 Time stepping

The pseudo time-stepping is carried out using the 5-stage Runge-Kutta scheme by Martinelli [73]. This can be expressed as follows:

$$Q_j^{(0)} = Q_j^n \tag{2.43}$$

$$Q_j^k = Q_j^n - \alpha_k \triangle t_j R_j^{(k-1)}, \qquad k = 1,2,3,4,5 \tag{2.44}$$

$$Q_j^{n+1} = Q_j^{(5)} \tag{2.45}$$

where

$$R_j^{(k-1)} = C_j(Q^{(k-1)}) - B_j^{(k-1)} \tag{2.46}$$

$$B_k^{(k-1)} = \beta_k D_j(Q^{(k-1)}) + (1-\beta_k) B_j^{(k-2)} \tag{2.47}$$

$C_j(Q^{(k-1)})$ is the convective (inviscid) contribution to R_j arising from the Galerkin approximation of the inviscid terms in equation (2.29) and $D_j(Q^{(k-1)})$ are the remaining (viscous) parts due to the source terms and dissipation, both numerical and physical. The coefficients α_k and β_k are given by

$$\alpha_1 = \frac{1}{4} \quad \alpha_2 = \frac{1}{6} \quad \alpha_3 = \frac{3}{8} \quad \alpha_4 = \frac{1}{2} \quad \alpha_5 = 1 \tag{2.48}$$

$$\beta_1 = 1 \quad \beta_2 = 0 \quad \beta_3 = \frac{14}{25} \quad \beta_4 = 0 \quad \beta_5 = \frac{11}{25} \tag{2.49}$$

This Runge-Kutta scheme is designed to be stable with low cost, as it does not calculate $D_j(Q^{(2)})$ and $D_j(Q^{(4)})$. A local time-stepping approach is used, and this is based on a combination of the viscous and inviscid time-steps,

$$\frac{1}{\triangle t_j} = \frac{1}{CFL} max\left(\frac{1}{\triangle t_j^I}, \frac{\varepsilon^V}{\triangle t_j^V}\right) \tag{2.50}$$

where CFL is the inviscid Courant–Friedrichs–Lewy (CFL) [29] number and $\varepsilon^V = 0.5$.

The inviscid time-step is based on a spectral radius upper bound on the Jacobians of the discrete inviscid operator, namely,

$$\frac{1}{\triangle t_j^I} = \frac{1}{V_j}\left(\sum_{i\in E_j}\rho(A_{ij}\triangle s_{ij}) + \sum_{i\in B_j}\rho(A_k\triangle s_k)\right) \tag{2.51}$$

here $\rho(A)$ is the spectral radius of the matrix $\partial \mathsf{F}^I/\partial Q$.

The viscous time-step is obtained from the quasi-linear form in terms of primitive variables of the viscous terms and based again on the maximum spectral radius of the Jacobian matrices [30]. If $|\mathbf{x}_i - \mathbf{x}_j|$ is chosen as representative of the geometric quantities for each edge (i, j), the viscous time-step is defined by

$$\frac{1}{\triangle t_j^V} = \frac{1}{V_j}\sum_{i\in E_j}\rho(B_{ij})\frac{1}{|\mathbf{x}_i - \mathbf{x}_j|}\triangle s_{ij} \tag{2.52}$$

where $\rho(B)$ denotes the spectral radius of the matrix $\partial \mathsf{F}^V/\partial Q$.

In Hydra, a block Jacobi preconditioner is used to improve the convergence rate without affecting the steady-state solution [78]. An element collapsing multi-grid approach, employing a four-step v-cycle, is used to accelerate convergence to steady-state [80]. The solver runs in parallel on both shared and distributed memory machines using domain decomposition.

Validation of Hydra can be found in [127] [109] [61] [23] and a summary of its use in various applications is given by Lapworth [100].

2.4.6 Optimisation

Once a CFD solution has been found and the performance of a particular component has been assessed, the next step is to optimise the geometry.

Optimisation is the process by which a design space (formed by the selected parameterisation method and the upper and lower limits imposed on each parameter) is explored to search for an optimum design (given a set of constraints). When carrying out an optimisation, an objective function is specified, with the optimum design having the highest (or lowest - depending on the objective) possible value of the objective function. There are many optimisation approaches available and these can generally be split into two categories, gradient-based and global methods [9].

Gradient based methods locally calculate the gradient of the objective function with respect to the design variables, to decide in which direction to search for improved designs. These approaches are local by nature, meaning only a limited portion of the design space is explored. This has the benefit of reducing the cost of the optimisation, but means that the optimum design will only be a local optimum, and there may be an improved design in a different region of the design space that will be missed.

Due to these reasons, the performance of gradient based methods can be very sensitive to the starting point of the optimisation.

Gradient-free methods are generally global methods. This means that, as opposed to local approaches, they aim to explore the whole design space and are usually far more computationally expensive. Examples of gradient-free methods are genetic algorithm (a process that imitates natural selection by repeatedly modifying a population of individual designs through breeding and mutation) [76] and response surface analysis (where a response surface is created from a sample of points to represent the design space, and this surface is searched to find the optimum).

The SOFT optimisation library (that forms part of the SOPHY system) comprises of various optimisation methods including genetic algorithm, adjoint, response surface methods and the Multi-point Approximation Method (MAM) [124]. In this work MAM is used as it provides an efficient method for exploring large design spaces.

MAM

The MAM method is used for all of the optimisation studies carried out in this work. It is a gradient based method that uses localised Design of Experiments (DoE) and trust regions to efficiently search through the design space.

When using MAM, an initial generation of simulations (chosen by DoE) is carried out around the start point. A response surface is constructed for this region and the sub-optimal point found. The search is then moved to this point, where a new generation is constructed and the process repeated until the search converges on the optimal design. The process is illustrated in Figure 2.3.

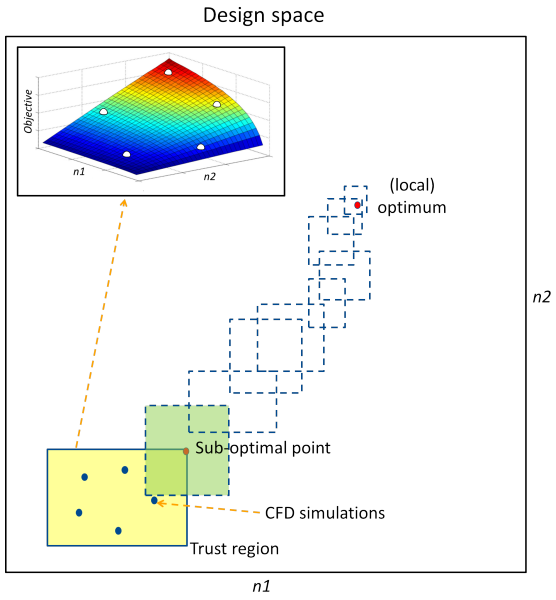

Fig. 2.3 Schematic of the MAM method (example for two design variables n1 and n2).

The MAM method has been successfully used for many optimisation cases by the Rolls-Royce CFD methods team. It has been shown to be the most efficient and consistent approach for a wide range of studies, working successfully for up to 100 design parameters.

2.4.7 Adjoint analysis

Adjoint analysis is used in this work to investigate the surface sensitivity of various designs. The adjoint is the linearisation of the Navier-Stokes equations. Surface sensitivities are the sensitivities of the flow to perturbations to the geometry, and can be used to inform which region of the geometry has the greatest potential to impact on design performance. The sensitivity of the flow field to changes in the surface geometry / design parameters can be computed using finite-differences, but this is very time consuming for a large number of design parameters. The adjoint approach [20] provides an alternative method to compute design sensitivities.

The objective function $I(U,\beta)$ is a function of the flow, U, and the design variables β. R is the Navier Stokes residual. To derive the adjoint equation we begin with the residual of the Navier Stokes equations, which is in theory zero:

$$R(U,\beta) = 0 \tag{2.53}$$

U is the discrete flow solution and β is the design variables. Differentiating with respect to β gives:

$$\frac{\partial R}{\partial U}\frac{\partial U}{\partial \beta} + \frac{\partial R}{\partial \beta} = 0 \tag{2.54}$$

which determines the change in U due to a change in β. Given some nonlinear objective function, $I(U,\beta)$, which, for example, may be a discrete approximation to the lift or drag on an airfoil, the derivative with respect to β (sensitivity of the objective function with respect to the design variables) is:

$$\frac{dI}{d\beta} = \frac{\partial I}{\partial U}\frac{\partial U}{\partial \beta} + \frac{\partial I}{\partial \beta} \tag{2.55}$$

Combining equations (2.54) and (2.55), the following is obtained:

$$\frac{dI}{d\beta} = -\frac{\partial I}{\partial U}\left(\frac{\partial R}{\partial U}\right)^{-1}\frac{\partial R}{\partial \beta} + \frac{\partial I}{\partial \beta} \tag{2.56}$$

Using the fact that the flow variables satisfy the Navier-Stokes equations, $R(U,\beta) = 0$, and introducing the adjoint variable, ν, the following is obtained:

$$\nu^T = \frac{\partial I}{\partial U}\left(\frac{\partial R}{\partial U}\right)^{-1} \tag{2.57}$$

Equations (2.55) and (2.57) can be rewritten as follows:

$$\frac{dI}{d\beta} = -\nu^T \frac{\partial R}{\partial \beta} + \frac{\partial I}{\partial \beta} \tag{2.58}$$

$$\left(\frac{\partial R}{\partial U}\right)^T \nu = \left(\frac{\partial I}{\partial U}\right)^T \tag{2.59}$$

where equation 2.59 is the adjoint equation. The advantage of the adjoint approach is that it depends only on the objective function and not the design parameters, thus allowing the sensitivity with respect to many design parameters to be computed from a single adjoint solution.

Hydra Adjoint [40] is used to provide the blade surface sensitivities in this work: A primal Hydra simulation is first used to provide the flow solution, followed by Hydra adjoint which calculates the flow-adjoint sensitivity and provides the sensitivity of the objective function to changes in the flow.

Once these two relatively expensive simulations are completed, the mesh sensitivities are then mapped onto the surface. This finds the relationship between changes in the flow to changes in the blade surface mesh. Combining these provides the sensitivity (gradient) of the objective function to perturbations of the blade surface, which gives the surface sensitivities.

2.4.8 Summary

This chapter has introduced the methods that will be used to carry out the investigations throughout this thesis. The next chapter (Chapter 3) investigates the use of a novel shaping method to improve transonic compressor blade design.

Chapter 3

Novel Compressor Blade Shaping Through a Free-Form Method

In this chapter, a high pressure ratio, transonic compressor with a strong shock wave and complex flow interaction is simulated and analysed. A novel parameterisation method is then used that allows more shaping freedom than traditional geometry control methods. This provides the ability for an optimiser to select a novel design that controls the shock behaviour and mitigates its impact, thus demonstrating a method for improving the efficiency of transonic compressors through novel shaping and shock control.

3.1 Compressor blades

3.1.1 Overview

The energy released by the combustion process in a gas turbine engine is proportional to the mass of air consumed and the pressure of the air [42]. Hence, to extract sufficient energy from the fuel to produce useful work at reasonable efficiencies, the air entering the gas turbine engine needs to be significantly compressed before the combustor. This is the role of the compressor.

Axial compressors are made up of several alternating stages of rotor and stator blades. Rotor blades add energy to the flow by turning and accelerating the air. The following stator blades then straighten out the flow (ready for the following row) and convert the increased velocity into static pressure. This results in a gradual increase in total pressure through the compressor. The flow path of the axial compressor decreases in cross-sectional area in the direction of flow as the volume of air reduces from stage to stage. These effects can be seen in Figure 3.1a.

Because the airflow in an axial compressor is generally diffusing, it is as a result very unstable. To maintain high efficiencies and stability many compressor stages are required, with a small amount of work done by each stage. In comparison to a turbine, a large number of stages are needed in a

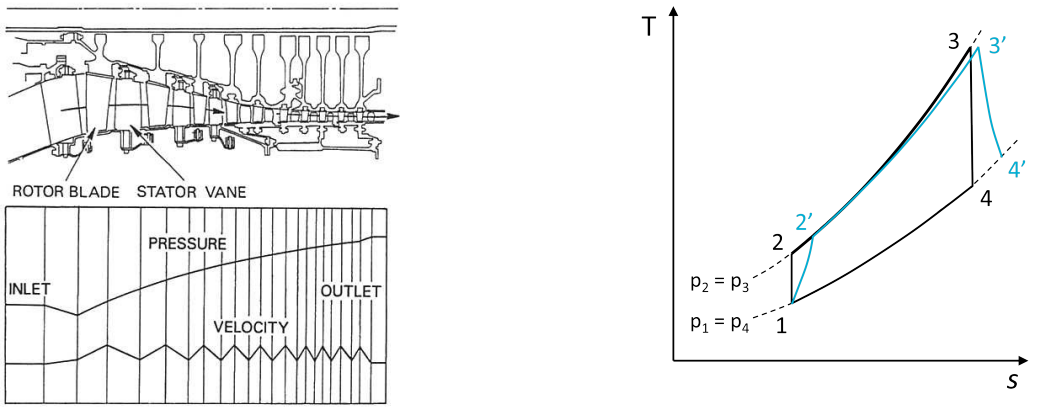

(a) (b)

Fig. 3.1 (a) Pressure/velocity variation through an axial compressor, (b) Brayton T-s cycle diagram.

compressor. This keeps the diffusion rate small and is also required due to the limited amount of turning possible through each stage without separation occurring.

To keep the flow attached and the compressor operating efficiently, the amount of turning that can be achieved is limited compared to turbines. The greater the number of compressor stages however, the greater the weight of the engine. Therefore the capability to increase the loading across a single row or stage while maintaining efficiency would be highly beneficial.

3.1.2 Efficiency

Jet engines operate on the principle of the Brayton cycle. An ideal Brayton cycle is depicted by the solid black lines in Figure 3.1b. In an ideal scenario, the pressure increases through the compressor isentropically (point 1 to 2 in Figure 3.1b). This is not realistic however. In reality viscous effects in the flow cause an increase in entropy, resulting in an additional increase in temperature through the compressor (point 1 to 2'). This results in more work being required by the compressor to produce the same pressure increase. More work is therefore extracted from the turbine to drive the compressor, and less is available for thrust, resulting in greater fuel burn to achieve the same thrust. The isentropic efficiency of a compressor (blade or stage) is calculated from the ratios of the ideal and actual total temperature rise between the inlet and the outlet, namely

$$\eta = \frac{T_{02s} - T_{01}}{T_{02} - T_{01}} \tag{3.1}$$

where η is the isentropic efficiency, T_0 is the total temperature, the subscript denotes the point in the cycle and s indicates the ideal value. It is this efficiency that blade designers aim to increase.

The ideal (isentropic) total temperature is calculated from the total pressure ratio:

$$T_{02s} = T_{01}\left(\frac{p_{02}}{p_{01}}\right)^{(1-\frac{1}{\gamma})} \tag{3.2}$$

where p_0 is the total pressure and γ is the ratio of the specific heats.

3.1.3 Losses in compressor blades

There are several main sources of loss for compressor blades that reduce their efficiency. Reducing the impact of, or eliminating any of these would improve the efficiency of the compressor and reduce fuel burn. These include boundary layer growth and separation, shock waves (shock waves are non-isentropic processes), tip leakage flows, secondary flows (e.g. radial / endwall flows) and mixing losses (e.g. in the wake). The aim of the optimisation process employed in this section is to reduce blade loss and increase efficiency by redesigning the blade geometry.

3.1.4 The need for novel shaping

Modern compressor blades are well designed, having been iteratively improved upon over the years. As the available margin of performance improvement using traditional methods begins to diminish, a more novel and flexible design approach may prove to be beneficial.

Compressor blades are traditionally designed firstly using 1D through flow analysis, then as 2D blade sections with the final 3D design constructed by radially stacking these. It is generally quite expensive to do in-depth 3D design work, and by that stage the room for design change is fairly limited. Blade designs therefore traditionally have limited 3D shaping, apart from the twist required from hub to tip to align the blade section with the relative velocity component that varies as the tangential velocity increases with radius.

3D design and optimisation has the potential to improve upon standard blade designs by finding geometries that a designer wouldn't have time to consider, or are outside of the traditional design limits. The work in this chapter uses a novel shaping approach combined with optimisation to demonstrate the benefit that can be achieved over other methods.

3.1.5 Transonic compressor blades

Transonic compressor blades are widely used in aircraft engines as they allow the opportunity to increase stage loading and achieve higher pressure ratios, therefore reducing the overall number of stages required and engine weight. A consequence of their increased loading and transonic behaviour is that the flow becomes complex and is dominated by the shock behaviour. The shock strength is highest towards the tip of the blade, as the rotational speed increases the relative flow velocity that the blade sees. Transonic rotors often have supersonic relative flow velocities in the tip region. The shock waves that occur in transonic rotors are irreversible processes that produce entropy and loss

which reduce the blade efficiency. The interaction between the shock and the secondary flows can also cause further detrimental effects. Shock-boundary layer interaction can trigger flow separation, and the interaction of the tip leakage flow and the shock can cause the tip leakage vortex to expand in size and even burst [45]. This is thought to be a critical feature of transonic rotor stall.

3.2 The case under investigation - NASA Rotor 37

In this chapter a novel shaping and optimisation method is used to redesign a highly loaded, transonic compressor blade. To allow comparison to previous work in this topic, and to benchmark the results, a well established test-case that has been studied intensively in the literature is selected. NASA Rotor 37 is an axial compressor rotor blade designed at the NASA Lewis Research Center in the 1970s. It was developed as part of a program aiming to achieve high pressure ratios using minimal stages, and is therefore a highly loaded blade, with a pressure ratio of 2.1. It is a well-documented case, having been extensively tested [93] [48] and simulated as part of a turbomachinery validation study [93] [39] [48] [24]. It is a transonic rotor with a tip speed of 454 m/s and a peak Mach number of approximately 1.4. Despite being just a single row with a circumferentially uniform inlet flow, it has historically been a real challenge for CFD simulation. The very high pressure ratio, strong shock wave-boundary layer interaction, large tip-leakage vortex and highly separated flow mean that it still poses challenges for modern turbomachinery solvers. Rotor 37 is such an interesting case that it has been the subject of review articles that highlight the complexity of matching experimental and computational measurements and the associated uncertainties [36] [32]. Figures 3.2a and 3.2b detail the Rotor 37 geometry.

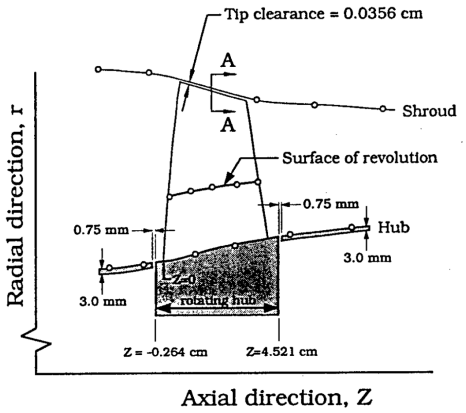

(a)

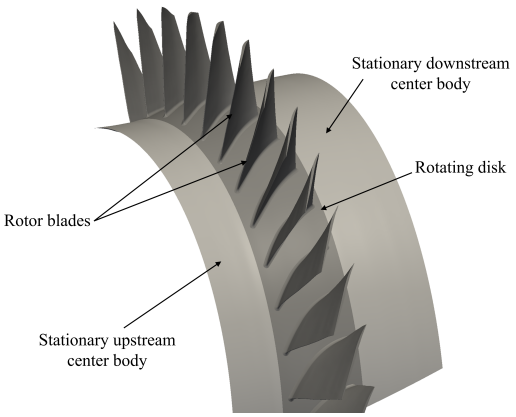

(b)

Fig. 3.2 (a) Meridional description of NASA Rotor 37, (b) 3D schematic.

Table 3.1 Comparison of some previous optimisation results for Rotor 37.

Paper	Setup				Results		
	Parameterisation method	**No. of parameters**	**Optimisation method**	**Constraint**	**Efficiency**	**Pressure ratio**	**Discussion**
Benini 2003 [11]	Hub, mid and tip profiles described by camber and thickness distributions	23	Multi-objective evolutionary algorithm	Fixed mass flow. Throat area constrained ±0.2%	+1.5%	0	Unchanged choke. Blade profile leaned significantly in direction of rotation
Duta, Shahpar and Giles 2007 [40]	Sweep, lean, skew, LE and TE re-cambering (at five heights) plus skew LE and TE re-cambering at root	28	Adjoint gradients with SQP optimiser	Fixed mass flow. ±2% PR	+1%	-2%	Efficiency increase across characteristic, apart from at choke where constant
Chen et al. 2007 [22]	Stagger angle, maximum camber location and maximum thickness location for hub, mid and tip, plus blade solidity	10	Response surface method with zooming algorithm	-	+1.73%	+2.476%	Compressor solidity changed
Wang et al. 2008 [125]	Camberline and thickness distribution at 27, 54 and 81% span	12	Response surface model searched with genetic algorithm (NSGA-II)	±0.5% mass flow	+0.7%	+0.66%	- LE and TE positions not moved, limited benefit +1.5% benefit available if accept PR decrease
Samad and Kim 2008 [102]	Sweep, lean and skew of blade	3	Response surface model with SQP, multi-objective	Fixed mass flow	+1.41% / +0.51%	-1.55% / +1.25%	Competing objectives of PR and efficiency
Shahpar, Polynkin, Toropov 2008 [111]	Sweep, lean, skew, LE and TE re-cambering at 6 control heights	30	MAM	±1% PR ±0.5% mass flow	+1.9%	-	Blade is highly leant from one side to the other from hub to tip. Choke point is reduced
Polynkin, Shahpar, Toropov 2010 [87]	Sweep, lean, skew, LE and TE re-cambering at hub plus skew, LE and TE recambering at tip	8	Metamodel based on genetic programming, ARMOGA	Von Mises stress constraint. ±0.5% PR ±0.5% mass flow	+1.75% / 1.05% with MDO	+0.5%%	Initial optimised result had large sweep at tip, stress constraint reduced this
Brooks, Forrester, Keane, Shahpar 2011 [17]	Sweep, lean, skew, LE and TE re-cambering (at five heights) plus skew LE and TE re-cambering at root	28	Response surface models constructed from combination of expensive high-fidelity data and cheap low-fidelity data	±1% PR ±0.5% mass flow	+2.34%	-0.98%	Large proportion of cheap CFD runs failed to converge because of coarse mesh. Unacceptable reduction in choke mass flow

3.3 Previous optimisation studies

Several previous studies have been carried out to optimise Rotor 37 [11, 40, 22, 125, 102, 17, 111, 87]. These have had reasonable but varying degrees of success. A wide range of parameterisation methods and optimisation techniques have been used and a summary of this previous work is given in Table 3.1. Several conclusions can be drawn from analysing the results presented in the table. It can be seen that blade parameterisation was mostly done through modifying either the camber line and thickness distributions, or using engineering parameters such as sweep, lean, skew and leading/trailing edge re-cambering. A similar benefit is achievable through the two approaches. Various numbers of parameters were used, with the maximum being about 30. A wide range of optimisation techniques were also used, and therefore a range of computational expense incurred. Some optimisations were multi-objective, aiming to increase both efficiency and pressure ratio. While some struggled to achieve this [11, 102], it was predicted by two studies [22, 125] that this is actually possible without using a multi-objective approach. Some optimisations resulted in highly leant blades that would likely fail a mechanical stress constraint [11, 111] and some also found their optimised design reduced the choke mass flow of the blade, restricting its usability [17, 111]. The maximum efficiency benefit predicted by the studies analysed was in the range 1.7-1.9% (without decreasing PR). The results provide a useful benchmark to compare optimisation results from the current study. Uncertainty in the varying accuracy of each simulation must also be taken into account however.

3.4 Parameterisation

As described in Section 2.2, the parameterisation method selected is critical to the outcome of an optimisation. In this study, the flexible Free From Deformation (FFD) approach is used. In this section, FFD is described and compared to a commonly used blade parameterisation approach, the use of engineering parameters.

3.4.1 Engineering Parameters

Often used engineering parameters include Sweep (movement along engine axis), Lean (circumferential movement), Skew (solid body rotation), leading edge re-cambering and trailing edge re-cambering [110]. Using engineering parameters to control blade shaping has the benefit that conclusions relating performance to design parameters can be made, as the parameters are intuitive to the designer. The design parameters can also be translated onto other geometries and this approach requires a smaller number of parameters than a free form approach.

3.4.2 Free Form Deformation

In Free-Form-Deformation (FFD) [28] the solid part to be perturbed is surrounded by a grid of control points that are linked to the solid shape. Moving the position of these points in turn perturbs the geometry. A one-to-one correspondence is established between points within the original and deformed volumes of space. Objects embedded within the original volume are thus deformed by mapping the point-set representing the object to their corresponding points in the deformed volume. Figure 3.3a shows an example 3x3 FFD grid and the shape deformation of a 2D section. It can be seen how it is a combination of the movement of all points that is responsible for the overall blade shaping. Figure 3.3b shows the movement of three FFD control grids (which are then interpolated radially) and the resultant 3D shape change. The Free-Form-Deformation approach has the benefit of providing more shaping freedom and greater design potential. It has a flexible design space that includes and extends upon other design spaces, and the potential to perturb the blade surface as well as the bulk shape. The drawback is the increased cost and non-intuitive nature of the approach.

Both parameterisation methods (FFD and engineering parameters) are available through use of the in-house PADRAM [110] design space. This carries out shape perturbation to 2D radial blade sections during construction of the CFD mesh, maintaining an identical topology to the datum design. The design parameters are specified at a number of radial positions and then smoothly interpolated through the span onto each blade section, making use of a cubic B-spline to create the radial stacking. Several improvements to the PADRAM FFD technique were added by the author during this project to improve its flexibility and control of the point positioning.

In this study, the FFD parameterisation approach was selected as the aim of the study is to investigate what benefit can be achieved through the use of a flexible blade design method. The main advantage of FFD parameterisation is that it provides a greater range of shaping flexibility.

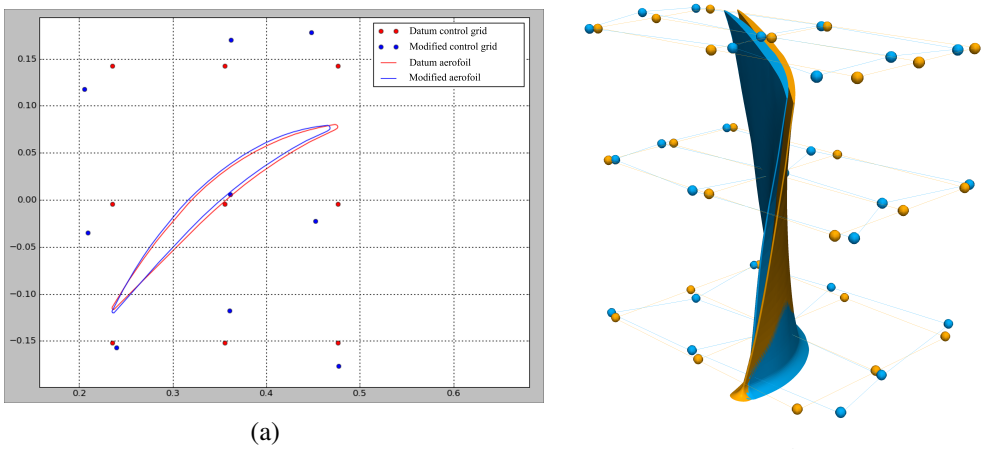

(a) (b)

Fig. 3.3 Examples of the FFD deformation approach (datum orange, perturbed blue) (a) 2D, (b) 3D.

3.5 Simulation

3.5.1 Mesh

The fine multi-block structured mesh (shown in Figure 3.4) was generated by PADRAM [110]. The y+ is of the order of 1. The relatively large number of cells (4.27 million with 30 cells in the tip gap, 200 around the blade perimeter, 25 across each passage and 128 in total radially) is used to provide a mesh independent solution (see Figure 3.5) and to ensure that the complex flow features are adequately captured.

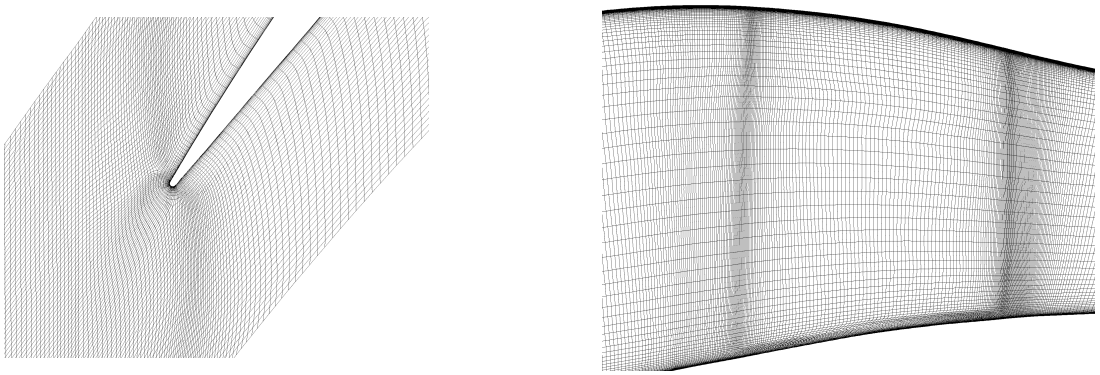

(a) (b)

Fig. 3.4 The fine PADRAM mesh used (a) Radial slice near LE, (b) Meridional view.

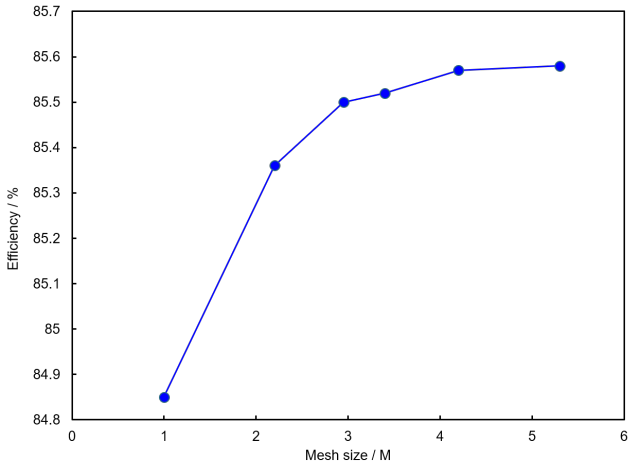

Fig. 3.5 Mesh independence.

3.5.2 CFD set up

The CFD setup is shown in Figure 3.6. At the inlet, a radial distribution of total pressure and temperature (based on the original experimental values [39]) is specified. At the outlet, a value for circumferentially mixed-out and radially mass-meaned capacity ($\frac{\dot{m}\sqrt{T_{tot}}}{P_{tot}}$) is used to ensure fair comparison between designs. Periodic boundaries are used for the front and back portions of the domain shown in Figure 3.6 to represent full annulus flow. Stationary walls are treated as adiabatic viscous walls and the rotational speed of the non-stationary portions of the domain is 1800.01rads^{-1}, as specified in the experiment.

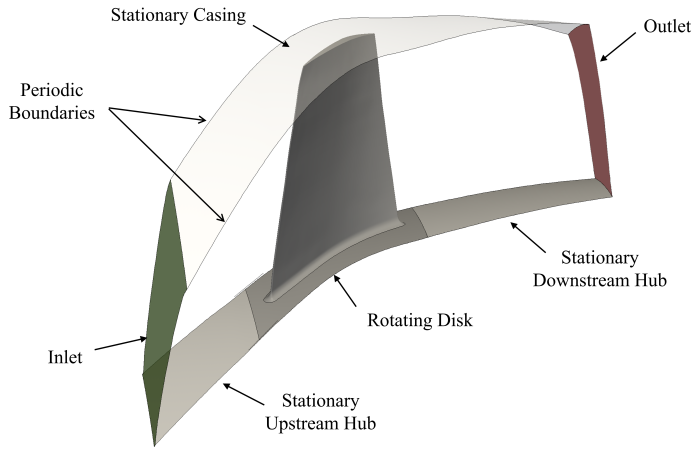

Fig. 3.6 The CFD domain used.

3.5.3 Validation

As previously alluded to, many studies have struggled when matching simulations of Rotor 37 to the experiment. A wide range of work has been undertaken to investigate the discrepancy found between simulation and experiment, with the primary work being the 1994 ASME/IGTI blind test case study, in which a wide range of codes were used to simulate the rotor, with no knowledge of the experimental values. A huge variation was seen between the predictions, prompting analysis by Denton [36]. Recent work has also been carried out by Chima [24] and Hah [48]. Most codes have been shown to over-predict the pressure and temperature ratios, with the efficiency prediction varying widely, depending on variations in those ratios. Seshadri et al. [104] looked at the range of previous simulations carried out and concluded that, in general, if the pressure ratio is well predicted, the efficiency values are under predicted compared to the experiment. These differences are sometimes attributed to uncertainty in the experimental measurements, but are also the result of the inability of CFD to fully resolve and correctly predict some of the complex physics.

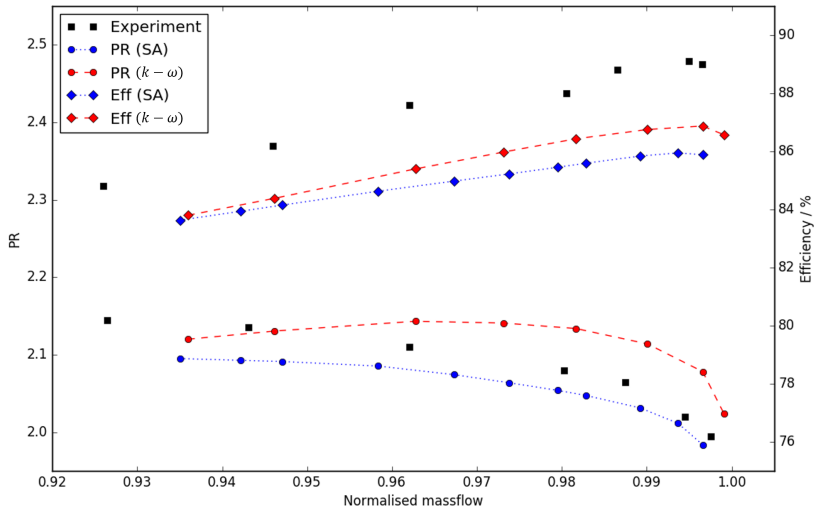

Fig. 3.7 Simulated characteristics vs experimental data.

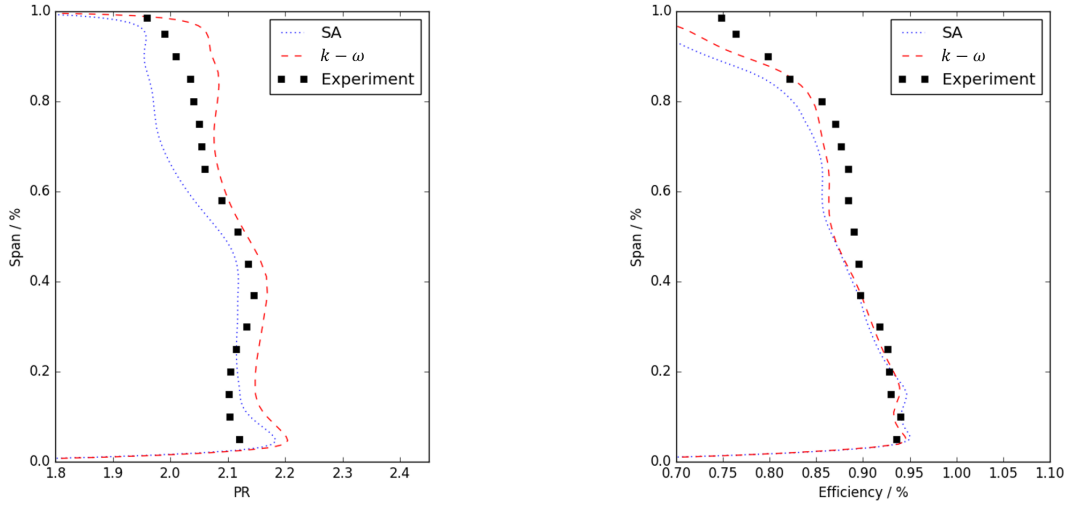

(a) (b)

Fig. 3.8 Radial profiles vs experimental data. (a) PR, (b) Efficiency.

It is important to demonstrate that the current simulation lies within the expected region of agreement with the experimental results. Figure 3.7 shows the simulated rotor characteristics versus experimental data for both the Spalart-Allmaras and k-omega-SST turbulence models.

It can be seen that the two turbulence models give quite different results. The k-omega-SST model over-predicts the pressure ratio, while the Spalart-Allmaras model matches this closer, but still under-predicts it. The k-omega-SST model more closely matches the efficiency, reflecting the 'trade-off' seen in a range of previous simulations [39], where the closer the prediction of one flow property, the further away the prediction of the other. The efficiency prediction is between 1.5 and 2% below the experimental values for both models. This comparison with experimental data is typical of those seen by other researchers.

Figure 3.8 gives the radial profiles of total PR and efficiency at 98% of simulated choke compared to the experimental values at 98% experimental choke. The radial trends have been captured fairly well, with some deficit from the experiment. Similar differences are seen between the two turbulence models as demonstrated in the characteristics. The results lie either side of the experimental data, demonstrating the variation that is found for this case. Seshadri et al. [104] demonstrated that a small hub leakage flow (often missing from simulations, including this study) was the culprit for the inability of CFD to capture the hub pressure deficit. The deficit at the tip can be attributed to the difficulty capturing the complex behaviour of the tip leakage vortex. An analysis of the uncertainties associated with predicting the tip leakage behaviour of Rotor 37 is given by Seshadri et al. [105]. The choke mass flow found in the simulations was 20.91kgs^{-1}, matching quite closely the experimental value of 20.93kgs^{-1}.

Figure 3.9 shows a comparison of the design point simulated relative Mach number contours at 70% span compared to the experiment. The shock and shock-boundary layer interaction can be seen to be well captured.

3.5.4 Flow field for the datum case

Figure 3.10 shows the most significant features of the rotor 37 flow field. As mentioned previously, a strong shockwave is present and shock induced boundary layer separation occurs. It can be seen how the strong shock of Rotor 37 causes complex shock-boundary layer interaction and a large shock-induced separation (this can be seen by the thickening of the boundary layer shown in Figure 3.10). At the point where the shock impinges on the suction surface, its interaction with the boundary layer causes it to separate and a large wake forms.

In Figure 3.10b the volume of reverse flow (orange isosurface of zero axial velocity) shows the 3D extent of this separation. It can be seen how streamlines seeded near the hub of the blade travelling through this separated region are immediately thrown radially by the rotation of the domain. The streaklines on the blade surface help to indicate the point of separation.

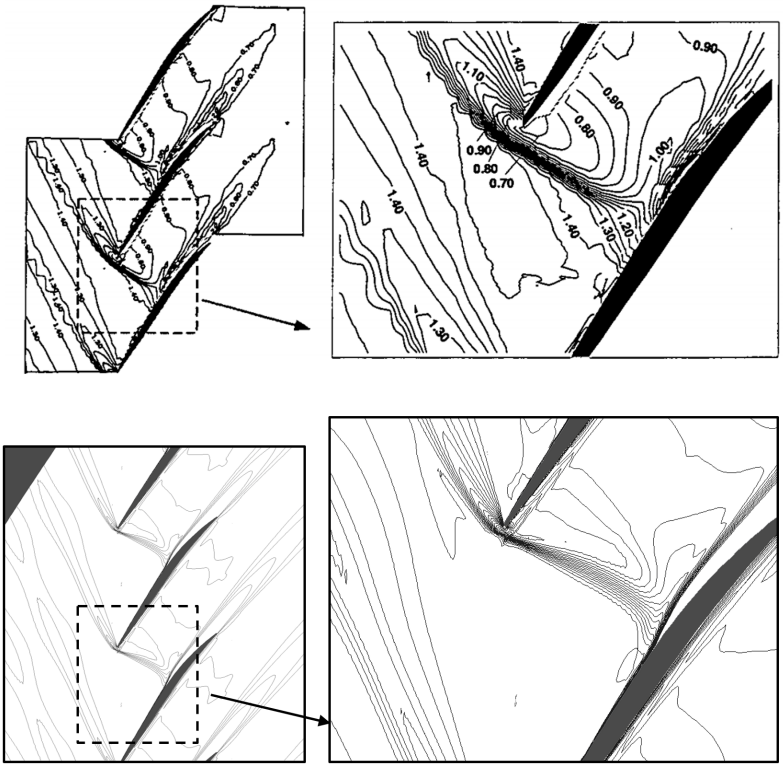

Fig. 3.9 Experiment [121] (top) vs simulated (bottom) relative Mach contours at 70% span.

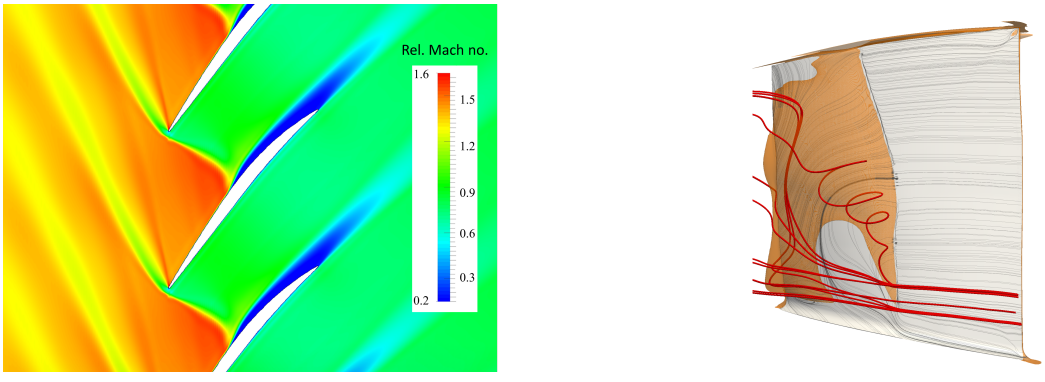

(a) (b)

Fig. 3.10 Shock induced separation on Rotor 37 (a) Relative Mach number contour at mid-span, (b) Suction surface streamlines and separation.

Boundary layer separation

The key feature of the flow is shock induced boundary layer separation. Boundary layer separation occurs when the boundary layer can no longer stay attached to the wall. Figure 3.11 shows a schematic demonstrating boundary layer separation.

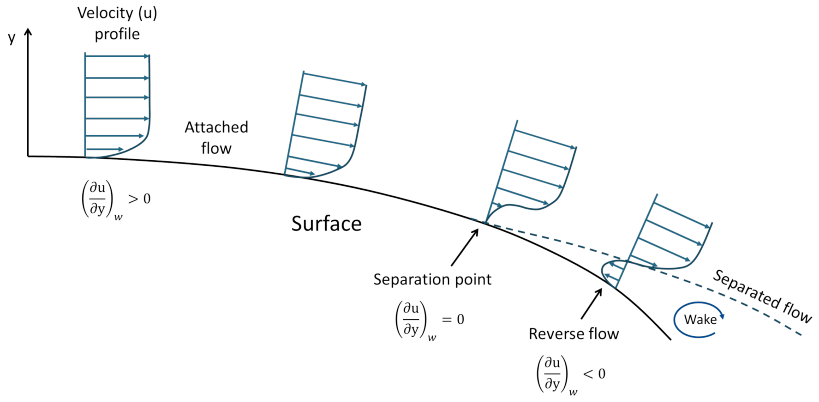

Fig. 3.11 Schematic of boundary layer separation.

Wall-bounded flows are characterised by the presence of a boundary layer where a velocity profile extending normal to the wall can be clearly distinguished. The flow at the wall has zero velocity with the velocity (u) increasing with increased distance from the wall (y) until it reaches the free-stream velocity outside of the boundary layer. When separation occurs, the flow right at the wall in the boundary layer stagnates and subsequently reverses, as can be seen in Figure 3.11. At the point of separation the velocity gradient at the wall becomes zero (as is highlighted in Figure 3.11). This can occur if the boundary layer becomes too thick, which can be caused when an adverse pressure gradient is present (i.e. pressure increasing in the direction of flow). This causes the velocity gradient at the wall to reduce until the point where the flow can no longer remain attached and separation occurs.

In shock-induced separation, it is the sudden increase in pressure and decrease in velocity across the shock that causes the boundary layer to separate. The shock creates a very strong adverse pressure that acts in the opposite direction to the flow. As a result, the velocity gradient near to the wall is decreased, and becomes low enough for the flow near the wall to reverse, leading to separation.

3.6 Adjoint sensitivity analysis

Adjoint sensitivity analysis is a useful tool that can be used to provide information on the sensitivity of an objective function to changes in the geometry. This is used here to analyse in more detail the Rotor 37 flow behaviour. Adjoint sensitivities are found by solving the adjoint equations [57], as described in Chapter 2.

Here, the 'adjoint sensitivity' is the sensitivity of efficiency (as a percentage) to surface point deformation (in mm) normal to the surface. This can be used to inform which regions of the blade will have the greatest impact when modified, and are therefore most important to control during an optimisation.

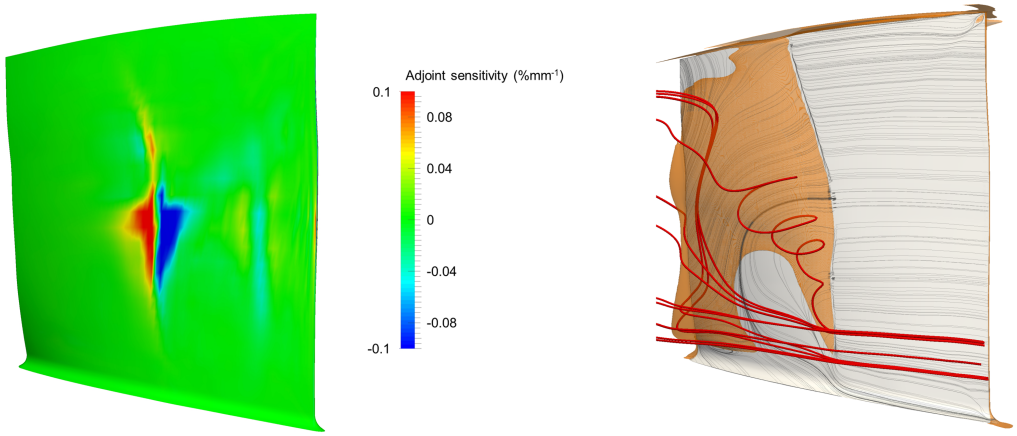

(a) (b)

Fig. 3.12 Suction surface: (a) adjoint sensitivity and (b) limiting surface streamlines with 3D streamlines (red) and reverse flow volume (orange).

Figure 3.12a shows the result of the adjoint surface sensitivity analysis. A positive value indicates that the surface should be moved outwards to improve the efficiency, and vice-versa for a negative value. It can be seen that the efficiency is most sensitive to positive and negative perturbations of the suction surface around mid-chord at mid-span. It makes sense that this highlighted region is located around the shock wave, as this is responsible for the creation of a large amount of loss. Figure 3.12b shows the complexity of the flow structure in the separated region. Any small surface change in this region is therefore likely to result in a significant change to the flow and in turn the blade efficiency. This is why it is highlighted by the adjoint sensitivities.

It is interesting to note that the most sensitive region appears isolated to the mid-span of the blade. One might expect that as the Mach number is highest towards the blade tip, and therefore the shock is strongest, it would be this region that efficiency sensitivity is greatest. This can be understood by analysing the suction surface static pressure distribution shown in Figure 3.13a. It can be seen that it is in fact near mid-span that the difference in pressure across the shock is greatest, and therefore the strength of the shock will be highest. Moving towards the tip of the blade, the pressure upstream of the shock is higher, and the shock strength therefore lower. A slice of entropy perpendicular to the flow and just downstream of the shock (looking upstream) in Figure 3.13b shows the radial variation in the separated flow that forms the wake. The greatest amount of separation clearly occurs around mid-span, which agrees with the prediction of greater shock-strength in this location. The increased

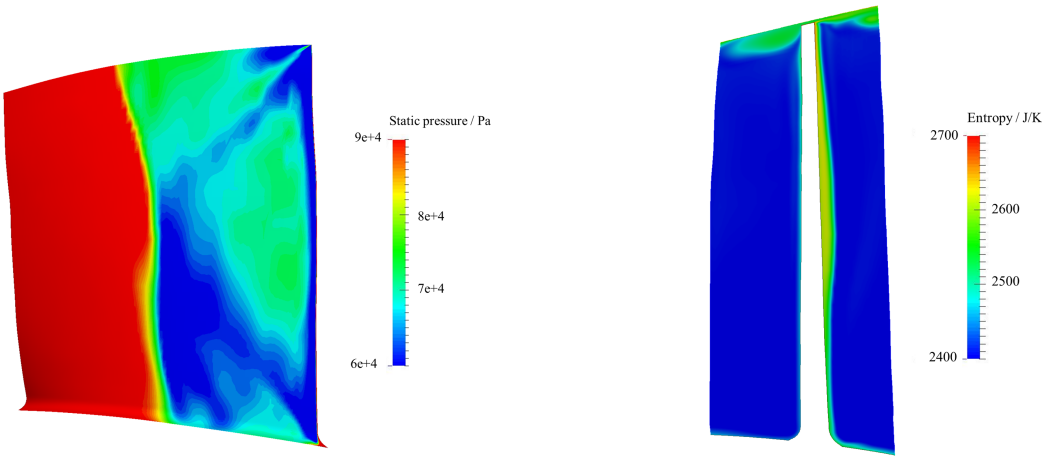

(a) (b)

Fig. 3.13 (a) Suction surface pressure distribution, (b) slice of entropy just downstream of the shock.

adjoint sensitivity around mid-span can therefore be explained by the higher shock strength and greater shock-induced separation occurring in this region, meaning changes to the blade surface in this location will have the greatest impact on efficiency.

3.7 Optimisation and results

3.7.1 Optimisation set up

The optimisation technique used for this work is the MAM method described in Seciton 2.4.6. The FFD optimisation discussed in this paper was carried out using an exit capacity that provided 98% of choke mass flow (the design point) for the datum design. The objective optimised was efficiency, with the pressure ratio constrained to ±1.5% of datum. Each MAM generation contains 60 simulations. The design parameters used for the optimisation consist of three 3 x 3 FFD grids controlled at 0, 50 and 100% span, resulting in a total of 54 parameters. The 2D movement of each point in the grid was limited to a quarter of the inter-point distance.

The computations were carried out on the Rolls-Royce CFMS high-performance linux cluster. Five to ten simulations, each using forty-eight cores, were run in parallel, depending on the cluster load. Each was converged to steady state, with the few that failed to reach the minimum residual threshold rejected. About 500 simulations were carried out in total for the optimisation, with it taking about two weeks to converge.

3.7.2 Optimisation history

The convergence history for the optimisation is shown in Figure 3.14. The objective was to minimise 100%–efficiency. The small number of failed runs have been removed for clarity. It can be seen that MAM quickly moves away from infeasible designs that violate the pressure ratio constraint, and within eight generations converges upon what was found to be the optimum design, performing impressively well considering the large number of design parameters.

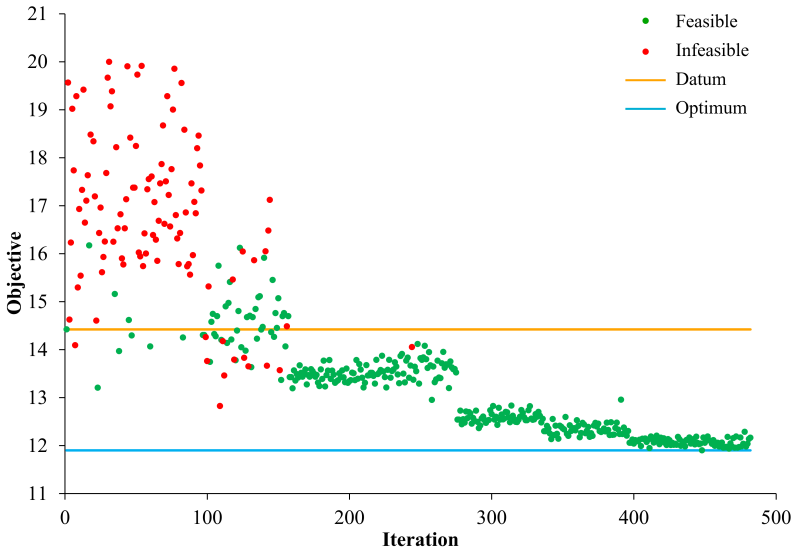

Fig. 3.14 The MAM optimisation history.

3.7.3 Optimised shape

The three-dimensional representation of the FFD control grids and blade shapes for the optimised and datum designs can be seen in Figure 3.15b. Figure 3.15a shows how the FFD control points have moved to create the 2D shape of the optimised design at mid-span.

The datum and optimised blade sections through the span are shown in Figure 3.16. There is very little visible change in the blade shape near the hub. The optimised shape is leant slightly in the direction of rotation through the rest of the span. The most obvious and significant change however is the s-shape present around mid-height, as shown in Figure 3.16c. This s-shape has been created via the optimiser's control of the mid-span FFD grid. It has then been interpolated from 50% span towards the hub and tip; and the s-shape can therefore also be seen in the 2D slices at 30 and 70% span. The LE region of the blade has been moved parallel to the datum but then around mid-chord the suction surface of the blade begins to slope away from the direction of rotation, forming the s-shape. It can be seen in Figure 3.15a that it is a combination of the individual movements of the control points that pull and push the blade shape into the new design. Both surfaces are perturbed by the

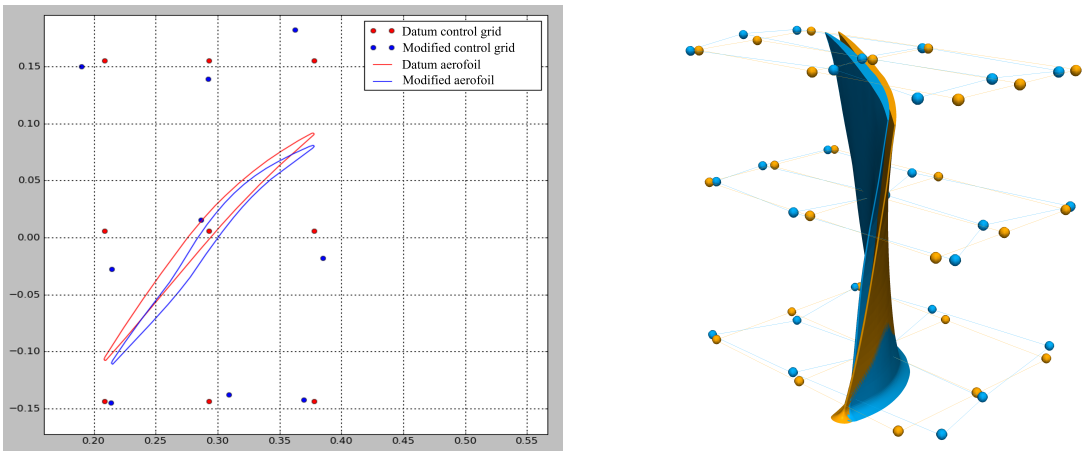

(a) (b)

Fig. 3.15 Optimised and datum FFD control grids and geometries (a) 2D grid at mid-span, (b) 3D FFD grids and geometries, datum orange, optimised blue.

movement of a single point, and therefore where the suction surface has moved outwards, the pressure surface has followed. This has its advantages, i.e. the thickness of the blade is maintained, meaning no significant addition of material or thinning of the blade (affecting the structural properties) has occurred.

3.7.4 Performance comparison

Table 3.2 shows the performance comparison of the datum and optimised designs at the optimised point. It can be seen that an impressive improvement in efficiency has been predicted. The magnitude of efficiency increase is significant compared to that found previously in the literature, and a pressure ratio increase is also predicted – something that other optimisation attempts struggled to achieve, even via a multi-objective approach.

Table 3.2 Comparison of the optimised and datum design performance.

	PR	**Difference in PR**	**Efficiency**	**Difference in Efficiency**
Datum	2.06	1.46%	85.45	1.99%
Optimised	2.09		87.44	

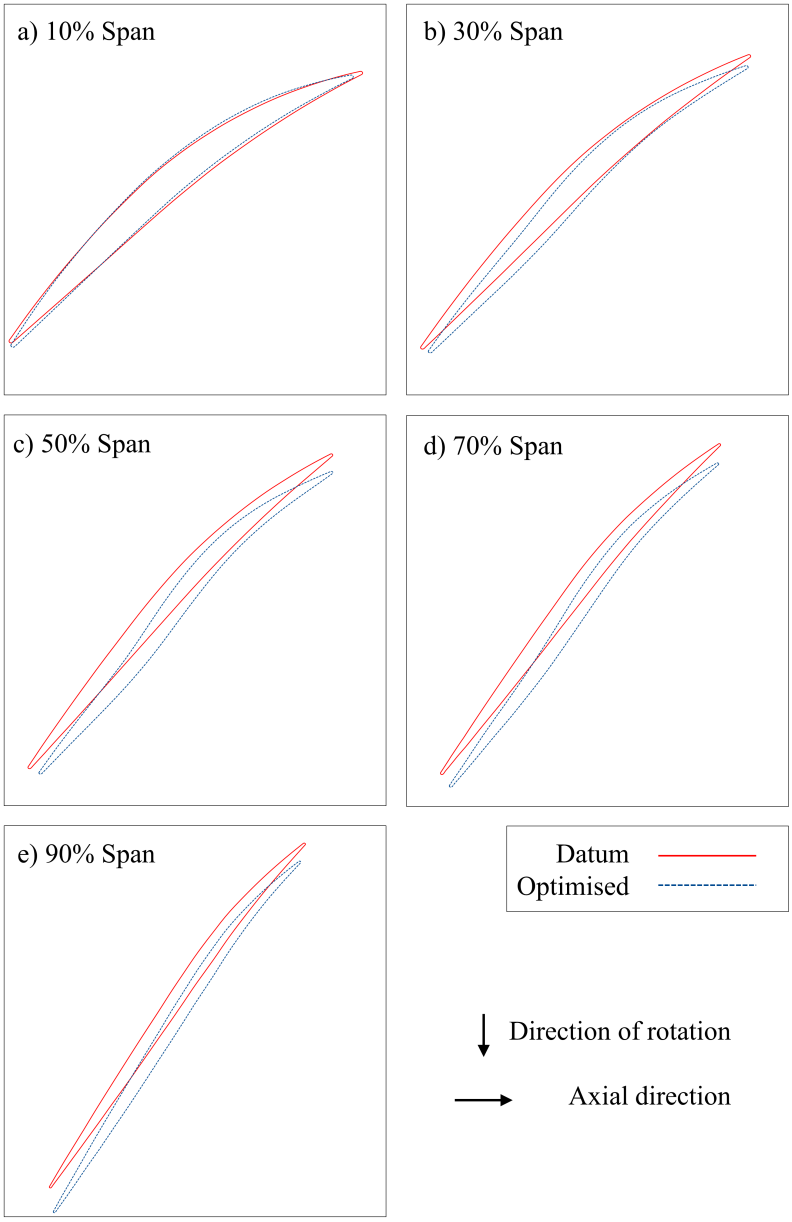

Fig. 3.16 Comparison of the 2D datum and optimised blade shapes.

3.8 Analysis of the s-shape design

The impact of the shaping produced through the FFD optimisation is discussed here in detail. The most notable feature generated by the optimisation is the s-shape of the blade around mid-span. This acts to control the shock and provides several key benefits. This is discussed below in detail.

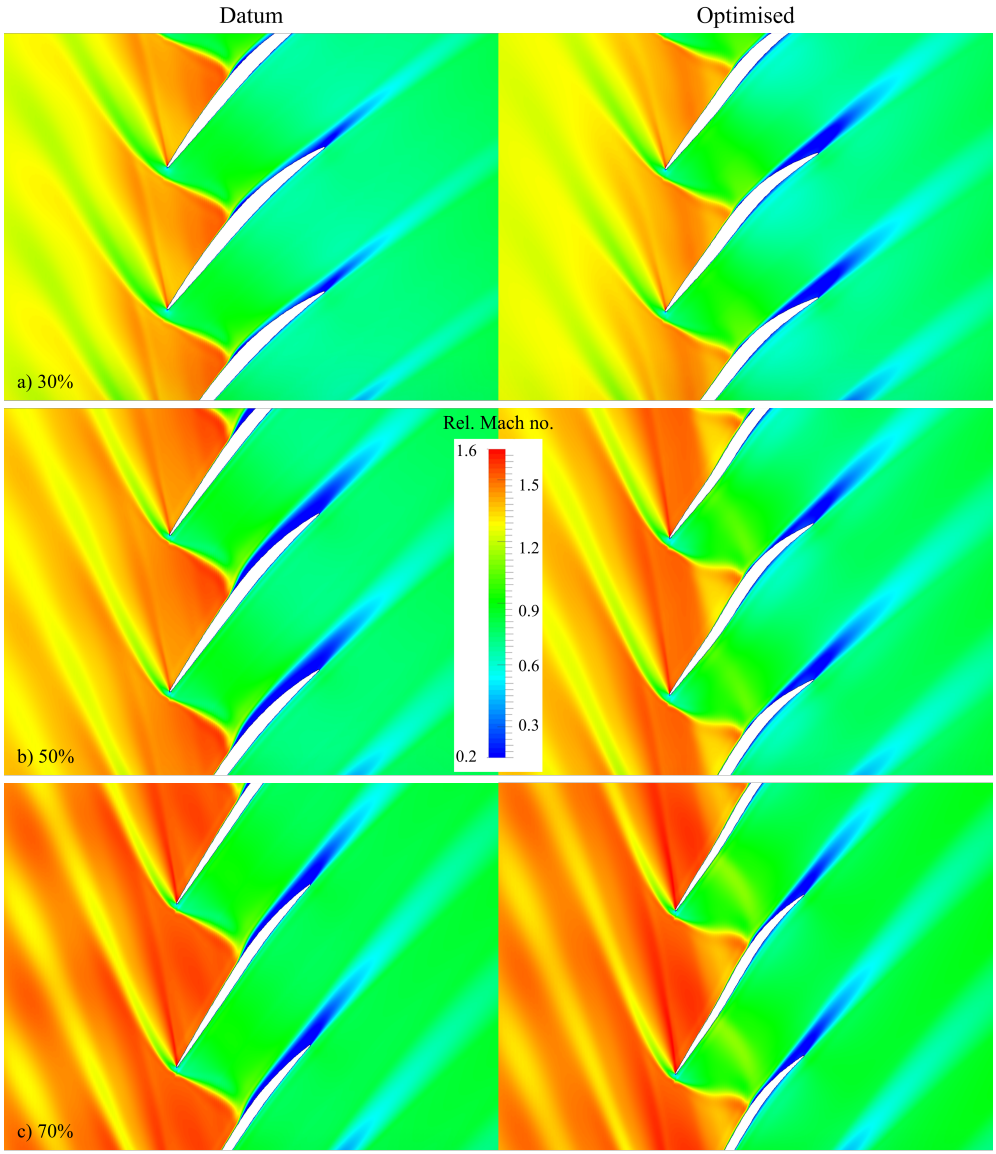

Fig. 3.17 Relative Mach number contour comparison at various span.

Figure 3.17 shows relative Mach number plots of the datum and optimised designs at 30, 50 and 70% span. First, the s-shape and its effect on the shock can clearly be seen. The shock impingement point on the blade has moved downstream and the shock shape has been significantly altered, making it more oblique. The point of the shock-induced separation that influences the blade wake has in

turn moved downstream, and the amount of separated flow has been reduced. Suction surface static pressure contours in Figure 3.18 show how the shock position has moved downstream for the majority of the span, increasingly towards the tip. This trend is reflected in the spanwise distribution of the efficiency benefit from the datum to optimised designs, shown later in Figure 3.22. This highlights the effect that changes to the shock have on the blade efficiency. The contribution of the s-shape produced by FFD can be clearly identified in Figures 3.17-3.19. The s-shape protrudes into the passage around the shock region. The effect of this is to form a 'pre-compression ramp' that causes the flow to slightly increase in pressure and decrease in velocity. Figure 3.18 shows the large region of higher pressure from around 30-100% span upstream the shock.

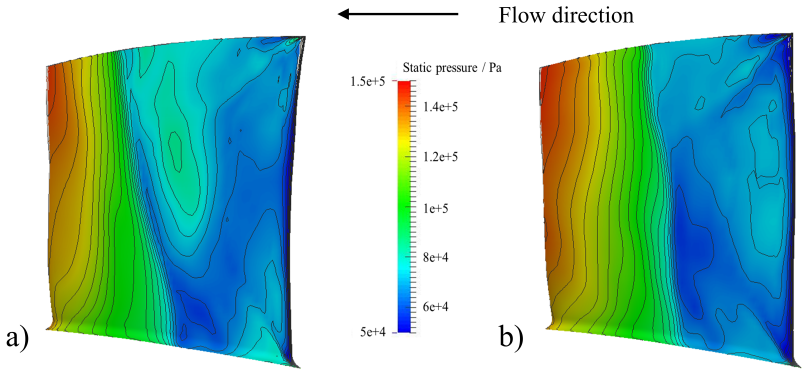

Fig. 3.18 Suction surface static pressure contours a) datum, b) optimum.

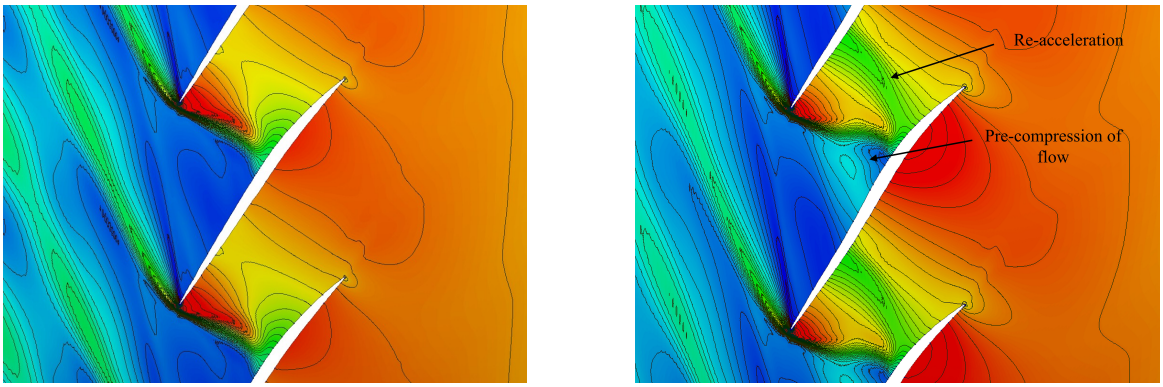

Fig. 3.19 (a) Datum and (b) optimised static pressure contours showing the pre-compression effect. Scale the same as Fig. 3.18.

Figures 3.19a and 3.19b show the static pressure contours at 70% of span. The pre-compression that occurs along the ramp produced by the s-shape can be seen and there is a steady increase in pressure shown by the contours towards the shock. The beneficial effect of the pre-compression is three-fold. It weakens and delays the terminating shock. The resulting Mach number / pressure

difference across the shock is therefore reduced, thus reducing the strength of the shock and therefore entropy production across it. The reduction in shock strength, combined with the change in position of the shock, also reduces the level of shock-induced separation, as seen in Figure 3.17.

3.8.1 Explanation of the shock control effect

Figure 3.20 shows a schematic that describes the shock control effect achieved by the optimised design. As can be seen in Figure 3.20, the effect of the s-shape is to change the almost normal shock seen in the datum design to an oblique (and therefore weaker) shock that terminates with a small normal shock leg. The novel geometry creates a series of isentropic pre-compression waves up to the shock, which gradually increase the pressure before the shock, reducing the pressure difference across the shock and delaying its impingement on the surface. Overall, the same pressure ratio can be achieved by the compressor, but less of it is achieved across the shock, reducing the losses generated. The ability of this geometry to carry out the pre-compression comes from the s-shaping of the blade where the suction surface protrudes into the flow. The blade surface protruding into the flow acts to locally compress the flow and increase its pressure, due to the fact that narrowing of a passage at M>1 increases pressure and reduces velocity. The relative Mach number contours at 50% span seen in Figure 3.17 previously, could give the impression that at the start of the pre-compression region a lambda shock is present. This is a feature seen for shock control bumps that do not maintain curvature continuity where they begin on a blade surface [18]. However, in this case the surface curvature is continuous and a lambda shock is not present. This is confirmed by the lack of a jump in entropy that would be present for a shock (not shown here).

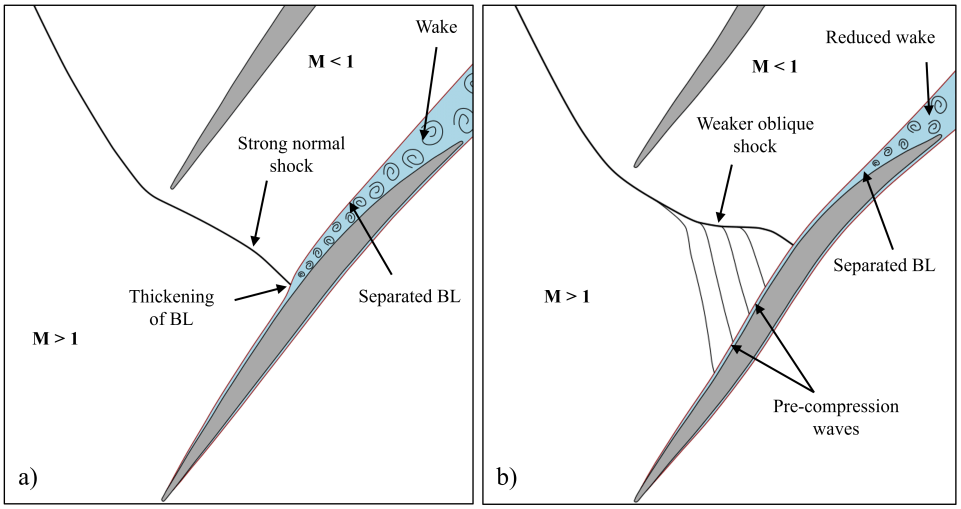

Fig. 3.20 Schematic of shock structures (a) datum, (b) optimised.

3.8.2 Further analysis

Evidence for the reduction in shock strength can be seen by looking at how the relative Mach number and entropy vary along the mid-passage line at 70% span, shown in Figure 3.21a. The shock structure has clearly changed significantly. The pre-compression effect can be seen in the gradual reduction of Mach number just before the shock (producing the 'knee shape' seen). It is interesting to note that this knee-shaped shock structure was also identified in optimised shock control bumps over transonic wings [91]. The drop in Mach number from this point across the shock is shown to be far smaller than for the datum design. The entropy increase across the shock can also be used as a measure of its strength. This has clearly been reduced by an impressive amount, showing nearly an 80% reduction.

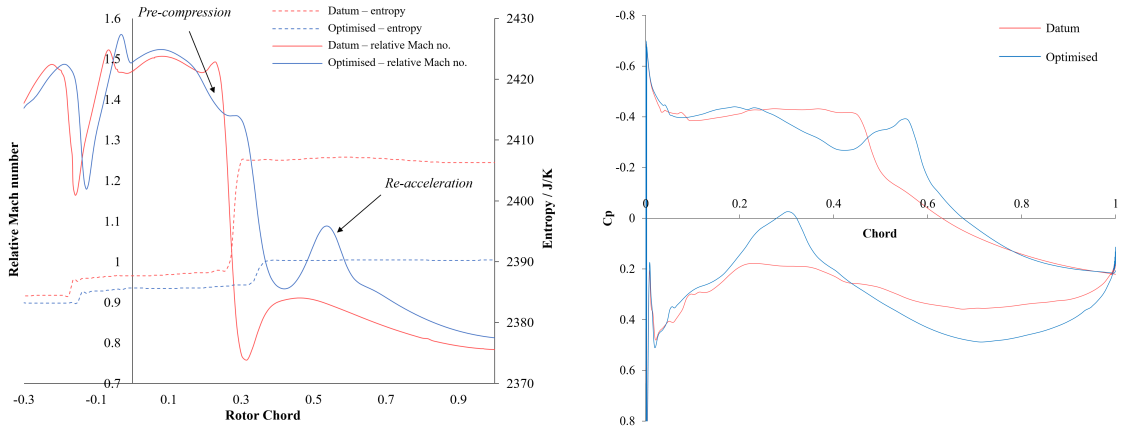

(a) (b)

Fig. 3.21 (a) Relative Mach number and entropy along the periodic boundary, (b) blade loading profiles at 70% span.

The lift plot in Figure 3.21b shows the impact of the shape change on the blade loading. The shock structure on the suction surface has been altered as described in Section 3.8, though the reduction in strength is not as significant as through the passage. The other clear change is the lower pressure region on the pressure surface of the optimised blade. The reason for this can be seen in Figure 3.19b - after the shock a re-acceleration occurs, due to narrowing of the passage caused by the s-shape. This is highlighted in Figures 3.19b and 3.21a. Whilst the original blade was highly front-loaded, the loading has shifted rearward on the optimised blade. A parameter that provides further insight and is known to be important for supersonic blades is the blade solidity, i.e. the blade chord over the spacing (distance between blades in theta direction). It has been shown that lowering the solidity (which reduces viscous losses) increases the blade loading and shock strength [103]. Through this optimisation the FFD parameterisation has slightly reduced the blade chord, as was seen in Figure 3.16. This has lowered the blade solidity, which would normally increase shock strength, but due to the novel shaping achieved, no negative impact of this has been realised.

3.8.3 Three dimensional shaping of the blade

Aside from the s-shaping of the blade discussed above, the FFD shaping of the optimised design has perturbed the 3D design of the blade. As could be seen in Figures 3.15 and 3.16, the blade has been leant in the direction of rotation slightly from hub to tip. This is a modification that was also predicted to provide benefit by previous researchers [111] [11].

Traditionally, rotor blades are constructed from radially stacked, 2D designed blade sections. Using sweep, lean or other three-dimensional shaping methods can provide benefit by modifying the 3D flow field. The three-dimensional shock structure is one of the key features of a transonic compressor flow field, and this is affected by the 3D shaping of the blade.

Previous research has noted that leaning transonic blades in the direction of rotation can lead to efficiency improvement. Bergner et al. [12] observed that the use of a stacking line curved in the tangential plane can induce a significant change in the shock pattern within the blade passage. Ahn and Kim [3] showed the possibility to increase the overall rotor efficiency by skewing the blade toward the direction of rotor rotation. The authors attributed this to reduction of shock intensity and change of shock shape and position at the blade tip. Benini [11] and Jang et al.[60] observed that using a 3D blade curvature can change the shock into two oblique shocks of lesser intensity. Biollo [13] systematically studied various amounts and combinations of lean and sweep for NASA Rotor 37. He concluded that, in short, the higher efficiency shown by a forward (in the direction of rotation) leaned blade can be mainly associated with a more efficient flow diffusion at the outer half span region, due to a more favorable three-dimensional shock structure, reducing the entropy generated by the shock. It was observed that this is partly due to a change in the flow incidence in consequence of an unusual flow pattern around the blade leading edge.

It can therefore be concluded that the 3D change in the shock structure seen can be attributed to both the 2D blade section shaping and the impact of lean modifying the 3D shock shape. It is a combination of the pre-compression effect of the 2D s-shape design and the 3D blade shaping that delays and weakens the shock and improves the rotor efficiency. This demonstrates one of the significant advantages of the FFD technique. Through a single parameterisation approach, global (3D) blade shaping has been carried out, as well as the re-design of the 2D aerofoil sections. In the previous investigations of the effect of lean (described above), the maximum efficiency benefit predicted was in the order of 1-1.5% efficiency benefit. The optimised design used here significantly surpasses that, demonstrating the benefit of the free-form approach.

One of the concerns with leaning blade geometries is that the stresses due to centrifugal loads are increased, and may violate a maximum allowable stress in the material. To prevent this in future work, a mechanical stress constraint could be applied during the optimisation process, or the freedom of movement of each FFD parameter limited to reduce the magnitude of lean that could occur. It is expected that applying either of these methods would reduce the efficiency benefit predicted, as

shown by Polyinkin et al. [87] who observed a reduction in efficiency from their optimisation once a stress constraint was added.

3.8.4 Radial distribution of performance

Figure 3.22 shows the radial distributions of efficiency for the datum and optimised blades. It can be seen that an increase in efficiency has been realised right through the span. Near the hub a small efficiency benefit is predicted, with around mid-span the efficiency benefit becoming larger as the s-shaping of the blade sections helps to weaken and delay the shock. The greatest efficiency benefit is around 70-80%, as the shock is delayed through a combination of pre-compression and the blade lean. It can be seen by comparing Figure 3.22 with Figure 3.18 that it is this delay to the shock that has the greatest impact on efficiency, with the difference between the datum and optimised efficiency profiles and static pressure contours matching through the span.

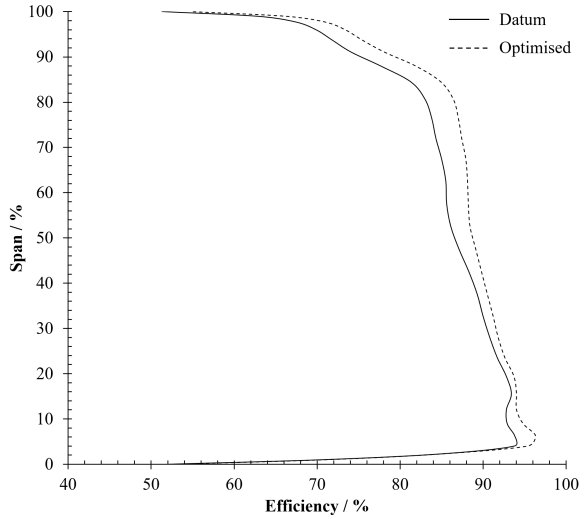

Fig. 3.22 Radial distribution of efficiency for datum and optimised designs.

3.8.5 Adjoint analysis

An interesting feature of the optimised s-shape design is that it conforms with the adjoint sensitivity prediction. The s-shape is localised to mid-span and around mid-chord (shown in Figure 3.12 to be the location most important for improving efficiency). Also of interest is that the direction of the surface sensitivities match; the strong red region in Figure 3.12 (indicating benefit from outward surface movement) is where the second half of the s-shape protrudes outwards into the flow. This demonstrates the usefulness of adjoint analysis to the designer. Adjoint surface sensitivity analysis carried out on the optimised blade is shown in Figure 3.23b. The movement of the shock rearward is clear and it can be seen how the magnitude of the sensitivity on the blade suction surface has been

reduced due to the successful weakening of the shock (though it clearly remains significant to the blade's efficiency). Analysing the peak magnitudes of adjoint sensitivity demonstrates the extent of the benefit predicted. Table 3.3 shows the reduction in surface sensitivity levels for the optimised design compared to the datum. A 30-40% reduction in the magnitude of surface sensitivity has been predicted. This is due to the optimised design having a weaker shock and therefore changes to the surface in the region around the shock would have less impact. This also demonstrates how the design has been improved, indicating there is now less benefit available through geometry change.

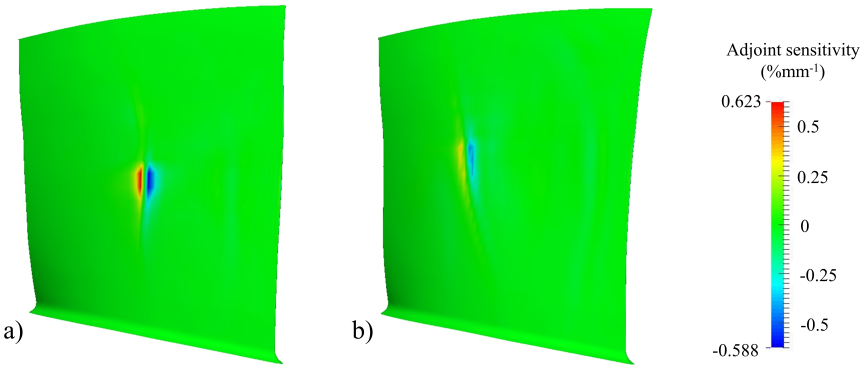

Fig. 3.23 Adjoint efficiency surface sensitivities for (a) datum and (b) optimised designs.

Table 3.3 Adjoint sensitivity difference between datum and optimised designs.

	Sensitivity of efficiency to surface perturbation $\%mm^{-1}$		
	Min	Max	Range
Datum	-0.588	0.623	1.211
Optimised	-0.417	0.384	0.801
Reduction	-29%	-38%	-34%

3.8.6 Use of s-shape in the tip region

An interesting question is whether the localisation of the s-shape to mid-span is due to this region being the most sensitive (as previously discussed), or whether it is in fact a manifestation of the FFD grid being controlled at three radial locations, and the optimiser only finding the s-shape for the mid-span set of design parameters. Would an s-shape for the whole blade have been more beneficial? To partially answer this, the mid-span FFD control grid was copied to the tip and this geometry simulated. This design performed worse than the optimum, giving similar efficiency to the datum. This is partially due to the limited translatability of the FFD parameters (the shape perturbation depends on the initial shape) but also suggests that the s-shape was particularly suited to mid-span.

3.9 Off-design analysis of the optimised design

3.9.1 Rotor characteristics

To properly assess the benefit provided by the optimised blade it is necessary to look at the performance away from the design point. A common concern with flow control designs, such as the s-shape, is that while it may provide benefit at a particular operating point, moving away from this could result in a significant drop in performance. Figure 3.24 shows how both the efficiency and pressure ratio for the optimised design have been increased by a significant margin, and the delta is encouragingly consistent across the range of flow rates. A clear reduction in the choke mass flow can be seen however, which could limit the usability of such a design. This reduction in the choke mass flow is due to a restriction on the passage area caused by the s-shape protruding into the passage, the effect of this can be seen in the re-acceleration after the shock seen in Figures 3.19b and 3.21a.

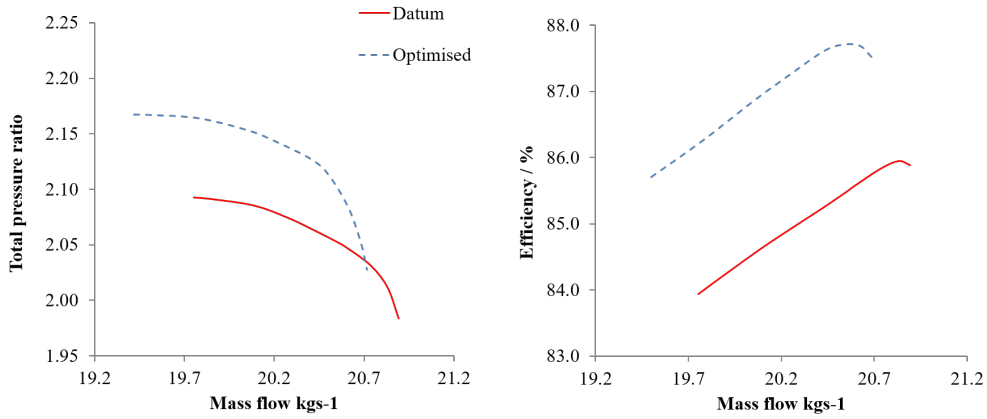

Fig. 3.24 Characteristics of the datum and optimised blades.

3.9.2 Regaining choke point

A reduction in the choke mass flow may cause issues with the take-off performance of an engine, where maximum mass flow through the core is required – rendering this design unusable. Fortunately this can be remedied. Applying a manual skew to the blade (a solid body rotation that modifies the blade design) around the centroid opens up the passage and increases the mass flow that can pass through. Figure 3.25 shows the result of skewing the blade by 0.1 degrees (at the design conditions). Skewing the blade clearly has the desired effect. The choke point has been shifted back to the original mass flow, with no apparent efficiency penalty, though a reduction in stall margin. This might have proved useful to a couple of the studies discussed in Table 3.1, whose final designs suffered from the limitation of a significant reduction in choke mass flow.

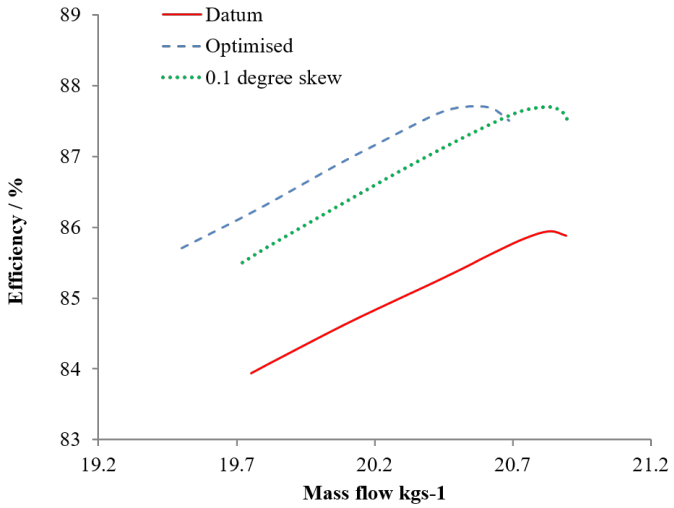

Fig. 3.25 The effect of skewing the optimised design.

3.9.3 Varying rotor speed

A more in-depth study of the off-design performance of the rotor includes assessing its behaviour at different shaft speeds. Taking the design speed from the experiment of 1800rads^{-1} and evaluating at ±5% gives another indication of how the geometry performs away from the design point. The original, optimised and optimised skewed blades have been included in Figure 3.26. An example working line for the datum blade has also been added, with corresponding operating points for each blade indicated.

It can be seen that the optimised (and skewed) designs perform well at both the datum and 105% speed. In fact, the optimised designs significantly outperform the datum at the higher speed. The skewed design has a higher choke mass flow (which is the most critical consideration for take-off condition), the efficiency and pressure ratio deltas are significant and the operating range has been increased. Although this optimisation was only carried out at a single point, it can be seen that one of the key objectives when re-designing a blade (optimising efficiency whilst operating on a specified working line) has been achieved. The skewed design is able to operate on a working line specified by the datum geometry while providing a significant performance benefit. The performance of the optimised designs at part-speed is satisfactory (here efficiency is not the primary concern) though this analysis does indicate that increasing the choke margin through further skewing may be beneficial. The pressure ratio at part speed is greater than datum for the majority of the characteristic and the stall margin has been maintained.

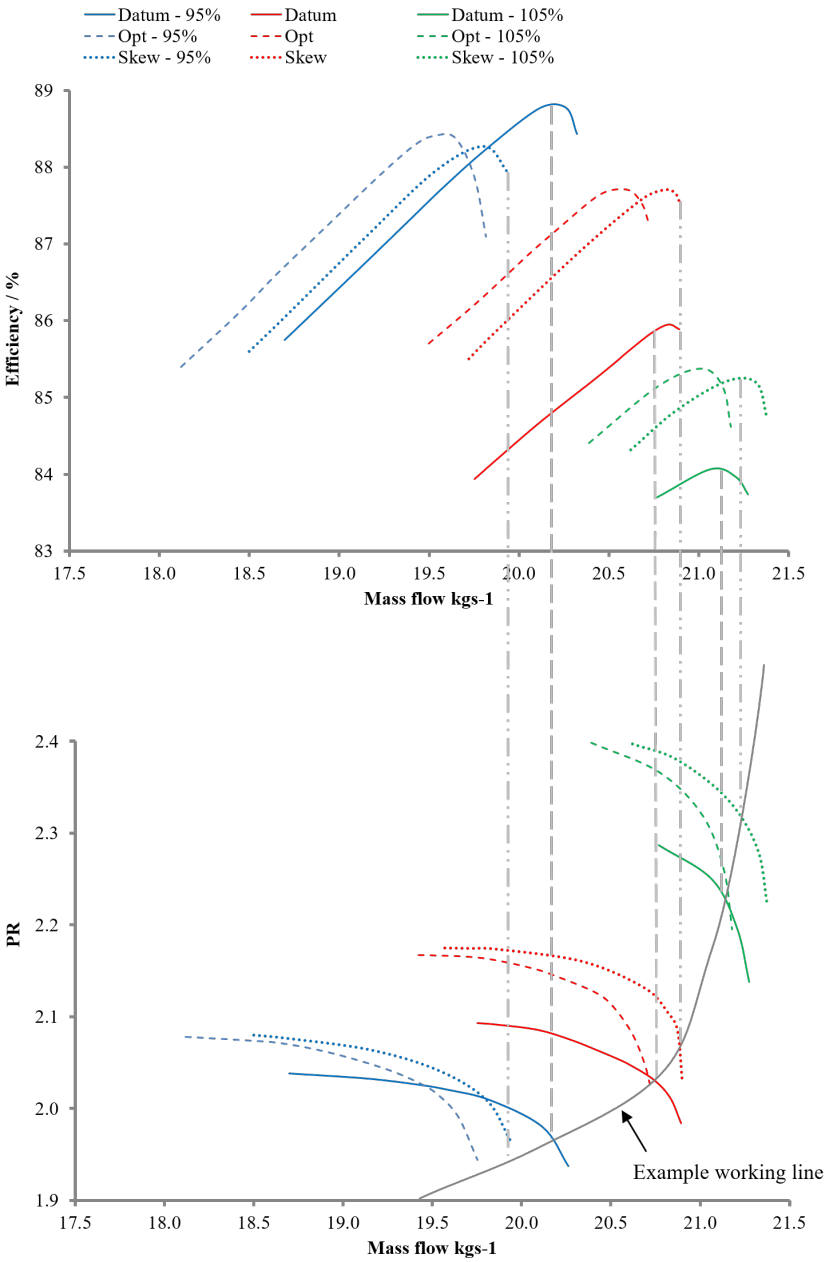

Fig. 3.26 Off-design performance analysis at various rotor speeds.

3.9.4 Practical use of pre-compression

The off-design analysis above shows no evidence of undesirable aerodynamic behaviour due to the use of pre-compression. To further eliminate these concerns, the characteristic for the optimised design was simulated down to 80% speed. Issues such as bi-stable operation / unstart, that plagued early attempts [3], were not present. The use of pre-compression was historically a concern because a shock positioned in a decelerating flow can be unstable. Here pre-compression is only used for a portion of the span, allowing the shock to be held in place by the surroundings. This suggests that parts of an airfoil may, due to 3D relief, employ large amounts of pre-compression without exhibiting undesirable behaviour.

It has been known for some time that reducing the pre-shock Mach number of transonic compressors (via pre-compression) can improve their efficiency [47]. It was clear to transonic compressor designers in the 70s and 80s that shock strength was increased by the amount of convex curvature on the suction side between the leading edge and the shock [34]. Nearly flat suction surfaces that minimised the expansion were therefore favoured, and some attempts were made to try pre-compression designs with concave curvature [34].

However, these early attempts showed undesirable results [1] (such as bi-stable operation), leading the design community to shy away from using pre-compression [51]. The early work aforementioned was conducted when 3D CFD was in its infancy, and left the community wary. It is therefore suggested that it may be time to revisit the previously held "rules of thumb" about pre-compression in the design of transonic rotors. This issue is re-addressed here. It is predicted that using modern simulation, optimisation and a 3D design, large amounts of pre-compression may be employed without the negative effects that plagued early attempts.

3.10 Conclusion

In this chapter, a flexible design parameterisation technique, namely Free-From-Deformation, has been used in the novel shaping of a highly loaded axial compressor blade. Through the efficient MAM optimisation method a novel blade shape was found that showed a significant efficiency benefit. The novel shape includes an ‘S’ in the blade sections in the mid-region of the span. The s-shape protrudes into the passage in the region near the shock, causing a pre-compression effect, delay in the shock impingement point and weakening of the shock. This results in later and smaller shock induced separation, and the entropy increase across the shock on the mid blade line is reduced by almost 80%. This shock-control feature is localised to the region where adjoint sensitivity analysis predicts changes in the blade surface will have the greatest impact on efficiency. The sensitivity of efficiency to changes in the surface was reduced by 30-40% for the optimised blade, giving an indication of the design improvement. Through the FFD parameterisation the blade was also slightly leant in the direction of rotation. This has been identified by previous researchers as a form of 3D

blade shaping with the potential to beneficially modify the 3D shock structure. This 3D shaping may have also contributed to delaying and weakening the shock, and it is a combination of this and the blade profile shaping that provides the benefit seen. The FFD geometry parameterisation technique was able to produce a design predicted to have increased benefit compared to that previously found in the literature. Through an understanding of the blade behaviour a manual skew of the optimised blade was carried out that ensures the design maintains the choke margin of the datum design. The novel design shows impressive performance gains at design and 105% speed, with satisfactory performance at 95% speed compared to the datum. Contrary to early work on the application of pre-compression to transonic rotors, the results given here suggest that these airfoils may employ large amounts of pre-compression (over part of the geometry) without exhibiting undesirable aerodynamic behaviour.

This work demonstrates the benefit that can potentially be achieved using a free-form parameterisation approach. It provides a greater flexibility than other parameterisation approaches and combines together 3D blade shaping with blade profile modification, which has predicted greater benefit than previous optimisation approaches. Delaying and weakening the shock for transonic rotor blades appears to be the area where the most benefit can be achieved. A weakened shock reduces entropy generation directly from the shock, as well as secondary effects, such as shock induced separation.

The uncertainty present in the results due to the difficulty validating the simulations must be considered however. Experimental testing would therefore be ideal to assess whether this kind of design perform as well as predicted when tested.

3.11 Summary

This chapter has demonstrated the benefit to transonic compressor design that can potentially be achieved through the use of a flexible parameterisation approach and shock control. Chapters 4 and 5 also study shock control, investigating alternative approaches and applications. Chapter 4 focuses on assessing the use of shock control bumps to improve engine intake performance.

Chapter 4

Using shock control bumps to improve engine intake performance and operability

Engine intakes must be designed to avoid separation across a wide range of operating conditions, including extreme angles of attack or high cross-winds. At high incidences, the acceleration of flow around the nacelle lip can cause a strong shock wave to form, leading to shock-induced separation. Due to the increased bypass ratio of future engines (as discussed in Chapter 1) the diameter of the engine inlets is increasing and therefore, to compensate for this and reduce weight and drag, engine intakes will become shorter. These shorter intakes are more agressive and have steeper diffuser angles, which exacerbates the problem of shock induced separation at off-design conditions. This chapter investigates the use of shock control bumps to mitigate the effect of nacelle lip shocks at high angles of attack, to see if the addition of shock control bumps can reduce shock induced separation and increase the safe operating range of engine inlets.

4.1 Background

The role of a jet engine intake is to control and diffuse the incoming flow before it reaches the fan. The engine intake/nacelle geometry must be carefully designed to decelerate and diffuse the flow with the minimum total pressure loss. Any inefficiencies in the duct result in successively magnified losses through other engine components. A 1% increase in inlet pressure loss can result in equivalent or greater loss in overall engine thrust [42]. If the flow on the nacelle lip separates at high angles of incidence, the effect on the fan is dramatic. The separated flow causes a large blockage to the fan, and will not only decrease engine performance but if great enough could cause the whole engine to stall.

4.1.1 Shock induced separation on engine inlets

When the freestream flow enters the engine inlet at large angles of attack, the acceleration of this flow around the nacelle lip causes a shockwave to form. This normal shock on the lower nacelle lip causes the boundary layer to thicken, and if the angle of attack is high enough and the shock strong enough, large separation can occur. A description of the phenomenon of shock induced separation was given in Chapter 3. Figure 4.1 shows the key features of intake flow behaviour at high incidence.

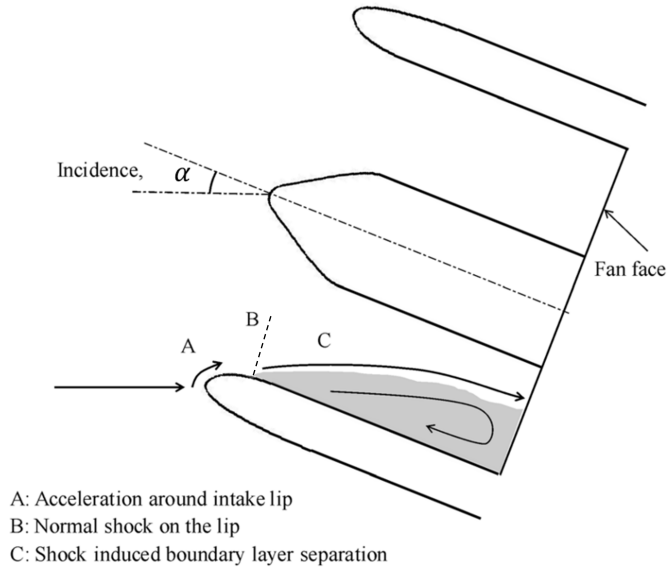

Fig. 4.1 Schematic of the flow physics around an intake operating under a typical off-design, high-incidence condition ([21]).

Previous work has studied engine inlet behaviour at high incidence and in crosswinds [84] and the hysteresis phenomena associated with flow separation and reattachment [49]. The impact of inlet distortion on the stability of the fan [44], and fan–intake interaction on the incidence tolerance [68] [14] has also been studied. These show that the suction of the fan is actually capable of slightly suppressing nacelle separation.

The redesign of short engine intake shapes has been carried out by Peters et al. [86]. This work focused on the re-design of the entire engine inlet geometry using a spline-based inlet design tool to reduce the likelihood of intake separation. A parametric design study of inlet lengths and diameters was carried out and it was concluded that the recommended inlet L/D is between 0.25 and 0.4, else losses in fan efficiency offset gains in propulsive efficiency. Christie et al. [25] used intuitive design variables based on class shape transformation (CST) curves to re-shape nacelle inlets. They were able to achieve similar performance to a mature inlet design in a rapid fashion using intuitive design variables. These investigations re-designed the entire inlet profile to find improved designs.

The work in this chapter focuses on the assessment of shock control bumps as a method of controlling and preventing shock-induced separation on the intake geometry. In this approach, the baseline nacelle geometry is unaltered, apart from the addition of bumps. The investigations are carried out on the isolated intake geometry, as the inclusion of the fan is beyond the scope of this project. It is understood that while including the fan in the simulations would affect the angle of attack at which separation occurs, a re-design that is beneficial in a nacelle only simulation will also be beneficial once the fan is included [113].

4.1.2 Shock control

To combat the adverse effects of shock waves (including shock-induced separation) several methods have been proposed to alleviate shock strength on aerodynamic surfaces. These include synthetic jets [92], slotting of aerofoil surfaces [115], porous surfaces [114] and shock control bumps. The aim of these methods, either active (requiring energy input) or passive (pure geometric modification) is to reduce the strength of the shock, i.e. the magnitude of the pressure/velocity jump across the shock, and therefore reduce the associated adverse effects.

Of those listed above, shock control bumps are one of the most promising methods, and are the focus of this work. They are cost effective, requiring a simple modification to the baseline geometry, and have a low skin friction drag penalty.

4.1.3 Use in the literature

Shock control bumps are bumps added to aerodynamic surfaces to alter the behaviour of the shock and improve aerodynamic performance. One of the early examples of shock control bump usage is in the design of the dromedaryfoil in the 1970s [122]. This was a modified supercritical aerofoil with a bump added in an attempt to increase its drag-divergence Mach number. The 'hump' was shown to weaken the shock wave when implemented in the right position. This also demonstrated the importance of shock control bump positioning, as, if the bump was misplaced, an increase in wave drag was seen.

Ashill et al. [6] found a significant reduction in drag could be achieved for a 2D aerofoil via the correct application of a shock control bump, however when the shock position changed severe drag penalties were incurred due to secondary shocks and separation being produced. Drela and Giles [38] carried out numerical studies into shock control in 1987, describing the behaviour of shock-induced separation. Sommerer et al. [116] optimised shock control bumps at various Mach numbers. They concluded that the bump height, width and position of the bump peak are the key parameters. Collins et al. [26] tested shock control bumps in a wind tunnel, and analysed the performance of shock control bumps at off-design conditions. They noted that at off design conditions even smooth bumps can create vortical flow between the individual bumps, and that large, separated regions occur.

The EUROSHOCK II project [118] began in 1996 and concluded that shock control bumps had the most potential out of a range of shock control devices tested. A large amount of research was carried out into shock control bumps, with both 2D and 3D analysis, although no optimisation was undertaken.

Qin et al. [91] first proposed 3D shock control bumps with a finite width, allowing additional design complexity. They showed that 3D bump configurations were more robust than 2D bump designs.

The off-design performance of shock control bumps is one of the key issues associated with their implementation. If they are not robust to shock position they can easily lead to increased drag or pressure loss away from the point at which they were designed.

4.1.4 Types of shock control bump

Two types of bump are generally favoured in the literature, these are catagorised here as smooth and ramp-type bumps. Both types of bump consist of a ramp or curve which compresses the flow, and then a rear part that returns the flow to the aerofoil surface.

Ramp-type bumps

Ramp bumps form at sharp angles to the baseline surface, and are ramp/wedge like in appearance. Some examples of these are given in Figure 4.2. These use the sudden change in geometry at the upstream edge of the ramp to trigger a lambda-shock structure (see Figure 4.3). The ramp deflects the flow away from the aerofoil and the flow experiences a compression up towards the crest. The strength of the lambda shock structure is much weaker than the datum, normal shock. The angled edges of these discrete bumps are also thought to create vortices between the bumps that energise the boundary layer and may help to reduce shock-induced separation [18].

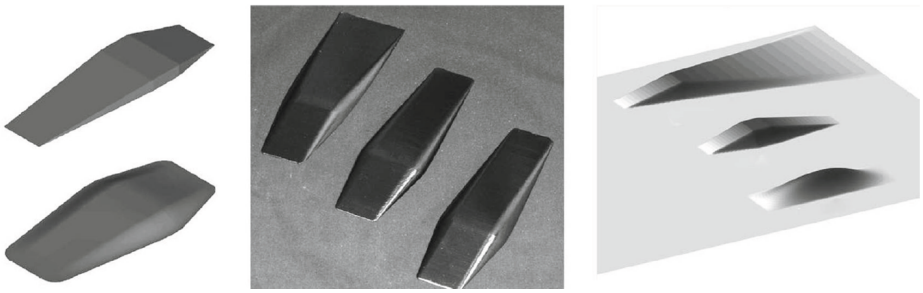

Fig. 4.2 Ramp/wedge type shock bumps from Ogawa et al. [83] (left), Bruce et al. [19] (middle) and Colliss et al. [26] (right). Image from Bruce [18].

.

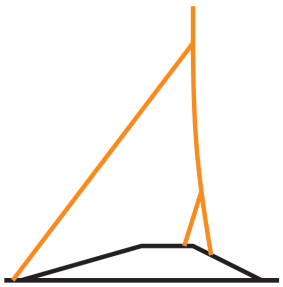

Fig. 4.3 The lambda shock structure created by an optimally positioned wedge bump. Image from Hinchliffe [52].

Smooth bumps

Smooth bumps have a smooth transition between the datum surface and the bump geometry, i.e. where the bump joins the airfoil surface the gradient is zero. These type of bumps have been produced through a variety of parameterisation methods, including piece-wise cubic curves [90], Hicks-Henne curves [123] or Class-Shape function Transforms (CST) [130]. The bumps protrude into the flow in the region where the shock forms, thus causing a pre-compression effect on the local supersonic flow. Rather than splitting the shock into a lambda structure as with the ramp-type bumps, the smooth bumps create isentropic pre-compression waves upstream of the shock that gradually increase the pressure before the shock. Using this approach, it is technically possible to completely eliminate the shock.

The central image in Figure 4.4 shows how the pre-compression effect works when the bump is optimally positioned. The shock has been reduced to a series of pre-compression waves, and a 'knee' like structure is formed [52]. It is also shown how critical the positioning of the shock control bump is. If the bump is placed too far downstream, it can trigger a second shock, and if it is placed too far upstream it can cause separation. The EUROSHOCK II project suggested that the peak of the bump should be approximately 5% of bump length downstream of the original shock location. Qin et al. [91] showed that the effectiveness of the bump primarily depends on the bump position and height. A larger height will cause greater pre-compression but may lead to separation after the bump.

Smooth bumps are used in this investigation as they are likely to result in less pressure loss than the ramp-type bumps. The reduction of total pressure loss produced is a critical concern for intake design.

Types of smooth shock control bumps

Various types of smooth bumps have been explored. The differences between these were studied by Wong [128]. The main types are the 2D style bump which extends along the span (referred to here as continuous) and the individual type, where separate bumps are connected smoothly in the spanwise

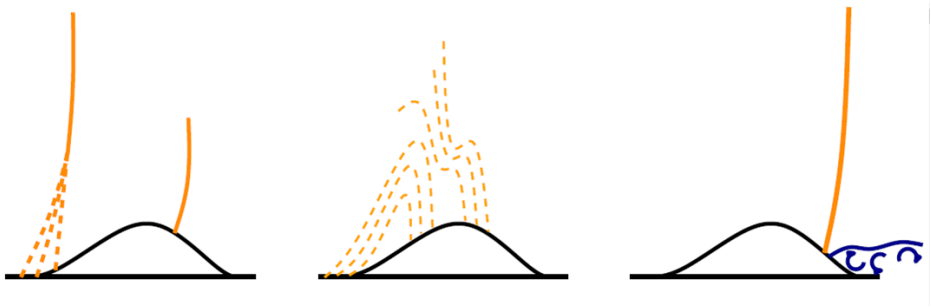

Fig. 4.4 The impact of the smooth bump on the shock when positioned too far downstream (left), optimally (middle), and too far upstream. Image from Hinchliffe [52].

direction. These two types are shown in Figure 4.5. Wong concluded that the 2D (continuous) bumps were not as robust as the individual bumps, and that the individual bumps needed to be larger than the continuous ones to have the same effect.

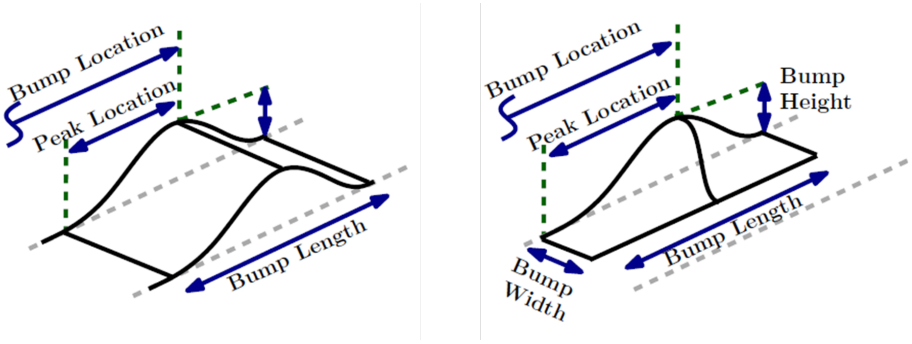

Fig. 4.5 2D type (continuous) bump (left) vs individual type (right). Image from Hinchliffe [52].

Using individual bumps adds complexity into the design study as the bump width and spacing in the spanwise direction need to be considered. In this work both the continuous and individual type bumps are investigated.

4.2 Aims and objectives

The aim of this work is to investigate whether the addition of shock control bumps to the engine intake lower lip has the potential to reduce shock strength and shock-induced separation, and in doing so extend the safe operating range of the inlet.

4.2.1 Objectives

The objectives of this work are therefore to:

- Simulate and understand the flow physics for the baseline case

- Assess whether the manual addition of shock control bumps offers any benefit
- Optimise shock control bump designs
- Assess the optimised designs and any benefit found - do they reduce separation, and can they allow the inlet to safely operate at a higher angle of attack?

4.3 Simulation

4.3.1 Geometry and parametrisation

The geometry under investigation is a Rolls-Royce engine intake design [96]. A CAD representation of the geometry was provided and this was tesselated before being converted into a series of 2D constant-θ slices, which are required for PADRAM [110]. These 2D slices can be individually perturbed to modify the 3D geometry (which is smoothly interpolated from the 2D sections during the meshing process). A python script was developed that can deform each θ section to produce smooth bumps on the nacelle surface. The datum nacelle geometry can be seen in Figure 4.6.

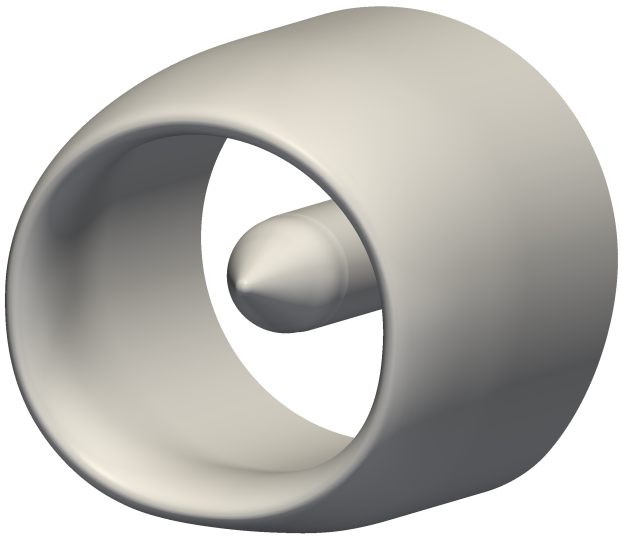

Fig. 4.6 The datum nacelle and spinner geometry generated in PADRAM [110].

The CST method

The Class-Shape Transformation (CST) method [66] is used in this work to define the bump geometries. The CST method uses several Bernstein polynomials to create smooth contour bumps. For this project 3rd order CST bumps are used. The bumps are defined by a start and end point on the datum

surface and the amplitudes for the four polynomials. These amplitudes control the bump height and asymmetry.

The CST method has two parts: the class function and the shape function. The class function determines the order of the geometry and the shape function is then employed to control it. Here, the class parameters N1 and N2 are set to 3. This sets the general class of shapes to follow a similar shape as a cubic-spline. The full description of the 2D CST bump with shape function is [52]:

$$\varepsilon(\phi) = \mathscr{C}^{3.0}_{3.0}(\phi).\sum_{i}^{n} A_i.\mathscr{S}_i(\phi) \tag{4.1}$$

where $\varepsilon(\phi)$ is the relative bump height that is added to the datum geometry, ϕ is the normalised distance between the bump start and end points, C is the class function, n is the b-spline polynomial order (here 3), A is the amplitude of each bump order and S is the shape function.

The class function is as follows:

$$\mathscr{C}^{N_1}_{N_2}(\phi) = \phi^{N_1}.(1-\phi)^{N_2} \tag{4.2}$$

The shape function is as follows:

$$\mathscr{S}_i(\phi) = \sum_{r=0}^{n} K_{r,n}\phi^r(1-\phi)^{n-r} \tag{4.3}$$

$$K_{r,n} = \frac{n!}{i!(n-i)!} \tag{4.4}$$

Figure 4.7 shows how the bump geometry (solid black line) is a sum of the four polynomials (coloured dashed lines). During geometry parameterisation the bumps are defined by a start and end point on the datum surface, and the amplitudes of the four polynomials which control the overall bump amplitude and asymmetry. The CST bump parameterisation provides a high degree of flexibility, enabling the generation of smooth, asymmetric bumps in 2D and 3D.

Continuous vs individual bumps

As mentioned above, two types of bumps are used in this work; bumps that are smoothly interpolated in the θ direction without returning to the baseline geometry (continuous) and individual bumps, where the surface retains the datum geometry in between the bumps in the θ direction.

For the continuous bumps, the 2D CST bump profile is specified at several positions in the θ direction and this is then smoothly interpolated using a cubic b-spline onto every θ section. For the individual bumps, the bumps are defined in a similar way, but then varied in the θ direction using an additional CST function. A demonstration of the difference between these bumps is given in Figure 4.8.

Some examples of these bumps on a nacelle geometry are shown later in Figures 4.25 and 4.28.

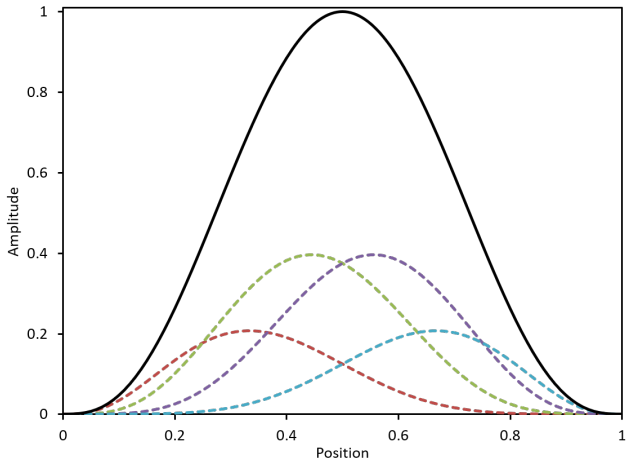

Fig. 4.7 Example 2D CST bump (solid line) and the four polynomials used to construct it (dashed lines).

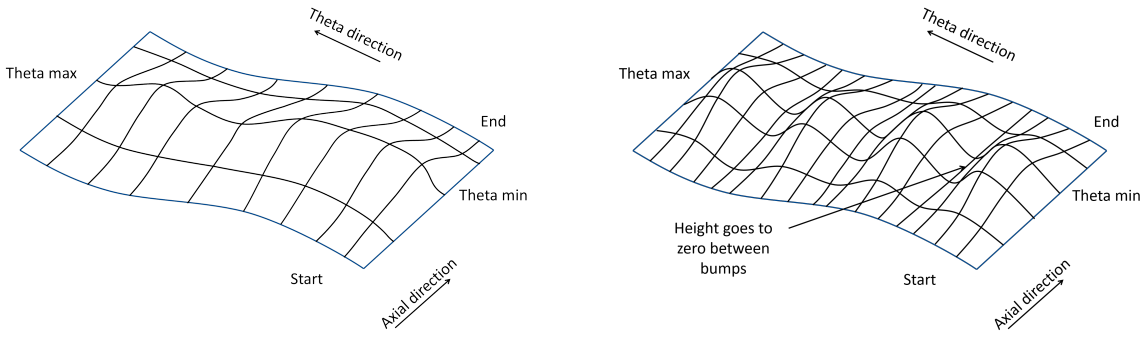

(a) (b)

Fig. 4.8 (a) Continuous vs (b) individual bumps with varying start/end points and bump parameters.

4.3.2 Simulation set up

Meshing

Meshing of the nacelle geometry is carried out in PADRAM [110], using the sub-module RApid Meshing for INtakes (RAMIN). A symmetry plane is used to reduce computational cost. A mesh dependence study was undertaken to select the appropriate mesh size to be used. The overall mesh size was gradually increased until the pressure loss measured at the fan-face changed by less than 1%. It was found however during this investigation that the region most sensitive to mesh density was around the shock location, on the upper side of the lower nacelle lip. Even with a globally mesh independent grid, if too few axial cells are used in the shock region it is not properly captured, resulting in a false prediction of the physics. The mesh in this region was therefore refined to ensure proper resolution of the shock. The effect of this can be seen in Figure 4.9, where the difference in shock resolution is shown on the symmetry plane using static pressure contours. The y^+ is of the order of 1 for the whole nacelle.

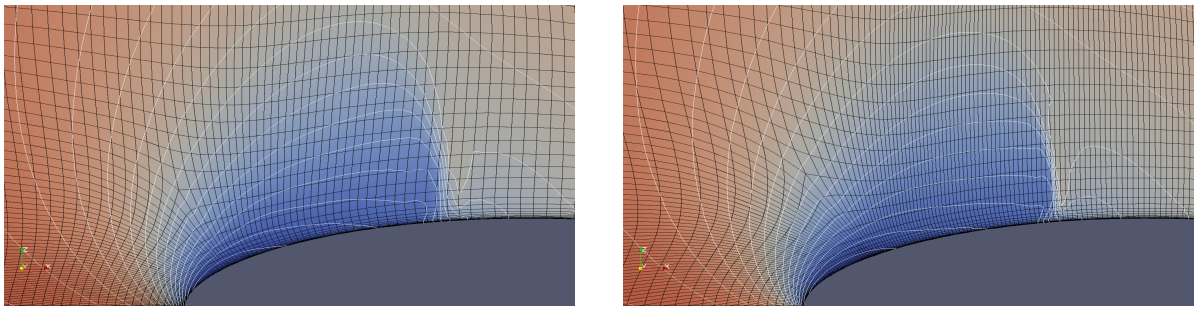

(a) (b)

Fig. 4.9 (a) Too coarse a grid not resolving the shock, (b) Refined grid better resolving the shock.

The mesh in the θ direction comprises of 130 cells, matching the resolution of the geometry file. The final mesh used can be seen in Figure 4.10.

The CFD domain

A schematic of the CFD domain can be seen in Figure 4.11, with more detail in Figure 4.12. The freestream is defined by a Mach number of 0.25, pressure of 53,592Pa, density of 0.73138kgm^{-3} and the angle of attack (α) under investigation. The outflow (engine inlet) boundary condition is set by a non-dimensionised mass flow (capacity ($\frac{\dot{m}\sqrt{T_{tot}}}{P_{tot}}$)). This is used to provide a fair comparison between different angles of attack or designs. If a fixed mass flow were used, this may force extra flow through the intake than would occur in reality, where the fan exerts a constant suction on the incoming flow and the mass flow achieved depends on the pressure loss over the intake lip. The non-dimensionalised mass flow value was iterated until 277.4kgs^{-1} was achieved at the engine inlet at $\alpha = 27^o$. This is known to be close to the separation boundary of the lip from previous work. The nacelle walls have

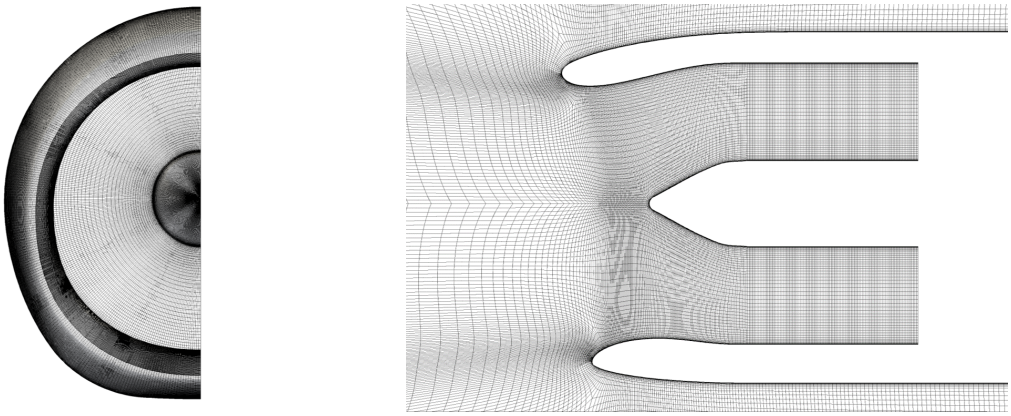

(a) (b)

Fig. 4.10 PADRAM mesh of the intake geometry (a) front view, (b) side view on the symmetry plane.

no slip boundary conditions applied. Because the nacelle geometry is symmetrical, a symmetry plane is used and only half of the geometry simulated to reduce computational cost.

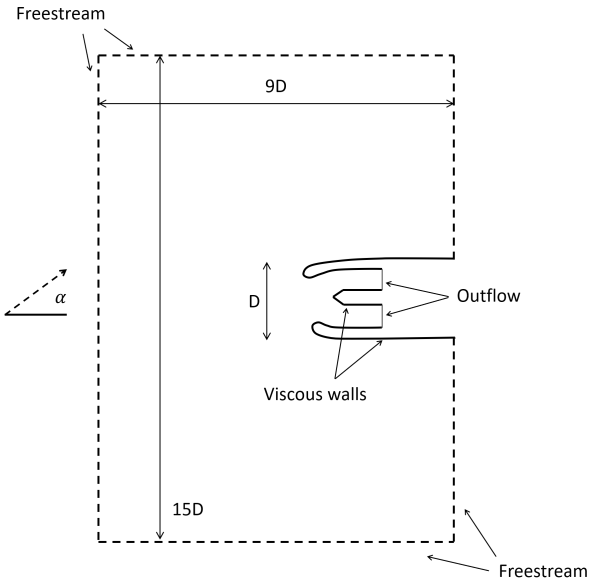

Fig. 4.11 Simulation domain and boundary conditions.

4.4 Analysis of baseline geometry

To analyse the flow features of the datum geometry, the nacelle was simulated at a range of angles of attack, from zero to 33 degrees. To understand the impact of the flow through the intake on engine performance, it is necessary to assess its impact on the fan. To do this, measurements are taken at the fan face position. Figure 4.12 shows the location of the fan face in the simulation domain, as well as

details of the geometry set up. The fan face is just a measurement plane and does not interact with the flow.

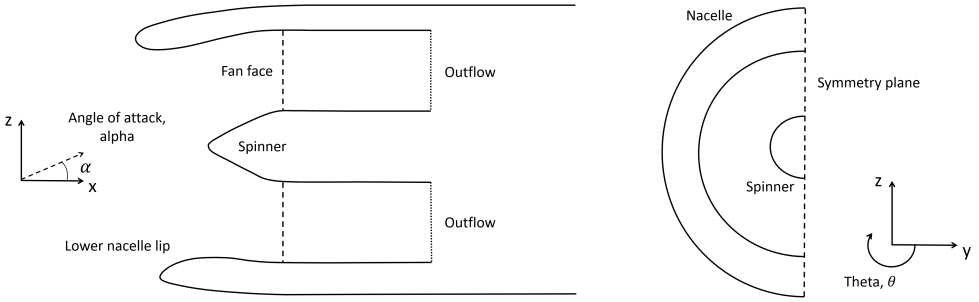

Fig. 4.12 Fan face position and symmetry plane.

4.4.1 Variation of flow features with angle of attack

Figures 4.13 and 4.14 show the flow features of the nacelle as the angle of attack is varied from $\alpha = 0^o$ to $\alpha = 33^o$. The spinner is not shown in the 3D images for clarity. It can be seen how the flow accelerates around the nacelle lip and forms a low pressure region. Static pressure contours show the shock formation and how it strengthens as the angle is increased. The shock-induced separation is highlighted by 3D isosurfaces of zero axial velocity (in yellow). These isosurfaces also highlight reverse axial velocity as the flow hits the underside of the lower nacelle lip and travels around it, however it is only the isosurfaces 'inside' of the nacelle that indicate separation and are of concern. The impact of this separation on the fan is shown via the contour of entropy at the fan face (dark regions indicate high entropy). The location of the fan face in the simulation domain is shown in Figure 4.12. An increase in entropy at the fan face indicates an increase in pressure loss and blockage, which would adversely affect the fan's operating performance. A small amount of separation begins to appear at 29 degrees, and gradually increases in size as the shock strength increases. At 33 degrees, the flow can no longer remain attached and a sudden increase in the amount of separation is seen.

4.4.2 Variation in entropy and mass flow with alpha

Figure 4.15 shows the variation in pressure loss and mass flow at the fan face for various angles of attack. These can be used to indicate how favourable the flow entering the fan is. Greater pressure loss is caused by boundary layer thickening, flow separation or increased shock strength. Due to the non-dimensionised mass flow boundary condition imposed at the engine inlet, increased pressure loss will result in a lower mass flow. A steady increase in pressure loss (and reduction in mass flow passing through the fan) can be seen as the angle is increased, due to greater shock strength and separation, up until $\alpha = 32^o$. At $\alpha = 33^o$ the intake becomes separated [1], and a sudden rise in pressure loss and drop

[1] In this work, when the intake is referred to as 'separated', it should be interpreted that the boundary layer is completely separated, with a large recirculation region present (such as $\alpha = 33^o$ in Figure 4.14). Therefore nacelles that have a small

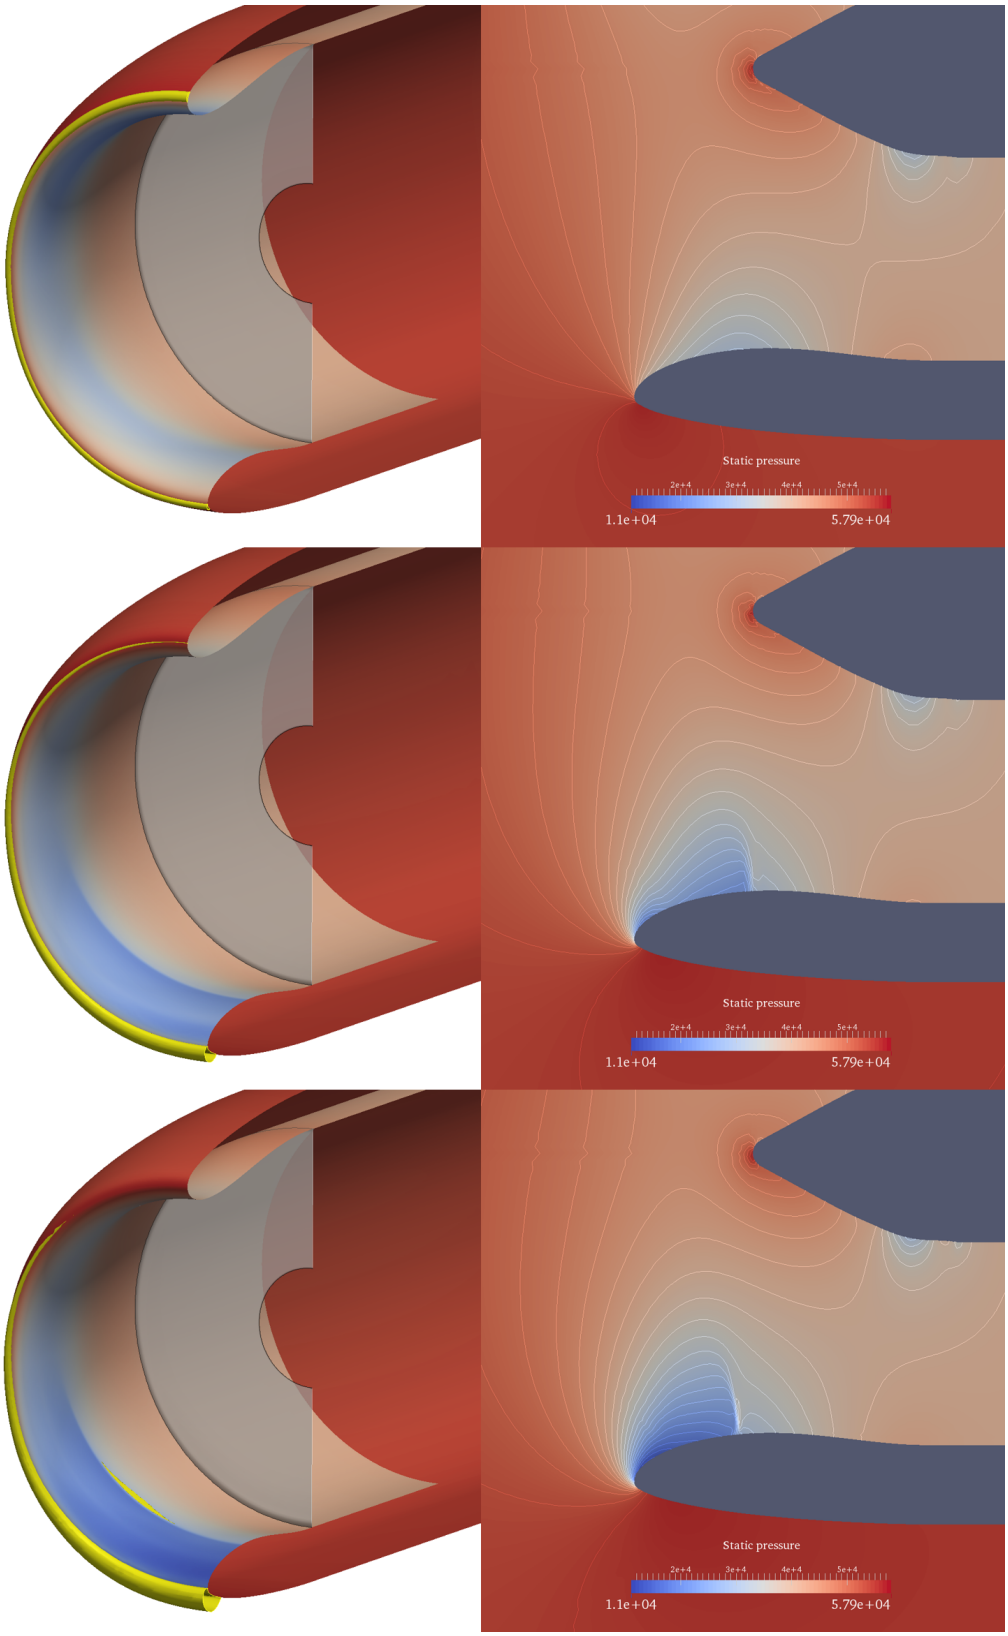

Fig. 4.13 Flow behaviour and separation (yellow) at $\alpha = 0^o$ (top), 20^o (middle) and 29^o (bottom).

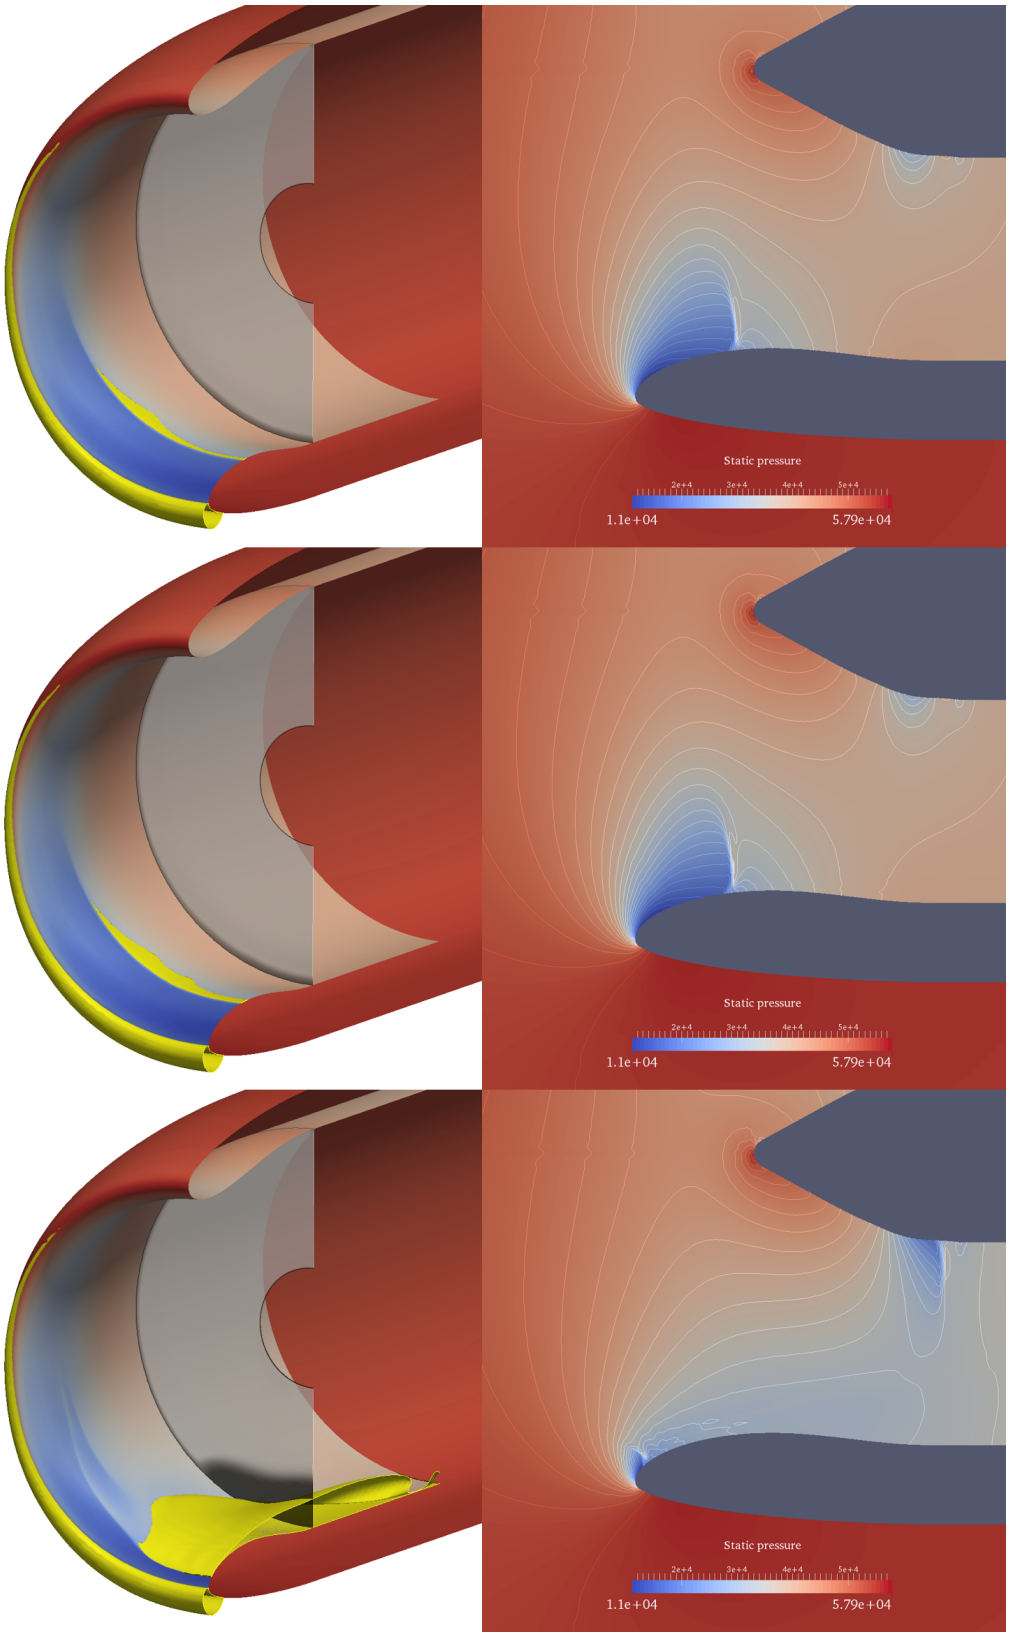

Fig. 4.14 Flow behaviour and separation (yellow) at $\alpha = 31^o$ (top), 32^o (middle) and 33^o (bottom).

in mass flow is seen. This is the critical point that must be avoided during engine operation, as the large disturbance to the fan would severely decrease performance and could even stall the fan. Hence the safe operating range for this intake is $\alpha = 32^o$ and less.

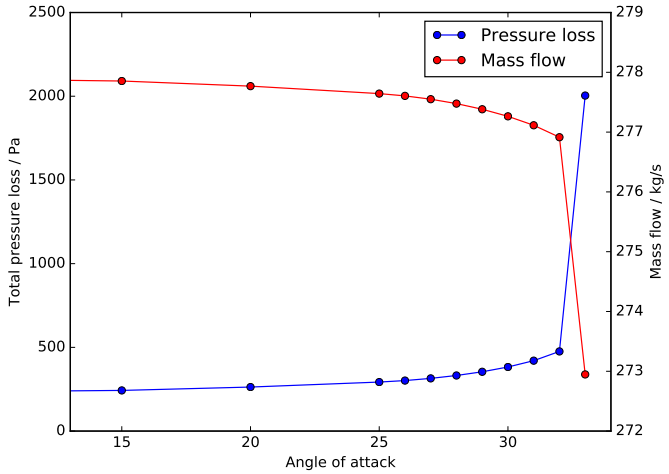

Fig. 4.15 Variation in area-averaged total pressure loss and mass flow at the fan face with angle of attack.

4.4.3 Shock-induced separation

The phenomenon of shock-induced-separation was described in Chapter 3. As can be seen in Figure 4.14, shock induced separation occurs from $\alpha = 29^o$ (indicated by the yellow regions of reverse flow). The process of this shock induced separation can be seen in Figure 4.16. This shows the Mach number along a slice at $\theta = 10^o$ for $\alpha = 31^o$. The process of the shock terminating at the nacelle surface, interacting with the boundary layer and triggering separation can be seen. The sudden increase in boundary layer thickness (and separation) can be seen by the large low velocity region after the shock impingement point.

4.4.4 Shock variation with incidence

Figure 4.17 shows the shock shape at $\theta = 0^o$ for various angles of attack. It can be seen that the difference between the up and downstream pressures of the shock becomes greater as α is increased. The upstream pressure reduces with angle of attack, as the flow accelerates more around the nacelle lip. It can be seen how the shock position shifts as α is increased. For clarity the profile at $\alpha = 29^o$ is shown in Figure 4.17a.

separation bubble, e.g. $\alpha = 32^o$ in Figure 4.14 are not classed as 'separated', despite some boundary layer separation occuring

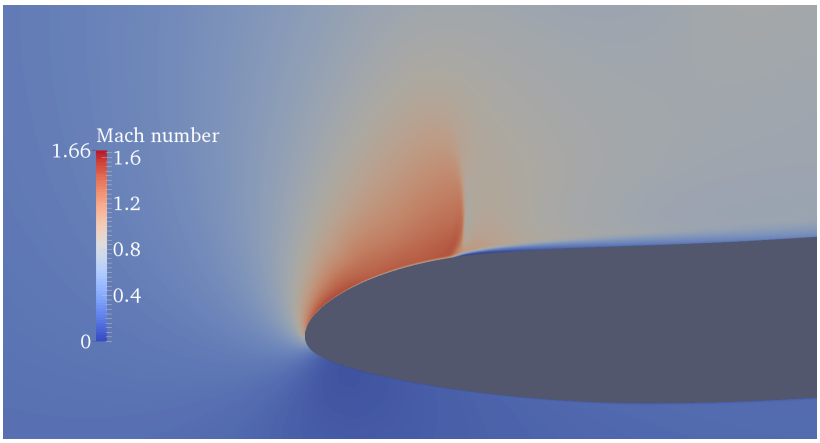

Fig. 4.16 Mach number contour at $\alpha = 31^o$, $\theta = 10^o$.

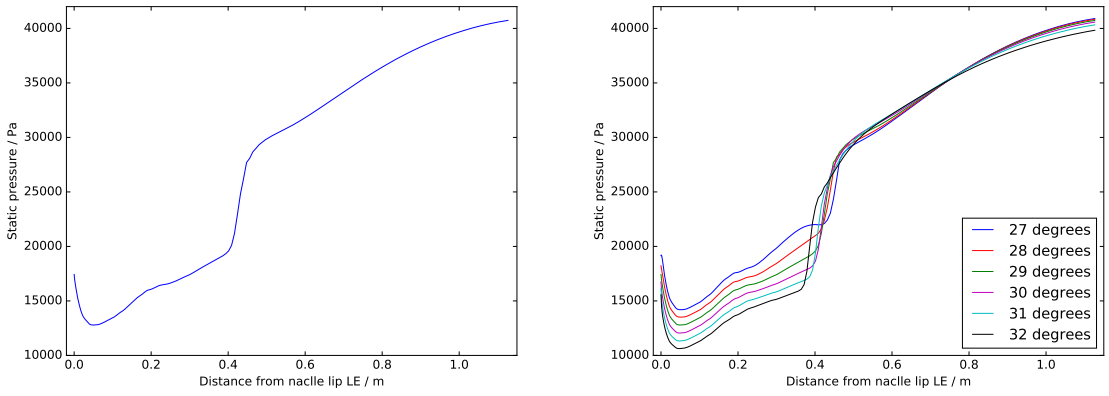

(a) (b)

Fig. 4.17 Pressure profiles along the lower nacelle lip at bottom-dead-centre for (a) $\alpha = 29^o$, (b) various values of α.

To further understand the shock, its behaviour away from the wall can be analysed. This gives a clearer picture of the shock behaviour. Figure 4.18 shows how the shock at $\theta = 0^o$ varies with α. This is plotted using Mach number, measured at 0.1m from the wall (along the white line shown in Figure 4.18b). It can be seen how the shock forms and its strength increases steadily with increasing angle of attack. The pre-shock Mach number increases with α, as greater acceleration of the flow around the nacelle lip increases its velocity. The post-shock Mach number is fairly similar for all α, hence the jump in velocity and pressure across the shock increases with incidence, increasing the adverse pressure gradient, which is the reason for the increased separation seen at higher angles of attack.

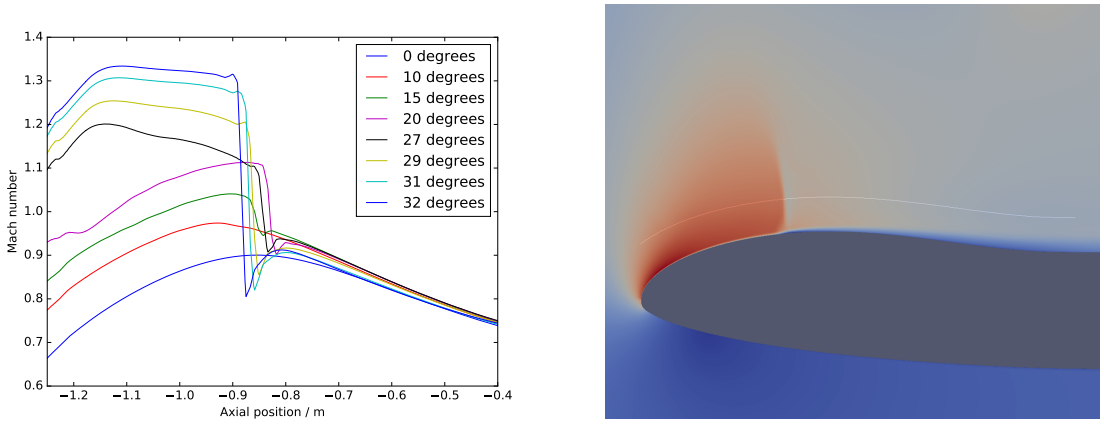

(a) (b)

Fig. 4.18 (a) Mach number profiles at bottom dead centre for various α along the constant wall-distance line shown in (b).

4.4.5 Flow variation in the theta direction

As can be seen in Figure 4.14, the amount of shock induced separation at high angles of attack varies in the θ direction. Separated flow is also isolated to between $\theta = 0^o$ and $\theta = 45^o$. This demonstrates that there is a variation in the flow behaviour occurring in the θ direction. Here, further analysis of the flow is carried out at $\alpha = 31^o$. Figure 4.19 shows the shock profiles at various θ positions from 0 to 50 degrees and it can be seen that the pre-shock flow varies with θ. The shock strength is difficult to judge from this figure however, so an alternative approach is needed.

The entropy increase from before to after the shock (along a flow streamline) gives a direct measure of the shock strength. For these results, the entropy jump across the shock at 0.1m distance from the wall was measured for each θ angle. The variation in shock strength with θ can be seen in Figure 4.19b. This shows that the strongest shock is around $\theta = 30^o$-40^o, which reflects the regions of highest separation seen in Figure 4.14.

As suggested by the above results, the pressure loss, and therefore mass flow, at the fan face is not evenly distributed in θ. To understand the θ distribution of these flow features, measurements can be

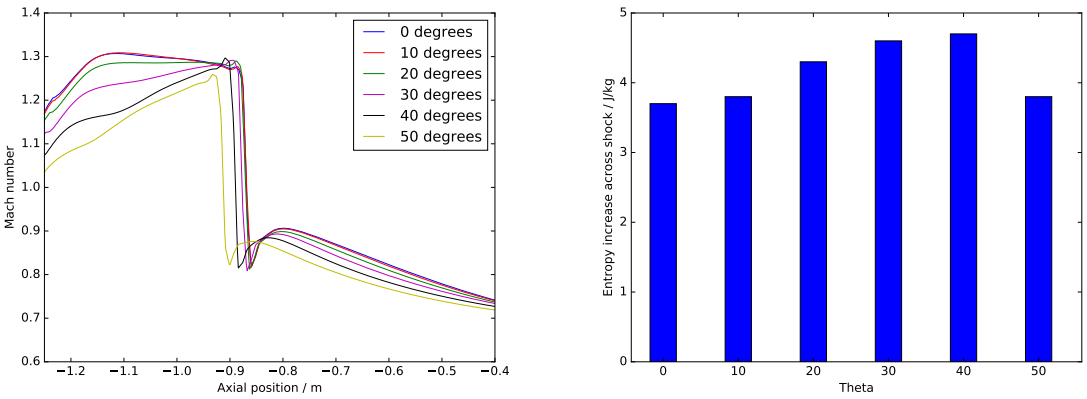

(a) (b)

Fig. 4.19 Shock variation with theta at $\alpha = 31^o$ (a) Mach number profiles, (b) shock strength measured by entropy increase.

taken at the fan face. Figure 4.20 shows the 2 degree sectors that are used to measure these variations. The flow properties are integrated (using area averaging) over each sector to give the θ distribution.

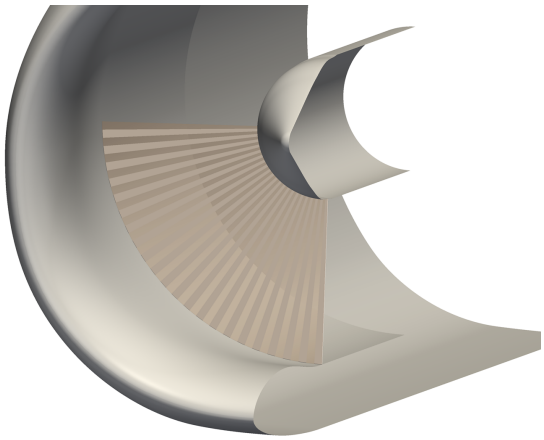

Fig. 4.20 The 2^o sectors used to calculate flow variations in the theta direction.

Figure 4.21 shows the variation in mass flow with θ at the fan face. The delta between zero and 31 degrees incidence (subtraction of the mass flow values at zero degrees from values at 31 degrees) shows how the mass flow is reduced by the increased angle of attack and the associated losses. It can also be seen that a slight increase in mass flow occurs above 60 degrees, as the flow is re-distributed due to the blockage near to bottom dead centre. The drop in mass flow, seen for the majority of the nacelle, is due to the increased pressure loss caused by shock and separation at increased incidence. The total pressure loss distribution (shown in Figure 4.22) can therefore be seen to reflect the mass flow distribution, as increased loss results in lower mass flow. The total pressure loss was calculated

by area-averaging the total pressure over each 2 degree sector and subtracting this value from the total pressure in the far field.

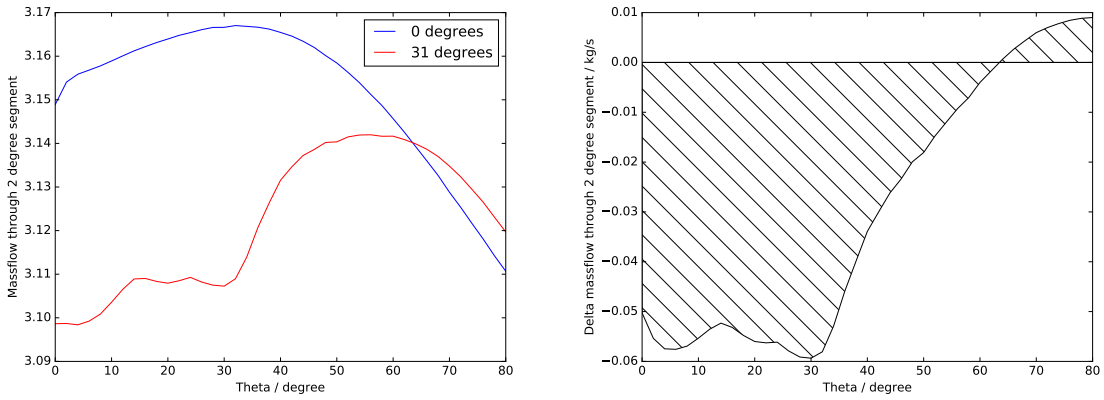

(a) (b)

Fig. 4.21 (a) Mass flow vs theta at $\alpha = 0^o$ and 31^o, (b) delta mass flow vs theta from $\alpha = 0^o$ to 31^o.

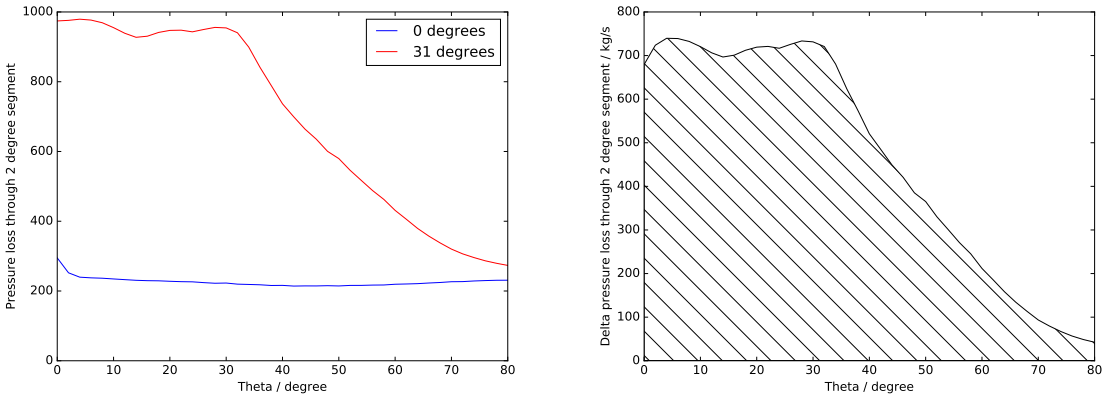

(a) (b)

Fig. 4.22 (a) Total pressure loss vs theta at $\alpha = 0^o$ and 31^o, (b) delta total pressure loss vs theta from $\alpha = 0^o$ to 31^o.

4.4.6 Explanation of the most sensitive region in theta

As shown above in Section 4.4.5, the shock strength and separation are greatest around $\theta = 30^o$, and not at $\theta = 0^o$, where one might expect the strongest shock to be, due to the greatest vertical component of the velocity being at this point. The reason for this is due to the geometrical features of this particular nacelle design [113].

As one traverses in the θ direction around the nacelle, the local incidence normal to the lip due to the global angle of attack reduces as θ increases. This local incidence is maximum at $\theta = 0^o$ and zero at $\theta = 90^o$. This suggests that the acceleration around the lip, and shock strength, will be strongest at lower values of θ. The nacelle geometry also varies in θ, however. Figure 4.23 shows the front view of the geometry and a diagram that explains the contraction ratio.

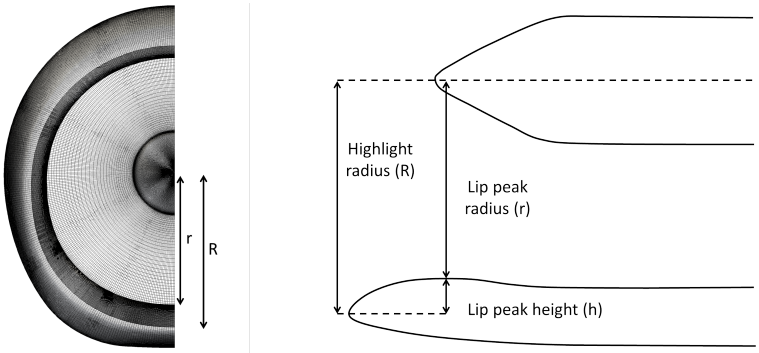

Fig. 4.23 Front view of the nacelle geometry (left) and a schematic defining geometric parameters of the nacelle (right).

The contraction ratio (CR) is defined as $CR = \frac{R^2}{r^2}$, where R is the highlight radius and r is the lip peak radius, as shown in Figure 4.23. For nacelle designs, the contraction ratio reduces as θ increases, as a higher contraction ratio (thicker lip) is more favourable near to bottom dead centre, but a lower contraction ratio (thinner lip) is preferred near to $\theta = 90^o$ due to crosswind considerations. For this particular nacelle, constraints on the nacelle internal volume mean that it is 'flat-bottomed', so the component of α remains quite large until higher values of θ than if the nacelle geometry was circular. Due to this shaping, the contraction ratio becomes lower at first as θ increases from bottom dead centre.

The acceleration around the nacelle lip and shock strength are related to the local incidence and the contraction ratio. For this geometry, there is a point in the θ direction where the combination of local incidence and the contraction ratio results in a stronger acceleration than at bottom dead centre, which pushes the strongest shock and greatest separation further around the nacelle away from bottom dead centre, as described above.

4.5 Application of shock control bumps

The objective of this study is to investigate what benefit can be achieved through the use of shock control bumps on engine inlets. Shock control bumps have the potential to weaken the nacelle lip shock, reduce separation and increase the safe operating range of the nacelle. If it were shown that shock bumps provided a higher limiting angle of attack, the benefit would probably be realised through the re-design of inlet geometries utilising bumps. This could allow thinner lip designs (reducing

weight and drag), or shorter intake lengths to be used while maintaining the same operating range as the standard design.

The desired benefit from the application of shock control bumps is an increase in the fully separated angle of attack. This is a difficult objective to design for however, as simulations at this limiting angle of attack do not converge for the baseline case, and so it would be difficult to carry out a systematic study of different designs, as the simulations of most designs will not converge. To demonstrate the benefit from the use of bumps, they will therefore be assessed at a reduced angle of attack where the simulation convergence is improved. If the bumps are able to provide benefit at this condition, they will potentially have a positive impact as incidence is further increased.

The benefit due to the application of shock control bumps at a non-separated incidence is expected to be seen through:

- A reduction in shock strength
- A reduction in pressure loss at the fan face
- An increase in the mass flow through the fan face

To assess the impact of adding shock control bumps to the nacelle surface, it was decided that their effect would be tested at $\alpha = 31^o$. This operating condition is close to the fully separated point, but being unseparated means the simulation convergence is still satisfactory. The approach of this study is therefore to try and reduce the separation and pressure loss at a single point, and to later assess whether improving the flow behaviour at the single point near to full separation allows the nacelle to safely operate at higher values of incidence. To assess any improvement through the use of bumps, the mass flow and pressure loss can be measured at the fan face. An increase in mass flow, or reduction in pressure loss, would be indicative of a more favourable design. Reducing the shock strength, separation and pressure loss produced at this point would potentially allow the redesigned geometry to successfully operate at greater angles of attack.

4.5.1 Parametric study

To see in the first instance whether shock control bumps provide any benefit, some example bumps are here placed on the nacelle lip. Both continuous and individual bumps are assessed. Previous work on wings [130] has suggested that individual bumps are more robust, so these were assessed first. The initial placement of the bumps was set so that 60% of the bump was downstream of the datum shock position. This follows the recommendation by Hinchliffe [52]. The bumps were positioned from 0 to 60^o θ, as this is the region where separation mostly occurs and the region where increased angle of attack is most detrimental to the flow (see Figure 4.22b). The region of bump placement for these studies is highlighted in Figures 4.24 and 4.27. For these investigations, the start and end positions of the bumps remains fixed and the four parameters controlling the bump height are kept constant

and modified as one variable, hence just one variable (the overall symmetric bump height) is studied. The CFD simulations are initialised from a converged solution for the datum geometry in order to accelerate the convergence.

Individual bumps

The individual bumps are controlled in the θ direction by the CST method (see Section 4.3.1). Figure 4.24 shows an example of the 5 individual bumps added to the datum surface, positioned to have 60% of bump length downstream of the datum shock position. Figure 4.25 shows three of the bump geometries tested. Only the overall bump amplitude varies between the designs. Due to the nature of the CST function, the amplitude values are dimensionless, but are approximately equivalent to millimetres. Many bump amplitudes were tested to find a suitable range. Here results for bump amplitudes of 10, 12.5 and 15 are given. All of the designs were simulated at $\alpha = 31^o$. Table 4.1 shows the performance of these designs compared to the datum geometry, and the delta values to the datum geometry at $\alpha = 0^o$ given for reference.

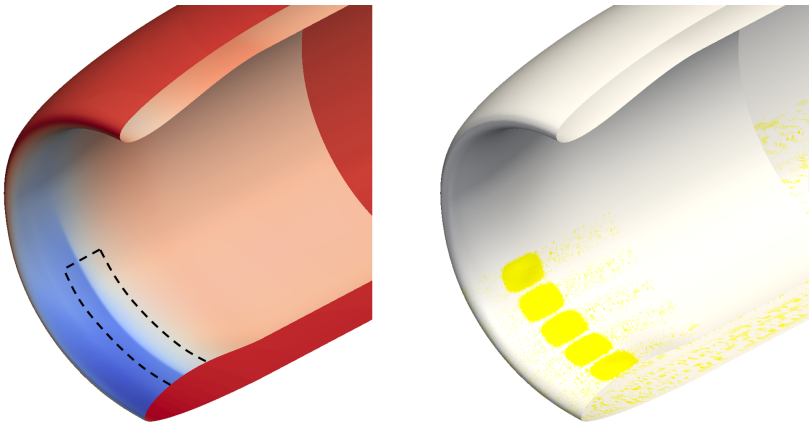

Fig. 4.24 The bump placement region (left) and an example set of individual bumps (right).

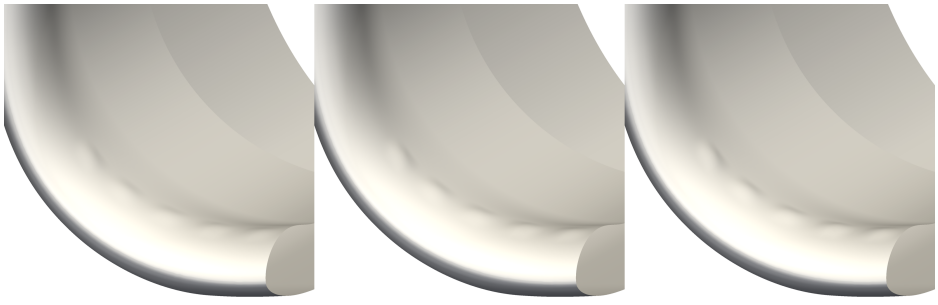

Fig. 4.25 Various amplitudes of individual bumps, with amplitudes of 10 (left), 12.5 (middle) and 15 (right).

Table 4.1 Comparison of the performance of the individual bumps tested.

	Mass flow / kg/s	**Delta mass flow**	**Benefit / %**	**Total pressure loss**	**Delta total pressure loss**	**Benefit / %**
Datum (@ α=0^o)	277.87	0.0	-	239.05	0.0	-
Datum (@ α=31^o)	277.11	-0.76	-	420.48	181.43	-
Bump 10	277.07	-0.8	-5.26	446.36	207.30	-14.26
Bump 12.5	277.13	-0.74	3.95	429.13	190.08	-4.72
Bump 15	277.14	-0.73	2.63	428.69	189.63	-4.52

The results show that the two largest bumps were able to slightly increase the mass flow compared to the datum design (at α = 31^o), but that the pressure loss was increased compared to the datum. The most promising design was the largest amplitude bump tested 'bump 15'. The impact on the flow for this geometry can be seen in Figure 4.26.

It can be seen that the use of the bumps has slightly delayed the shock and reduced the separation for some regions of θ, but around θ = 40^o (the most separated point for the datum design), the separation is increased. This shows that the bumps have potential to provide benefit, but that their geometry needs to be carefully designed to avoid increasing separation in this key region (around θ = 30^o-40^o).

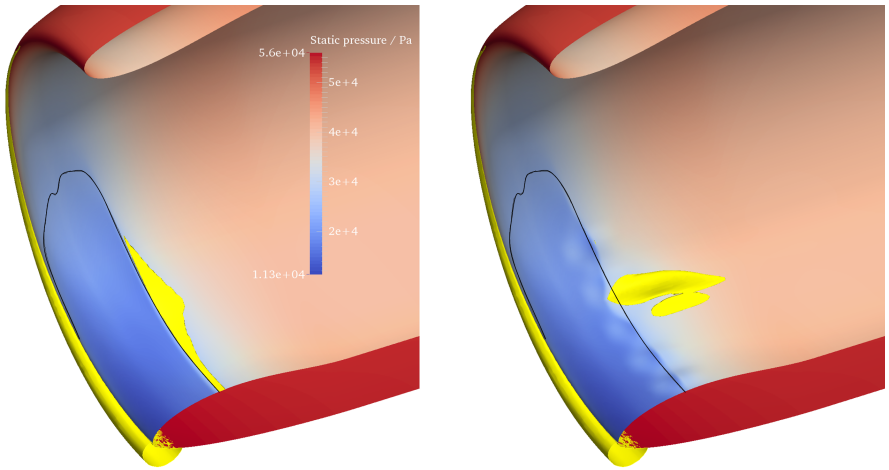

Fig. 4.26 The separation (yellow) and shock position for the datum geometry (left) and for individual bumps with amplitude 15 (right). The datum shock position is indicated by the black line.

Continuous bumps

To compare the impact of continuous to individual bumps, continuous bumps were generated based on the amplitudes used for the individual geometries. Figure 4.27 shows the region of bump application for the continuous designs. The two geometries studied are shown in Figure 4.28.

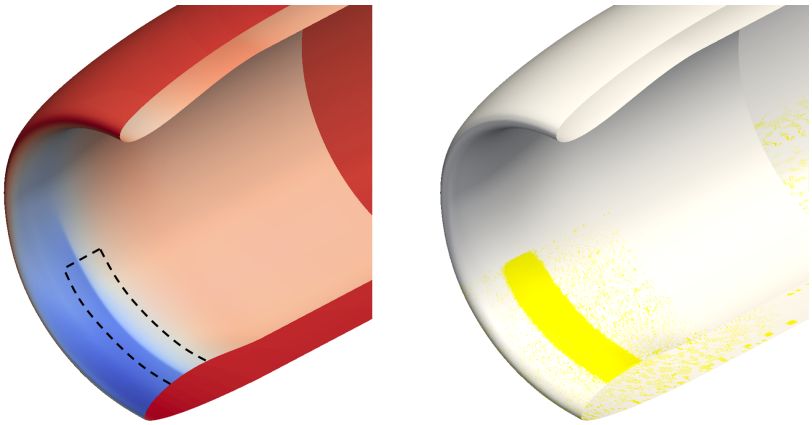

Fig. 4.27 The bump placement region (left) and an example continuous bump (right).

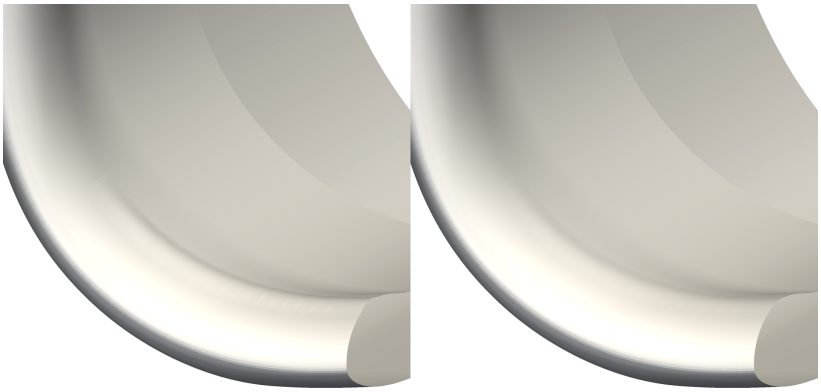

Fig. 4.28 Different amplitudes of continuous bump, with amplitude 15 (left) and 7.5 (right).

Table 4.2 shows the flow properties of the tested geometries. It can be seen that a continuous bump with magnitude 15 causes a large reduction in mass flow and increase in pressure loss. A magnitude of 7.5 however shows an increase in mass flow as well as a reduction in pressure loss compared to the datum design, thus making it the most promising of the designs tested.

The shock position and separation features for this continuous bump compared to the datum geometry can be seen in Figure 4.29. It can be seen how the shock position is delayed for part of the nacelle, and how this is smoother than the effect seen for the individual bumps. However, it can again be seen that the separation has been increased at the most sensitive region (around $\theta = 30^o$ - 40^o).

Table 4.2 Comparison of the performance of the continuous bumps tested.

	Mass flow / kg/s	**Delta mass flow**	**Benefit / %**	**Total pressure loss**	**Delta total pressure loss**	**Benefit / %**
Datum (@ α=0^o)	277.87	0.0	-	239.05	0.0	-
Datum (@ α=31^o)	277.11	-0.76	-	420.48	181.43	-
Cont. bump 15	273.38	-4.49	-490.79	1604.21	1365.15	-652.44
Cont. bump 7.5	277.20	-0.67	11.84	407.67	168.62	7.06

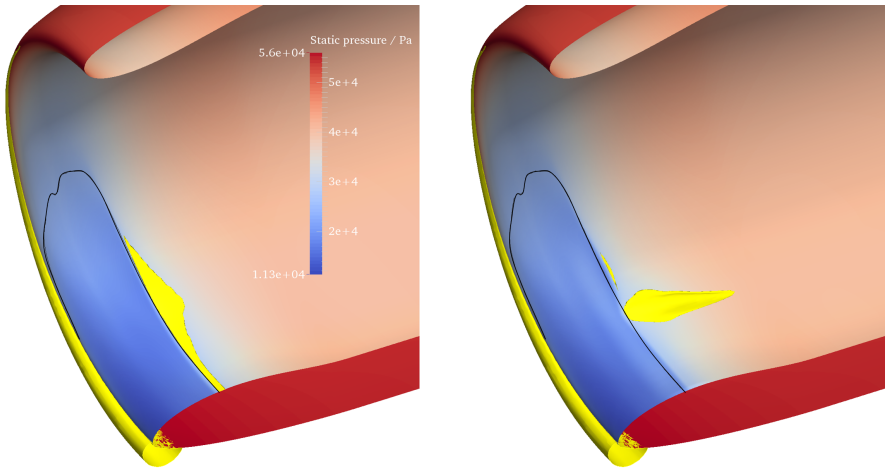

Fig. 4.29 The separation (yellow) and shock position for the datum geometry (left) and for a continuous bump with amplitude 7.5 (right). The datum shock position is indicated by the black line.

Figure 4.30 shows the mass flow histories with iteration for the best individual and continuous bumps tested. It can be seen how both these bump designs provide increased mass flow compared to the datum design, especially the continuous bump. However, the effect of the increased separation on the convergence of the solution is clear. The separated flow for the continuous bump behaves in an unsteady manner, and this is reflected in the mass flow variation seen in our steady simulation. Therefore an improved nacelle bump design is needed to avoid this unsteady separation.

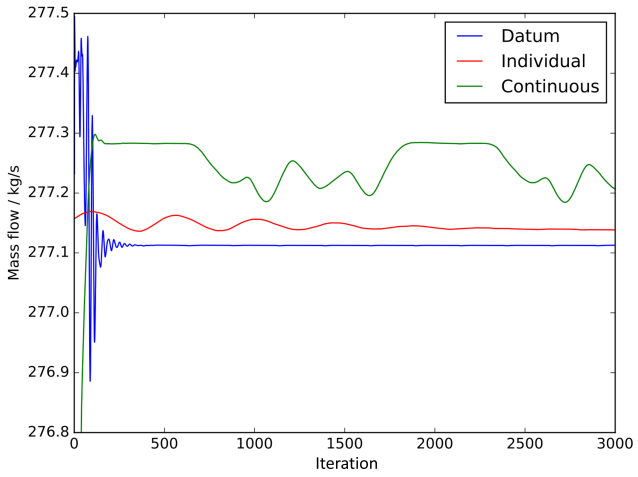

Fig. 4.30 Simulation mass flow histories of the datum, individual and continuous bump designs at $\alpha = 31^o$.

4.6 Optimisation of shock control bumps

The results of the individual and continuous bump parametric studies show that some benefit is achievable through the use of shock control bumps. The geometries tested only had simple, symmetrical designs, and the bump peak position and CST amplitudes were not modified however. To gain increased benefit from the shock bumps, their designs need to be optimised to find the most beneficial bump crest position and bump amplitudes. In this section the bump designs are optimised to assess the increase in benefit that can be achieved. The optimisation method used for this work is the MAM method, as described in Chapter 2.

4.6.1 Optimisation of individual bumps

For the optimisation of individual bumps, five bumps were positioned as shown above in Figure 4.24. The bump start and end positions are fixed and the four CST parameters are allowed to vary for each of the five bumps, thus resulting in 20 parameters. The objective for the optimisation is to maximise the mass flow through the fan face, but as the optimiser can only minimise functions, the

objective function is set to 280kg/s - the mass flow. Some of the designs perform worse than the datum geometry, as poorly designed bumps can cause separation and generate increased pressure loss, especially at the high incidence operating condition of 31^o at which the optimisation is carried out. When this separation occurs, the convergence of the simulation becomes poor, hence, when calculating the objective function, the last few hundred iterations of each simulation were averaged. The reason for this, as well as the large number of simulation iterations used, is clear from Figure 4.31, where the mass flow convergences for a few designs in the optimisation are shown. The unsteadiness of some of the results obtained, and how they can under-perform compared to the datum, can clearly be seen.

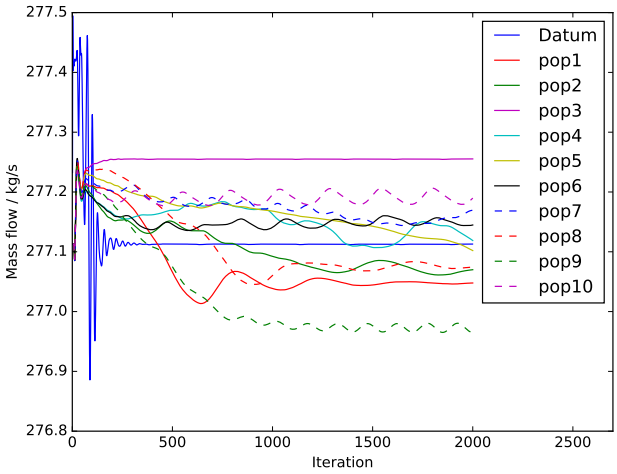

Fig. 4.31 Simulation mass flow histories for 10 populations of optimisation generation 7.

The optimisation convergence history is shown in Fig 4.32. As can be seen, many of the initial designs explored have a worse performance than the datum. The optimiser is able to move away from these designs however and converges on an improved design.

The optimised individual bump geometry is shown in Figure 4.33. The flow features of the optimised and datum geometries are shown in Figure 4.34. It can be seen that the design successfully delays the shock, but that separation behind the bumps limits the benefit predicted.

Table 4.3 shows the mass flow increase and pressure loss reduction predicted for the optimised design. To further understand these effects, the θ distribution of delta mass flow and pressure loss compared to the datum design can be seen in Figure 4.35. These results show that benefit has clearly been predicted through the application of the bumps, but this benefit oscillates with the bump positions, suggesting that the individual nature of these bumps may be detrimental to their overall performance by causing separation.

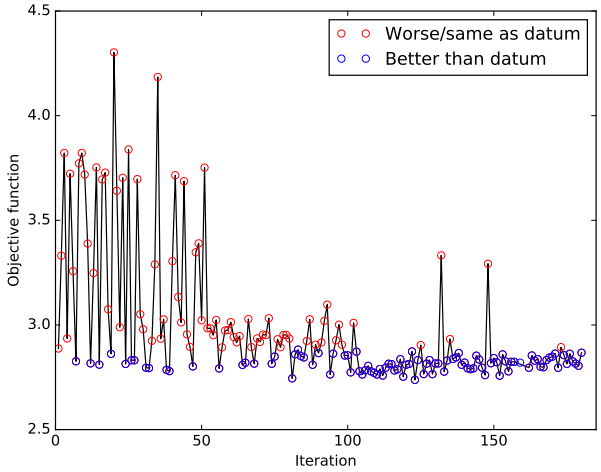

Fig. 4.32 The convergence history of the individual bump optimisation.

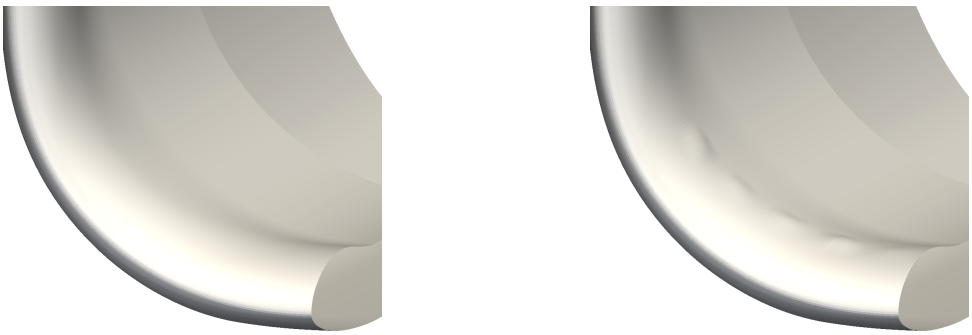

(a) (b)

Fig. 4.33 (a) Datum geometry, (b) Optimised individual bump geometry.

Table 4.3 Comparison of the performance of the optimised individual bump with the datum geometry.

	Mass flow / kg/s	**Delta mass flow**	**Benefit / %**	**Total pressure loss**	**Delta total pressure loss**	**Benefit / %**
Datum (@ α=0^o)	277.87	0.0	-	239.05	0.0	-
Datum	277.11	-0.76	-	420.48	181.43	-
Optimised individual	277.26	-0.61	19.74	388.85	149.79	17.44

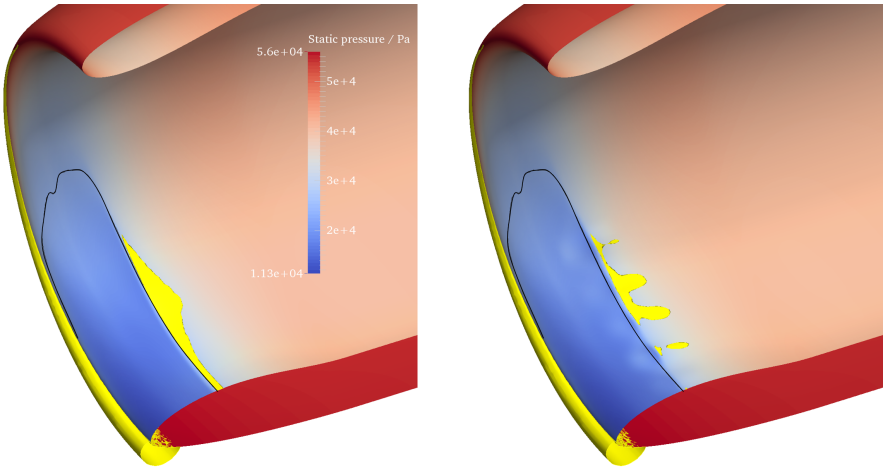

Fig. 4.34 Shock position and separation (indicated in yellow) for the datum geometry (left) and the optimised individual bump geometry (right). The datum geometry shock position is indicated by the black line.

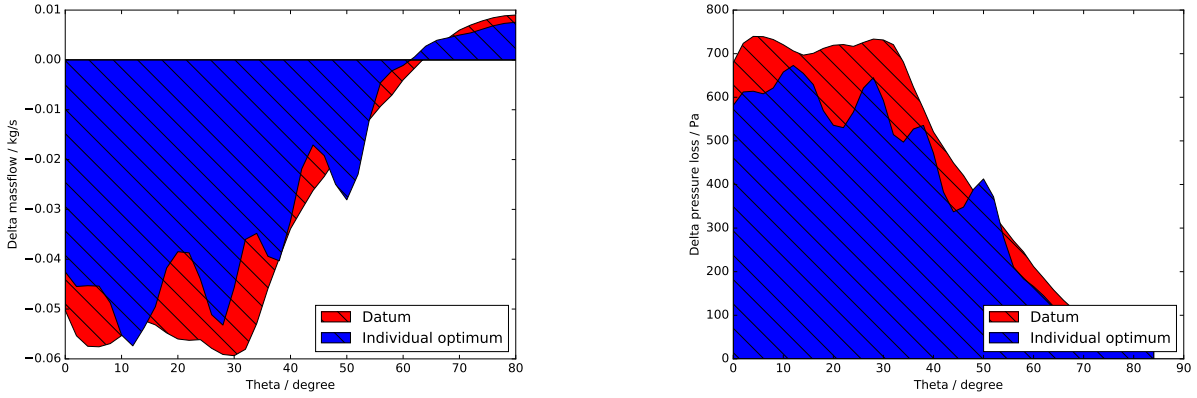

(a) (b)

Fig. 4.35 The effect of optimised individual bumps on (a) mass flow and (b) pressure loss with theta.

4.6.2 Optimisation of continuous bumps

An optimisation using continuous bumps is here carried out to assess how the result compares to the optimised individual bumps. Previous results from Section 4.5.1 suggest that the continuous bumps can out-perform the individual designs. Here, the continuous bumps are optimised using a similar set-up to the individual ones (where the bump is specified using five θ control positions) but the geometry is smoothly interpolated in the θ direction using a cubic spline.

The convergence history of the continuous bump optimisation is shown in Figure 4.36. Again, many of the bump designs under-perform compared to the datum, but the optimiser is able to converge on an improved design.

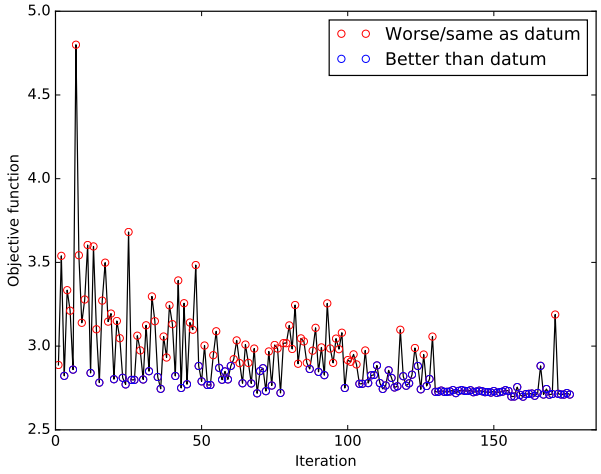

Fig. 4.36 The convergence history of the continuous bump optimisation.

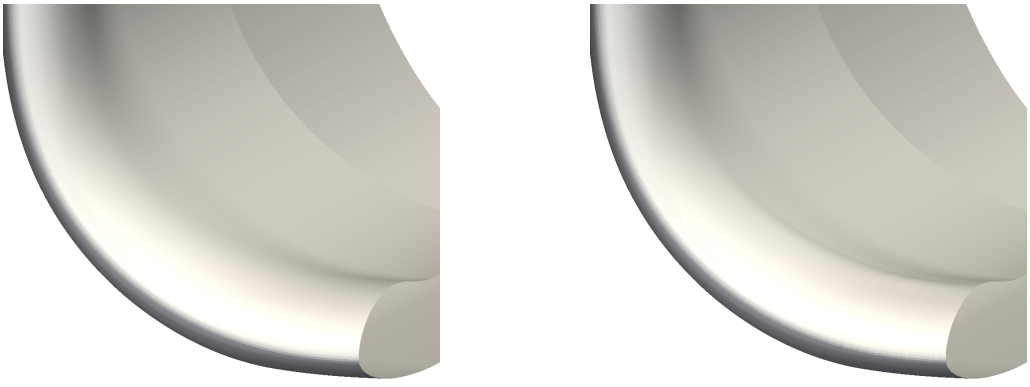

(a) (b)

Fig. 4.37 (a) Datum geometry, (b) Optimised continuous bump geometry.

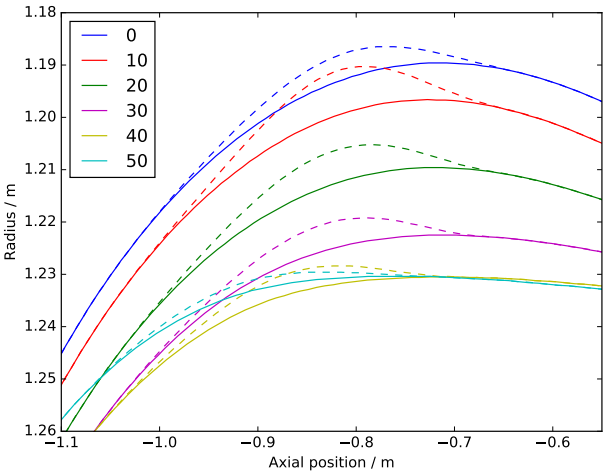

Fig. 4.38 Constant-theta geometry contours of the optimised (dashed) and datum (solid lines) continuous bump geometries.

The optimised continuous bump geometry is compared to the datum in Figures 4.37 and 4.38. The impact of the bumps on the shock and separation can be seen in Figure 4.39. A significant delay in the shock position and elimination of the separation has been predicted at lower θ positions. At higher values of θ the shock has still been delayed, but the bumps have not been able to reduce the separation in this region compared to the datum geometry.

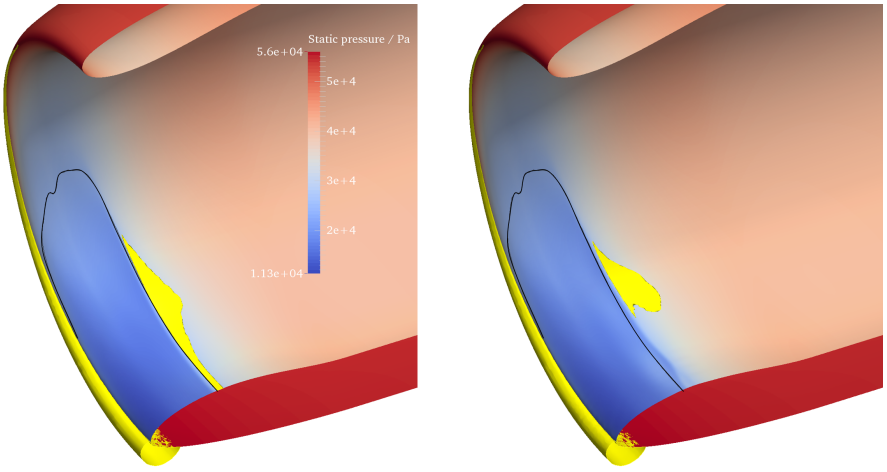

Fig. 4.39 Shock position and separation (indicated in yellow) for the datum geometry (left) and the optimised continuous bump geometry (right). The datum geometry shock position is indicated by the black line.

A comparison of this design's performance with the optimised individual bump and the datum geometry is given in Table 4.4. It can be seen that the optimised continuous design has increased the

mass flow and reduced the pressure loss compared to both the datum and the optimised individual bump design. The mass flow difference to the datum geometry at zero incidence has been reduced by 25%, and the pressure loss by 22%.

Table 4.4 Comparison of the optimised continuous bump with datum.

	Mass flow / kg/s	**Delta mass flow**	**Benefit / %**	**Total pressure loss**	**Delta total pressure loss**	**Benefit / %**
Datum (@ α=0^o)	277.87	0.0	-	239.05	0.0	-
Datum (@ α=31^o)	277.11	-0.76	-	420.48	181.43	-
Optimised individual	277.26	-0.61	19.74	388.85	149.79	17.44
Optimised continuous	277.30	-0.57	25.00	379.86	140.80	22.39

The distribution of delta mass flow and delta pressure loss with θ shows the benefit provided by the optimised geometry (see Figure 4.40). The reduction in pressure loss and mass flow increase are both maximum from $\theta = 0 - 30^o$, with reduced benefit after this point. This matches the result from Figure 4.39, which showed a reduction in separation only for lower values of θ. The region around θ = 40^o-50^o appears to be the most difficult to reduce separation, as was also shown in the parametric studies in Section 4.5.1.

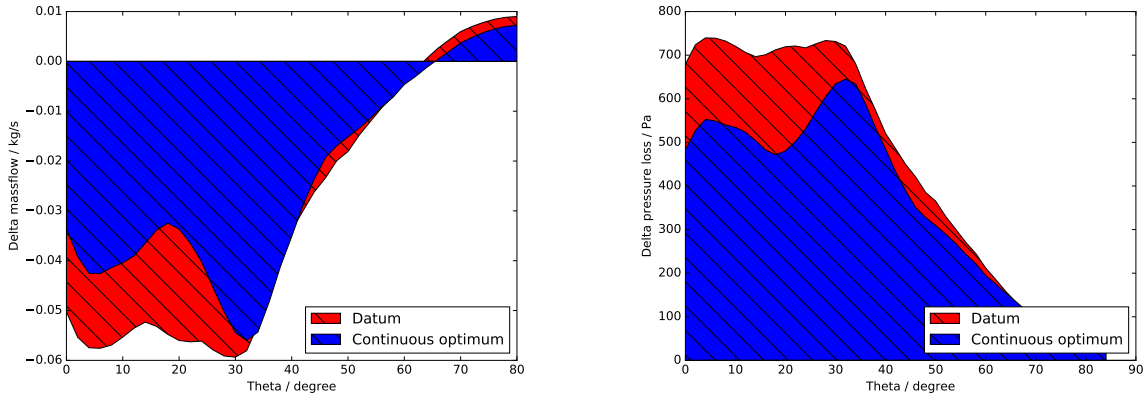

(a) (b)

Fig. 4.40 The delta (a) mass flow and (b) pressure loss distributions with theta for the datum and optimised continuous bump geometries.

4.7 Varying the bump start/end positions

The above work demonstrates the potential benefit that can be achieved through the use of shock bumps (with the continuous bumps being most promising), but also how this design approach offers

limited benefit in certain θ regions. To improve the performance of the continuous optimised bump design, the axial start and end positions are here allowed to vary. This offers more flexibility in the bump geometry and the potential to find a design that performs well across the range of θ positions.

Figure 4.41 and table 4.5 shows the permitted variation in bump start and end positions.

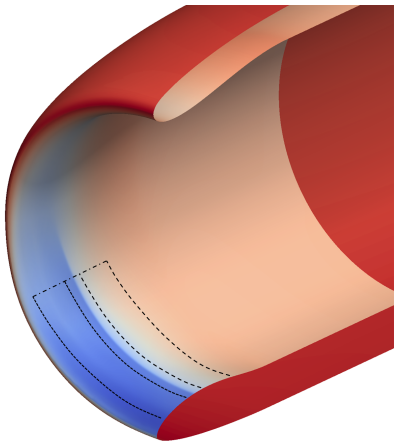

Fig. 4.41 The bump start (dots) and end (dashes) position regions.

Table 4.5 The parameter ranges for the varying start/end point optimisation.

Theta control position / ^{o}	**Start min axial position / m**	**Start max axial position / m**	**End min axial position / m**	**End max axial position / m**
0	-1.15	-0.95	-0.725	-0.525
15	-1.165	-0.965	-0.74	-0.54
30	-1.18	-0.98	-0.755	-0.555
45	-1.205	-1.005	-0.78	-0.58
60	-1.23	-1.03	-0.805	-0.605

With a greater number of parameters and more variation between designs, the moving start/end position optimisation initially struggled to converge, with a large number of poor designs produced. This can be seen in Figure 4.42a. To improve this, a constraint was added to the MAM optimisation process that designs should have a mass flow greater than the datum. Although the objective function meant the optimisation was already searching for designs with greater mass flow than the datum, adding this constraint allowed MAM to narrow down the search more effectively. A comparison of the two optimisation convergences, with and without this constraint, can be seen in Figure 4.42. It can be seen that the first generation is identical, but after this point the constrained optimisation begins to converge on improved designs. An extended view of the constrained optimisation is shown in Figure 4.43 where the MAM generations and the successful optimisation convergence can clearly be seen.

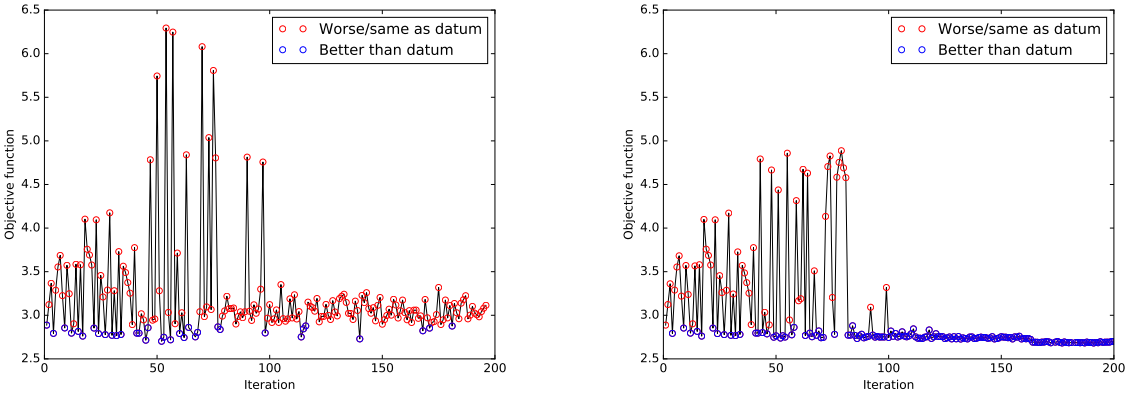

(a) (b)

Fig. 4.42 The MAM convergence histories (a) unconstrained and (b) constrained to perform better than datum.

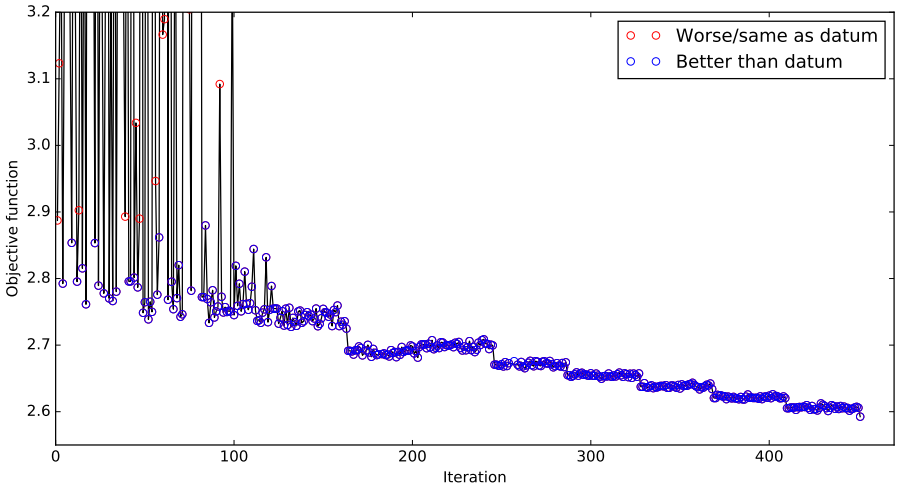

Fig. 4.43 Zoomed view of constrained MAM convergence.

With the constraint added, the optimisation managed to converge well, reaching a design far better than achieved without moving the bump start and end positions. The effect of this optimised geometry on the flow is shown in Figure 4.44. It can be seen how the optimised design has completely eliminated the separation and delayed the shock position along the region where the bump is placed.

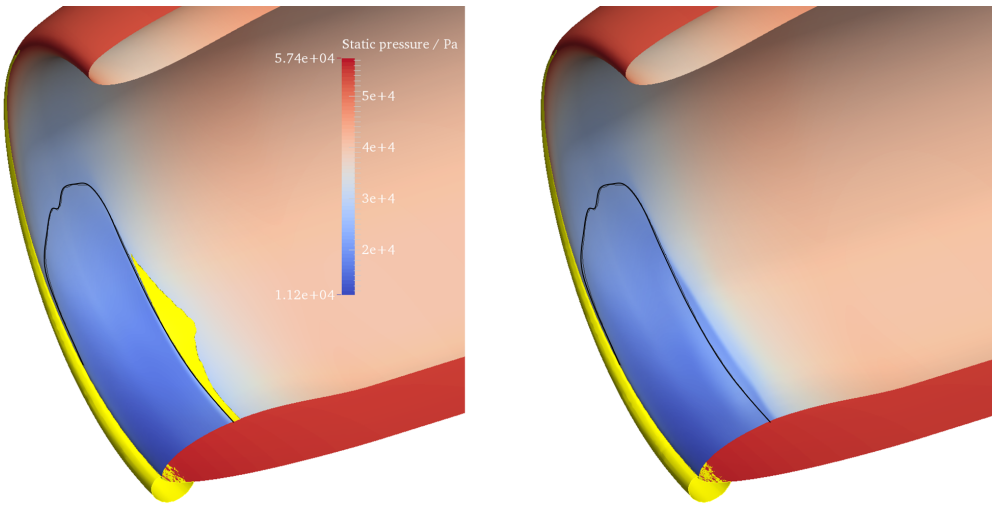

Fig. 4.44 Shock position and separation (yellow) for the datum (left) and optimised (start/end position varying) (right) geometries. The datum shock position is indicated by the black line.

Table 4.6 gives a performance comparison with the datum design of this optimised geometry. From here on, this geometry, designed at $\alpha = 31^o$, is referred to as Design-31. It can be seen that Design-31 results in a 37% reduction in the difference in absolute total pressure loss compared to the datum geometry.

Table 4.6 Comparison of the optimised Design-31 bump with datum.

	Mass flow / kg/s	**Delta mass flow**	**Benefit / %**	**Total pressure loss**	**Delta total pressure loss**	**Benefit / %**
Datum (@ α=0^o)	277.87	0.0	-	239.05	0.0	-
Datum (@ α=31^o)	277.11	-0.76	-	420.48	181.43	-
Optimised individual	277.26	-0.61	19.74	388.85	149.79	17.44
Optimised continuous	277.30	-0.57	25.00	379.86	140.80	22.39
Optimised Design-31	277.41	-0.46	39.47	353.20	114.15	37.08

The Design-31 geometry is compared to the continuous-optimised and datum geometries in Figure 4.45. The Design-31 bumps are in a similar position to the continuous design bumps at $\theta = 0$ and 10^o, but have moved their end position significantly downstream at greater θ positions. This shows that the initial bump positions used in the first set of optimisations were restricting the achievable benefit, and the ability of the bump start/end positions to move is key to creating a good design.

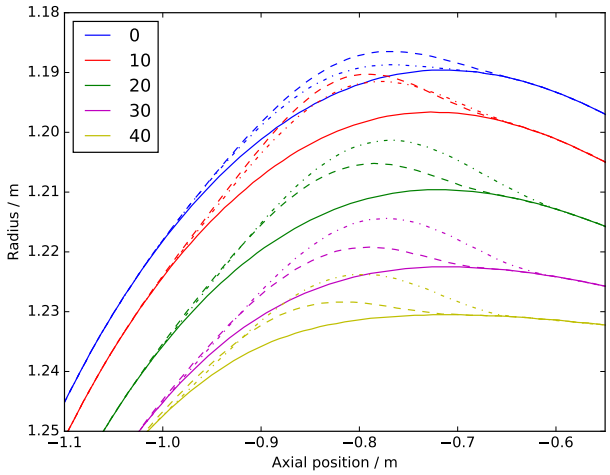

Fig. 4.45 Constant-theta geometry contours of the optimised-continuous (dash), datum (solid) and Design-31 (dash-dot) bump geometries.

4.7.1 Further analysis of the most promising design (Design-31)

Design-31 is the most beneficial design found. To understand where the benefit comes from, detailed analysis of the impact of the bumps on the flow are carried out in this section.

The above results show the benefit predicted through delaying the shock and reducing separation and pressure loss. To demonstrate how the bumps delay the shock and reduce its strength, the pressure contours on a constant-θ slice near to bottom dead centre for the datum and optimised geometries can be compared (see Figure 4.46). The pre-compression upstream of the shock and the knee-shape structure formed (reminiscent of that seen in Figure 4.4) can be seen. The pressure contours where the shock meets the surface spread out, reducing the pressure/velocity gradient across the shock.

Mach number plots (measured along a line 0.1m from the wall, (shown in Figure 4.18a) show the change in the shock structure caused by the bumps at various θ positions.

It can be seen how the shock is weakened by the presence of the bump. The pre-compression effect gradually reduces the Mach number upstream of the shock. The magnitude of the jump in velocity/pressure across the shock is significantly reduced for the optimised design. To clearly assess the reduction in shock strength, the entropy change across the shock can be used. This is given in Figure 4.48 for various θ positions. It can be seen that the shock strength has been significantly weakened by up to 90%. At θ=50^o the benefit becomes reduced, as the design begins blending into the datum geometry.

Figure 4.49 shows the theta distributions of delta pressure loss and mass flow from the datum geometry at θ=0^o. The benefit is clearly seen and shows how moving the start and end positions has improved the result, especially around θ=30-40^o.

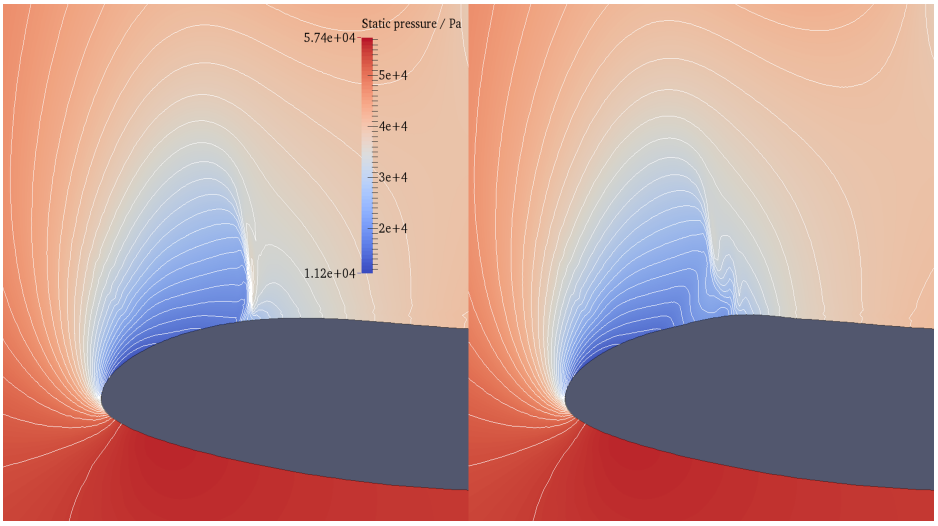

Fig. 4.46 Static pressure contours for the datum (left) and optimised continuous geometry (right) at $\theta = 20$.

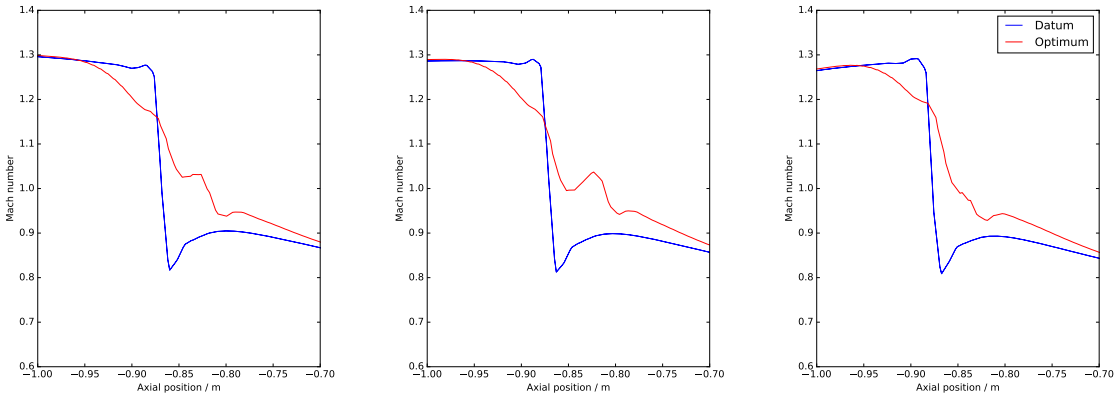

(a) (b) (c)

Fig. 4.47 Mach number profiles across the shock at 0.1m wall distance at $\theta =$ (a) 10^o, (b) 20^o and (c) 30^o for Design-31.

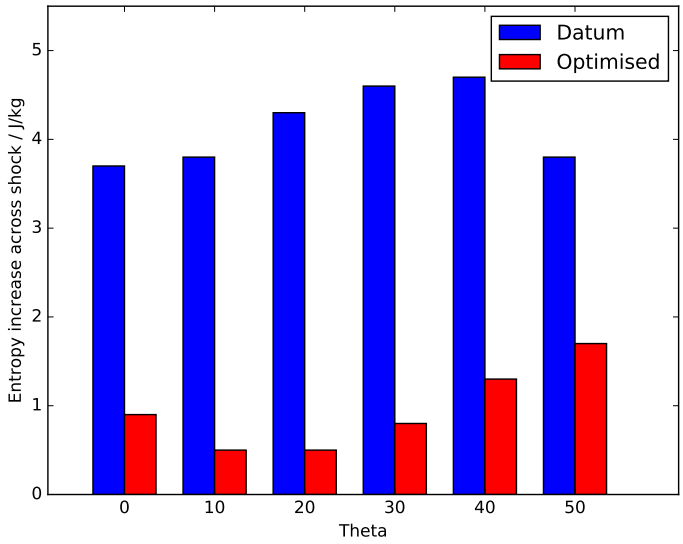

Fig. 4.48 Entropy increase across the shock for the datum and Design-31 geometries.

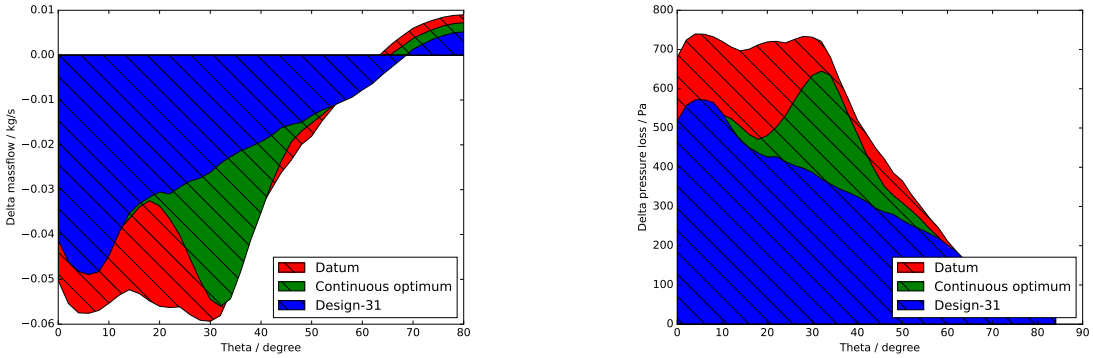

(a) (b)

Fig. 4.49 The (a) mass flow and (b) pressure loss distributions with theta for the datum and optimised Design-31 geometries compared to the best continuous bump geometry (with fixed start/end positions).

4.7.2 The effect of turbulence model

To ensure that the turbulence model used in this research did not adversely influence the results, the datum and Design-31 geometries were also simulated using the k-omega SST model instead of SA.

It can be seen in Fig 4.50 that although differences exist in the magnitudes of separation predicted for the datum design at $\alpha = 31^o$, the maximum separation θ position and distribution is the same for both cases. Complete elimination of the separation is also seen in both SA and k-omega simulations of the optimised design. Table 4.7 shows how the total pressure loss reduction and mass flow increase are similar for both.

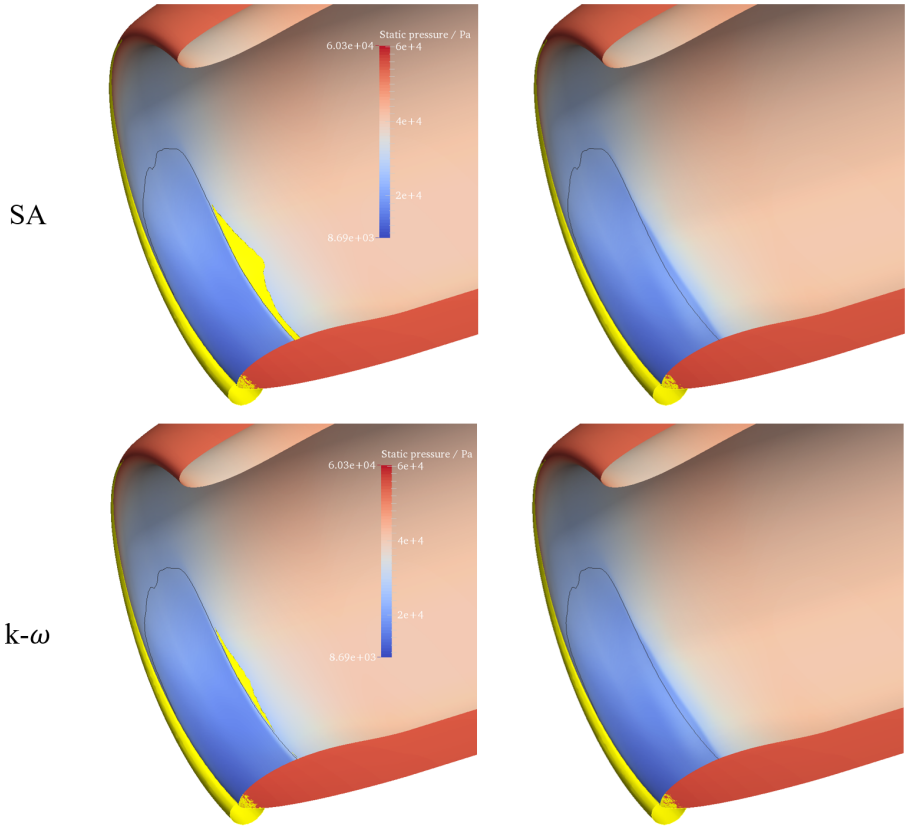

Fig. 4.50 Datum and Design-31 flow features (separation in yellow) with SA and k-omega turbulence models.

These results indicate that the optimised design is beneficial regardless of the turbulence model used, and the results can be trusted.

Table 4.7 Comparison of the two turbulence models tested

	Mass flow / kg/s	**Difference in mass flow**	**Total pressure loss**	**Difference in total pressure loss**
Datum SA	277.114		420.5	
Design-31 SA	277.366	0.251	363.9	-56.6
Datum k-ω	277.097		413.5	
Design-31 k-ω	277.341	0.245	359.3	-54.2

4.7.3 Performance at various angles of attack

The two primary objectives of this work were to assess whether bumps can be used to weaken the shock and reduce separation at a non-fully-separated angle of attack, and then to see whether this design can be used to extend the maximum unseparated angle of attack of the intake. The first of these objectives was clearly achieved above through Design-31.

The next step is to assess whether this optimised bump, that can weaken the shock and reduce separation at $\alpha = 31^o$, is capable of increasing the unseparated angle of incidence that the nacelle can achieve. To investigate this, the optimised geometry (Design-31) was run at a range of incidences. The resulting total pressure loss at various angles of attack for Design-31 and the datum geometry are shown in Figure 4.51.

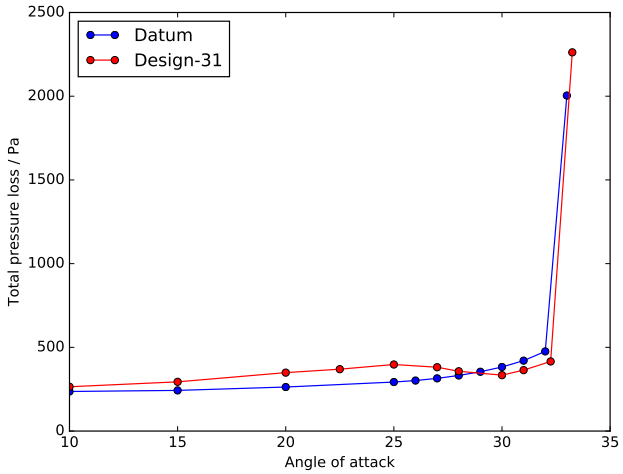

Fig. 4.51 Variation in total pressure loss at the fan face with angle of attack for the datum and Design-31 geometries.

It can be seen that around 31 degrees (the point at which it was designed) Design-31 out-performs the datum geometry by increasing the massflow and reducing the entropy generated. It can also be seen that the operating range of the nacelle has been increased slightly, by 0.25 degrees.

At lower angles of attack Design-31 performs worse than the datum however, generating greater pressure loss. This is due to the bump placement being mis-aligned with the shock at these incidence angles, possibly increasing the shock strength or leading to regions of separation. The addition of the bump has also slightly increased the surface length area, resulting in greater friction losses. At cruise condition ($\alpha = 0^o$) the pressure loss increase is around 3%.

This drop in performance towards low angles of attack is not a major concern however, as the main objective and potential benefit of this work comes from increasing the operating range of the engine (which demonstrates that a more aggressive inlet may be used with the same operating range, therefore reducing drag).

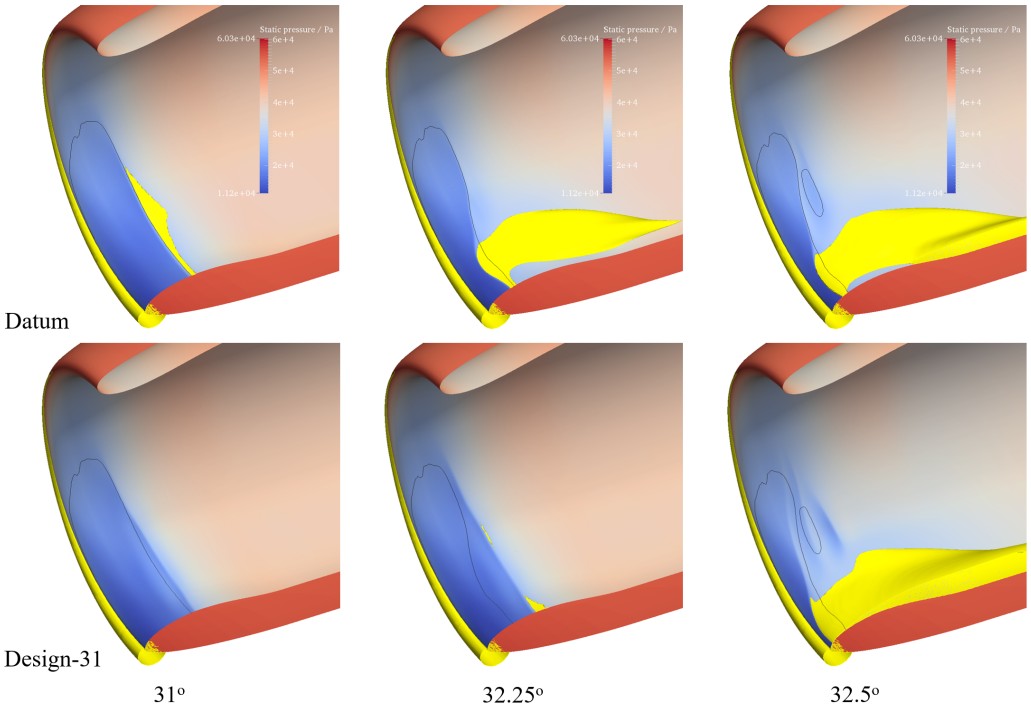

Fig. 4.52 Variation in separation (in yellow) and shock position for datum and Design-31 near to the fully separated angle of attack.

Figure 4.52 shows how the optimised design reduces the separation at $\alpha = 31^o$ and extends the unseparated range of the intake to $\alpha = 32.25^o$ (the last attached point for the datum design is $\alpha = 32^o$). The improvement in operating range of 0.25 degrees is unsatisfactory however, as it does not provide significant benefit.

It has been predicted through these simulations that designing the bump geometry at an unseparated angle of attack (31 degrees) can provide significant benefit at the operating point, and slightly extend the unseparated angle of attack, but only by a small amount. Further work is therefore needed to improve the design.

4.8 Designing bumps at a separated angle of attack

It was demonstrated above that designing the bump at an unseparated angle of attack (31 degrees) only marginally increases the unseparated angle of incidence. This is likely due to the bump geometry not being correctly aligned with the shock position once the incidence is increased to the separated angle (i.e. $\alpha > 32^o$).

The issue with attempting to optimise the bump at a separated angle of attack to reduce this separation (i.e. optimising at $\alpha > 32^o$) is that the majority of geometries tested are fully separated. When the geometries fully separate, there can be little difference between the objective functions of each bump design, making it difficult for the optimiser to recognise beneficial designs and struggle to converge. This was the reason that the approach of optimising at a non-separated angle of incidence was initially adopted.

4.8.1 Optimising at $\alpha = 33^o$

It is clear from the above results that a different approach may be needed to increase the stable operating range of the nacelle. An optimisation was attempted at $\alpha = 33^o$ with varying start/end positions, as carried out for Design-31. The seed point of the optimisation (i.e. the initial design that MAM begins its search from) was the datum nacelle geometry with no bump. For this optimisation none of the desgins tested in the first couple of generations were unseparated and therefore the optimisation failed to converge or find any useful designs.

This process was then repeated, but using Design-31 as the seed-point of the optimisation. This moves the initial search of the optimisation to designs that are related to Design-31. Design-31 is separated at $\alpha = 33^o$, but as can be seen in Figure 4.53, some of the designs related to it are not (i.e. mass flow around 277kg/s).

The massflow convergences shown here indicate the difficulty with this approach; the majority of the designs are separated, and therefore the optimisation struggles to converge. The binary nature of designs that are separated or unseparated can clearly be seen. This does demonstrate that there are designs capable of operating un-separated at 33 degrees incidence however, providing an increase in operating range of 1 degree over the datum geometry.

4.8.2 Optimising at $\theta = 34$ and 35^o

The same approach was repeated at $\alpha = 34$ and 35^o. While the optimisations generally struggled to converge in the same way that the optimisation at $\alpha = 33^o$ did, leaving the optimisation to continue running and searching through many designs allows it to move the search towards unseparated geometries. This is demonstrated in Figure 4.54, which shows the optimisation history for $\alpha = 35^o$. It can be seen that almost all of the first 100 designs tested are fully separated (i.e. objective function > 8), but due to the optimiser finding a couple of non-separated designs, it is able to move the search in

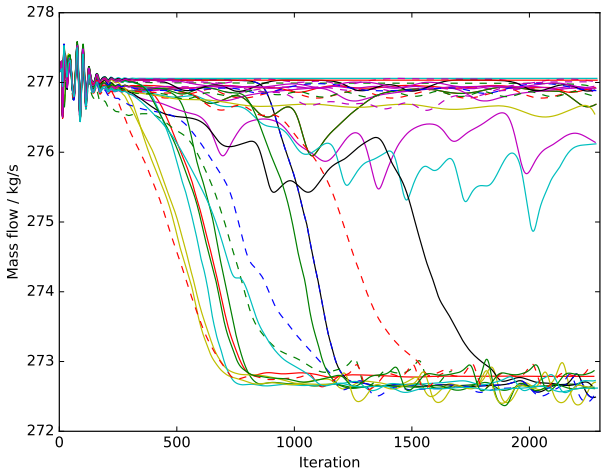

Fig. 4.53 The simulation mass flow history for each design in generation 1 of the Design-33 optimisation.

a beneficial direction. By the end of the optimisation, it can be seen that the majority of the designs are unseparated.

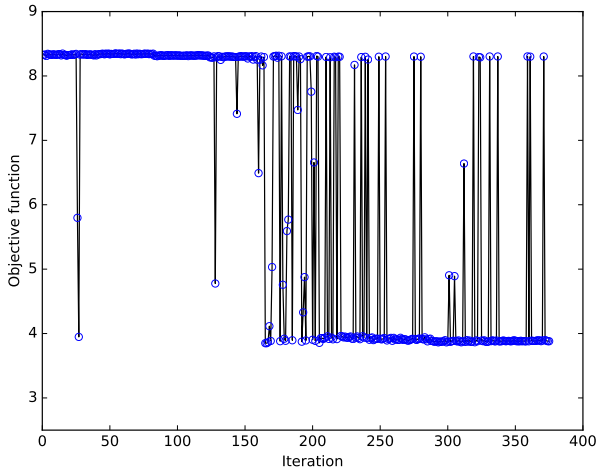

Fig. 4.54 The performance history for the Design-35 optimisation.

Figures 4.55, 4.56, 4.59 and 4.60 show the flow features of the optimised designs compared to the datum at $\alpha = 34$ and 35^o. It can be seen that in both cases the optimised designs have prevented the large separation from occuring, and the blockage at the fan face (shown by the entropy contour) is minimised. Some small regions of separation are present but these are far less than if the design were fully separated. The presence of the bumps has stabilised the shock position; dramatically reducing the separation.

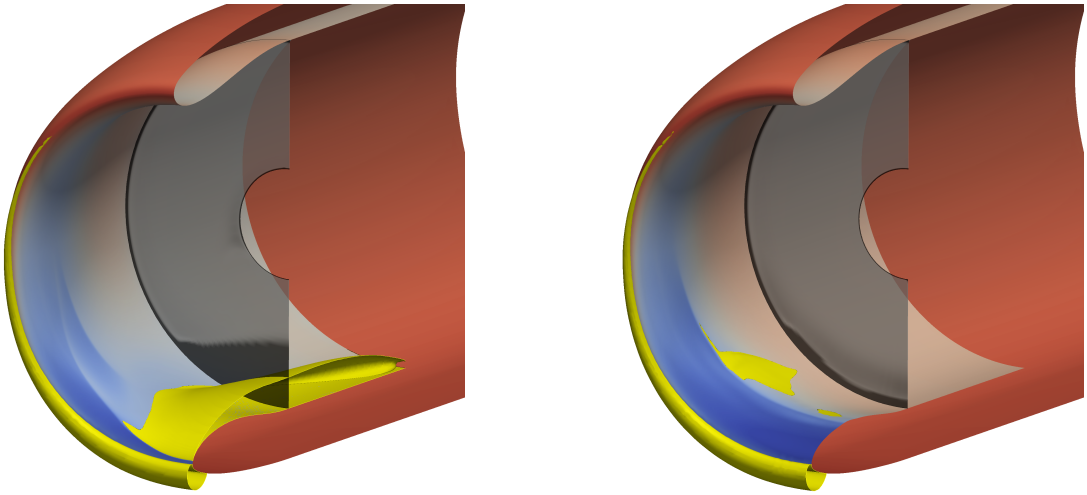

(a) (b)

Fig. 4.55 The shock position (static pressure contours), separation (yellow regions) and fan face blockage (dark regions indicate high entropy) for the (a) datum and (b) optimised Design-34 geometries at $\alpha = 34^o$.

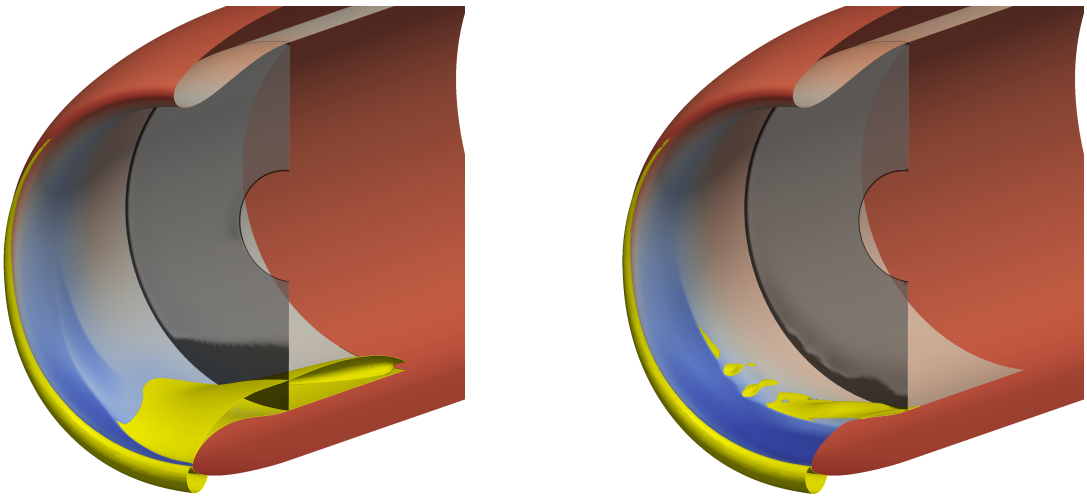

(a) (b)

Fig. 4.56 The shock position (static pressure contours), separation (yellow regions) and fan face blockage (dark regions indicate high entropy) for the (a) datum and (b) optimised Design-35 geometries at $\alpha = 35^o$.

Figures 4.57 and 4.58 show the total pressure loss versus angle of attack for the Design-34 and Design-35 geometries versus the datum. The increase in stable operating rage of 2 degrees for Design-34 and 3 degrees for Design-35 can clearly be seen. It can also be seen how the optimised geometries perform best around the region at which they were designed, but under-perform compared to the datum at lower incidence angles.

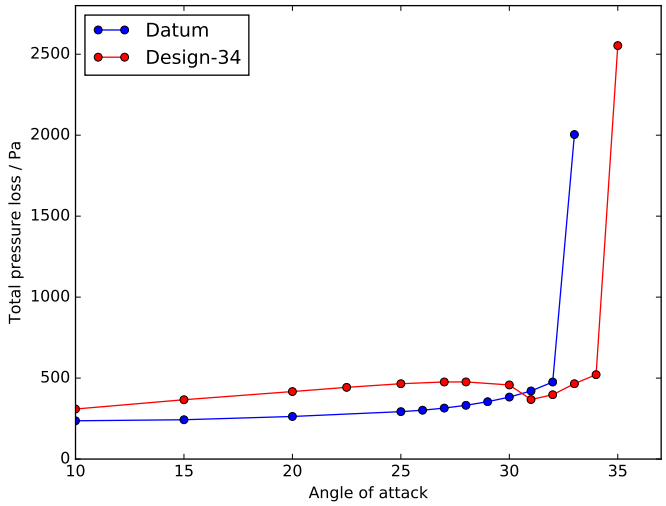

Fig. 4.57 Total pressure loss versus incidence angle compared to the datum geometry for a) Design-34.

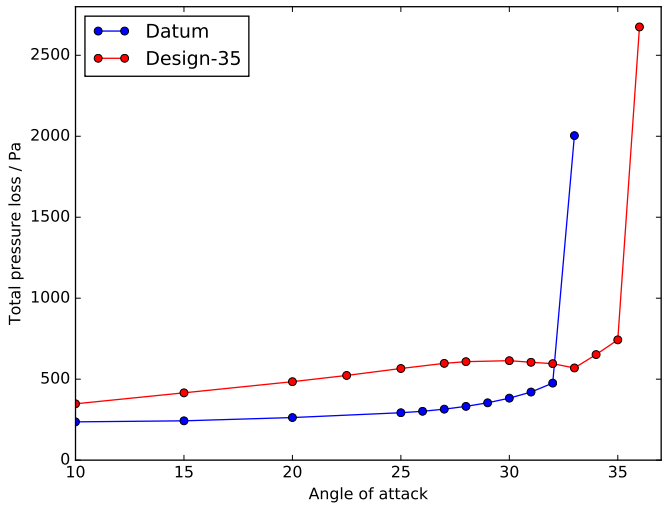

Fig. 4.58 Total pressure loss versus incidence angle compared to the datum geometry for a) Design-35.

Figure 4.61 shows the Design-34 and Design-35 geometries compared to the datum and Design-31. It can be seen that there is not much difference between Design-31 and Design-34 apart from at

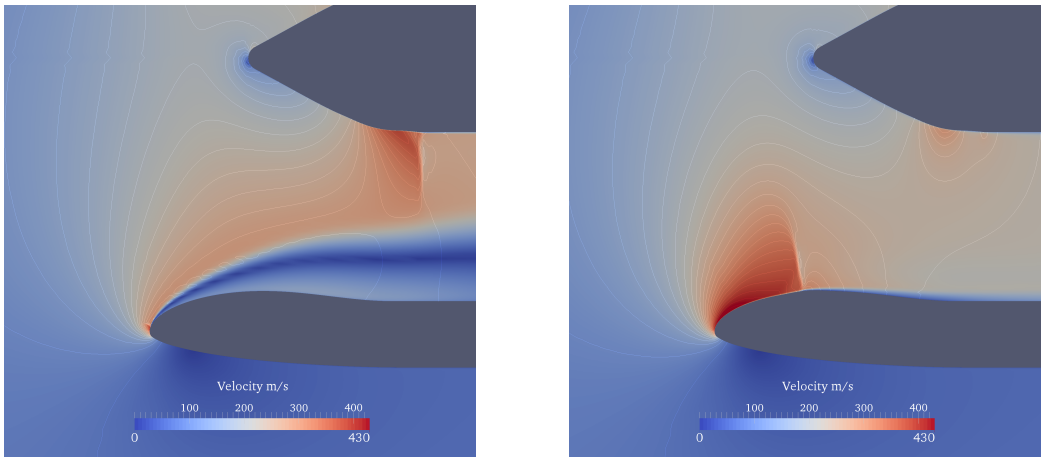

(a) (b)

Fig. 4.59 Velocity magnitude contour at theta = 0, $\alpha = 34$ degrees for the a) datum and b) Design-34 geometries.

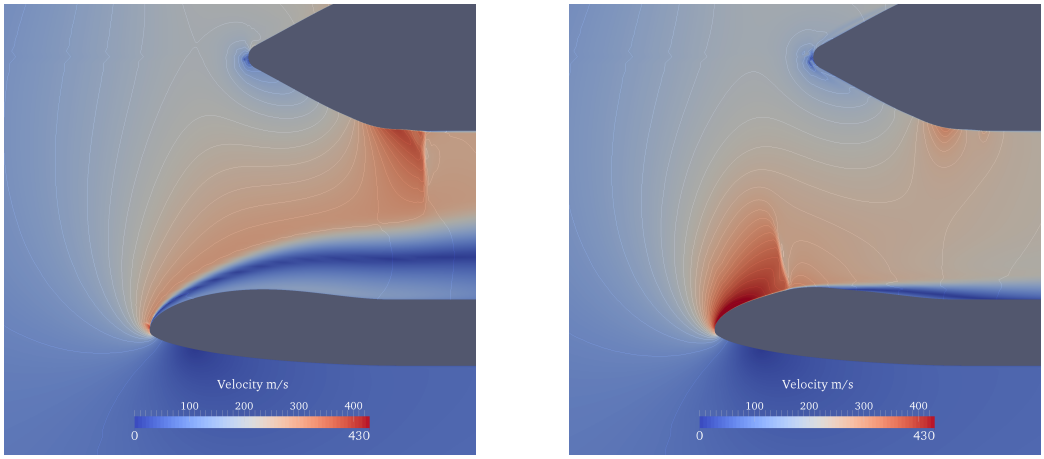

(a) (b)

Fig. 4.60 Velocity magnitude contour at theta = 0, $\alpha = 35$ degrees for the a) datum and b) Design-35 geometries.

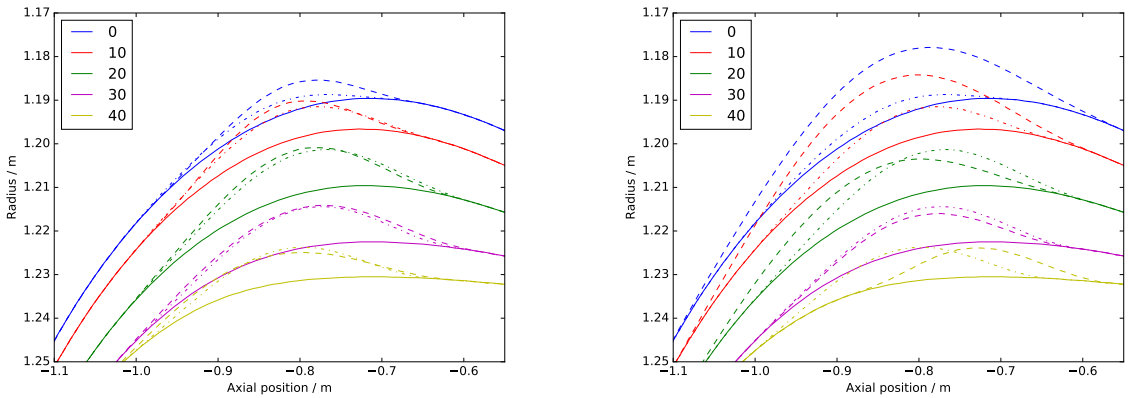

(a) (b)

Fig. 4.61 Constant theta-slices for a) Design-34 and b) Design-35 (dash) compared to the datum geometry (solid) and Design-31 (dot dash).

bottom dead centre. This indicates how critical this region is for controlling separation, but also indicates the high sensitivity of the flow to small changes in the geometry. Design-35 has bumps of far greater amplitude. It is likely that this design is not close to the optimum design, but is simply the best that was discovered due to the difficulty had by the optimiser (caused by many designs being separated). The very large bumps cause some excessive separation, as shown in Figure 4.56. Although it is also possible that this magnitude of bump is necessary at this high incidence.

The axial position of the shock (for the lower θ values) for the datum geometry near to separation is around -0.86 m. Figure 4.61 shows how the bump geometries are all designed with their peak shortly after this point. This was demonstrated previously in the literature to be the best practice for reducing shock strength, and is confirmed again here.

4.8.3 Summary of various designs

Figure 4.62 shows a comparison of the various optimised designs' performance at a range of incidence. It can be seen that the higher the angle of attack the geometry was optimised at, the greater its predicted operating range, but the greater its pressure loss at lower incidence angles. As mentioned previously however, it may be advantageous to trade-off a small increase in pressure loss at lower incidence for the opportunity to use a design that can safely operate at higher incidence, or re-design the intake to be more aggressive (shorter/thinner and therefore lighter and with less drag) but using bumps that allow it to maintain the same safe operating range.

It is key to note that the work carried out in the first part of this chapter to optimise the bump geometry at 31 degrees was necessary (despite this design not improving the unseparated angle of attack) to be able to find designs that can operate at higher incidences. Without an appropriate initial

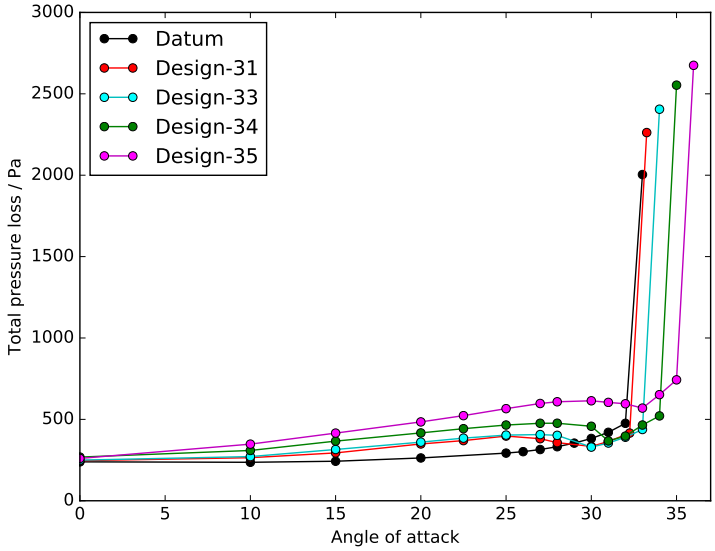

Fig. 4.62 Variation in total pressure loss at the fan face with angle of attack for the various optimised geometries.

design, the localised DoE search of the optimisations would not have been able to find beneficial geometries.

4.9 Conclusions

Several conclusions can be drawn from this research, these are summarised here:

- Continuous, smooth bumps are predicted to out-perform individual bumps
- An optimised shock control bump has the potential to completely eliminate shock-induced separation near to the limit of the intake's operating range
- The best bump designs had the peak of the bump positioned just downstream of the starting shock position
- Designing the bump geometry at an un-separated angle of incidence has the potential to increase the unseparated angle of attack, but only a small amount
- It is critical that the bump geometry is allowed as much flexibility as possible. Allowing the length of the bump to change was necessary for the optimum design to be reached
- To create designs that have an increased operating range, they must be designed at their maximum desired incidence angle

- The greater the incidence angle, the larger the bump amplitude required
- Optimising at separated incidences is difficult, and for gradient-based optimisation this will only work if the start point of the optimisation is a reasonably good design, e.g. a design that has been optimised at a non-separated incidence
- Shock control bumps have the potential to increase the un-separated angle of incidence of the nacelle by at least 3 degrees
- Bumps designed at a particular incidence will generally out-perform the datum geometry near to the incidence angle at which they were designed, but will usually under-perform at lower incidence

Due to the fact that only steady RANS simulations (with no model of the fan included) were carried out, uncertainties remain in the results presented here. It would be desirable to carry out more advanced simulations and possibly experimental testing to validate the benefit predicted for the optimised designs.

4.10 Summary

This chapter has demonstrated the benefit that can potentially be achieved through the application of shock control bumps to engine inlets. A patent application based on the results of this work has been submitted.

Chapter 3 demonstrated that shock control can also be beneficial in transonic compressor blade design. The following chapter therefore investigates the use of shock control bumps on turbomachinery compressor blades. This will assess whether some of the benefit of using shock control bumps, demonstrated in this chapter, can be replicated for transonic blades.

Chapter 5

Using shock control bumps on turbomachinery compressor blades

5.1 Motivation

It was demonstrated in Chapter 3 how the shock on transonic compressor blades has a significant influence on their performance. It was also predicted that alleviation of this shock through flexible blade shaping may improve the blade efficiency by delaying and weakening the shock and reducing separation. Chapter 4 showed how shock control bumps can potentially be used to reduce shock strength and separation on an engine intake. The focus of this chapter is on whether shock control bumps, similar to those used in Chapter 4, can be added to fan and compressor blade geometries to improve their performance.

Shock control bumps have the potential to control the shock without modifying the entire blade geometry (as was done in Chapter 3). Two cases are investigated; the transonic compressor blade NASA Rotor 37 from Chapter 3, and a modern, low speed [1], transonic Rolls-Royce fan blade (RR-FAN-2). In this chapter, the impact of shock control bumps on both blade efficiency and stall margin is investigated.

5.2 Previous work

As described in Chapter 4, shock control bumps have been tested extensively for external wing geometries. Their potential for use on turbomachinery fan/compressor blades is relatively unknown however. The only source found in the literature is by Mazaheri and Khatibirad [74], who investigated using CFD a shock control bump on a 2D (mid-span) section of the NASA rotor 67 geometry.

[1] In this context, the term 'low speed' indicates that these modern blades have lower rotational speeds than their pre-decessors, not that their velocity is actually 'low', which might suggest they would not be transonic.

In this work they added a bump modelled using the Hicks–Henne function [50]. It was shown how the interaction of the bump with the original wave structure resulted in a more desirable pressure gradient, with weaker compression waves and a more isentropic compression. The bump design was optimised and was predicted to reduce the separation area in an off-design condition. They described how this has the potential to improve the stall properties of the blade section. Two optimisations were carried out, one at the design condition and another at 4% higher rotational speed. Optimal bumps were produced for each condition, with an increase in efficiency of 0.67% for the on-design case and 2.9 % in the off design case reported. The optimised geometry for the design condition is shown in Figure 5.1.

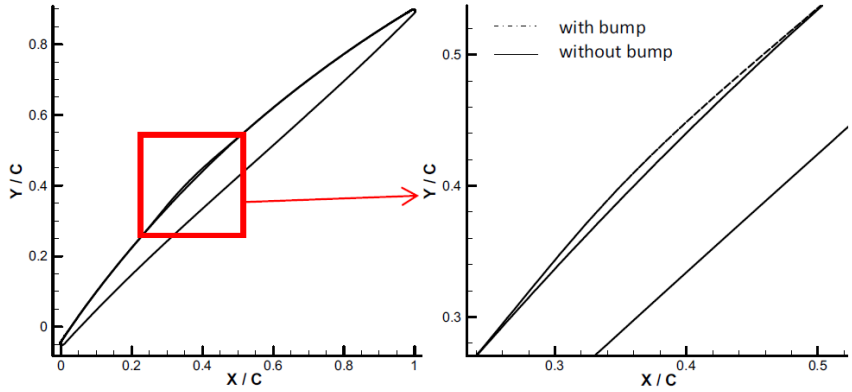

Fig. 5.1 Datum and Optimised shock control bumps on the mid section of NASA Rotor 67 [74].

This work demonstrated the benefit that bumps may provide at both on and off-design conditions, and their potential to improve stall margin. The simplified 2D analysis lacks accuracy however as the complex behaviour of radial and 3D separated flow cannot be predicted. For a thorough understanding of the potential for the use of shock control bumps, 3D analysis is needed to truly assess their effect. This is carried out here.

5.3 Compressor blade shock control

As described in Chapter 3, NASA Rotor 37 has a very strong shock wave which causes large separation, decreasing the blade efficiency. It was predicted in Chapter 3 how an s-shaped design produced a pre-compression geometry and, combined with 3D shaping of the blade, this was able to delay the passage shock, reduce separation and provide a significant increase in efficiency.

One of the questions remaining from this work was - how much of the efficiency benefit came from the s-shaped, pre-compression design, and how much from the 3D (e.g. leaning in the direction of rotation) shaping of the blade? These two effects are hard to decouple, so the testing of shock control bumps will inform how much benefit is achievable purely through control of the shock. It is also desirable to demonstrate a method of controlling the shock without modification to the whole

three-dimensional blade geometry. The use of shock control bumps offers the potential to answer both of these questions.

5.3.1 Bump parameterisation

The same, CST based approach is used here as in Chapter 4 for the definition of the bump geometries. To generate the geometries, a CST bump parameterisation technique was developed inside the PADRAM [110] meshing software. The technique modifies each 2D radial section of the blade geometry, adding a bump. The properties of these 2D bumps are smoothly interpolated in the radial direction from control sections. Both continuous and individual CST bumps are tested, controlled by the bump start and end positions, the four Bernstein polynomial amplitudes and the span-wise distribution. Examples of the blade with individual and continuous bumps added is shown in Figure 5.2.

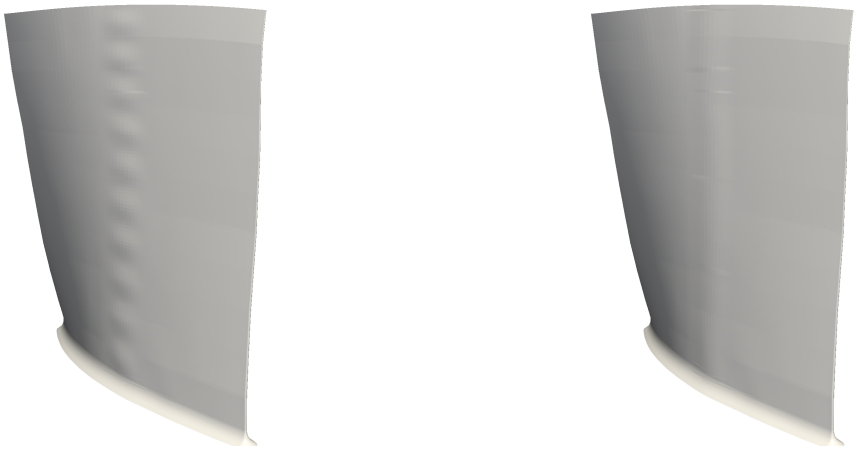

(a) (b)

Fig. 5.2 a) example individual bump geometry and b) example continuous bump geometry.

Figure 5.3 shows the flow features of the datum NASA Rotor 37 at design point. The shock position is highlighted by static pressure contours and the shock-induced separation is shown using an orange isosurface of zero axial velocity. For the optimisations the bumps' positions were allowed to vary as is described in Table 5.1. The design used to initialise the optimisation process had bumps positioned with approximately 60% of the bump downstream of the datum shock, as is known to be beneficial [118].

5.3.2 Optimisation

The same CFD set up is used here as in Chapter 3, using the same mesh and boundary conditions. Some parametric studies were carried out to identify the range of bump positions and amplitudes to be tested. These results are not shown here for brevity. As was demonstrated in Chapter 4, to

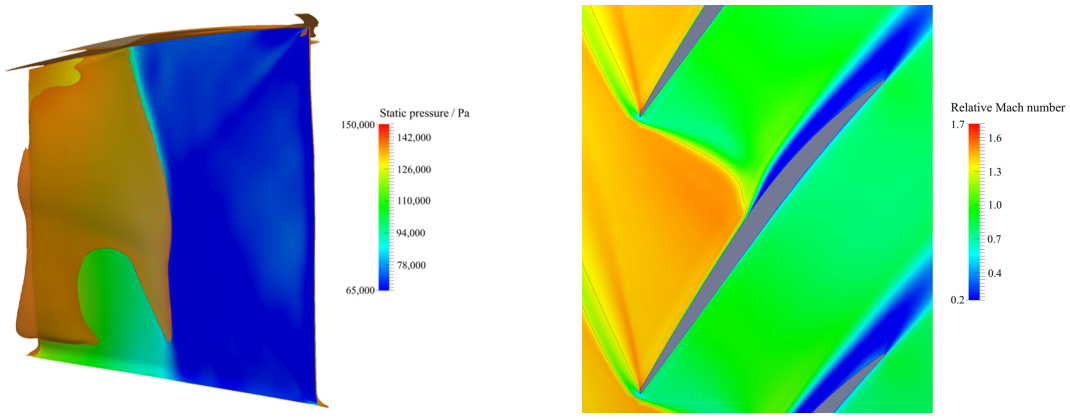

(a) (b)

Fig. 5.3 (a) 3D separation on the datum geometry (orange), (b) Rel. Mach contour at 60% span.

achieve significant benefit via the use of shock control bumps their geometries need to be optimised. The objective function used is efficiency. Two optimisations were carried out, using individual and continuous bumps.

Individual bump optimisation

For the individual bump optimisation, 10 bumps were specified in the radial direction. This was based on typical numbers of spanwise bumps used by previous researchers [52]. To reduce the number of parameters in the optimisation these were controlled at 5 radial heights. The start and end points of the bumps are allowed to vary as shown in table 5.1. Towards the tip, the bump placement and movement range are increased in chord-wise position as the shock is sat further downstream at the tip. The MAM optimiser was used with the objective function set as blade efficiency.

Table 5.1 The Rotor 37 bump parameter ranges.

Spanwise bump position	**Minimum start point / chord**	**Maximum start point / chord**	**Minimum end point / chord**	**Maximum end point / chord**
0-0.6	0.25	0.55	0.55	0.85
0.8	0.25	0.59	0.59	0.9
1.0	0.25	0.63	0.63	0.9

The optimisation convergence histories are not given for brevity. The optimised, individual bump design and the effect on the flow is shown in Figures 5.4 and 5.5. The white line on the optimised geometry shows the datum position of the shock. It can be seen that the bumps have delayed the shock, but that the shock position is irregular.

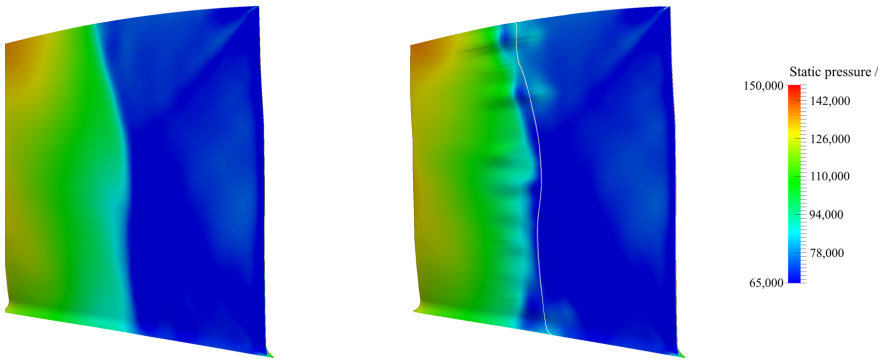

Fig. 5.4 Datum (left) and individual bump optimum (right) static pressure contours.

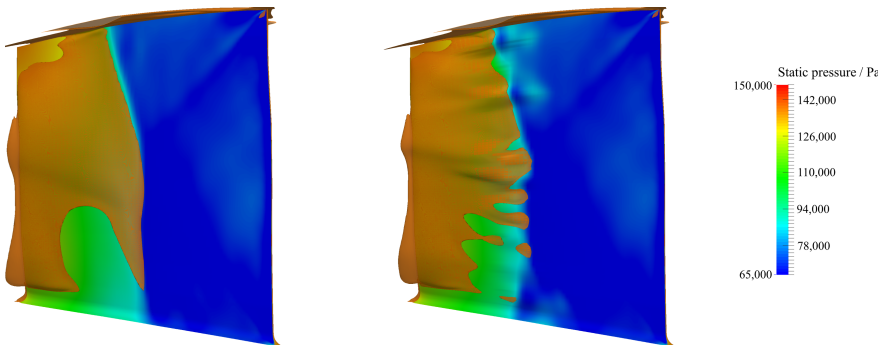

Fig. 5.5 Datum (left) and individual bump optimum (right) static pressure contours with separation shown by zero axial velocity isosurface.

Figure 5.5 shows the change to the separation produced by the optimised design. It can be seen that the bumps have delayed separation and affected the shock, but the separation behind some of the bumps has been increased. The bump amplitudes are large enough that their individual effect on the separation (sometimes increasing it) can be seen.

The efficiency difference is given in table 5.2. It can be seen that efficiency has increased by 0.76% and the pressure ratio has also been slightly increased.

Table 5.2 Individual optimised bump performance comparison.

	Pressure ratio	**Pressure ratio benefit / %**	**Efficiency / %**	**Efficiency benefit / %**
Datum	2.06		85.45	
Individual	2.07	0.5	86.21	0.76

The individual bumps have increased the efficiency, and also the pressure ratio, but the bumps are clearly disrupting the flow and in some regions increasing the separation. To create the benefit seen the amplitudes of the individual bumps have become excessively large. This could be restricted in the optimisation set up, but at the cost of reduced efficiency improvement. In Chapter 4 it was demonstrated that the individual bumps also required larger amplitudes than the continuous bumps to have a positive effect. It is investigated next whether the continuous bumps can be shown to be more beneficial here.

Continuous bump optimisation

For the continuous bump optimisation the bump parameters are specified at the same radial control heights as the individual bumps, but the geometry is smoothly interpolated radially using a cubic b-spline.

The optimised, continuous bump design is compared to the datum in Figure 5.6. The datum shock position is shown via a white line on the optimised geometry. It can be seen how the use of bumps has delayed the shock. The benefit is also greater than produced by the individual bumps.

The reduction in separation caused by the optimised design can be seen in Figure 5.7. The delay of the shock position has delayed the separation point and the volume of separated flow is therefore reduced.

Comparing this to the individual bump design it can be seen that it has delayed the shock and reduced the separation further. The continuous bumps also do not seem to cause localised increases of separation, as was shown for the individual ones.

Table 5.3 shows a comparison of the performance of the datum, optimised continuous and optimised individual bumps. It can be seen that the predicted efficiency benefit is greatest for the

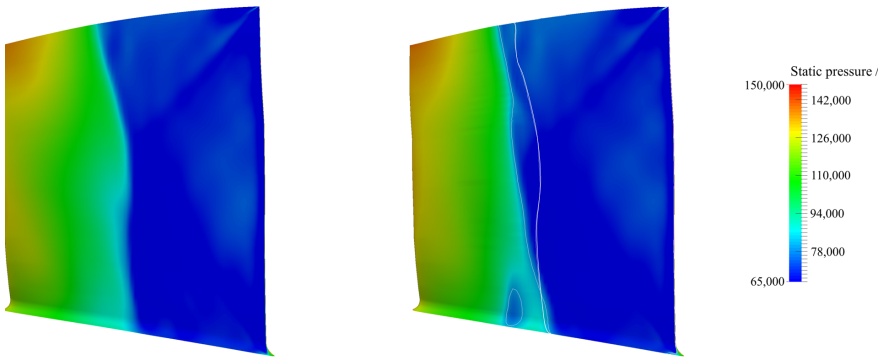

Fig. 5.6 Datum (left) and continuous optimised (right) static pressure contours.

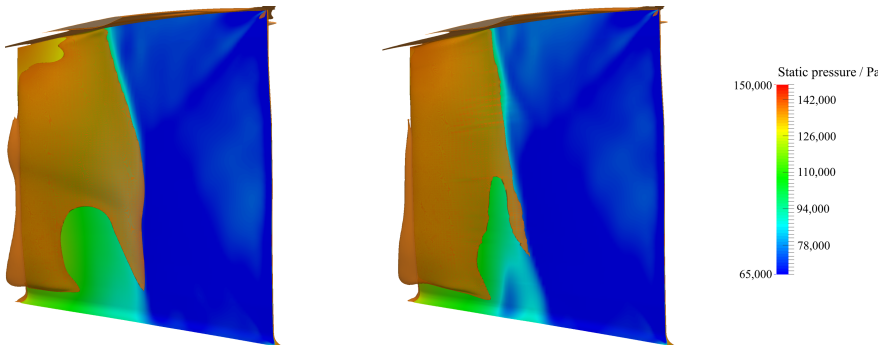

Fig. 5.7 Datum (left) and continuous bump optimum (right) static pressure contours with separation shown by zero axial velocity isosurface.

continuous bump design. The efficiency is increased by 1.48%, a significant amount considering the only change to the geometry is the addition of the bump.

Table 5.3 Continuous optimised bump performance comparison.

	Pressure ratio	**Pressure ratio benefit / %**	**Efficiency / %**	**Efficiency benefit / %**
Datum	2.06		85.45	
Individual	2.07	0.5	86.21	0.76
Continuous	2.10	1.84	86.93	1.48

5.3.3 Analysis of the optimised bump design

The most beneficial bump design is the continuous bump geometry seen in Figure 5.6. Here the bump geometry and impact on the flow is analysed. The geometry of the shock bump can be seen in more detail in Figure 5.8. The 3D geometry compared to the datum is shown in Figure 5.9.

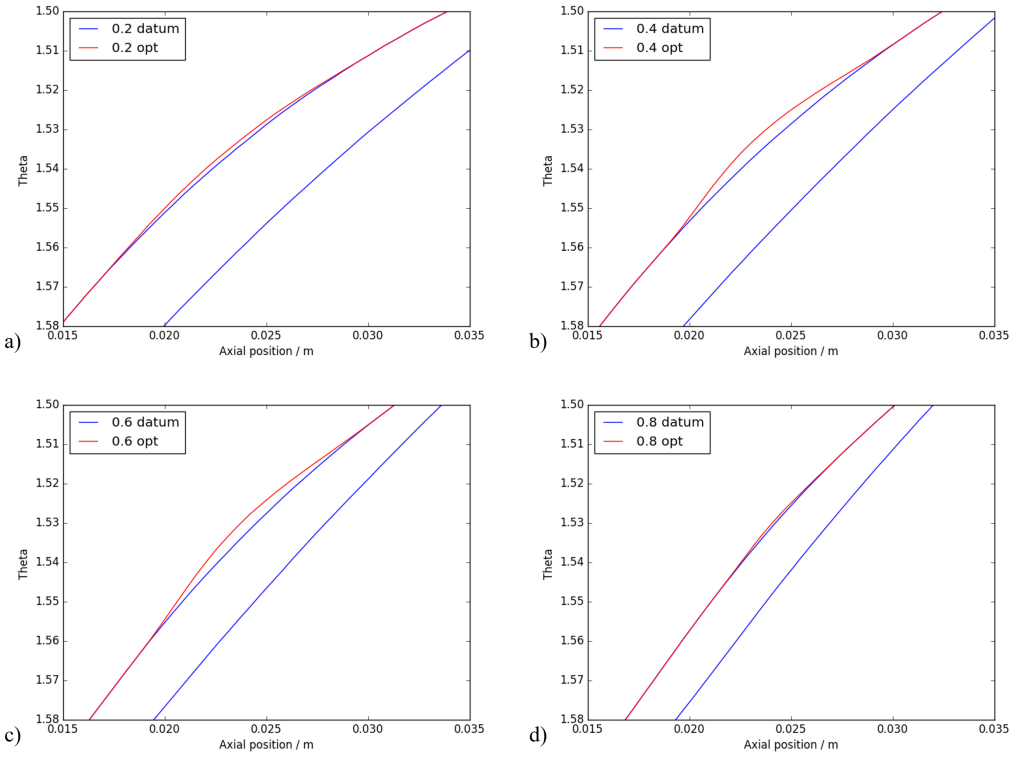

Fig. 5.8 Spanwise slices of the datum and optimised geometries at (a) 20%, (b) 40%, (c) 60% and (d) 80% span.

It can be seen how the bump applied to the datum geometry varies radially, with the maximum height of the bump localised around mid-span. This makes sense as the strongest shock location,

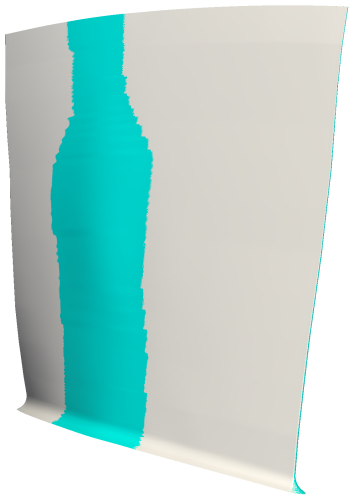

Fig. 5.9 Optimised continuous bump (blue) added to the datum blade geometry (grey).

greatest separation and adjoint sensitivity occur around mid-span for Rotor 37, as discussed in Chapter 3.

Figures 5.10 and 5.11 show the passage flow for the datum and optimised geometries at 50% span. The effect of the bump weakening and delaying the shock can be seen, with the datum shock position shown by the black line. The Mach number contour is lower just upstream of the shock, and the low velocity wake region has clearly reduced in size.

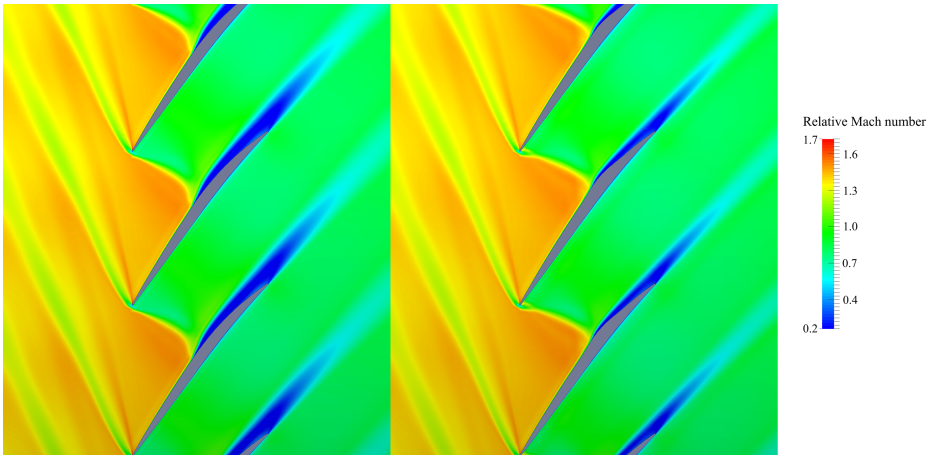

Fig. 5.10 Datum (left) and continuous optimised (right) flow features at 50% span.

Figure 5.12 shows the datum and optimised lift plots. It can be seen how the shock has been delayed through the span. At 60% span that the shock has been weakened, with some pre-compression caused by the bump.

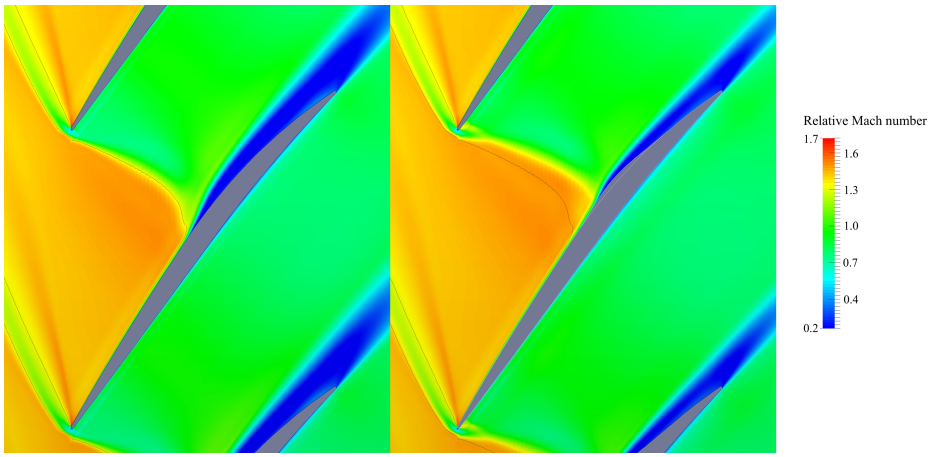

Fig. 5.11 Datum (left) and continuous optimised (right) flow features at 50% span (zoomed view).

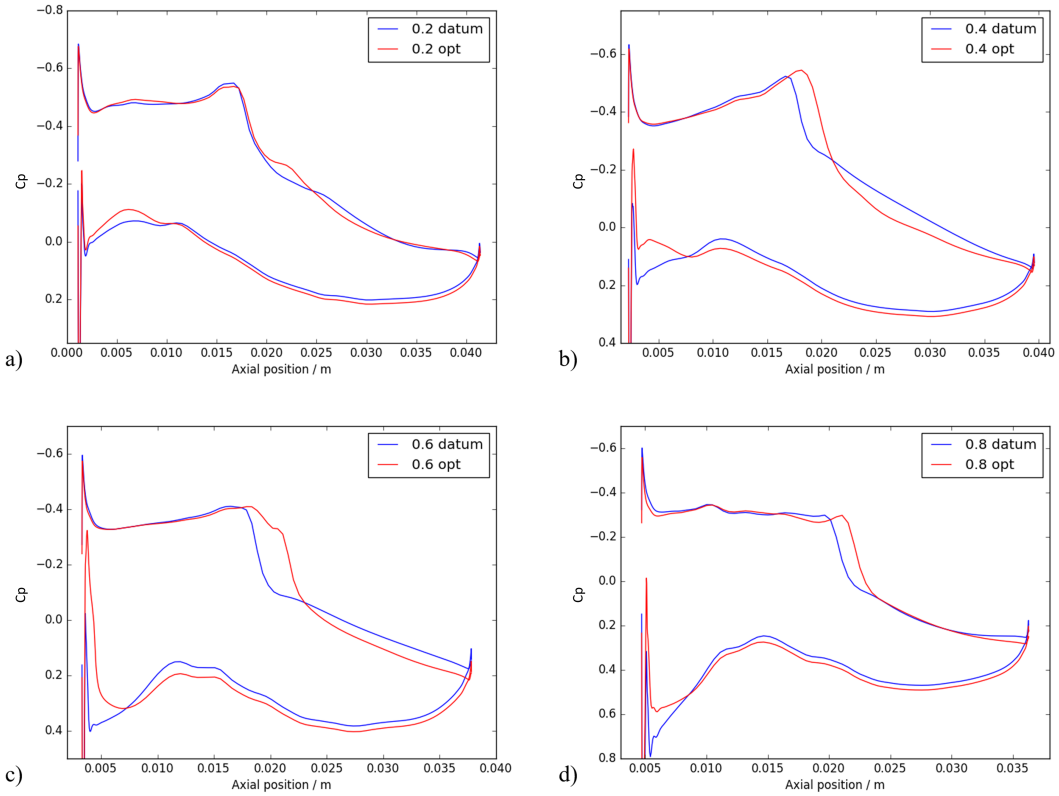

Fig. 5.12 Lift plots for the datum and optimised geometries at (a) 20%, (b) 40%, (c) 60% and (d) 80% span.

5.3.4 Comparison to the FFD optimised design

Figure 5.13 shows datum, FFD optimised and continuous bump optimised NASA rotor 37 geometries.

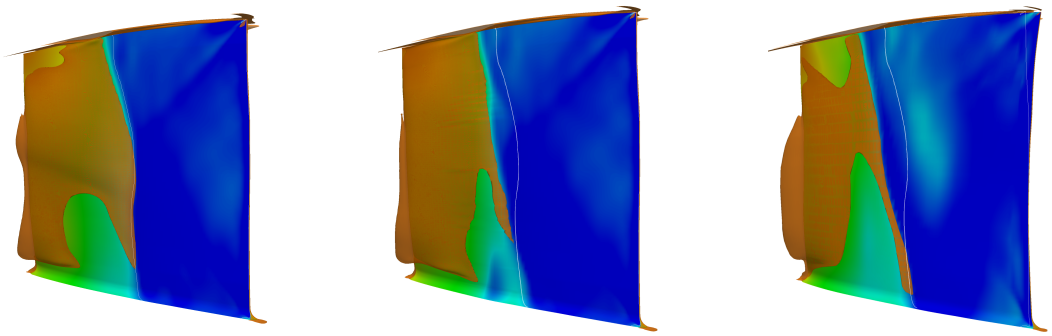

(a) (b) (c)

Fig. 5.13 (a) Datum geometry, (b) continuous bump geometry and (c) s-shaped geometry with separation in orange.

It can be seen that the s-shape and bump geometries both delay the shock in a similar way, but the larger extent of s-shape geometry modification allows it to have a larger impact. Table 5.4 compares the performance benefits.

Table 5.4 Comparison of performance of the three designs.

	Pressure ratio	**Pressure ratio change / %**	**Efficiency / %**	**Efficiency change / %**
Datum	2.06		85.45	
Continuous bump	2.10	1.84	86.93	1.48
S-shape	2.09	1.46	87.44	1.99

The FFD design offers greater efficiency improvement. The continuous optimised geometry has made a comparatively significant increase in efficiency, considering only the local blade surface was modified, whereas the FFD parameterisation was able to perturb the entire geometry.

5.3.5 Off-design performance

The off-design performance is a key consideration. The characteristics for the datum and optimised designs are shown in Figure 5.14.

It can be seen that an efficiency and pressure ratio increase has been predicted across the characteristic. The choke mass flow does not appear affected, although it is possible that the choke margin has been modified at other rotor speeds due to the throat area being reduced.

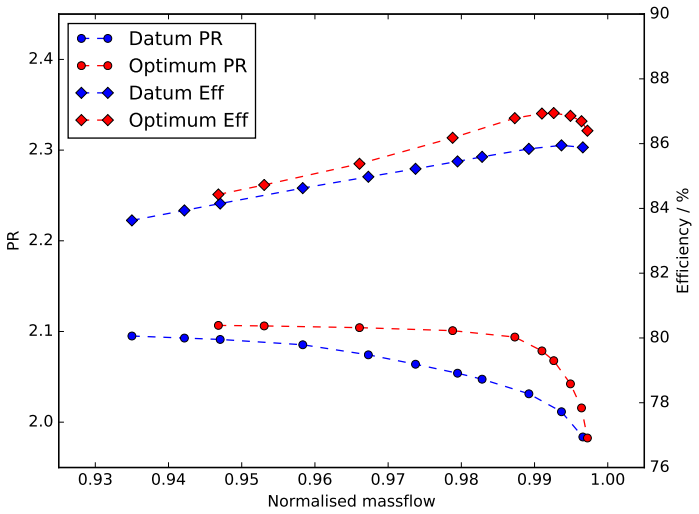

Fig. 5.14 R37 optimised characteristic vs datum.

The simulation results suggest a reduction in stall margin for the optimised design. This is likely due to the shock bump being mis-placed at conditions away from where it was designed, leading to increased separation and thus reduced stall margin. This is a concern and could render such a design using shock control bumps unusable, as stall margin is a key consideration for blade design.

This section has demonstrated the benefit that can potentially be achieved by applying shock control bumps to a compressor blade without modifying the entire blade geometry. This shows that a significant benefit may come from applying geometry modifications through bumps just in the shock region, but their performance away from the design point could limit their usage.

5.4 Fan blade shock control

In this section, the benefit that shock control can provide in extending the stall margin of fan/compressor blades is investigated. The case used is a modern, very high bypass ratio, low speed fan blade.

5.4.1 Fan blade design

The fan blade used in this investigation is known as RR-FAN-2. It is the latest generation of high bypass ratio, low speed fan designs that are being developed by Rolls-Royce as they work towards designing the fan blades that will be used in the Ultrafan engine. Further investigations, with more detail of the CFD set up and validation of this and another fan blade are provided in Chapter 6.

5.4.2 Set up

The blade is simulated in rotor only format, with a downstream splitter geometry and separate exit boundaries for the core and bypass flows. An outline of the CFD domain is shown in Figure 5.15. For confidentiality reasons the actual geometry of this fan blade cannot be shown. At the inlet a radial distribution of total pressure and temperature is specified and at the exit boundaries radially averaged mass-meaned non-dimensionalised flow rate (capacity ($\frac{\dot{m}\sqrt{T_{tot}}}{P_{tot}}$)) is specified.

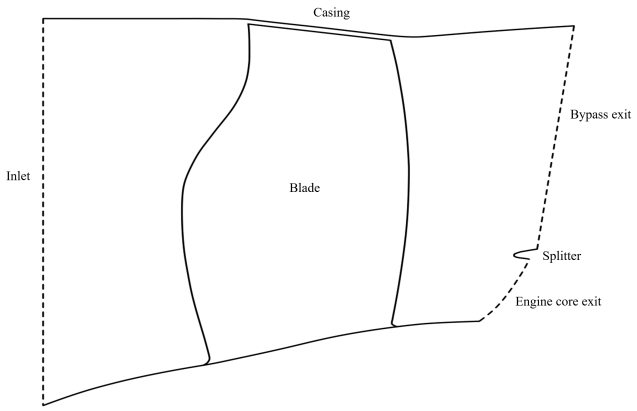

Fig. 5.15 The domain used for RR-FAN-2 (not representative of the actual fan geometry).

5.4.3 Blade flow features

To understand the behaviour of this blade and select a point at which to optimise the design, the flow behaviour for a range of flow rates was studied (see Figures 5.16 and 5.18). The blade characteristic for a range of flow rates at cruise speed is shown in Figure 5.16. The design point is point G. A stalled point (point A) is included and indicated by the dashed line. Point B is therefore the last stable point before stall. As this blade is still in development, there is no experimental data for validation. A related blade, RR-FAN, is used in Chapter 6, and is validated against experimental data (see section

6.5). The set up is very similar between the two cases and hence the simulation accuracy can be trusted.

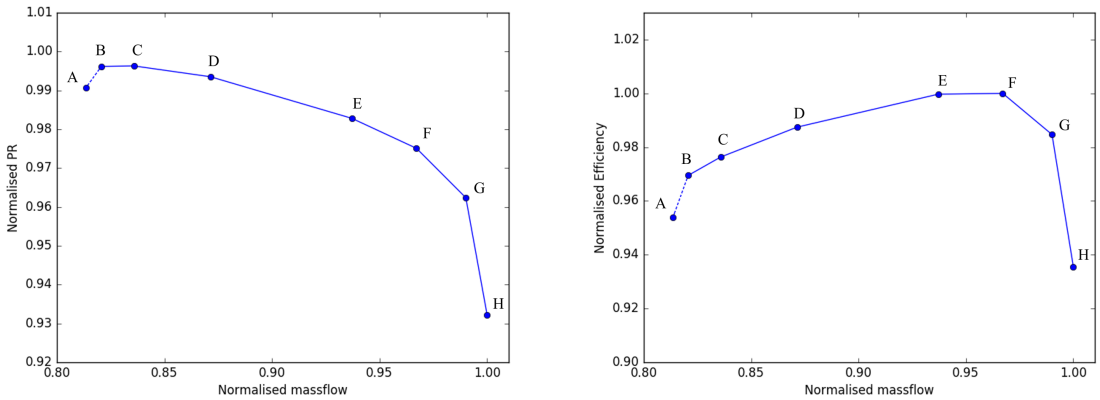

Fig. 5.16 RR-FAN-2 PR (left) and efficiency (right) characteristics.

For proprietary reasons the whole blade geometry of RR-FAN-2 cannot be shown, hence, flow behaviour in just the region of interest is shown in the following figures. Figure 5.17 describes this region of interest.

Figure 5.18 shows the flow features of the blade design as the flow rate is varied (as shown in Figure 5.16). The shock position on the blade surface moves towards the LE at stall. As the pressure ratio increases and flow rate becomes lower, the strength of the shock increases and separation is caused towards stall. It is this separation (in orange) that contributes to the full stall of the blade. It can be seen that the shock-induced separation increases in magnitude and radial extent as the flow rate is lowered, until full separation eventually occurs.

These near-stall operating points are a promising area to investigate the benefit of shock control bumps. It is the shock-induced separation that is responsible for limiting the operating range of the blade, and if this separation can be reduced then it may be possible to extend the stable working range of this fan.

Adjoint sensitivity analysis

The adjoint sensitivity analysis for RR-FAN-2 at point D (see Figure 5.16) is given in Figure 5.19. It can be seen that the most sensitive regions of the geometries are focussed around the shock on the suction surface. This indicates that geometry changes in this region will have a significant impact on the blade efficiency, and therefore if shock control bumps are applied here some benefit should be found.

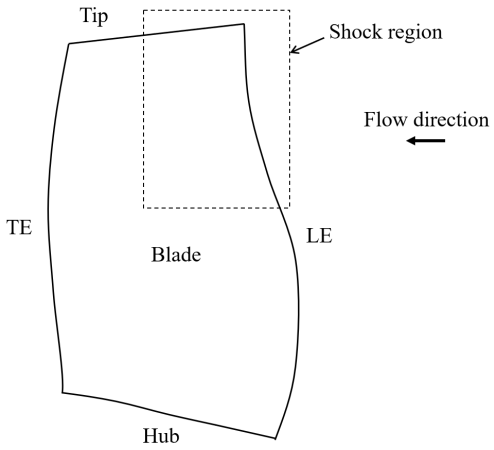

Fig. 5.17 The region of interest (shock region) presented in further figures (not representative of actual geometry).

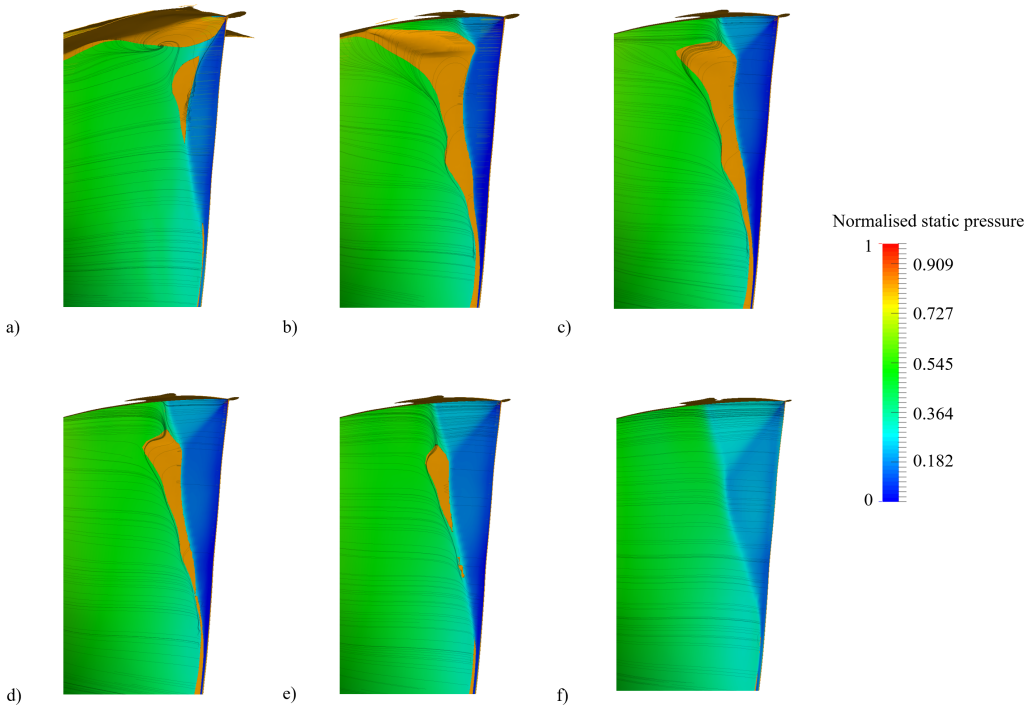

Fig. 5.18 Shock region flow features at points a) A, b) B, c) C, d) D, e) E, f) F. Flow direction right to left. Separation shown by zero axial velocity isosurface (orange).

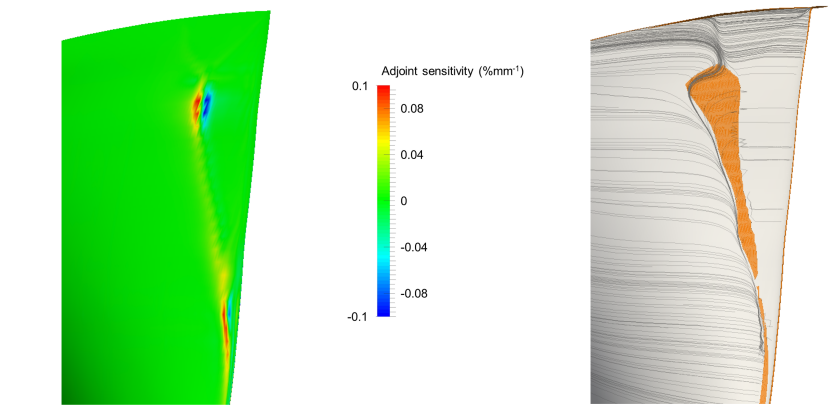

Fig. 5.19 RR-FAN-2 shock region separation shown by zero axial velocity isosurface (orange) and adjoint surface sensitivity near to stall (point D). Flow direction right to left.

5.4.4 Optimisation of shock-control bumps

The optimisation approach for this fan blade is similar to that used in the Rotor 37 study described above, except in this case the shock bumps are only used in the top 40% of the span. This is because the shock (that is critical to the blade aerodynamics near stall) is located here. The region within which the bumps are added is shown in Figure 5.20. These are positioned so that they can control the shock, and potentially reduce shock-induced separation. Both individual and continuous bumps were tested.

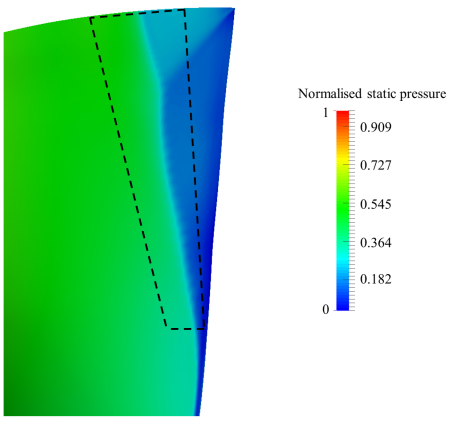

Fig. 5.20 Static pressure contour on the suction surface and the region within which bumps are positioned. Flow direction right to left.

Individual bumps

Five control heights were used for the optimisation of the individual bumps, at 60-100% span in 10% intervals. The intitial bump positions were set with 60% of the bump downstream of the shock, and could vary as is shown in table 5.5.

Table 5.5 The RR-FAN-2 bump parameter ranges.

Spanwise bump position / Span	**Minimum start point / chord**	**Maximum start point / chord**	**Minimum end point / chord**	**Maximum end point / chord**
0.5	0	0.1	0.1	0.25
0.6	0	0.15	0.15	0.3
0.7	0.05	0.22	0.22	0.4
0.8	0.1	0.25	0.25	0.5
0.9	0.1	0.25	0.25	0.5
1.0	0.15	0.36	0.36	0.6

The optimisation was carried out using the MAM optimiser at condition D shown in Figures 5.16 and 5.18. The objective function is the blade efficiency. The flow at this operating point can be seen in Figure 5.19. It is at a mass flow slightly higher than stall conditions so that the flow is stable, but close enough that shock-induced separation is present and any improvement at this design point should in theory improve stall margin. The optimised bump geometry and the effect of the bumps on the shock and separation is shown in Figure 5.21.

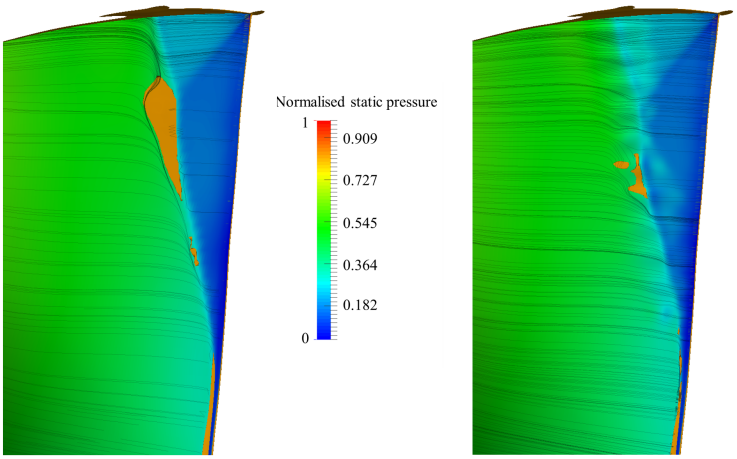

Fig. 5.21 The datum (left) and optimised individual (right) geometries with separation shown in orange. Flow direction right to left.

Figure 5.21 shows some benefit is predicted from the optimised bumps. It can be seen however that the individual nature of the bumps leads to an inconsistent benefit, and in some regions an increase in separation.

Continuous bumps

The optimisation using continuous bumps uses the same set up as the individual bumps, except a smooth radial interpolation approach is used between the specified bump sections. The result of the optimisation at operating point D is shown in Figure 5.22.

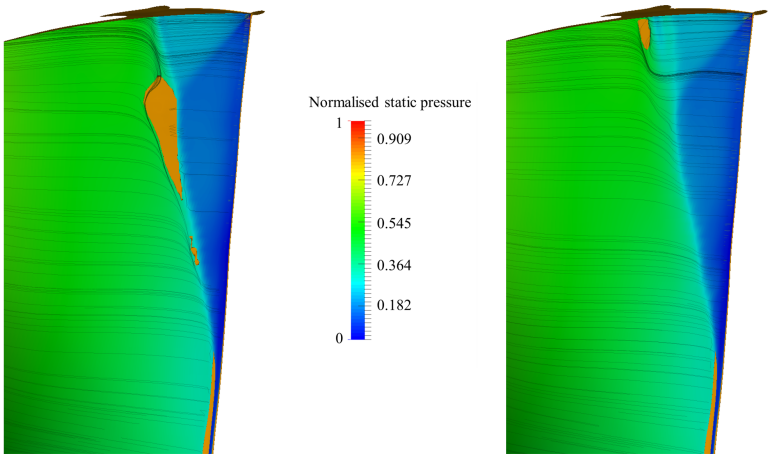

Fig. 5.22 The datum (left) and optimised continuous (right) geometries with separation shown in orange. Flow direction right to left.

It can be seen that the bumps significantly reduce the shock induced separation, almost eliminating it. Towards the tip a small separation region is present. This is likely due to the optimiser getting stuck in a local minima and not finding a globally optimum design. This is one of the limitations of gradient-based optimisation methods. The efficiency benefit for the optimised continuous geometry is greater than for the individual bump geometry, and therefore this is analysed here in more detail.

Figure 5.23 shows slices at 80% span of the datum and optimised continuous bump geometries. It can be seen how the addition of the bump has weakened the contours across the shock, and reduced the dark blue, low-momentum regions of the boundary layer, separation and wake.

It can be seen in Figure 5.24 how the wake of the blade at this height has been reduced in width, by over 15%. The maximum relative Mach number deficit has also been reduced. The result is that the optimised design has lower loss generation due to the wake, increasing its efficiency. This can be seen in Figure 5.26.

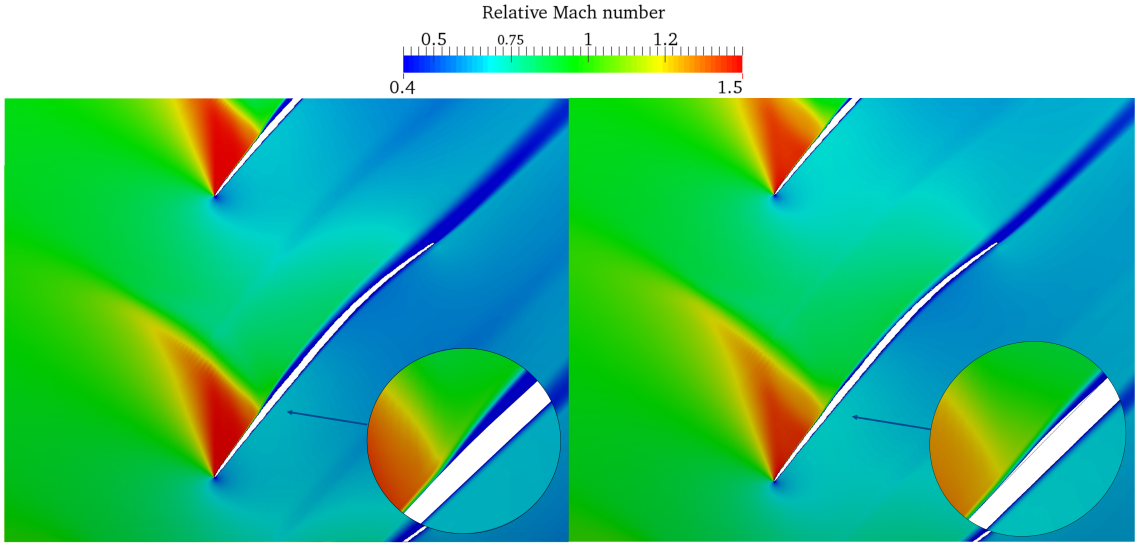

Fig. 5.23 Relative Mach number contours at 80% span showing the datum (left) and optimised continuous (right) geometries at operating point D.

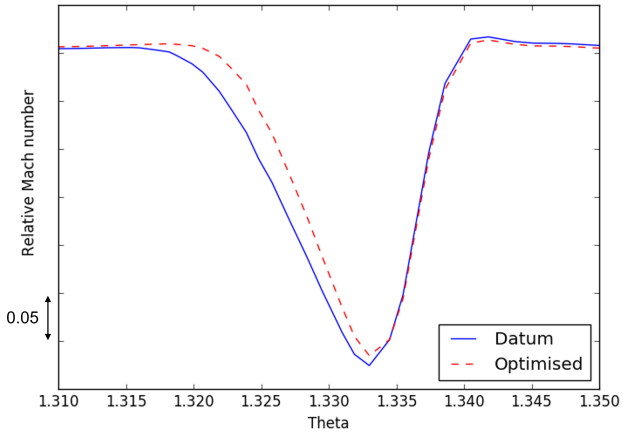

Fig. 5.24 Blade wakes for the datum and optimised geometries measured at 80% span and 0.1 chord downstream of the trailing edge.

Analysis of the stall margin and characteristic

The application of shock control bumps was predicted to reduce shock induced separation at the near-stall condition (at which the optimisation was carried out). The objective for this case was to increase the stall margin via the use of these bumps. To assess whether this has been achieved, the optimum design was simulated at a variety of flow rates. The variation in blade behaviour at a range of flow rates near to stall is shown in Figure 5.25. It can be seen that the size of the separation is reduced at each of these operating points for the optimised design.

As shown in Figure 5.26, the presence of the bump allows an operating point to the left of B to remain un-stalled, whereas full stall occurred at this point for the datum geometry. The bump designs reduce separation and delay the point at which the blade stalls. The impact of the added bumps for the rest of the fan operating range must also be assessed however (see Figure 5.26).

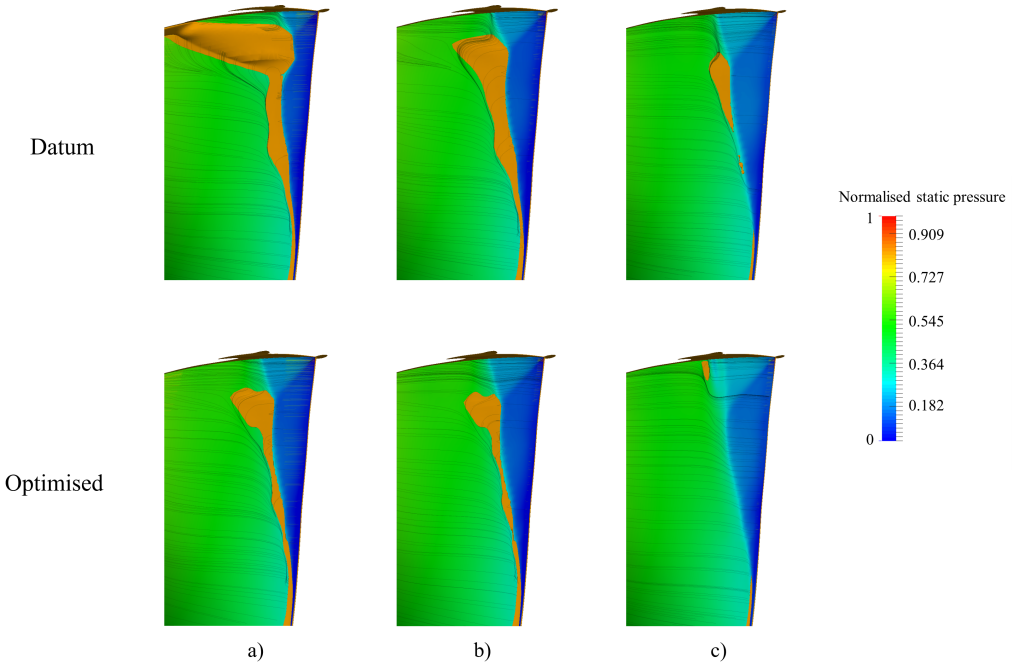

Fig. 5.25 Flow separation near stall for the datum and optimised continuous designs at a) B, b) C and c) D operating conditions. Flow direction right to left.

The optimised design has increased efficiency, and also stall margin at the left of the characteristic, but unfortunately in doing so has adversely affected the behaviour near to the design point. The bump that is designed to operate successfully near to stall results in unwanted flow behaviour at other flow rates. As the shock position moves (as the flow rate varies), the bump is no longer well aligned with the shock and has a negative impact. This effect can be seen in Figure 5.27 on the choke side of the characteristic. The bump causes a region of extra acceleration resulting in a stronger passage shock compared to the datum geometry. This increases entropy generation and results in lower efficiency on

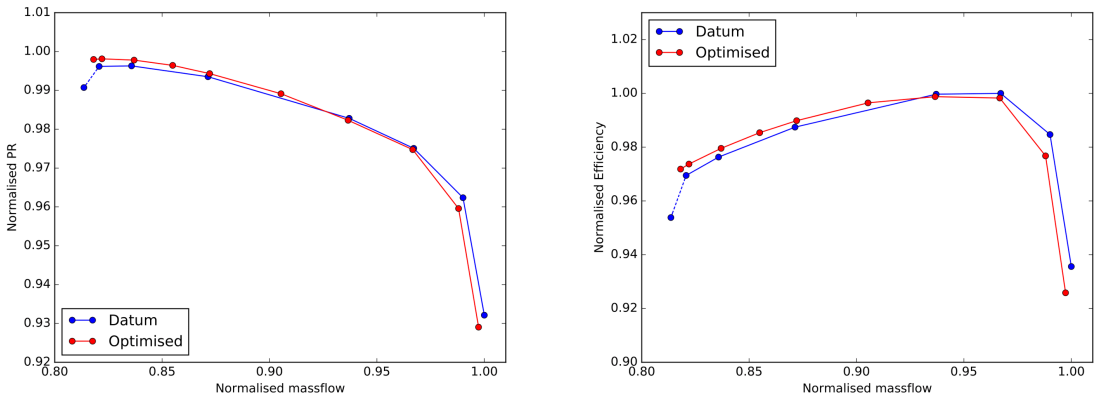

Fig. 5.26 PR (left) and efficiency (right) characteristics for the datum and optimised designs.

the right hand side of the characteristic. The flow capacity on the choke side of the characteristic is also affected (as can be seen in Figure 5.26) due to the impact of the bump at this flow rate.

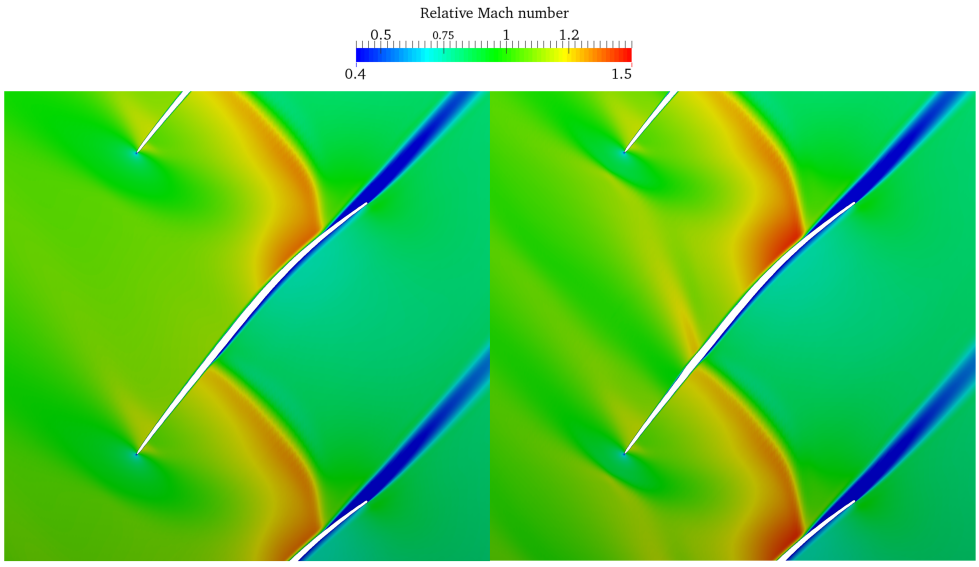

Fig. 5.27 Relative mach number contours at 80% span (at operating point F), showing the datum (left) and optimised (right) geometries.

To assess whether further benefit in the near stall region could be achieved, the optimisation process was repeated at point B on the RR-FAN-2 characteristic. This point is far closer to stall for the datum geometry, and as a result, the optimised design produced outperformed the datum and previous optimised design in the stall region. This can be seen in Figure 5.28. As occurred previously however, improving the performance in the stall region had a negative impact nearer to the design point. The choke mass flow rate is reduced, and the efficiency also, for points on the right hand side

of the characteristic compared to the datum. This highlights the impact that the operating point of the optimisation has on the resulting design behaviour.

It must be noted that assessing the impact of these designs on fan stall margin using steady-state analysis has limitations, due to the truly unsteady nature of stall-inception. It has been demonstrated here that the application of shock control bumps can reduce the separation that appears to trigger fan stall, but an assessment using unsteady simulations would be desirable to verify this benefit.

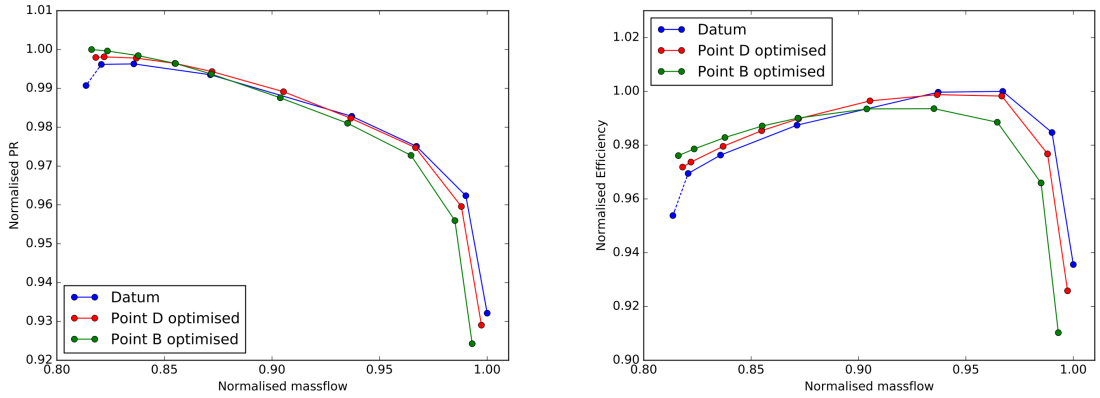

Fig. 5.28 Characteristics for the datum, point D optimised and point B optimised geometries.

5.5 Conclusions

This work has demonstrated how shock control can be used to reduce the loss generated by shock-induced separation for a highly loaded compressor blade and reduce the shock induced separation that appears to trigger the stall of a modern fan blade.

These results highlight the benefit that can potentially be achieved through the use of shock control bumps on rotor blades, but also the limitations with their usage. Single-point design of these bumps leads to geometries that perform well under one condition, but whose performance can deteriorate as the operating condition changes. This may limit their potential real-world usage, as the negatives associated with poor off-design performance outweigh the benefit of their design. Further investigation would therefore be beneficial. Multi-point optimisation may help to find designs with acceptable performance across the blade operating range.

5.6 Summary

This chapter has demonstrated that shock control bumps have potential to improve transonic fan/compressor blade performance, but their usage may be limited by poor off-design performance. Chapter 6 investigates another key aerodynamic feature of fan blade aerodynamics, the tip clearance.

Chapter 6

Fan blade tip aerodynamics with realistic, operational casing geometries/clearances

In this chapter, a detailed understanding of the tip flow physics for a modern fan blade is developed. The tip leakage flow structure, tip leakage mass flow distribution and tip vortex properties are analysed. For the datum geometry, it is shown that an unusual tip leakage distribution exists, which has a large influence on the local aerodynamics. The passage shock sitting on the pressure side of the blade causes the tip leakage flow to partially reverse, travelling from suction to pressure surface for part of the chord. The variation in blade performance and tip flow features due to uniform changes in tip clearance are then discussed in detail.

The impact of varying casing geometries with different tip clearances, trench depths and sharp steps is then investigated using a parametric model. The analysis allows a detailed understanding to be developed of the impact that realistic casing geometries have on the tip region flow. The effect of cuts in the casing profile compared to cropping of the blade tip is investigated at a range of effective clearances. Casing geometries generated using a range of permutations of the model parameters are simulated. Response surface analysis of the results is then able to predict the relationship of each parameter to blade performance. The importance of each parameter is found, allowing an understanding of the effects of trenches and the various clearance geometries that exist in real engines to be developed.

6.1 Background and motivation

Fan blade tip clearances vary significantly throughout each flight cycle and over an engine's lifetime. Several mechanisms lead to this: differing expansion rates of components during start-up, the non-

circular shape of the casing during flight, and gusts, which cause the core and nacelle to shift relative to each other. As a result, the fan tip and casing regularly make contact. To maximise the fan blade performance and reduce the engine fuel consumption, tip clearances are designed to be as small as possible. This reduces leakage flow through the gap between the tip of the blade and the stationary casing which, upon entering the blade passage, generates a significant amount of loss and reduces blade performance [37, 120, 112]. To allow minimum clearances while also preventing damage to the blade due to contact with the casing, a 'tip rubbing' solution is used, where a sacrificial abradable liner is added to the casing. At certain points during the engine cycle, the liner becomes worn by the blade tip. At the extreme points of the blade's movement, steps in the casing liner are created, forming a trench with non-uniform clearance above the blade. The result of these effects is a casing profile with steps just upstream of the leading edge and just downstream of the trailing edge, and clearances larger than initially set by the designer.

Kammer et al. [62] described how the repeated contact of the blades with the fan casing is analogous to the teeth of a milling machine, cutting away at the liner. The mechanical implications of the complex rubbing process (with friction-induced wear, heating and the dynamic coupling of the blade and the elastic casing) are well discussed in the literature from an abrasion and dynamic perspective [72, 43, 70]. The effect of this casing rubbing on blade aerodynamic performance is not well analysed however. Realistic casing profiles, with steps up and downstream of the blade are not generally investigated. The trend of increasing overall clearance reducing blade performance and operability is well known [37, 120, 112], however work in this area focuses on varying clearance with simple crops of the blade tip. This chapter aims to develop a detailed understanding of a modern fan blade's tip aerodynamics for datum clearance, before investigating the impact that variations to the casing profile caused by rubbing have on its performance.

6.2 Objectives

The aim of this chapter is to provide an understanding of modern fan blade tip aerodynamics so that design rules for setting the build clearances of future Rolls-Royce engines can be based on a more sound footing than in the past. When an engine is first put into operation, the liner must be bored by a boring machine to set the cold build clearance. This is the clearance between the abradable liner and the blade tips when the engine is not running (in cold condition). As the engine runs up for the first time and over its many cycles, the blade tips cut into the abradable liner. The starting, cold build clearance that is set therefore determines how much rubbing of the liner will occur. Setting a tighter cold build clearance will generally reduce the operating clearance, however this increases the depth of rub into the liner and, as a result, the depth of trenches that are produced. A lifetime cost analysis is therefore required to determine the optimum cold build clearance to be used. This analysis requires

an understanding of the drop in fan efficiency (and therefore increase in engine SFC) that will occur as the fan clearance increases and trenches are developed.

Previously, the prediction of fan performance drop with engine lifetime due to clearance increase has been based on a very simple relationship. This did not account properly for the true variability of the clearance geometry and the grooves that are cut into it. This work aims to produce an improved model based on a detailed understanding of the fan blade tip aerodynamics that can be used to more accurately predict the drop in efficiency due to variations in engine lifetime clearance geometry.

The objectives of this work are therefore to:

- Understand the tip clearance aerodynamics of a modern fan blade
- Understand the impact of trenches on the tip clearance aerodynamics
- Develop a model that can predict fan efficiency variation due to changes in clearance size and geometry

6.3 Literature review

6.3.1 Tip leakage flow

The tip leakage vortex is common to all rotating blades with a clearance. The basic mechanism involves flow from the pressure side of the blade being driven through the clearance by the pressure gradient to the suction side. As this leakage flow exits the clearance and interacts with the passage flow, the difference in velocities causes a vortex to form. The vortex usually begins at the leading edge and passes across the suction side passage towards the adjacent blade. This low momentum flow causes blockage in the tip region of the passage, reducing the mass flow that can pass through. As the vortex mixes out it generates significant losses, and hence tip clearances are designed to be as small as possible. Figure 6.1 shows the process of tip leakage vortex development.

A number of researchers have investigated tip clearance effects for fans and compressors. Adamczyk et al. [2] demonstrated the role of tip clearance in high speed fan stall, showing the reduction in stall margin, pressure ratio and efficiency caused by an increase in clearance. They also described how the interaction of the leakage vortex with the in-passage shock plays a major role in determining the fan flow range. Denton [37] developed a model for predicting the loss generated through turbomachine clearances based on the properties of the leakage and passage flows. He described how the entropy generation by the tip leakage vortex is proportional to the difference in streamwise velocity of the two flows. A CFD study of compressor efficiency variation from vanishing to large clearance is provided by Sakulkaew et al. [101]. They showed how the loss variation with tip gap is linear for clearances between 1 and 3%. Seshadri et al. [105] and Beheshti et al. [8] also discussed the linear variation of blade efficiency with tip clearance for axial compressors. Some examples of blade performance vs

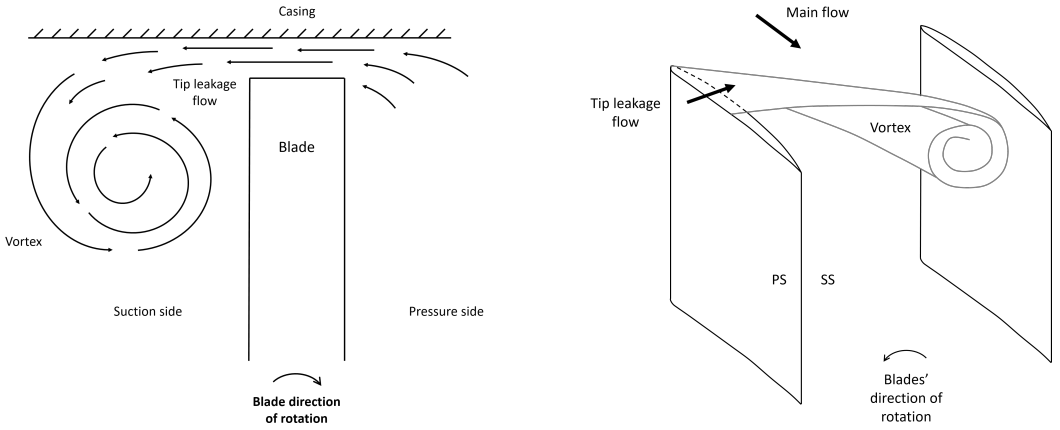

(a) (b)

Fig. 6.1 Tip leakage vortex formation (a) view from upstream, (b) 3D view.

clearance magnitude (showing linear trends) are given in Figure 6.2. Both Seshadri et al. [105] and Sakulkaew et al. [101] found in their simulations however that close to zero clearance the behaviour is non-linear and an 'optimum' clearance is found. This is due to increased shear losses at the casing as the clearance becomes small [101]. These studies all focused only on cropping the blade tip uniformly with no changes to the casing streamline.

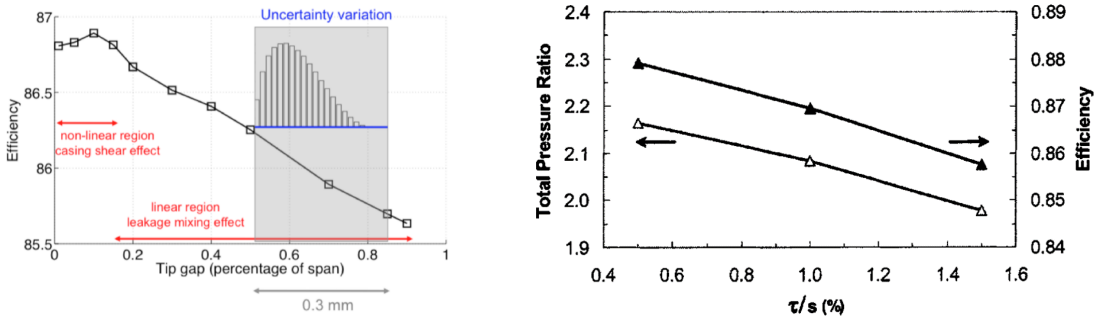

(a) (b)

Fig. 6.2 Linear performance variations with clearance (a) Seshadri et al. [104], (b) Beheshti et al. [8].

The main discussion of casing geometry variations in the literature is related to passive casing treatment for improved blade performance. These investigations use several discrete axial or circumferential grooves in the casing above blade tips to 'trap' the leakage flow and reduce its impact, improving stall margin or efficiency [106] [55]. Qin et al. [89] demonstrated how casing groove depths and angles can be optimised to provide greater performance benefit. Their optimised configuration of casing grooves can be seen in Figure 6.3. These casing grooves for passive tip treatment are far larger than the blade clearance however (unlike trenches which are of similar magnitude), and the behaviour of these multiple discrete grooves is not applicable to this work.

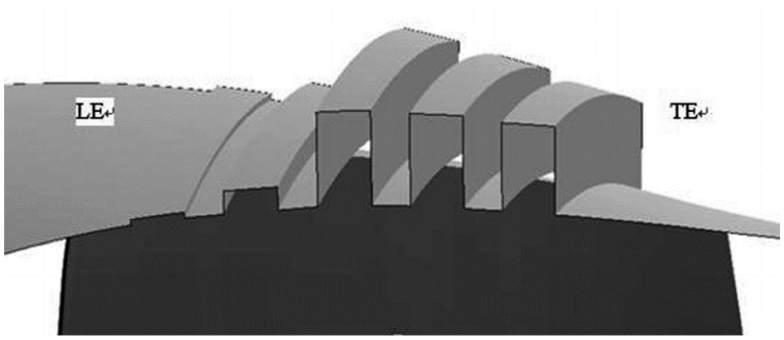

Fig. 6.3 Optimised casing grooves for NASA Rotor 37 (from Qin et al. [89]).

Investigations into the effect of single, circumferential trenches are provided by a few researchers for compressor blades. Robideau et al. [95] state that a performance benefit can be achieved by allowing the airfoil tips to sit above the full height of the flow path. This is achieved using a trench and allowing the blades to extend further radially (see Figure 6.4a). The idea is to move the leakage flow out of contact with the passage flow, and thus reduce losses. Korting et al. [64] tested the impact of 'trenched' clearances above the tip of a compressor blade in more detail. Experimental results describe the impact of various depths of trench and upstream and downstream shapes. It was shown that an important feature was the forward facing step at the trailing edge side of the trench. Increased depth of the trench produced more loss because of the increased forward facing step. The trench only provided an efficiency benefit when the step at the trailing edge was replaced by a sloped geometry. It was concluded that trenches did not alter the sensitivity of the compressor to changes in clearance.

Nezym [81] investigated the impact of 'entire annular recess casing treatment' on the stability of compressor rotors. It was found that an increased stable operating range of the compressor was obtained due to the presence of the annular entire recess. The larger the recess depth and volume, the greater the improvement. The improvement was especially effective when the blade tip coincided with the baseline casing. It was also shown how a recess of triangular configuration and relatively small dimensions may be useful for efficiency improvement. The range of geometries tested by Nezym can be seen in Figure 6.4b.

Beheshti et al. [8] conducted a numerical parametric study of tip clearance coupled with casing treatment. Figure 6.5 shows how the casing treatment proposed resembles the trenches under investigation in this work. Such a casing treatment provides a means for fluid to exit the flow path where the blade loading is severe, migrate circumferentially, and re-enter the flow path at a location where the pressure is more moderate [8]. This leads to stability improvement since the flow relocation relieves the locally severe blade loading. Using this device, the authors showed the possibility to improve both the efficiency and stall margin.

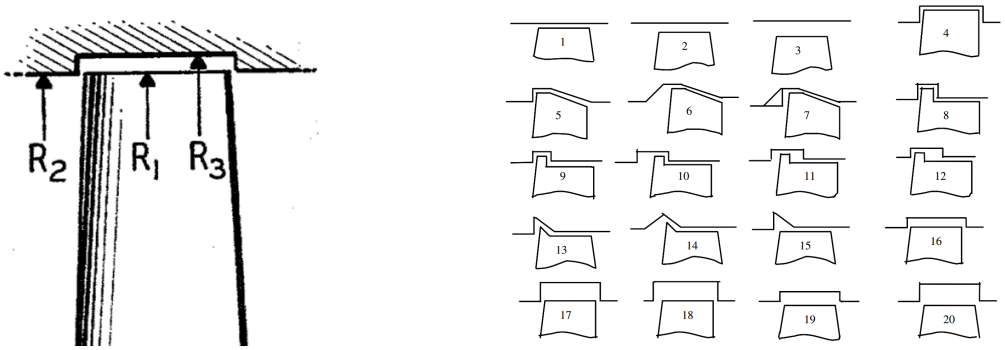

(a) (b)

Fig. 6.4 Some trench geometries previously investigated (a) casing treatment by Robideau et al. [95], (b) Recessed clearances tested by Nezym [81].

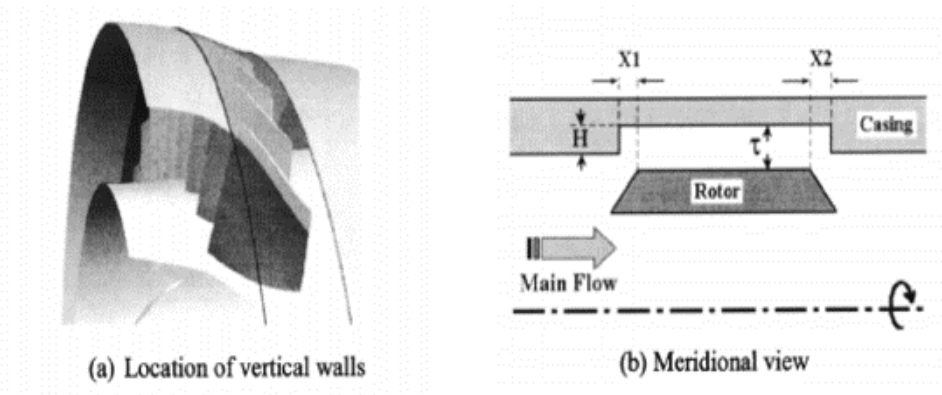

Fig. 6.5 Tip clearance casing treatment by Beheshti et al. [8].

The trenches described above are of far greater spanwise extent than those rubbed by the tip of a modern fan blade. The effects of clearance also differs in a multistage compressor compared to fan blades as tested here. These results are therefore informative, but the conclusions may not be directly applicable to the investigation carried out here.

6.4 Problem definition

A demonstration of the variation in the relative blade and casing positions that cause rubbing of the abradable liner can be seen in Figure 6.6. This shows how the blade's position can vary in the axial and radial directions and the resulting liner profile that is formed after many cycles. In reality the trench step edges will be slightly rounded, but testing of abraded liners from Rolls-Royce engines has confirmed the steps formed have 90 degrees corners and are reasonably 'sharp', so it is appropriate to model them this way. The position of the blade at cruise and a comparison of these features to the datum liner profile can also be seen.

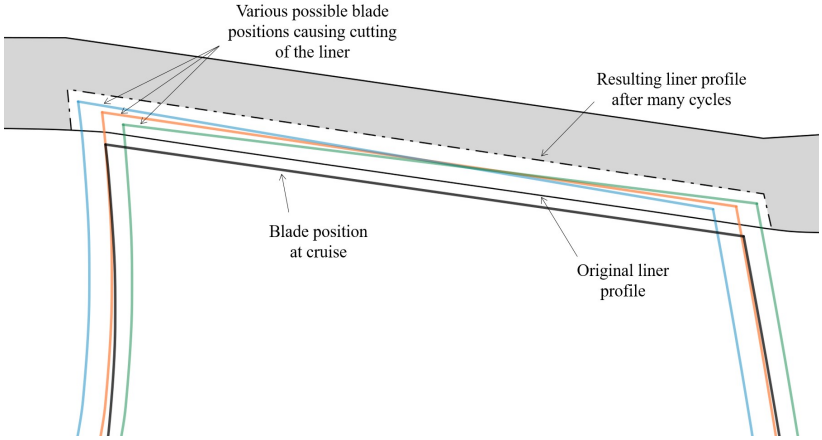

Fig. 6.6 Casing liner trench formation.

To understand the realistic size and topology of liner trenches that form in operational engines over their lifetime, an in-house tip-rubbing prediction tool was used. The tool input parameters include the blade and casing geometries and the number of engine cycles. The various blade tip and casing positions are predicted and, through calculating when they intersect, the new casing profile is built up cycle by cycle. This tool uses a whole engine model to give the predicted positions of components based on their expansion under varying rotational loads at conditions through each engine cycle. The untwist of the fan blades under load (which causes the leading and trailing edges to move radially different amounts) is included, as is the variation of the axial position of the blade. Together these effects are responsible for producing the steps at either end of the trench. Additionally, a gust prediction model that can simulate the sudden shift of the core radially within the nacelle is used. This uses data on the likelihood of gusts occurring each engine flight and adds the resulting displacement

into the tip-rubbing prediction tool. Gusts are the most significant contributor to casing liner trenches' depths.

For this study, the tip-rubbing prediction tool was used to predict the range, depth and variety of trench geometries likely to exist in an engine throughout its lifetime. This gives reasonable values of maximum depth of cut and the upstream and downstream positions that the trenches extend to. These are used to define the boundaries for the geometries to be investigated. The maximum overall clearance investigated in this study is about 1% span.

It is this tip rubbing prediction tool that forms the basis of the cost analysis used to specify cold build fan clearances that will use the results from this study. The rubbing tool predicts the depth of the clearance and trenches at various axial positions. The efficiency drop from datum will then be calculated based on the model developed in this work, and this can then be converted into a change in specific fuel consumption. The overall lifetime cost of any cold clearance can be found by integrating the SFC over an engine's lifetime.

6.4.1 Topologies under investigation

To understand the effects of casing liner trenches three types of tip clearance must be modelled. These are shown in Figure 6.7. Firstly, a blade with a standard clearance is simulated, and the effect on the fan aerodynamics of increasing the clearance through cropping the blade tip understood. This is the method usually used to vary blade clearance in simulations, and also represents the result of shortening a blade design to prevent any rubbing occurring, should this be a concern. In this study, the datum casing streamline is always kept constant (outside of the trench region). This is important as attempting to vary the clearance by adjusting the entire casing streamline would modify the inlet and exit areas of the fan, adding extra variables into the computations that would confuse the results. The second geometry that requires simulating is trenches above the blade tip, with the blade tip remaining outside of the trench and below the datum casing line. This allows an understanding of the impact of a trench appearing above the blade tip for an unchanged blade geometry to be developed, as well as a comparison to increasing clearance through cropping the blade. Finally, the blade should be allowed to extend inside of the trench, to see whether allowing a blade to extend further radially would provide any benefit, and whether interference from the trench steps has any impact. The clearance will be varied for each of the geometry types.

6.5 Simulation set up

The fan blade analysed in this work is a modern Rolls-Royce fan blade (rig scale) known as RR-FAN. RANS CFD simulations of the blade are carried out using Hydra [67] and the SA model. The blade is simulated in rotor only format (i.e. without the bypass or engine core stators), with a downstream splitter geometry and separate exit boundaries for the core and bypass flows. A schematic of the

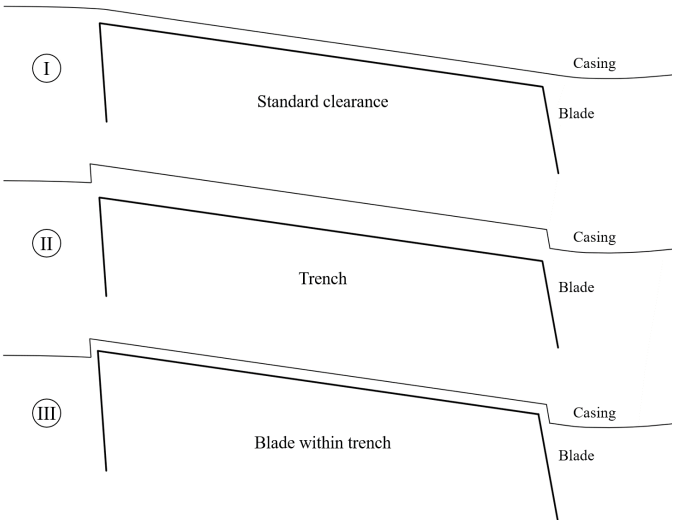

Fig. 6.7 Representation of the tip topologies under investigation.

CFD domain showing the position of the trench steps can be seen in Figure 6.8. At the inlet a radial distribution of total pressure, temperature and flow angle is used and at the exit boundaries radially averaged mass-meaned non-dimensionalised flow rate (capacity ($\frac{\dot{m}\sqrt{T_{tot}}}{P_{tot}}$)) is specified.

Measurement planes at -0.1 and 1.75 chord are used to calculate performance variables (total pressure ratio (PR), efficiency and total temperature ratio (TR)) as shown in Figure 6.9. Circumferentially mixed out, radially area-meaned total PR, and mass-meaned TR are used to calculate efficiency as these best match the approach used in the experiment. Each CFD simulation is carefully converged until no further change in performance values occurs (within 0.01%). This is important due to the small performance variations seen between clearance increments.

The turbomachinery meshing system PADRAM [110] is used to create the structured multi-block meshes. A range of mesh sizes was tested until mesh independence was indicated for both the overall mesh count and the number of radial cells in the tip gap (see Figure 6.10). The selected mesh has a total mesh size of 4.4M cells with 40 cells in the tip gap for the datum clearance. An extra mesh block is created above the blade tip when a trench is present to properly model the flow (see Figure 6.11). When tip clearance is varied in this study (through a trench or cropping the blade tip) the number of radial cells in the tip gap is varied proportionally, while maintaining those in the remainder of the span. The mesh across the tip gap is well aligned with the leakage flow direction. Figure 6.12 shows an example of the trench and clearance mesh used at the blade tip. The y^+ measured along the blade surface at mid-span and along the casing above the blade tip at mid-chord is of the order of one.

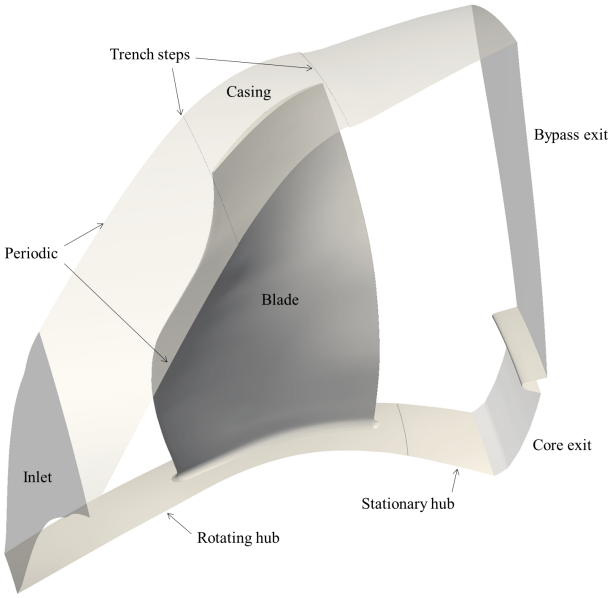

Fig. 6.8 Location of the trench within simulation set up.

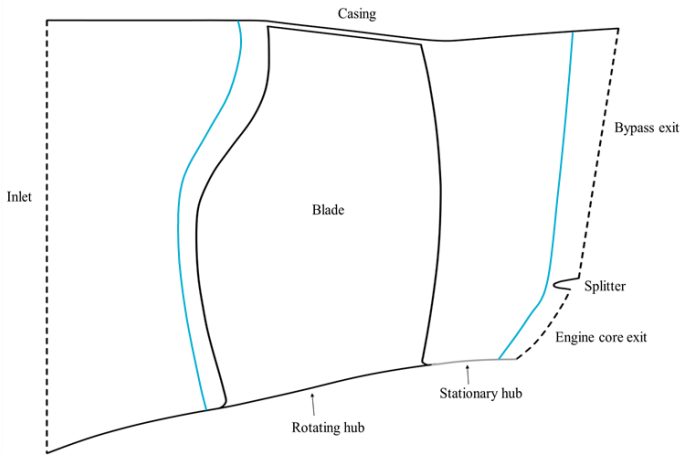

Fig. 6.9 Performance measurement planes and meridional view of setup.

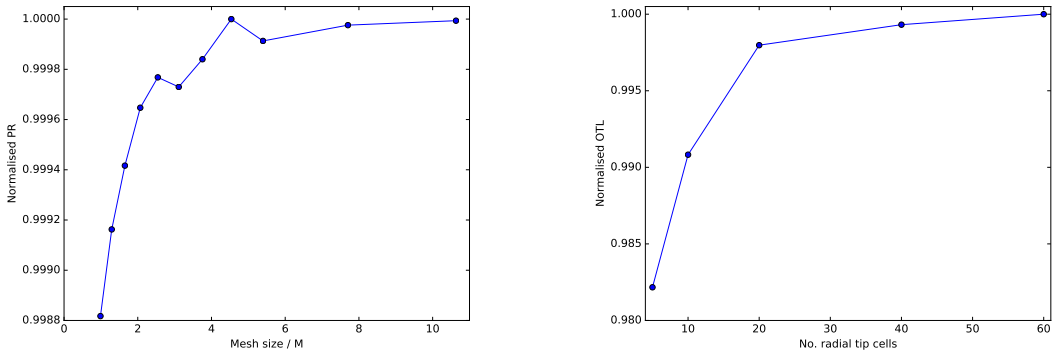

(a) (b)

Fig. 6.10 Mesh independence studies: (a) total mesh size, (b) radial cells in the clearance.

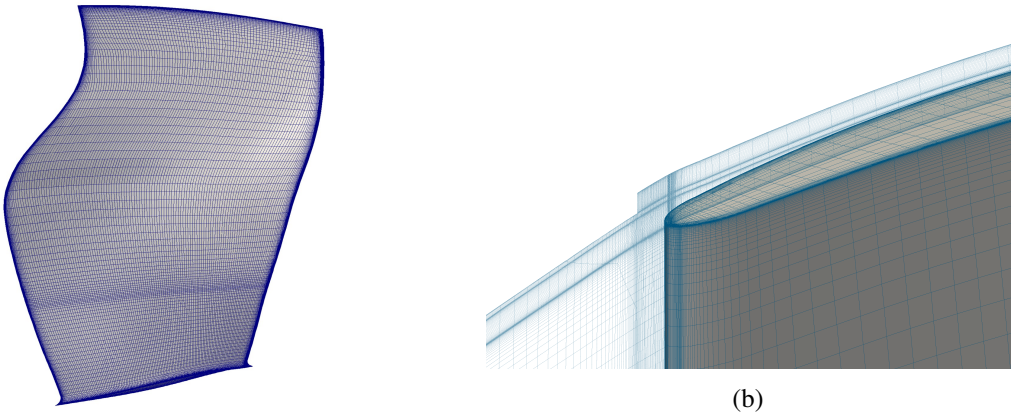

(a) (b)

Fig. 6.11 CFD mesh used (a) Blade surface mesh (not to scale) (b) Mesh across the tip gap.

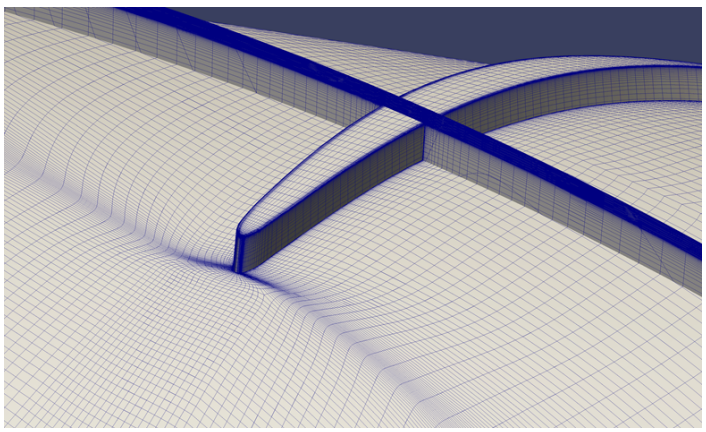

Fig. 6.12 An example PADRAM mesh across a clearance.

6.5.1 Validation

A comparison of the simulated blade performance with experimental data can be seen in Figures 6.13 and 6.14. Both the pressure ratio and efficiency curves match the experimental data well, though there is a slight offset to the overall values and stall margin. The radial curves show good comparison to experimental data, although the radial variation in efficiency is under-predicted compared to the experiment. Overall, the simulation compares well, lying within 1% across the range of flow rates.

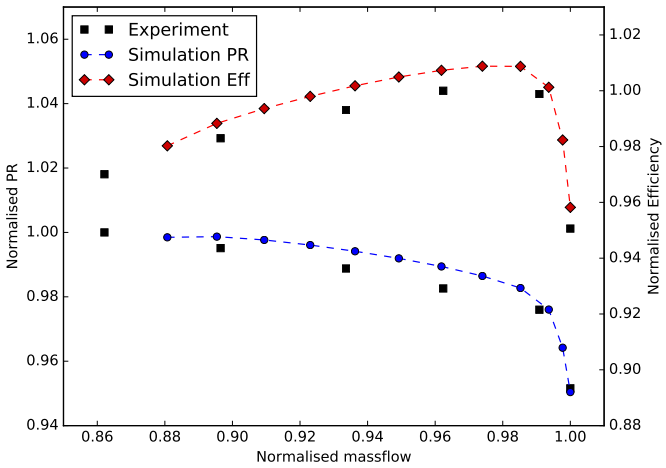

Fig. 6.13 Comparison of the simulation results with experimental data.

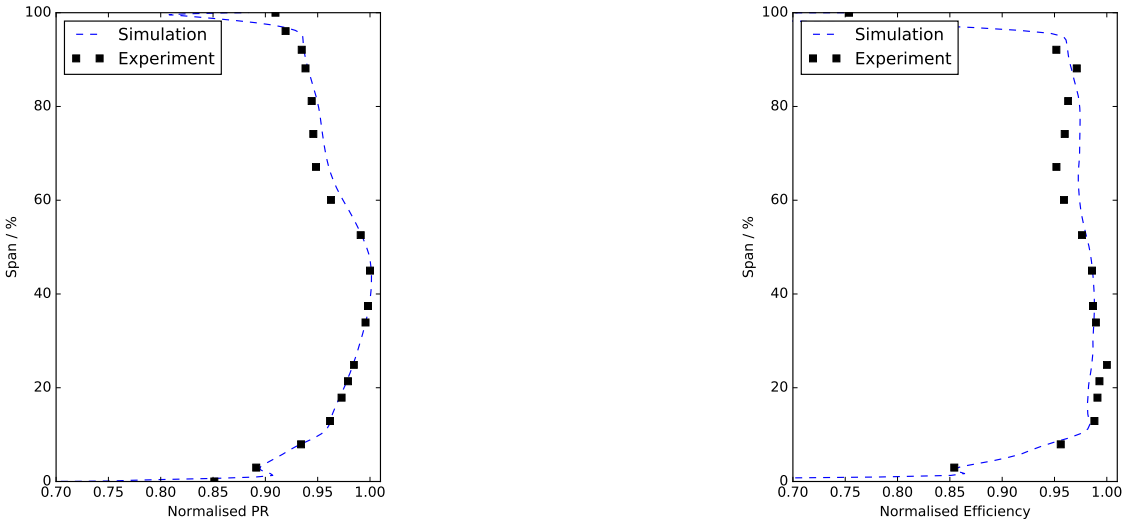

(a) (b)

Fig. 6.14 Radial profiles vs experiment at design point (a) Pressure ratio, (b) Efficiency.

6.6 Tip aerodynamics for the datum blade

Figure 6.15 shows 3D streamlines that highlight the vortex produced by the tip leakage flow for this case. The loss and blockage generating nature of the vortex can be seen in the 2D axial slices of entropy that highlight the path the vortex takes. The tip leakage vortex would normally be expected to begin at the leading edge of the blade, before increasing in size and progressing further across the passage as leakage flow along the remainder of the chord adds to the vortex. As can be seen however in Figure 6.15, for this case the tip vortex begins at around 0.3 chord. The reason for this is described below.

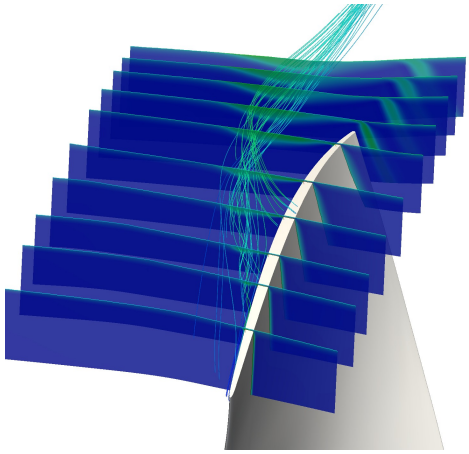

Fig. 6.15 Tip leakage vortex highlighted by 3D streamlines and slices coloured by entropy.

A visualisation of the flow streamlines within the tip gap is useful to understand the formation of the tip clearance vortex. A constant span slice taken halfway between the blade tip and the casing can be seen in Figure 6.16a. The image is coloured by relative Mach number and faint streamlines added to show the flow directions in the plane of the slice. It can be seen how the flow from the pressure side turns perpendicularly as it passes over the tip of the blade. It then exits the suction side and becomes entrained in the vortex. The path of the vortex at this height can be clearly seen in the streamlines that pass diagonally from near the LE towards the adjacent blade. Figure 6.16b gives a schematic of the vortex formation method.

It is interesting to observe that the formation of the main vortex does not start at the very LE of the blade. A small vortex can be seen emanating from the LE, but shortly after this there appears to be no flow streamlines passing from the blade gap. In fact, the streamlines actually reverse in this location, before (at roughly 0.3 chord) reverting to the expected direction and forming the main vortex that continues into the passage.

The cause for this unusual leakage behaviour is the passage shock that impinges on the pressure surface of the blade at around 0.2 chord. The low pressure region just upstream of the shock creates a positive pressure gradient from suction side to pressure side, which tries to drive the air back across

the tip gap around this location. This effect is a critical feature of the tip aerodynamics when a passage shock is present, and has an impact on the fan clearance behaviour studied throughout this work.

Figure 6.17 shows the chordwise tip leakage mass flow distribution for the datum case at design point. The data has been normalised by the maximum value at datum clearance. The tip leakage flow is measured through sections of a radial plane created along the camber line within the tip clearance and is normalised by the measured area. Some tip leakage occurs at the LE of the blade but then drops to almost zero, before at 0.2 chord returning to more typical behaviour. This is where the main vortex forms.

The importance of the chordwise distribution of tip leakage has not been discussed in detail in the literature. Previous researchers have acknowledged that chordwise variations in leakage flow occur and that it is the tip pressure field that controls this [120]. However the impact that this can have on the tip leakage aerodynamics has not been highlighted.

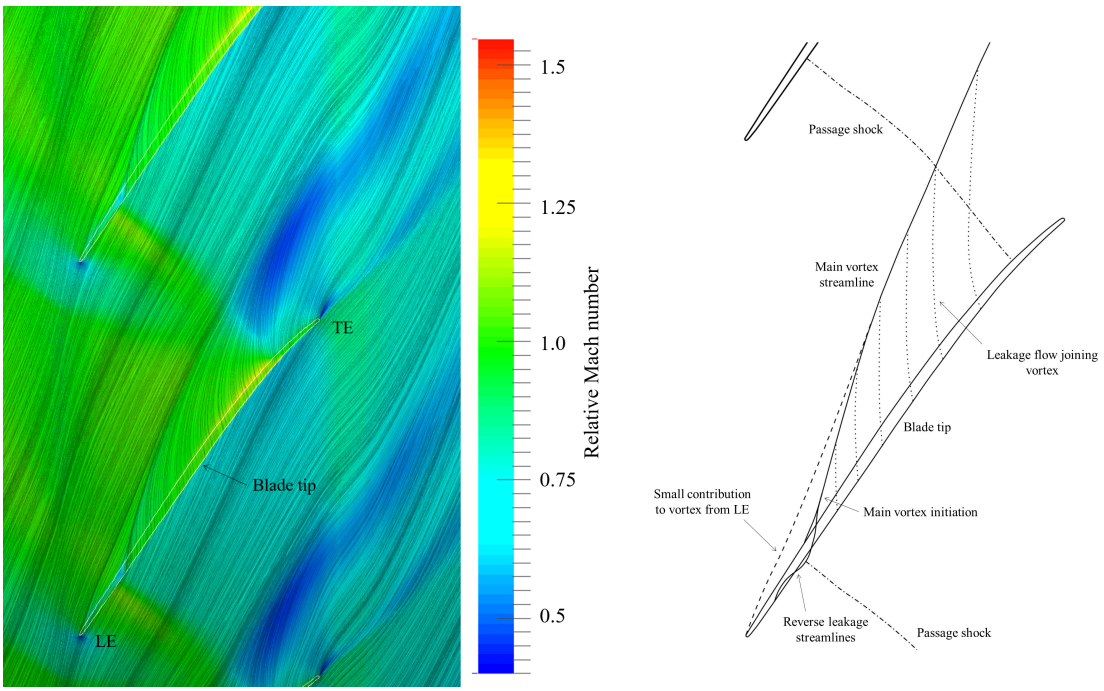

(a) (b)

Fig. 6.16 Clearance vortex formation method: (a) streamlines and rel Mach no. contours, (b) schematic of vortex formation.

6.7 Variation in performance due to cropping the blade tip

This section discusses the impact on blade performance of increasing clearance through cropping the blade tip. The significance of the tip leakage distribution defined by the shock location is again apparent from Figure 6.17. Increasing clearance can be seen to only have a significant effect after 0.2

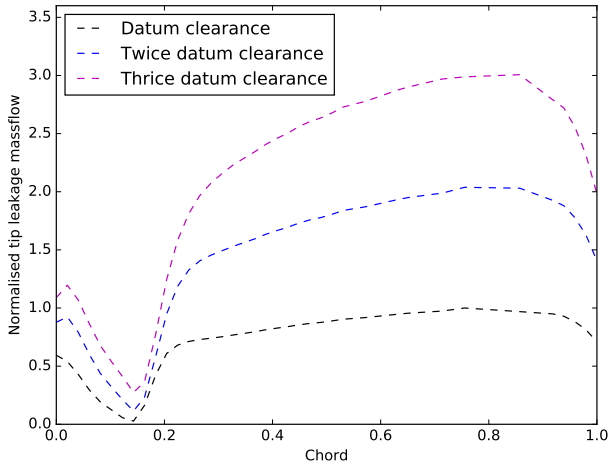

Fig. 6.17 Tip leakage distributions for various clearances.

chord, whereas before this point the lower pressure difference across the tip means that the increase in clearance area does not give the same change in mass flow. The shock position fixes the lowest leakage point, and dictates the leakage distribution. Towards the TE of the blade, a proportional relationship between clearance and leakage flow can be seen, but near the LE this is not the case.

It can be seen in Figure 6.18 how the leakage develops as the gap increases. Increased clearance results in the leakage flow progressing further across the passage. The leakage vortex cores are highlighted by contours of high Q-criterion in Figure 6.19. Q-criterion is a measure of vorticity, and hence these contours surround and reveal the vortex cores. The vortex core is larger and moves further across the passage as the clearance increases. The vortex also forms at a greater angle to the blade as the greater momentum of the increased leakage competes with the passage flow.

Figure 6.20 shows the increase in tip region entropy and accompanying axial velocity reduction (blockage generation) associated with cropping the blade tip. The plots are area averaged circumferentially and the measurements are taken at the downstream plane shown in Figure 6.9. Blockage is where the low momentum, high entropy flow 'blocks' part of the passage, reducing the mass flow that can pass through that region and re-distributing the flow in the remainder of the span. The increased tip leakage mass flow and the resulting larger vortex means there is greater mixing with the passage flow and generation of losses. The increased maximum entropy in the tip region and the radial extent of the losses due to leakage can be seen.

Figure 6.21 shows the delta in blade performance from zero clearance as the gap changes. A blade with no clearance (but maintaining correct rotating boundary conditions) was simulated to give the reference point. The gradient of efficiency drop with clearance magnitude is almost linear. This is similar to results found by previous researchers [101] [8] [105]. To assess whether an 'optimum'

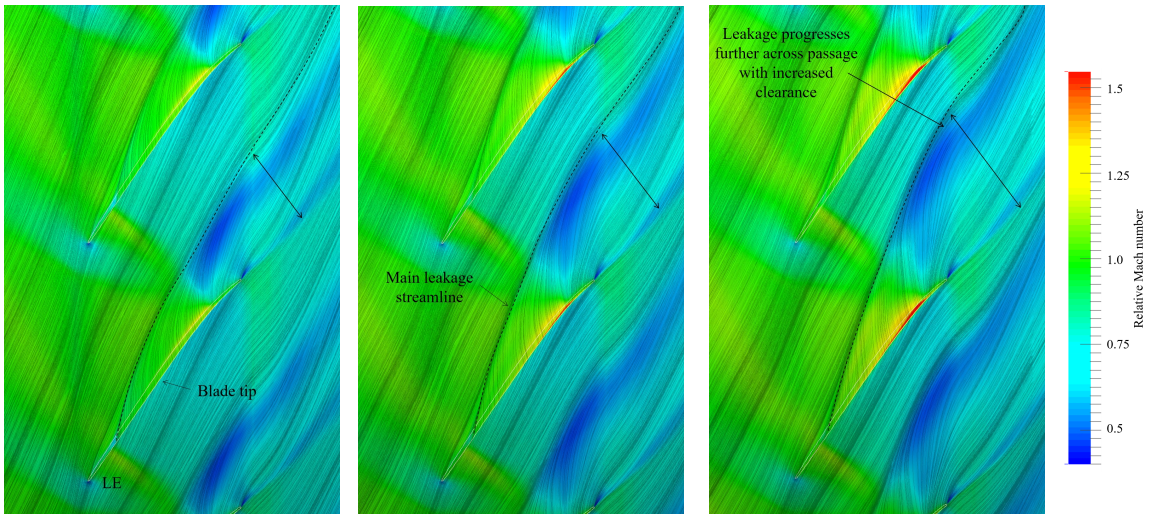

(a) Datum (b) Twice datum (c) Thrice datum

Fig. 6.18 Tip streamlines at various clearances with relative Mach number contour.

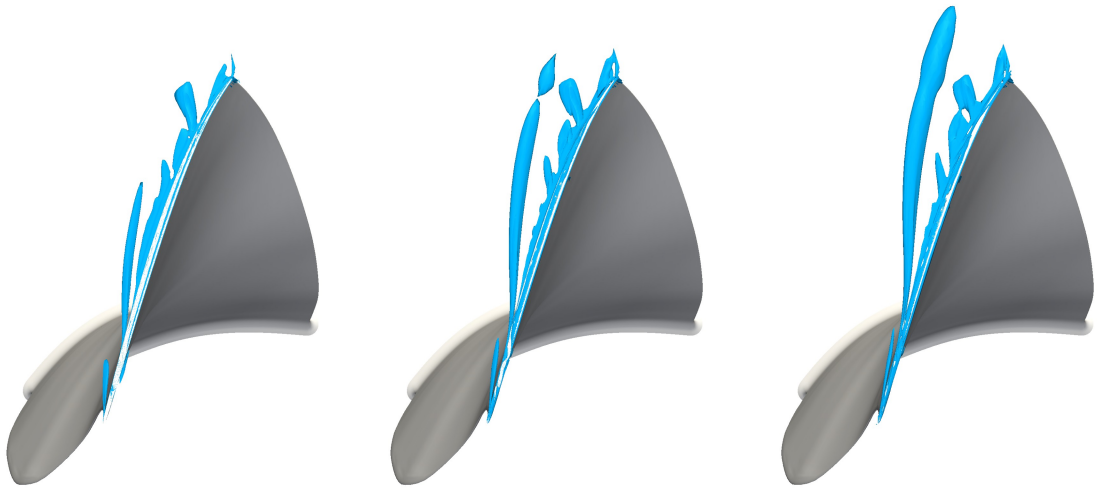

(a) Datum clearance (b) Twice datum (c) Thrice datum

Fig. 6.19 Q-criterion isocontours (10^8) showing change in vortex with clearance.

clearance occurred at a non-zero value, several gap sizes were simulated between zero and datum clearance. As can be seen in Figure 6.21a, no evidence of this was found.

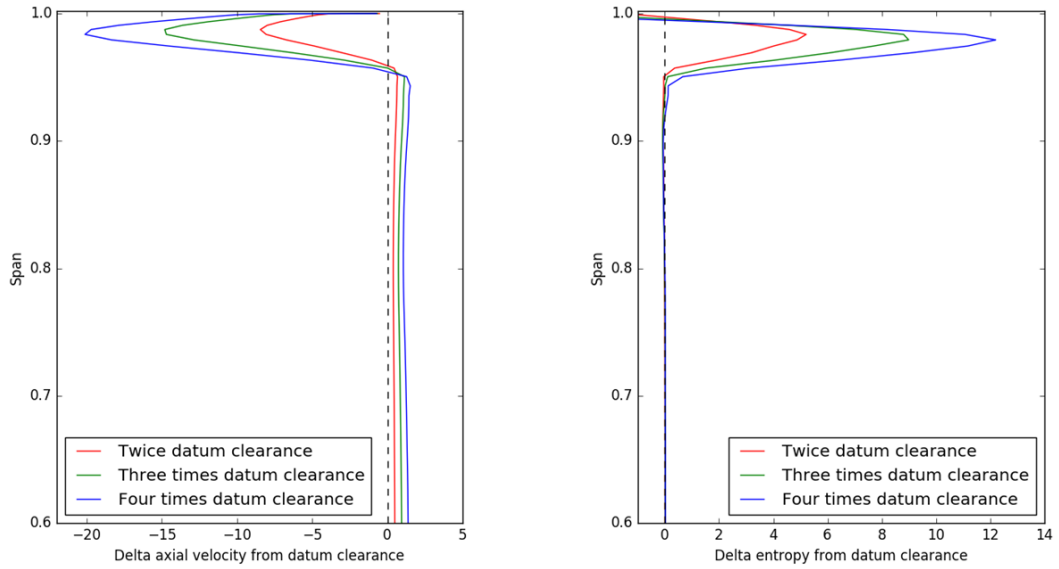

(a) (b)

Fig. 6.20 Radial distributions of delta (a) axial velocity and (b) entropy in the tip region from datum clearance.

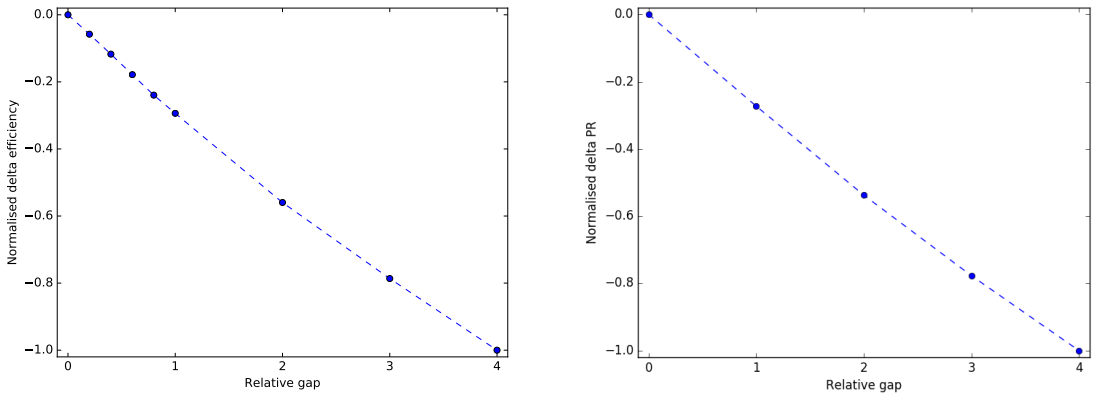

(a) (b)

Fig. 6.21 The variation of (a) efficiency and (b) pressure ratio with tip clearance magnitude.

6.8 Tip flow features at various operating points

It is important to not only understand the tip flow features at the design point but also how they vary across the blade characteristic at various mass flow rates. Hence, three flow conditions are assessed; the aerodynamic design point (ADP), a point with lower mass flow (near stall) and a point where the blade is almost choked (near choke). Figure 6.22 shows the tip leakage streamlines at the various conditions. The main vortex core, passage shocks and shock-vortex interaction have been highlighted. Figure 6.23 clearly shows the vortex core differences. It can be seen that the operating point has a significant effect on the tip flow structure. Near stall, the leakage vortex begins at the very LE of the blade because the shock is outside of the passage, in contrast to the design point. The vortex is also at a greater angle to the blade and the streamlines pass over the next blade's tip, meaning 'double-leakage' is present. Near choke, the tip leakage flow is similar to the design point but, with the shock further inside the passage and even greater acceleration upstream of it, the features are more enhanced. Near choke, reverse leakage flow can be seen from around 0.2-0.3c and the vortex is not initiated until 0.5c.

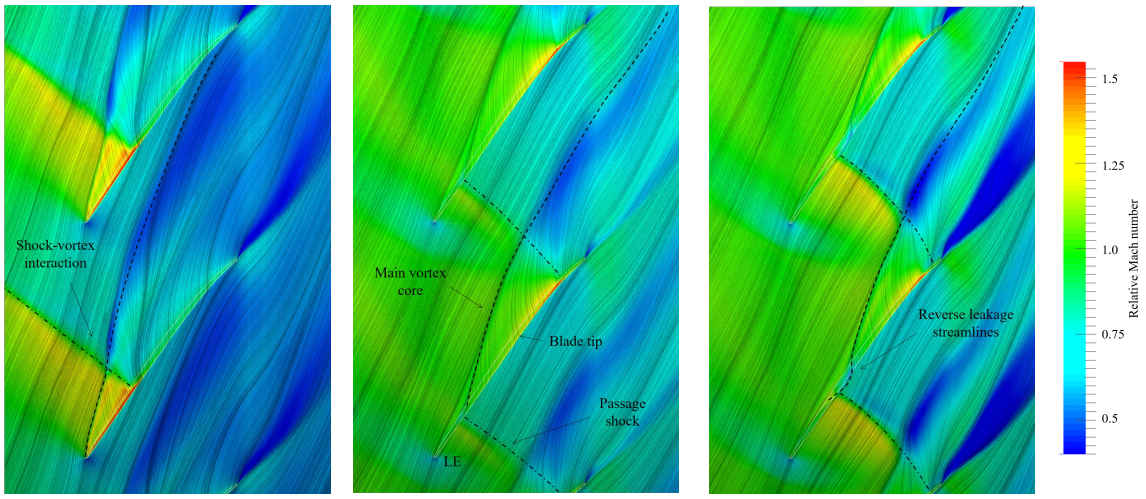

(a) Near stall (b) ADP (c) Near choke

Fig. 6.22 Tip leakage streamlines with relative Mach number contour at various mass flows.

These differences in the leakage behaviour can be clearly described using the tip leakage mass flow distributions shown in Figure 6.24a. The difference in the leakage flow and the dependence on shock position is clear. Near stall, the leakage flow is more consistent along the chord. It is maximum from 0-0.3c before the passage shock, as here the pressure difference across the blade is greatest. After this point, the tip leakage flow levels out to a value similar to those at other operating points. The reduction in leakage flow near choke around 0.2c is more significant than at the design point, and the net effect from around 0.15-0.3c is 'negative' leakage flow across the tip. From these results it is

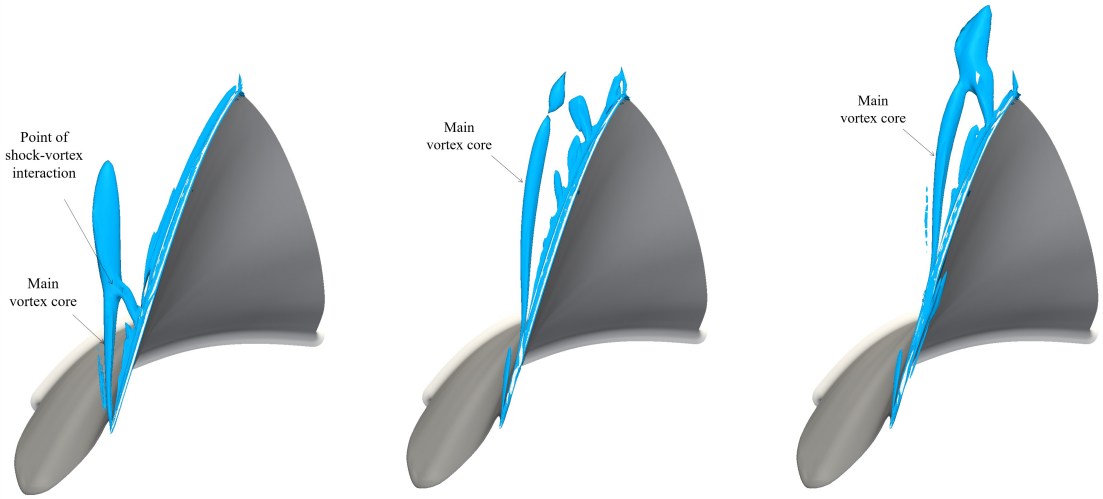

(a) Near stall (b) Design point (c) Near choke

Fig. 6.23 Q-criterion isocontours (10^8) showing change in vortex with flow rate.

clear that the primary feature controlling the leakage distribution is the positioning of the passage shock.

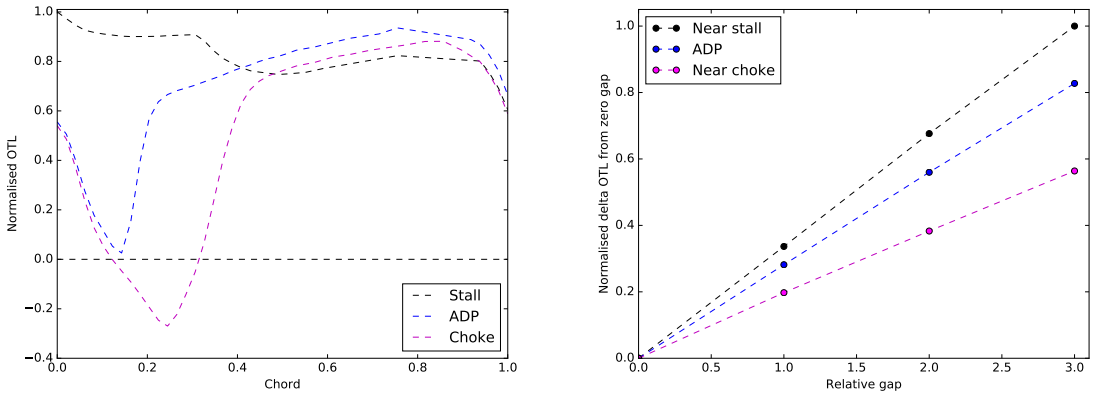

(a) (b)

Fig. 6.24 (a) Tip leakage distirbution and (b) overall tip leakage mass flow at various operating points.

The overall tip leakage mass flow (shown for a range of tip clearances) for each operating point is given in Figure 6.24b. Near stall, the overall tip leakage flow is significantly greater than at the design point, which again is greater than near choke. While the overall PR and blade loading (which usually has the greatest influence on leakage flow) does vary for each case, it is only small compared to these changes, which are almost entirely due to the unusual leakage distributions limiting the overall mass flow that can pass across the tip.

6.8.1 Performance variation with clearance at different flow rates

As shown above, the overall tip leakage mass flow is heavily influenced by the shock position and the resulting leakage distribution. The variation in the overall OTL is given again in Figure 6.25a. Figure 6.25b gives the corresponding efficiency variation with clearance. It can be seen that the variations in efficiency are not directly related to the variation in clearance mass flow as might be expected. The trends are all linear, however the case with the steepest gradient of leakage flow with clearance, 'near stall', has the shallowest gradient of efficiency with clearance. This suggests that the variation in efficiency with clearance is not entirely dependent on the leakage magnitude, and some secondary effects are having an impact.

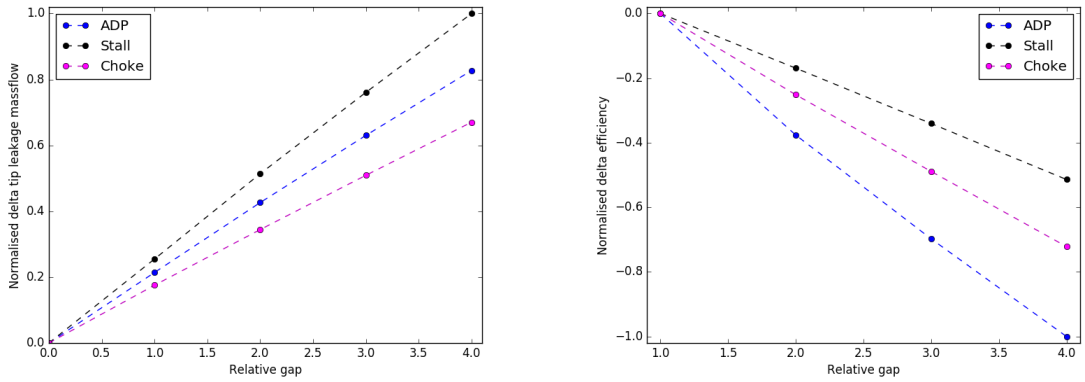

(a) (b)

Fig. 6.25 Variation in (a) OTL and (b) efficiency with clearance at various flow rates.

Figure 6.26 shows the radial distribution of the difference in entropy from four times to datum clearance for the three flow rates studied. It can be seen that in the tip region similar behaviour occurs at the different flow rates, with increased clearance (and leakage mass flow) locally increasing the entropy in the tip region. Away from this tip region however, an additional effect can be seen for the near-stall case. As the clearance is increased (from datum to four-times datum), the entropy between 0.85 and 0.94 span decreases. It is this decrease, counteracting increased loss in the top 5% of span, that has caused the unexpected efficiency gradient with clearance described above. The cause of this entropy decrease from datum to four times clearance in this region is the volume of separated flow that can be seen in Figure 6.26b. As the clearance is increased, and the tip region is blocked up by the low momentum leakage flow, the flow is redistributed, leading to an increase in axial velocity further down the span (similar to that shown in Figure 6.20a). For the near-stall case, this increase in axial velocity has a positive effect on the suction surface separation, locally reducing the angle of attack, reducing the size of the separation and therefore the entropy generated. This is the reason the gradients of efficiency versus clearance in 6.25 do not follow the expected ranking based on the gradients of leakage flow versus clearance for the different flow rates.

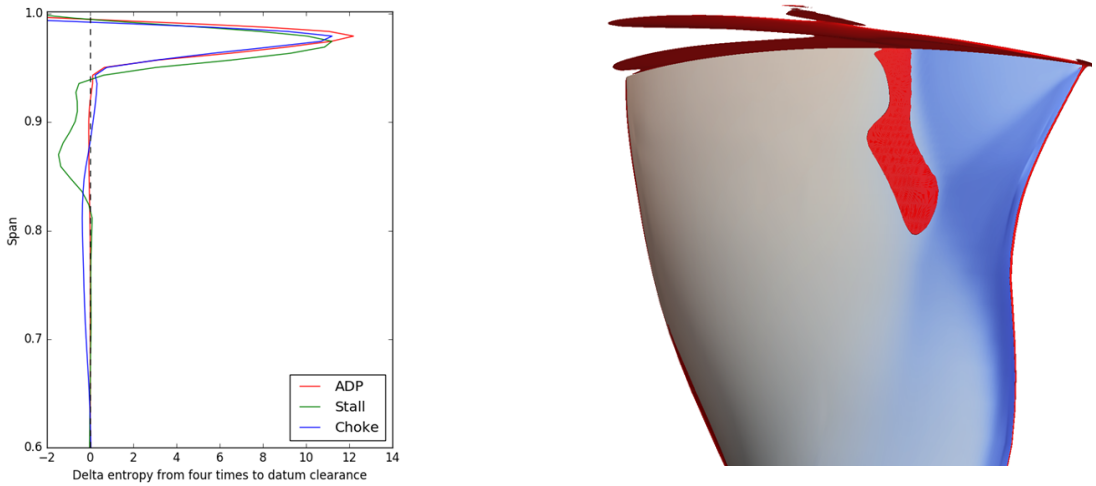

(a) (b)

Fig. 6.26 (a) Delta entropy from four times to datum clearance and (b) Separation on the suction surface near to stall.

6.8.2 The impact of clearance on stall margin

The impact of varying the clearance on the fan blade characteristics is given in Figure 6.27. It can be seen how the variation in efficiency with clearance matches that shown in Figure 6.25b, with the greatest performance drop focused around the middle of the characteristic, and more limited near to choke and stall.

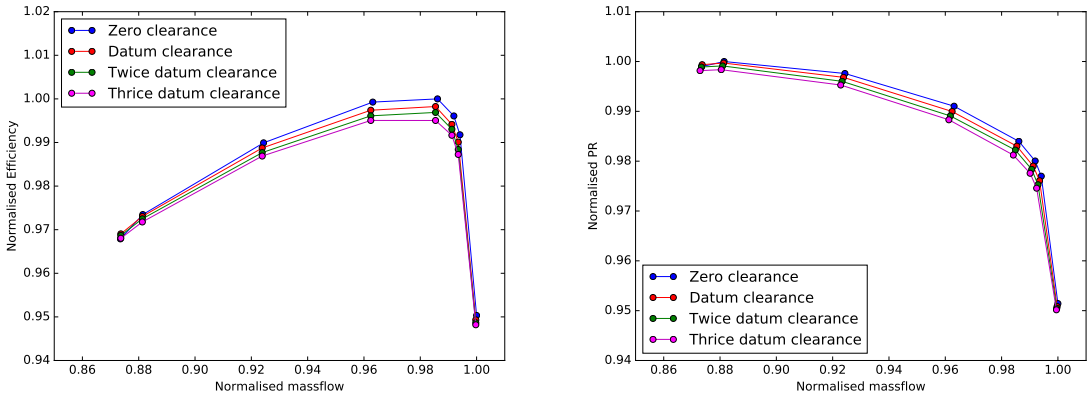

(a) (b)

Fig. 6.27 The impact of clearance on the (a) Efficiency and (b) Pressure Ratio characteristics.

According to the literature, one of the key implications of increased clearance is its impact on stall margin. As noted by Adamczyk et al. [2], the increased blockage associated with greater tip leakage

can reduce the amount of flow passing through the blade passage and, at lower mass flow rates, trigger stall. This reduction in stall margin would be observable on the left side of the characteristics, with increased clearance cases having a higher minimum stable flow rate. As shown in Figure 6.27 however, no evidence of this was detected. The last stable point for each clearance magnitude is shown, and there is no variation in the mass flow of this point, suggesting there is no difference between the stall margins of the different blade clearance geometries.

A contributing factor to this may be the impact of increased clearance on the suction side separation, as shown in Figure 6.26 and described above. The fact that increased leakage flow reduces separation and blockage slightly lower down the span means that these two effects may partially cancel each other out, thus vastly reducing the impact of increased leakage flow on the stall margin.

6.8.3 RR-FAN-2

As part of this work an additional fan blade (the same as used in Chapter 5 - RR-FAN-2) was simulated and the influence of clearance on its performance assessed. Most of the results for this blade have not been included for brevity, but here the variation in blade performance across the characteristic is assessed to see if a similar impact on stall margin is seen.

RR-FAN-2 is the next iteration of the RR-FAN fan blade. It is designed to perform a similar duty, but has a higher hub-tip ratio and a greater bypass ratio. A very similar tip flow structure is observed, as shown in Figure 6.28, with the main difference being a slightly more swallowed shock for RR-FAN-2, due to the design point being further to the right of the characteristic.

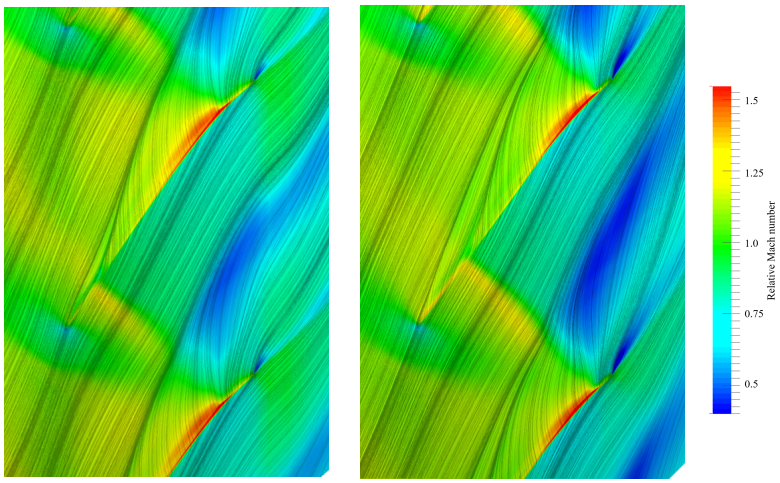

(a) (b)

Fig. 6.28 Comparison of (a) RR-FAN and (b) RR-FAN-2 tip leakage streamlines and rel. Mach number contours.

Figure 6.29 shows the variation in PR for RR-FAN-2 across the characteristic. The last stable point for each geometry is shown on the stall side. It can be seen how in this case the stall margin appears to

increase slightly with clearance; the last stable point is further to the left on the characteristics as the tip gap is increased. This is quite an unusual finding, as it is usually expected that increased clearance reduces blade stall margin.

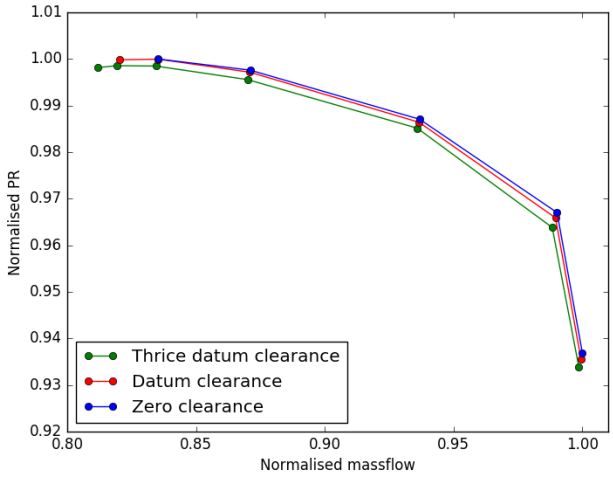

Fig. 6.29 The impact of clearance on stall margin.

The reason for this result is related to the separation on the blade surface. Figure 6.30 shows how the separation changes with flow rate for different tip clearance magnitudes. Only part of the blade geometry can be shown for proprietary reasons. It can be observed that it is this separation that leads to stall of the blade (as also described in Chapter 5), and how the size of this separation is reduced when the clearance is increased, due to spanwise flow re-distribution.

The results for both blades suggest that in these cases, contrary to that shown in other work, the clearance magnitude does not appear to greatly influence stall margin and may even have a positive impact. It must be noted however that the steady RANS simulations carried out here are not able to fully capture the physical phenomena that lead to fan blade stall. In a real fan, unsteady features and physical effects, such as rotating and spike-induced stall [82] can lead to stall occurring at a different point than might be predicted by a steady simulation.

6.9 Non-uniform variations in tip clearance

During the operation of a real engine, fan tip clearances are rarely uniform along the chord. As the blade and casing positions vary, the casing liner becomes rubbed. Due to the way the blade untwists under varying loads and at different rotational speeds, sometimes the LE rubs more than the TE and vice-versa. The result is that in reality the fan will often experience non-uniform tip clearances, with variations from LE to TE. As an initial investigation into the impact of non-uniform clearances, various sloped clearances are here investigated. To investigate the impact of this, several geometries

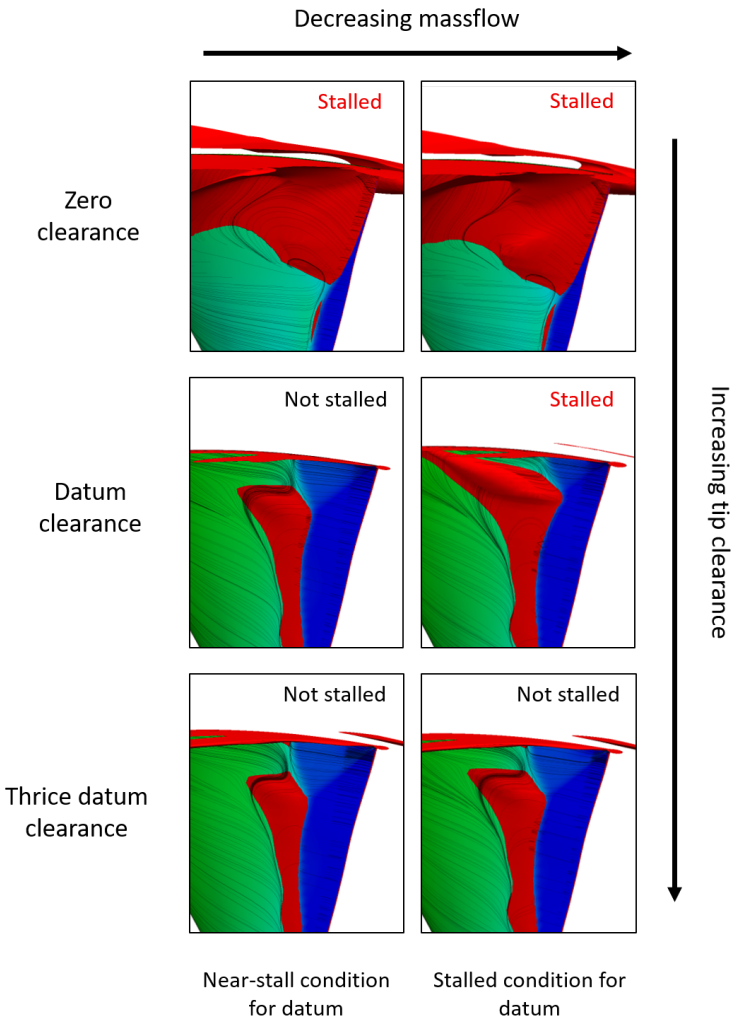

Fig. 6.30 The impact of clearance on stall margin - separated flow indicated in red.

have been simulated. Figure 6.31a gives examples where the LE/TE tip clearances have been changed by ±50%. The geometries here are referred by their relative tip clearances to uniform; i.e. '0.5-1.5' has 50% of datum tip clearance at the LE and 150% datum clearance at the TE. Several variations have been tested up to a 75% increase/decrease. The tip clearance of each of these designs was then uniformly increased to understand the impact of a sloped clearance at different average radial gaps between tip and casing. This indicates which regions have the biggest impact, and how critical the variations are compared to the overall average tip clearance.

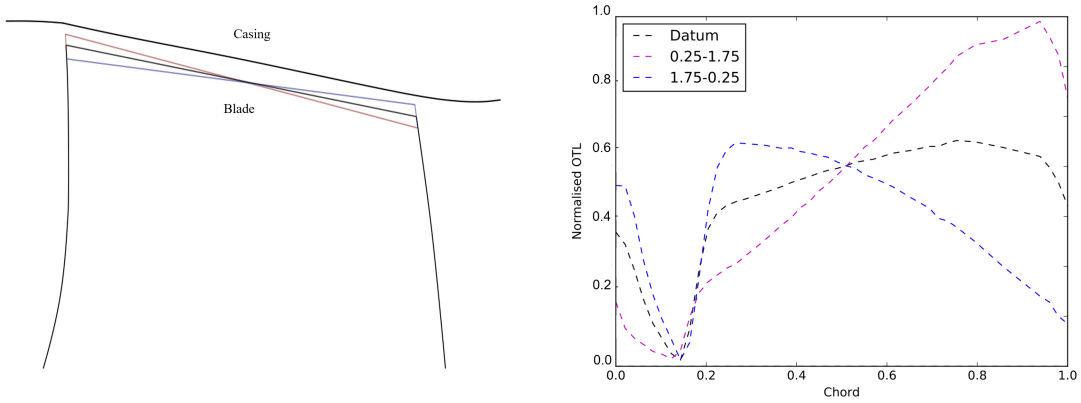

(a) (b)

Fig. 6.31 (a) Example clearances with variation LE to TE, (b) OTL distributions for varying clearances.

Figure 6.32 shows the differences in efficiency and tip leakage mass flow between a design with no tip clearance and the maximum varying clearance tested, ±75%. A small offset in both efficiency and mass flow is seen between the geometries. This stays quite consistent as the average tip clearance is increased, though is so small it is doubtful whether experimental testing would pick this kind of trend up, and therefore how significant this is. Increasing the LE tip clearance (shown in blue) is interestingly shown to decrease both efficiency and tip leakage mass flow by a small amount.

Figure 6.33 highlights how across the range of sloped tip clearances analysed, a decrease in LE clearance increases efficiency. A consistent correlation is seen across all levels of tip clearance bias, and the greater the variation between LE and TE clearances, the greater is the effect. This is reflected with decreased TE tip clearance having the opposite effect. It is interesting to note that the impact of the varied tip clearance is the same for both tip leakage mass flow and efficiency. This is counter-intuitive, as typically an increase in tip leakage mass flow leads to a reduction in blade efficiency and vice versa. There is normally a negative correlation between tip leakage mass flow and blade efficiency. The opposite is observed here for chordwise tip clearance variations.

Figure 6.31b shows the impact that a sloped tip clearance has on the chordwise tip leakage flow distribution. Once again, the impact of the unusual leakage distribution is shown to be important. The low pressure gradient near 0.2c fixes the minimum leakage near this point and at a low value.

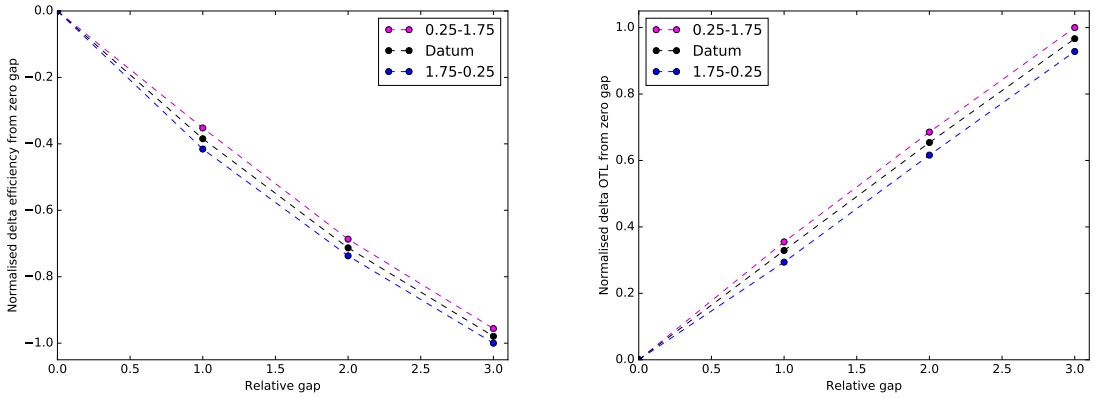

(a) (b)

Fig. 6.32 (a) efficiency and (b) OTL variation with clearance for ± 75 % sloped gaps and datum.

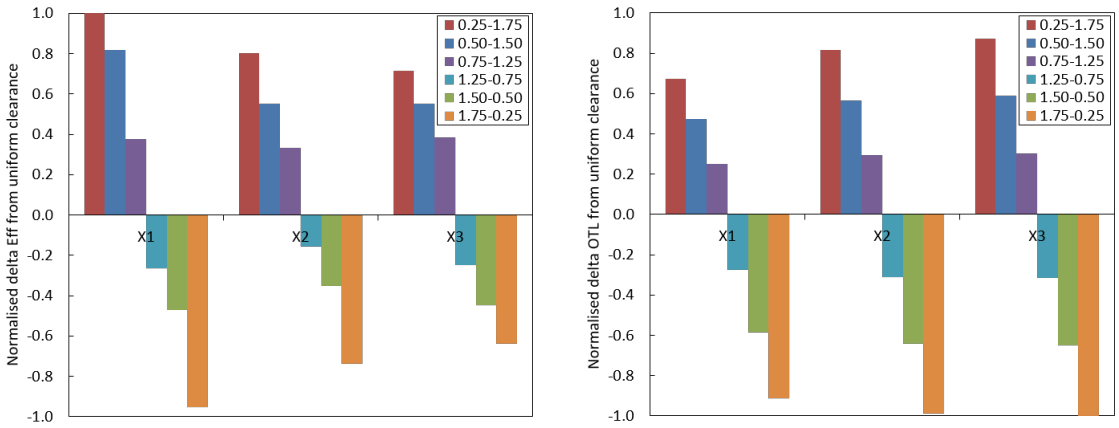

(a) (b)

Fig. 6.33 Sloped gap (a) difference in efficiency and (b) difference in OTL from uniform clearance.

The change in the leakage flow caused by the varying tip clearances can be seen. In the TE half of the chord the variation in tip clearance size results in a large change in the leakage flow, whereas towards the LE the same variation in tip clearance has a lesser effect. The result is an overall increase in leakage when the tip clearance is biased towards the TE, and vice-versa. The reduced pressure gradient across the tip towards the LE means that an increase in tip clearance has very little effect on the leakage flow there.

The above results suggest that the addition of leakage flow further downstream is not as significant as the leakage flow nearer the LE. To understand why increased TE clearance increases efficiency, despite increasing the overall tip leakage mass flow, the radial performance distributions must be analysed.

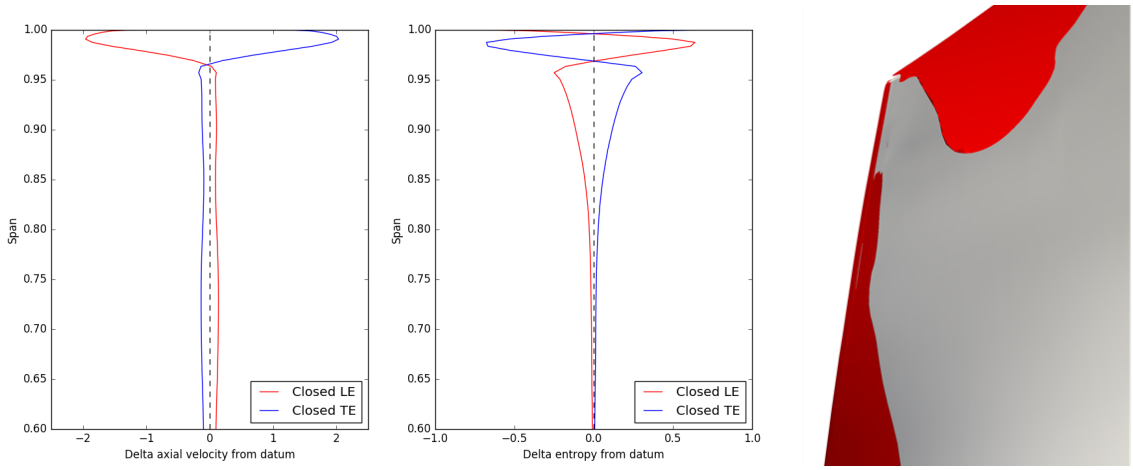

(a) (b) (c)

Fig. 6.34 (a) delta axial velocity and (b) delta entropy from uniform clearance, (c) separation near to the blade TE.

As can be seen in Figure 6.34, once again it is flow redistribution due to increased tip blockage that is responsible for the relationships seen. As the clearance is biased towards the TE, the increase in OTL causes a small increase in axial velocity lower down the span that reduces separation and has a significant effect on the entropy generated. TE separation is reduced when clearance is biased towards the TE, and this causes an overall increase in efficiency.

These results demonstrate the impact that variations in clearance geometry can have, and that the radial distribution of flow can play an important role. To understand further, a more detailed analysis of the impact of various non-uniform clearance geometries, as well as the impact of trenches, follows.

6.10 Modelling trenched casing geometries

To understand the impact on the tip aerodynamics of trenches above the blade tip, the three types of geometry shown in Figure 6.7 are simulated. Each case is simulated with a range of clearances

between the blade tip and casing. This was achieved by either 'cropping' the length of the blade for the case with no trench (type I in Figure 6.35), raising the depth of the trench above a constant blade (type II), or by increasing the length of the blade inside of a large trench (type III). The clearances are all referred to by their relative magnitudes to datum clearance.

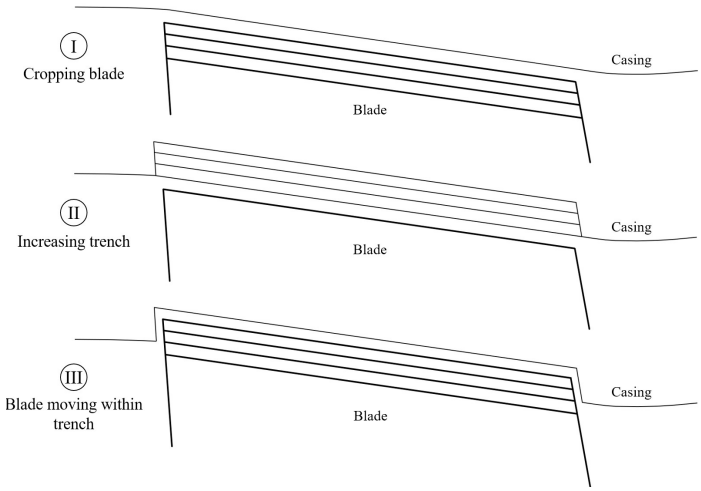

Fig. 6.35 Representation of the geometry types under investigation.

6.10.1 Geometry parameterisation and meshing

The trench geometries are created using the PADRAM 'trench' technique. This technique was originally designed to modify the casing streamline and produce 'trenches' upstream of the LE near the inlet and downstream of the TE near the exit. To adapt this technique to be used here, it was combined with modifying and raising the casing streamline, so that once parts of the casing (up and downstream of the blade) were removed, the desired trench was produced, with the casing outside of the trench region matching the original casing geometry.

A python script was developed that could parametrically modify the casing streamline (in the *Blade Definition file*) and the trench and variable tip gap settings in PADRAM to produce the desired geometry. The radial mesh inside the trench was generated in the same way as the tip clearance mesh, thus ensuring consistent mesh density regardless of clearance and trench depths. The radial mesh in the remainder of the span was kept constant.

Developing the capability to parametrically define trench geometries in this way was one of the most demanding aspects of the work in this chapter. The original version of the PADRAM technique had many faults and limitations therefore for some time BOXER [35] was used to produce trench geometries. Limitations in the method of defining and passing the geometries into BOXER, and the lower quality of mesh produced for the same cell count, meant that this was not practical

however. Some development was then carried out in various stages in PADRAM with the help of the Rolls-Royce 3DGEMO team that meant the trench technique could eventually be used as intended.

6.10.2 Comparison of different trench topologies

The variation in performance due to cropping the blade tip has been discussed above and is given in Figure 6.21. To compare the effect on blade performance of cropping the tip to varying the clearance with trenches, four clearances were simulated for each geometry type. Representations of the various clearances simulated can be seen in Figure 6.35 and the obtained results are given in Figures 6.36 and 6.37. The result have been normalised by the datum clearance value.

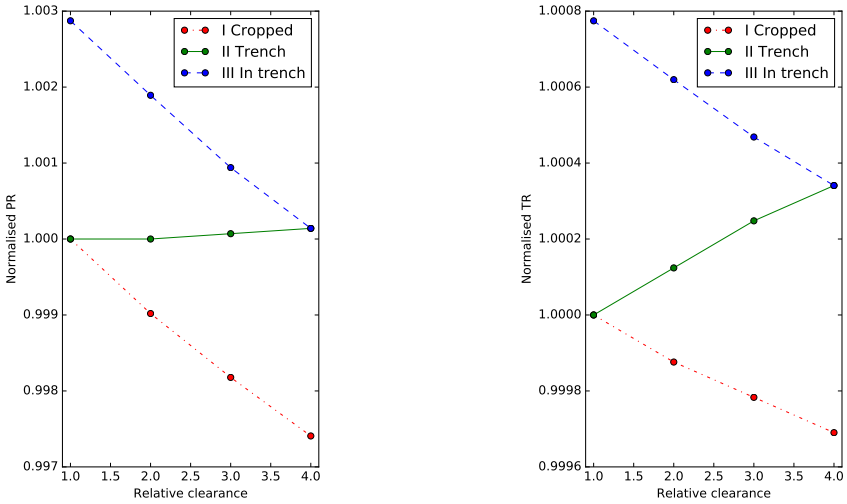

(a) (b)

Fig. 6.36 (a) PR and (b) TR variation with clearance.

It can be seen in Figure 6.36a how there is a linear relationship between pressure ratio and clearance for each of the geometry types. As the blade is cropped for type I the PR decreases. However a different behaviour is seen for the trench case. As the trench is introduced above the blade and increased in depth, the overall pressure ratio remains almost constant. For the in-trench case (type III), extending the blade inside of the trench follows the same linear gradient of PR as the case with no trench. The impact of varying the trench depth above the blade tip clearly has a different effect to cropping/uncropping the blade.

Figure 6.36b shows the variation of temperature ratio with clearance. It can be seen that the gradient of TR with clearance does not follow that of PR, as shown in Figure 6.36a. This is due to the increased losses associated with increased clearance, caused by extra tip leakage mass flow and mixing of the tip leakage vortex. These viscous effects result in additional, non-isentropic temperature generation which increases the temperature ratio, despite little variation in pressure ratio for the trench

case. For the cropped and in-trench cases, the variation in TR is caused by a combination of the PR variation and the increase in viscous heating. The overall TR decreases with clearance for these cases because the variation in PR (and therefore isentropic TR) is greater than the TR change due to extra tip leakage effects.

To understand the impact of the three different cases on performance, the variation in blade efficiency can be analysed. This is given in Figure 6.37. It can be seen that the gradient of efficiency with clearance is steeper for the cropped (I) and the in-trench (III) cases than for the trench case. The geometries where the blade is being cropped are most sensitive to clearance, and very similar to one another, suggesting the performance gradient is not sensitive to whether the blade tip is sat in or outside of the trench. It is the difference between cropping the blade or varying the clearance through a trench that is most significant.

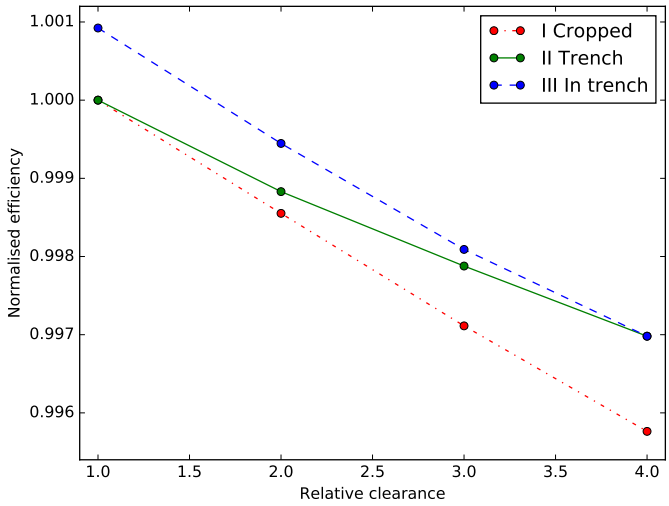

Fig. 6.37 Efficiency variation with clearance.

To understand the performance differences between the three cases, geometries with the same clearance for each case can be directly compared. Figure 6.37 shows a difference in efficiency for geometries with twice datum clearance from each case. The radial profiles of loss for these are given in Figure 6.38. These were taken at 0.9 chord ('within' the trench in the axial direction), to ensure the effect of tip region features could be clearly seen. It is shown that the extent of tip region loss is lower for the blade within the trench compared to the other two cases. The trench case also has less tip region loss than the cropped blade. This shows that lower tip region loss is produced when the clearance is moved further radially via a trench. It can also be seen how the region of the span affected by the tip loss has been shifted radially, moving the leakage out of the main passage flow. The effects seen match the explanation by Robideau et al. [95], who described how moving the clearance outside of the main flow path is beneficial.

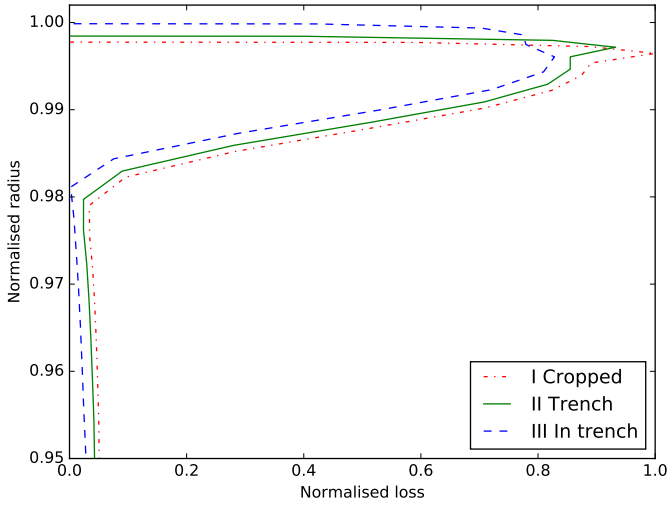

Fig. 6.38 Tip region loss for the different cases at twice datum clearance, measured at 0.9 chord.

6.10.3 Flow visualisation

To further understand the differences between the tip region behaviour and losses for the three cases described above, visualisation of the tip flow features is given in Figures 6.39 and 6.40. These confirm that seen in Figure 6.38. Figure 6.39 demonstrates how raising the clearance through a trench shifts the leakage vortex radially. The spanwise slices shown are at a constant height just below the datum casing streamline, and show how the vortex size and intensity at this height is reduced by raising the clearance via a trench. The blue, low velocity region that highlights the vortex core shows how at this height in the main passage, the impact of the vortex is reduced.

Figure 6.40 shows slices of entropy just downstream of the trailing edge step. These highlight the vortex size and intensity for each case. It can be seen how lower vortex intensity and entropy is present where the clearance has been raised radially via a trench. This shows the impact of the trench on the main passage flow (outside of the trench region), and indicates that shifting the leakage flow out of the main passage reduces the entropy generated.

6.10.4 The influence of trench steps

An additional objective of this work was to understand the importance of the trench steps themselves and whether they have any direct influence on the tip flow behaviour. Figure 6.41 shows isosurfaces of zero axial velocity (showing separated flow) in the LE and TE tip regions. The separation caused by the LE Backward Facing Step (BFS) and the TE Forward Facing Step (FFS) can be seen, as well as the tip leakage flow. It can be seen that the LE trench step separation does not directly interact with the tip leakage flow. The extent of the separation is not sufficient for this to occur and also the

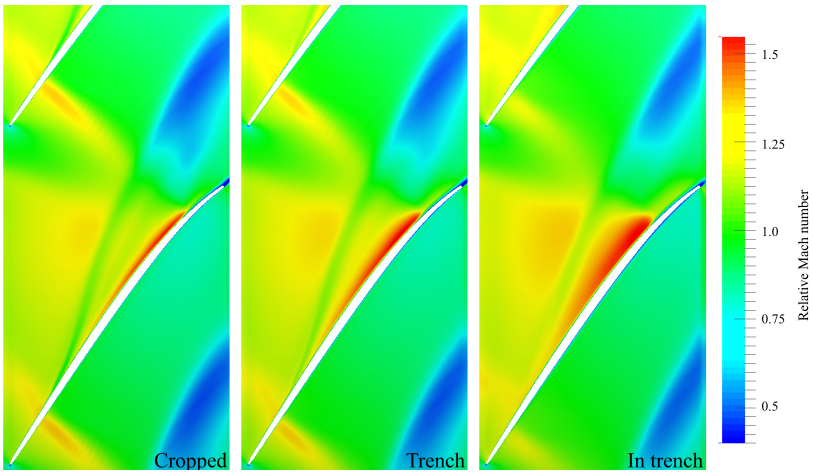

Fig. 6.39 Relative Mach number contours showing the variation in vortex size at a constant radius for twice datum clearance.

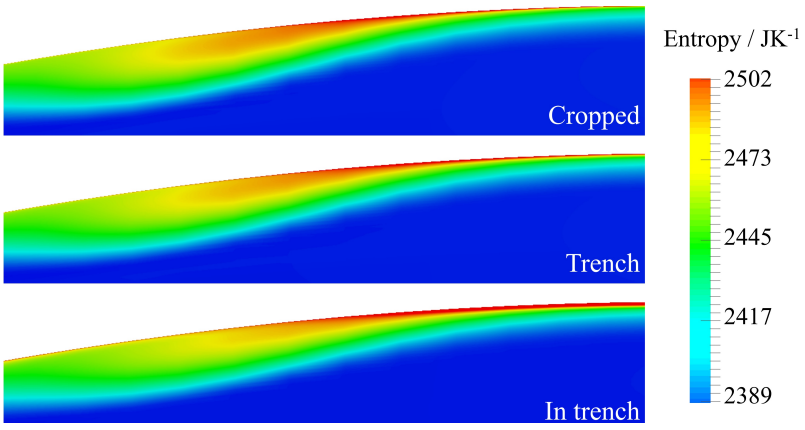

Fig. 6.40 Variation in blockage intensity downstream of trailing edge step for each case at twice datum clearance.

tip leakage vortex is initiated away from the blade LE, as discussed previously. The TE trench step appears more significant. The leakage flow can be seen to directly interact with the TE step, and it is shown that separation at the step is reduced on the suction side, where the leakage vortex washes over it. The separation at the TE step enters the main passage flow, and is therefore likely to have greater impact than separation from the leading edge step which is contained within the trench. The increase in entropy due to this separation can be seen as the dark red contour along the top right of the in trench image in Figure 6.40.

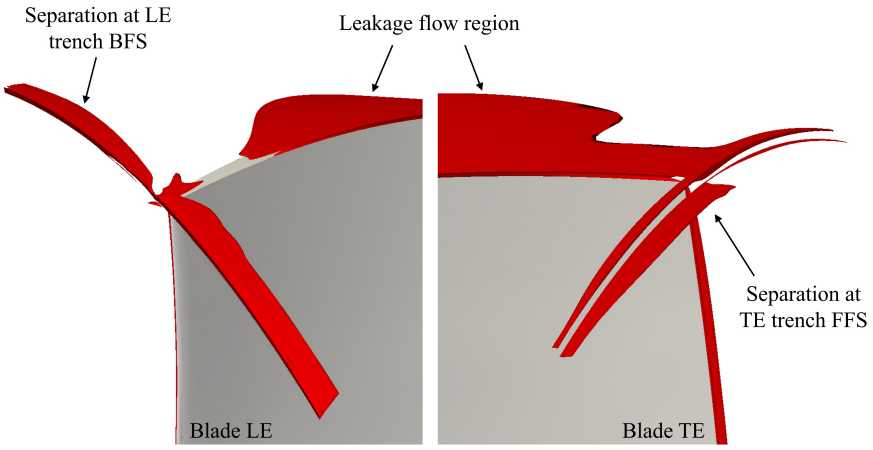

Fig. 6.41 Separated flow around the trench steps.

6.10.5 Radial profiles

Aside from the differences in the tip region behaviour between the three cases, the gradients of performance with clearance were also shown to vary. Figure 6.36a showed how for the trench case the PR appeared to stay almost constant while when cropping the blade tip a drop in PR is seen. Figure 6.42 shows the delta in radial PR profiles from twice datum clearance to datum clearance for each of the three cases.

It can be seen that there are differences in the tip region blockage between the three topologies, reflecting the differences described above in Section 6.10.2. What can also be seen is the difference slightly further down the span. The PR delta for the trench case is different to the cases where the blade is cropped. The effect of this is that the PR deficit in the tip region, caused by increasing clearance, is almost cancelled out by the changes lower down the span, which is the reason the PR gradient appears nearly constant in Figure 6.36a. This suggests that the overall blade length has an influence on the blade performance, and also shows how the differences between these cases are not isolated to the tip region.

This can be understood, as, in the 'trench' case, the blade length is unaltered, therefore a redistribution of flow increases PR further down the span. In the other two cases, the blade length is

shortened, and this results in a reduction in the overall PR that the blade achieves; i.e. there is less blade surface to do work and the PR is lowered further down the span.

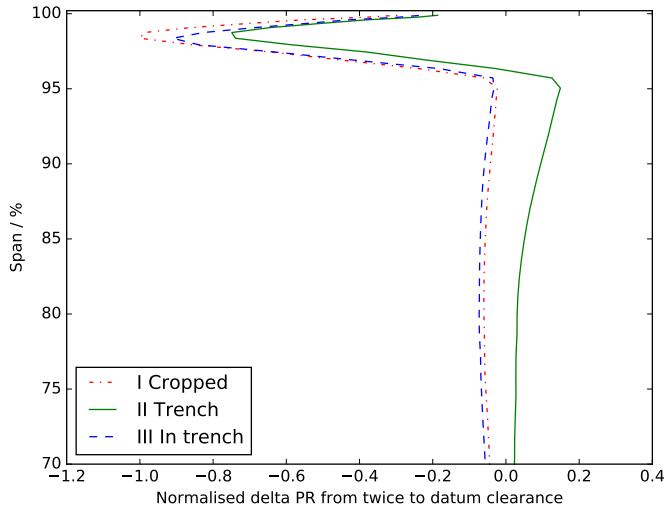

Fig. 6.42 Delta PR radial profile from twice to datum clearance for each case.

The results demonstrate the differences between cropping the blade tip and varying clearance through a trench, and the mechanisms that lead to differences between the tip region flows and overall performance values. To gain a more complete understanding and investigate which regions of the clearance and trench are most critical to control performance, a more complete parameterisation is needed.

6.11 Full parametric study

To fully understand the effect of casing liner trenches and the impact of the large variety of possible casing shapes, a detailed parametric model was developed. The parametric model is able to generate casing liner profiles and varying tip clearances to be used within the PADRAM geometry and meshing system. The parametric model consists of eight parameters, allowing varying blade clearances and trench depths from LE to TE, with different up and downstream trench step heights and axial positions. Figure 6.43 shows the parameters controlled in the model.

The parametric model allows chord-wise variation of both the casing and blade geometries. The blade tip and casing profiles are generated through a linear interpolation of the specified parameters. It is necessary to model chord-wise variation of the trench depth, as this occurs in an engine as the blade cuts differently at the leading and trailing edges. Uneven blade tip clearances from LE to TE can also occur in an engine due to the varying untwist of the blade during expansion.

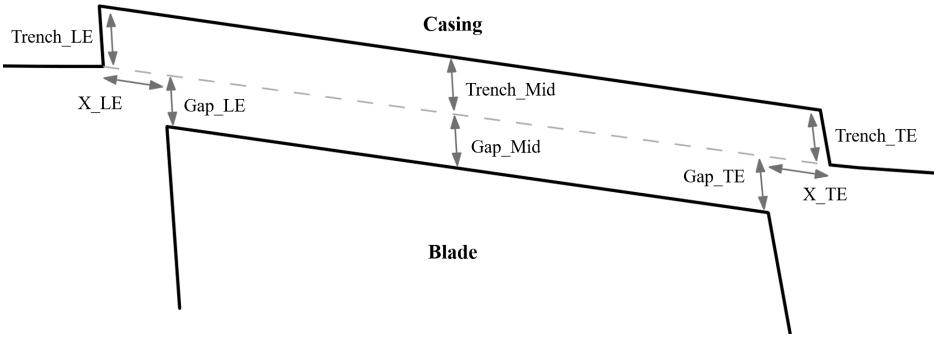

Fig. 6.43 Representation of the parametric definition of the geometry.

6.11.1 Design of experiment

To understand the influence of each parameter on the blade performance, simulations of various combinations of the parameters were carried out. To plan the set of design points to simulate, a design of experiments (DoE) approach was used. Through the Rolls-Royce optimisation toolbox SOFT [107] a latin hypercube [75] design of experiments was set up. Latin hypercube sampling is a method for generating a near-random distribution of parameter values for a multidimensional system. Unlike truly random sampling, where there is no guarantee of the distribution of sample points, latin hypercube sampling attempts to ensure a good representation of the real variability by ensuring the design space is evenly covered [75]. This was used to maximise the relevance of the results obtained through the surrogate modelling approach to be used.

The trench and gap parameters were allowed to vary up to three times the datum clearance, while the axial step positions could vary up to five times. Fifty sample points from the design space were selected to be simulated through latin hypercube sampling. The objective function extracted from each simulation and to be used in further analysis was efficiency.

6.11.2 Surrogate model

To find the dependence of blade performance on each parameter, a polynomial response surface analysis method (RSM) was used [16]. This is a surrogate modelling method that allows an empirical model to be built from the results of the simulations carried out, allowing the response (blade efficiency) to be predicted given the input variables (geometric parameters). A first-order polynomial response surface was fitted to the data. This allows a coefficient for each parameter to be found, allowing an equation to be formed that relates the blade performance to each of the input parameters.

Figure 6.44 shows a comparison of the first-order RSM with the data to which it was fitted for each point in the DoE. It can be seen that the model fits the data well, accurately reflecting the variation in efficiency of each design. This gives confidence that the model is reflecting the physics captured

by the CFD and that the linear trends predicted are a good representation of the real effects of each parameter.

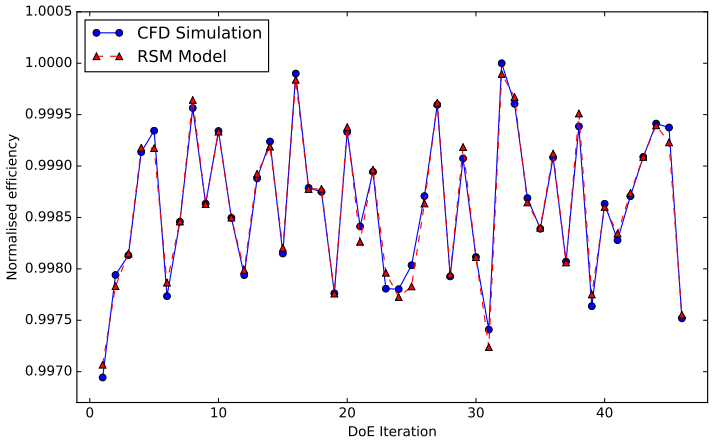

Fig. 6.44 Comparison of RSM fit to CFD data.

The form of the generated first order polynomial is given in Equation 6.1:

$$f = X_1 \times \mathrm{C}_1 + X_2 \times \mathrm{C}_2 + ... + X_n \times \mathrm{C}_n + const \tag{6.1}$$

where f is the blade efficiency, X represents a parameter and C a coefficient. The calculated coefficients (C) for each parameter found via the response surface method are given in Table 6.1. A negative coefficient means blade efficiency reduces with increased magnitude of that parameter. The constant in the equation is the efficiency value that a geometry with zero clearance has, equivalent to setting all gap and trench parameters to zero.

Figure 6.45 shows a comparison of previously obtained simulated results from Figure 6.37 versus predictions from the response surface model. This can be used to assess the ability of the polynomial response surface to represent the real effect of each parameter. The specific points shown here were not sampled during the DoE analysis and therefore provide a chance to further assess the accuracy of the RSM. It can be seen that the polynomial equation obtained predicts well the gradient of efficiency with clearance for the two cases.

The coefficients for the polynomial equation found via the RSM analysis indicate the importance of each parameter. It can be seen in Table 6.1 that the axial position of the trench steps is of far less importance than the clearance magnitudes. While the separation at the steps may have an influence, the position of the steps relative to the blade LE/TE is not that critical. The relationships in Table 6.1 also indicate that the mid-trench and mid-blade clearance parameters are the most important. However, to understand which region has the most influence on the blade performance, the relationship of the clearance area to the parameters in the model must also be considered. If every parameter

Table 6.1 Model coefficients for each geometric parameter.

	Performance sensitivity %/mm		
Parameter	LE	Mid	TE
Trench	-0.008	-0.150	-0.117
Gap	-0.103	-0.307	-0.028
Step position	0.004	-	-0.009

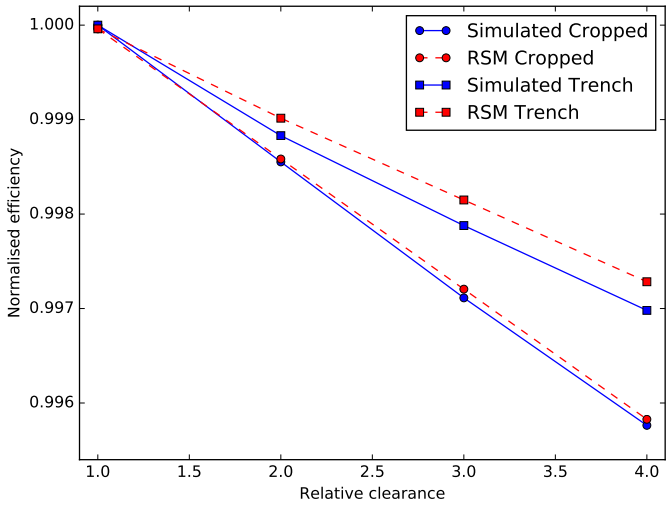

Fig. 6.45 Comparison of simulated and predicted efficiency variation with clearance.

in the geometric model were set to zero, then the LE and mid-trench values increased in turn, it would be clear that increasing the mid value increases the clearance area by twice as much as the LE parameter, see Figure 6.43. Hence, the mid values (of both blade and trench) geometrically have twice the influence on tip leakage as the LE or TE in this model. To understand the effect of equivalent variations in clearance at the LE, mid and TE to form a clearer understanding, the RSM coefficients must be normalised by the clearance area that they influence. The result of this is given in Table 6.2.

6.11.3 Explanation of the coefficients

Table 6.2 gives a comparison of the influence on blade performance of different regions of the tip/trench clearance. It can be seen that overall the sensitivity of performance is greater due to cropping the blade tip than increasing the depth of trench, reflecting the results found earlier in section 6.10. The chord-wise distribution of this sensitivity varies however between the gap and trench.

The sensitivity to gap clearance is biased towards the LE half of the blade, between the LE and mid-chord. This can be explained by understanding the method of tip leakage vortex formation for

Table 6.2 Trench and gap coefficients adjusted by their influence on clearance area and normalised.

	Adjusted performance sensitivity		
Parameter	LE	Mid	TE
Trench	-0.05	-0.49	-0.76
Gap	-0.67	-1.00	-0.18

this case. Due to the nature of the passage shock controlling the tip leakage distribution, the tip leakage vortex is initiated at about thirty percent of chord. It is therefore around this point that the leakage mass flow will have the greatest impact on the vortex formation, and therefore changes to clearance in this area will have the greatest impact on blade performance. The negative influence of increasing the gap is also reduced at the TE, due to the fact that this can redistribute the flow through the span and partially alleviate separation, as was shown in Figure 6.33.

The sensitivity of performance to the trench clearance magnitudes is reduced compared to the gap sensitivities, due to the partially positive effect that adding a trench and raising the clearance has. This can be seen in the individual sensitivities in Table 6.2. The LE trench clearance is far less critical than the LE gap clearance, due to the benefit that it provides by moving the critical region of vortex formation radially, out of the passage flow. The TE trench parameter can be seen to be significant however, and the impact of the trench is biased towards the TE. This shows the difference between the impact of the LE and TE trench steps. The TE trench step acts as a FFS, and the separation due to this generates losses and enters the passage flow. The LE BFS separation does not have the same impact, any separation is confined to the corner of the LE trench step, remaining out of the flow path, and does not have any influence on the downstream flow.

6.11.4 Additional impact of clearance variations

Aside from variations in the tip region, an additional effect of varying clearances and trench geometries that must be noted is the influence of the tip flow behaviour on flow in the remainder of the span. This was also noted in Section 6.7. Changes in blockage and local velocities at the tip can slightly influence the radial distribution of mass flow and velocities for the rest of the blade. An impact of this can be seen in Figure 6.46.

This shows that separation at the blade TE is affected by different clearance geometries and sizes. The effect of this can also be seen in Figure 6.42, where localised changes at the tip cause an additional effect from 95% span and below. This will result in a secondary impact on the blade performance, aside from changes localised to the tip region, and the performance variations with clearance discovered for each trench parameter also include these effects.

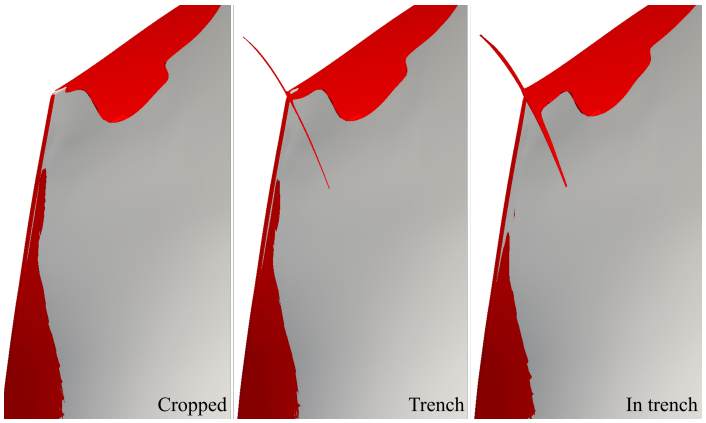

Fig. 6.46 Variation of trailing edge separation due to different clearance types.

The result of this is that the performance trends found are influenced by the overall blade design and aerodynamic behaviour away from the tip region. This behaviour is specific to this fan blade design, and therefore the relationships found may not be directly applicable to blades with different designs. This must be considered when trying to translate tip performance trends found for one blade to another, or develop generic relationships.

6.12 Guidance for fan clearance design

The primary objective of this work was to develop an understanding of clearance effects for a modern fan blade of similar design as that to be used in the Ultrafan. It was discovered that it is beneficial to allow the clearance to be formed as much as possible by the trench, reducing the clearance below the datum streamline. The recommended cold build clearance is therefore as tight as possible, as this will maximise trench depth and reduce the clearance below the datum casing. The model developed can be used in a detailed study of lifetime impacts of various cold build clearances to understand these effects in more detail.

6.13 Conclusion

This chapter has investigated the tip region aerodynamics of a modern fan blade. The tip flow features, tip leakage distribution, clearance vortex size and tip loss have all been studied for the datum design, as well as the effect of increasing clearance through cropping the blade tip. It was shown how an unusual tip leakage distribution exists due to the passage shock controlling the chord-wise vortex initiation location.

It was demonstrated that the tip leakage flow behaviour is significantly influenced by the operating point on the characteristic. As the shock position varies, it controls the tip leakage flow and can even cause reverse tip leakage mass flow, from suction to pressure side.

The tip leakage behaviour can influence the fan behaviour in different parts of the span. Blockage in the tip region re-distributes flow through the span and this can alleviate separation on the blade surface. This was predicted to influence the fan performance variation with clearance at different operating points and for sloped clearance variations and can also affect the fan stall margin.

A parametric model was used to compare the effect of increasing clearance through cropping the blade tip, creating a trench above the blade tip and modifying the blade length inside of a trench. It was predicted that a benefit is achieved by increasing the clearance via a trench compared to cropping the blade tip. This is due to lower tip region loss being generated when the clearance is raised radially out of the main passage flow.

A response surface method was used to fit a polynomial equation to a set of sample points selected by Latin hypercube design of experiments. The polynomial equation was shown to have accurately captured the dependence of blade efficiency to each parameter. It was confirmed that blade performance is more sensitive to cropping of the blade tip than increasing the clearance through a trench. It was predicted that for the blade clearance the key region is between the LE and mid-chord, as this is where the tip leakage vortex is initiated. It was demonstrated that the axial position of the trench steps is generally of less importance than the clearance magnitudes. The TE of the trench is of greater importance than the LE, due to the extra loss caused by the FFS into the passage flow.

This work has discovered performance relationships for tip geometries and clearances that can be used to influence fan blade design and understand engine performance degradation in-service.

It is apparent from this work however that these relationships depend on the specific tip aerodynamics of the blade under investigation. The presence of a passage shock, the tip leakage distribution, shock-vortex interaction and overall blade design influence the tip flow behaviour and associated performance drop with increased clearance. Uncertainty is also present due to the steady nature of the simulations used. Repeating similar analysis for another fan blade would therefore give more confidence in the model developed, as would experimental testing to validate the model, if this were possible.

The results of this study will be used to provide engine lifetime performance improvement through the optimum setting of fan blade clearances. In fact, Rolls-Royce have already adopted a version of the model presented to set the cold-build clearances of their future engines.

6.14 Summary

This chapter has investigated the tip region behaviour of a modern fan blade. The following chapter summarises all of the work carried out in this thesis, and suggests potential future work.

Chapter 7

Conclusions

7.1 Summary

Various investigations have been carried out in this thesis. The major findings from each of these are presented here and the opportunities for future work discussed.

7.2 Major findings

The aim of this work was to investigate novel methods to improve aero-engine performance. The work carried out has successfully demonstrated potential methods to:

- reduce engine fuel consumption through compressor blade efficiency improvement (via novel blade shaping)
- optimise engine build tip clearances to improve engine lifetime fuel consumption
- improve blade stall margin through novel shaping and shock control
- improve engine operability through novel engine intake design, and reduce drag by allowing thinner, lighter intake designs thanks to shock bumps reducing separation

It was shown in Chapter 3 that novel compressor blade shaping, through the use of a free-form parameterisation approach, offers the potential to achieve increased blade efficiency over alternative methods. Combining 3D blade shaping with the re-design of aerofoil sections allows the maximum benefit to be found. An s-shaped design was discovered that alleviates shock-induced separation and increases efficiency. It was demonstrated that the off-design performance of the optimised blade was satisfactory, meeting the same choke mass flow (once skewed) and achieving reasonable performance at different rotational speeds. The uncertainty present in the results due to the difficulty validating the simulations must be considered however.

Chapter 4 investigated the use of shock control bumps on engine intakes. It was concluded that shock control bumps are capable of reducing the shock strength on engine inlets and the resulting separation. It was predicted that shock control bumps can completely eliminate shock-induced separation near to the limiting angle of incidence and, crucially, extend the un-separated angle of incidence by at least three degrees. Due to the simplified simulation approach that enabled optimisation to be conducted, uncertainties remain in the results presented. It would be desirable to carry out more advanced simulations and possibly experimental testing to validate the benefit predicted for the optimised designs.

The use of shock control bumps on transonic blades was also predicted to improve their performance in Chapter 5. It was demonstrated that shock control can delay and reduce separation for a transonic compressor blade, increasing its efficiency without the need for 3D deformations of the blade. The benefit of shock control bumps in reducing shock-induced separation was also demonstrated for a fan blade, predicting a small increase in stall margin. This chapter also highlighted the difficulties with using shock bumps at off-design conditions however. The negatives associated with poor off-design performance may outweigh the benefit of their novel design. Further investigation would therefore be beneficial to improve these designs.

The tip clearance and associated issues for modern fan blades were investigated in detail in Chapter 6. The influence of the trenches that become worn in the abradable liner was analysed and a model developed that can predict fan blade efficiency for combinations of non-uniform tip clearances and liner trenches. This model allows designers to predict fan efficiency variation over engine lifetime and decide what cold build clearances to select for optimal lifetime performance. It was also shown how tip leakage behaviour can influence the flow and separation in other regions of the span, and the influence that the shock position has on the leakage aerodynamics. It is apparent from this work that these relationships depend on the specific tip aerodynamics of the blade under investigation, repeating similar analysis for another fan blade would therefore give more confidence in the model developed, as would experimental testing to validate the model.

Through the use of the approaches demonstrated and knowledge developed in this thesis, engineers have the potential to improve future engine designs in terms of efficiency, fuel consumption and also engine operability.

7.3 Recommendations for future work

In this section recommendations are given for future work relevant to each of the investigations carried out in this thesis.

7.3.1 Novel compressor blade shaping

The novel s-shape produced was shown to be beneficial, but a further consideration could be the investigation of the aero-elastic behaviour of the optimised design. Any instability of the weakened shock could potentially cause amplified forcing or flutter at lower rotational speeds – full multi-disciplinary analysis would be needed to assess this.

It would also be interesting to investigate the impact of the FFD grid 'resolution' on the result of the optimisation. In this work a 3x3 grid was used, controlled at three radial heights. Would a 4x4 or 5x5 grid provide any additional benefit? What is the optimum number of parameters that should be used? The use of an adjoint based optimisation approach would provide significant computational cost savings for research in this area. Adjoint optimisation uses the surface sensitivities to guide the optimisation and means that any increase in the number of design parameters would not significantly increase the computational cost, even with hundreds, or thousands of FFD grid points.

Finally, due to the uncertainty present in the results due to the difficulty validating the simulations, experimental investigation would be useful to assess how well this kind of design performs when tested.

7.3.2 Engine inlet shaping using shock control bumps

This work has demonstrated some promising results, and there are several areas where investigations could continue.

One issue with the approach used here was that the gradient-based optimisation struggled to find beneficial designs at a separated angle of attack, unless it was provided with an already-optimised starting geometry. To avoid the need to optimise at an un-separated incidence first, a global optimisation approach could be employed. Such an optimisation method, such as genetic algorithm [76], would allow the optimisation of shock bumps directly at the separated condition from the start.

A multi-objective approach might aid the discovery of designs that can reduce separation at the design incidence, whilst also keeping to a minimum the increase in pressure loss at lower incidence.

Varying the number of individual bumps in the theta direction might help discover individual bump designs that perform as well as, or better than, the optimised continuous bump designs.

It would be useful to assess the bump designs that have shown benefit through this analysis using a more accurate simulation approach; for example unsteady CFD or including a full fan model (that can represent the suppressive effect of the fan on intake separation) to gain a fuller understanding of their potential. If possible, experimental testing would also be useful to validate the findings.

The application of bumps to aid intake performance during crosswind would also be interesting. During crosswind on runway conditions, the wind passing over the nacelle from the side can lead to extremely high effective incidences for the side lip of the nacelle, leading to very strong shocks and separated flow worse than seen in this work. This has the potential to stall the engine during taxi.

The application of bumps on the side of the inlet would therefore have the potential to alleviate the catastrophic engine blockage that can occur in these conditions.

7.3.3 Compressor blade shaping using shock control bumps

This work demonstrated the potential benefit of shock control bumps on fan/compressor blades, but also how their use can adversely affect performance away from the point at which they are designed. To find improved benefit from the use of shock control bumps, a multi-point optimisation approach could be used, where the optimisation balances the performance between the design and off-design conditions. This would ensure that any benefit achieved at the design point does not significantly deteriorate the performance elsewhere.

Unsteady analysis of the optimised and datum RR-FAN-2 design would make it possible to assess whether reducing the shock-induced separation on the stall side of the characteristic does indeed improve the stall margin of the fan, as suggested by the steady simulations which have limited ability to accurately predict stall margin.

7.3.4 Fan blade casing liner investigations

It would be interesting to apply the methodology developed in this work to other fan/compressor blades to understand how specific the trends presented are to this blade design, and whether generic relationships can be obtained.

Experimental testing of the RR-FAN blade with different clearances or trench depths would help to validate the model developed.

An additional feature of abradable casing liners (whose impact on the tip aerodynamics is not well understood) is the surface roughness. It would be useful to understand whether the roughness level of the abradable liner has an impact on the results presented. To understand the impact of abradable liner roughness, measurements would need to be taken of various worn liners from in-service engines. This work was begun, but due to various difficulties and time constraints was stopped. To continue this research, the roughness measurement method would need to be improved and made more robust, as currently large discrepancies are seen between the hand-held and the white light interferometer approaches. A key feature of the abradable liner surface studied during these initial investigations was the hexagonal sub-structure of the abradable material. It would be necessary to assess whether this is a common feature of all fan abradable liners and critically what impact this has on the effective roughness values. Can the simple roughness correlations currently used in CFD simulations properly replicate the effect of this structure? The Hydra roughness model would need further extensive validation and also possibly improvement to be able to accurately simulate the effects of abradable liner roughness.

7.3.5 Novel tip shaping

Studies carried out by the author (not included in this thesis) have shown that there is some efficiency benefit achievable through novel blade tip shaping.

The benefit is small however, and the cost of manufacturing highly complex blade tip geometries would possibly outweigh the efficiency benefit achieved. To become viable, novel tip shaping would need to provide an additional benefit. A potential avenue for this, and relating to the previous investigations of casing liner wear, is in the design of tip shapes that provide a benefit when cutting abradable casing liners.

With the latest Rolls-Royce fan blades being manufactured from composite materials, there is the concern that the blade tips might be susceptible to damage when rubbing (through delamination). This could be avoided through increasing the clearance until the blade tips do not rub, though the cost of this in terms of fan efficiency and performance is apparent from Chapter 6 (the model and coefficients provided could help to inform this decision).

An alternative method would be to use a titanium insert on the tip of the blade. This can allow rubbing to occur (as was shown to be beneficial in Chapter 6) without the concern of damage to the blade tip. When adding this additional tip insert, it might be possible to shape this in a way that could reduce tip leakage and provide a performance benefit, with the cost of this addition already partially covered due to its planned use as a tip insert for rubbing.

An alternative approach (which could also be machined onto metallic compressor blades) could be to simply slope the blade tip from pressure to suction side, which would ensure that the pressure (leading) side is always the part of the tip that rubs, avoiding the possibility of the tip being 'peeled' apart and delamination occurring.

Hence, there are a few avenues for investigation:

- analysing the cost benefit of using a tip insert versus cropping the blade tip to avoid rubbing
- carrying out more analysis of novel tip designs to select a design that would be feasible to manufacture, whilst still providing benefit
- investigating sloped gaps, and studying whether sloping the gap has any benefit/loss from an aerodynamic perspective, considering its mechanical benefit

References

[1] Adamczyk, J. (2015). Private communication.

[2] Adamczyk, J., Celestina, M., and Greitzer, E. (1991). The role of tip clearance in high-speed fan stall. In *ASME 1991 International Gas Turbine and Aeroengine Congress and Exposition*, pages 109–115. American Society of Mechanical Engineers.

[3] Ahn, C. and Kim, K. (2003). Aerodynamic design optimization of a compressor rotor with navier—stokes analysis. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of power and energy*, 217(2):179–183.

[4] Aircraft-Engine-Historical-Society (May 2017). Online - http://www.enginehistory.org/members/conway.php.

[5] Anderson, M., Katz, V., and Wilson, R. (2009). *Who gave you the epsilon? and other tales of mathematical history*. The Mathematical Association of America, USA.

[6] Ashill, P. and Fulker, J. (1992). A novel technique for controlling shock strength of laminar-flow aerofoil sections. In *The First European Symposium on Laminar Flow, Hamburg, March 1992*, pages 175–183.

[7] BBC (Online - accessed May 2017). Frank Whittle (1907 - 1996).

[8] Beheshti, B. H., Teixeira, J. A., Ivey, P. C., Ghorbanian, K., and Farhanieh, B. (2004). Parametric study of tip clearance—casing treatment on performance and stability of a transonic axial compressor. *Journal of turbomachinery*, 126(4):527–535.

[9] Belegundu, A. D. and Chandrupatla, T. R. (2011). *Optimization concepts and applications in engineering*. Cambridge University Press.

[10] Bellis, M. (Online - accessed May 2017). The history of the jet engine. 'https://www.thoughtco.com/history-of-the-jet-engine-4067905'.

[11] Benini, E. (2004). Three-dimensional multi-objective design optimization of a transonic compressor rotor. *Journal of propulsion and power*, 20(3):559–565.

[12] Bergner, J., Hennecke, D. K., Hoeger, M., and Engel, K. (2003). Darmstadt rotor no. 2, ii: Design of leaning rotor blades. *International Journal of Rotating Machinery*, 9(6):385–391.

[13] Biollo, R. (2008). *Systematic Investigation On Swept And Leaned Transonic Compressor Rotor Blades. Scuola di Dottorato di Ricerca in Ingegneria Industriale Indirizzo Energetica*. PhD thesis.

[14] Boldman, D. R., Iek, C., Hwang, D., Larkin, M., and Schweiger, P. (1993). Effect of a rotating propeller on the separation angle of attack and distortion in ducted propeller inlets. NASA-TM-105935.

[15] Boussinesq, J. (1877). Théorie de l'écoulement tourbillant. *Mem. Présentés par Divers Savants Acad. Sci. Inst. Fr*, 23(46-50):6–5.

[16] Box, G. E. and Wilson, K. (1992). On the experimental attainment of optimum conditions. In *Breakthroughs in Statistics*, pages 270–310. Springer.

[17] Brooks, C. J., Forrester, A., Keane, A., and Shahpar, S. (2011). Multi-fidelity design optimisation of a transonic compressor rotor. In *In: 9th European Turbomachinery Conference, Istanbul, Turkey, 21st-25th March*.

[18] Bruce, P. and Colliss, S. (2015). Review of research into shock control bumps. *Shock Waves*, 25(5):451–471.

[19] Bruce, P. J., Colliss, S., and Babinsky, H. (2014). Three-dimensional shock control bumps: effects of geometry. In *AIAA 52nd Aerospace Sciences Meeting, National Harbor, Maryland*, number - AIAA 2014-0943.

[20] Campobasso, M., Duta, M., and Giles, M. (2001). Adjoint methods for turbomachinery design. In *ISABE Conference*. Citeseer.

[21] Cao, T., Vadlamani, N. R., Tucker, P. G., Smith, A. R., Slaby, M., and Sheaf, C. T. (2017). Fan–intake interaction under high incidence. *Journal of Engineering for Gas Turbines and Power*, 139(4):041204.

[22] Chen, N., Zhang, H., Xu, Y., and Huang, W. (2007). Blade parameterization and aerodynamic design optimization for a 3d transonic compressor rotor. *Journal of Thermal Science*, 16(2):105–114.

[23] Chilla, M., Hodson, H., and Newman, D. (2013). Unsteady interaction between annulus and turbine rim seal flows. *Journal of Turbomachinery*, 135(5):051024.

[24] Chima, R. (2009). Swift code assessment for two similar transonic compressors. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, number - AIAA 2209-1058.

[25] Christie, R., Heidebrecht, A., and MacManus, D. (2017). An automated approach to nacelle parameterization using intuitive class shape transformation curves. *Journal of Engineering for Gas Turbines and Power*, 139(6):062601.

[26] Colliss, S. P., Babinsky, H., Nübler, K., and Lutz, T. (2016). Vortical structures on three-dimensional shock control bumps. *AIAA Journal*, 54(8):2338–2350.

[27] Connors, J. (2010). *The Engines of Pratt & Whitney: A Technical History: As Told by the Engineers Who Made the History: Birth of the Two-Spool Turbojet. pp. 215-256*. American Institute of Aeronautics and Astronautics.

[28] Coquillart, S. (1990). Extended free-form deformation: a sculpturing tool for 3d geometric modeling. In *SIGGRAPH '90, In Computer Graphics*, volume 24, pages 197–196.

[29] Courant, R., Friedrichs, K., and Lewy, H. (1928). Über die partiellen differenzengleichungen der mathematischen physik. *Mathematische annalen*, 100(1):32–74.

[30] Crumpton, P., Mackenzie, J., and Morton, K. (1993). Cell vertex algorithms for the compressible navier-stokes equations. *Journal of Computational Physics*, 109(1):1–15.

[31] Crumpton, P., Moinier, P., and Giles, M. (1997). An unstructured algorithm for high reynolds number flows on highly stretched grids. *Numerical methods in laminar and turbulent flow*, pages 561–572.

[32] Cumpsty, N. (2010). Some lessons learned. *Journal of turbomachinery*, 132(4):041018.

[33] Cumpsty, N., Alonso, J., Eury, S., Maurice, L., Nas, B., Ralph, M., and Sawyer, R. (2010). Report of the independent experts on the medium and long term goals for aviation fuel burn reduction from technology. *International Civil Aviation Organization, Doc*, 9963.

[34] Cumpsty, N. A. (1989). *Compressor aerodynamics*. Longman Scientific & Technical.

[35] Dawes, W. (2006). Eliminating serial bottlenecks not just in the flow solving but also in mesh generation and geometry management. In *ASME Turbo Expo 2006: Power for Land, Sea, and Air*, pages 1355–1362. American Society of Mechanical Engineers.

[36] Denton, J. (1997). Lessons from rotor 37. *Journal of Thermal Science*, 6(1):1–13.

[37] Denton, J. D. (1993). Loss mechanisms in turbomachines. In *ASME 1993 International Gas Turbine and Aeroengine Congress and Exposition*, pages 621—656. American Society of Mechanical Engineers.

[38] Drela, M. and Giles, M. B. (1987). Viscous-inviscid analysis of transonic and low Reynolds number airfoils. *AIAA journal*, 25(10):1347–1355.

[39] Dunham, J. (1998). Cfd validation for propulsion system components (la validation cfd des organes des propulseurs). Technical report, DTIC Document.

[40] Duta, M. C., Shahpar, S., and Giles, M. B. (2007). Turbomachinery design optimization using automatic differentiated adjoint code. In *ASME Turbo Expo 2007: Power for Land, Sea, and Air*, pages 1435–1444. American Society of Mechanical Engineers.

[41] Electric, G. (1979). *Seven decades of progress*. Aero Publishers Inc. p.76.

[42] Farokhi, S. (2014). *Aircraft propulsion*. John Wiley & Sons.

[43] Fois, N., Watson, M., and Marshall, M. (2016). The influence of material properties on the wear of abradable materials. *Proceedings of the Institution of Mechanical Engineers, Part J: Journal of Engineering Tribology*, 231:240–253.

[44] Freeman, C. and Rowe, A. L. (1999). Intake engine interactions of a modern large turbofan engine. In *ASME 1999 International Gas Turbine and Aeroengine Congress and Exhibition*, number - ASME 99-GT-344. American Society of Mechanical Engineers.

[45] Furukawa, M., Inoue, M., Saiki, K., and Yamada, K. (1999). The role of tip leakage vortex breakdown in compressor rotor aerodynamics. *J. Turbomach*, 121(3):469–480.

[46] General, . S. (2010). Assembly resolutions in force (as of 8 october 2010). Technical report, International Civil Aviation Organization.

[47] Ginder, R., Calvert, W., et al. (1987). The design of an advanced civil fan rotor. *Transactions of the ASME. Journal of Turbomachinery*, 109(3):340–345.

[48] Hah, C. (2009). Large eddy simulation of transonic flow field in nasa rotor 37. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, number - AIAA 2009-1061.

[49] Hall, C. A. and Hynes, T. P. (2006). Measurements of intake separation hysteresis in a model fan and nacelle rig. *Journal of propulsion and power*, 22(4):872–879.

[50] Hicks, R. M. and Henne, P. A. (1978). Wing design by numerical optimization. *Journal of Aircraft*, 15(7):407–412.

[51] Hield, P. (Compressor Aerodynamics Specialist Rolls-Royce) (2015). Private communication.

[52] Hinchliffe, B. L. (2016). *Using Surface Sensitivity for Adjoint Aerodynamic Optimisation of Shock Control Bumps*. PhD thesis, University of Sheffield.

[53] Hirsch, C. (1988). *Numerical computation of internal and external flows*. Wiley.

[54] Hirsch, C. (2007). *Numerical computation of internal and external flows: The fundamentals of computational fluid dynamics*. Butterworth-Heinemann.

[55] Houghton, T. and Day, I. (2012). Stability enhancement by casing grooves: The importance of stall inception mechanism and solidity. *Journal of Turbomachinery*, 134(2).

[56] Jameson, A. (1981). Transonic airfoil calculations using the euler equations. In *IMA Conference on Numerical Methods in Aeronautical Fluid Dynamics, Reading*, pages 289–308.

[57] Jameson, A. (1989). Aerodynamic design via control theory. In *Recent advances in computational fluid dynamics*, pages 377–401. Springer.

[58] Jameson, A. (2017). Origins and further development of the Jameson–Schmidt–Turkel scheme. *AIAA Journal*, 55(5):1487–1510.

[59] Jameson, A. and Mavriplis, D. (1986). Finite volume solution of the two-dimensional euler equations on a regular triangular mesh. *AIAA journal*, 24(4):611–618.

[60] Jang, C.-M., Samad, A., and Kim, K.-Y. (2006). Optimal design of swept, leaned and skewed blades in a transonic axial compressor. In *ASME Turbo Expo 2006: Power for Land, Sea, and Air*, pages 1279–1288. American Society of Mechanical Engineers.

[61] John, A., Shahpar, S., and Qin, N. (2017). Novel compressor blade shaping through a free-form method. *Journal of Turbomachinery*, 139(8):TURBO–16–1228.

[62] Kammer, A. S. and Olgac, N. (2016). Blade/casing rub interaction in turbomachinery: Structural parameters' influence on stability. *Journal of Propulsion and Power*, 32(1):1–10.

[63] Kharina, A. and Rutherford, D. (2015). Fuel efficiency trends for new commercial jet aircraft: 1960 to 2014, retrieved from the international council on clean transportation, 'http://www.theicct.org/fuel-efficiency-trends-new-commercial-jetaircraft-1960-2014'.

[64] Korting, P. and Beacher, B. (1989). Improved compressor performance using recessed clearance (trenches). *Journal of Propulsion and Power*, 5(4):469–475.

[65] Krein, A. and Williams, G. (2012). Flightpath 2050: europe's vision for aeronautics. *Innovation for Sustainable Aviation in a Global Environment: Proceedings of the Sixth European Aeronautics Days, Madrid, 30 March-1 April, 2011*, page 63.

[66] Kulfan, B. M., Bussoletti, J. E., et al. (2006). Fundamental parametric geometry representations for aircraft component shapes. In *11th AIAA/ISSMO multidisciplinary analysis and optimization conference*, volume 1, pages 547–591.

[67] Lapworth, L. (2004). Hydra-cfd: a framework for collaborative cfd development. In *International Conference on Scientific and Engineering Computation (IC-SEC), Singapore, June*, volume 30.

[68] Larkin, M. J. and Schweiger, P. S. (1992). Ultra high bypass nacelle aerodynamics inlet flow-through high angle of attack distortion test, nasa cr-189149. (NASA CR-189149).

[69] Lee, D. S., Fahey, D. W., Forster, P. M., Newton, P. J., Wit, R. C., Lim, L. L., Owen, B., and Sausen, R. (2009). Aviation and global climate change in the 21st century. *Atmospheric Environment*, 43(22):3520–3537.

[70] Legrand, M., Batailly, A., and Pierre, C. (2012). Numerical investigation of abradable coating removal in aircraft engines through plastic constitutive law. *Journal of Computational and Nonlinear Dynamics*, 7(1):011010.

[71] Luo, H., Baum, J. D., and Lohner, R. (1994). Edge-based finite element scheme for the euler equations. *AIAA journal*, 32(6):1183–1190.

[72] Ma, H., Yin, F., Guo, Y., Tai, X., and Wen, B. (2016). A review on dynamic characteristics of blade–casing rubbing. *Nonlinear Dynamics*, 84(2):437–472.

[73] Martinelli, L. (1987). Calculations of viscous flows with a multigrid method. Technical report, Princeton Univ., NJ (USA).

[74] Mazaheri, K. and Khatibirad, S. (2017). Using a shock control bump to improve the performance of an axial compressor blade section. *Shock Waves*, 27(2):299–312.

[75] McKay, M. D., Beckman, R. J., and Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61.

[76] Mitchell, M. (1998). *An introduction to genetic algorithms*. MIT press.

[77] Moinier, P. (1999). *Algorithm developments for an unstructured viscous flow solver*. PhD thesis, Oxford University.

[78] Moinier, P., M-uacute, J.-D., ller, and Giles, M. B. (2002). Edge-based multigrid and preconditioning for hybrid grids. *AIAA journal*, 40(10):1954–1960.

[79] Morgan, K., Peraire, J., Peiro, J., and Hassan, O. (1991). The computation of three-dimensional flows using unstructured grids. *Computer Methods in Applied Mechanics and Engineering*, 87(2-3):335–352.

[80] Müller, J.-D. and Giles, M. B. (1998). Edge-based multigrid schemes for hybrid grids. In *Numerical Methods for Fluid Dynamics, VI, ICFD*. Citeseer.

[81] Nezym, V. (2005). Parametric investigation of entire annular recess casing treatment on compressor stable operation. *Experimental thermal and fluid science*, 29(2):209–215.

[82] Nishioka, T. (2013). Inception mechanism and suppression of rotating stall in an axial-flow fan. In *IOP Conference Series: Materials Science and Engineering*, volume 52, page 012002. IOP Publishing.

[83] Ogawa, H., Babinsky, H., Pätzold, M., and Lutz, T. (2008). Shock-wave/boundary-layer interaction control using three-dimensional bumps for transonic wings. *AIAA journal*, 46(6):1442–1452.

[84] Oriji, U. R. and Tucker, P. G. (2014). Modular turbulence modeling applied to an engine intake. *Journal of Turbomachinery*, 136(5):051004.

[85] Orszag, S. A. (1970). Analytical theories of turbulence. *Journal of Fluid Mechanics*, 41(02):363–386.

[86] Peters, A., Spakovszky, Z. S., Lord, W. K., and Rose, B. (2015). Ultrashort nacelles for low fan pressure ratio propulsors. *Journal of Turbomachinery*, 137(2):021001.

[87] Polynkin, A., Toropov, V., and Shaphar, S. (2010). Multidisciplinary optimization of turbomachinary based on metamodel built by genetic programming. In *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, number - AIAA 2010-9397.

[88] POON, R. (Online - accessed May 2017). Rolls-royce at a transitional altitude. 'http://www.aspireaviation.com/2015/08/31/rolls-royce-transitional-altitude/'.

[89] Qin, N., Carnie, G., Wang, Y., and Shahpar, S. (2014). Design optimization of casing grooves using zipper layer meshing. *Journal of Turbomachinery*, 136(3):031002.

[90] Qin, N., Monet, D., and Shaw, S. (2002). 3d bumps for transonic wing shock control and drag reduction. In *CEAS Aerospace Aerodynamics Research Conference*. Royal Aeronautical Society.

[91] Qin, N., Wong, W., and Le Moigne, A. (2008). Three-dimensional contour bumps for transonic wing drag reduction. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 222(5):619–629.

[92] Qin, N., Zhu, Y., Ashill, P., and Shaw, S. (2000). Active control of transonic aerodynamics using suction, blowing, bumps and synthetic jets. In *AIAA 18th Applied Aerodynamics Conference*, number - AIAA 2000-4329.

[93] Reid, L. and Moore, R. D. (1978). Performance of single-stage axial-flow transonic compressor with rotor and stator aspect ratios of 1.19 and 1.26, respectively, and with design pressure ratio of 1.82. *NASA-TP-1659*.

[94] Reynolds, O. (1894). On the dynamical theory of incompressible viscous fluids and the determination of the criterion. *Proceedings of the Royal Society of London*, 56(336-339):40–45.

[95] Robideau, B. A. and Niiler, J. (1980). Blade tip seal for an axial flow rotary machine. US Patent 4,238,170.

[96] Rolls-Royce (June 2017b). Online - https://www.rolls-royce.com/products-and-services/civil-aerospace/airlines/trent-700.aspx.

[97] Rolls-Royce (May 2017a). Online - https://www.rolls-royce.com/about/our-technology/research/research-programmes/sustainable-and-green-engine-sage-itd/open-rotor-technology.aspx.

[98] Rolls-Royce (May 2017c). Online - https://www.rolls-royce.com/products-and-services/civil-aerospace/airlines/trent-xwb.aspx/.

[99] Rolls-Royce (May 2017d). Online - http://www.rolls-royce.com/innovation/advance-and-ultrafan.

[100] Rolls-Royce presentation at ICFD meeting, Reading University, . (2008). The challenges for aero-engine cfd - http://www.icfd.rdg.ac.uk/icfd25/talks/llapworth.pdf. In *ICFD meeting, Reading University*.

[101] Sakulkaew, S., Tan, C., Donahoo, E., Cornelius, C., and Montgomery, M. (2013). Compressor efficiency variation with rotor tip gap from vanishing to large clearance. *Journal of Turbomachinery*, 135(3, 031030):TURBO–12–1132.

[102] Samad, A. and Kim, K.-Y. (2008). Multi-objective optimization of an axial compressor blade. *Journal of Mechanical Science and Technology*, 22(5):999–1007.

[103] Schmidt, J., Moore, R., Wood, J., and Steinke, R. (1987). Supersonic through-flow fan design. In *23rd Joint Propulsion Conference*, number - AIAA-1987-1746.

[104] Seshadri, P., Parks, G. T., and Shahpar, S. (2014a). Leakage uncertainties in compressors: The case of rotor 37. *Journal of Propulsion and Power*, 31(1):456–466.

[105] Seshadri, P., Shahpar, S., and Parks, G. T. (2014b). Robust compressor blades for desensitizing operational tip clearance variations. In *ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*, number - GT2014-26624. American Society of Mechanical Engineers.

[106] Shabbir, A. and Adamczyk, J. J. (2005). Flow mechanism for stall margin improvement due to circumferential casing grooves on axial compressors. *Journal of turbomachinery*, 127(4):708–717.

[107] Shahpar, S. (2001). Soft: A new design and optimisation tool for turbomachinery. *ROLLS ROYCE PLC-REPORT-PNR*.

[108] Shahpar, S. et al. (2005). Sophy: An integrated cfd based automatic design optimization system. *ISABE-2005-1086*.

[109] Shahpar, S., Giacche, D., and Lapworth, L. (2003). Multi-objective design and optimisation of bypass outlet-guide vanes. In *ASME Turbo Expo 2003, collocated with the 2003 International Joint Power Generation Conference*, pages 591–601. American Society of Mechanical Engineers.

[110] Shahpar, S. and Lapworth, L. (2003). Padram: Parametric design and rapid meshing system for turbomachinery optimisation. In *ASME Turbo Expo 2003, collocated with the 2003 International Joint Power Generation Conference*, pages 579–590. American Society of Mechanical Engineers.

[111] Shahpar, S., Polynkin, A., and Toropov, V. (2008). Large scale optimization of transonic axial compressor rotor blades. In *49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, number - AIAA 2008-2056.

[112] Sieverding, C. (1985). *Tip clearance effects in axial turbomachines: April 15-19, 1985*. Von Karman Institute for Fluid Dynamics.

[113] Smith, A. (Engine intake aerodynamics expert, Rolls-Royce) (2017). Private communication.

[114] Smith, A., Babinsky, H., Dhanasekaran, C., Savill, M., and Dawes, B. (2003). Computational investigation of groove controlled shock wave/boundary layer interaction. In *41st Aerospace Sciences Meeting and Exhibit*, number - AIAA 2003-0446.

[115] Smith, A., Holden, H., Babinsky, H., Fulker, J., and Ashill, P. (2002). Control of normal shock wave/turbulent boundary layer interactions using streamwise grooves. In *40th AIAA Aerospace Sciences Meeting & Exhibit*, number - AIAA 2002-0978.

[116] Sommerer, A., Lutz, T., and Wagner, S. (2000). Design of adaptive transonic airfoils by means of numerical optimisation. In *Proceedings of ECCOMASS 2000, Barcelona*.

[117] Spalart, P. R. and Allmaras, S. R. (1992). A one equation turbulence model for aerodinamic flows. *AIAA journal*, (- AIAA 1992-439).

[118] Stanewsky, E. (2002). *Drag Reduction by Shock and Boundary Layer Control: Results of the Project EUROSHOCK II. Supported by the European Union 1996-1999*, volume 80. Springer Science & Business Media.

[119] Stokes, G. G. (1845). On the theories of internal friction of fluids in motion. *Trans. Camb. Philos. Soc.*, 8:287–305.

[120] Storer, J. and Cumpsty, N. (1994). An approximate analysis and prediction method for tip clearance loss in axial compressors. *Journal of Turbomachinery*, 116(4):648–656.

[121] Suder, K. L. (1998). Blockage development in a transonic, axial compressor rotor. *Journal of Turbomachinery*, 120(3):465–476.

[122] Tai, T. C. (1977). Theoretical aspects of dromedaryfoil. Technical report, DTIC Document.

[123] Tian, Y., Liu, P., and Li, Z. (2014). Multi-objective optimization of shock control bump on a supercritical wing. *Science China Technological Sciences*, 57(1):192–202.

[124] Toropov, V., va, F., Markine, V., and d, H. (1996). Refinements in the multi-point approximation method to reduce the effects of noisy structural responses. In *6th Symposium on Multidisciplinary Analysis and Optimization*, number - AIAA 1996-4087.

[125] Wang, X., Wang, S., and Han, W. (2008). Multi-objective aerodynamic design optimization based on camber line and thickness distribution for a transonic compressor rotor. In *ASME 2008 International Mechanical Engineering Congress and Exposition*, pages 25–32. American Society of Mechanical Engineers.

[126] Warwick, G. (Online - accessed May 2017). Snecma prepares for crucial open-rotor tests. 'http://aviationweek.com/commercial-aviation/snecma-prepares-crucial-open-rotor-tests'.

[127] Wilson, A. (2001). Application of CFD to wake/aerofoil interaction noise-a flat plate validation case. In *7th AIAA/CEAS Aeroacoustics Conference and Exhibit*, number - AIAA 2001-2135.

[128] Wong, W. S. (2006). *Mechanisms and Optimisations of 3D Shock Control Bumps*. PhD thesis, University of Sheffield.

[129] Xu, L., Bo, S., Hongde, Y., and Lei, W. (2015). Evolution of rolls-royce air-cooled turbine blades and feature analysis. *Procedia Engineering*, 99:1482–1491.

[130] Zhu, F. and Qin, N. (2015). Using mesh adjoint for shock bump deployment and optimisation on transonic wings. In *53rd AIAA Aerospace Sciences Meeting*, number - AIAA 2015-1488.