High Pressure Turbine Blade Platform Cooling and Feed Architecture



James A. Parker

Department of Engineering

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

University College

Hilary Term 2021

Abstract

In this thesis, novel cooling systems for gas turbine blade platforms were developed and assessed experimentally. A new highly modular facility was designed, manufactured, and commissioned to support investigations of a variety of different platform cooling systems. The facility is a high pressure, high temperature, transonic, linear cascade. A high degree of engine fidelity is achieved by: operating at scaled engine conditions; using test blades with full platform and fir tree root geometry; and including engine-representative internal hub seal structures and coolant feeds. The inclusion of hub seals, realistic coolant feeds, and full test blades allowed engine-feasible systems to be tested and representative leakage flows to be reproduced.

Two novel platform cooling systems were developed, using low-order numerical models, which utilise internal cooling features embedded directly into the platform surface. These systems are: a convective cooling passage design which draws coolant from the front seal, passes it through the platform, and exhausts into the rear seal; and an impingement jet array design which draws coolant from the shank cavity and similarly exhausts into the rear seal. Each system was integrated into cascade geometry rotor blades, manufactured in titanium via additive manufacturing, and used to experimentally determine metal effectiveness distributions across mainstream platform surfaces. Metal effectiveness distributions are determined from IR measurements of surface temperature. The cooling performance of each concept was quantified and compared to a baseline case with no cooling features, cooled solely by leakage flows.

The film cooling performance of front seal purge flow was investigated experimentally, to determine adiabatic film effectiveness, at a range of purge flow angles and mass ratios. The aerodynamic impact was assessed using a downstream area traverse system to determine total pressure loss coefficient distributions. A practical design to alter the purge flow angle in engine architectures was proposed and the viability of the system demonstrated using numerical modelling. A new experimental data processing technique was developed to correct adiabatic film effectiveness distributions determined from experiments with a non-uniform inlet temperature profile, such as that caused by a time-varying inlet thermal boundary layer.

"Sometimes you gotta run before you can walk."

—Tony Stark

To Mum and Dad

Acknowledgements

I must begin by expressing my most sincere gratitude to my supervisor Tom Povey, who not only welcomed me into the dreaming spires of Oxford but also unwaveringly supported me through my many years as his student. Tom's technical expertise is surpassed by none and his infectious enthusiasm and drive always pushed me above and beyond what I thought I was capable of. His confidence in me has always been incredibly motivating and it has been a privilege to develop as an engineer under his tutelage.

I am deeply grateful to everyone in Team Blowdown for their friendship and help. I would particularly like to thank my long-time friend Dan Burdett, who has been there and put up with me since the start. Thank you also to Dominic Harris for his invaluable contributions, without which none of my experimental work would have been possible, and for patiently teaching me the practical aspects of engineering. I am also indebted to Nafiz Chowdhury for his help and willingness to spend long hours running experiments.

A great big thank you to everyone in the Osney Thermofluids Institute, both current and former, who made my time here a thoroughly enjoyable experience. To my CDT companions, Anna Bru-Revert, Daniel Fahy, and Mathias Michaud who know how much I appreciate them. To the lads who helped me out over the years and kept everything running smoothly in the lab: Bill, Craig, Greg, Hal, Jason, Paul, Phil. To the EPSRC CDT in Gas Turbine Aerodynamics for making this work possible and for introducing me to a motley crew of reprobates, who may now be spread far and wide but will never be forgotten.

Thank you also to Rolls-Royce for the technical and financial support they provided for the project and in particular to Ignacio Mayo, Ed Romero and Alex Matthews for their technical insights, genuine interest, and engagement in my work.

To Emilia, for her love and support.

Finally, and most importantly, I would like to thank my parents and sister who have always supported me and shown me endless love, for which I cannot truly express the extent of my gratitude.

Contents

Nomenclature **x**

1 Introduction **1**

1.1 Research objectives 3

1.2 Thesis outline 4

2 Literature review **6**

2.1 The increasing need for platform cooling 6

2.2 The high pressure turbine rotor aerothermal field 11

2.2.1 Rotor aerodynamics 11

2.2.2 Rotor platform heat transfer 17

2.3 Quantifying cooling system performance 20

2.3.1 Metal effectiveness 20

2.3.2 Non-dimensional mass ratio 21

2.3.3 Total pressure loss coefficient 22

2.4 Platform cooling techniques 22

2.4.1 Film cooling 23

2.4.2 Platform geometric modifications 27

2.4.3 Internal cooling 29

2.5 The effect of rotation 33

2.5.1 External coolant films 33

2.5.2 Internal flows 34

2.6 Experimental techniques 35

2.6.1 Surface temperature measurement techniques 35

2.6.2 Experimental facilities 38

3 The Metal Effectiveness Rotor Cooling (MERC) facility **43**

3.1 Facility overview and operating conditions 44

3.1.1 Operational considerations 47

3.1.2 Pressure rating and facility safety 49

3.2 Scaling to engine conditions . . . 50
3.3 Main stack assembly . . . 53
3.4 Test section . . . 53
3.4.1 Turbulence grid . . . 55
3.4.2 Inlet contraction . . . 56
3.4.3 Fast response heater mesh . . . 58
3.5 Instrumentation and data acquisition . . . 60
3.5.1 Coolant orifice plate . . . 62
3.5.2 Downstream traverse cassette . . . 64
3.5.3 Infra-red hardware and capture setup . . . 68
3.6 Linearised rotor blades . . . 70
3.7 Modular blade cassette . . . 73
3.7.1 3D printed SAS insert . . . 75
3.7.2 Rotor blades and disk holder . . . 76
3.7.3 Front and rear seal blocks . . . 77

4 Design of novel integrated platform cooling systems 79
4.1 General design strategy and discussion . . . 79
4.2 Convective cooling passage concept overview . . . 81
4.3 Modelling convective cooling passages . . . 82
4.3.1 Compressible flow through passages with friction and heat addition . . . 83
4.3.2 Low-order model solver . . . 86
4.3.3 Low order model results . . . 91
4.3.4 Limitations of the low-order model . . . 96
4.4 Integrated blade design 1: convective passages . . . 97
4.5 Convective passage experimental capacity . . . 99
4.6 Impinging jet array concept overview . . . 103
4.7 Modelling jet impingement arrays . . . 105
4.7.1 Low-order model solver . . . 108
4.7.2 Low-order model results . . . 113
4.8 Integrated blade design 2: impinging jet array . . . 119

5 MERC facility aerothermal commissioning and processing techniques 121

5.1 Typical operating conditions 121

5.1.1 Inlet temperature profile 125

5.1.2 Heater mesh characteristic 129

5.1.3 Coolant capacity calibration 130

5.2 Downstream area traverse 132

5.2.1 Transonic probes 132

5.2.2 Midspan traverse demonstrating passage periodicity 136

5.2.3 Area traverses with pressure and temperature probes 137

5.3 IR calibration and determination of wall temperature 140

5.3.1 Blackbody calibration 141

5.3.2 Greybody calibration 145

5.4 Determination of metal effectiveness 150

5.4.1 Recovery temperature 150

5.4.2 Determination of compressible mainstream recovery ratio 151

5.5 Summary of IR image processing 154

5.6 Uncertainty of wall temperature and metal effectiveness 156

5.7 Recommendations for improvements to the MERC facility 160

6 Method for determining adiabatic film effectiveness in presence of thermal boundary layer 164

6.1 Abstract 164

6.2 Introduction 166

6.3 Related literature 169

6.4 Theory 172

6.5 Solution of the equations 174

6.6 Experimental demonstration of technique 176

6.6.1 Facility overview and instrumentation 176

6.6.2 Application of the proposed method 180

6.6.3 Stability of the process 187

6.7 Uncertainty with two timestamps 188

6.7.1 Effect of mainstream-to-coolant temperature ratio 190

6.7.2 Effect of timestamp selection . . . 191
6.7.3 Effect of value of equivalent mainstream effectiveness . . . 194
6.8 Method to reduce uncertainty by using multiple timestamps . . . 195
6.9 Conclusions . . . 200
6.10 Appendix A: Extension to compressible flow . . . 202
6.11 Appendix B: Correction for non-ideal adiabatic surface behaviour . . . 206

7 Front seal purge flow angle and platform cooling 211
7.1 Front seal purge flow angle in a rotating frame . . . 212
7.1.1 Practical method to alter purge flow angle . . . 216
7.2 Experimental setup and testing procedure . . . 220
7.3 Adiabatic film effectiveness results . . . 222
7.3.1 Typical film cooling characteristics of front seal purge flow . . . 223
7.3.2 Comparison of blade passage adiabatic film effectiveness distributions . 226
7.3.3 Comparison of axial platform adiabatic film effectiveness profiles . . . 233
7.3.4 Comparison of area-averaged platform adiabatic film effectiveness values 236
7.4 Downstream total pressure loss measurements . . . 237
7.4.1 Comparison of total pressure loss area distributions . . . 237
7.4.2 Comparison of spanwise and tangentially averaged loss coefficient profiles 239
7.4.3 Comparisons of maximum and mass-weighted average total pressure loss coefficient . . . 242
7.5 Conclusions . . . 244

8 Metal effectiveness of platform cooling concepts 247
8.1 Experimental setup and boundary conditions . . . 247
8.2 Baseline blades: results and discussion . . . 250
8.3 Impinging jet array: results and discussion . . . 254
8.4 Convective passage: results and discussion . . . 257
8.5 Comparison of all cooling configurations . . . 259
8.5.1 Comparison of area-averaged metal effectiveness . . . 260
8.5.2 Comparison of minimum metal effectiveness . . . 263
8.6 Conclusions . . . 264

9 Summary and recommendations for future work **267**

9.1 Summary of facility development and cooling system design work 267

9.2 Summary of adiabatic film effectiveness experiments 268

9.3 Summary of metal effectiveness experiments 269

9.4 Recommendations for future work 270

References **273**

Nomenclature

Roman Symbols

A	Area, m^2
A_{duct}	Impinging jet duct area, m^2
A_{pf}	Front seal slot flow area, m^2
A_{cell}	Numerical model discrete cell area, m^2
$A_{Fl}, B_{Fl}, m_{Fl}, n_{Fl}$	Jet impingement Nusselt number correlation coefficients, –
a, b	Blackbody calibration coefficients, –
Bi	Biot number, –
C_r	Raw count of IR camera sensor, –
c, c_p	Specific heat capacity; specific heat capacity at constant pressure, J $kg^{-1}K^{-1}$
D	Convective passage diameter, m
D_h	Hydraulic diameter, m
D_J	Impingement jet hole diameter, m
f_d	Darcy friction factor, –
$f_{d,trb}$	Turbulent flow Darcy friction factor, –
f_{pf}	Fraction of front seal purge flow supplied from NGV, –
H_D	Duct height, m
H_{exit}	Duct exit height, m
h, h_1, h_2	Heat transfer coefficient, W $m^{-2}K^{-1}$
h_{pf}	Front seal slot radial height, m
h_{trb}	Turbulent convective passage flow heat transfer coefficient, W $m^{-2}K^{-1}$
I	Circuit current, A
K_L	Pressure resistance coefficient, –
k	Thermal conductivity, W m^{-1} K^{-1}
k_f	Fluid thermal conductivity, W $m^{-1}K^{-1}$
k_w	Wall thermal conductivity, W $m^{-1}K^{-1}$
L	Platform length; Convective passage length, m
L_p	Perimeter length, m
M	Mach number, –
MR	Non-dimensional mass ratio MR = $\dot{m}_c\sqrt{T_{0c}}/\dot{m}_m\sqrt{T_{0m}}$, –
$\dot{m}_{air}$	Mass flow rate of air processed by gas turbine, kg s^{-1}
$\dot{m}_{Cf}$	Crossflow mass flow rate, kg s^{-1}
$\dot{m}_c$	Coolant mass flow rate, kg s^{-1}
$\dot{m}_J$	Impingement jet mass flow rate, kg s^{-1}
$\dot{m}_m$	Mainstream mass flow rate, kg s^{-1}
$\dot{m}_{pf}$	Front seal purge mass flow rate, kg s^{-1}
$\dot{m}_{total}$	Total impingement duct mass flow rate, kg s^{-1}
N	Number of experiments; number of timestamps, –
Nu	Nusselt number, –
n_c	Number of discrete cells used in numerical model, –
n_J	Number of impingement jets, m
P	Power dissipated by heater mesh, W
P_{Jx}	Impingement jet hole pitch/spacing in streamwise direction, m
P_{Jy}	Impingement jet hole pitch/spacing in tangential direction, m
P_J	Impingement jet hole pitch/spacing, m

Pr	Prandtl number, –
p	Static pressure, Pa
p_0	Total pressure, Pa
p_1, p_2, p_3, p_4	Brayton cycle station pressures, Pa
$p_{0,\text{inlet}}$	Supply total pressure, Pa
p_{0c}	Coolant total pressure, Pa
p_{0m}	Mainstream total pressure, Pa
p_{exit}	Exit static pressure, Pa
p_{pf}	Front seal purge pressure, Pa
$p_{01,\text{m}}$	Mainstream inlet total pressure, Pa
$p_{02,\text{m}}$	Mainstream exit total pressure, Pa
$p_{2,\text{m}}$	Mainstream exit static pressure, Pa
$\dot{Q}_{\text{in}}$	Heat input from fuel combustion, W
q, q_1, q_2, q_{12}	Heat flux, W m^{-2}
q^*	Surrounding radiance, W $\text{sr}^{-1}\text{m}^{-2}$
q_{cam}	Radiance incident on IR camera, W $\text{sr}^{-1}\text{m}^{-2}$
q_{obj}	Radiance from object of interest, W $\text{sr}^{-1}\text{m}^{-2}$
q_{opt}	Radiance from objects along optical pathway such as IR window, W $\text{sr}^{-1}\text{m}^{-2}$
q_{ref}	Radiance of surroundings reflected by object of interest, W $\text{sr}^{-1}\text{m}^{-2}$
R	Individual gas constant, J kg^{-1} K^{-1}
R_{a}	Surface roughness, m
R_{c}	Compressible mainstream recovery ratio, –
Re	Reynolds number, –
r	Compressible recovery factor, –
r_{p}	Mean platform radius, m
s	Entropy, J K^{-1}
T	Temperature, K
T^*	Equivalent mainstream temperature, K
T_0	Total temperature, K
T_1, T_2, T_3, T_4	Brayton cycle station temperatures, K
$T_{0,\text{pf}}$	Front seal purge total temperature, K
$T_{0c,\text{inlet}}$	Supply total temperature, K
T_{2s}, T_{4s}	Ideal isentropic Brayton cycle station temperatures, K
T_{Cf}	Crossflow temperature, K
T_{J}	Impingement jet temperature, K
T_{R}	Recovery temperature, K
T_{w1}	External mainstream gas path wall temperature, K; Upstream wall reference temperature, K
T_{w2}	Internal platform wall surface temperature, K
T_{01}	Core-flow temperature, taken as a proxy for the mainstream temperature, K
T_{02}	Coolant supply total temperature, K
$T_{\Re}$	Recovery and redistribution temperature, K
T_{0c}	Coolant total temperature, K
T_{0m}	Mainstream total temperature, K
$T_{\text{bb, factory}}$	Factory-calibrated blackbody temperature, K
$T_{\text{bb, lab}}$	Laboratory-calibrated blackbody temperature, K
$T_{\text{bb,TC}}$	Blackbody temperature at copper patch locations, K

T_{bb}	Blackbody temperature, K
T_c	Coolant temperature, K
T_m	Mainstream temperature, K
T_{opt}	IR window temperature, K
T_s	Effective temperature of the surroundings of a target surface for IR imaging, K
$T_{w,ad}$	Adiabatic surface temperature, K
$T_{w,TC}$	Surface temperature at copper patch locations measured with thermocouple, K
T_w	Surface temperature, K
Tu	Turbulence intensity, –
t	Time, s
t'	Non-dimensional time, referred to heated period, –
t_1	Start time of heated period, s
t_2	End time of heated period, s
t_w	Platform thickness, m
U	Rotor tangential velocity $U = r_p\omega$, m s^{-1}
U_Φ	Absolute uncertainty of generic quantity Φ expressed as $\Phi \pm U_\Phi$
V	Voltage drop across heater mesh, V
v	Velocity, m s^{-1}
$v_{1,abs}, v_{2,abs}, v_{3,abs}$	Seal flow, NGV exhaust jet, and mixed purge velocity in the absolute reference frame, m s^{-1}
$v_{1,rel}, v_{2,rel}, v_{3,rel}$	Seal flow, NGV exhaust jet, and mixed purge velocity in the rotor relative frame, m s^{-1}
v_{abs}	Velocity in the absolute reference frame, m s^{-1}
v_{cross}	Impinging jet duct crossflow velocity, m s^{-1}
v_J	Impinging jet velocity, m s^{-1}
v_{rel}	Velocity in the rotor relative frame, m s^{-1}
v_T	Tangential component of velocity in the rotor relative frame, m s^{-1}
v_{T1}, v_{T2}, v_{T3}	Seal flow, NGV exhaust jet, and mixed purge tangential component of velocity in the rotor relative frame, m s^{-1}
v_x	Axial component of velocity, m s^{-1}
v_{x1}, v_{x2}, v_{x3}	Seal flow, NGV exhaust jet, and mixed purge axial component of velocity, m s^{-1}
$\dot{W}_{comp}$	Compressor power consumption, W
$\dot{W}_{net}$	Gas turbine net power output, W
$\dot{W}_{turb}$	Turbine power output, W
W	Platform width, m
W_J	Platform width suitable for impingement cooling, m
w	Passage spacing, m
x	Surface position, m; Streamwise distance, m
Y_p	Total pressure loss coefficient, –
y	Tangential distance, m
z	Spanwise distance, m

Greek Symbols

α	Flow angle in the absolute reference frame, $^\circ$

$\alpha_1, \alpha_2, \alpha_3$	Seal flow, NGV exhaust jet, and mixed purge flow angle in the absolute reference frame, $^\circ$
β	Flow angle in the rotor relative frame, $^\circ$
$\beta_1, \beta_2, \beta_3$	Seal flow, NGV exhaust jet, and mixed purge flow angle in the rotor relative frame, $^\circ$
Γ	Capacity, $\Gamma = \dot{m}\sqrt{T_0}/p_0$, kg K$^{1/2}$ s^{-1} Pa^{-1}
Γ_c	Total coolant capacity of coolant setting plate, kg K$^{1/2}$ s^{-1} Pa^{-1}
Γ_h	Individual hole capacity in coolant setting plates, kg K$^{1/2}$ s^{-1} Pa^{-1}
γ	Ratio of specific heat capacities, –
ε	IR emissivity, –
$\varepsilon_\mathrm{emit}$	Blackbody emitter surface emissivity, –
$\varepsilon_\mathrm{paint}$	Black paint emissivity, –
η	Adiabatic film effectiveness
η^*	Equivalent mainstream effectiveness, –
η'	Corrected adiabatic film effectiveness, –
η_comp	Compressor isentropic efficiency, –
η_cycle	Cycle efficiency, –
η_turb	Turbine isentropic efficiency, –
$\overline{\eta^*}$	Weighted average equivalent mainstream effectiveness from multiple timestamp method, –
$\overline{\eta'}$	Weighted average corrected adiabatic film effectiveness from multiple timestamp method, –
$\overline{\eta'}_\mathrm{BE}$	'Best estimate' equivalent mainstream effectiveness after correction for non-adiabatic behaviour, –
θ	Metal effectiveness, –
ρ	Density, kg m^{-3}
ρ_pf	Front seal purge flow density, kg m^{-3}
σ	Standard deviation, [various units]
σ_bb	Stefan-Boltzmann constant, W m^{-2}K^{-4}
τ	IR transmittance, –; time constant, s
τ_f	Frictional shear stress, Pa
ω	Rotor angular velocity, s^{-1}

Subscripts

A	Evaluated at timestamp A
ASSUMED	Value assumed for modelling surface response (esp. in non-adiabatic situation)
B	Evaluated at timestamp B
c	Referred to compressible flow conditions
INFERRED	Values obtained through application of proposed methods (esp. in non-adiabatic situation)
R	Referred to compressible recovery condition
TRUE	Values specified as inputs to model or assumed to be true from supplementary experiment

Other Symbols

$\Re$	Recovery and redistribution parameter, –

Acronyms / Abbreviations

1D	One-Dimensional
2D	Two-Dimensional
3D	Three-Dimensional
BSI	British Standards Institution
CAD	Computer Aided Design
CFD	Computational Fluid Dynamics
DC	Direct Current
DMLS	Direct Metal Laser Sintering
GF	Glass-Filled
HP	High Pressure
IR	Infra-Red
LP	Low Pressure
LVDT	Linear Variable Differential Transformer
MERC	Metal Effectiveness Rotor Cooling
NGV	Nozzle Guide Vane
NI	National Instruments
NOx	Nitrogen Oxides
PR	Pressure Ratio
PS	Pressure Side
PSP	Pressure Sensitive Paint
RPM	Revolutions Per Minute
RQL	Rich-Quench-Lean
RSS	Root Sum of the Squares
SAS	Secondary Air System
SS	Suction Side
TCT	Transonic Calibration Tunnel
TET	Turbine Entry Temperature
TLC	Thermochromic Liquid Crystal
TR	Temperature Ratio
TSP	Temperature Sensitive Paint

Chapter 1

Introduction

Modern gas turbines, such as those found in commercial jet engines, operate with extremely high combustion temperatures in excess of 2000 K. The high pressure turbine is situated immediately downstream of the combustor and is subjected to some of the most extreme engine temperatures, which exceed the melting point of the metal alloy used for turbine components by several hundred degrees. Constant cooling of turbine components is therefore necessary for them to survive in such a harsh environment and allow the gas turbine to operate safely for a prolonged service life. These cooling systems must reduce peak metal temperatures to safe limits and minimise thermal gradients, which induce thermal stresses within components and reduce component life. Demand for more powerful and efficient gas turbines has driven the development of increasingly intricate cooling systems to deal with ever higher gas temperatures and minimise coolant usage.

Turbine cooling systems utilise high pressure, relatively cool air as a coolant. This is drawn from the compressor and bypasses the combustor. The cycle efficiency of the gas turbine is negatively affected by the removal of this air from the compressor. Additional aerodynamic losses are incurred by the reintroduction of this air into the main gas path within the turbine (in typical cooling systems such as film cooling) which further reduces the cycle efficiency. There is therefore a great deal of interest for research into novel cooling systems which can adequately cool turbine components whilst minimising the amount of required coolant.

Significant costs are associated with testing engine components in complete gas turbines. This limits the rate at which new novel systems, which may present some design risk and uncertainty, can be developed and introduced into commercial engines. An essential part of the modern gas turbine development process is the use of comparatively low-cost experimental facilities which can simulate engine-like conditions and be used to test and validate early-concept cooling systems. Advanced experimental facilities aim to replicate a higher fidelity engine-like

environment and produce results which can be easily scaled from relatively low temperature and pressure laboratory experiments to realistic engine conditions. This allows for new concepts to be developed rapidly within a lab-based setting, reducing costs and development time, before being integrated into engine components for final validation.

Turbine blade platform cooling systems have typically been less complex, and have been less well researched, than those employed within the main turbine blades due to the platforms typically having a lower thermal load. Current aircraft engine combustor designs produce hotter gas temperatures towards the midspan and reduced temperatures towards the endwalls leading to the most demanding cooling requirements being the blade lifting surfaces towards the midspan. However, engine technologies such as ultra lean-burn combustors, of increasing interest due to restrictions on exhaust emissions, produce flatter temperature profiles with increased gas temperatures towards the endwalls. The high pressure turbine blade platform is a region that is therefore likely to face increasingly challenging thermal loads, which require more advanced cooling systems, in future engine architectures.

The platform represents a significant portion of the metal surface area exposed to hot combustion gases as it sits at the base of each turbine blade and forms the hub-side endwall of the mainstream gas path. The platform also forms part of the front and rear hub seals which sit upstream and downstream of the rotor blade row respectively. Gaps in the hub endwall and internal cavities exist between stationary and rotating parts of the turbine to allow for motion. The hub seals are pressurised with air, also drawn from the compressor, to prevent hot gas ingestion into these spaces. The purpose of the platform is to constrain the mainstream flow and prevent ingestion of hot gas into the internal wheelspace cavities in the engine core. Preventing this is paramount as failing to do so would cause overheating of internal components leading to reduced service life or disastrous engine failure. The platform is a critical part of the turbine that must be adequately cooled. The platform and hub seals are intrinsically linked both in design and operating conditions. Any cooling features associated with the platform therefore must work within the context of the internal seal systems.

New, and better, platform cooling systems will be essential in allowing future aircraft engines to meet emissions targets whilst producing required power outputs and operating as efficiently as possible.

1.1 Research objectives

The aim of the work presented in this thesis was to investigate potential platform cooling systems and develop an experimental framework to design, test, optimise, and validate these novel systems. The key objectives of this research project, and a brief description of how they were achieved are listed below.

- Identify the most promising potential methods of cooling the rotor platform. A comprehensive survey of relevant literature and patents was conducted to find proposed new methods, identify current cooling systems employed elsewhere in the gas turbine which are suitable for integration into the platform, and discover opportunities for novel systems.

- Develop a number of novel platform cooling systems. Three concepts were developed using low-order numerical models to predict cooling behaviour and optimise design parameters. Two of these concepts (using convective cooling passages and impingement jet arrays) were integrated into 3D blade geometry and made via additive manufacturing.

- Design, manufacture, and commission an experimental facility capable of simulating representative engine conditions to allow for thermal assessment of a wide variety of platform cooling systems. The Metal Effectiveness Rotor Cooling facility (MERC) is a new experimental rig that was installed at the University of Oxford.

- Develop and refine the necessary data acquisition and processing methods to be able to determine useful performance parameters of cooling systems. An infra-red thermography system and number of processing techniques were used to determine metal effectiveness distributions from surface temperature measurements. A new method was developed for processing adiabatic film effectiveness values in the presence of a time-varying inlet thermal boundary layer.

- Experimentally assess the cooling performance of novel platform cooling systems. A method of enhancing the film cooling performance of front seal leakage flows was assessed by conducting a parametric set of experiments, varying leakage flow angle and mass ratio, and comparing adiabatic film effectiveness and total pressure loss coefficient. Two cooling systems integrated into the blade platform were tested at a variety of conditions and metal effectiveness distributions used to compare performance to a baseline, uncooled case.

1.2 Thesis outline

In Chapter 2, the increasing need for high pressure rotor platform cooling is discussed. A brief overview is given of relevant literature including: platform heat transfer; blade aerodynamics; existing and proposed cooling technologies; and experimental methods for assessing platform cooling.

In Chapter 3, the design of a new experimental facility (the Metal Effectiveness Rotor Cooling facility) used to assess platform cooling systems is detailed. The facility operates at scaled engine conditions and is a highly modular, high pressure, high temperature, transonic, linear cascade. This facility was designed, manufactured, and commissioned as part of the present work.

In Chapter 4, the designs of two novel platform cooling systems utilising convective cooling passages and impinging jet arrays are discussed. Low-order numerical models are detailed which were written to inform the design process of the the two cooling systems. Semi-optimised systems were then integrated into cascade blade geometry and manufactured, via direct metal laser sintering in titanium, to be used in experiments.

In Chapter 5, aerothermal commissioning data from the MERC facility are presented. The primary data acquisition systems for infra-red thermography and downstream aerodynamic traverses are described, along with processing techniques which allow for surface metal effectiveness distributions to be determined.

In Chapter 6, a journal paper is presented which covers a new processing technique to correct

values of adiabatic film effectiveness obtained from experiments in facilities with non-uniform inlet temperature profiles. The technique is demonstrated on data obtained from the MERC facility for the specific case of a time-varying inlet thermal boundary layer. Additional corrections are outlined to account for compressible recovery temperature and non-ideal adiabatic surface behaviour.

In Chapter 7, a method of altering the flow angle of front seal purge flow to improve its film cooling performance is described. Results from a parametric investigation of the film cooling performance of front seal purge flow at a range of flow angles and mass flow rates are analysed. Surface adiabatic film effectiveness distributions and downstream area traverses of total pressure loss coefficient are used to rank test conditions and demonstrate the viability of a method to enhance cooling performance.

In Chapter 8, metal effectiveness distributions across two blade platforms for each of the novel cooling system designs covered in Chapter 4 are presented. Cooling performance is compared between the two designs and a baseline case, with no integrated cooling features, at a range of coolant mass flow conditions. The viability of both systems and potential for further improvement to the designs are discussed.

In Chapter 9, the key findings and contributions to knowledge of the work presented in this thesis are summarised. Recommendations for future work are given.

Chapter 2

Literature review

In this chapter context is given regarding the motivation of the current research. It is argued that there is an increasing demand for rotor platform cooling systems and that future engines will require sophisticated platform cooling. The complex aerothermal flow field of the rotor passage and internal seal systems is briefly discussed as well as metrics which are used to assess cooling system performance. Following this a brief summary of current platform cooling technologies and research is presented to highlight aspects most pertinent to the content of this thesis. Additional reference to relevant material is included in later chapters. This chapter concludes by examining various experimental techniques and facilities used in cooling research and presents an argument for the need of a high degree of engine similarity to adequately assess platform cooling systems.

2.1 The increasing need for platform cooling

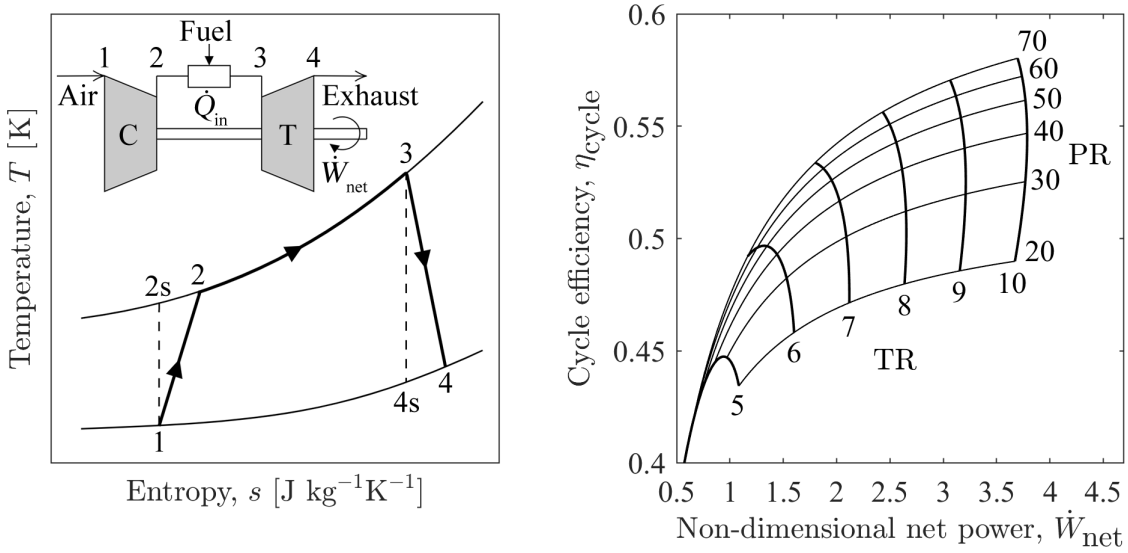
Active cooling of various turbine components is required in modern gas turbines due to exposure to high temperatures from combustion gases which exceed the melting point of component materials. Efforts to improve engine performance have led to increasingly high turbine entry temperature (TET) and therefore higher component cooling requirements. Turbine rotor platforms have historically had a reduced thermal load compared to the blades themselves due to the characteristic *hot-streak* radial temperature profile produced by RQL (rich-quench-lean) combustor configurations. There is however increasing interest in ultra-lean premixed combustors, due to environmental awareness of combustion gas emissions, which have a flatter radial profile leading to increased thermal loading of rotor platforms. As more and more powerful gas turbines are required the trend of increasing TET is likely to continue and this, combined with more usage

of ultra-lean burn combustors will increase the rotor platform cooling requirements in future engine architectures. The reasons why increasing TET leads to better engine performance are now discussed.

A typical gas turbine in an aircraft engine can be represented by the Brayton (or Joule) cycle. A schematic of an simplified gas turbine and a temperature-entropy plot of the Brayton cycle are presented in Figure 2.1a. The cycle consists of four stages: (1–2s) isentropic compression of air; (2s–3) isobaric addition of heat via burning fuel; (3–4s) isentropic expansion of combustion gases through a turbine used to run the compressor and provide aircraft thrust; and (4s–1) an isobaric cooling stage closes the ideal cycle by returning combustion gases to ambient conditions. Although the ideal cycle is closed, which includes the final cooling stage to the initial air state, aircraft engines typically operate in open cycles where fresh air is continuously drawn into the system and exhausted at higher-than-ambient temperature. In real engines, losses occur in both turbine and compressor resulting in the compression and expansion stages not being truly isentropic. The ideal isentropic stages are represented as dashed lines in Figure 2.1. The non-ideal operating lines are the solid lines and real operating stations are denoted 2 and 4, as opposed to 2s and 4s. Isentropic efficiencies for compressor and turbine are introduced, equation 2.1, to quantify the actual useful work consumed or extracted compared to the ideal isentropic case. A typical value of isentropic efficiency for both compressor and turbine is 0.9.

$$\eta_{\text{comp}} = \frac{T_{2\text{s}} - T_1}{T_2 - T_1} \qquad , \qquad \eta_{\text{turb}} = \frac{T_3 - T_4}{T_3 - T_{4\text{s}}} \tag{2.1}$$

The performance of a gas turbine represented by the Brayton cycle can be quantified by the non-dimensional net power output ($\dot{W}_{\text{net}}/\dot{m}_{\text{air}} c_p T_1$) and cycle efficiency ($\eta_{\text{cycle}}$). The net power output is the difference between the power output by the turbine, and the power required to run the compressor and, for an aircraft engine, represents the available power for propulsion. The non-dimensional net power output is normalised by the inlet properties of the working fluid ($\dot{m}_{\text{air}} c_p T_1$). For simplicity, we assume that the working fluid is air with constant specific heat capacity (c_p) and ratio of specific heats (γ). We also assume that the mass flow rate of fuel is negligible, compared to that of air, and that combustion gases can be treated as air. The



(a) (b)

Fig. 2.1 (a) Idealised gas turbine schematic and temperature-entropy plot of the Brayton cycle. (b) Contours of engine non-dimensional net power against cycle efficiency plotted for temperature ratios between 5 and 10 and pressure ratios between 20 and 70, with assumed $\eta_{comp} = \eta_{turb} = 0.9$ and air as working fluid.

non-dimensional net power output can then be expressed as in equation 2.2 as a function of the isentropic efficiencies of the compressor and turbine, the temperature ratio across the compressor and combustor (T_3/T_1), and the pressure ratio (p_3/p_1). The cycle efficiency is defined as the ratio between the net work output ($\dot{W}_{net}$) and the heat input into the system from fuel combustion ($\dot{Q}_{in}$). Cycle efficiency is expressed in equation 2.3.

$$\frac{\dot{W}_{net}}{\dot{m}_{air} c_p T_1} = \frac{\dot{W}_{turb} - \dot{W}_{comp}}{\dot{m}_{air} c_p T_1} = \eta_{turb} \frac{T_3}{T_1} \left(1 - \left(\frac{p_3}{p_1} \right)^{\frac{1-\gamma}{\gamma}} \right) - \frac{1}{\eta_{comp}} \left(\left(\frac{p_3}{p_1} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right) \tag{2.2}$$

$$\eta_{cycle} = \frac{\dot{W}_{net}}{\dot{Q}_{in}} = \frac{\eta_{turb} \frac{T_3}{T_1} \left(1 - \left(\frac{p_3}{p_1} \right)^{\frac{1-\gamma}{\gamma}} \right) - \frac{1}{\eta_{comp}} \left(\left(\frac{p_3}{p_1} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right)}{\frac{T_3}{T_1} - 1 - \frac{1}{\eta_{comp}} \left(\left(\frac{p_3}{p_1} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right)} \tag{2.3}$$

Increasing both non-dimensional power output and cycle efficiency is advantageous. Increasing the non-dimensional net power output means that the required power for flight can be generated by a smaller engine core which processes a reduced mass flow rate of air. The

reduction in core size reduces aircraft weight and, for a given engine nacelle size, allows for a larger bypass ratio. This increases propulsive efficiency—the efficiency with which fuel energy is converted into kinetic energy of the aircraft—and reduces specific fuel consumption [80]. Increasing cycle efficiency also reduces specific fuel consumption as a larger proportion of the heat from combustion is used for propulsion. Reduced specific fuel consumption means that, for a given flight, less total fuel is required. Using less fuel is advantageous as this decreases aircraft weight, raw fuel costs, and emissions of environmentally hazardous combustion gases.

Cycle efficiency is plotted as a function of non-dimensional net power output for a number of temperature ratios (TR) in the range 5–10 and pressure ratios (PR) in the range 20–70 in Figure 2.1b. Here γ has been taken as 1.4 and the isentropic efficiency of both compressor and turbine as 0.9. The non-dimensional net power output increases significantly with temperature ratio. For a given temperature ratio there are two different optimal pressure ratios which maximise either the non-dimensional net power output or the cycle efficiency. The cycle efficiency and the peak cycle efficiency attainable, at the higher of the two optimal pressure ratios, increase with temperature ratio. The reduction in non-dimensional net power output, for a given temperature ratio, at higher pressure ratios (closer to but lower than the value which maximises cycle efficiency) is relatively small. For this reason higher pressure ratios are favoured to maximise cycle efficiency however are predominantly limited by the increased weight, size, and reduction in operating life associated with them. At very high pressure ratios, greater than the peak efficiency value, the cycle efficiency decreases as the power required to compress air out-paces the extracted turbine work. Typical pressure ratio values for modern aircraft engine gas turbines are 40–60 and temperature ratios are 6–7.

Higher engine performance can therefore be attained by maximising the temperature ratio, which results in very high TET and an increased need for component cooling. Cooling typically involves using high pressure air, drawn from the compressor, as a coolant in the turbine. This practice comes at a cost to engine cycle efficiency for two reasons. Firstly, work is done by the compressor to increase the pressure of the coolant flow and no work is subsequently extracted from this coolant within the turbine. Secondly, the reintroduction of coolant into the main gas path and subsequent mixing with mainstream flow causes additional aerodynamic losses within

the turbine. There is therefore a trade-off between higher TET and coolant mass flow usage; a gas turbine with very high TET and coolant usage may have a lower cycle efficiency than a design with relatively low TET and coolant usage. Over time, cooling systems have improved both in capability to reduce metal temperatures and in using less coolant mass flow. This, in addition to the development of new materials with higher melting points, has led to a consistent rise in TET in commercial aircraft engines. The TET is typically increased to the very limit of what new technology allows in pursuit of increased performance. This is demonstrated in Figure 2.2 which charts historical Rolls-Royce engine operating TET and material capabilities.

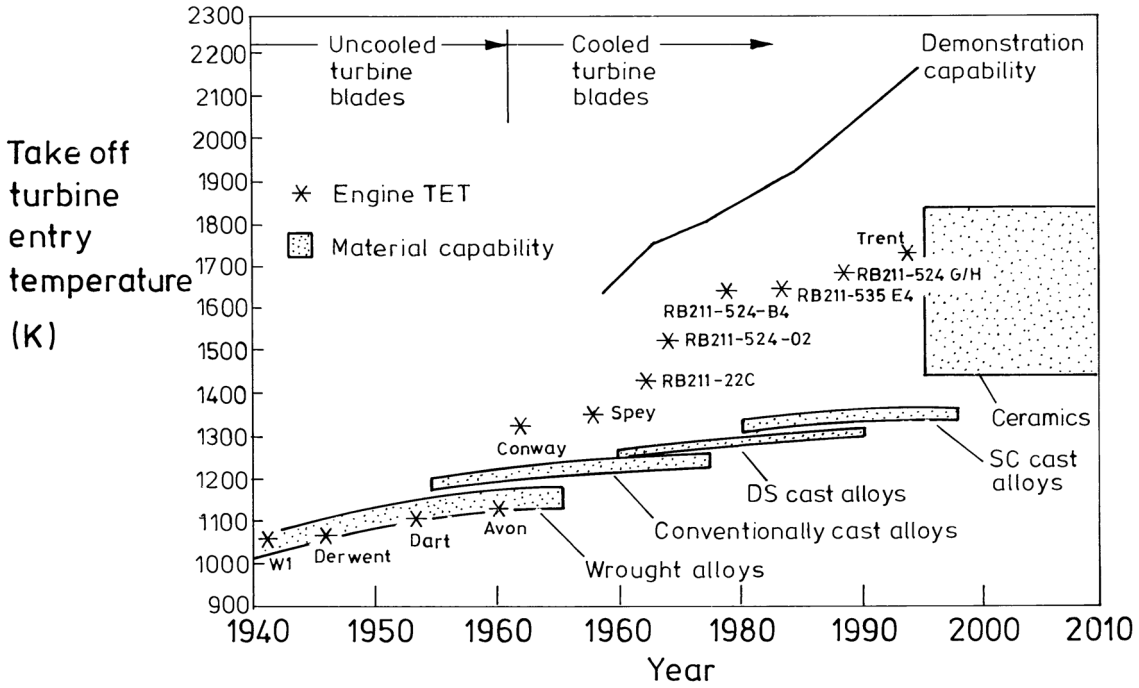


Fig. 2.2 Plot of historical turbine entry temperature (TET) for Rolls-Royce engines also showing material capability, from Cumpsty [15]

Awareness of the environmental impact of combustion gases has led to efforts, both regulatory and as a matter of principle, to reduce emissions from aircraft engines. Whilst the primary focus has been carbon dioxide (CO_2) production there are increasingly strict limits being imposed on other greenhouse gases such as nitrogen oxides (NOx). A consequence of increased combustion temperatures has been an increase in NOx production. There is an inherent trade-off between NOx and CO_2 production as higher efficiency engines (which require less fuel and therefore produce less CO_2) require higher TET and produce more NOx.

This has led to interest in technologies such as ultra-lean premixed combustors which have more complete combustion leading to reduced emission gases and fuel consumption. These combustors operate with very high air-fuel ratios and utilise highly swirling flows within the combustor to premix air and fuel and stabilise the flame in the combustion chamber. The spatial temperature profile of aircraft engines, which enters the high-pressure turbine is typically non-uniform and characterised by a 'hot streak' emerging from the centre of the combustor, surrounded by regions of relatively cool air. For rich-burn and RQL combustors the hot streak persists through the HP turbine rows and consequently the highest thermal load, and therefore highest cooling requirement, has been more stringent towards the mid-span of the turbine vanes and blades. Gas temperatures towards the endwalls are relatively lower and thus platform cooling has been of secondary importance. The distorted velocity profile associated with the highly swirling flow of lean-burn combustors also persists through the downstream nozzle guide vanes and causes significant redistribution of the hot combustion gases [72]. The result of this is radial redistribution of hot combustion gases, with an overall 'flatter' inlet profile to the HP turbine rotor with increased gas temperatures toward the endwalls. The cooling requirements of rotor platforms in engines with these kind of combustors are therefore higher, even when peak combustion temperatures are reduced compared to current RQL combustors.

2.2 The high pressure turbine rotor aerothermal field

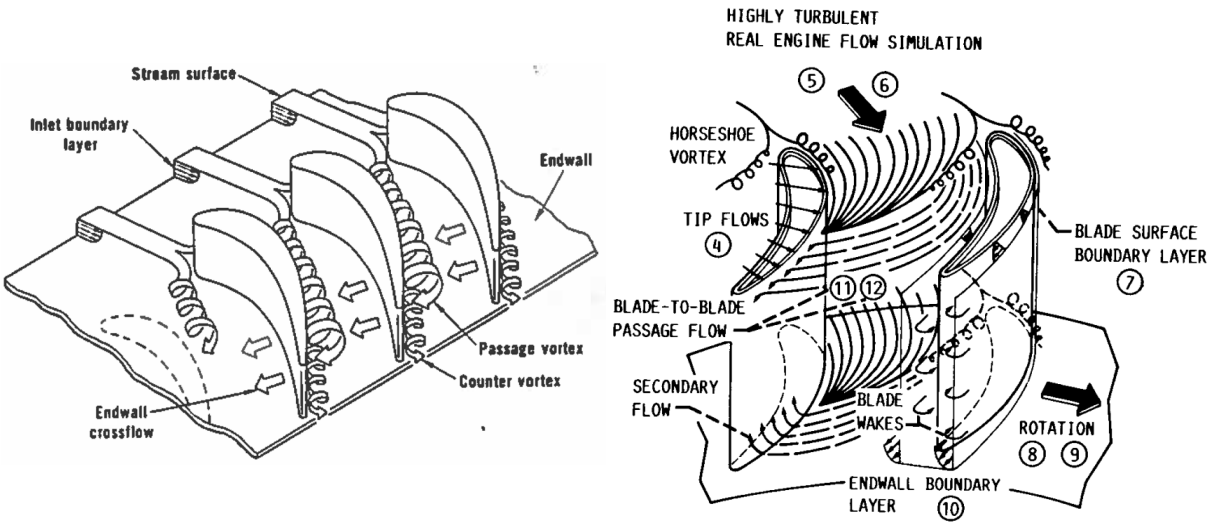
2.2.1 Rotor aerodynamics

Mainstream rotor passage aerodynamics

The aerodynamic flow field of the mainstream gas through the HP turbine blade passages is highly complex and three-dimensional in nature. The flow is unsteady, typically transonic, and has non-uniform boundary conditions with spatial velocity and temperature distributions caused by upstream combustor design and nozzle guide vanes. In addition to the bulk flow of the mainstream gas, a number of secondary flows develop as a consequence of the complex pressure field and which have a high degree of interaction with other flows. These secondary flows are

of great importance as they account for up to 50% of the aerodynamic loss in the turbine stage [78] or an approximate efficiency drop of 2–4% [49]. Heat transfer between the mainstream flow and rotor surfaces is highly dependent on the secondary flows and therefore an understanding of them is essential to understand turbine blade heat transfer and the cooling systems employed to protect the blades during operation.

The rotor flow field has been thoroughly investigated and a number of reviews exist such as Langston [52] which document the progression of understanding of the key flow features found within blade passages. The main secondary flow features within the blade passage are shown in Figure 2.3. At the leading edge of the blade the relatively slow endwall boundary layer stagnates and rolls up into two vortex 'legs' which migrate around the leading edge to the pressure and suction sides. Due to the shape of the two legs wrapping around the leading edge, this is known as the horseshoe vortex.



(a) (b)

Fig. 2.3 Visualisation of the secondary flows in a turbine rotor passage (a) from Langston [51], (b) from Gladden and Simoneau [33].

The blade turning and expanding flow induces a strong pressure gradient across the passage from the blade pressure side to the suction side. This pressure gradient drives the pressure side leg of the horseshoe vortex and the low-velocity boundary layer flow across the passage towards the suction surface. The latter is referred to as passage crossflow and forms a dominant passage vortex across the entire blade passage in the streamwise direction. The passage vortex interacts

with the suction-side leg of the horseshoe vortex, which attaches to and flows along the suction surface. The pair of counter-rotating vortices lift off from the platform and convect along the suction surface before being shed at the trailing edge. Despite the large amount of research into secondary flows, due to the difficulty in fully measuring their characteristic features, there are discrepancies between some models of these secondary flows. Some models put the suction side leg below the passage vortex instead of wrapping around it, and others consider the pair of vortices to be part of the same structure. Additionally some models introduce additional corner vortices where the endwall boundary layer crossflow encounters the suction surface and rolls-up similarly to the formation process for the horseshoe vortex (not shown in Figure 2.3.)

There is a small clearance gap between the rotor tip and case endwall. The pressure gradient between the pressure and suction sides of the blade draws flow through this tip clearance gap, causing highly turbulent tip leakage flows. Within the tip gap, the flow is highly turbulent and, due to the large pressure gradient and small flow area, the velocity of the tip leakage is significant. The flow can become supersonic and form shocks introducing additional aerodynamic losses. The tip leakage flow is reintroduced into the mainstream at the suction surface and at an oblique angle to the passage mainstream flow. The tip leakage flow separates from the endwall into a vortex, rotating in the opposite direction to the passage vortex, and is subsequently convected downstream. Losses associated with tip leakage and the tip leakage vortex are significant. In typical rotor cascades the circulation of the tip leakage vortex is significantly larger than the passage vortex [50]. Some shrouded turbine blades attempt to mitigate tip leakage and the associated losses by introducing a labyrinth seal at the blade tip.

Models of secondary flows are generally derived based on flows through linear cascades. Additional complexities are introduced by radial pressure gradients in annular systems. In these systems the pressure towards the hub is reduced. Investigations of radial pressure gradients in NGV passages by Sieverding et al. [79] and Povey et al. [70] showed that whilst the overall secondary flow structures remain similar, the endwall, low momentum boundary layers were strongly affected. Results indicated that secondary flows towards the casing are strengthened whilst those at the hub were reduced by the radial pressure gradient.

Secondary air system aerodynamics

In addition to the mainstream flow through the gas turbine there is a system of internal air flows referred to as the secondary air system (SAS) which performs a number of important roles required for safe operation. The SAS is responsible for: supplying cooling air to components; supplying sealing air, to prevent hot gas from the mainstream from entering the wheel space (internal cavities within the engine); and controlling the axial load on bearings. Additionally, the SAS is responsible for providing air for various aircraft systems outside of the engines such as for cabin ventilation and de-icing [48]. The SAS is highly complex with multiple air sources from bleeds within the compressor, numerous pathways to various areas of the gas turbine, a large number of varied seal types in numerous locations, and interactions at multiple positions with the mainstream flow. A schematic drawing, showing the main gas path and SAS, of a typical engine arrangement is presented in Figure 2.4 with arrows indicating SAS flows.

Cavities, known as the wheel spaces, exist between the rotating turbine discs and stationary components. As clearance is required to allow for the rotor motion, gaps exist along the mainstream hub endwall which may act as a leakage path for the hot mainstream gases into the wheel spaces. The high temperature of the mainstream gas would cause disastrous damage to internal components therefore a number of seal structures are used and the wheel spaces are pressurised with relatively cool, high pressure air from late stages within the compressor. There are many types of seal used within the SAS and all work principally by created resistance to flows by restricting flow area or forming convoluted pathways via shaped metal components. The supply of high pressure air into these seals assists to minimise, or altogether eliminate flow from the mainstream into the wheel space.

Effective sealing of the internal wheel space is critical to engine operation and therefore understanding the mechanisms which cause gas ingestion, and the performance of various seal types have been areas of intense interest for researchers. The primary drivers of gas ingestion, described in detail in reviews by Scobie et al. [76] and Chew et al. [12], are: rotationally-induced ingestion due to the radial pressure gradient caused by air rotating within the wheel space; pressure-induced ingestion due to circumferentially varying pressure distributions in the main

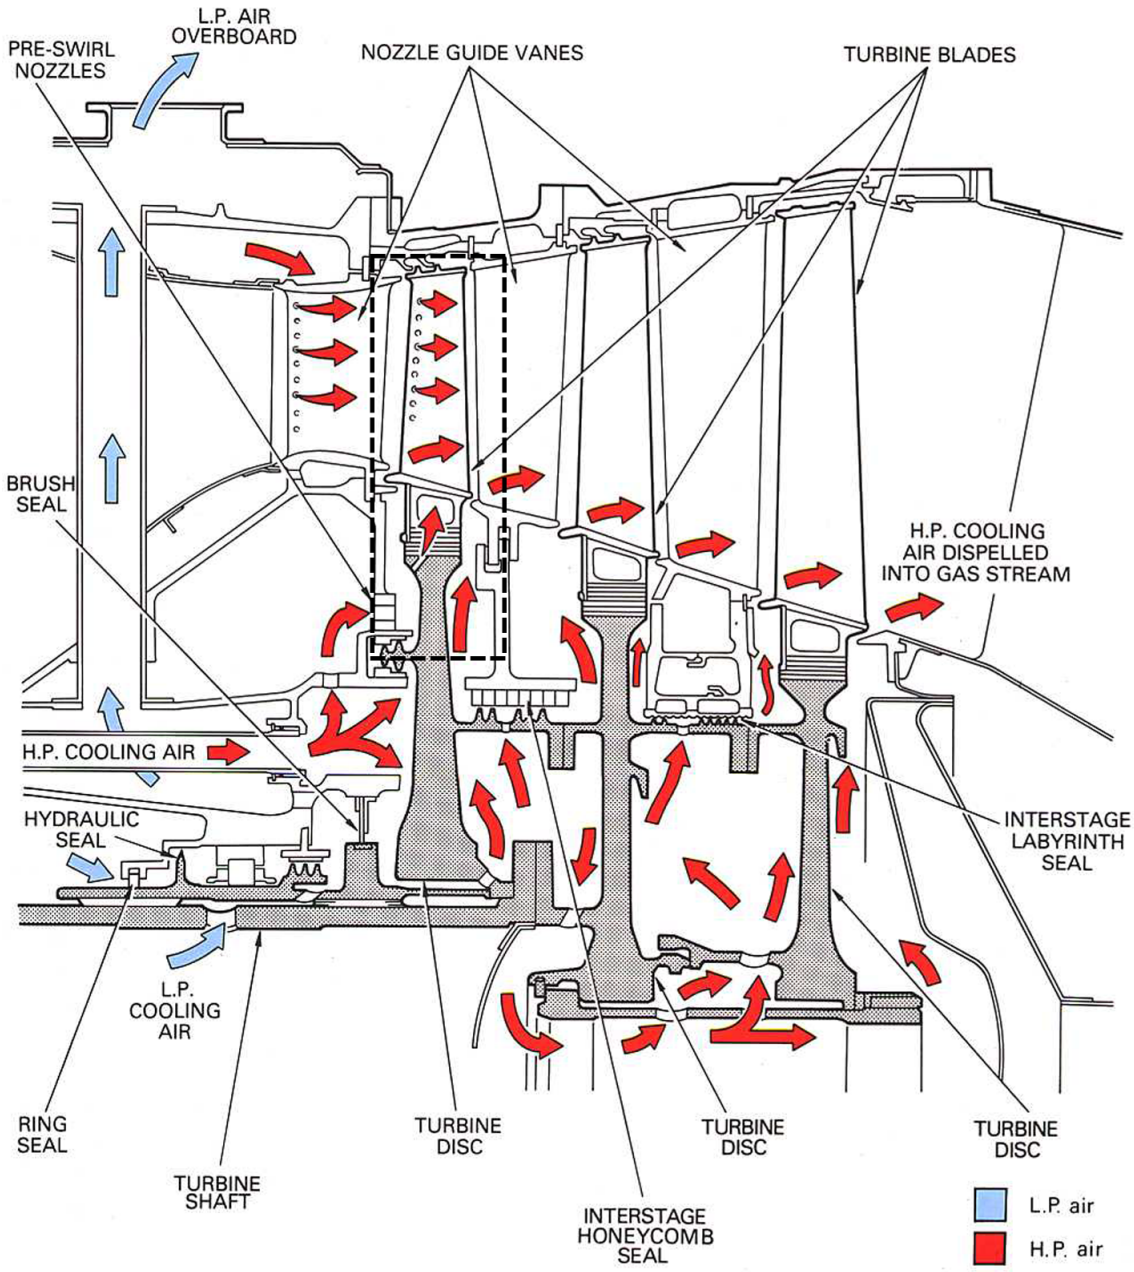


Fig. 2.4 A typical turbine cooling and sealing arrangement demonstrating the secondary air system flows, from [75]. High pressure rotor and relevant seals indicated within dashed-box area. H.P: high pressure; L.P: low pressure.

gas path due to NGV wakes and rotor blade potential fields. These have been experimentally investigated under different mainstream flow configurations by, for example, Bru Revert et al. [8] who looked at sealing performance in a rotating facility. Their set up considered purely rotationally-induced ingestion and combined ingestion, with both rotationally-induced and pressure-induced, ingestion in the presence of NGVs.

Experimental correlations are typically used to estimate the minimum amount of sealing air required for effective sealing of the wheel spaces. Sufficient air must be supplied for adequate sealing however, as this sealing air is bled from the compressor, this reduces the cycle efficiency of the engine. Accurate prediction is therefore important to minimise the required sealing air mass flow rate and the impact on performance. Whilst the sealing performance of hub seals present both upstream and downstream of the rotor blade is outside of the scope of the present work, a high-level knowledge of the complex interactions between the sealing flow and mainstream is paramount. The design requirements of these seals are especially important in understanding the role of the rotor blade platform.

The first stage HP rotor is located in the region within the bounding dashed line box in Figure 2.4. The rotor blades are individually cast and held fixed within the rotor disk by fir-tree root geometry. The main gaspath hub endwall is formed by the rotor platform which acts as part of the seal structure for both the front seal between the upstream NGV and the rotor, and the rear seal between the rotor and downstream NGV. The front and rear seal wheel spaces are pressurised with sealing air from the compressor. In some engines, including that shown in Figure 2.4, the sealing flow passes through pre-swirl nozzles to reduce temperature and pressure and to impart a rotational velocity to assist entry into the rotating wheel space flow [75]. This sealing air flows radially and eventually exhausts into the mainstream flow. Cavities, located below the platform and above the edge of the rotor disk, also exist between adjacent blades. These shank cavities (or shank pockets) are similarly pressurised with sealing air to prevent mainstream gas ingestion between the gaps between blades.

Seal flows and internal SAS architecture design have a great impact on the design of the platform as it is an integral part of the seal structures protecting the inner wheel space. In addition, any platform cooling system must be supplied with coolant from one of these surrounding flows.

2.2.2 Rotor platform heat transfer

Heat transfer between the hot mainstream gas and blade platform is a forced convection process which is highly dependent on the flow structures present and the state of the platform boundary layer. A number of boundary layer regimes (laminar, turbulent, transitional, separation *etc.*) and their associated heat transfer characteristics are described in Lakshminarayana [49]. Heat transfer characteristics are driven largely by the boundary layer velocity profile and are described by the Fourier conduction equation (2.4) and Newton's law of cooling (2.5), which introduces a *heat transfer coefficient* (h) to characterise the velocity profile. In general, a thicker boundary layer has a lower heat transfer for a given flow than a thin boundary layer since the temperature gradient at the wall is reduced. Stagnation regions typically have high rates of heat transfer due to very thin boundary layers. Laminar boundary layers have a lower heat transfer than turbulent boundary layers for the same thickness due to the latter having significant mixing and a steep near-wall velocity gradient.

$$q = k \left. \frac{\delta T}{\delta z} \right|_{z=0} \tag{2.4}$$

$$q = h(T_{\mathrm{R}} - T_{\mathrm{w}}) \tag{2.5}$$

Here q is the heat flux, k is the thermal conductivity at the wall surface, T is the fluid temperature, z is the direction normal to the surface, T_{R} is the recovery temperature (discussed in detail in Chapter 5) of gas at the surface, and T_{w} is the surface metal temperature.

The rotor blade passage boundary layers are influenced by secondary flows and the interactions between them, which strongly influences the heat transfer to the platform. Seven distinct regions of heat transfer on the rotor platform were identified by Friedrichs [25]. These are shown in figure 2.5 and are: the leading edge; suction surface shoulder; pressure surface corner; trailing edge wake; downstream of the passage vortex lift-off line; downstream of the throat; and the suction surface corner. These are now briefly discussed.

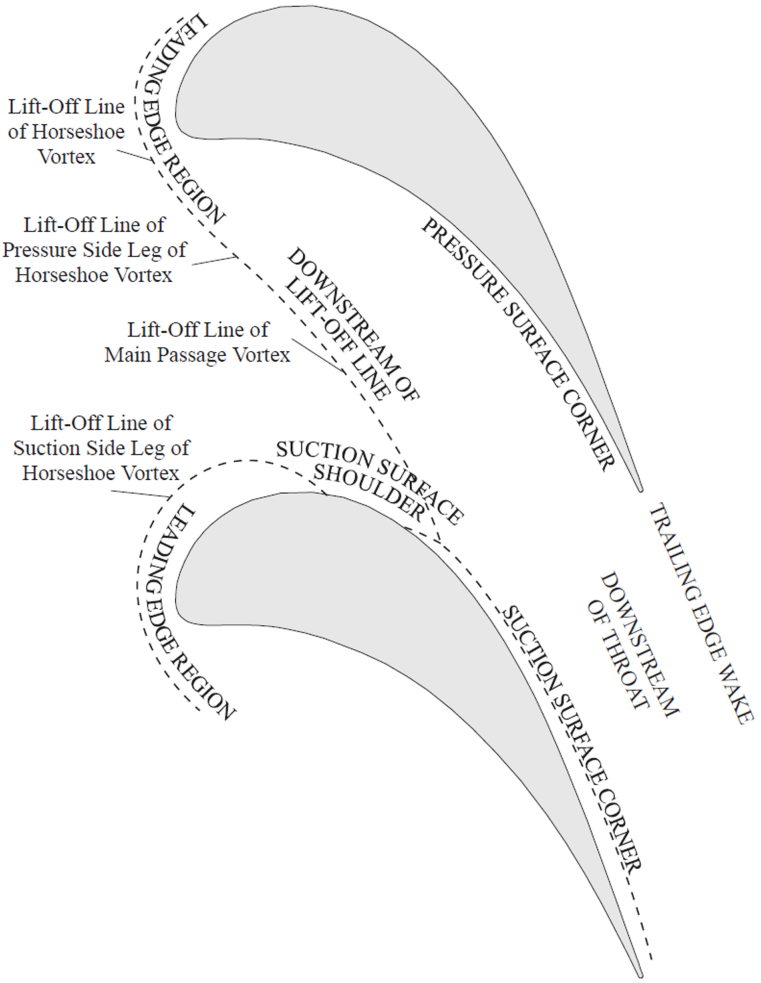


Fig. 2.5 Lift-off lines and regions of potential high heat transfer on the rotor platform. From Friedrichs [25].

The leading edge region, behind the lift-off line of the horseshoe vortex, has very high heat transfer due to the newly formed thin boundary layer and hot mainstream gas being swept down towards the platform by the vortex structure. At the suction surface shoulder region the passage vortex and pressure-side horseshoe vortex interact and roll up onto the blade surface. Consequently, the heat transfer in this shoulder region on the platform is high, as boundary layer fluid is entrained within the vortices resulting in fresh, thin boundary layers. The thick, turbulent inlet boundary layer is swept up by the main passage vortex and downstream of the lift-off line new laminar boundary layer growth occurs. This new boundary layer is thin which, along with the downwash of the passage vortex, leads to high heat transfer. The mainstream flow accelerates through the blade passage leading to high velocities and so new boundary layers remain thin as they develop along the platform. At the pressure side corner, downstream of the passage vortex

lift-off line, high rates of heat transfer are typically observed which may be due to high flow velocity and thin boundary layer.

High levels of turbulence and large velocities in the trailing edge wake result in another region of high heat transfer along the platform, downstream of the trailing edge. Hot gas from the mainstream core flow mixes within the wake and is convected down towards the platform. The suction surface corner and platform downstream of the throat can exhibit high levels of heat transfer but the magnitudes are more dependent on specific blade flow characteristics and Reynolds number effects. In the suction surface corner, the interaction between the counter-rotating main passage and corner vortices leads to high heat transfer. Downstream of the throat was found to have high heat transfer at high Reynolds numbers due to the thin boundary layer and high flow velocity present.

The HP rotor platform heat transfer is thus characterised by several distinct regions of high heat transfer. These 'hot spots' result in high metal temperatures and may require targeted cooling in order to bring these temperatures within acceptable limits for engine operation and component life. Secondary flows are the features which most strongly determine the location and magnitude of the high heat transfer regions and also cause the majority of aerodynamic losses within the turbine stage. Efforts to reduce the severity of secondary flows, or to control their behaviour would be beneficial to both engine efficiency and platform thermal loading. The wall temperature at each of the described locations scales with the temperature of the mainstream fluid which the platform is subjected to and will therefore increase with both higher TET and higher endwall temperatures characteristic of lean-burn combustors.

Cooling systems employed to reduce platform temperatures must reduce the peak platform metal temperatures in the hot spot locations to an acceptable value, and should aim to create as uniform a temperature distribution as possible. Thermal gradients within the metal of the platform lead to increased component fatigue during thermal cycling which reduces the service life. Cooling systems may also directly interact with both the mainstream flow (and secondary flows) and internal SAS flows which must be considered as part of the design and performance estimates.

2.3 Quantifying cooling system performance

A number of performance metrics are used to assess turbine cooling systems both during the industrial design process and in academic literature. These metrics typically allow for the performance of a cooling system to be quantified for comparative analysis of many designs, or for optimisation of a single system. The most pertinent performance metrics depend on specific factors driving the design and requirements of the cooling system. For the present work, the primary goals of the platform cooling are to reduce peak platform metal temperatures, achieve a uniform surface metal temperature distribution, and to maximise cycle efficiency by minimising the required mass flow rate of coolant and aerodynamic losses within the turbine stage. Platform metal temperatures are monitored by the non-dimensional metal effectiveness, θ whilst non-dimensional mass ratio, MR, and total pressure loss coefficient, Y_p, are used to assess coolant mass flow requirements and losses through the blade passage respectively.

2.3.1 Metal effectiveness

Cooling performance is quantified by the non-dimensional metal effectiveness, θ, (sometimes referred to as overall cooling effectiveness) and is defined in equation 2.6. It is the ratio of the difference between the driving gas temperature at the wall (the recovery temperature T_{R}) and surface metal temperature (T_{w}) normalised by the maximum temperature difference in the system, the difference between total mainstream ($T_{0\mathrm{m}}$) and coolant ($T_{0\mathrm{c}}$) inlet temperatures. A metal effectiveness of unity indicates a perfectly cooled surface where the maximum potential cooling of the supply coolant is being utilised, resulting in the wall temperature being the same as that of the coolant. A metal effectiveness of zero indicates that the wall temperature is the same as the driving hot gas temperature at the wall surface and therefore no cooling of the surface is being achieved.

$$\theta = \frac{T_{\mathrm{R}} - T_{\mathrm{w}}}{T_{0\mathrm{m}} - T_{0\mathrm{c}}} \tag{2.6}$$

Metal effectiveness is useful as, for a fluid of constant properties, it is independent of temperature and thus values can be scaled between different experiments and to engines at different operating temperatures. Typically, the metal effectiveness distribution is calculated from measured temperatures in relatively low-temperature experiments and then used with estimated (from CFD or other), or measured, engine temperatures to determine the expected engine component metal temperature distribution. Limits of acceptable peak metal temperature within an engine can be used to compute an acceptable lower limit of metal effectiveness via equation 2.6 and this value is then transportable between different temperature environments.

2.3.2 Non-dimensional mass ratio

Within this thesis, when comparing different cooling systems which utilise air as coolant in a number of ways (discussed below), the *non-dimensional mass ratio* MR is used to quantify the amount of required coolant. Here $\dot{m}_{\mathrm{c}}$ and $\dot{m}_{\mathrm{m}}$ are the mass flow rates of coolant and mainstream flow respectively and $T_{0\mathrm{c}}$ and $T_{0\mathrm{m}}$ are their respective total temperatures. This quantity was previously proposed and used by Kirollos [47] as a useful parameter that depends only on pressure ratio and can be expressed as

$$\mathrm{MR} = \frac{\dot{m}_{\mathrm{c}}\sqrt{T_{0\mathrm{c}}}}{\dot{m}_{\mathrm{m}}\sqrt{T_{0\mathrm{m}}}} = f\left(\frac{p_{0\mathrm{c}}}{p_{0\mathrm{m}}}\right) \tag{2.7}$$

An advantage of using the non-dimensional mass ratio, as defined above, is that it can be relatively simply matched between experimental and engine operating conditions (assuming the external Mach number distribution is matched). During the engine design process, coolant mass flow rates are commonly specified as a fraction of the mainstream mass flow rate ($\dot{m}_{\mathrm{c}}/\dot{m}_{\mathrm{m}}$). Matching engine temperature ratios in experimental facilities is difficult and as a consequence of this the ratio of coolant-to-mainstream mass flow rate cannot be matched. However, matching the non-dimensional mass ratio between engine and experiment allows the pressure ratio and $\dot{m}_{\mathrm{c}}/\dot{m}_{\mathrm{m}}$ to be determined. The mass flow rate of coolant directly affects the engine cycle efficiency as previously discussed. The non-dimensional mass ratio can be used to determine the required engine coolant mass flow rate and this can subsequently be used with correlations to estimate

expected cycle efficiency penalties. In general, the higher the mass ratio the more coolant is being used, and therefore the larger the drop in cycle efficiency which is incurred.

2.3.3 Total pressure loss coefficient

The total pressure loss coefficient is used to quantify aerodynamic losses through a blade passage due to boundary layers, flow separation, secondary flows such as the passage vortex, and shock waves. Aerodynamic losses cause a rise in entropy through the turbine and therefore a reduction in turbine isentropic efficiency and engine cycle efficiency. For turbines, the total loss coefficient is defined as the difference between the inlet and exit total pressure of the flow, normalised by the isentropic exit dynamic pressure, equation 2.8 [17]. Larger values of total pressure loss coefficient indicate larger increases in entropy through the blade passage and therefore larger penalties on cycle efficiency. Area-averaged and mass-averaged values of loss coefficient calculated from distributions surveyed downstream of the bladerow are commonly used to quantify differences between different blade and cooling system designs and operating conditions.

$$Y_p = \frac{p_{01,\mathrm{m}} - p_{02,\mathrm{m}}}{p_{01,\mathrm{m}} - p_{2,\mathrm{m}}} \tag{2.8}$$

2.4 Platform cooling techniques

There are a wide range of methods which can be implemented to cool the blade platform and a combination of several techniques is usually preferred. Comprehensive literature reviews of platform cooling techniques and relations to governing flow parameters have been compiled by Han et al. [36]; Wright et al. [88]. Cooling systems can broadly be categorised into three groups: external methods, such as film cooling, which limit the heat transfer between the mainstream hot gas and platform; internal methods which actively cool the platform or enhance heat transfer on the underside of the platform to the shank cavity; and surface geometry modification methods which alter and control secondary flow structures to reduce heat transfer on the platform.

Most external and internal cooling systems described below have seen widespread use in cooled rotor blades. Very few internal systems have been implemented into rotor platforms due

to the historically lower thermal loading on the platform, in addition to manufacturing difficulties in forming single-crystal metal structures with complex geometric features in two planes (the blade and platform are perpendicular which is challenging for this manufacture process even when the platform is solid metal). An example schematic of an externally and internally cooled turbine blade and platform is presented in Figure 2.6.

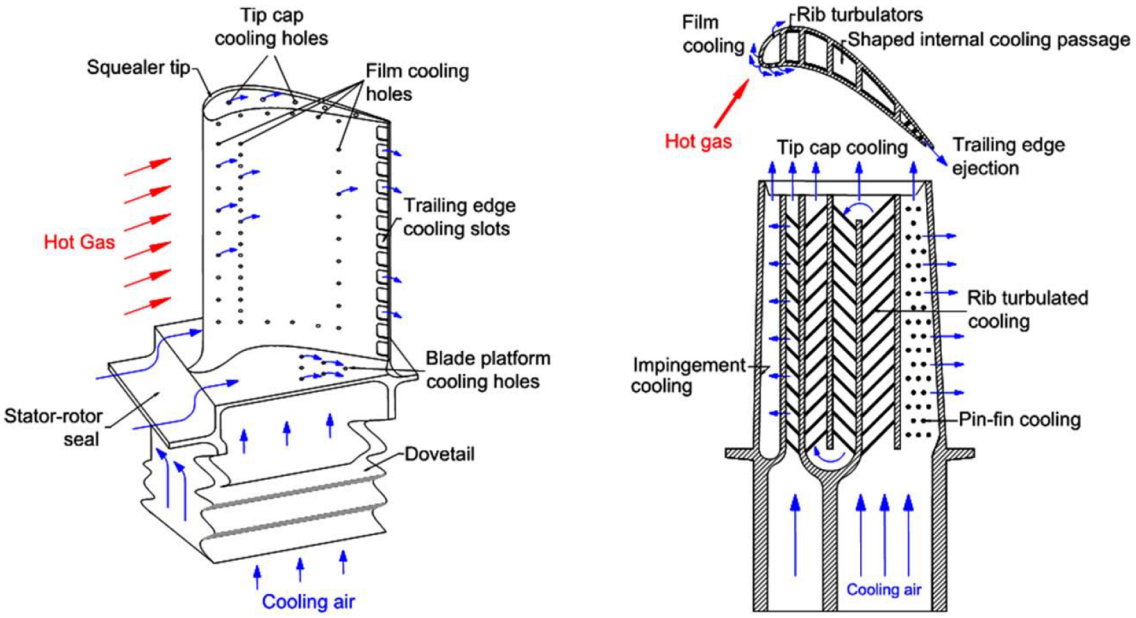


Fig. 2.6 Example schematic of a modern, cooled gas turbine rotor blade showing both external and internal cooling methods. From Han and Rallabandi [37]

2.4.1 Film cooling

Film cooling is a technique in which relatively low-temperature air is ejected onto a surface and forms a thin protective layer limiting the heat transfer from hot mainstream fluid. As coolant flows along the surface, it acts as a barrier that gradually increases in temperature and mixes with mainstream gas. Film cooling results in a penalty to cycle efficiency for two reasons. Firstly, due to the coolant being used typically being drawn from late stages of the compressor to fulfil pressure requirements. Secondly, the reintroduction of this flow into the main gas path as a film, and subsequent mixing with the mainstream flow, results in aerodynamic losses downstream of the ejection location. It is therefore important to maximise the cooling performance of introduced films and minimise the amount of coolant required. The performance of a film cooling system

is commonly quantified by the adiabatic film effectiveness,η, (a similar parameter to metal effectiveness) as in equation 2.9.

$$\eta = \frac{T_{0\mathrm{m}} - T_{\mathrm{w,ad}}}{T_{0\mathrm{m}} - T_{0\mathrm{c}}} \tag{2.9}$$

The adiabatic wall temperature ($T_{\mathrm{w,ad}}$) is the same as the coolant supply temperature ($T_{0\mathrm{c}}$) for an adiabatic film effectiveness value of unity. This represents the maximum possible reduction in heat transfer between the hot gas and the wall. An adiabatic film effectiveness of zero means that the wall temperature is the same as that of the hot mainstream gas $T_{0\mathrm{m}}$ and therefore the film has mixed out or increased in temperature such that it no longer provides any useful cooling on the surface. In an adiabatic system, the adiabatic film effectiveness is analogous to a mixing mass fraction of the hot mainstream gas and cold coolant gas at a particular location, resulting in a gas temperature at the wall of $T_{\mathrm{w,ad}}$. For a non-adiabatic system, the film effectiveness and surface heat transfer coefficient can be used with numerical models to estimate metal temperatures.

Hub seal purge flows

Film cooling on rotor platforms occurs as a consequence of leakage flows from the hub seals and clearance gaps between adjacent rotor blades. The front and rear hub seals are pressurised to prevent hot gas ingestion and must therefore be able to overcome the maximum mainstream flow pressure at the inlet and exit of the blade row. As the pressure field in the gas path is non-uniform, due to features such as wakes and potential field of the blade, regions of relatively low pressure occur along the annulus. This results in sealing flow, typically air from the same source used as coolant in blades, from within the seal to leak out onto the mainstream platform surface through an annular slot between the upstream NGV trailing edge and rotor platform. Similar leakage occurs between adjacent blades due to the reduction in pressure through the rotor passage, resulting in leakage from the shank cavity towards the rear of the platform.

Flow which leaks onto the mainstream platform surface is sometimes referred to as *purge flow*. Purge flow acts as a cooling film and whilst largely unintentional, it can effectively cool part of the blade platform. The cooling benefit of purge flows have been investigated by many

research groups. Early work by Blair [5] investigated the use of purge flow from a slot upstream of a set of turbine vanes. Adiabatic film effectiveness measurements were taken at a range of purge mass flow rates. They showed that the ejected purge flow is quickly swept towards the suction surface by the passage cross flow and that adiabatic film effectiveness drops after the passage vortex lift-off line. An increase in the ejected mass flow was shown to slightly increase both the area and values of adiabatic film effectiveness on the platform. However the strong effect of crossflow within the passage limited the added benefit of additional coolant as all of the purge was swept towards the blade suction surface. Other groups have demonstrated similar findings and showed that purge from the front seal can only partially cool the front half of the platform [3].

Popović and Hodson [69] identified a leakage vortex within the purge flow slot caused by mainstream gas ingestion. This vortex is formed near the pressure side and travels down the purge slot before ejecting near the suction side. This biases the purge flow to being ejected towards the suction surface, further reducing the area cooled before the coolant is swept onto the suction surface by passage crossflow. Higher purge mass flow ratios ($\dot{m}_{\mathrm{pf}}/\dot{m}_{\mathrm{m}}$) of 1.5% were shown to push out ingested air more effectively. The increased mass flow rate and diminished leakage vortex resulted in improved adiabatic film cooling effectiveness.

Geometry of the front seal slot was shown to influence the cooling effectiveness and spread of purge flow by Piggush and Simon [68] and Thrift and Thole [84]. Piggush and Simon [68] showed that having a backward-facing step before the injection slot led to a thickening of the boundary layers and an increase in cooling effectiveness across the whole platform. A forward-facing step was shown to have the opposite effect and increased heat transfer from the hot mainstream due to inlet boundary layers becoming thinner. Thrift and Thole [84] studied the effect of the angle at which the slot coolant is injected onto the platform. When the coolant is injected perpendicular to the axial direction and at large angles (65°) a vortex is formed above the slot, in front of the leading edge, which results in lower cooling effectiveness across the platform. At lower angles of injection the boundary layer is energised by the purge flow and cooling effectiveness enhanced as separation is prevented as it meets the leading edge.

Purge flows from the shank cavity typically occur towards the rear of the platform due to reduced mainstream pressure. Similarly to front seal purge flow, the coolant is rapidly swept across the passage towards the suction surface by the passage crossflow. The flow in this region is transonic which results in large mixing losses. Slashface leakage is generally minimised for this reason by introducing damper plates which assemble between adjacent blades and act as a seal. These plates also assist in tuning vibrations of the blades during operation which gives them their name.

Discrete hole film cooling

Film cooling through discrete holes is the primary way in which turbine blades are cooled and as such have been heavily investigated and are well understood. Exhaustive reviews of film cooling research are available by Gilboy and Gan [30] and Bogard [6]. Many parameters have been shown to affect the cooling performance of films including flow parameters such as coolant-to-mainstream mass flux ratio (also referred to as blowing ratio) and momentum flux ratio, and geometric parameters such as hole size, shape and angle of inclination.

Film cooling via holes has been implemented and investigated in platform cooling to supplement purge flow cooling. Holes are placed near to the pressure side of the blade and in the downstream half of the platform, which is unprotected by the purge flow due to being swept towards the suction surface and entrained in the passage vortex. Experiments by Chen et al. [11] demonstrated the potential of a combined purge flow and cooling hole system to cool the majority of the blade platform. Coolant ejected from film cooling holes is similarly drawn by the passage crossflow towards the suction surface, limiting the cooled area.

A series of experiments by Narzary [62], who tested a number of film cooling hole shapes and configurations, showed that the layout of the holes could be used to change local platform hotspots. However the difference in averaged effectiveness was insignificant. They also demonstrated that fan-shaped holes outperformed cylindrical holes in both coverage and values of adiabatic film effectiveness. The performance increased with blowing ratio from 0.5 to 2 for fan-shaped holes whereas for cylindrical holes the optimum was unity due to jet lift-off.

Regions of the platform near to the pressure side and towards the rear, which are unprotected by leakage films, are problematic areas to cool using film cooling. The high pressure towards the pressure side of the blade requires very high coolant supply pressures to allow for films to eject onto the surface, and these would then be expected to have low momentum flux ratios increasing aerodynamic losses. Towards the rear of the platform, the mainstream flow velocity is high, which also increases aerodynamic losses due to mixing with the introduced film cooling jets.

2.4.2 Platform geometric modifications

Endwall contouring involves altering the surface geometry of the platform in order to control secondary flows and their impact on heat transfer. Many configurations have been proposed which work on the principle of favourably altering the pressure gradients in the passage to diminish the strength of secondary flows and thus reduce heat transfer into the platform. Research by Piggush and Simon [67] showed that this could also reduce aerodynamic losses within the blade passage associated with secondary flows. A non-axisymmetrically contoured endwall was shown by Lynch et al. [57] to be able to reduce the strength of the passage vortex and reduce the heat transfer near the pressure side compared to that of an axisymmetric endwall.

Two different endwall contours were studied by Erickson et al. [20] who looked at both the aerodynamic and heat transfer effects of the profiles. These consisted of a gradual slope 'shark nose' profile and a steep slope 'dolphin nose' profile at the leading edge of the platform along which the purge flow was ejected. Both profiles improved the adiabatic film cooling effectiveness across the front half of the platform compared to a sharp step, especially towards the suction side. The steep profile was shown to outperform the shark nose profile both in terms of heat transfer and aerodynamic losses. The somewhat unexpected result was that increasing the mass flow of the coolant from the purge slot, from 0.5% to 1.5% of the mainstream, did not lead to any significant increase in aerodynamic losses but did increase the cooling effectiveness on the first half of the platform.

The impact of endwall contouring on cooling effectiveness and distribution also depends on the location of any film cooling ejection. It was shown by Thrift et al. [83] that purge flow

cooling was positively impacted by the addition of an endwall contour however film cooling on the second half of the platform was negatively affected. Any attempt to implement platform contouring as part of a cooling scheme must therefore not only take the changes in aerodynamics into account, but must also tailor the contour to enhance the film cooling performance along the platform. An example of a proposed system of a contoured platform is given in Figure 2.7a. This design has an elevated bulb, near the pressure surface leading edge, which forms a channel along the pressure surface to accelerate flow and reduce static pressure and hence reduce passage cross flow. Advantages may be to reduce aerodynamic losses associated with passage vortices and could potentially increase the adiabatic film effectiveness of purge flow film cooling near the pressure surface. The bulb however creates a significantly thicker region of metal which may be difficult to adequately cool.

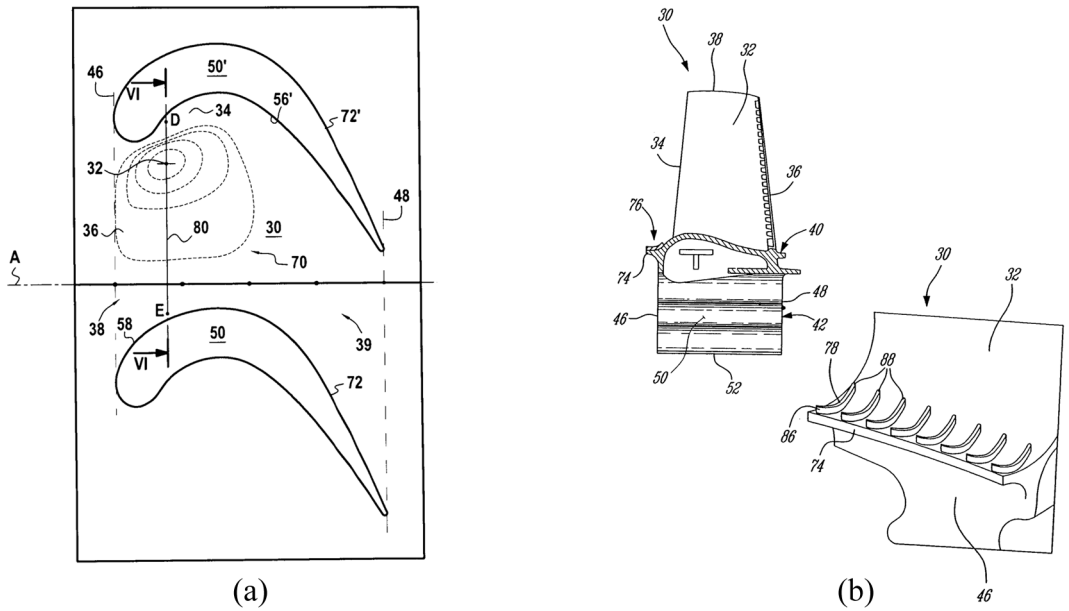


Fig. 2.7 Example patents of platform geometric modifications to alter heat transfer (a) contoured platform with bulb introduced near the pressure surface leading edge to inhibit passage crossflow [35] (b) introduction of front seal deflectors to guide purge flow [32].

In addition to contouring the platform surface, other modifications to the platform geometry have been considered such as the introduction of leading edge fillets and purge slot flow deflectors. The introduction of a prominent fillet at the leading edge between the blade and platform was studied by Han and Goldstein [38] and was shown to impede the formation of the horseshoe vortex. Corner vortices were instead observed to form behind the fillet and led to increased heat

transfer when compared to a blade with no fillet. Attempts to optimise the design by Mahmood et al. [58] led to overly large fillets which posed a significant cooling problem.

Leading edge deflectors, which alter the purge flow onto the platform, have also been proposed in a number of patents. The purpose of these deflectors would be to turn the purge flow and increase its momentum in the direction of the main gas flow. This aims to reduce aerodynamic mixing losses and improve film coverage and effectiveness. An example of these deflectors is shown in Figure 2.7b. A more simple variation of this is a blockage deflector which blocks sealing flow from exiting near the suction surface and diverts the purge flow to exit the slot nearer to the pressure side of the blade. In reality such systems would be impractical as a finite clearance gap is required for rotor motion and the effectiveness of such low aspect ratio features in the front seal cavity is dubious. Expanding the front seal cavity to accommodate such features would likely also introduce sealing issues.

2.4.3 Internal cooling

The cooling systems detailed above all focused on reducing the heat transfer between the mainstream hot gas and external platform surface. On the rear side of the platform heat transfer also occurs between the metal and coolant used to pressurise the front and rear seals, and shank cavity. Internal cooling systems are desirable as they can potentially incur less cycle efficiency penalty than external systems due to not reintroducing coolant into the mainstream, causing mixing losses. Proposals have been made to implement technology common to blade internal convective cooling, such as roughened channels and impinging jets, to additionally cool the platform. Implementing internal cooling poses a significant challenge in terms of manufacturability and coolant supply. However with the increasing cooling requirement of the platform these systems may prove to be as essential for the platform as they currently are for blades and NGVs.

There are many patents which demonstrate clear interest in cooling schemes integrated directly into the platform surface. There is however very little published material experimentally demonstrating such concepts. This is likely due to a lack of viability due to manufacturing and

component strength, and the desire to protect any intellectual property which overcomes these challenges. Nonetheless, such cooling systems are of interest at an academic level and may become commercially viable with future developments in manufacturing techniques such as additive manufacturing.

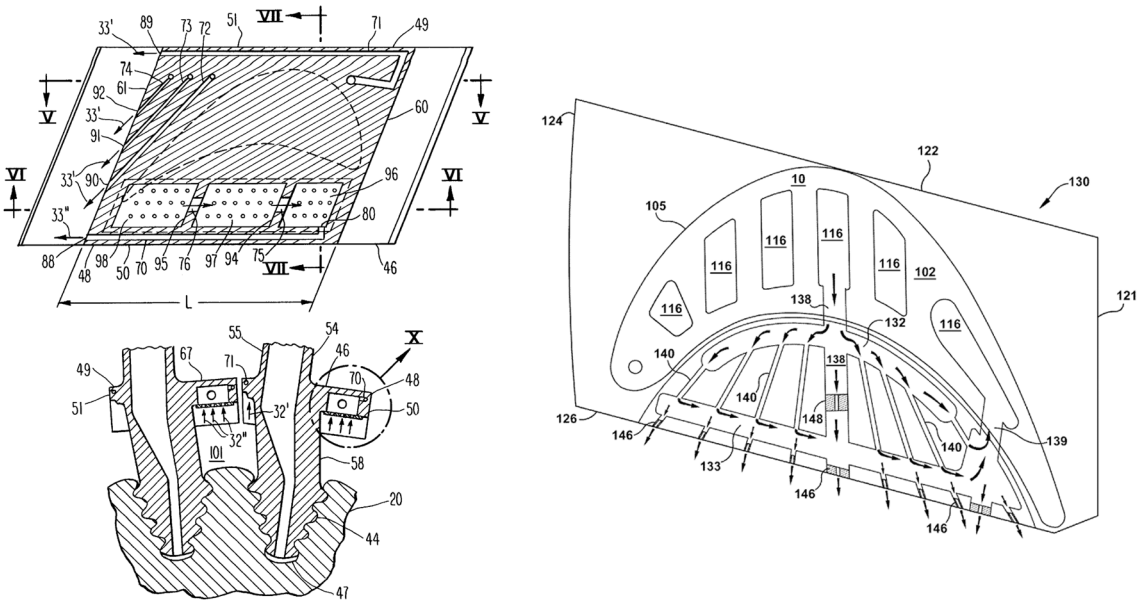
Convective cooling

Convective cooling systems all work on the same basic principle demonstrated by a flowing heat exchanger. Coolant flows through an internal pathway, cooling the bounding walls by picking up heat and convecting it out of the system. The temperature of the component as a whole is reduced as heat is conducted towards the cooler internal surfaces and into the coolant flow. In gas turbines a large number of variations of convective cooling passages are used to cool turbine rotor blades and NGVs. Simple passages with coolant flowing radially from hub to case are widely used, alongside more complicated multipass systems in which the flowpath includes several turns through large angles within the blade body. Convective passages are commonly used alongside film cooling, with coolant first convectively cooling inside the blade and then being exhausted into the mainstream to act as a cooling film.

Internal heat transfer along cooling passages can be enhanced by introducing a number of additional features within the flow path such as flow turbulators and pedestals, or pin-fins. Turbulators, typically angled ribs or other features such as dimples, along the walls of the flow path act to disrupt the cooling flow and enhance heat transfer by increasing coolant turbulence and mixing. Heat transfer is increased by the introduction of secondary flows caused by the turbulators and also by disruption of the boundary layers. Pedestals are generally arranged in arrays, span the entire coolant flow path, and are shaped to similarly disrupt the flow and enhance heat transfer. Additional cooling benefit is provided due to pedestals joining surfaces and allowing conduction between between them to redistribute heat. Pedestals also have large surface areas to convectively pass heat from the joined metal surfaces into the coolant. One other advantages of pedestals are the mechanical strength they provide by connecting surfaces.

Convective cooling systems have been exhaustively studied in the context of multi-pass cooling channels in turbine vanes and blades and thorough reviews of past literature and current

understanding of these systems can be found in Han et al. [36]; Lakshminarayana [49]. Though these systems are vital aspects of blade cooling they have had very limited application to platform cooling primarily due to the difficulty in manufacturing. Proposals in filed patents have mostly involved the introduction of cooling circuits within the platform similar to those found in blades. Figure 2.8b shows an example of a system where high pressure coolant is taken from the blade supply and passed through a network of cooling passages integrated within the platform. The coolant is then used to cool the platform slashfaces by exhausting it between neighbouring platforms. Excess coolant which has reduced pressure through the network is returned to the blade trailing edge supply. A number of similar systems exist in published patents which generally all draw coolant from the blade supply, which would therefore have to have a larger amount supplied from the SAS incurring additional cycle efficiency penalty.



(a) (b)

Fig. 2.8 Example patents of integrated internal platform cooling systems (a) jet impingement array from the shank cavity [43] and (b) convective cooling passages drawing coolant from the blade multipass [19].

Impingement cooling

Impingement cooling involves directing a high velocity jet of coolant, typically formed by a small-diameter hole in an internal plate, to impinge on a target surface to be cooled, such as the rear of the platform. As the flow impinges upon the target surface a stagnation region of high

heat transfer is formed. The jet flow is forced to turn and flows laterally along the target surface, convectively cooling an area around the stagnation region. Multiple jets can be used, arranged in an array, to effectively cool a large surface area. Adjacent jets in arrays interact with one another and this interaction may positively or negatively affect the cooling performance depending upon a range of factors such as jet spacing, hole/nozzle diameter, distance to target surface, and mass flow rate of each jet. Jets are commonly supplied from a high pressure source into a cavity beneath the target surface. These cavities are often bounded by walls which constrain flow in one of more directions. Spent jet coolant, which has impinged on the target surface, is therefore forced to flow towards any available exhaust area and may interact with a large number of jets along the available flow path. The migration of spent jet coolant perpendicular to the direction of jet flow is referred to as jet crossflow, not to be confused with passage crossflow which refers to the migration of fluid within the blade passage due to the passage pressure gradient.

Impingement systems of varying complexity have been considered for platform cooling. Bunker and Bailey [9] proposed a system which used a single jet, supplied from the front seal cavity, to cool the platform near the leading edge. Spent jet coolant would then flow along an internal surface of the platform and supply film cooling holes. More complex systems using various configurations of impingement jet arrays integrated into the platform are proposed in many patents. One such example is given in Figure 2.8a which includes three chambers integrated into the platform, each with an array of jets. In most proposed patents, coolant is supplied from the shank cavity and is also used to feed the film cooling holes on the platform surface, or between adjacent blades to cool the platform slashfaces. In some systems the impingement cavities are machined or cast directly into the platform surfaces, whereas others form the jet arrays and cavities by bonding a secondary plate to the rear of the platform. Both approaches are highly challenging from a manufacturing and component stress perspective.

2.5 The effect of rotation

2.5.1 External coolant films

The effect of rotation on film cooled blade platforms has been investigated by Pau et al. [65] and Suryanarayanan [81] in fully rotating facilities. Pau et al. [65] tested the aerothermal impact of purge flow at two rotational speeds (at design speed of 6500 revolutions per minute (RPM) and off-design 4500 RPM) and found that the purge flow contributes to a passage blockage effect which impacts the flow field of both the rotor and upstream stator. Higher purge mass flow rates were found to decrease incidence on the rotor leading to increased hub loading.

Suryanarayanan [81] tested both upstream seal purge flow and discreet film cooling holes at three rotational speeds (2400 RPM, design-point of 2550 RPM and 3000 RPM). They found that platform cooling effectiveness and coverage increased, at each rotational speed tested, when a larger purge mass flow rate was used. Reductions in cooling effectiveness, compared to the nominal design point and for each of the tested purge mass flow rates, were observed for speeds both higher and lower than the nominal design point. This was attributed to changes in incidence and the associated adverse effect on the secondary flows.

In both sets of experiments the same general flow patterns described previously, of the purge flow rapidly being swept towards the suction surface, were observed. The effect of higher rotation speeds was to decrease cooling effectiveness values and coverage as the purge flow was more rapidly swept towards the suction surface. The feature which dominated the effectiveness of purge-flow in these experiments remained the passage vortex, with very little coverage occurring downstream of the lift-off line. A combination of purge flow and downstream cooling holes was found to adequately protect much of the platform, though the pressure side remained relatively uncooled. The rotation speeds tested were far lower than typical aircraft engine speeds (a representative engine speed at the maximum take off condition is 13,000 RPM) and therefore the cooling effect of front seal purge flow would be expected to be further diminished in engine operating conditions.

Further investigations in a rotating transonic turbine stage by Pau et al. [66] identified regions of hot gas ingestion into the disc cavity and showed that these could be reduced by increasing the purge mass flow rate. The purge flow was shown to improve cooling effectiveness by reducing the strength of the passage vortex since it was entrained by the counter-rotating leg of the horseshoe vortex at the leading edge. The purge flow was again shown to only effectively cool the first half of the platform. CFD models of the experimental setup showed good agreement, however the predicted adiabatic film effectiveness values were shown to be overestimated compared to experimental values.

Purge flow from the front seal in rotating engine geometries has a tangential velocity component, that differs from the rotational speed of the rotor. This tangential velocity component is related to the rotor speed but is reduced by the shearing forces, between stationary and rotating surfaces within the front seal cavity, leading to purge flow being ejected at various swirl angles depending on the rotor speed. Barigozzi et al. [3] and Chen et al. [11] attempted to simulate the tangential velocity component in stationary cascades by including internal guide vanes within the purge slot. The guide vanes were designed to impart swirl to the coolant flow within the cavity prior to being ejected onto the platform surface. They demonstrated that the tangential velocity component of the coolant negatively impacts the cooling effectiveness. Reducing the tangential velocity component, resulting in more axially aligned flow, was shown to improve the area coverage, increase cooling effectiveness values, and reduce aerodynamic losses. At higher coolant mass flow rates, the reduction in tangential velocity led to purge coolant reducing the strength of the passage vortex.

2.5.2 Internal flows

Rotational effects have a strong impact on the velocity profiles of flow through internal convective systems used in blades [36]. Convective cooling channels have altered flow and therefore heat transfer characteristics, due to rotation depending on whether the flow is radially outwards or inwards. This is due to the Coriolis Effect which depends on the radial flow direction and direction of rotation. In typical multi-pass channels where the flow changes direction, flowing

first radially outwards towards the tip and then looping back and flowing inwards towards the hub, this causes adjacent channels in the multipass to have skewed velocity profiles. The associated secondary flow and heat transfer features within the convective passages are therefore also affected by rotation.

Rotational buoyancy additionally occurs due to temperature gradients in coolant flowing through rotating convective channels. As heat is transferred to the coolant its temperature increases and it becomes less dense than the fresh coolant being supplied. The relatively dense coolant is more strongly affected by centrifugal inertia pushing it radially outwards. The combined effects of Coriolis and rotational buoyancy on the flow within multi-pass channels are shown in Figure 2.9. The impact of rotational buoyancy is to flatten off the skewed velocity profile caused by the Coriolis effect. For platform cooling systems, which are employed at a relatively constant radius, the impact of Coriolis and rotational buoyancy is expected to be less pronounced. However, very little experimental research has been conducted exploring convective cooling in rotor platforms due to the expense and complexity of rotating facilities.

2.6 Experimental techniques

A number of techniques for obtaining surface temperature measurements have previously been developed and successfully applied to quantitatively assess the cooling performance of various cooling systems. In this section, the most common methods for measuring surface temperatures, which are subsequently used to determine cooling effectiveness parameters, are discussed. Prior research into turbine cooling systems has been conducted in a large number of experimental facilities of varying complexity. Broad categories of these experimental facility types are described and their relative costs and benefits discussed.

2.6.1 Surface temperature measurement techniques

Experimental investigation of cooling systems requires accurate measurements of surface temperatures from test components, from which metal effectiveness can then be calculated. A number

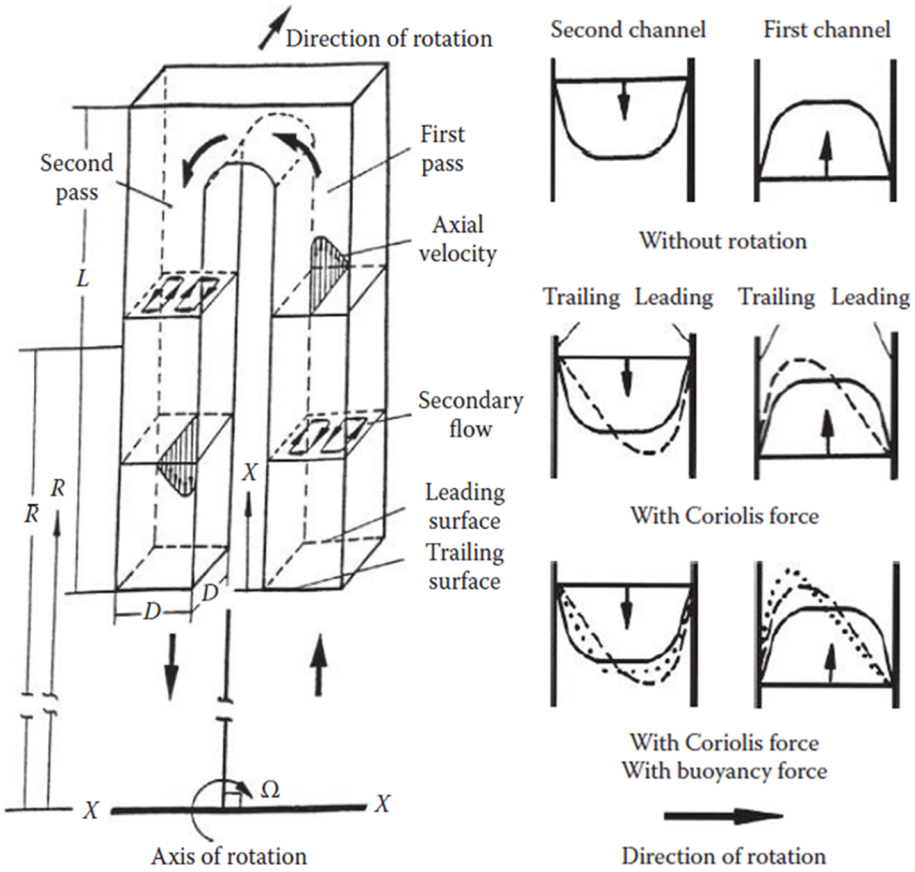


Fig. 2.9 Effect of rotationally-induced Coriolis and buoyancy forces on coolant flow velocity profiles through internal channels in a turbine blade multipass system. From Han et al. [36]

of widely-used methods exist and are used to obtain highly accurate point and area distribution measurements of surface temperature.

Point measurements are taken by thermocouples, which come in a range of specifications, which measure a change in voltage across the junction of two metals and calculate temperature via the thermoelectric effect. Point temperature measurements can also be taken by thin-film gauges which use the change in resistance of the gauge, due to a change in temperature, at a constant current. Heat flux measurements can be taken by using hot film gauges which measure the current required to keep a gauge at constant temperature. Film gauges are particularly useful in gas turbine research as many gauges can be printed onto the same film and can be easily applied to curved rotor blade surfaces.

Area coverage methods allow measurements of an entire surface rather than discrete points. Temperature sensitive paints (TSPs) are coatings which can be applied to complex surface

geometry and allow surface temperature to be inferred from light emitted by luminescent molecules. These molecules are excited by absorbing photons and then return to their ground state by emitting a photon of a longer wavelength. The emission intensity of molecules at high temperatures is lower than at relatively low temperatures and hence the surface temperature can be inferred from recording the emitted light intensity and a known paint response [87]. Similar pressure sensitive paints (PSPs) are also frequently used for film effectiveness measurements and in these the luminescent molecules are sensitive to partial pressure of oxygen. In experiments with PSP different gases are used for coolant and mainstream flows and the concentrations of these at each surface location are inferred from the paint response. This in turn allows for adiabatic film effectiveness and hence adiabatic wall temperature to be determined. The responses of PSPs are sensitive to temperature and must therefore be calibrated at the same temperatures used in experiments [13].

Thermal paint changes its colour permanently when exposed to very high, engine-representative temperatures. The change in colour depends on absolute temperature and the length of exposure. A number of thermal paints with differing bands of temperature sensitivity are applied to engine components in the same test to increase the accuracy of results. Due to the cost and complexity of testing at engine conditions, thermal paint tests are usually reserved for the final engine design stages. Costs associated with full-engine testing means that 'rainbow' configurations are often used which include different variations of cooled components in the same tests so that the best can be taken forward for extended endurance testing and production [53].

Full coverage temperature measurements in experimental facilities are commonly taken using thermochromic liquid crystals (TLCs) in both steady and transient experiments. TLCs are painted onto test surfaces and reversibly change colour, in the visible spectrum, when exposed to temperatures within a specific range. The relationship between the colour of emitted light and the temperature of the TLCs are monotonic which allows for accurate surface temperature measurements from prior calibrations by simply recording the TLC colour with a standard digital camera. The level of accuracy of TLCs and the range of available bandwidths and operating temperatures for laboratory experimental ranges gives them an advantage over PSPs and TSPs.

Infra-red (IR) thermography is a non-intrusive technique in which an infra-red camera is used to record light intensity emitted from a test surface. The emitted light intensity is a function of the local surface temperature. Highly accurate, full-area measurements can be recorded over a wide range temperatures and across complex surface geometries in both steady and transient experiments. Test components are typically coated in inexpensive high-emissivity paints which can be easily reapplied as necessary. IR thermography is ideal for testing in high temperature and pressure environments and when testing a large number of components, as surface preparation is simple. IR thermography was used throughout the experiments reported in this thesis and the setup and processing techniques are covered in more detail in Chapter 5.

2.6.2 Experimental facilities

Recent advances in computational fluid dynamics (CFD) have dramatically altered the way in which gas turbine research is conducted. Modern computers are capable of simulating increasingly complex systems of fluid flows and heat transfer at good levels of accuracy. The use of CFD is now a routine part of the design process of gas turbines which has reduced developmental time scales and allowed for novel concepts to be relatively inexpensively explored. Despite these advances, modelling highly complex conjugate systems such as that of the turbine rotor remains challenging. Accurate prediction of blade platform metal temperatures, essential for component service life and operation, relies on modelling the mainstream flow and associated secondary flow structures, SAS cooling flows and features, and conduction through the blade. It is generally acknowledged that current conjugate CFD modelling uncertainties are not yet satisfactory for the prediction of these critical component temperatures within reasonable computational times.

Experimental measurements therefore continue to be an essential aspect of heat transfer and cooling systems research. The role of experimental research has in more recent years shifted to become symbiotic with CFD studies, providing essential validation of numerical codes and providing appropriate system boundary conditions. Matching computational results to experiments is the primary way in which CFD uncertainties are currently reduced. Experimental studies which simulate more engine-realistic conditions are increasingly preferred for direct

component validation and fundamental research. Many types of experimental test facilities, of varying test conditions and complexity, have been used to analyse turbine rotor cooling. In general, an increase in fidelity of representing engine conditions comes at the expense of increased facility complexity, cost, and restricted measurement capability.

Noting the complexity of the rotor passage flow and heat transfer discussed above, the minimum requirement for experiments in which rotor platform cooling schemes are studied would be: representative mainstream inlet flow; correct rotor aerodynamic field; realistic SAS cavity reproduction; realistic leakage flows; good similarity to engine conditions (Mach number, Reynolds number, Biot number *etc.*); similar shank and fir-tree design to simulate blade root thermal boundary condition. With this level of engine fidelity replicated in experiments, it becomes possible to study the overall cooling effectiveness of the part in a meaningful manner. The Metal Effectiveness Rotor Cooling (MERC) facility was developed for this project in order to accomplish this, and is discussed in detail within Chapter 3. A brief summary of different types of experimental facility and their main advantages and disadvantages are now described.

Fully annular rotating facilities

Rotating facilities, which simulate one or more stages of a turbine, are extremely complex and expensive to build and operate but offer the most accurate simulation of engine flow conditions. They are the only means of testing rotation-induced effects on cooling systems. Engine geometry test components, with highly 3D designs, can be tested and internal SAS seals and flows can be replicated, making them ideal for testing late-stage designs at a reduced cost when compared to full engine testing.

As well as cost, the primary drawback of these facilities is the difficulty in instrumentation and obtaining high quality data from rotating components. Optical access is difficult to establish. High-speed cameras and complex processing are required to composite together measurements from multiple passes of rotating surfaces and correct for motion blur. Any physical instrumentation must be able to either record to an on-board data logger or communicate wirelessly during operation. Due to bandwidth constraints, this typically means that less detailed measurements

can be taken. There is therefore a trade-off between the high accuracy of modelled rotating flow and the level of measurement detail which can be obtained.

Annular cascades

An alternative to rotating facilities are stationary annular cascades which operate at conditions scaled into the rotor relative frame. These facilities model the mainstream flow through a ring of test components bounded by concentric hub and case endwalls, representative of engine geometry. They are well suited to tests in which 3D flow effects are of primary interest since actual engine geometry components can be used. The flow field is highly representative of the engine and includes radial pressure gradients, perfect periodicity, and all 3D flow features resultant from test components. Radial pressure gradients play an important role in the redistribution of cooling films in blade cooling schemes, however have lesser impact on endwall films which are dominated instead by passage vortices and crossflow. Annular cascades do not simulate rotational inertia effects and therefore flows are different to the true rotating situation. Nonetheless cascades have been used extensively for turbine rotor research and have been demonstrated to provide good representations of flows through rotor passages [39]. Disadvantages of annular cascades are the cost and complexity. Whilst less expensive than rotating facilities, the needs to manufacture a complete ring of test components and supply large mass flows of air lead to significant expenses. Modifications to test components are typically time-consuming and expensive as the entire annular ring must either be rebuilt or replaced.

Annular sector cascades, such as the Annular Sector Heat Transfer Facility at the University of Oxford, are a compromised solution to the expense of fully annular facilities. They consist of a segment of the annulus and so manufacture and operating costs are reduced whilst the 3D flow similarity is maintained. Annular sector cascades have typically found use when existing engine components are available to use and the desired experiments pertain to specific engine hardware. Achieving periodicity within an annular sector cascade is difficult as only a finite number of blades are included, and are bounded by endwalls. The 3D nature of the flows introduces strong radial and circumferential pressure gradients which cause significant challenges in conditioning the flow and establishing an acceptable level of periodicity.

Linear cascades

Linear cascades are a further simplification of the gas turbine annular flow path. They consist of a finite number of blades of constant profile, typically a linearised two-dimensional profile extruded across the entire span. The flow path can be considered to be a sector of the turbine annulus which has been 'unrolled' into a rectangular section. This means that radial pressure gradients cannot be replicated. Linear cascades have been extensively used to study a wide variety of fundamental flow phenomena. They have been shown to provide excellent two-dimensional approximations of the rotor flow field and are well suited to fundamental research and parametric studies of early design concepts. They are relatively inexpensive to build and run, and typically have good access for instrumentation.

The main drawbacks of linear cascades are that they cannot replicate the radial pressure gradients and rotational effects present in engines, nor use more complex 3D blade geometries. Similarly to annular sector cascades it can be difficult to establish periodicity in the blade passages due to the finite number of blades and bounding endwalls. True periodicity in a linear cascade would require an infinite set of blades. Typically fewer than 10 blades are used, limited by mass flow requirements and geometric constraints. Periodicity can however be achieved within linear cascades by using flow conditioning features such as exit tailboards, passages exhausting as a free jet into a downstream contraction, and boundary layer suction along the endwalls upstream of the bladerow.

Although generally considered to be 2D approximations of the flow path, the finite span and presence of endwalls means that all of the secondary flows previously discussed also develop within linear cascades (assuming certain geometric conditions such as tip gaps are replicated). The flow within linear cascades is therefore not truly 2D. This is advantageous as it opens up the possibility to study secondary flows with a reasonable degree of engine similarity and means that endwall heat transfer experiments, which are fundamentally reliant on simulation of secondary flows, can be conducted. Fundamental studies of secondary flow aerodynamics and heat transfer of, for example, film cooling have largely been conducted in low speed linear cascades as these are simple and cost effective to build and operate. The endwall boundary layers in these low

velocity flows are large and so super-sized blades, much larger than engine components, are typically used to try and diminish their effect.

Whilst low speed cascades can operate at Reynolds numbers similar to that in an engine, they cannot replicate the transonic flow typical of high pressure turbine rotor passages which is essential in more accurate replication of engine-like endwall boundary layers and the development of secondary flows. Experiments with a higher degree of engine similarity, though lacking radial pressure gradients and rotational effects, can be achieved in high pressure facilities which operate at matched engine Mach numbers and Reynolds numbers. These facilities require high operating pressures and are significantly more expensive and time-consuming to set up, not dissimilar to annular sector cascades, however they can provide the most engine-representative flow conditions achievable in a linear cascade. A high pressure, high-fidelity linear cascade has several advantages over an annular sector cascade for conducting platform cooling research. As engine parts cannot be used, test components can be simplified reducing manufacturing cost and making it easy to alter geometry for comparative studies, whilst retaining most essential features. Optical techniques used for IR capture are also easier to implement as platforms can be modelled as flat planes at a fixed focal distance without any angular distortion.

Chapter 3

The Metal Effectiveness Rotor Cooling (MERC) facility

The MERC facility is a high pressure, high temperature, transonic linear cascade which was designed, built, and commissioned to conduct the experimental work detailed in this thesis. The facility operates in a steady, blow-down configuration at scaled engine conditions (in the rotor relative frame), with a high degree of similarity to engine components to allow detailed experimental analysis of rotor platform cooling systems. The test section contains four rotor blades, including platform and fir-tree root geometry, and profiled endwalls forming five rotor blade passages. Also included is a representative secondary air system architecture including hub seals which allow for realistic coolant supply and for engine leakage flows to be replicated. A modular design approach throughout allows experiments to be conducted at relatively low cost and short timescales, making the facility ideal for parametric studies of novel cooling systems. A photograph of the facility is given in Figure 3.1.

The facility is designed to obtain high accuracy IR measurement of two blade platform surface temperature distributions. Metal effectiveness and adiabatic film effectiveness distributions are determined from measured surface temperature distributions in experiments using 3D-printed titanium and glass-filled nylon test components respectively. Downstream area traverse measurements are possible for aerodynamic loss studies. Detailed traverse measurements can be obtained during a single experiment due to the long run times (up to 6 minutes steady-state) possible in the facility.

In this chapter the overall facility is described, with a focus on novel and interesting aspects of the design. Over the course of the commissioning process and subsequent test campaigns a large number of incremental improvements were made to the facility. The description given here

is for the build of the facility at the time of writing (2021). Suggestions of further alterations to improve results or quality of life during facility operation are made following commissioning results in Chapter 5.

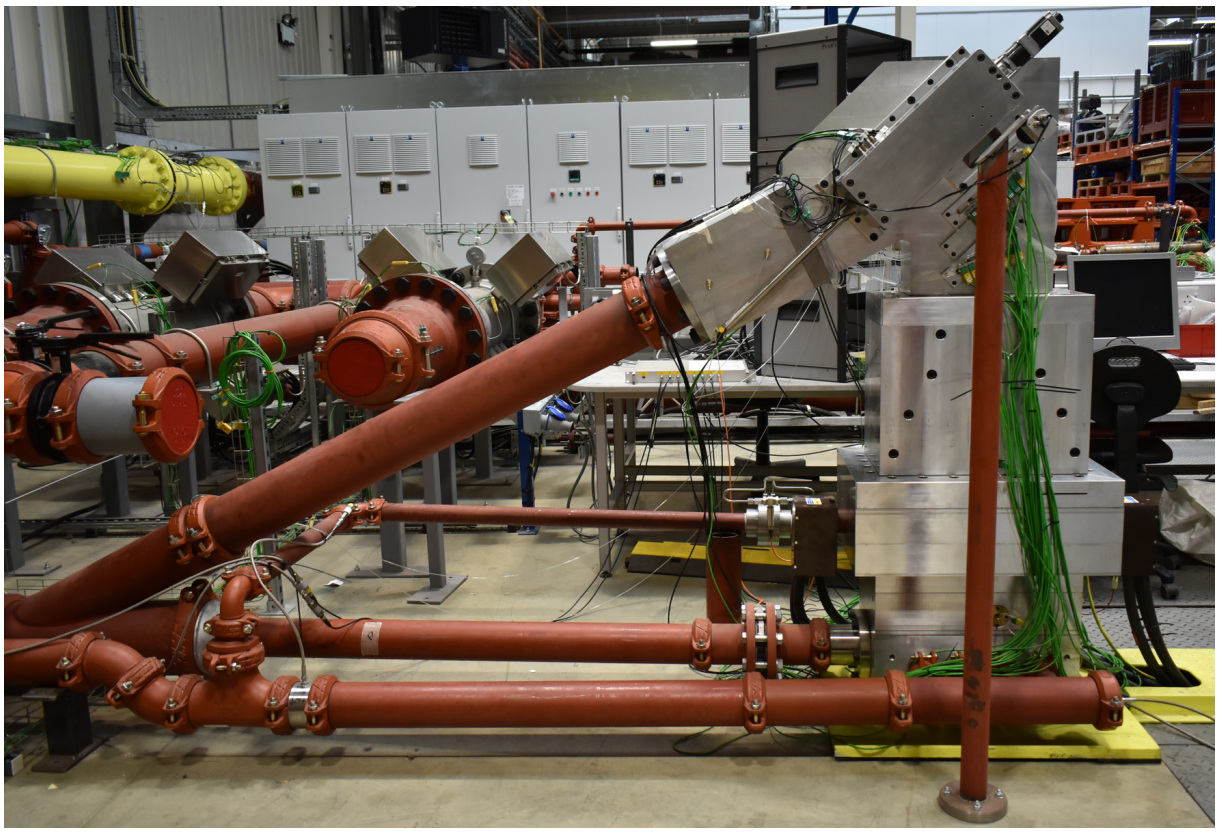


Fig. 3.1 Photograph of the MERC facility at the Osney Thermofluids Institute, University of Oxford.

3.1 Facility overview and operating conditions

The overall MERC facility and operation is now described; detailed descriptions of key components are given in later sections. A labelled blueprint of the main assembly is given in Figure 3.2, and a flow schematic is given in Figure 3.3.

Air is supplied to both mainstream and coolant lines from an air tank (approximately 60 m^3 pressurised up to 27.5 bara) through an early *Oxford Flow* prototype regulator which reduces the feed total pressure to 15 bara. Regulating the flow allows the facility to operate at a constant operating point for the full duration of a run. Run times of up to six minutes can be achieved before the tank supply pressure drops close to 15 bara and the regulator can no longer maintain

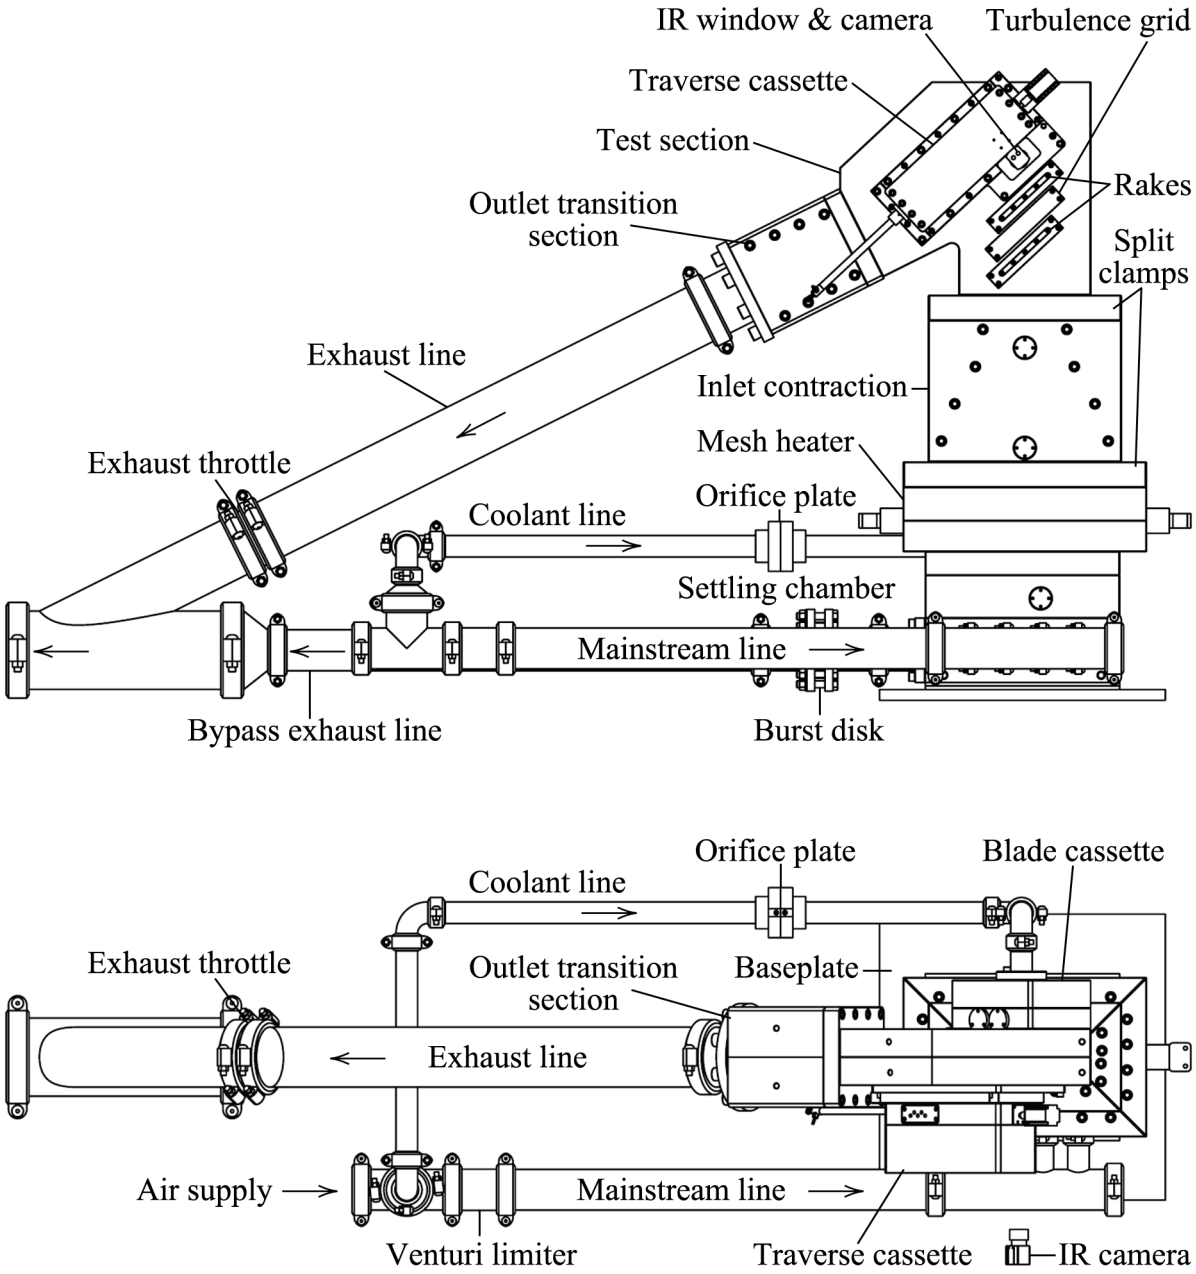


Fig. 3.2 Labelled assembly diagram of the MERC facility.

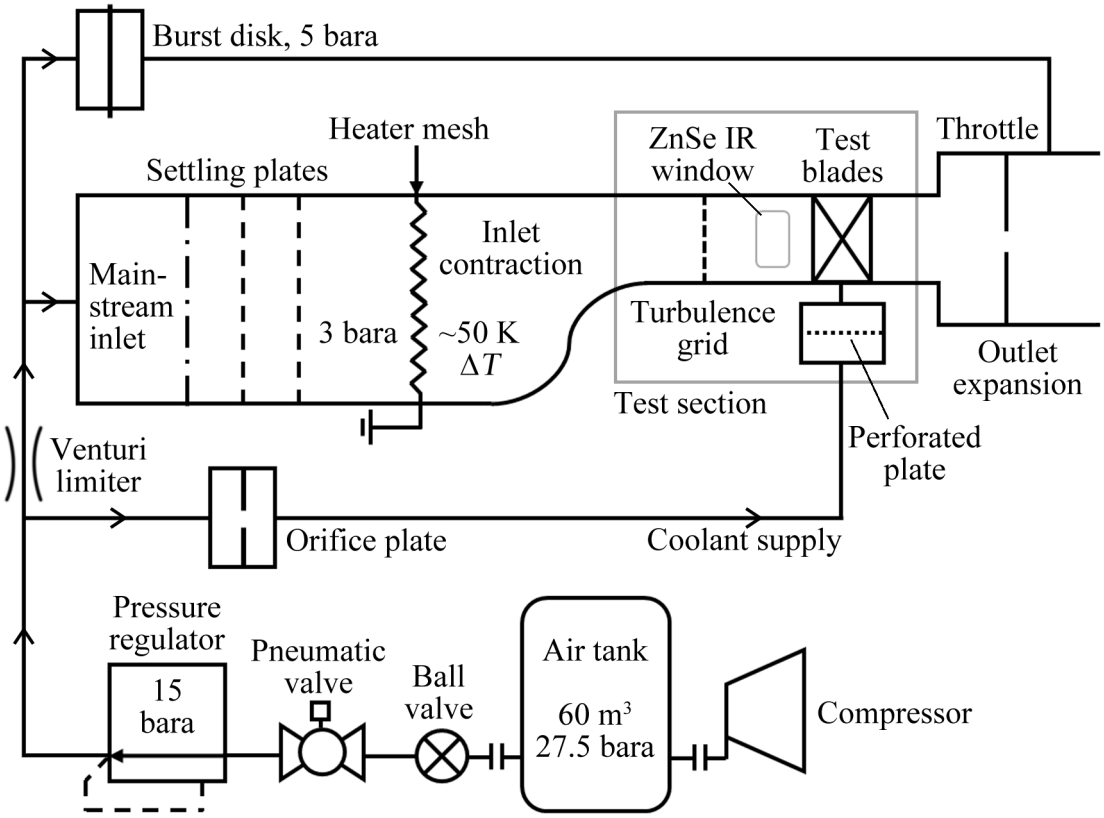


Fig. 3.3 Flow schematic of the MERC facility.

a constant outlet pressure. However, runs are typically limited to two minutes to increase the number of conditions tested and reduce the time between runs for tanks to recharge. The mainstream air is metered with an interchangeable, choked venturi nozzle (a diameter of 28 mm was used for all experiments), which acts to set a constant mass flow rate into the facility during the run period. The mainstream air then enters the main facility stack through a set of four perforated nozzles and passes through a series of settling plates to homogenise the flow. Once an uniform flow is achieved in the settling chamber, the flow is passed through a mesh heater which operates at a power of up to 150 kW. The hot flow is then accelerated through an inlet contraction situated just upstream of the main test section. The inlet profiles to the test section were designed to be uniform in both temperature and pressure however, for reasons discussed later, there are slight non-uniformities in both.

The coolant line is drawn off downstream of the regulator at 15 bara, and enters a small settling chamber before passing through a set of choked holes in a machined plate. The plate

performs two functions: it splits coolant into three flows that are distributed to individual plenums (supplying representative front seal, shank and blade, and rear seal), and acts as a metering device for the three streams. Total coolant mass flow is measured by an orifice plate assembly mounted along the coolant supply line.

Immediately downstream of the test section is a carefully designed expansion, into an area approximately four times the combined blade throat area. A transition section smoothly transitions the exit duct from the rectangular area to a 2 m long 200 mm diameter pipe. This design is based on the design of the miniature supersonic cascade by Povey et al. [71] who showed that good periodicity can be achieved by this free-jet exit configuration without the need for tailboards. This configuration lowers the minimum pressure that can be achieved at exit to sub-atmospheric pressures, allowing for higher operating Mach numbers to be achieved, using the self-ejector pumping effect. The back pressure is raised to higher values by varying an adjustable limiting nozzle at the exit of the ejector pipe. At the current operating point the back pressure is approximately 1.7 bara. Exhaust flow vents to atmosphere.

3.1.1 Operational considerations

The facility operating point (Mach number, Reynolds number, coolant-to-mainstream mass flow rate, and coolant-to-mainstream temperature ratio) can be varied by adjusting: the area of the mainstream venturi nozzle; the calibrated, choked coolant setting plate; the regulator set point; the downstream adjustable limiter; and the heater set voltage. Mach number and Reynolds number can be independently set by altering the inlet total pressure (by changing the fixed inlet mass flow boundary condition) and downstream pressure (by changing the adjustable limiter flow area). The facility is generally run at the blade design Mach number and Reynolds number but can easily be run at off-design conditions for both parameters, if required. Nominal operating conditions are tabulated in Table 3.1.

Tests are conducted by a pair of operators and are controlled remotely from a control room for safety. The operating point can be adjusted from the control room between runs by using a hand-adjusted regulator valve located next to the main control panel. This regulator sets the

Table 3.1 Nominal operating conditions of the MERC facility indicating achievable ranges.

Parameter	Symbol	Value	Unit
Inlet total pressure	$p_{01,\mathrm{m}}$	2.75	bara
Inlet turbulence intensity	Tu	6.0	%
Blade exit Mach number	M	0.81	–
Blade exit Reynolds number	Re	$1.1 - 1.72 \times 10^6$	–
Mainstream mass flow rate	$\dot{m}_\mathrm{m}$	2.25	kg s^{-1}
Coolant mass flow rate	$\dot{m}_\mathrm{c}$	0 – 0.112	kg s^{-1}
$\dot{m}_\mathrm{c}\sqrt{T_{0\mathrm{c}}}/\dot{m}_\mathrm{m}\sqrt{T_{0\mathrm{m}}}$	MR	0–5	%
Inlet total temperature	$T_{0\mathrm{m}}$	290–350	K
Temperature difference	$T_{0\mathrm{m}} - T_{0\mathrm{c}}$	0–60	K
Mainstream-to-coolant TR	$T_{0\mathrm{m}}/T_{0\mathrm{c}}$	1–1.2	K
MERC-to-engine Biot number ratio	$\mathrm{Bi}_\mathrm{rig}/\mathrm{Bi}_\mathrm{eng}$	≈ 1	–

pressure in a pilot tank, located near to the facility, which is in turn connected to the main facility regulator. The main facility regulator matches the downstream pressure to this pilot tank pressure. Wall mounted gauges display the pressure downstream of the regulator, in the supply tanks, and of the pneumatic valve system for monitoring during operation. Airflow and power to the mesh heater are both switched on during the run from the main control panel. A computer located next to the control panel is used to remotely start data acquisition.

Test duration is typically two minutes, during which time pressure and temperature measurements are taken throughout the facility and associated pipework. A typical short-duration run is conducted in a '10-30-40-30' second schedule where: 10 s of data are acquired with no air flow to log initial ambient conditions; 30 s with cold flow through the cascade at the operating point; 40 s heated period with the mesh heater turned on and hot flow; and 30 s with cold flow to cool down the facility, primarily the heater mesh.

The facility has a high density of monitoring instrumentation throughout and can simultaneously record these alongside IR camera and downstream area traverse measurements. For highly detailed downstream area traverses (low traverse speed and high spatial resolution of data points) the test duration is increased up to the maximum 6 minutes, limited by tank air pressure and the ability of the regulator to maintain a constant inlet pressure. Traverses are conducted with cold-flow as the heater mesh cannot be powered for longer than 1 minute at a time. The

turnaround time between tests can be as low as 15 minutes (time to recharge tanks) but may be longer if, for example, isothermal conditions are desired before the start of the following run.

3.1.2 Pressure rating and facility safety

The facility was designed as a pressure vessel, rated to 10 bara. The main body of the facility acts as a single pressure vessel but is split into sub-sections consisting of many separate components bolted together. Large diameter (4–8 mm) nitrile O-rings seal gaps between all components. Component stress distributions and displacements under load (displacements as small as 0.5 mm may lead to O-ring seal failure) were analysed using *SolidWorks* and resulted in using thick component walls and necessitated 10 mm radii fillets in the corners of rectangular sections to reduce stress concentrations. A safety factor of 3 on the Von Mises yield criteria was used on all components. The assembly was designed to minimise shear bolt loading and so the failure mode would be the hardened steel bolt stripping the thread of the relatively softer aluminium hole. Bolt stresses were assessed analytically to determine the safety factors against stripping hole threads. Safety factors greater than 3 are thus achieved throughout by using a large number of thick diameter (M8-M16), high-strength bolts. Supporting pipework is all suitably rated for the maximum 27.5 bara (400 PSI) of the air supply.

The facility was pressure tested using water up to a pressure of 15 bara. No component deformation was observed however significant weeping from several O-ring seals occurred. Due to several sections of the facility (inlet contraction, test section, transition section) having to be manufactured in two halves bolted together, and to upstream and downstream sections, there are a number of 3-way, perpendicular O-ring seals in use in the facility. At sufficiently high pressures the bonded O-ring joints deform enough to allow leakage flow through the seal. It was exclusively from these types of seal that weeping occurred. The nominal working pressure of the facility is 3 bara and hydraulic testing at 7.5 bara (2.5 times the working pressure) showed no weeping from any seals.

Prior to any experimental run, and before pipework pressurisation, a visual inspection of the facility is carried out by both operators to check pipework is correctly installed (including

regulator, venturi, and orifice plate) with victaulic clamps properly tightened and the facility main stack and all cassettes are assembled and secured. Whilst assembly work is being conducted on the facility the air supply is isolated by a pair of hand-operated ball valves, securely padlocked closed, and an automatic pneumatic valve, also locked to prevent opening even in the case of a misfire, located upstream of the regulator. Locks are removed and manual ball valves opened only immediately prior to experiments and are immediately put back in place once complete. During assembly, where supporting pipework may be disconnected, the main facility stack is prevented from being hit and toppling over, an unlikely event but possible due to the forklift path running directly next to the facility, by being bolted to a large 1" thick, 1×1 m ground steel baseplate.

During operation the facility is controlled remotely so that nobody is near the cascade. Facility over-pressure is prevented, in the case of user error or damage (such as the regulator failing which would result in the maximum 27.5 bara being connected to the inlet line), by a 4" bypass exhaust line and 6 bara burst disk connected to the inlet settling chamber, and the flow restrictions which would choke on each air line. These limit the mass flow rate into the facility and allow it to vent fully to atmosphere at pressures below the rated limit.

3.2 Scaling to engine conditions

The operating point of the facility has been determined by the requirement to match engine values of a number of non-dimensional groups relevant for aerodynamic and heat transfer similarity. The nominal conditions are those of a typical (conceptual) two stage high pressure turbine. Conditions in the rotor frame of reference have been taken. Mach number is matched by setting inlet total pressure, by specifying an inlet mass flow condition, and back pressure via a variable downstream nozzle. Inlet turbulence intensity is matched to a typical rotor inlet value by an upstream turbulence grid designed to give an inlet turbulence intensity of 6%.

Reynolds number, based on axial chord length, can also be exactly matched by setting the inlet total pressure to 4 bara. This requires a large mass flow rate of air which has two key disadvantages: the maximum mainstream temperature rise achievable with the heater mesh

is reduced, due to the maximum output of the power supply, which increases the uncertainty in IR temperature measurements (discussed in Chapter 5); the maximum facility runtime is reduced as the supply air tanks have a fixed, finite volume. The facility thus operates at a reduced pressure of 2.75 bara, requiring a lower mass flow rate, leading to higher possible mainstream temperatures and a small difference in cascade-to-engine Reynolds number commensurate with similar facilities of this type.

Coolant-to-mainstream momentum flux ratios of leakage and film cooling flow can be matched by scaling the coolant mass flow rate. Due to the achievable temperature differential the density ratio cannot be exactly matched. All coolant feeds have a specified mass flow rate boundary condition. The coolant mass flow rates can be independently varied by means of a coolant setting plate, held within the blade cassette, which has a number of choked holes machined into it. The total number and diameter of the holes, and the upstream total pressure, determines the mass flow rates into each downstream feed and can therefore be changed easily and rapidly by swapping plates.

Thermal similarity, important for metal effectiveness to be calculated from surface temperature measurements, is achieved by using a low-thermal conductivity material (titanium alloy Ti6Al4V, with $k \approx 6.6$ W m^{-1} K^{-1}) and slightly supersizing the blades (geometrically by a factor of 1.2). This allows Biot number, relating the surface to body resistances of heat transfer, to be matched between cascade and engine [46]. Increasing the size of test blades is useful not only for matching Biot number but also to make it generally easier to 3D print details and to instrument. The trade-off of increasing size is increasing the required mainstream mass flow and component cost.

Periodicity is achieved in the facility by introducing a sudden passage expansion at the trailing edge, which allows the blade passages to eject as a free jet downstream. CFD was conducted prior to the design of the facility to check periodicity and find the minimum number of passages required, to reduce both air consumption and component cost. An example midspan Mach number contour plot of the CFD domain used is presented in Figure 3.4a with the exit plane annotated. CFD exit Mach number characteristics are presented in Figure 3.4b for three, five and seven blade passages. Domains were created as 3D models and then meshed fairly

coarsely with an unstructured mesh of approximately 7.8 million elements. Simulations were run in *Ansys CFX* with $k-\varepsilon$ used as the turbulence model. These results demonstrate that the passage periodicity is reasonable for the centre-most passages in each case. Due to the high turning angle a pressure gradient exists in the tangential direction causing deviations in Mach number distribution in the outer passages, which is more pronounced for higher numbers of blade passages. A small separation at the pressure surface endwall can be seen (at the bottom of Figure 3.4a). Five blade passages were selected as the best compromise between acceptable air consumption and periodic flow conditions in blade passage of interest (those observed by IR camera) whilst allowing for two full platform areas to be observed. Moreover, the two observed platforms are flanked on either side by complete blade geometry, so that the internal cooling flows would match. CFD of the internal cooling flows was not completed during the timeframe of the present work.

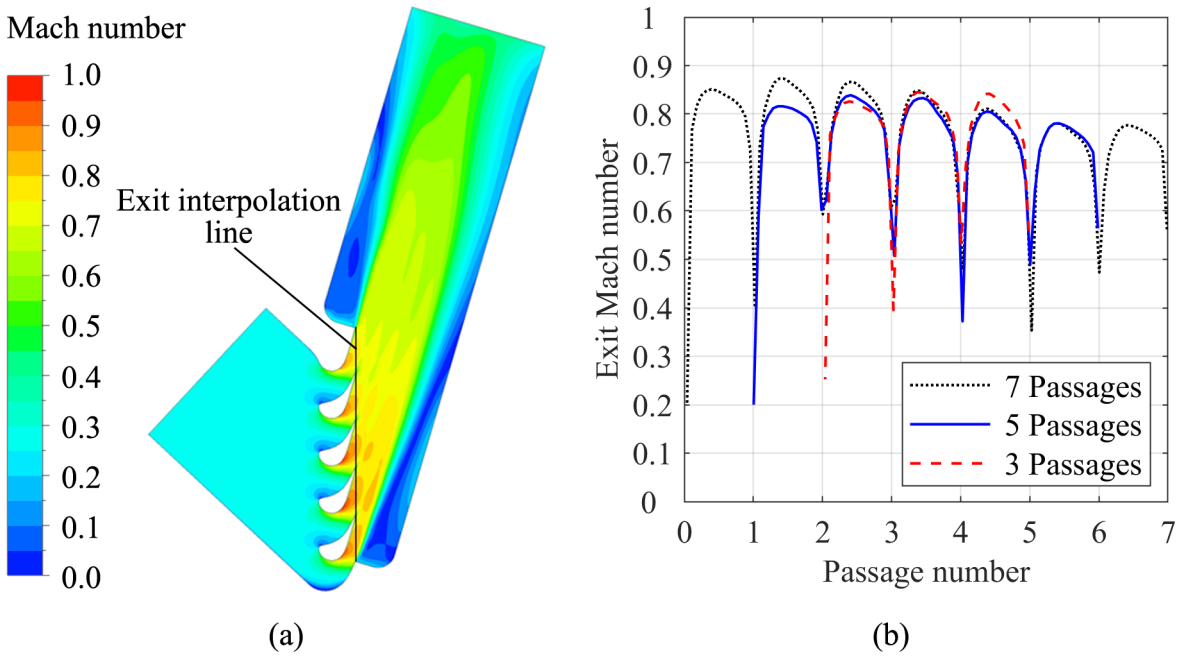


Fig. 3.4 CFD results from 3D simulations of the MERC facility test section ejecting into a sudden downstream expansion (a) midspan Mach number contour plot of 5-blade passage domain; (b) midspan exit Mach number distributions cascades with 7, 5, and 3 blade passages showing periodicity.

Simulations of a number of exit designs indicated that periodicity is established best when the expansion occurs immediately downstream of the trailing edge, with the centre of expansion

being the centre of the exit plane. That is, the area expands in both spanwise and tangential directions with the blade exit plane being centred in the newly expanded area. Due to the blade cassette being bolted from the outside and including a rear seal geometry, the expansion in the spanwise direction could not happen at the trailing edge as this would be into the rear seal. The expansion therefore occurs in two stages: a tangential expansion occurs at the trailing edge; a secondary spanwise expansion occurs further downstream, at the closest position mechanically permitted by the width of the blade cassette and required pressure vessel wall thickness.

3.3 Main stack assembly

The facility is arranged in a vertical stack configuration which accommodates the large rotor flow turning angle, reduces the cascade footprint, and has a number of assembly advantages. A CAD render of the main stack and cross section to view the internal flowpath is shown in Figure 3.5. Each major component can be added to the assembly sequentially using an overhead crane, with no complex alignment issues. Components of the main pressure vessel are bolted together, either directly or using two split-clamp rings. Nitrile O-ring seals are present between components. The bolt sequence was carefully designed so that the stack can be easily separated at several locations for rapid access to instrumentation and to remove and repair the mesh heater. The vertical arrangement and overall height of the facility (approximately 1.7 m) make it easy to access and work on any section. The modularity of the cascade, utilising a number of removable cassettes, means that changing components or instrumentation, and repairing the mesh heater can be achieved, in ideal conditions, within a single day.

3.4 Test section

The test section of the MERC facility was designed to be modular allowing for rapid refit and turnaround of test components and instrumentation. The main test section block was machined in halves out of two blocks of high strength aluminium alloy, *Dural* (Al 1080). The inlet and exhaust flow angles match those of the test blades for a total turning angle of 116°. Two replaceable

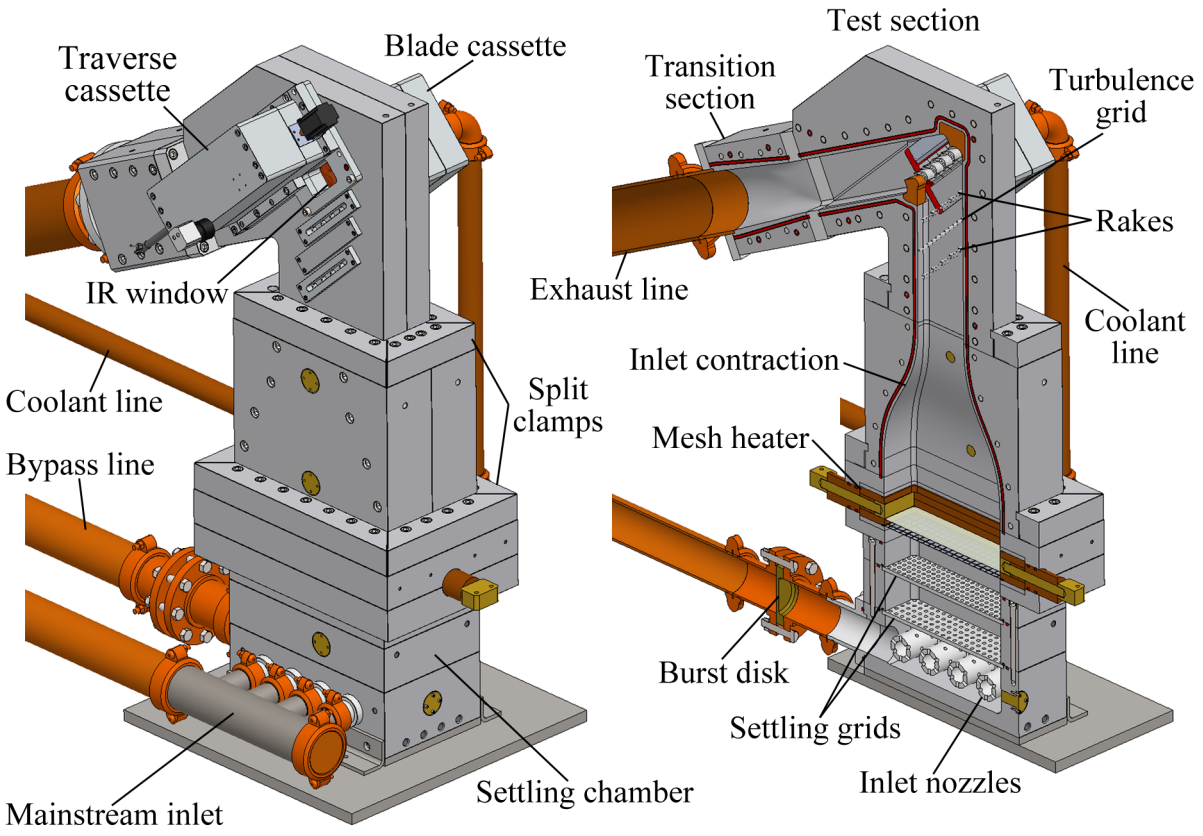


Fig. 3.5 3D CAD render of the MERC facility main stack assembly and cross section showing mainstream flowpath and key components.

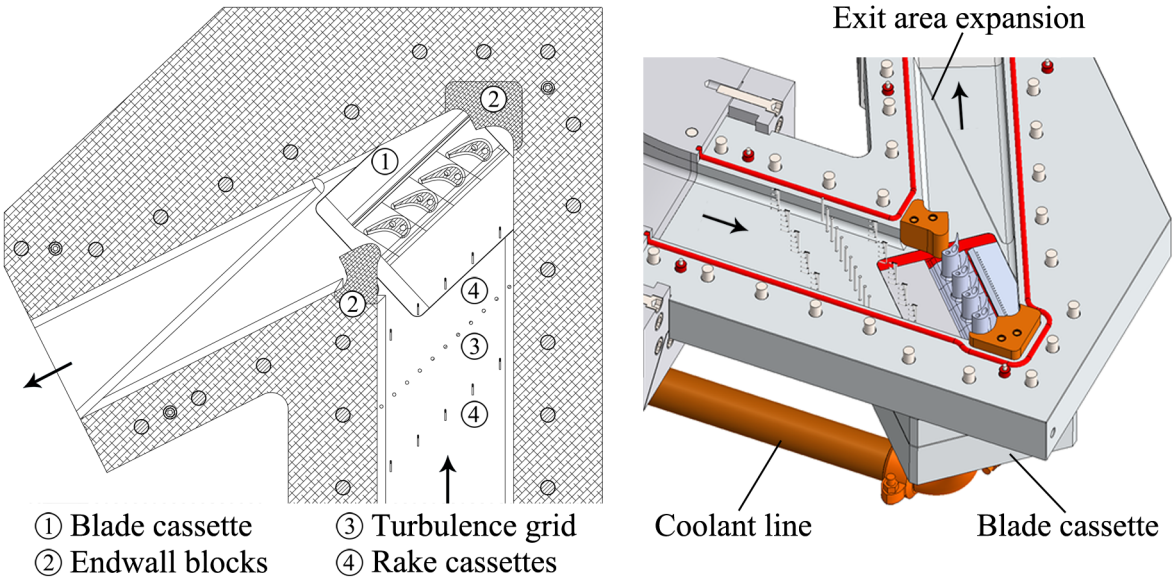


Fig. 3.6 Schematic diagram and 3D CAD render of the test section mainstream flowpath showing modular components and cassettes.

aluminium blocks form the endwalls and are machined with test blade upper and lower surface profiles. These blocks have cutaways at the trailing edge to form the exit expansion. There are no seals around the endwall blocks to prevent mainstream flow from leaking around them into the downstream expansion. This leakage flow is expected to be small due to the tight tolerances on these components. However, it is also hypothesised that this flow path, from relatively high to low pressure, may act as a boundary layer bleed along both endwalls and assist in achieving blade passage uniformity.

A number of self-contained cassettes, which contain instrumentation and test components, slot into the test section from the outside. Test components are held in a modular blade cassette which is bolted into the main gaspath and attached to the coolant line externally. The four blades held in the blade cassette, together with the contoured endwall blocks, form five blade passages. Upstream of the test blades are three cassettes, mounted on the opposite side to the blade cassette, containing two sets of probe rakes and a turbulence grid. A traverse cassette also attaches opposite the blade cassette and contains the traverse system and IR window module. Each of these cassettes can be assembled individually on a bench prior to being added to the facility.

3.4.1 Turbulence grid

An upstream turbulence grid, built into a removable cassette, provides an engine-representative turbulence intensity value of 6% (value supplied by Rolls-Royce via private correspondence), and corresponding length scale of 4.9 mm, at the leading edge of the test blades. The grid is designed according to the guides and correlations given by Roach [74] and is made up of a single row of round, 4 mm diameter, steel bars across the span. The porosity was recommendedto remain above 50% to avoid flow instability caused by the turbulence grid. Also according to [74] the wakes from these bars can be considered to have mixed out to a homogeneous turbulent flow at a downstream distance equal to at least ten times the bar separation. The turbulence intensity decays further downstream and so the grid is aligned with the inlet axial plane, rather than being mounted perpendicular to the inlet duct flow, to provide a uniform turbulence distribution at the

inlet of each blade passage. These requirements led to the bars being spaced by 15 mm and the grid installed 150 mm upstream of the blade leading edges. Round rods were used as they are insensitive to incidence angle to allow for this skewed (in relation to the duct) orientation. It should be noted that the turbulence intensity has not been directly measured within the facility and values have been assumed based on these correlations.

3.4.2 Inlet contraction

An inlet contraction section of the facility smoothly transitions the flow area from the heater mesh to the test section and is designed to create an uniform inlet flow distribution. The contraction assembly, a CAD render of which is given in Figure 3.7, is machined from two blocks of aluminium, bolted together and sealed with O-rings (displayed in red). Split clamp rings bolt the contraction to the upstream heater section and downstream test section. Brass monitoring ports (for measuring total pressure, static pressure, and temperature) are located at in the inlet and exit of the contraction on both sides.

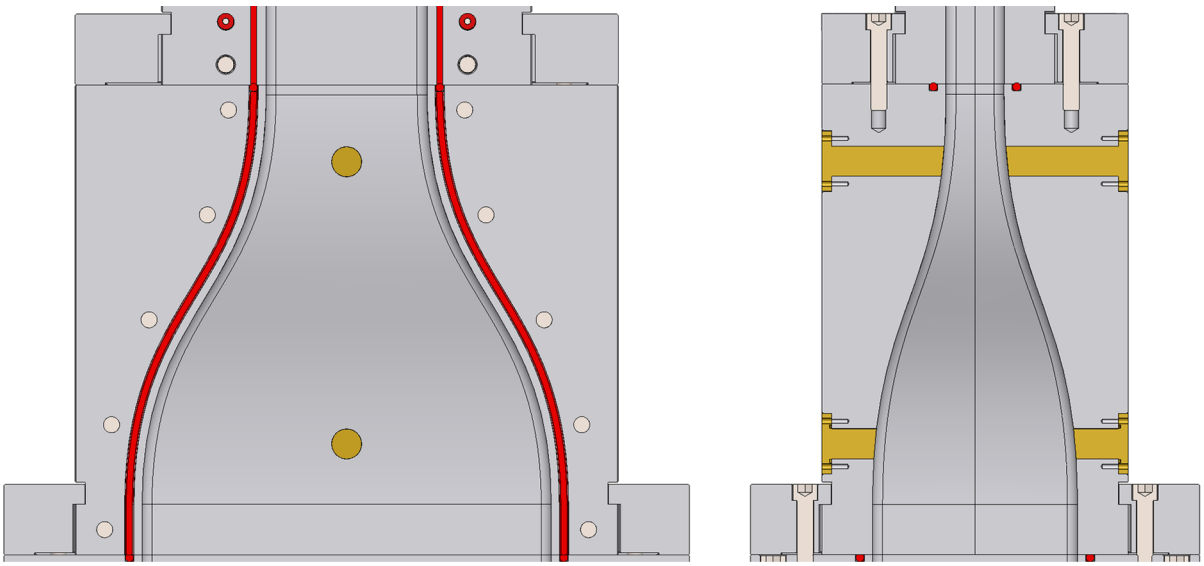


Fig. 3.7 CAD render of the two principle axis cross sections of the inlet contraction section assembly showing the main flowpath.

The contraction bounding walls are contoured to minimise adverse pressure gradients at the inlet and outlet of the section, which are both attached to constant-area rectangular cross-section ducts. These adverse pressure gradients must be minimised in order to prevent boundary layer

separation at the inlet to the test section. For a given area ratio (heater-to-test section ratio of approximately 8) the adverse pressure gradients can be reduced by increasing the streamwise length of the contraction [60], and by altering the contoured walls to more smoothly guide the flow.

Increasing the contraction length leads to more gradual flow turning by the curved endwalls which reduces the end pressure gradients. There is an associated increase in both required space and material component cost for longer contractions. Longer endwalls also result in the development of thicker boundary layers, somewhat mitigated by the acceleration of flow to an inlet Mach number of 0.22. The design challenge was therefore having a minimal contraction streamwise length that still produced clean inlet flow to the cascade with no boundary layer separation.

Two endwall contours were designed to transition between the respective long and short sides of the rectangular flow areas according to the methodology described by Cattafesta et al. [10]. Each contour is made by matching a pair of power-law arcs at a specified inflection point along the streamwise length of the contraction. Third-order polynomials were used since they were found by Watmuff [86] to minimise the inlet and exit pressure gradients. The polynomials are fit to geometric end points corresponding to the walls of the adjoining facility flow sections. At these end points the first and second derivatives are specified as zero so that streamlines are straight and parallel to the endwalls. Gradients are matched at the inflection point to create a smooth, continuous line.

For simplicity, the inflection point was set in the middle of the streamwise length which was 472 mm. Moving the inflection point towards either side of the contraction decreases the pressure gradient at that side and can thus be used to preferentially suppress either side. Rudimentary 3D CFD simulations in *ANSYS CFX*, using an unstructured mesh with inflation layers on the endwalls to resolve boundary layers and $k-\varepsilon$ turbulence model, demonstrated no separation of flow for this design and it was decided that no further optimisation was warranted.

3.4.3 Fast response heater mesh

Mainstream air is heated within the facility using a single mesh heater located immediately upstream of the inlet contraction. The heater consists of a finely woven stainless steel mesh suspended across the full 200×400 mm flow path, soldered to brass bus bars on either side through which a large current (up to 1500 A) is delivered from a 150 kW DC power supply. Power delivered to the mesh heats the metal due to its electrical resistance and this heat is then dissipated convectively into the mainstream flow as it passes through the mesh. The mesh wire is 130 μm in diameter with 180 μm apertures resulting in 34% open area fraction. Turbulence downstream of the mesh is expected to be very low and spatially uniform. The mainstream flow passes through the mesh at approximately 10 m s^{-1} which results in what can be considered a step-change in gas temperature (both in time and space) across the mesh when power is delivered. The temperature rise across the mesh is expected to be highly uniform, with previous research finding non-uniformities to be within 3.5% of the temperature rise [44].

This style of mesh heater was originally proposed by Gillespie [31] and has been used in a large number of experimental facilities. The design used here, for high pressure and large mass flow rate application, is based on the pair used within the Annular Sector Heat Transfer facility described in detail by Luque Martinez [56]. The mesh is built into an assembly, a CAD model of which can be seen in Figure 3.8, that is put together externally and then placed into an aluminium surround within the facility. The mesh and brass bus bars are mounted within an insulating material with high dielectric strength (tufnol, a synthetic resin bonded fabric). Layers of tufnol sandwich the mesh in place and isolate it from a pair of aluminium plates which are bolted together through the assembly, acting as clamps. Bolts are electrically isolated from the mesh by nylon sheaths. A filter mesh, of the same specification as the heater, and support grid are integrated within the assembly sandwiched between the upstream aluminium plate and a layer of tufnol. This filter prevents any small particulates from upstream pipework from getting caught on the heater mesh which could cause failure.

Each of the brass bus bars are made from three pieces. A rectangular cross-section bar is soldered onto the 200 mm edge of the steel mesh along one flat face. A round, 1 inch diameter, bar

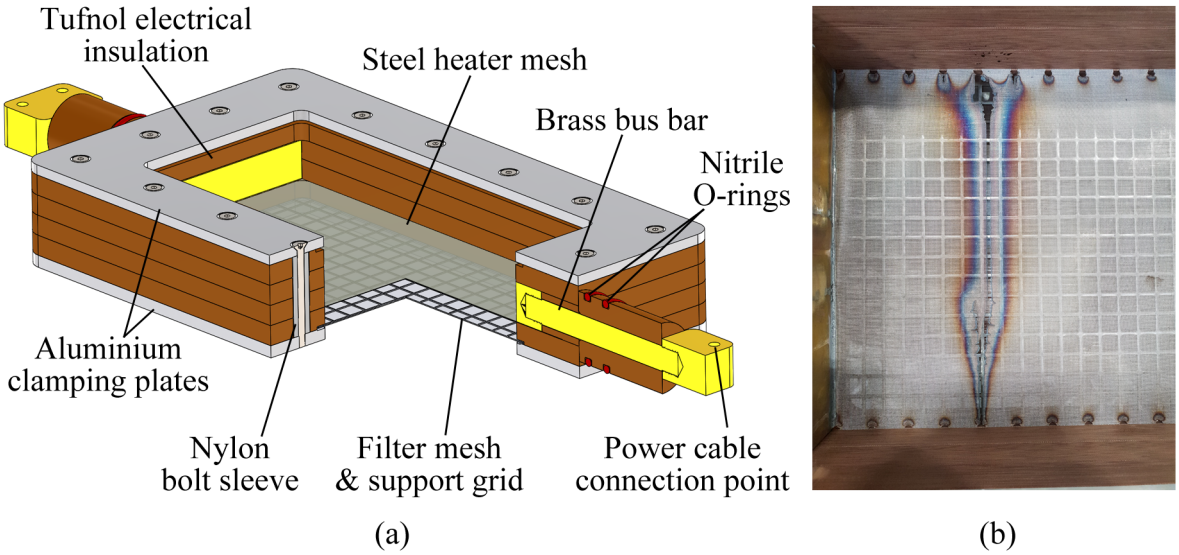


Fig. 3.8 (a) 3D CAD render of the heater mesh assembly which is assembled externally and then placed within an aluminium casing in the main stack of the facility;(b) photograph of failed burnt-out heater mesh

with threaded ends screws into the soldered piece and an additional rectangular block mounted externally to the facility. This external block is attached via through-bolts to cables from the power supply and is isolated within a tufnol box surround during facility operation. The round bar runs through a sealed tufnol sheathe which has as a pair of outer O-rings to seal the mesh assembly in the facility. The bus bars have a relatively large cross-sectional area to prevent overheating during operation which could damage the tufnol insulation.

Power delivered to the mesh is set by output voltage at the DC power supply which can be varied continuously between 0 and 110 V. The voltage across the mesh, which is lower than the power supply set voltage due to electrical resistance of cables and bus bar attachment points, is monitored during operation via a readout on the main facility control panel. A pressure transducer interlock system prevents power to the heater in low flow or no flow situations which would cause overheating and failure of the mesh.

The heater mesh design performs well and is reasonably robust despite the extreme power dissipation required. Failure typically occurs after 40–60 high powered runs, each having a 30–60 second heated period. The failure mode is characterised by a burned tear, perpendicular to the current direction, as shown in Figure 3.8b. It is hypothesised that failure begins as a localised burn which near-instantaneously propagates across the mesh in a runaway process

due to increased local electrical resistance. These tears always occur between two of the mesh cut-outs clamped within the tufnol layers. Cut-out features are present to prevent current from flowing through clamped sections of mesh. Mesh burn may begin at these cut-out edge due to conduction of heat from the nearby active part of the mesh causing overheating due to no air flow to cool the edges. A significant advantage of the facility being designed as a separable stack, and the heater assembly being easily removable, is that the heater can be repaired and replaced in a relatively short period of time when this infrequent failure occurs.

3.5 Instrumentation and data acquisition

Overall the MERC facility is instrumented with 78 pressure and 44 temperature measurements, the majority of which are located in the test section. Surface temperature measurements, of two complete blade platforms, are acquired via an IR camera setup. Downstream aerodynamic measurements are also acquired by an area traverse system. These are described in detail in subsequent sections. Pressures are measured via stainless steel tubes mounted either as wall taps or pitot probes and which are connected to banks of 5 bara *First Sensor CTE8000* transducers via urethane *Scanivalve* tubing. High pressure pipework is measured by connecting 35 bara transducers via stainless steel braided cables rated to 400 psi and standard fittings. Temperature measurements are taken by K-type beaded thermocouples mounted in various ways. Traverse multi-hole probe pressures are measured by a set of five 5 bara *Kulite ETM-506-375M* transducers.

Inlet conditions are monitored for uniformity using two sets of fixed upstream rakes mounted in cassettes, shown in Figure 3.9. Both cassettes can contain up to five rakes, each with six measurement points, evenly spaced across the full span, for either temperature or total pressure. The rakes are made up of small-diameter steel tubes, aligned with the openings perpendicular to the incoming flow and soldered together for structural strength. For temperature measurements, the same rake design is used but with beaded K-type thermocouples protruding from the end. Rakes are sufficiently far upstream (at least three axial chord lengths) and are only 1.6 mm in diameter to avoid distorting the inlet flow profile. A line of 19 evenly spaced surface taps

provides static pressure measurements downstream of the blade trailing edge. A pair of wall taps within the 3D printed housing monitors coolant cavity pressure.

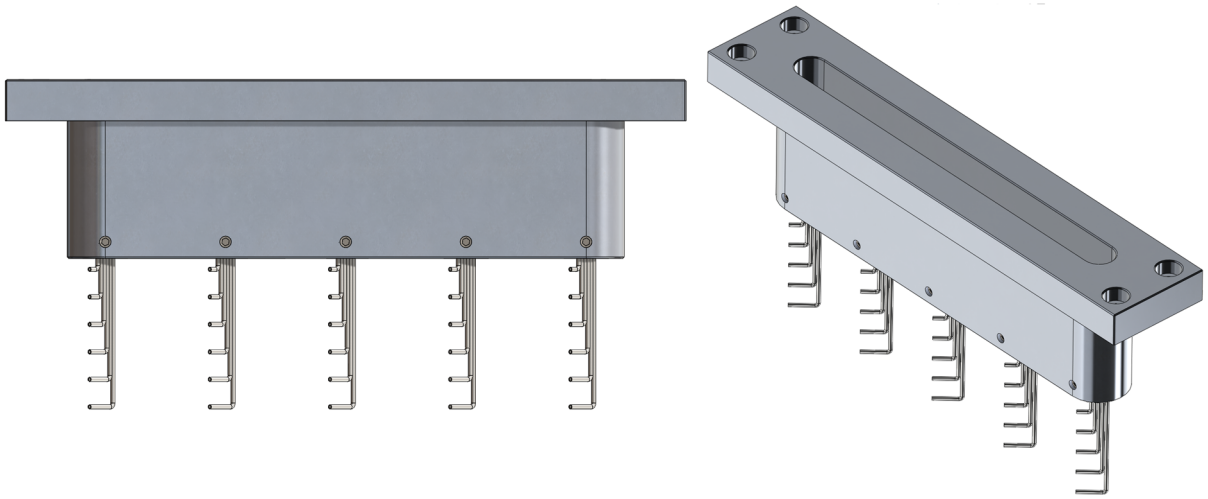


Fig. 3.9 CAD render of rake cassettes located in the MERC facility test section. Each of the two cassettes contain 5 probes each with 6 measurement points for either temperature or pressure.

Pressure and temperature conditions are measured at seven locations upstream of the test section via brass monitoring ports which each contain a wall tap, total pressure, and temperature probes. These ports are located at: the inlet and exit of the inlet contraction; upstream of the heater mesh; and in the first part of the settling chamber. Ports are typically arranged in pairs on opposite sides of the facility, the exception being the inlet to the settling chamber due to space constraints.

Pressures within the supporting pipework system are measured at key locations: the supply tank, downstream of the regulator, across the coolant line orifice plate, coolant supply plenum, and at several points in the downstream ejector pump line.

Data acquisition is controlled from a *LabVIEW* program which is manually started before a test run. The program collates and records voltage signals from pressure transducers, measured by a *National Instruments* (NI) 80-channel *NI-6225 PXI* card, and from thermocouples by a 68-channel *NI-6221 PXI* card. Thermocouple cables are plugged into two *NI TC-2095* terminal blocks, which are connected to two *NI SCXI* 32-channel *1102B* amplifier modules. Data is acquired generally for 30 seconds longer than run durations at a sample rate of 500 Hz. IR

camera, traverse, pressures, and temperatures are all acquired simultaneously. Recorded data files are post processed using custom *MATLAB* scripts.

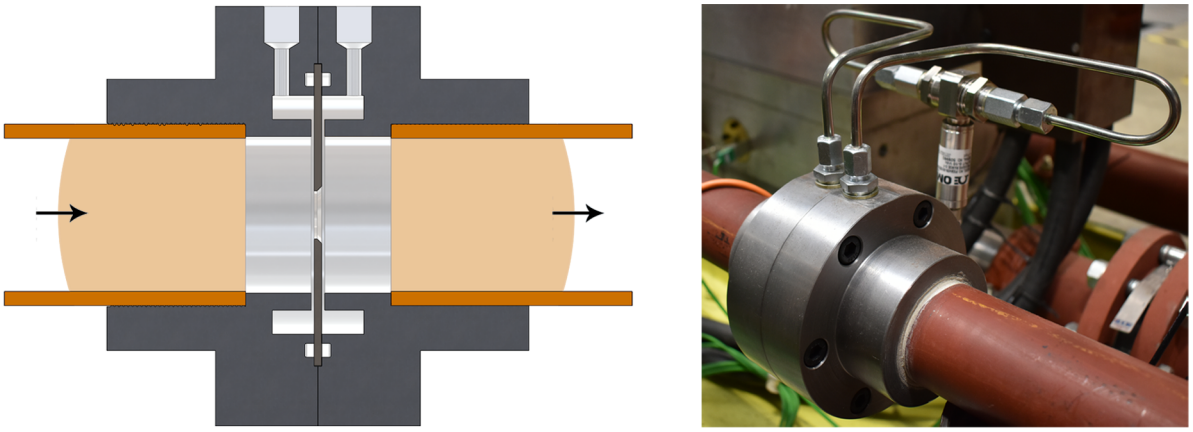
Data is acquired at 500 Hz for both pressure and temperature measurements. As the facility operates at steady conditions data is filtered with a low-pass filter with 20 Hz cutoff frequency and is down sampled to the Nyquist sampling limit of 40 Hz. This removes high frequency noise from measurements. Pressure transducers are periodically calibrated using an in-lab calibration system. Thermocouple measurements are normalised to an isothermal condition recorded at the start of any test day when the facility has had a minimum of 12 hours to settle to ambient temperature.

3.5.1 Coolant orifice plate

Total coolant mass flow rate is measured along the coolant line using an orifice plate assembly designed according to British Standards Institution (BSI) specifications [7], and shown in Figure 3.10. The assembly consists of a 2.5 mm thick, 100 mm diameter stainless steel orifice plate disk clamped between two machined stainless steel carriers. These carriers are threaded to connect to 2" diameter pipework with no flow area change. The orifice plate has a 16 mm diameter hole, chamfered on the downstream side by 45°. An O-ring seal prevents leakage around the orifice plate disk and sealant was used on pipe threads to seal the carrier assembly.

Annular slots, 1 mm wide, in the carrier at both faces of the disk act as corner tappings to measure the upstream and downstream pressures. These slots open into larger volume annular chambers which act to dampen any high frequency pressure changes and prevent a large transient pressure differential during pressurisation, which would damage the attached pressure transducer. Pipe wall tappings, 5 mm in diameter, which connect via stainless steel tubes to each side of a differential pressure transducer, break through from the outside of the carrier into the annular chambers.

The coolant line is characterised by a high line pressure of 15 bara but relatively low mass flow rate and hence, relatively small pressure difference across the unchoked orifice plate. This is an unusual case requirement for differential transducers however a suitable model (*Omega*



(a) (b)

Fig. 3.10 (a) Schematic cross section of the orifice plate assembly along the coolant line of the MERC facility and (b) photograph of the orifice plate assembly in-situ with attached differential pressure transducer.

PXM459), which is rated for a maximum line pressure of 35 bara and has a 1 bara differential pressure measurement range and 0–10 V DC output, was found. The proof pressure, the maximum differential pressure across the transducer terminals before damage occurs, is 4 bara and care was taken with the design to ensure that no transient pressure spike greater than this occurs during facility pressurisation. The transducer is mounted next to the orifice plate assembly and is connected via stainless steel tubes and standard fittings. The flow paths from both sides of the orifice plate are identical to prevent small pressure losses between the orifice plate and transducer from affecting measurements.

The orifice plate assembly is mounted on the longest straight-pipe of the 2" diameter coolant line. The upstream and downstream straight sections of pipework are approximately 19 and 12 diameters long in accordance with BSI guidelines to ensure clean and stable flow. As there are multiple elbows, in more than one plane, in both upstream and downstream pipework then even with these lengths of straight pipe there is an added 1% uncertainty to the flow coefficient calculated from this arrangement. A simple *MATLAB* script was used to calculate the mass flow rate from the measured pressure difference and estimate the uncertainty. Plots of coolant mass flow rate against pressure difference across the orifice plate and uncertainty in mass flow rate measurement are given in Figure 3.11. Typical maximum and minimum values of coolant mass flow rate used in the MERC facility are indicated by red dashed lines. Uncertainty on the

measured mass flow rate is larger for smaller quantities of coolant however expected values below 3% are achieved using this orifice plate arrangement.

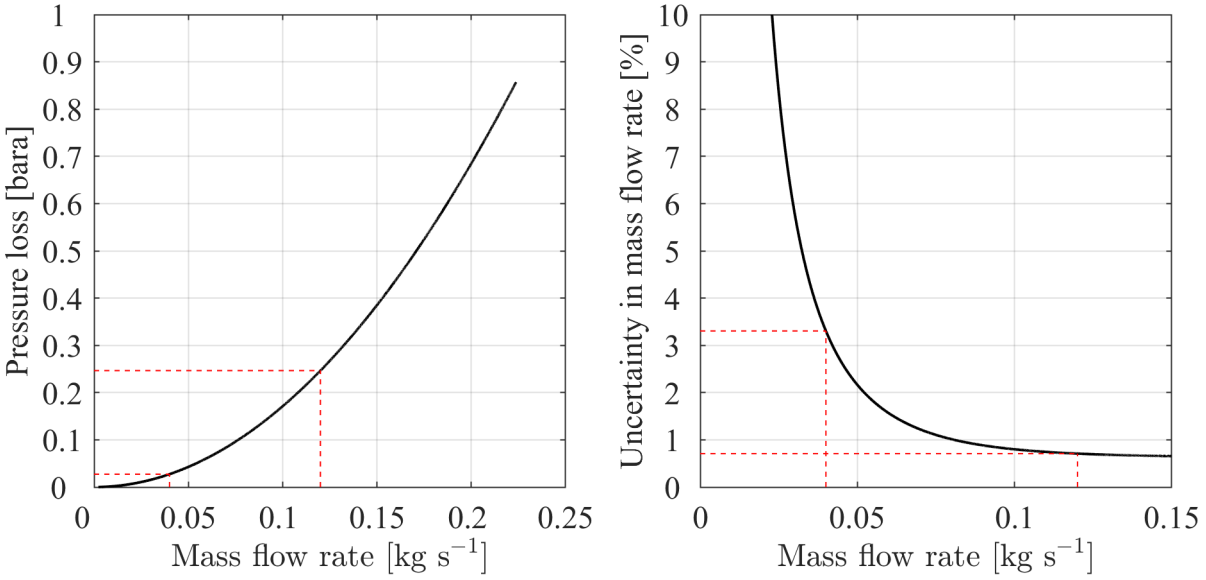


Fig. 3.11 Characteristic pressure loss across orifice plate and expected mass flow rate uncertainty plotted against mass flow rate for the MERC facility coolant line orifice assembly. Typical maximum and minimum coolant mass flow rate values indicated by red dashed lines.

3.5.2 Downstream traverse cassette

An area traverse system and IR window for optical access are contained within a traverse cassette mounted into the main test section. The cassette is assembled in two subassemblies attached to a thick aluminium baseplate and includes a pressure vessel in which the traverse system is housed, and an IR window module inserted into the test section. The pressure vessel is split into two halves which are bolted together and connected to the baseplate by a set of split clamps. This arrangement allows access to the traverse system when installed in the facility without removing the entire cassette and also reduces the weight, making it easier to handle. A CAD render of the traverse cassette, the traverse system, and a cross-section schematic are given in Figure 3.12. Photographs of the assembly mounted in-situ in the facility are given in Figure 3.13.

The traverse is a two-axis system in tangential and spanwise directions at an exit plane, just downstream of the blade trailing edge (approximately 1 mm downstream) and aligned with the bladerow. The system is capable of traversing the entire span for all five blade passages,

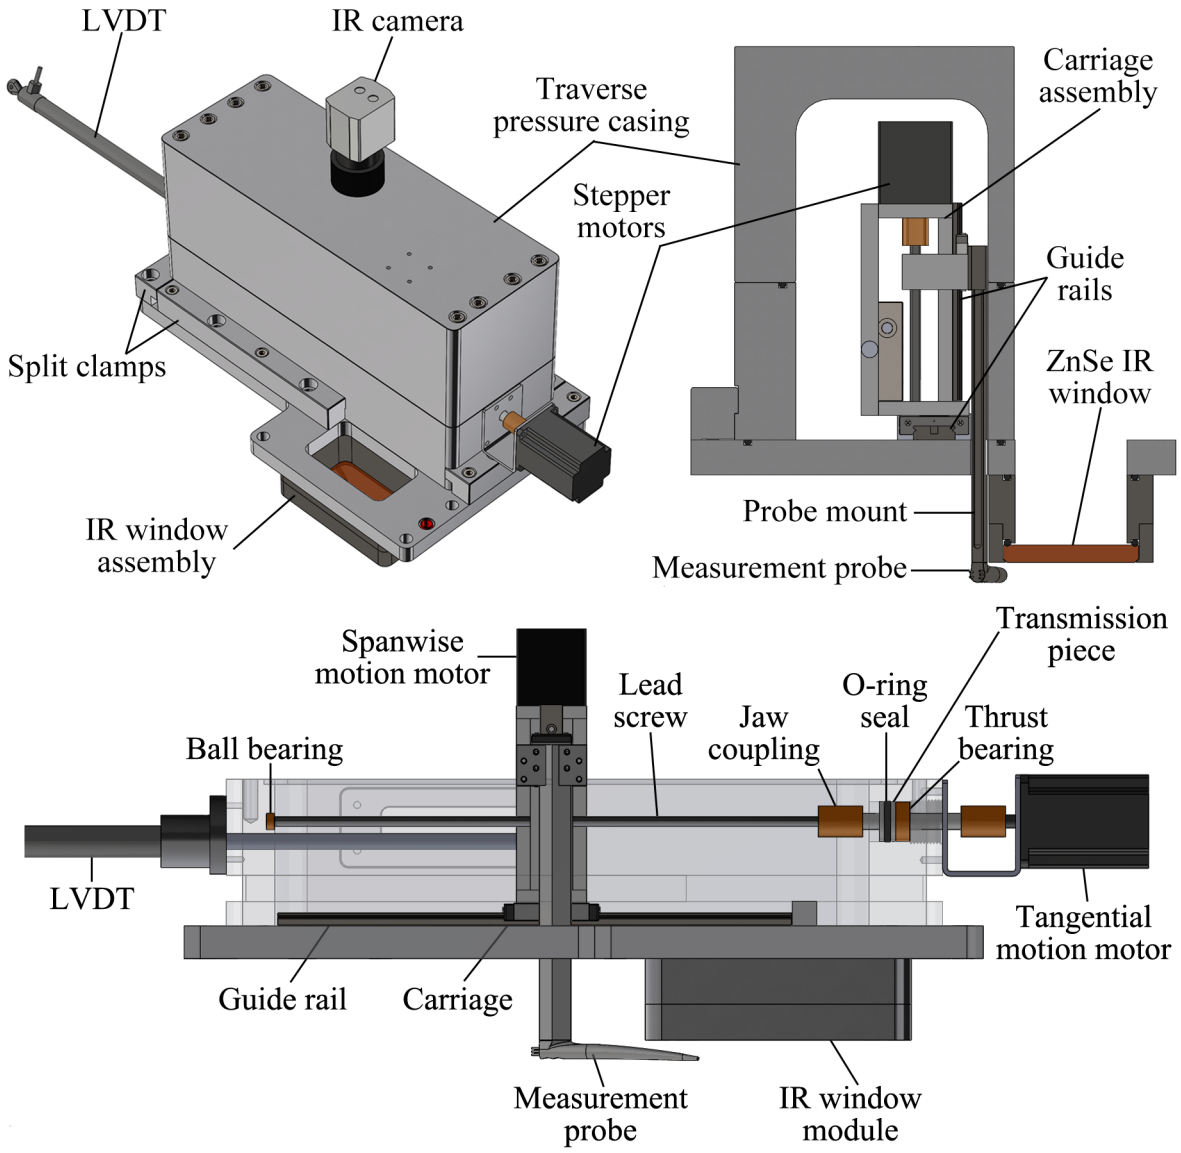


Fig. 3.12 3D CAD render and schematic cross section of the traverse cassette containing traverse system and IR window module for measurements in the MERC facility test section.

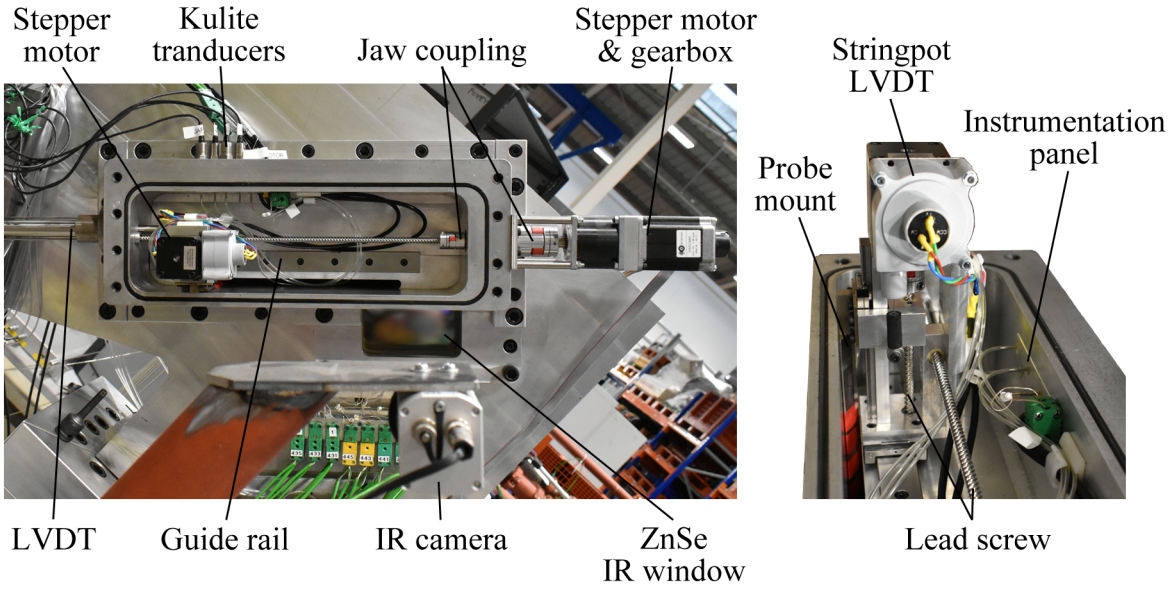


Fig. 3.13 Photographs of the traverse cassette installed in-situ with pressure vessel cap removed to show traverse system.

though typically only the middle two blade passages are measured. Accurate probe positioning is achieved using a pair of stepper motors and lead screws to move carriages along precision *HepcoMotion* rails. The main carriage moves tangentially and has an assembly, including motor, leadscrew, rail, and position sensor, built onto it to control spanwise motion. The probe is mounted to this second carriage assembly and enters the test section through a slot along one wall of the pressure vessel. The probe mount has a wedge-shaped fairing to reduce aerodynamic loading.

Motors were selected primarily to provide sufficient holding torque during facility operation. The possible loads on the system, due to the transonic flow into which the probe is entered, may be large and must not cause the probe to slip. The spanwise motor (*McLennan HT17-271D*) is capable of producing 0.24 N m of torque at 5 rev s^{-1} and was the highest torque model that could reasonably fit within the pressure vessel. The tangential motion motor (*McLennan 23HSX306E*) is mounted outside of the pressure vessel and is connected to a 10:1 gearbox allowing it to produce more than 2 N m of torque at 15 rev s^{-1}. Motors are run in half-step mode with a resolution of 400 steps per revolution and both drive lead screws with a 12.7 mm lead resulting in possible spatial resolutions of 31.8 µm for the spanwise direction and 3.18 µm in the tangential direction (due to the gearbox). The tangential motor is connected to the internal lead screw by

a transmission piece sealed with a shaft-mounted O-ring to prevent leakage. The transmission piece slots into a hole from the inside and is captive so that it cannot be ejected from the facility when pressurised. To prevent the system from locking up when pressurised, a thrust bearing is used to allow the transmission piece to freely rotate whilst pressed against the end cap.

Motors are connected to a pair of *Applied Motion Systems ST5-Q-NN* controllers. These power the motors and control the rotation speed and number of steps. The controllers are connected to the same computer as the data acquisition system and communicate with a *LabVIEW* program set up to automatically traverse a predetermined area. The position of the probe is monitored by a pair of linear position sensors (linear variable differential transformer, LVDT) which output a voltage proportional to the displacement from some reference position. The tangential position is monitored by a *POSITEK LIPS P117* slim-line LVDT which uses a rod connected to the main carriage and inductive sensor. The rod length determines the maximum displacement that can be measured and therefore the LVDT had to be mounted outside of the pressure vessel. A crush O-ring seal, built into a mounting bracket, seals the rod as it passes into the cassette. Inside the pressure vessel, the spanwise position is measured by a *TE sensor solutions SP1* compact string pot LVDT which uses a cable wound around a spool and rotational sensor. The output from each of the LVDTs is a 0–10 V DC voltage signal which is recorded by the same NI system described above. These signals are used to calculate instantaneous probe position during operation and are used by the *LabVIEW* program to calculate and correct motor movements.

External connections to motor control units, power supply, and probe measurements are via a removable instrumentation panel on the side of the traverse cassette. For multi-hole aero probe measurements, a set of 5 high-frequency *kulite* pressure transducers are mounted directly into the outside of the the instrumentation panel to minimise lengths of flexible tubing between the probe tip and transducers. This minimises the fill volume of the tubing and increases the frequency response of the probes.

The IR window module attaches opposite to the traverse system. It is made up of two machined steel rings which hold the IR window (full specification included in the next section) in position and include O-rings to prevent leakage from the test section out of the viewing port.

The window is held in a floating arrangement, with great care taken to ensure that there is no direct metal contact to any part of the window. A small 2 mm chamfer on the window rests on a lip on the bottom steel ring and prevents the window from falling into the test section when unpressurised. A layer of liquid silicone sits between the metal and window on the retaining lip and around the window perimeter. A 5 mm diameter, soft silicone O-ring seals the top face of the window. When the facility is pressurised, the window is pushed up against this O-ring compressing it and forming the seal whilst simultaneously preventing contact with the upper steel ring.

The overall cassette design was largely determined by the need to minimise the wall thickness between the IR window and the slot through which probes enter the test section. The traverse covers a plane in-line with the bladerow and, due to the large exit angle, small downstream axial displacements of the probe slot result in long probe lengths to reach the traverse plane. The challenge was designing a cassette in which the whole blade platform could be seen unobstructed through the window, and have a slot close enough to the trailing edge that probe lengths would be reasonable (approximately 10 cm) to minimise tip deflection and vibration from aerodynamic loading. This had a number of design consequences: the window module was included as part of the traverse assembly instead of being a separate cassette; the window module was machined from steel to reduce the required wall thickness to withstand pressure loading; the pressure vessel wall thickness was minimised along the slot-side resulting in the two halves being bolted together only at the ends; and split clamps were used to bolt the pressure vessel to the baseplate removing the need for additional wall thickness for a second set of bolts.

3.5.3 Infra-red hardware and capture setup

A single *Micro-Epsilon TIM 640* camera was used to record data of mainstream blade platform surfaces. The camera was mounted directly over the centre of the test section viewport, normal to the blade platform surfaces. Two interchangeable lenses were available: an 18.7 mm, 33° field of view lens; and a telephoto 41.5 mm, 15° field of view lens. The camera was primarily equipped with the telephoto 41.5 mm lens and was mounted to a steel bracket welded to a steel

pipe support and fixed in place with grounding bolts into the concrete floor. Early testing had the camera configured with the 18.7 mm lens mounted close to the IR window by an aluminium bracket attached to the cap of the traverse cassette. The configuration was changed due to the out-of-focus blade tips obscuring a significant portion of the blade platform surfaces, and also to eliminate potential camera issues due to facility vibration during operation. The camera mounting setup and a raw temperature measurement image from the camera can be seen in Figure 3.14 below.

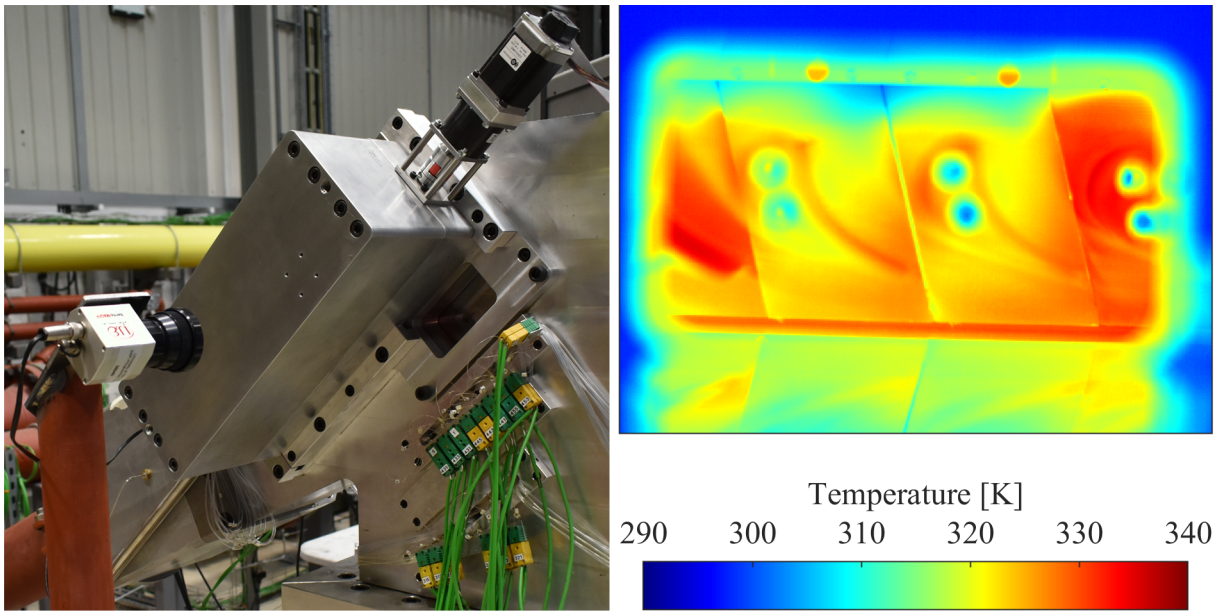


Fig. 3.14 (a) Photograph of IR camera mounted in-situ in the MERC facility and (b) image of measured temperatures by the camera, output from recording software during a heated facility run.

Camera focus is manually set by adjusting the lens and is fixed with no digital focus possible. Prior to mounting the camera by the facility the focus is set on a bench by measuring a pair of calibration disks. These are highly reflective, and optionally heated, steel disks with sharp edges to make setting the focus easy. The distance between camera and disks is carefully measured to be the same as the camera is from the blade platforms, when mounted.

The camera is connected via USB cable to the same computer which records data from the facility NI data acquisition system, and is recorded by *Micro-Epsilon's TIM Connect* software. Video is captured at a resolution of 640×480 pixels and framerate of 32 Hz (the maximum

capability of the camera) for the full run duration. An *AutoHotkey* script is called and executed by the DAQ system *LabView* script in order to automatically trigger the IR camera recording at the start of each run.

Optical access to the facility is via a zinc-selenide (ZnSe) window, contained in a module attached to the traverse cassette described above, and mounted flush with the endwall opposite the blade platforms. The ZnSe window (of dimensions $128 \times 77 \times 10$ mm) was manufactured by *Crystran* and has been treated with an anti-reflective coating, optimised to give maximum transmittance in the 8–14 μm range of the camera. The thickness of the window reduces transmissivity but is required to survive the internal pressure loading [14]. The camera has a clear view of two complete blade platforms (orientation is: flow inlet towards the top of Figure 3.14b), and approximately half of each of the adjacent two blade platforms. A strip, approximately 5 mm wide, of the front seal block can be seen and is instrumented with a number of thermocouples for in-situ calibration, discussed later.

3.6 Linearised rotor blades

The test blade design used within the facility is based on a typical aircraft engine high pressure turbine rotor blade. A schematic drawing of the test blades is given in Figure 3.15. Each component includes the lifting surfaces, platform, and fir-tree root. The main blade is a 2D profile extruded to the full span height of 57.6 mm (with 0.4 mm tip clearance). The profile is based on the 15% span profile of an engine geometry (close to the hub but above the fillet between blade and platform), which was unrolled from the radial design onto a plane surface. A summary of the blade geometric and flow parameters are given in Table 3.2.

The blade is mostly uncooled with the exception of two internal passages running radially, representative of a basic internal cooling system. The cooling channels run through the platform and will therefore have some cooling effect on it in engine designs. It was therefore decided to have these nominal basic internal blade cooling flow features to simulate this effect. Coolant exhausts from these cooling passages at the tip through two tapped holes. Brass inserts are threaded into these holes to change the flow area, and therefore the amount of coolant flowing

Table 3.2 Geometric parameters, inlet and outlet flow conditions of cascade test blades.

Parameter	Value	Unit
Blade axial chord	33.6	mm
Blade pitch	43.8	mm
Blade span	57.6	mm
Inlet relative Mach number	0.23	–
Exit relative Mach number	0.81	–
Inlet relative whirl angle	43.0	°
Exit relative whirl angle	73.5	°

through the blade. This configuration allows for fine-tuning of blade coolant mass flows, however for the work in this thesis only two insert types were used: inserts with 5 mm diameter central holes, and blanks which fully sealed the blades and prevented coolant from flowing through them. A squealer tip design was added with a cutaway towards the trailing edge of the suction surface to try and establish a clean downstream path for exhausted coolant. The aerodynamic and thermal performance of the blade tips were not assessed in the present work, however it should be noted that the MERC facility design is well suited for such investigations in future work.

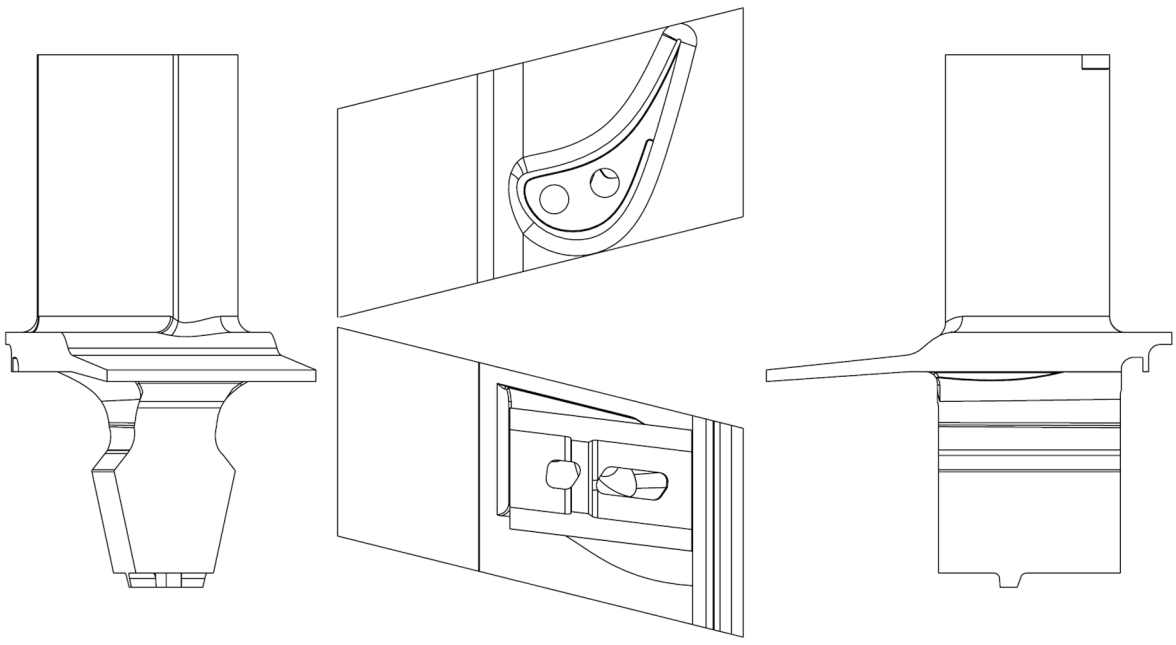


Fig. 3.15 Design drawings of the linearised rotor blades used in the MERC facility including complete blade, platform, and fir-tree root geometry.

The platform was similarly unrolled to create a flat surface. The fillet between blade and platform, which plays an important role in the development of secondary flow features, was preserved. The platform has a long front lip, which forms part of the hub labyrinth seal and which was significantly extended to increase internal overlap and allow for additional features to be added (guide vanes, discussed below) to influence coolant purge flow. A smooth curve connects the front seal lip to the mainstream flat platform surface, which is level with the upstream and downstream endwall surfaces, at the blade leading edge. There is also a platform trailing edge overhang and rounded slot on the rear which allows for sealing features (faceplates in an engine but an O-ring strip in the cascade). The sides of the platform (the slashfaces) are solid and flat, with no cooling or sealing features. When sealing of the slashfaces is required, liquid silicone sealant is used between blades during assembly.

The platform is angled by 14° with respect to the axial direction whereas the fir-tree root is angled by 7°. The internal surface of the platform is connected to the fir-tree root via a highly 3D loft, representative of an engine design and which is shaped to cope with extreme component structural stresses. It was decided to include this complicated loft connection as it covers a large portion of the rear platform surface and was expected to have significant effects on the platform temperature distribution. The loft design also limits where and how cooling features can be integrated onto the rear of the platform. The fir-tree root profile was simplified for ease of manufacture and assembly. The characteristic teeth which give the fir-tree root its name were filled in, resulting in a simple 'cubist' block design of approximately the same volume.

Sets of blades were manufactured via direct metal laser sintering (DMLS) in titanium alloy (Ti6Al4V) and 3D printed plastic (glass-filled nylon). To reduce overall project costs, only the two visible blades through the IR window are made of titanium during metal effectiveness experiments and are flanked on either side by plastic blades. Two sets of four plastic blades, which are significantly cheaper to manufacture, were made for use in adiabatic experiments (one spare set to replace damaged components). Metal blades are sand-blasted by the manufacturer to reduce surface roughness and are subsequently polished in-lab using a handheld rotary drill. Figure 3.16 shows photographs of a baseline solid metal blade that has been polished. The coolant exhaust holes at the blade tip are also tapped in the lab and brass inserts fitted, using

PTFE sealing tape to seal the threads. All test components are also coated with a thin layer of high-emissivity black paint.

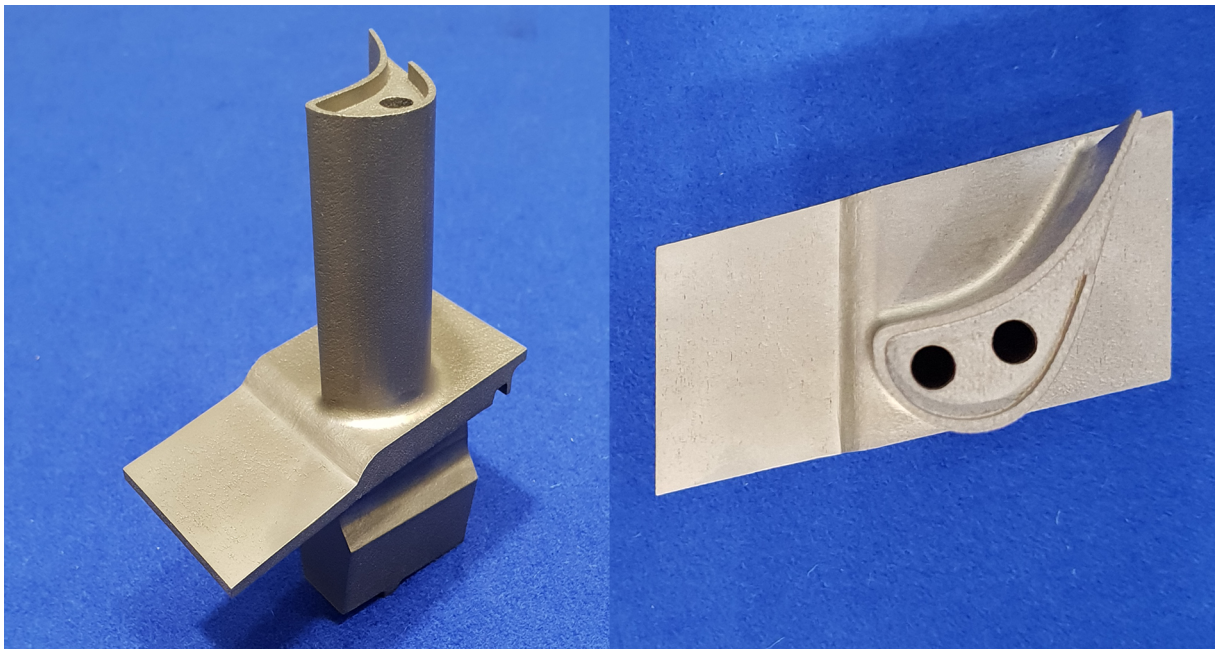


Fig. 3.16 Photographs of the baseline uncooled, linearised rotor blade manufactured via 3D printing in titanium alloy (Ti6Al4V).

3.7 Modular blade cassette

Test components are assembled within a modular cassette that contains a set of linearised rotor blades and representative secondary air system, local to the rotor blade, composed of several interchangeable parts. A schematic diagram showing a cross-section of the blade cassette is presented in Figure 3.17. The cassette is assembled on a workbench before being lifted and bolted into the facility from the outside. The modularity and easy access to remove and replace the cassette allows for rapid turnaround of test configurations, with alterations (e.g. replacing instrumentation or changing a single component) taking as little as half an hour to complete to a testable state.

The cassette splits into two main sections: the top half which contains test components and instrumentation; and the lower half containing a high pressure plenum which distributes coolant to the SAS. Each half has a thick-walled aluminium surround through which the cassette is

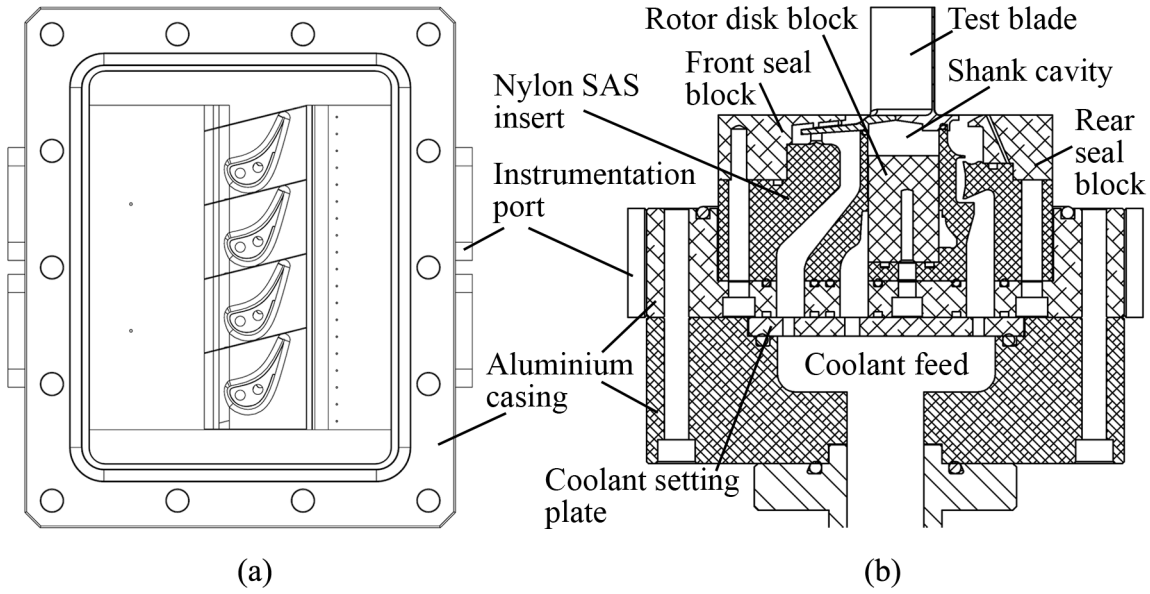


Fig. 3.17 Schematic diagram of the modular blade cassette which contains test components in the MERC facility.

bolted into the main stack. The top half is fully assembled on a bench before being inserted into the facility and is loosely held by dowel pins and retaining bolts. The lower half is then bolted through the aluminium casing of the top half, fully securing it in place, directly to the main test section block. The high pressure plenum is connected via a 2" steel flange to the coolant air supply. A steel plate is sandwiched between the two aluminium casings and contains a number of small-diameter choked holes which set the mass flow rate of coolant into three plenums in the representative SAS. Seals are integrated throughout the cassette to isolate each plenum from one another, allowing for independent control of coolant into each.

A range of coolant setting plates, with differing numbers of holes were manufactured for different mass flow configurations. The cassette assembly design means that the plate can be very quickly changed for back-to-back testing as complete disassembly is not required. The large through-bolts can be loosened to allow for the high-pressure half to slide back along them, leaving the top half inserted into the facility. The coolant setting plate can then be swapped out, like changing a cartridge, before re-tightening the bolts for the next experimental run.

The main body of the cassette, the top half, consists of a number of parts assembled within the aluminium surround. These are: a plastic insert representing the bulk of the SAS; a steel block representative of the rotor disk which holds test blades; aluminium front and rear seal

blocks which form the top third of their respective seal structures; and four instrumentation ports. These parts are all bolted together through the plastic insert, to secure them in place whilst preventing wear due to the bolting metal directly into plastic. Each of these parts are now described.

3.7.1 3D printed SAS insert

The majority of the SAS seal pathways are formed by a glass-filled nylon insert which was 3D printed in a single piece, as shown in Figure 3.17. These pathways are loosely based on real engine architecture but were made 'blockier' to increase strength for printing in plastic and to prevent damage during facility operation. Having the insert 3D printed, and in a single piece, allowed for the complex SAS structure to be relatively well preserved. This gives realistic coolant flows around the central test blades and also made sealing the three pathways relatively simple and achievable using O-ring strips and liquid silicone.

The front seal (left pathway in Figure 3.17b) is a pseudo labyrinth seal which directs flow radially outwards from the supply and then through a U-turn before ejecting as a purge flow along the platform. A solid wall divides the front seal from the shank cavity of the assembled blades. In a real engine this wall is formed by faceplates between adjacent blades. A set of round columns, designed to have minimal impact on coolant flow, support the extended leading edge lip of the test blades near the top of the seal.

The blade and shank pocket are both fed coolant from the middle pathway. Clearance gaps allow coolant to flow radially along the faces of the fir-tree roots and disk block into the shank cavity. Coolant also flows along the bucket groove, in the disk block, and enters the multipass system of each blade from below. The rear seal pathway also contains a labyrinth seal and exhausts into the mainstream to form a secondary purge flow downstream of the bladerow. A dividing wall separates the rear seal from the shank cavity and has an O-ring sealing strip along the top which seals the cavity just upstream of the blade trailing edge overhang. A pressure tap in the wall and flexible tubing routing through the insert allow for the shank pressure to be measured.

In addition to the coolant pathways, there are four internal, flow-isolated routes for instrumentation from within the cassette to outside of the facility via aluminium instrumentation ports.

3.7.2 Rotor blades and disk holder

A set of four linearised test blades, described above, are assembled in a wire-cut steel block representative of a segment of the rotor disk which has been unrolled. The rotor disk block has cutouts into which blades are inserted and which also forms the bucket groove supplying coolant to the blades. The fir-tree root of each blade slides into the block in a similar way to the assembly of rotor disks in an engine. As the fir-tree design has been simplified, and there is no rotation to push the blades radially outward, a pair of set-screws are used to raise each blade and hold them securely in position. As the fir-tree root and loft to the base of the platform have been preserved, a shank cavity is formed between adjacent blades. This cavity, between the bottom of the blade platforms and the top of the disk block is pressurised with the same coolant supply as the bucket groove and blade. The slashfaces between blades can either be left unsealed, resulting in leakage flow from the shank cavity, or sealed with liquid silicone sealant during assembly.

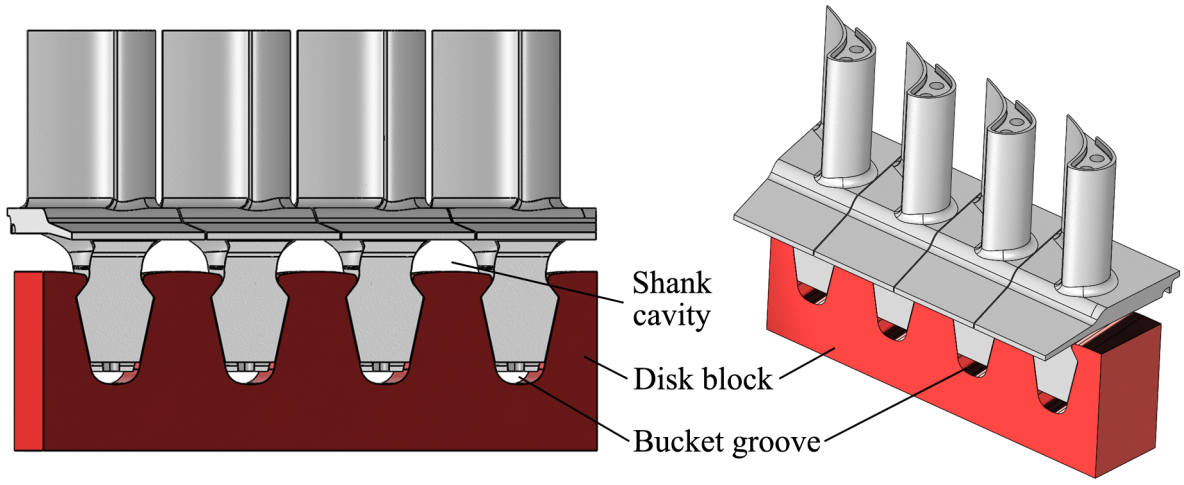


Fig. 3.18 3D CAD render of the test blade assembly including four linearised test blades assembled via fir-tree root in a block representative of an unrolled sector of the rotor disk.

3.7.3 Front and rear seal blocks

The top portion of the front seal is a machined aluminium block which forms part of the upstream main gas path endwall and has an overhang which overlaps with the blade leading edge lip. A set of guide vanes are machined on the underside of this overhang, to impart a tangential swirl component to purge coolant as it exits onto the platform surface. This is done to try and simulate the shearing between stationary and rotating parts in an engine. A number of these front seal blocks were manufactured with different guide vane angles, representing different engine speeds, and are discussed in greater detail alongside results in Chapter 7.

Each front seal block is instrumented with surface thermocouples located along the edge of the overhang on the mainstream side. These are used during IR camera calibration and processing to measure the wall surface temperature. There are two pairs of thermocouples with different mounting styles. One pair of bare beaded K-type thermocouples were epoxied into shallow pits drilled into the surface. These were however found to be difficult to distinguish on camera and were not sufficiently insulated from the surrounding metal, leading to erroneously low measured temperatures. A pair of improved thermocouple assemblies, designs of which were proposed by [59], were implemented. Each assembly consists of a 5 mm diameter, 0.5 mm thick copper disk with underside K-type thermocouple and 1 mm thick insulating layer of Rohacell foam, which has a very low conductivity ($k = 0.03$ W m^{-1}K^{-1}). These assemblies measure an accurate wall gas temperature as they are essentially adiabatic. The blocks also have a pressure tap and K-type thermocouple on the internal face allowing for temperature and static pressure measurements inside the front seal, immediately upstream of the guide vanes.

A band of the front seal, which includes the surface mounted thermocouples, is visible in the IR camera field of view and is used as a calibration strip for IR image processing. A number of round reflective markers, made from aluminium tape, are also adhered to the mainstream surface along the visible edge and were used only to check camera focus.

A similar aluminium rear seal block forms the top portion of the rear seal and features an angled leading lip to direct purge coolant flow downstream along the endwall. A line of 19 surface pressure taps, spaced at 5 mm increments, are located in a line along the block, parallel

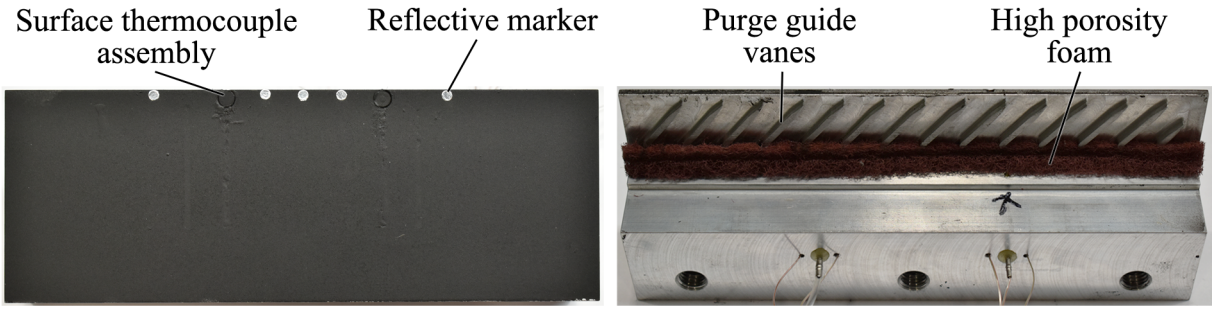


Fig. 3.19 Photographs of an instrumented front seal block including surface thermocouples and purge flow guide vanes.

with the blade trailing edges. These pressure taps are angled and use flexible tubing to run through the instrumentation pathways from a machined slot on the bottom of the block.

Chapter 4

Design of novel integrated platform cooling systems

This chapter describes the design of two novel platform cooling systems which were integrated directly into the platform geometry of the linearised test blades used in the MERC facility. The two cooling designs are: parallel convective passages running through the platform from the front seal and exhausting into the rear seal; and a jet impingement array targeting the platform from the shank cavity and exhausting spent coolant to the rear seal. Low-order numerical modelling was used to estimate average metal effectiveness and required coolant mass flow rate, and hence quantify expected cooling performance of both systems over a range of design parameters. Results from these numerical models, written as *MATLAB* code, were used to inform the design of a first iteration of each cooling system, which were then manufactured in titanium alloy and used in experiments.

4.1 General design strategy and discussion

The intent of the present design work was to produce viable cooling systems that can adequately cool problematic hot spots on the rotor platform, discussed in Chapter 2, whilst minimising the cycle efficiency penalty incurred. The rotor platform is inextricably linked to the SAS seal structures present in gas turbines and therefore integrated cooling systems must consider the impact, and opportunities afforded by these surrounding systems. Each of the systems developed and tested in the present work attempt to mitigate the cycle efficiency penalty by making dual-use of coolant for direct cooling and assisting in sealing the downstream rear seal. Spent coolant is exhausted into the rear seal and mixes with make-up coolant supplied to fully seal the downstream

wheelspace. Coolant used within the platform cooling systems has a significant pressure margin allowing it to effectively seal the rear seal, which is located downstream of the blade row, where pressure is relatively low. Whilst thermal pickup through the platform increases the temperature of the coolant, this is not expected to be significant enough to cause damage or overheating of the rear seal structure. The heat picked up by coolant will largely be offset by expanding into the seal and mixing with fresh make-up coolant.

The two cooling schemes selected, convective passages and jet impingement, are well-established technologies used extensively in NGV and blade cooling. Despite their prevalence in blade cooling, and a number of proposed systems in filed patents, they have seldom, if ever, seen actual use in platform cooling and experimental analysis. The designs of each system here are purposefully simple implementations of the technology, optimised to first-order by numerical modelling. The aim of the designs is to demonstrate the viability of such cooling schemes and highlight key design considerations. Models developed to understand the factors affecting each systems performance can be easily extended and used to design more complicated systems in future design efforts.

The use of additive manufacturing to have metal blades made with cooling features directly integrated affords a lot of creative flexibility. Manufacturing methods capable of producing printed engine-components that can withstand the extreme operational stress loads are still early in development but promise to revolutionise engine component design. Qualitative consideration to the mechanical loads on rotor blades was afforded during the design process, however detailed analysis was not performed. For experimental components, additive manufacturing allows for intricate novel cooling designs to be implemented and rapidly manufactured, reducing overall test times and associated costs. Thermal similarity is achieved by careful selection of the printed material however the manufacturing process also has implications (surface roughness, minimum features sizes, *etc.*) that affect component design compared to current engine-design methods. These were also briefly considered and a number of validation tests of printed components was undertaken.

4.2 Convective cooling passage concept overview

Rotor platforms are an ideal candidate for convective cooling features as it is formed of a pseudo flat-plate geometry in which features can be easily implemented, and there is readily available access to multiple high pressure coolant supplies from the internal seal structures and blade supply. The two primary candidates to feed coolant to convective cooling passages integrated within the rotor platform are: the blade cooling supply, either directly from the blade passages or from the bucket groove via the shank cavity; and the front hub seal.

The blade pseudo-multipass system is fed from the disk bucket groove and passes radially through the platform into the blade. Platform passages could feasibly originate from this radially-flowing multi-pass system. As coolant from the multipass is used to film-cool the blade surface, the pressure of supplied coolant is high enough to overcome the significant static pressure on the blade leading edge. This gives a large potential pressure margin for any platform system to make use of. The drawback of such a system is that the platform cooling design would be intrinsically linked with the main blade cooling design and any coolant drawn from the multipass would have an impact on the blade cooling. Adding the additional balancing of a connected platform cooling system to the already highly complex blade multipass design process would likely lead to increased development time and reduced design freedom.

The front seal is pressurised with coolant in order to prevent hot gas ingestion into the wheelspace. This is another high-pressure cavity since pressure in the front seal must be able to overcome the highest mainstream pressure, near the blade leading edge, at the inlet of the rotor. Platform cooling passages could originate from the internal rotor-side of the hub seal and make use of this high pressure coolant. Coupling the platform cooling system to the front hub seal is less problematic than to the blade multipass system as, although the flow within the front seal can be highly complex, the number of flow paths are limited (supplied coolant either remains within the front seal cavity or leaks into the mainstream via the gap at the rotor-stator interface) and the overall design focusses on overcoming the high rotor inlet pressure, whilst minimising leakage into the mainstream. It would therefore be relatively simple to negate the impact of drawing coolant from the seal, on seal performance by increasing the amount of coolant to maintain

required seal pressure. In addition to this, there is a large pressure difference between the front and rear hub seals giving an ideal flow path for convective passages.

The rear seal is similarly pressurised with coolant but is required to overcome much lower static pressures at the outlet of the rotor. Coolant could thus be drawn from the front seal, used to convectively cool the platform via internal passages, and then be exhausted into the rear seal. An added benefit of this would be that coolant used to cool the rotor platform, as long as the thermal pickup of the coolant through the system resulted in acceptable exhaust temperatures, could be reused to help pressurise the rear hub seal. Reusing coolant in this way has the potential to offset some of the engine cycle efficiency penalty, caused by taking coolant from the compressor, since a lower amount of additional coolant would need to be supplied to the rear seal in order to establish the required seal pressure. This front-to-rear seal flowpath was therefore selected for development.

The design space for such a system is broad and a number of constraints were chosen to reduce design parameters. For the initial design process it was decided to keep the system simple, with a small number of parameters which could be investigated and optimised with low-order numerical models. This led to the simplifying constraints:

1. Constant-area passages of circular cross section.
2. Passages run unidirectionally through the platform (i.e. no u-turns or backtracking).
3. Passages are evenly-spaced and continuous, with no merging or crossing paths.
4. Passages run at a constant depth below the mainstream platform surface.

Design parameters that could be altered given these constraints were the passage diameter, spacing, and depth into the platform.

4.3 Modelling convective cooling passages

In order to understand the impact of various design parameters, and aid in the design process, a numerical model was written in MATLAB to simulate flow through the convective passages and

estimate the cooling performance. The system was simplified by approximating the platform as a flat plate with parallel passages running internally from one side to the other. As passages are evenly spaced along the entire length, a periodic domain was drawn and the model therefore only needed to simulate a single passage to estimate overall system performance. The passage is split into a number of equally-sized discrete cells with changes in flow properties across each cell calculated incrementally from the inlet. Boundary conditions were specified to model experimental conditions in the MERC facility, with the ultimate aim to tune the model with experimental data before applying it to engine conditions. Due to the limited experimental testing this was not completed and so this model was used exclusively to aid in the design process. A diagram showing the simplified platform geometry, periodic domain, and discrete passage cell is given in Figure 4.1. Despite the relative simplicity of this system, modelling the passage is complex as the flow is compressible, friction is non-negligible, and heat is transferred from the mainstream through the metal surface and into the coolant.

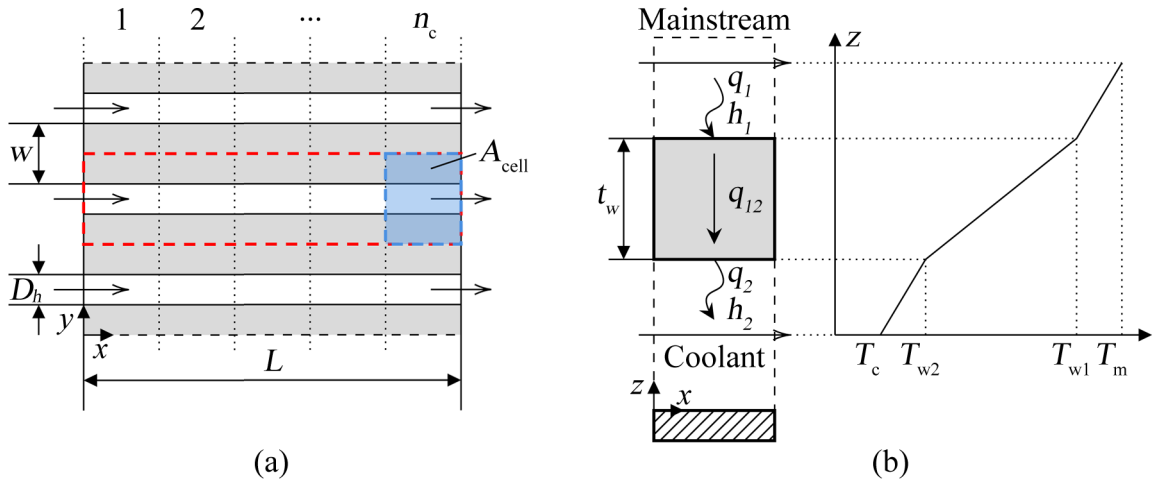


Fig. 4.1 (a) Schematic of convective passage numerical model domain and (b) discrete cell heat fluxes and thermal profile

4.3.1 Compressible flow through passages with friction and heat addition

Expressions describing steady, 1D flow of a perfect gas through a passage are derived by considering an element of a generalised system, as shown in Figure 4.2. For completeness, a passage with a change in area is considered which has heat entering the system and a friction-

induced shear force (τ_f) acting upon the perimeter (L_p). Conservation of mass, momentum and energy across this element can be expressed as differential equations.

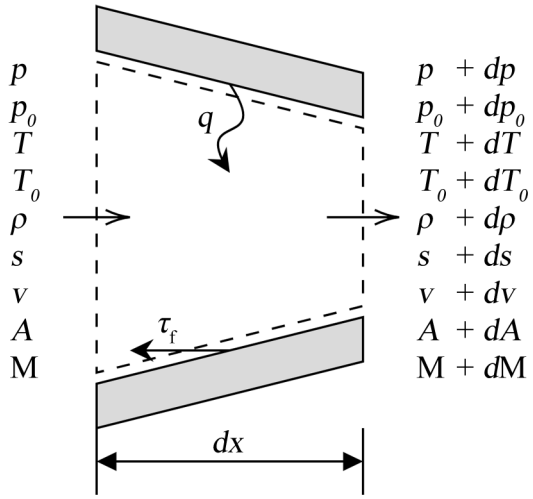


Fig. 4.2 Schematic of an element of a generalised compressible passage flow with friction and heat addition.

Mass conservation:

$$\frac{d\rho}{\rho} + \frac{1}{2}\frac{dv^2}{v^2} + \frac{dA}{A} = 0 \tag{4.1}$$

Momentum conservation:

$$\frac{dp}{p} + \frac{\tau_f}{p}\frac{L_p}{A}dx + \gamma \mathrm{M}^2 \frac{dv}{v} = 0 \tag{4.2}$$

Energy conservation:

$$\frac{q}{c_p T} - \frac{(\gamma - 1)}{2}\mathrm{M}^2 \frac{dv^2}{v^2} - \frac{dT}{T} = 0 \tag{4.3}$$

The frictional shear force can be expressed in terms of the dimensionless Darcy friction factor (f_d). The perimeter on which friction acts can also be expressed in terms of the hydraulic diameter (D_h) of the passage

$$f_d = \frac{4\tau_f}{1/2\rho v^2} \tag{4.4}$$

$$D_h = \frac{4A}{L_p} \tag{4.5}$$

Substituting equations 4.4 and 4.5 into 4.2 gives

$$\frac{dp}{p} + \frac{\gamma \mathrm{M}^2}{2}\frac{f_\mathrm{d} dx}{D} + \gamma \mathrm{M}^2 \frac{dv}{v} = 0 \tag{4.6}$$

Using these conservation equations, the ideal gas law, and definition of Mach number, the change in Mach number across the element can be expressed as

$$\frac{d\mathrm{M}^2}{\mathrm{M}^2} = \frac{\left(1 + \frac{\gamma - 1}{2}\mathrm{M}^2\right)}{(1 - \mathrm{M}^2)} \left[\gamma \mathrm{M}^2 \frac{f_\mathrm{d} dx}{D} + (1 + \gamma \mathrm{M}^2)\frac{dT_0}{T_0} - 2\frac{dA}{A}\right] \tag{4.7}$$

This differential equation allows us to calculate Mach number at any location down the passage, given some starting inlet value. It also gives us insight into how each factor (friction, heat transfer, and area change) affect the flow down the passage as each term in the rightmost bracket is the respective contribution from each factor. The denominator $1 - \mathrm{M}^2$ indicates that the behaviour changes, either increasing or decreasing Mach number along the passage, depending on if the flow is initially subsonic ($\mathrm{M} < 1$) or supersonic ($\mathrm{M} > 1$). Expressions for other flow properties can be similarly found and expressed in terms of M and $d\mathrm{M}$ and hence can also be determined at any point along the passage [77].

In the classical case of an isentropic flow with area change (for example, a nozzle), the first two terms in the rightmost bracket of equation 4.7 are zero since there is no friction and it is an adiabatic system. A converging area, $dA < 0$, results in an increase in Mach number for subsonic flow, and decrease in Mach number for supersonic flow. Conversely a diverging area, $dA > 0$, results in a decrease in Mach number for subsonic flow, and increase for supersonic flow. A constant-area passage with $dA = 0$ can only have a change in flow properties when either friction, heat transfer, or both occur.

Fanno flow refers to flow through an adiabatic, constant-area passage with friction. In this case the 2nd two terms in the rightmost bracket of 4.7 disappear. Friction causes an irreversible rise in entropy, with corresponding pressure drop, along the passage. Air is an approximately calorically perfect gas (constant c_p) at subsonic and low-supersonic speeds and it can be shown mathematically that for such a flow the maximum entropy occurs when $\mathrm{M} = 1$. From equation

4.7 we can see that for subsonic flow, friction causes an increase in Mach number and therefore drives the flow towards M = 1. Friction in supersonic flow causes a decrease in M which similarly tends towards M = 1. This means that friction alone cannot cause a subsonic to supersonic transition and in a case of a sufficiently long passage, with subsonic inlet M, the flow will choke at the exit. The maximum passage length which can support a given passage mass flow and inlet M can be found by integration of 4.7. Passages longer than the maximum supported length will choke at the exit, causing a reduction in mass flow. Supersonic inlet flow in a sufficiently long passage will result in a normal shock within the passage, reducing the flow to subsonic speeds. The Mach number, post-shock, will then increase again along the passage due to friction and in a long enough passage will also choke at the exit.

Rayleigh flow refers to flow through a frictionless, constant-area passage with heat transfer. In this case only the middle term in the rightmost bracket of 4.7 remains. Heating of the flow along the passage, $dT_0 > 0$, causes an increase in entropy and therefore both subsonic and supersonic flows tend towards M = 1. Similarly to Fanno flow, this means that heating alone cannot transition flow from subsonic to supersonic speeds and so subsonic flows will, given sufficient heat input, choke at the exit. Heating of supersonic flow will eventually result in a normal shock along the passage, transition to subsonic speed, and subsequent increase in M along the remaining length of the duct.

Fanno and Raleigh flow approximations are useful in a wide range of applications as when only one term exists in 4.7 the expression can be exactly integrated. When two or more of friction, heat transfer, and area change occur, the differential equation must instead be numerically integrated.

4.3.2 Low-order model solver

A *MATLAB* script was written to model flow down one convective passage, split into n_c discrete cells as shown in Figure 4.1. The effects of both friction and heat transfer, from the mainstream flow into the convective passage flow, are modelled by using differential equations of the form in 4.7 to calculate changes in flow properties across each cell along the passage in a forward Euler

method, in space. Boundary conditions for the flow and geometric parameters for the passage are specified as inputs, summarised in Table 4.1.

Table 4.1 Specified boundary conditions and geometric values for numerical convective passage solver.

Parameter	Symbol	Value	Unit
Supply total pressure	$p_{0,\text{inlet}}$	275000	Pa
Exit static pressure	p_{exit}	170000	Pa
Coolant inlet total temperature	$T_{0c,\text{inlet}}$	291	K
Mainstream temperature	T_{m}	340	K
Mainstream surface heat transfer coefficient	h_1	1200	W m^{-2}K^{-1}
Passage diameter	D_{h}	0.4–4.0	mm
Passage spacing	w	0.4–2.0	mm
Passage roughness	R_{a}	10	μm
Passage length	L	50	mm
Platform total width	W	40	mm
Platform thickness	t_{w}	2.5	mm
Platform thermal conductivity	k_{w}	6	W m^{-1}K^{-1}
Air thermal conductivity	k_{f}	0.026	W m^{-1}K^{-1}
Air specific heat capacity	c_p	1004	J kg^{-1}K^{-1}
Air Prandtl number	Pr	0.71	–
Number of discrete passage cells	n_{c}	10^3–10^6	–
Inlet resistance coefficient	K_{L}	0.5	–

Solution of the passage begins at the inlet with an initial guess of the inlet Mach number, taken as a low subsonic value. From this, and the specified supply pressure and temperature, the inlet velocity and density are computed and entry pressure loss calculated using equation 4.8 assuming an inlet resistance coefficient, K_L, of 0.5 for a flush, sharp-edged entrance [42]. Each cell along the duct is then solved sequentially. Flow properties into each cell ($p, p_0, T, T_0, \rho, s, v$, M, Re) are computed from values in the cell immediately upstream and the calculated changes in properties across that cell. In order to use equation 4.7 values of friction factor, f_{d}, and dT_0 must be determined.

$$\Delta p = K_{\text{L}} \frac{\rho v^2}{2} \tag{4.8}$$

The Darcy friction factor is determined from equations 4.9 and 4.10. The Colebrook-White equation (4.9) is an empirical relation expressing Darcy friction factor as a function of Reynolds number and relative roughness for turbulent flow through passages with a Reynolds number

greater than 4000. The relative roughness is the ratio of the average height of surface roughness to the hydraulic diameter of the passage (R_a/D) and is useful as a means of comparing the roughness of passages of differing size. For laminar flow (Re $<$ 2300), a simple function of Reynolds number is used to determine f_d. Between these two thresholds the discontinuity in expressions for laminar and turbulent flows is smoothed as shown in equation 4.10, which assists in solver convergence. The uncertainty in f_d within this region is expected to be large however the majority of passage flow cases examined in the present work had Reynolds number greater than 4000. Cases below this threshold typically occur only for very small diameter passages (less than 0.4 mm) which are impractical to manufacture via DMLS and could not be used.

$$\frac{1}{\sqrt{f_{\text{d, trb}}}} = -2\log\left(\frac{R_a}{3.7D} + \frac{2.51}{\text{Re}\sqrt{f_{\text{d, trb}}}}\right) \tag{4.9}$$

$$f_d = \begin{cases} 64/\text{Re} & \text{for Re} \leq 2300 \\ \left(\frac{64}{\text{Re}}\right)\left(\frac{4000-\text{Re}}{4000-2300}\right) + f_{\text{d, trb}}\left(\frac{\text{Re}-2300}{4000-2300}\right) & \text{for } 2300 < \text{Re} < 4000 \\ f_{\text{d, trb}} & \text{for Re} \geq 4000 \end{cases} \tag{4.10}$$

The change in the total temperature of the flow (dT_0) depends on the mass flow rate and specific heat capacity of the fluid, and the amount of heat transferred into the flow as in equation 4.11. Mass flow rate is constant along the passage and can be directly calculated from flow properties. For the range of fluid temperatures in the present study, the specific heat capacity of air is approximately constant and was taken as 1004 J kg^{-1}K^{-1}. For systems with larger temperature ranges than considered here, the solver could easily be extended to calculate c_p from a polynomial expression using JANAF coefficients.

In this model the periodic domain includes half of the regions of the platform surface between neighbouring parallel passages, as shown in Figure 4.2a by the red dashed area. In each cell, heat is transferred into the platform across the entire mainstream surface cell area, A_{cell}. We assume here that the system has reached a steady state and that all of the heat into the platform across this area is subsequently transferred into the passage fluid. We also assume that small differences in platform thickness, due to the curvature of the passage surface are negligible. We can therefore

consider each cell as an element as shown in Figure 4.1b. The heat transferred into the passage fluid is determined by considering the heat fluxes shown on this element which can be described by expressions 4.12–4.14. Convective heat transfer (q_1, q_2) occurs between the mainstream flow (T_m) and platform surface (T_w1) and from the passage wall (T_w2) to coolant flow (T_c). The convective heat transfer is driven by the temperature difference between the surfaces and fluids, and the local heat transfer coefficients ($h_1,\ h_2$). Heat conducted through the platform (q_{12}) is driven by the temperature difference between the mainstream platform surface and the passage wall surface and depends on the thermal conductivity (k_w) and thickness of the platform (t_w). For a system at steady-state, the heat fluxes in each cell are equal, leading to the expression 4.15. We therefore have three equations (4.12–4.14) and 5 unknown parameters ($q, h_1, h_2, T_{w1}, T_{w2}$).

$$q = (1/A_\mathrm{cell})\dot{m}\, C_p\, dT_0 \tag{4.11}$$

$$q_1 = h_1(T_\mathrm{m} - T_\mathrm{w1}) \tag{4.12}$$

$$q_{12} = (k_\mathrm{w}/t_\mathrm{w})(T_\mathrm{w1} - T_\mathrm{w2}) \tag{4.13}$$

$$q_2 = h_2(T_\mathrm{w2} - T_\mathrm{c}) \tag{4.14}$$

$$q_1 = q_2 = q_{12} = q \tag{4.15}$$

It should be noted that in this model the lower passage wall is considered adiabatic for simplicity. This was deemed acceptable by the author as, for the current system, the temperature gradient between mainstream platform and passage surfaces was expected to be dominant and the thickness of the platform below each convective passage is greater than that above the passages, meaning a greater thermal resistance. The thermal energy transferred into the coolant and convected out of each cell, downstream to the next cell, is therefore expected to be much greater than the heat that would be transferred into the metal of the platform below the passage, and ultimately into the shank and hub seal coolant flows. Lateral conduction of heat in the platform between neighbouring cells is also neglected in this model. For a large number of discrete cells, the temperature gradient between any pair of neighbouring cells is small compared to the driving temperature gradient from the mainstream platform to passage surfaces and therefore the proportion of heat conducted is expected to be small by comparison.

The external heat transfer coefficient (h_1) is estimated as a constant value of 1200 W m^{-2}K^{-1}, which is representative of expected values on the platform surface and is based on CFD simulations of engine components provided by Rolls-Royce. The internal heat transfer coefficient (h_2) is determined using equations 4.16 and 4.17 where, similarly to 4.10, different expressions are used for laminar and turbulent flow cases, and for the transition range between threshold Reynolds numbers. The expression for convection with laminar flow (Re < 2300) and uniform heat flux in circular passages is taken from Bergman et al. [4]. The Gnielinski correlation, equation 4.16, can be used for turbulent flows with $3000 \leq \mathrm{Re} \leq 5 \times 10^5$ through circular rough passages. It is a function of flow properties (fluid thermal conductivity k_f, Reynolds number, and Prandtl number), Darcy friction factor (determined as described above), and hydraulic diameter.

Using the piecewise equation 4.17 to determine h_2 and an estimate of the external heat transfer coefficient h_1, we are left with three unknown parameters and three equations 4.12–4.14. We can therefore solve these three equations to simultaneously determine $q, T_{\mathrm{w}1}$, and $T_{\mathrm{w}2}$. This in turn allows a value of q to be substituted into 4.11 and dT_0 to be determined. We then have values for all the required terms to directly calculate the change in Mach number (and similarly, other flow properties) across the cell using equation 4.7 which gives the inlet properties for the next cell. The process is repeated in each cell along the passage until the end of the duct.

$$h_{trb} = \frac{k_f}{D}\left(\frac{(f_d/8)(Re-1000)\mathrm{Pr}}{1+12.7(f_d/8)^{1/2}\left(\mathrm{Pr}^{2/3}-1\right)}\right) \tag{4.16}$$

$$h_2 = \begin{cases} \dfrac{48}{11}\dfrac{k_f}{D_h} & \text{for } \mathrm{Re} \leq 2300 \\ \left(\dfrac{48}{11}\dfrac{k_f}{D_h}\right)\left(\dfrac{3000-\mathrm{Re}}{3000-2300}\right) + h_{trb}\left(\dfrac{\mathrm{Re}-2300}{3000-2300}\right) & \text{for } 2300 < \mathrm{Re} < 3000 \\ h_{trb} & \text{for } 3000 \leq \mathrm{Re} \leq 5 \times 10^5 \end{cases} \tag{4.17}$$

In the final cell, the passage exit, two checks are performed to validate the guess of inlet Mach number. Firstly, the static pressure computed in the final cell is compared to the specified exit static pressure. If the values are not approximately the same (to within ± 0.1 Pa) then the Mach number is checked to see if it is approximately unity (to within ± 0.001) which would mean the

flow has choked at the exit. If neither of these conditions are met then the inlet Mach number used is not physically valid. The guess of the inlet Mach number is therefore updated, by increasing the value by a small increment (0.001), and the solver re-started from the inlet. This occurs until the above checks are satisfied, indicating the correct inlet Mach number, and therefore valid passage solution, has been determined.

4.3.3 Low order model results

The low-order solver was run for passages of varying diameter and spacing, limits of which are tabulated in Table 4.1. The aim of modelling passages for these ranges of values was to help select sensible values to use in a system implemented in the cascade blade geometry. The lower limits on both these parameters were therefore set by manufacturer recommendations for minimum hole diameter and wall thickness for metal DMLS parts.

Example outputs of flow parameters along a converged passage, with $D_\mathrm{h} = 0.8$ mm and $w = 3$ mm, are plotted in Figure 4.3. The distance along the passage has been normalised so that the passage inlet is indicated by zero, and the exit indicated by unity. Total pressure (shown on 4.3a) along the passage decreases gradually due to friction and heat addition. The inlet total pressure is lower than the supply total pressure, from the front seal, due to entry losses. The static pressure also decreases along the passage until it matches the exit pressure of the rear seal. The inlet Mach number (shown on 4.3c) for this example passage was determined to be 0.32 and increases to 0.48 at the exit. The gradient of the Mach number is higher towards the exit and rapidly increases as the Mach number increases. This is because the change in Mach number is a function of M^2 as shown in equation 4.7.

The total temperature (shown on 4.3b) increase along the passage as heat is transferred from the mainstream into the passage coolant. The coolant static temperature increases along the duct, but as the Mach number begins to increase more rapidly the rise in temperature levels off towards the exit of the duct. In longer passages, in which higher Mach numbers occur, the passage temperature may actually decrease. This is beneficial as it increases the heat transfer, enhancing cooling performance. As the coolant temperature rises, the difference between the mainstream

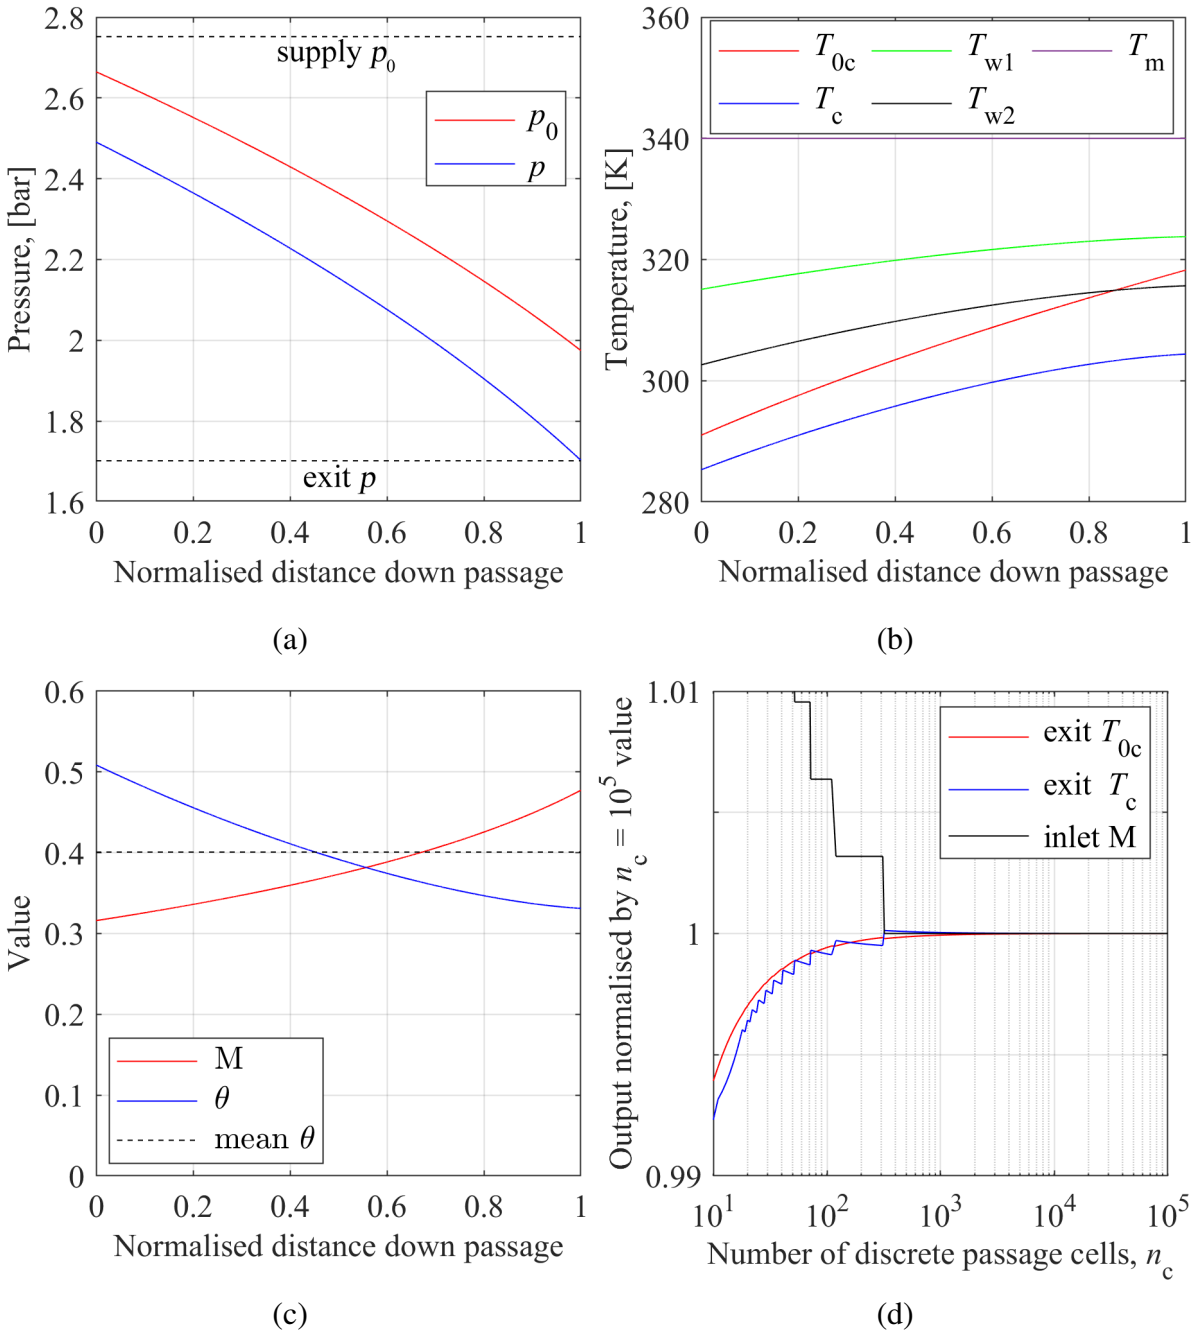


Fig. 4.3 Example results from convective cooling passage low-order model with boundary conditions specified in Table 4.1 and $D_h = 0.8$ mm and $w = 3$ mm. (a) Pressures, (b) temperatures, (c) Mach number and metal effectiveness along passage, (d) stability of output values for increasing numbers of discrete cells.

and coolant decreases (as the mainstream is assumed to be a constant temperature) and therefore both the internal and external platform surfaces have higher temperatures towards the exit of the passage. We can calculate the expected metal effectiveness from the mainstream platform surface temperature, mainstream temperature, and supply coolant temperature, this is plotted 4.3c. Here, we use coolant supply temperature, $T_c(x = 0)$, as we are interested in the performance of the whole passage, rather than local effectiveness from the coolant temperature in each cell (values of which would be different to those plotted.) The metal effectiveness can be seen to decrease along the duct as the coolant temperature rises. The area-average metal effectiveness is indicated by the black dashed line and we use this value below to compare passages of differing diameter and spacing.

As the model solves the passage by discretising the continuous differential equations governing the flow, the accuracy of any given solution depends on the number of discrete cells used to model the passage. For a passage of fixed length, increasing the number of cells also decreases the cell length which makes another of our assumptions—that lateral conduction of heat is low due to low temperature gradients between neighbouring cells—more accurate. The downside of using a large number of cells is increased computation time. The example passage described above was run with values $10^1 < n_c < 10^5$ and indicative values of exit total and static temperature, and inlet Mach number were compared to see how stable the solution was (indicated by how much difference in the values there were for different numbers of cells.) These are plotted in 4.3d where all values have been normalised by the corresponding $n_c = 10^5$ values, which are expected to be the most accurate. It can be seen that at low numbers of discrete cells, less than 10^2, there are differences in the passage solutions of these values of up to 0.8%. For $n_c = 10^3$ the maximum difference is less than 0.01%. For a standard office computer the passage solver, with $n_c = 10^4$, takes only a couple of seconds to complete and has a negligible difference in solution from higher cell counts. For these reasons this value was used when running many comparative passage cases.

The two performance metrics selected to compare convective passage systems were the average metal effectiveness achieved (taken as the area-average along one passage) and the total required coolant mass flow rate. An effective system is one which maximising the achieved metal

effectiveness and minimises the required coolant mass flow rate. Passages of differing diameter and spacing cooled areas of differing size; the periodic domain areas varied between convective passage cases. Comparing the amount of coolant used per passage would therefore not be a fair comparison as the same amount of coolant might achieve a higher metal effectiveness simply due to cooling a smaller area (this does not affect the average metal-effectiveness results as these are already area-averaged). The total mass flow rate required to cool an area representative of a blade platform (Taken as $W \times L = 40 \times 50$ mm) was therefore used. This was calculated by multiplying the single-passage mass flow rate by the area ratio between the platform representative area and periodic passage domain.

Results of area-averaged metal effectiveness and total coolant mass flow rate are plotted in Figure 4.4 for ranges of passage diameter, D_h, and spacing,w. As the passage diameter increases the coolant mass flow through each passage (shown in red on 4.4b) increases. For a given heat input the bulk temperature increase of a larger mass flow rate of coolant is lower, as can be seen from equation 4.11. This means that the temperature of the fluid remains lower along the duct when higher mass flow rates occur and thus the temperature gradient between the coolant and mainstream, which drives heat flux into the coolant, is larger at all points along the duct (than compared to a case with lower $\dot{m}_\mathrm{c}$). The net result of this is that passages with higher mass flow rates pick up more heat and can achieve consistently lower surface metal temperatures, leading to increased area-average metal effectiveness values. At very high coolant mass flow rates, the temperature rise of the coolant along the passage becomes negligible and an approximately constant heat-flux along the length is achieved. For this reason, the gradient in area-averaged metal effectiveness begins to gradually level off above a passage diameter of 1 mm.

As the spacing between adjacent passages increases, each passage is required to cool a larger platform surface area which increases the heat flux into the coolant. This in turn leads to a more rapid rise in fluid temperature and reduced cooling performance further along the duct. The effect of this can be seen in Figure 4.4d which shows an approximately linear decrease in area-averaged metal effectiveness (for the same passage diameter) for increasing passage spacing. The optimal system, from a cooling perspective, would be to have no spacing between passages resulting in one continuous cavity under the platform outer surface. This would however

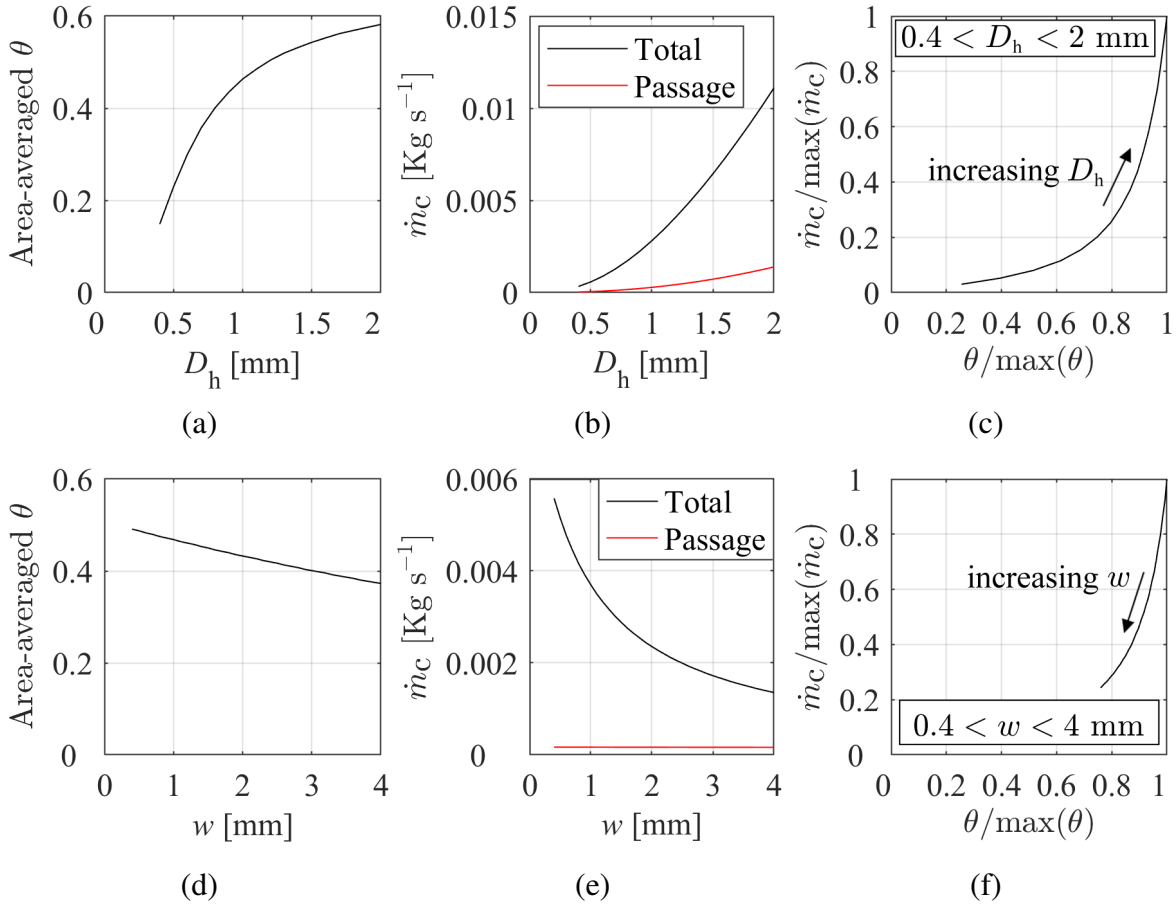


Fig. 4.4 Low-order model results of metal effectiveness and total required mass flow for varying values of passage diamater, D_h and spacing, w. (a), (b), (c) have $0.4 < D_h < 2.0$ mm and $w = 3$ mm. (d), (e), (f) have $0.4 < w < 4.0$ mm and $D_h = 0.8$ mm.

be impractical for mechanical reasons and would require an excessive amount of coolant. As the periodic cooled-area increases (due to larger w), fewer passages are used to cool the same representative platform area and therefore the total mass flow rate of coolant decreases, with a gradient dependent on the ratio D_h/w.

Figures 4.4c and 4.4f show the total mass flow of coolant plotted as a function of the area-averaged metal effectiveness for varying passage diameter and spacing. Values have been normalised by the maximum value of each axis, for the range of values examined. These plots are useful as they allow for the relative trade-off between increasing area-averaged metal effectiveness and minimising coolant mass flow rate to be visualised. Passages with small diameters use low amounts of coolant and achieve relatively low average metal effectiveness values. Increasing passage diameter, moving from left to right along the curve in 4.4c, leads to

an increase in coolant flow area and therefore the amount of coolant used. Increases in metal effectiveness at smaller passage diameters are large compared to the additional coolant used. At larger passage diameters the amount of coolant used increases exponentially however the added metal effectiveness is small.

An increase in coolant mass flow rate of 22% (from 3 to 25% of maximum) leads to a corresponding increase in area-averaged metal effectiveness of 55% (from 25 to 80% of maximum). Past this point, the increase in cooling performance is small compared to the increase in required coolant. The final 20% increase in $\theta/\max(\theta)$ requires a 75% increase in coolant mass flow rate. The turning point in the gradient of this plot corresponds to the range of passage diameters at which the area-averaged metal effectiveness line begins to level-off (between 0.8 and 1.1 mm). A similar trend, requiring a large additional mass flow rate of coolant to gain small increases in area-averaged metal effectiveness, can be seen for varying passage spacing (4.4f).

These plots were used to guide the selection of $D_{\mathrm{h}} = 0.8$ and $w = 3$ for an integrated platform design within the cascade blade geometry. This passage diameter was chosen as it is the value at which the area-averaged metal effectiveness curve begins to level off and is therefore efficient in terms of the amount of coolant mass flow rate used to achieve a reasonable metal effectiveness. The passage spacing selection was somewhat arbitrary, and driven largely by the geometry of the real platform, discussed in detail below. As this value is large compared to the diameter it has a relatively lower coolant mass flow rate requirement and only somewhat reduced area-averaged metal effectiveness compared to systems with smaller spacing.

4.3.4 Limitations of the low-order model

The low order model described in this section was developed as a means to explore the impact of various design parameters and rapidly compare ranges of possible values with some performance estimates. To look in detail at the ranges of parameters explored here with high fidelity techniques such as conjugate CFD of fully meshed, 3D designs would be very prohibitive in terms of setup and computational time. This led to a number of simplifications and reduction of the problem into a flat-plate geometry and relatively simple solution method. The solver could be described

as a 1.5-dimensional model as we model the flow along one passage, and the heat flux from the mainstream, perpendicular to the direction of passage flow. The model is not truly 2D as we neglect lateral conduction effects in the solid of the platform.

We model the passages with a constant mainstream surface heat transfer coefficient and driving gas temperature. In reality both of these vary across the surface of the platform. A possible extension to the model to account for these would be to take surface distributions from CFD and impose them as boundary conditions for the length of the passage. This would allow for more specific informed choices to be made for design parameters and would require different performance metrics to be used. Here, the area-averaged metal effectiveness was used to quantify performance however depending on the specifics of the application it may be more appropriate to look at the maximum and minimum values of metal effectiveness, which would give an indication of the maximum metal temperature (important for component life) and regions which are over-cooled.

This simple extension to using variable surface boundary conditions would not increase the computational time of the model by a significant amount. More complex extensions, towards true 2D or higher behaviour would be much more difficult and time-consuming to implement and solve but would likely still be much faster to iterate than fully 3D CFD.

4.4 Integrated blade design 1: convective passages

The convective passage blade design, shown in Figure 4.5, was designed based on the numerical work described and consists of 14 evenly-spaced passages which originate under the platform within the front seal, and exhaust to the rear seal. The passages have a constant circular cross-sectional area, with diameter 0.8 mm. The passages are printed directly into the platform during manufacture, rather than being implemented at a later stage via electrical discharge machining, or other.

The geometry of the cascade test blade is more complex than the simplified flat-plate used for numerical modelling and therefore there were additional challenges in integrating the convective passages. The primary difficulty to overcome was the presence of the radial cooling passages

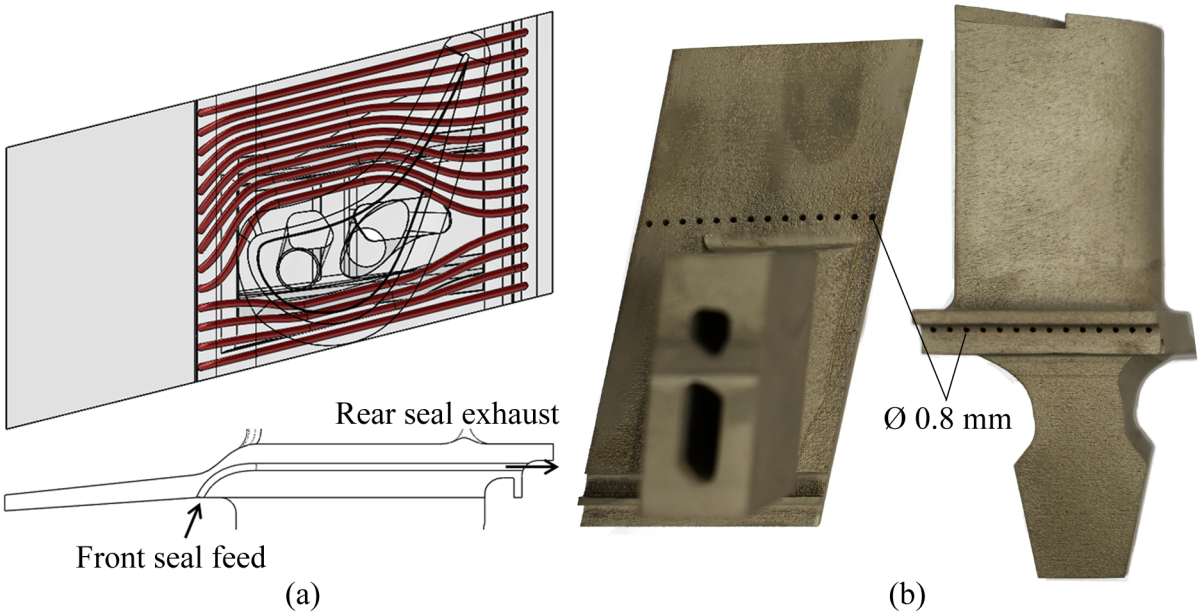


Fig. 4.5 (a) Schematic diagram of convective passage design integrated in cascade blade. (b) Photographs of titanium DMLS test components.

through the blade. The convective passages had to be routed around these in order to not interfere with the blade cooling system. The convective passages all begin evenly spaced with 3 mm between them at the inlet. One group of 4 passages then diverges from the other 10 and routes around the suction side of the blade. All of the passages converge slightly, whilst remaining evenly spaced, to a minimum spacing of 1.8 mm before diverging again. At the exit, all of the passages are evenly spaced with 3 mm between them. The convective passage paths can be seen shaded in red in the schematic diagram Figure 4.5a.

The inlet of the passages had to be below the leading edge section which forms the rotor-side of the front seal in order to guarantee coolant supply and prevent any hot gas ingestion into the passages. As a consequence of this, the passages must turn 90 degrees before running parallel to the flat platform surface. The curvature of this turn was kept large, with a radius greater than 10 times the passage diameter, in order to minimise associated pressure losses.

At the exit, the convective passages exhaust into the downstream rear seal just below the trailing edge overhang. This platform overhang is only 2.5 mm thick and in order to preserve mechanical strength it was decided not to route the passages through this relatively thin region. Having the exit of the passages below the overhang also potentially protects them from interaction with mainstream flow structures, such as the rotor wake and passage vortex, which may otherwise

change the exit pressure and therefore affect the performance of the passages. This positioning may also allow the coolant form the convective passages to be more effectively used to help seal the rear cavity from mainstream ingestion, however sealing performance was not assessed.

4.5 Convective passage experimental capacity

Manufacturing metal test components with intricate, small-diameter passages such as the convective passage blades via DMLS can be challenging. Although 3D printing technology and techniques are becoming increasingly accurate, there are a large number of variables which can affect the final component such as: print layer thickness, powder diameter, build speed, build orientation, part size, and wall thickness. Of particular note is the shrinkage that occurs as the component cools during the printing process. This significantly affects components such as these rotor blades since the parts are solid and the build volume large compared to the convective passages. Shrinkage of the passages would significantly affect their performance, as shown by the previous numerical modelling, and in the worst case would cause partial or complete blockage of one or more passages.

Prior to having the full test blades manufactured a pair of prototype platforms were printed by the manufacturer to assess how accurately the convective passages could be made. All printer settings were recorded so that the same values (and same machine) could be used for the complete blade prints. The prototypes were simplified to reduce print volume (and therefore cost of manufacture) by removing the fir tree and blade leaving just the platform with integrated convective passages. Two prototypes were manufactured with nominal passage diameters of 0.5 and 0.8 mm, these are shown in Figure 4.6a. These prototype prints were used to assess: passage continuity; inlet and exit diameter; and flow capacity at a range of pressure ratios.

Passage continuity was checked for each of the 14 passages on both prototype platforms by blowing air from a high pressure (100 PSI) line through a short length of 1.6 mm diameter scanivalve tubing into the inlet of each passage. In all cases a strong jet of air was observed from the corresponding passage outlet and no unusually low flow jets, which would indicate partial blockage, were observed. During the printing process, cracks or cavities could feasibly

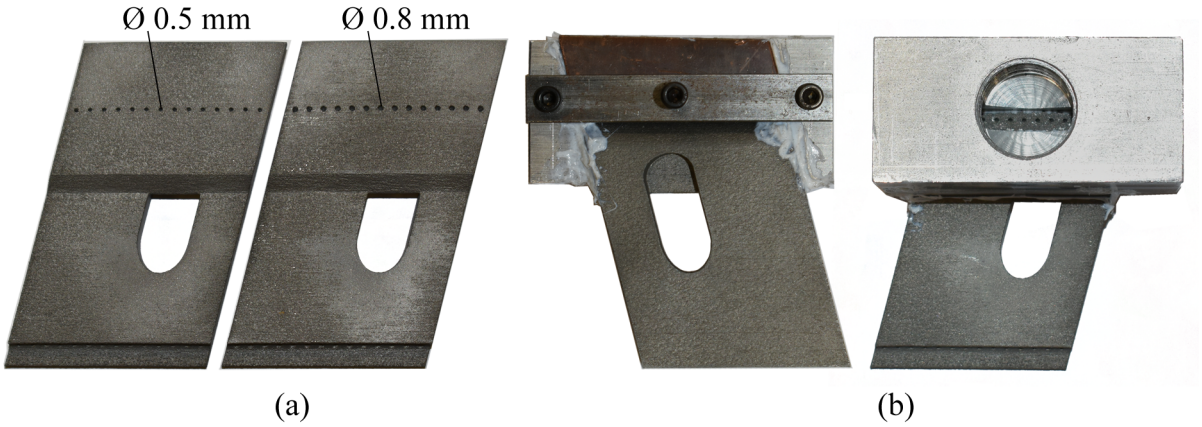


Fig. 4.6 (a) Photographs of simplified prototype platform titanium DMLS test components and (b) component capacity test jig.

form resulting in leakage paths between passages. These could be detected by x-ray imaging the components however due to the low-risk of leakage paths this was deemed prohibitively expensive and time consuming. A simple check was therefore carried out at the same time as the continuity tests by using a fine thread, easily moved by tiny mass flows of air, to try and detect any flow through neighbouring passages. No leakages, which would be indicated by the thread moving when positioned at neighbouring passage exits, were detected via this method.

A large number of aerospace components have previously been manufactured by the laboratory group and from this prior experience it was noted that cooling channels, such as the convective passages, typically shrink by 0.3 mm compared to the CAD geometry. For this reason, the CAD geometry of the two prototypes had passage diameters of 1.1 mm and 0.8 mm which were expected to result in prototype platforms with passage diameters of 0.8 mm (the design value) and 0.5 mm (which would result from printing CAD geometry with the design diameter, not modified to account for shrinkage). The inlet and exit diameters of each passage in the printed prototypes were measured using drillbits of known sizes, and size increments of 0.05 mm. Drill bits of increasing diameter were inserted approximately 2 mm into each opening until a diameter which could no longer fit. The mean diameters were determined as 0.78 ± 0.05 mm and 0.50 ± 0.04 mm, which matched well with expectations.

The flow capacity of the prototype platforms were measured experimentally. A relatively simple jig was manufactured, shown in Figure 4.6b to hold each of the test components and allow

for high pressure air to be blown through the convective passages. The jig was composed of a machined aluminium block with a tapped half-inch BSI hole in the bottom and groove in which the prototype platform could rest. The prototype platform was clamped to the aluminium block by a steel bar, copper shim, and three M4 bolts retaining it in position while the assembly was pressurised. Liquid silicone was used to seal all leakage paths. A short (0.5 m) flexible half-inch BSI hose, and ball-valve fitting, were used to connect the jig to a large 0.875 m^3 volume pressure vessel.

The ball valve between the tank and hose connected to the jig was closed and air supplied at 100 PSI to the pressure vessel until a tank pressure of 2 bar was established. The air supply was then turned off and the tank allowed to settle to ambient temperature. Once the pressure vessel had settled the valve was opened and the tank allowed to vent to atmosphere through the convective passages of the prototype. The mass flow rate through the prototype platform was small and therefore the vessel remained approximately stagnant throughout testing. Pressure and temperature inside the vessel, and atmospheric pressure, were recorded at 100 Hz whilst the vessel drained. From these measurements and the known volume of the vessel the mass of air in the tank at each recorded timestamp was calculated from the ideal gas law. The difference in the mass of air between each time stamp indicated the mass flow rate through the prototype platforms and hence the capacity at a range of pressure ratios.

Each test took between 25 and 45 minutes for the tank to vent to atmospheric pressure and each prototype was tested twice. The recorded data was down-sampled to 1 Hz to smooth out noise in calculated capacity, these are shown as the scattered data points in Figure 4.7 for both of the prototype platforms. Also shown on Figure 4.7, as solid lines, are output capacity curves from the numerical model. The model was run for the same range of pressure ratios as the experiments and for a number of passage diameters close to the measured means of the prototype platforms. The calculated capacity from the numerical model was scaled by 14 to match the number of passages in the prototype platforms. Each output capacity line for the range of passage diameters was checked against the experimental data to find the diameters which most closely matched. These were determined as 0.81 mm and 0.53 mm, well within the measured uncertainty of the prototype platforms.

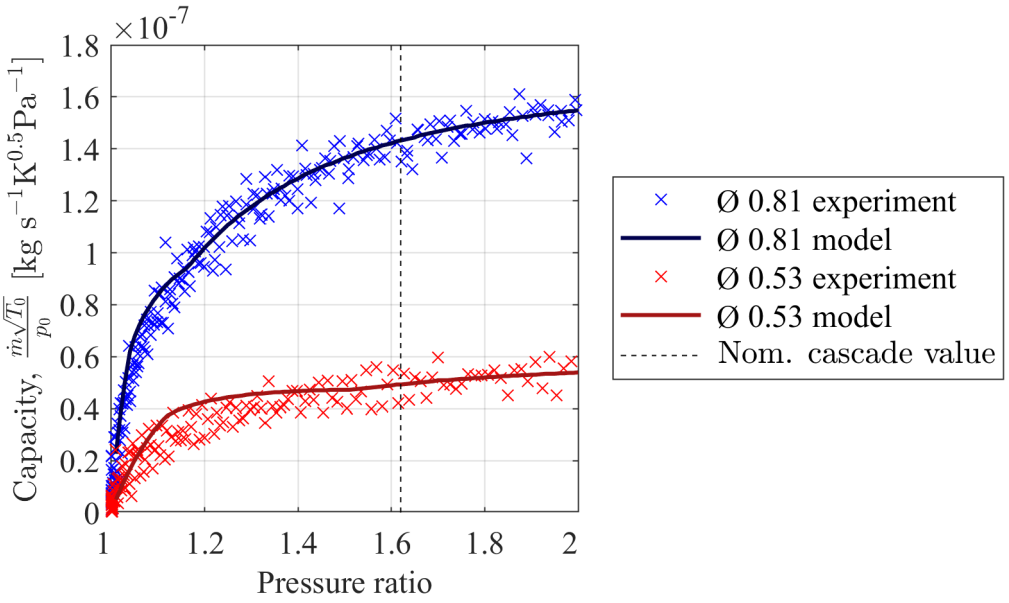


Fig. 4.7 Plots of experimental capacity measurements for two prototype convective passage platforms with nominal $D_h = 0.5, 0.8$ mm and predictions from numerical model.

At higher pressure ratios, greater than 1.4, the numerical model capacity lines match excellently with the experimental data lending confidence to the aerodynamic solution of the model. At lower pressure ratios the model can be seen to deviate from the experimental data, with the most visible example being the red line of the smaller passage diameter prototype platform, between pressure ratios of 1.1 and 1.4. These deviations correspond to transition Reynolds number ranges between laminar and turbulent flow regimes within the passages. It can be seen that the piecewise equations used in this region perform poorly at modelling the flow. It was found that using the Colebrook-White equation for passage flows with Reynolds numbers below the 4000 threshold matched the experimental data more closely (not shown). This suggests an early transition to turbulent flow within the passages at low Reynolds numbers. It is likely also the case that the passages are not uniform in diameter along the entire length due to uneven cooling during manufacture, and the statistical nature of surface roughness.

Results from the prototype platforms showed that the convective passage design could be manufactured. A pair of convective passage blades, for use in cascade tests, were therefore manufactured using the same supplier, machine and settings (Photographs of one of the blades are in Figure 4.5b). These blades were tested for passage continuity and diameter in the same way described above. Measured passage diameters were 0.81 ± 0.04 for each of the blades. The

nominal cascade pressure ratio between front and rear seals is 1.62, indicated by the black dashed line in Figure 4.7. At pressure ratios close to this value the model matches the experimental data well. When conducting experiments with these blades the pressure ratio between front and rear seal was therefore used to determine the mass flow rate of coolant through the blades, as opposed to leaking onto the platform as purge, by using the capacity curve generated by the numerical model.

4.6 Impinging jet array concept overview

Impinging jets are used extensively for cooling surfaces within modern gas turbine blades and NGVs. Cooling is provided by high-velocity jets of air which are directed onto the target surface. As the jet flow reaches the target surface it stagnates and flows radially outwards from the centre of the impingement region, providing additional convective cooling over a wide area. These jets are usually formed by supplying high-pressure air through small diameter holes, or more complex shaped nozzles, within an internal bounding wall. The rotor platform cooling design considered here, shown in Figure 4.8, utilises coolant from the pressurised shank cavity to supply a number of evenly-spaced impingement jets directed at the rear platform surface. A cavity is integrated within the platform body and jets formed by adding small holes between this and the shank cavity. Jets impinge on the rear of the platform within this cavity and spent coolant then flows downstream along the same cavity to an opening at the trailing edge, exiting into the rear seal.

The available platform area which can be cooled in such a way is limited due to the fir tree attachment loft as shown in Figure 4.9. Only a relatively small area of the rear platform surface is suitable for impingement holes to be located without putting holes within the fir tree loft, which may jeopardise it's structural integrity. More intricate impingement cooling systems exist in gas turbine blades which use a 'double-skin' arrangement where there is an additional internal wall separating two internal cavities through which coolant flows. Coolant is supplied to the inner-most cavity and holes in the dividing wall form impinging jets. For this platform cooling application the advantage of this would be the ability to target the majority of the platform

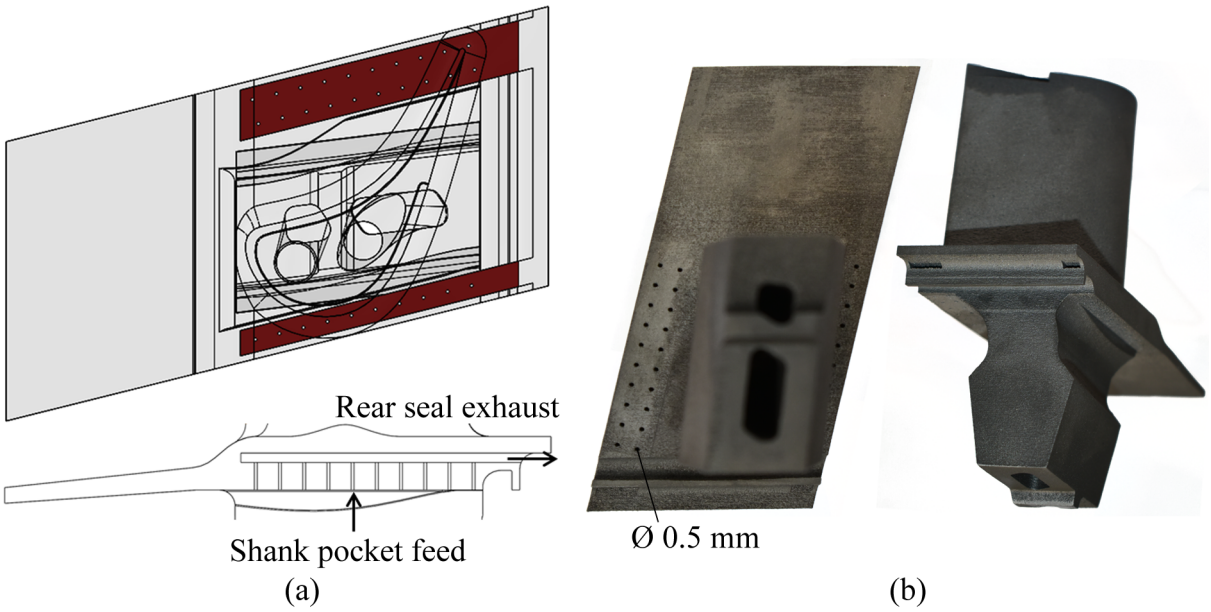


Fig. 4.8 Schematic diagram of impingement array design integrated in cascade blade (a). Photographs of titanium DMLS test components (b).

area with impingement cooling as the inner supply cavity could run through the entire platform. Whilst this kind of complex system is promising it was decided to focus initial design work on a simpler system which could more easily be modelled, manufactured, and tested.

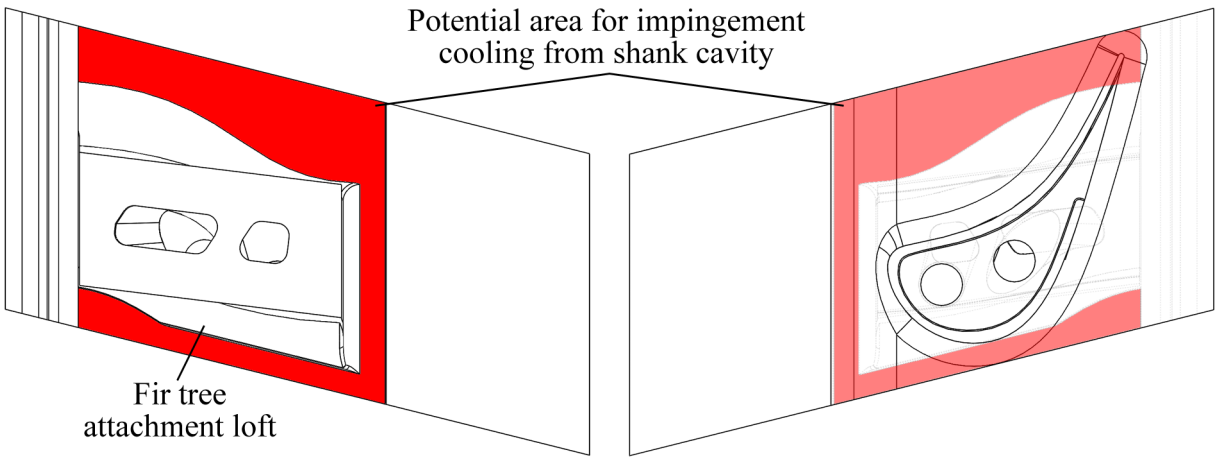


Fig. 4.9 Schematic of cascade test blade geometry with platform areas suitable for impingement cooling, from the shank cavity, indicated by shaded red regions.

Factors affecting the cooling performance of an impinging jet array, with a cavity acting as a duct through the platform, are numerous and consist of geometric and flow parameters which interact in complex ways. Jet diameter, velocity, and mass flow rate are determined by hole size, shape, and pressure ratio and affect the direct cooling performance and area cooled

by each jet. Jet spacing and arrangement (inline or staggered lines of holes), and duct height affects how adjacent jets interact with one another which can either enhance or deteriorate cooling performance. As the system is bounded, with coolant flow down the duct to a common exhaust, jets at downstream duct positions interact with increasing mass flow rates of crossflow. The extent that jet-crossflow interaction affects cooling is discussed in more detail below and depends on the relative mass flows of each stream, largely reliant on the total number of jets. Despite the simplicity of the cooling design, the sheer number of parameters which can be varied and their interaction means that the design process and modelling to estimate performance is more complex. The duct flow as a whole must be considered in order to estimate the cooling performance of each jet, and hence the overall system.

4.7 Modelling jet impingement arrays

The flow associated with a single impinging jet consists of a number of distinct regions, described in much greater depth by Zuckerman and Lior [90] but which can be summarised as: flow within the forming hole or nozzle structure, for round, tube-shaped holes this flow is identical to flow through pipes; ejection as a free jet into the downstream duct, which expands and may entrain additional surrounding fluid; a stagnation region as flow decelerates towards the target surface and begins to turn; flow along the target surface, radially outwards from the stagnation region. When multiple jets are present, and are close enough to directly interact, then a fountain region may form between them. This region is caused by flows parallel to the target surface, flowing in opposite directions, interacting to form a second stagnation region and causing the flow to turn away from the target surface. Flow within the fountain region then mixes out with surrounding fluid and forms a crossflow. These interactions are very difficult to model and require high-resolution CFD to accurate simulate. When considering only the expected cooling performance of an impinging jet system, fully modelling the aerodynamics is unnecessary as the impact of each of these features can be replaced by a single empirical relation.

A numerical model was written in *MATLAB* to solve the aerodynamic flow field and estimate the cooling performance down a duct with a number of impinging jets and subsequent crossflow.

The purpose of this model was to understand how various design parameters affect the cooling performance and hence use this insight to integrate an impingement cooling system in the cascade blade platform geometry. The system was simplified to a flat plate of fixed length, width, and thickness ($50 \times 12 \times 2.5$ mm), representative of the available platform area which could be cooled with impinging jets from the shank cavity. Within the flat plate, the cavity containing impingement jets and crossflow is split into equal-width ducts each consisting of a single row of jets. The number of ducts is variable and depends on specified geometric properties; the maximum number of ducts that can fit within the overall flat plate area is used. Each duct is subsequently split into a discrete number of cells, each containing a single jet. A schematic of the numerical domain and discrete duct cell is given in Figure 4.10.

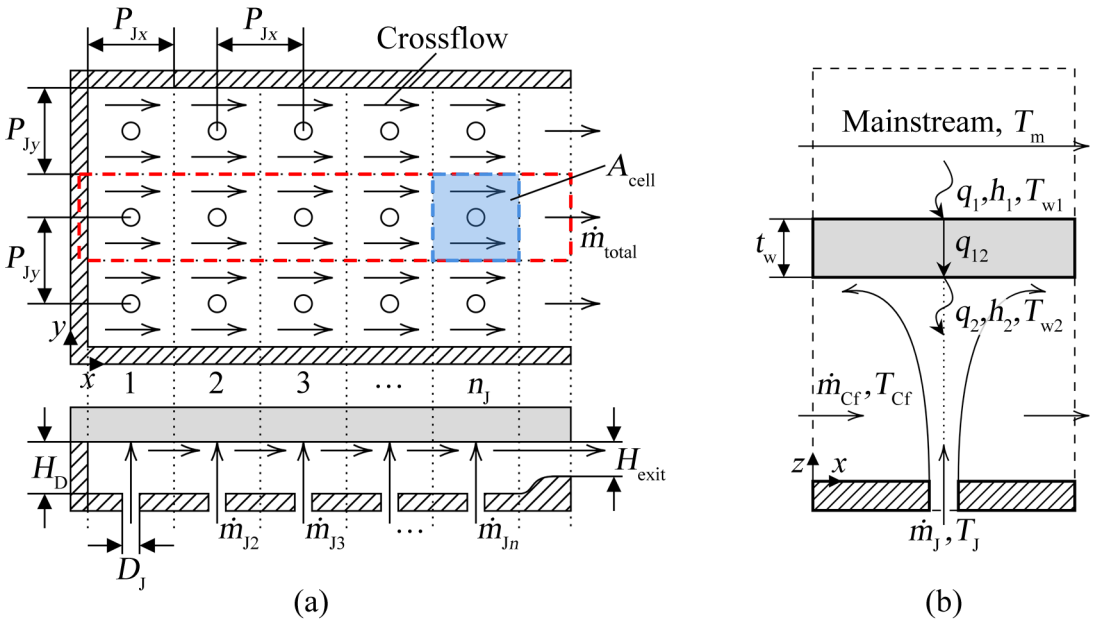


Fig. 4.10 Schematic of impingement array numerical model domain (a) and discrete cell flows and heat fluxes (b)

A single duct, outlined in red in Figure 4.10a, is modelled for a given set of input parameters. Key assumptions in the model are that each jet interacts only within its cell in an area of the platform (A_{cell}) defined by the spacings between adjacent jets (P_{Jx}, P_{Jy}) and that spent coolant from each jet flows only down its corresponding duct (ie no coolant flow in the y direction). Whilst the jets are assumed to interact only within their corresponding cell, the impact on cooling performance due to interaction with surrounding jets is modelled by using an empirical

correlation. This means that the jet is assumed to cool the entire cell area by some average, uniform amount and means that the complicated jet flow does not need to be modelled. These assumptions also mean that each duct can be considered periodic, and there is no distinction between ducts bounded by other ducts, or by sidewalls. We can therefore use the solution from one duct to estimate the cooling of the entire surface with multiple ducts, significantly reducing the computation time required.

Such a system has a large number of design parameters which influence the overall performance and which have a high degree of interaction between one-another. The purpose of this numerical modelling was to explore the impact of key parameters on the overall system performance so that a semi-optimised choice could be made for a design integrated in the cascade blade. This required running the model a large number of times with different sets of input parameters. The number of cases, defined by unique sets of input parameter values, to be run scales dramatically with both the number of parameters and the granularity of the explored ranges of values. The model, described below, was not particularly computationally efficient with a single case taking up to 30 s to solve and process. Depending on the number of parameters varied, and the increments of values over the tested ranges, the number of cases could be many hundreds of thousands to millions. This requires excessively long runtimes even with the solver running in parallel to process multiple cases simultaneously. To reduce the number of parameters and therefore the number of cases, the impingement array was limited to use only square, inline grids such that $P_{\mathrm{J}x} = P_{\mathrm{J}y} = P_{\mathrm{J}}$. Provision was made within the script to allow for an expansion (or contraction) at the duct exit however this was also constrained to reduce the number of cases. Each duct therefore has a constant flow area, including the exit aperture and $H_{\mathrm{D}} = H_{\mathrm{exit}}$. A change in flow area at the exit was found in preliminary testing to have very little effect on the overall system performances when compared to other parameters.

Four parameters which were varied were: the spacing between jets (jet pitch), P_{J}; the height of the duct, H_{D}; the number of jets along the duct, n_{J}; and the total coolant mass flow rate, $\dot{m}_{\mathrm{c}}$. The total coolant mass flow rate is the amount of coolant required to cool the entire available platform area (to be cooled with impingement). The total mass flow rate is split equally between

the total number of ducts. Ranges of values and boundary conditions for the solver are tabulated in Table 4.2.

Table 4.2 Specified boundary conditions and geometric values for numerical impingement array solver.

Parameter	Symbol	Value	Unit
Coolant supply total pressure	$p_{0,\text{inlet}}$	275000	Pa
Coolant supply total temperature	$T_{0c,\text{inlet}}$	291	K
Coolant total mass flow rate	$\dot{m}_c$	(0.45–4.5) $\times 10^{-3}$	kg s^{-1}
Exit static pressure	p_{exit}	170000	Pa
Mainstream temperature	T_m	340	K
Mainstream surface heat transfer coefficient	h_1	1200	W m^{-2}K^{-1}
Jet hole diameter	D_J	0.4–2	mm
Number of jets	n_J	4–15	–
Duct height ratio	H_D/D_J	1–3	–
Jet pitch ratio	P_J/D_J	5–8	–
Platform length	L	50	mm
Platform width available to cool with jets	W_J	12	mm
Platform thickness	t_w	2.5	mm
Platform thermal conductivity	k_w	6	W m^{-1}K^{-1}
Air thermal conductivity	k_f	0.026	W m^{-1}K^{-1}
Air specific heat capacity	c_p	1004	J kg^{-1}K^{-1}
Air Prandtl number	Pr	0.71	–

4.7.1 Low-order model solver

In this section, the methods used in the model to calculate the impingement duct flow and thermal performance are described. The duct flow problem is difficult to initialise and solve due to the nature of the system having multiple mass flow inlets and heat transfer into the duct. Several iterative loops are required as the system cannot be directly determined even when a large number of simplifying assumptions are made. First, the duct aerodynamics are iteratively solved to find pressure, velocity, and mass flow distributions along the duct length. Cooling performance of each jet and heat transfer into the duct flow is then estimated, which changes the aerodynamic solution. The overall duct solution is iterated until the thermal and aerodynamic solutions are stable and all flow properties along the duct are determined. Boundary conditions are set with fixed exit static pressure, and a common supply of coolant at a fixed total pressure

and temperature for all jets. Each jet is fed from the same coolant supply and the supply plenum, representative of the shank cavity, is assumed to be large compared to the jet hole diameters, and stagnant. The total mass flow rate of the system is also specified, and is split equally between ducts. The duct mass flow rate increases along the length as more jet flow is introduced and jets are not constrained to each have the same mass flow rate.

Flow properties (pressure, temperature, mass flow) vary along the duct. In order to solve the duct flow the introduction of each jet into the duct and subsequent interactions must be simplified. Each duct cell, drawn around one of the duct jets as shown in Figure 4.10b, is considered to have constant flow properties and acts as a black box for changes in properties due to the introduction of the jet. The jet flow is driven by the pressure difference between the supply, and the local static pressure within the duct cell. The jet Mach number and mass flow rate can be determined from the pressure ratio using the isentropic flow equations 4.18 and 4.19. We assume that the entire dynamic head of the jet is lost during impingement and subsequent mixing of the jet mass into the duct crossflow. We also assume that the jet fully mixes with the duct crossflow within the cell so that flow travelling down the duct, towards the exit, is uniform across the duct cross-sectional area.

$$\frac{p}{p_0} = \left[1 + \frac{\gamma - 1}{2}\mathrm{M}^2\right]^{\left(-\frac{\gamma}{\gamma-1}\right)} \tag{4.18}$$

$$\dot{m} = A\frac{p_0}{\sqrt{T_0}}\sqrt{\frac{\gamma}{R}}\mathrm{M}\left[1 + \frac{\gamma - 1}{2}\mathrm{M}^2\right]^{\left(-\frac{\gamma-1}{2(\gamma-1)}\right)} \tag{4.19}$$

The cooling performance of each jet, within its cell, is quantified as a heat transfer coefficient (h_2) calculated from an experimental correlation. This heat transfer coefficient is used, along with equations 4.11 – 4.15 which model the heat flux from the mainstream flow through the platform surface and into the duct flow, to determine the duct wall temperature ($T_{\mathrm{w}2}$), mainstream platform surface temperature ($T_{\mathrm{w}1}$) and amount of heat transferred into the duct flow (q_2). In reality, the heat transfer coefficient distribution varies across the impingement area, with a peak

generally at the stagnation point in the centre of the jet impingement. Here however we assume a fixed average heat transfer coefficient and coolant temperature across the entire cell surface.

A large number of experimental correlations, which detail area-averaged Nusselt numbers (from which heat transfer coefficient can be calculated) of various impinging jet configurations and for varying jet properties are available within open literature. Of the correlations reviewed by Zuckerman and Lior [90], the most relevant to this problem of an array of round jet nozzles with crossflow constrained to exit in one direction, is that of Florschuetz et al. [23][24]. They present correlations for area-averaged Nusselt number for impingement arrays defined in terms of a number of characteristic geometric ratios (duct height-to-jet diameter, jet pitch-to-diameter) and flow properties including the crossflow-to-jet velocity ratio. The correlations used in this model are given in equations 4.21–4.25. The characteristic length used to calculate heat transfer coefficent from the reported Nusselt number is the length of each duct cell, P_J.

$$h_2 = \frac{\mathrm{Nu}k_f}{P_\mathrm{J}} \tag{4.20}$$

$$\mathrm{Nu} = \mathrm{A_{Fl}Re^{m_{Fl}}Pr^{1/3}}\left(1 - \mathrm{B_{Fl}}\left[\frac{H_\mathrm{D}}{D_\mathrm{J}}\frac{v_\mathrm{cross}}{v_\mathrm{J}}\right]^{\mathrm{n_{Fl}}}\right) \tag{4.21}$$

where

$$\mathrm{A_{Fl}} = 1.180\left(\frac{P_{\mathrm{J}x}}{D_\mathrm{J}}\right)^{-0.944}\left(\frac{P_{\mathrm{J}y}}{D_\mathrm{J}}\right)^{-0.642}\left(\frac{H_\mathrm{D}}{D_\mathrm{J}}\right)^{0.169} \tag{4.22}$$

$$\mathrm{B_{Fl}} = 0.437\left(\frac{P_{\mathrm{J}x}}{D_\mathrm{J}}\right)^{-0.095}\left(\frac{P_{\mathrm{J}y}}{D_\mathrm{J}}\right)^{-0.219}\left(\frac{H_\mathrm{D}}{D_\mathrm{J}}\right)^{0.275} \tag{4.23}$$

$$\mathrm{m_{Fl}} = 0.612\left(\frac{P_{\mathrm{J}x}}{D_\mathrm{J}}\right)^{0.059}\left(\frac{P_{\mathrm{J}y}}{D_\mathrm{J}}\right)^{0.032}\left(\frac{H_\mathrm{D}}{D_\mathrm{J}}\right)^{-0.022} \tag{4.24}$$

$$\mathrm{n_{Fl}} = 0.092\left(\frac{P_{\mathrm{J}x}}{D_\mathrm{J}}\right)^{-0.005}\left(\frac{P_{\mathrm{J}y}}{D_\mathrm{J}}\right)^{0.599}\left(\frac{H_\mathrm{D}}{D_\mathrm{J}}\right)^{1.04} \tag{4.25}$$

Input parameters to the model were specified in the ratio forms used in this correlation and over the range of values the correlation is valid for. As an example, the height of the duct would be specified by $H_\mathrm{D}/D_\mathrm{J} = 2$ and the actual height of the duct in millimetres would be computed during the solving procedure from the jet diameter. The correlation is valid for: jet Re from 2500

to 70,000; $v_{\text{cross}}/v_{\text{jet}}$ from 0.0 to 0.8; $H_{\text{D}}/D_{\text{J}}$ from 1 to 3; $P_{\text{Jx}}/D_{\text{J}}$ from 5 to 15; $P_{\text{Jy}}/D_{\text{J}}$ from 4 to 8. As the impingement array in this solver was limited to square grids, the minimum range of jet spacing, from 5 to 8, was examined.

Jet diameters are not directly specified but are constrained to be the same value for all jets, both for ease of manufacture and for similarity to the correlation experimental setup. The total mass flow rate was specified, and total duct mass flow rate determined from this, but due to the differences in pressure along the duct each jet has a different mass flow rate. Individual jet mass flow rates and the diameter of all jets are determined by iteratively solving the duct. Initial jet diameters (set arbitrarily as 1 mm) and mass flow rates are set and the duct solved to find the pressure distribution. In each subsequent iteration the mass flows of each jet are adjusted and the jet diameters required to supply the mass flow rate calculated from the partially-solved duct pressures. Eventually, a stable duct solution is found where each jet has the same diameter and the mass flow of each jet is known, and physically valid for the local pressure ratio between the duct and supply plenum.

Solution of the duct begins at the exit, where the mass flow rate is always known (as it is specified) and the exit area is calculated from an initial guess of jet diameter and specified $H_{\text{D}}/D_{\text{J}}$ and $P_{\text{J}}/D_{\text{J}}$. Equations 4.18 and 4.19 are used to determine the Mach number and total pressure of the duct flow at the exit. In the first iteration, the duct flow total temperature is assumed equal to the supply total temperature. The mixed flow within the duct is assumed to be incompressible and so the density is calculated from the ideal gas law. The Mach number of the jet closest to the exit is then calculated from the known pressure ratio and equation 4.18. Upstream of the final jet, the mass flow in the duct is simply equal to the mass flow rate of the final jet subtracted from the specified total mass flow. It is therefore possible to calculate the duct crossflow velocity (and hence dynamic pressure) at that point from the continuity equation $\dot{m} = \rho A_{\text{duct}} v_{\text{cross}}$. The pressure drop caused by mixing of the jet is estimated from the momentum difference across the duct cell and hence the total and static pressure upstream of the final jet can be determined. This process is repeated for each jet, moving progressively up the duct, to give distributions of pressure along the duct and the Mach number of each jet.

Each jet is initially set to have an equal mass flow rate of $\dot{m}_c/n_J$. From this, and the solved duct pressures, it is possible to calculate a jet diameter (from 4.19) that would result in the set mass flow rate, which differs from the initial jet diameter specified. The jet diameters calculated in this way will not be equal. Each individual jet diameter is compared to the mean value of all jets to inform the required change to the mass flow rate of that jet. If an individual jet diameter is larger than the mean then the mass flow must be reduced to reduce the jet diameter and bring it closer to the mean on the next iteration. Conversely, if a jet diameter is smaller than the mean, then the mass flow rate must be increased. The total mass flow rate is kept constant and more and more fine adjustments of jet mass flow rates are made, and iterations of the duct flow solved, until all jet diameters match (stopping criteria when each jet diameter is within 0.05%). This process can take a large number of iterations because as the jet diameters change, the duct flow area, jet spacing and more change. For iterations where the jet diameters are not equal, the average diameter is used to calculate duct dimensions from the specified input ratios. During this process, the width of the duct changes and therefore the number of ducts which can fit below the available target platform area may also change. In these cases the duct mass flow in each duct changes, since it is the total overall massflow of the system that is specified, which requires the duct to be reinitialised. This is handled automatically but also leads to increased computational time.

Once a solution for the duct is determined, it is however only one half of the solution procedure as there is heat transfer occurring on the impingement surface. The flow properties of the solved duct are used to estimate the heat transfer by using the correlation 4.21 and equations 4.11 – 4.15. From these, the change in total temperature and other flow properties along the duct can be calculated. The duct is therefore iteratively re-solved until convergence is achieved, defined by changes in exit total temperature between iterations less than 0.05%. The final output of the model is an array of values tabulating all flow properties in each cell down the duct, and estimates of surface temperature for the duct impingement surface and external platform surface, from which the metal effectiveness of the system can be calculated. Example results are given in the next section.

4.7.2 Low-order model results

The low-order model was run for many cases over the valid ranges of design parameters: H_D/D_J and P_J/D_J in increments of 0.25; n_J for integers between 4 and 15; and $\dot{m}_c$ in increments of 4.5^{-5}kg s^{-1}. All combinations of these values were tested, with each case being defined by a unique set of input design parameters. Each case is equivalently the design of a jet impingement array and the purpose of modelling this large number of cases was to find which designs perform best and, by comparing results of each system, how the design parameters influence the performance. In each case the duct is solved as described above and the converged solution, including all jet and duct flow properties, recorded. Examples of key properties in one such case is given in Figure 4.11. This example case shares design parameters with the system later integrated into the cascade blade geometry, discussed below.

There is a reduction in duct pressures (Figure 4.11a) across each jet cell, as can be seen from the step-like pattern, due to the introduction and mixing of the impingement jet with the crossflow. At the exit of the duct (normalised distance = 1) the static pressure matches the specified exit pressure. The crossflow velocity increases along the duct as more mass is added to the flow by each jet leading to an increase in the difference between total and static pressure along the duct. Temperatures along the duct also increase in steps across each duct cell (Figure 4.11b) as the platform surface is cooled, with heat transferred into the duct coolant. The step size in temperature of the duct coolant flow, T_c, reduces along the duct as the duct massflow increases. The largest step occurs at the start of the duct when the mass flow of only a single jet is picking up the heat from the platform. As the coolant picks up heat along the duct, the platform surface temperature also increases due to the difference between coolant and driving mainstream hot gas temperature decreasing.

Metal effectiveness along the duct (Figure 4.11c) is calculated from the external platform wall temperature, T_{w1}, and the specified mainstream and supply coolant temperature. The supply coolant temperature is used, rather than the local temperature of coolant in the duct, as we are interested in the cooling performance of the duct system as a whole. The metal effectiveness reduces along the duct for two reasons: firstly due to the increase in duct flow temperature and

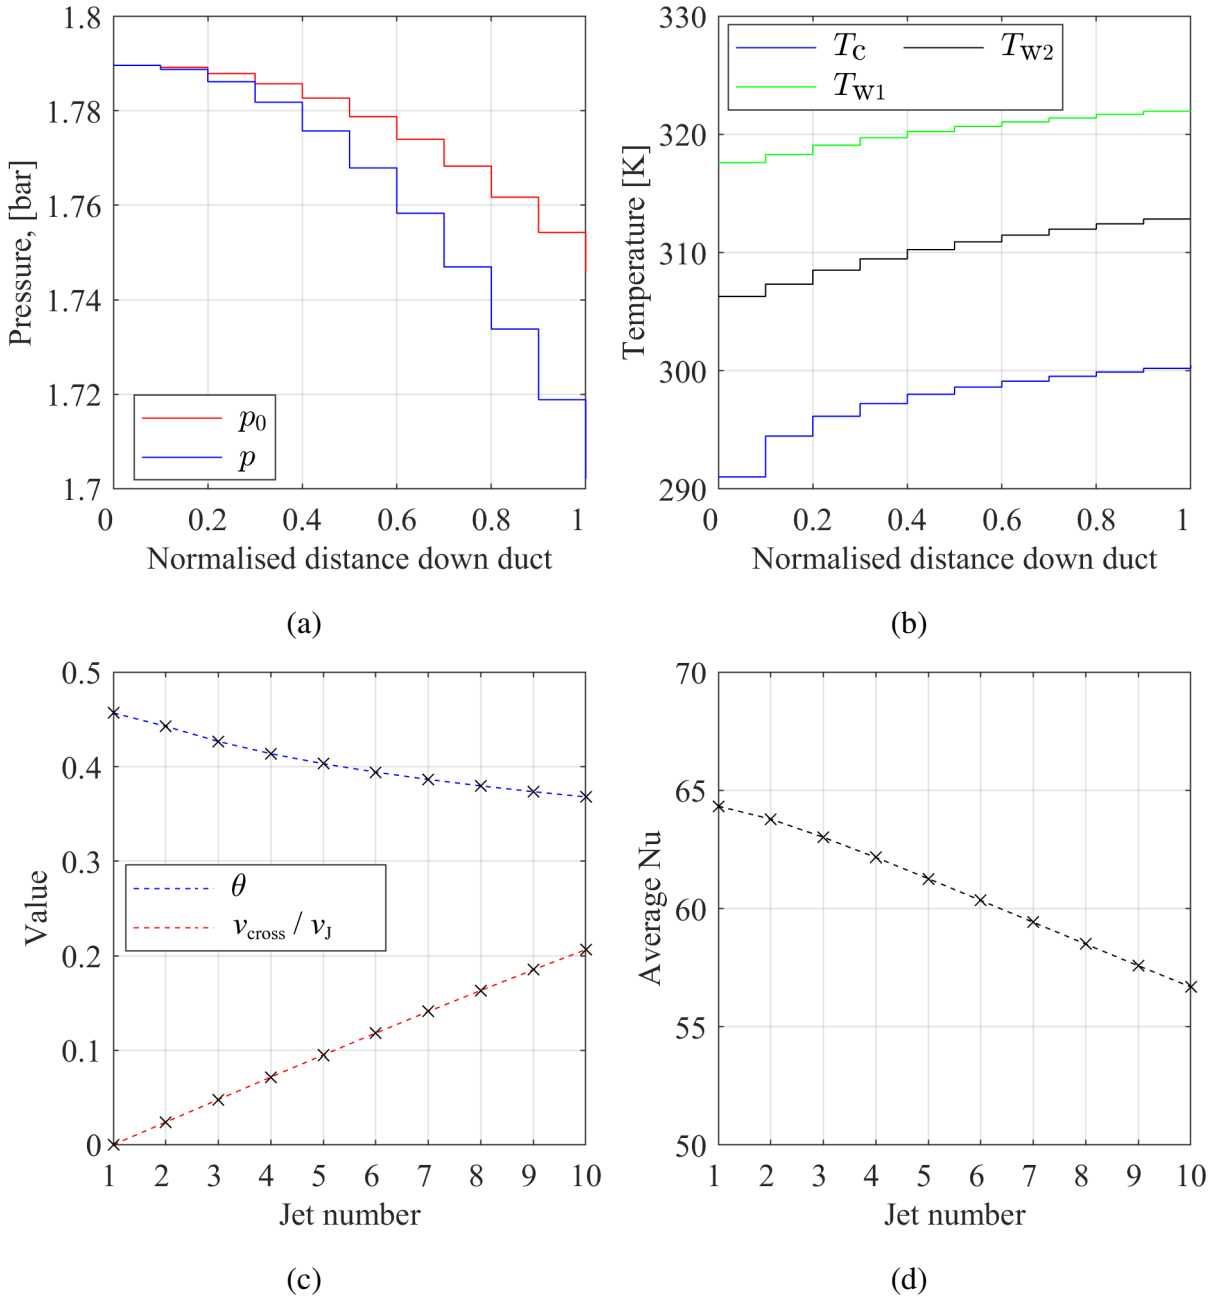


Fig. 4.11 Example results of one case from impingement jet array low-order model with boundary conditions specified in Table 4.2 and $H_D/D_J = 3$, $P_J/D_J = 8$, $n_J = 10$, $\dot{m}_c = 4.5 \times 10^{-3}$ kg s^{-1} . (a) Duct pressures, (b) duct temperatures, (c) metal effectiveness and crossflow-to-jet velocity ratio for each jet cell, (d) average nusselt number for each jet cell.

secondly due to the increase in duct crossflow velocity reducing the cooling performance of later jets. The crossflow-to-jet velocity ratio is also plotted on Figure 4.11c and shows that by the final jet the crossflow velocity is 20% of the jet velocity and significant interaction occurs. This is reflected in the reduction of the average Nusselt number along the duct (4.11d), calculated from the correlation 4.21.

The performance of each test case was assessed by the area-averaged metal effectiveness, which gives an indication of the average cooling performance of each system. There are two distinct values of area-averaged metal effectiveness used for each case: that of an individual duct; and that of the entire available platform area to be cooled. The individual duct value for area-averaged metal effectiveness gives a good indication of how design parameters influence the performance of an ideal duct allowing us, for example, to see what height ratio typically gives higher performance. This value does not however tell us how well the system will perform for our platform application as each duct is a different length and width and covers a different area. Consider an example of two cases where one duct is significantly larger than the other and has a higher area-averaged metal effectiveness. Let us assume that the dimensions of the larger duct is such that only a single duct can fit within the available platform area, and that it covers 60% of the area. Then let us assume that the smaller duct has dimensions such that several ducts can fit within the same available area, covering 90% of the area. We will consider the areas not covered by the ducts to be totally uncooled, with $\theta = 0$. Therefore, depending on the specific values of area-averaged metal effectiveness, the smaller duct may be the better option to achieve a higher average value when the entire available platform area is considered. By taking $\theta = 0$ for any area not covered by some multiple of ducts in each case, we can calculate an expected area-averaged metal effectiveness specific to our platform area, specified in Table 4.2.

The duct-specific and platform-specific area-averaged metal effectiveness values were calculated for all cases and are both examined in the following analysis. To be able to examine trends in the data we perform a dimensionality reduction on the multi-variable space of the total number of cases, which consists of quality parameters (area-averaged θ) spread across combinations of four design parameters. We do this for each of the four design parameters by looping through each specified incremental value and examining all cases (every combination of the other three

parameters) with that value and taking the case which has the highest area-averaged metal effectiveness. As an example, we fix say height ratio of 2.5 and then look at each case with that height ratio, the number of which would be the number of combinations of the increments of the other 3 parameters: pitch ratio, number of jets, total mass flow rate. The maximum area-averaged metal effectiveness value in that subset of cases, and the other parameters of the case are recorded. We then do this for each increment in the range of height ratio to observe the trend in how height ratio affects the maximum achieved cooling performance. We repeat this for each parameter for both the duct-specific and platform-specific area-averaged metal effectiveness. The results of this dimensionality reduction are plotted in Figure 4.12. A similar process can be performed to look at interactions between 2 or more design parameters, plotted as contours or volumes in 3D space for example however this is not demonstrated here. Each of these figures in now discussed.

It is immediately apparent that the highest area-averaged metal effectiveness of single-ducts (shown in black in Figure 4.12) are significantly higher than the highest achieved values for the available platform area. This is due to the ducts with the highest performance having geometric dimensions which allow for only partial cover the available area. As the pitch ratio increases, the duct area-averaged metal effectiveness decreases. This occurs due to the spacing between jets increasing, and therefore each jet is effectively required to cool a larger area. For pitch ratios lower than the range examined (which would not be covered by the correlation for Nu used) we would expect a decrease in area-averaged metal effectiveness as the jet spacing decreased and interaction between adjacent jets increased. The pitch ratio has a large influence on the geometric properties of the duct. There is a peak at approximately $P_J/D_J = 7$ in the platform area-averaged metal effectiveness which corresponds to an optimal tradeoff between the size of the duct (and therefore area coverage) and the reduced per-duct performance.

Both duct and platform values of area-averaged metal effectiveness increase for systems with a larger height ratio (4.12b). As the height ratio is increased, the duct flow area increases and therefore the passage crossflow velocity is lower at all points along the duct. The interaction between the crossflow and jets is therefore diminished and the cooling performance of each jet is increased, leading to higher effectiveness values. The trend appears to be asymptotic and for very large height ratios we would expect the interaction between crossflow and jets to be negligible.

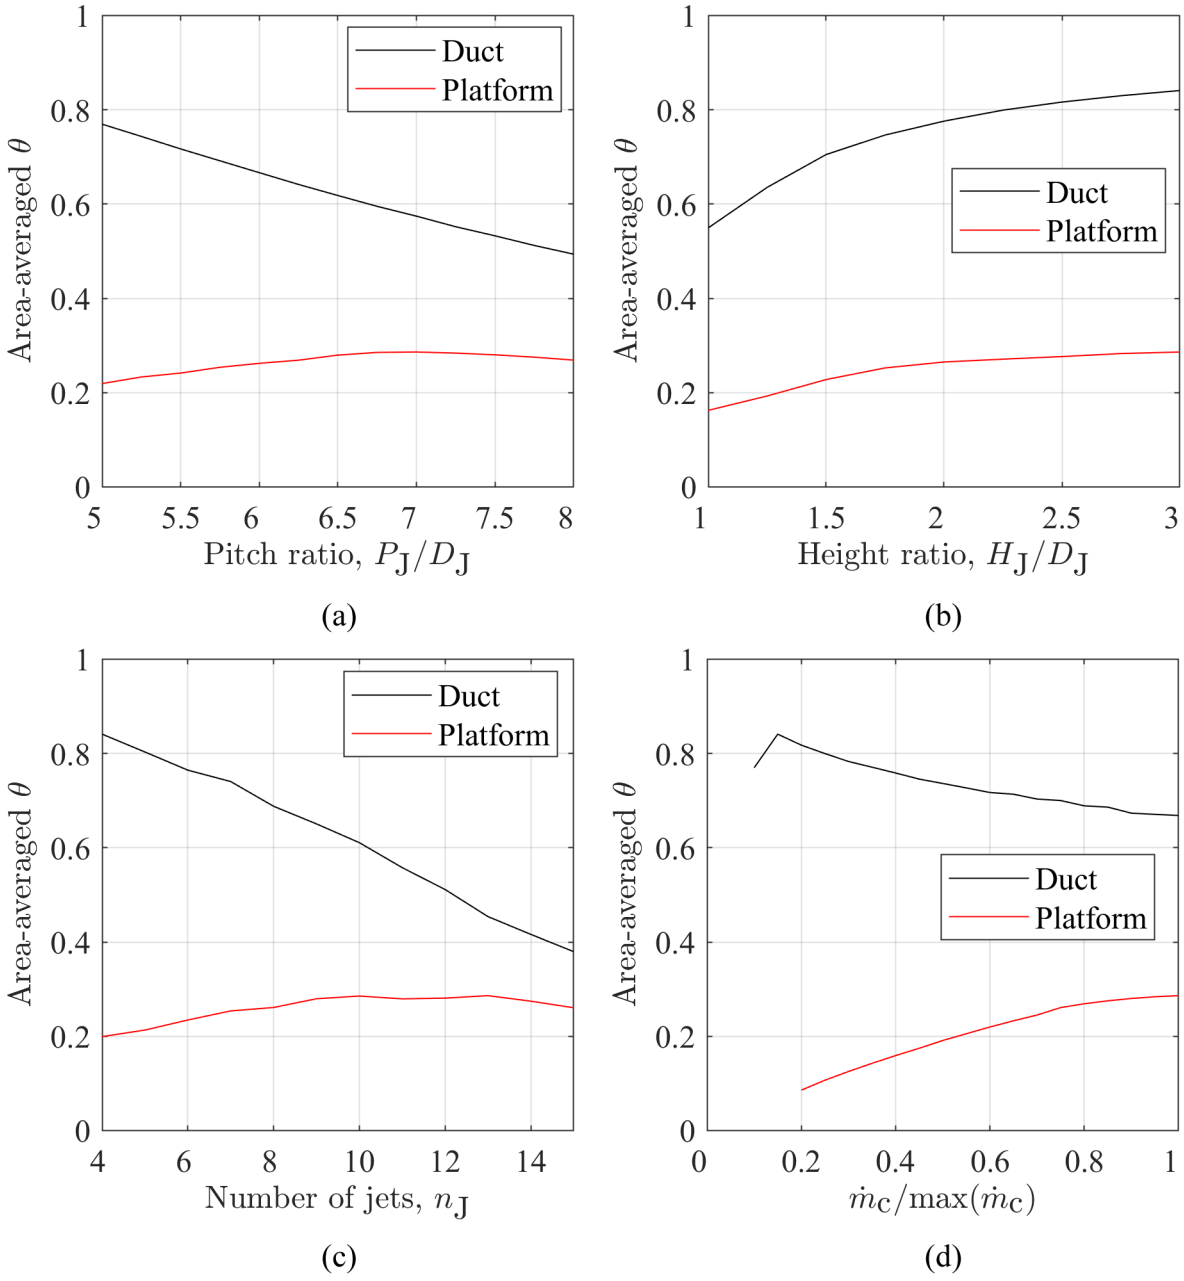


Fig. 4.12 Jet array low-order model solutions of area-averaged effectiveness against input design parameters (a) pitch ratio, (b) height ratio, (c) total number of jets, (d) normalised total mass flow rate. Lines show values of area-averaged effectiveness for cases with the highest achieved area-averaged effectiveness for the indicated design parameter value (x-axis). *Duct* refers to highest values from a single duct solution. *Platform* refers to highest values of the total cooled area with multiple ducts and uncooled regions.

A similar trend to the pitch ratio is observed for the number of jets (4.12c), with the duct performance decreasing with a larger number of jets and a peak occurring (for $10 < n_J < 13$). As the number of jets increases, the mass flow of each jet is reduced and therefore the jet velocities are reduced and the crossflow has a stronger interaction with each jet. This causes the reduction of duct area-averaged metal effectiveness for higher numbers of jets. The length of each duct is determined by the pitch ratio and the number of jets. The total area covered therefore changes with the number of jets. The peak in values for the platform area-averaged metal effectiveness shows cases where the duct length is close to the total platform length and the larger covered area offsets the reduction in individual jet cooling performance caused by higher crossflow interaction.

Somewhat surprisingly, the duct area-averaged metal effectiveness reduces as more total mass flow is used (4.12d). This is explained by the fact that the cases with the highest duct area-averaged metal effectiveness are typically ones which use the minimum number of jets to cool a small area. When the mass flow rate is increased for a given number of jets, the crossflow velocity increases and all but the first jet will have reduced performance. When metal effectiveness values are area-averaged across the platform the trend is reversed and using more mass flow leads to increased cooling performance. This is because when considering the entire platform, cases which cover the most area are preferentially selected. In these cases increasing mass flow increases jet velocity and increases the total amount of heat the duct flow can pickup, reducing the temperature rise along the duct.

The plots of 4.12 allow us to see the effect of the design parameters examined and allow us to pick out values which we would expect to give optimal performance in an system integrated into a real platform geometry. It should be noted that the way in which the dimensionality reduction was conducted means that the highest value of metal effectiveness (separately for the duct and platform lines) refers to the same case—the same set of input design parameters. We can therefore simply pick out this singular optimal case from all of the cases which were run and use those design parameters.

4.8 Integrated blade design 2: impinging jet array

An impingement array was integrated into the platform of the cascade blade geometry, based on the numerical modelling described above. The design which was printed and tested was based on values from an earlier data set from a version of the numerical model which was subsequently found to have errors. Namely the stopping condition for the calculation of the duct temperatures and initial splitting of mass flow rates not always working as intended. The design parameters chosen are thus not optimal values indicated in Figure 4.12 but the blades were manufactured prior to these errors being discovered and fixed. The impingement array design is shown schematically in Figure 4.8 alongside photographs of the titanium 3D-printed components which were manufactured for testing in the MERC facility. A summary of design parameters are also included in Table 4.3.

Table 4.3 Impingement array blade design parameters

Parameter	Symbol	Value	Unit
Jet hole diameter	D_J	0.5	mm
Number of jets	n_J	10	–
Total mass flow rate (cascade)	$\dot{m}_\mathrm{c}$	4.5×10^{-3}	kg s^{-1}
$\dot{m}_\mathrm{c}\sqrt{T_{0\mathrm{c}}}/\dot{m}_\mathrm{m}\sqrt{T_{0\mathrm{m}}}$	–	1.1	%
Duct height ratio	$H_\mathrm{D}/D_\mathrm{J}$	3	–
Duct height	H_D	1.5	mm
Spanwise pitch ratio	$P_{\mathrm{J}x}/D_\mathrm{J}$	8.0	–
Spanwise pitch	$P_{\mathrm{J}x}$	4.0	mm
Streamwise pitch ratio	$P_{\mathrm{J}y}/D_\mathrm{J}$	8.0	–
Streamwise pitch	$P_{\mathrm{J}y}$	4.0	mm

Parameters were selected by first choosing a sensible total mass flow rate for the cooling system to use. The previous results showed that the overall area-averaged metal effectiveness increased with total coolant mass flow rate however, as previously discussed, this comes at the cost of lower cycle efficiency. The total mass flow rate was therefore, somewhat arbitrarily, chosen to be 4.5×10^{-3} kg s^{-1}, or $\dot{m}_\mathrm{c}\sqrt{T_{0\mathrm{c}}}/\dot{m}_\mathrm{m}\sqrt{T_{0\mathrm{m}}} = 1.1\%$, which was the top end of the range examined (and therefore high expected performance), but also approximately matches with the design point for the amount of coolant supplied to the front seal. This was chosen so

that experiments could be run to easily compare the metal effectiveness achieved using the same amount of coolant, in two different cooling systems: the internal integrated impingement array; and the incidental film cooling from the front seal purge.

The impingement array design consists of two rectangular cross-section ducts within the platform which run parallel along either of the blade slashfaces and which have 10 evenly spaced jet holes of 0.5 mm diameter. Coolant is supplied from the shank cavity. The cooling jets impinge on the platform internal duct surface and the spent coolant then flows down the duct to a common exit. Each duct exhausts into the rear seal from openings of equal area to the ducts, located just below the platform trailing edge overhang.

Each jet hole is located along the flat rearside portions of the blade platform and none interfere or run through the attachment of the fir tree root. The flat area where jet holes could be located is limited, with enough space for ducts with only 1 and 2 parallel lines of impinging jets. The jets of both ducts are arranged in a square grid pattern, with equal spanwise and streamwise pitch ratios of 8.0 resulting in jet spacing of 4 mm. The height ratio of both ducts is 3.0 and so each duct has a constant height of 1.5 mm.

A pair of the impingement array design blades were manufactured via DMLS using the same manufacturers, machine and machine settings as for the original baseline blades and convective passage blades. The jet holes were resized to 0.8 mm in the CAD model sent for printing to account for shrinkage during component cooling. Due to limited testing and preparation time, the printed holes were only crudely measured by inserting small-diameter drillbits of known sizes, and size increments of 0.05 mm until a diameter was found which no longer fit. Each of the holes were found to be 0.5 ± 0.05 mm. No capacity tests were performed on the impingement array blades however during experiments the shank cavity was completely sealed such that the only flow path was through the impinging system. The mass flow rate through the system could therefore be directly measured during each test using the mass flow setting plate in the blade cassette.

Chapter 5

MERC facility aerothermal commissioning and processing techniques

In this chapter, aerothermal commissioning data is presented to demonstrate the operation of the MERC facility and key measurement capabilities. Example data sets of aerodynamic and temperature surveys are presented for the downstream area traverse, obtained using bespoke transonic probes. The IR thermography capture and processing techniques, used throughout this thesis to determine platform surface temperature and effectiveness values, are detailed. Full area surface temperature measurements of two blade platforms were obtained by recording the external surfaces with an IR camera during facility operation. Images are processed to account for blackbody and greybody parameters to yield highly accurate surface temperature distributions. These temperatures are then used to determine distributions of metal effectiveness and adiabatic film effectiveness. The uncertainty in metal effectiveness is quantified (adiabatic film effectiveness uncertainty is covered in Chapter 6). Finally a brief list of suggested improvements to the current facility configuration are given which could improve operation and data quality.

5.1 Typical operating conditions

The operating point of the facility is primarily fixed to the design conditions of the test rotor blades, with changes limited to the mass flow rates of coolant and mainstream-to-coolant temperature ratio. Figure 5.1 shows typical traces of facility pressures, mass flow rate, Mach number, and temperature during a standard 2 minute experiment. Figures 5.1e and 5.1f show typical temperature traces for runs at two different mainstream-to-coolant temperature ratios. These figures are now discussed in turn.

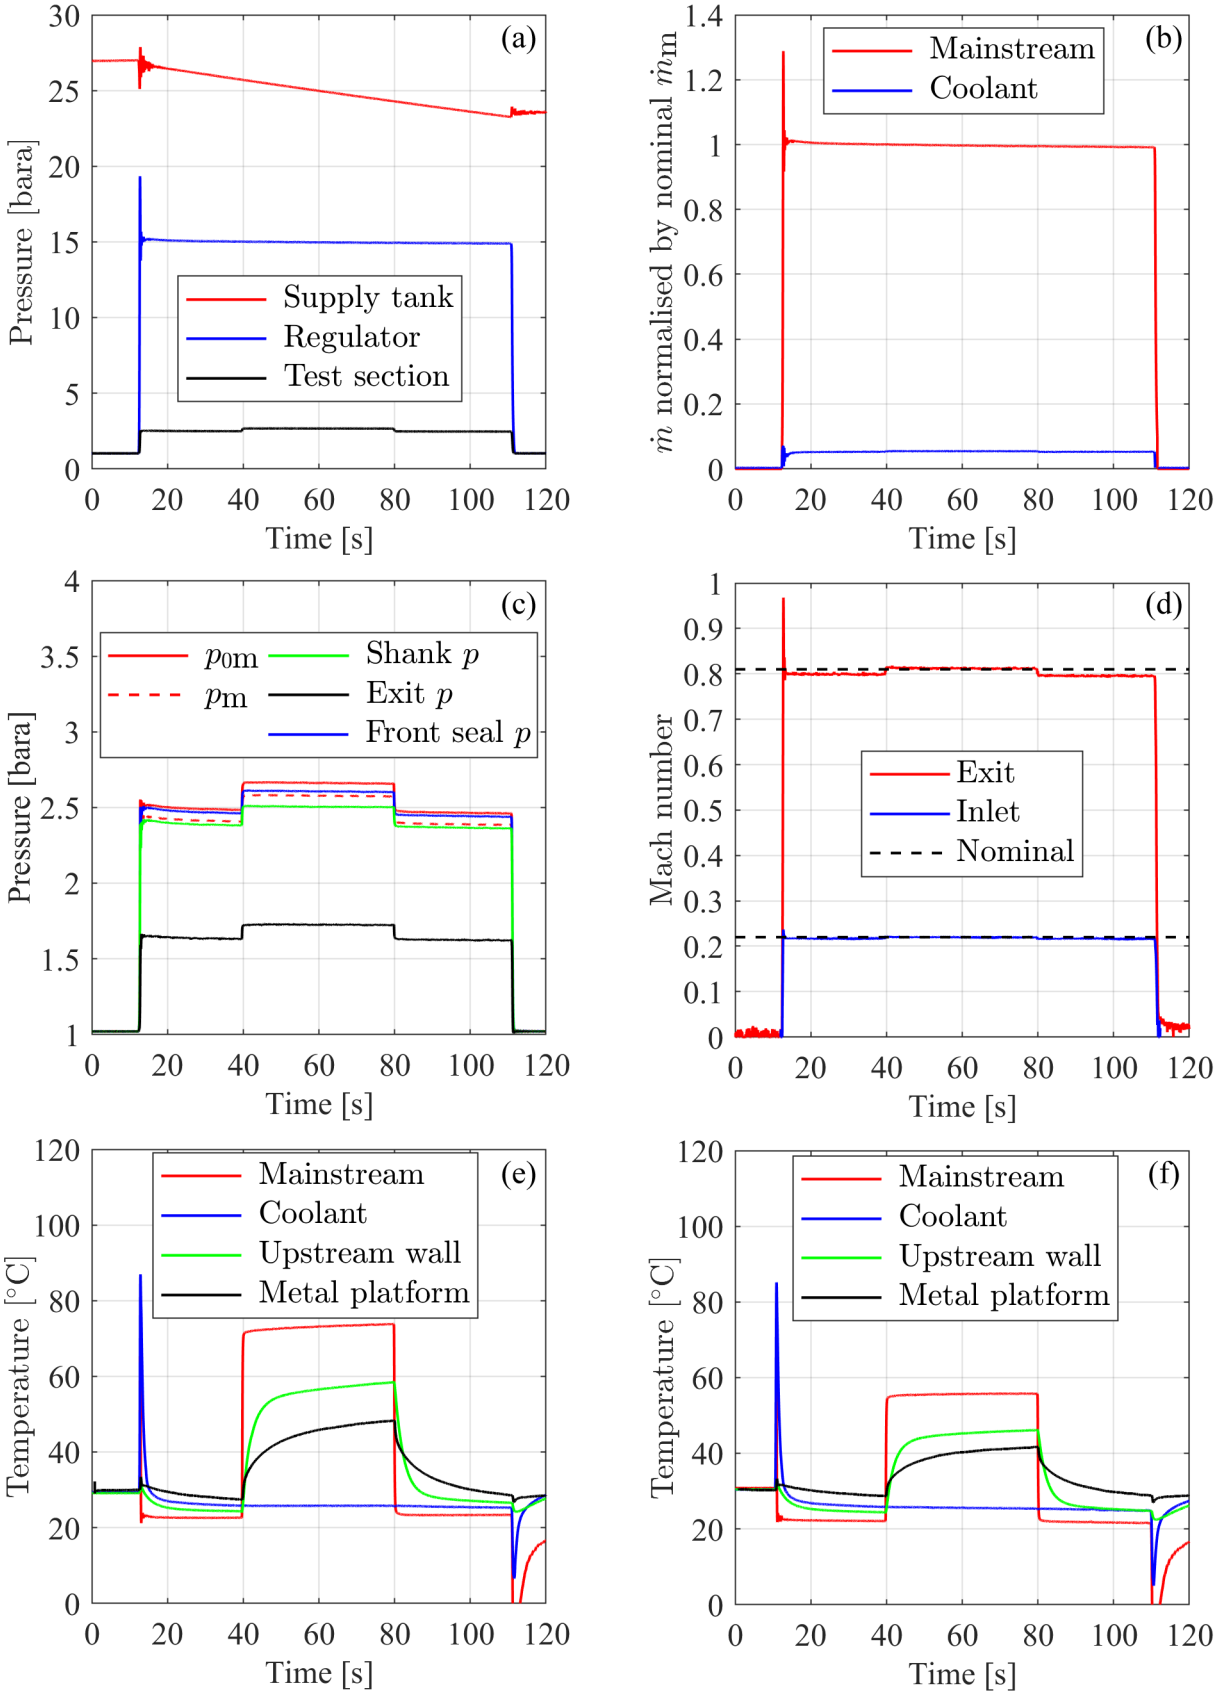


Fig. 5.1 Typical experiment traces of (a) line pressures, (b) mass flow rate (c) test section pressures, (d) Mach number, (e) and (f) temperature for high and low heater power runs respectively.

Figure 5.1a shows pressures of the air supply tank, downstream of the regulator, and the inlet of the test section. The air tanks are typically charged to their maximum pressure of 27.5 bara before tests and drain over the course of a run. Flow into the facility is initiated at 12 s by opening the inlet pneumatic valve (remotely from a control room). There is a small transient overshoot in pressure as the regulator settles to condition, and then steady conditions are maintained throughout the run until 110 s when the pneumatic valve is once again closed. The regulator pressure is set by connecting it to a small-volume pilot tank, which is independently charged to the desired pressure (15 bara) and controlled by a secondary hand-operated regulator in the control room. The main facility regulator maintains this same pressure for the inlet flow and this arrangement allows for minor adjustments of the inlet pressure during a run. As the tank pressure reduces during operation and approaches the regulator set point, the regulator cannot maintain the required downstream pressure which effectively limits the maximum run time of the facility. For this reason experiments are capped at approximately 6 minutes duration, when started with the tanks at maximum pressure.

The regulator effectively sets a mass flow boundary condition by way of a choked venturi nozzle (28.4 mm diameter throat, assumed discharge coefficient of 0.99) upstream of the cascade. The regulator sets a steady 15 bar inlet pressure to the nozzle that has been calibrated to deliver 2.25 kg s^{-1} to the cascade. The coolant air line is split upstream of the choked nozzle, at 15 bara, and a mass flow boundary condition is also set by a choked perforated plate within the blade cassette. Figure 5.1b shows the steady mainstream and coolant mass flow rates, normalised by the nominal mainstream operating point. In this example the total coolant mass flow is approximately 5% of the mainstream (0.11 kg s^{-1}) but this can be varied by changing the choked setting plate, discussed in a later section.

Figure 5.1b shows the pressures of mainstream and coolant within the test section. Mainstream pressures ($p_{0\mathrm{m}}$, p_{m}) are the average of the inlet rakes and monitoring ports. Coolant pressures are measured with wall taps within the blade cassette. The initial pressure overshoot is less significant in the test section and quickly settles to a steady condition. A small step in pressure (from approximately 2.5 bara to 2.75 bara in the mainstream total pressure) occurs between 40 s and 80 s as a result of the heater mesh being turned on and a mainstream tempera-

ture rise. As the mass flow rate into the cascade is fixed, and the capacity, Γ (equation 5.1), is approximately constant, the rise in mainstream temperature results in a corresponding change in pressure. The coolant feeds show a similar increase in pressure during the heated period. They similarly have a mass flow boundary condition and so the seal pressures arise as a consequence of the coolant-to-mainstream pressure ratios necessary to maintain that mass flow rate through the seal. The rise in mainstream pressure therefore causes a corresponding rise in the coolant pressures. Pressures in each seal cavity are measured by single wall taps and therefore pressure distributions within each cavity are not measured. During testing it was observed by the IR camera recordings that coolant is not evenly distributed in the front and rear seals due to the pressure gradient across the whole test section. This problem was solved for the front seal flow by introducing a high-porosity foam within the seal, immediately upstream of the guide vanes. This introduces a pressure drop in the coolant path and forces a more even mass flow distribution within the cavity.

$$\Gamma = \frac{\dot{m}\sqrt{T_0}}{p_0} \tag{5.1}$$

Figure 5.1d shows the inlet and exit Mach numbers which are maintained approximately constant (to within 1% of the nominal values of 0.22 and 0.81) at the blade design conditions during the heated period. The inlet Mach number is measured by total and static probes on monitoring ports at the inlet to the constant-area test section. Exit Mach number is calculated from the upstream total pressure and average of the downstream static wall taps in the rear seal block.

Figures 5.1e and 5.1f show temperature within the test section from relatively high and low-powered (heater mesh dissipated power) runs resulting in mainstream temperature rises of approximately 50°C and 30°C. During initial facility pressurisation, at 12 s, there is a transient spike in recorded temperature, which quickly re-settles to ambient conditions as the pressurisation shock passes through the facility. Coolant bypasses the mainstream heater mesh and is maintained at ambient temperature throughout the run. At 40 s the heater mesh is switched on and imparts a sudden temperature rise into the mainstream flow, taken as the average of

upstream rake measurements. Due to the low thermal mass of the mesh, once power is cut (at 80 s) it rapidly cools down and the mainstream flow returns to ambient temperature. At 110 s the flow is shut off and the facility vents to atmosphere causing a drop to sub-ambient temperature as the pressure decreases. The example metal platform traces are taken as the average of a 5×5 pixel area on the surface recorded by the IR camera. A non-uniform inlet temperature profile occurs as a consequence of upstream surfaces being heated during the run. This is discussed in greater detail below but is demonstrated by the upstream wall trace, measured by surface mounted thermocouples immediately upstream of the test blades, which gradually rise over the heated period. Test components also gradually heat up over the heated period, time constants vary across the surface and depend on local heat transfer coefficients. For this particular location (just upstream of of the leading edge) the measured time constants of the platform surface temperature for the two runs shown were 6.3 s and 6.7 s. Data used to determine metal effectiveness is typically taken from the final second of the 40 s heated period, by which time the metal temperatures are approximately steady.

5.1.1 Inlet temperature profile

The metal sections which make up the MERC facility have large thermal masses and, as a result, surfaces upstream of the test section gradually heat up during the heated period of the run. A consequence of this is that a relatively large thermal boundary layer develops, as heat passes from the heated mainstream air into the surrounding walls, leading to a non-uniform inlet temperature profile in the test section. The facility does not currently have an upstream traverse system which would allow detailed surveys of the inlet profiles however the two sets of rakes within the test section provide a coarse measurement across the majority of the inlet plane. Pressure plots are omitted here as none of the rake measurements are close enough to endwalls to measure aerodynamic boundary layers. Total pressure measurements are consistently within 0.5%, the majority of deviation is expected to be due to slight angular misalignments of the probes.

The inlet aerodynamic boundary layer within the cascade has not been measured or quantified, though this could be achieved in the future using a flat-head probe to traverse close to the endwall

surface. Within the upstream NGV row in an engine, the hub boundary layer is swept across the passage and entrained within a passage vortex, in a similar process to within the rotor bladerow. Downstream of the NGV passage vortex lift-off line, fresh boundary layer growth occurs and it is this which is incident on the rotor blades in an engine. The inlet boundary layer to the bladerow in the cascade would therefore be expected to be much thicker than that in a realistic engine environment as the upstream surface along which it grows is much larger, and no boundary layer bleed is present upstream of the test section or blade cassette. The cascade boundary layer grows initially at low velocity in the upstream chambers of the cascade stack however it is expected to become thinner as flow is accelerated through the inlet contraction. The extent to this is unknown and requires additional measurements, but it is still expected to be much thicker than in an equivalent engine when it reaches the bladerow. This is important to consider as the formation of secondary flows within the bladerow is largely driven by the roll up of the inlet boundary layer at the blade leading edge and hence the formation of the horsehoe and passage vortices. Although the formation of these vortices will not fundamentally change, the relative amount of flow within them and therefore their relative strength will be affected.

Figure 5.2a shows temperatures recorded by the upstream thermocouple rakes against normalised spanwise position at a time halfway through the heated period (60 s) for the run shown in Figure 5.1. There are five thermocouple rake probes in total, spaced evenly across the tangential distance and each with 6 evenly spaced measurement points across the span. The temperature profile is approximately flat between span fractions of 0.2 and 0.8 but tapers off towards the hub (span fraction of 0) and case (span fraction of 1) endwalls where temperatures are approximately 7 K lower than the midspan core. Probes are numbered based on their position in the tangential direction with Probe 1 being nearest to the suction surface endwall and Probe 5 being nearest to the pressure surface endwall (Probes 1 and 5 each have a missing data point due to thermocouples being damaged during assembly). A slight temperature gradient can be seen across the span as Probe 1 measures consistently higher temperatures, and Probe 5 consistently lower temperatures, across the entire span. This may be due to the pressure gradient across the 5 blade passages biasing more of the incoming hot mainstream flow towards the suction surface endwall. Probes 2–4 show excellent uniformity in measurements, with the exception of one midspan data

point for Probe 2 which was suspected to be damaged during installation. These three probes cover the area upstream of the two metal blades of interest.

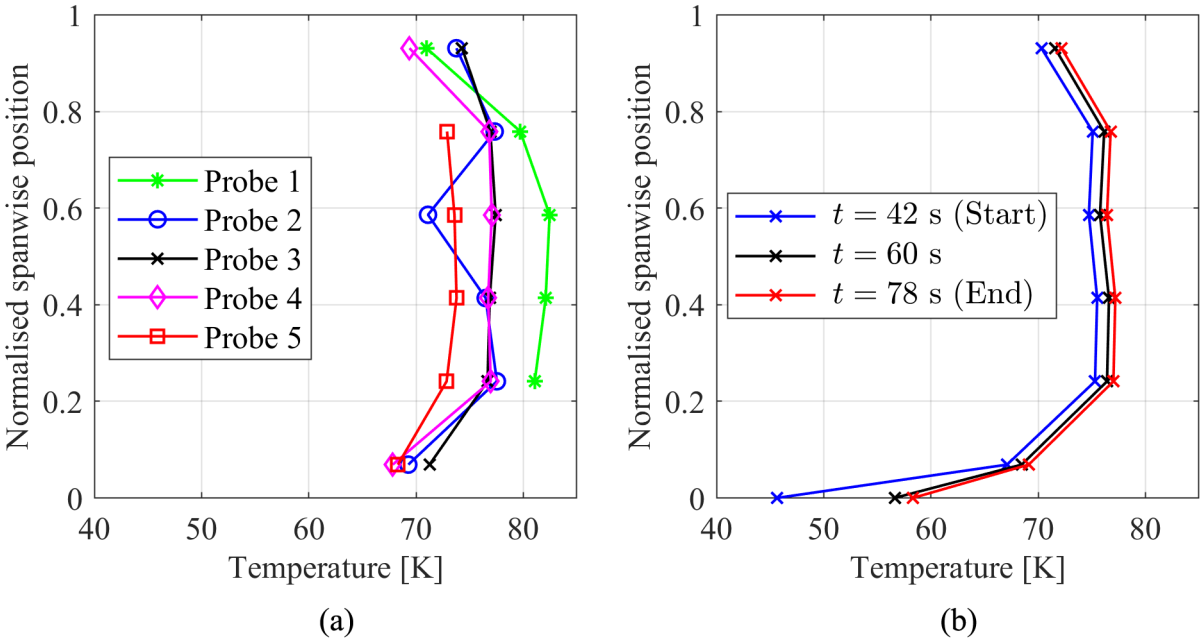


Fig. 5.2 (a) Temperatures measured by test section inlet rakes against span fraction for 5 rake probes at a time midway through the heated period in run ($t = 60$ *s*) and (b) spanwise inlet temperature profile at three times ($t = 42$ s, $t = 60$ s, and $t = 78$ s) corresponding to start, mid, and end times of the heated period during a run.

The inlet temperature profile develops over a long time as upstream surfaces heat up. This time period (greater than 10 minutes) far exceeds the possible run time of the heater and therefore a true steady-state inlet temperature profile cannot be achieved. Figure 5.2b shows the average of Probes 2–4 at three times near the start, middle, and end of the heated period (42 s, 60 s, and 78 s respectively). Measurements from surface-mounted thermocouples, located on the front seal, provide temperatures along the hub endwall and show the full extent of the thermal boundary layer. The temperature at the hub 2 s after the heater mesh is turned on ($t = 42$ s) is significantly lower, by 29.3 K, than the core midspan temperature. As upstream surfaces heat up less heat is drawn from the mainstream flow and hence the hub temperature increases rapidly from 45.7 K at $t = 42$ s to 56.6 K at $t = 60$ s and then more gradually rises to 58.3 K at $t = 78$ s. The approximate time constant of this temperature rise is 2.5 s and a similar trend is expected on the opposite case endwall. The 0.2 to 0.8 span fraction temperature profile remains relatively flat across each timestamp but the core temperature increases slightly (moving profiles to the right

in Figure 5.2b) from 75.5 K to 77.2 K as warmer upstream surfaces draw less heat away from the total flow. The hub wall temperature remains significantly lower than the core temperature throughout the heated period. It would take an impractically long time for the facility to fully heat up to a truly steady-state condition, far longer than the heater mesh could be safely operated for. Due to the lack of external insulation the internal surfaces would not heat up to the same temperature as the core flow, even with an extremely long heated period, and a non-uniform inlet temperature profile would still occur.

The non-uniform inlet temperature profile caused by this inlet thermal boundary layer affects values of metal effectiveness and adiabatic film effectiveness determined from measurements in the MERC facility. When calculating effectiveness values a hot gas reference temperature, normally the recovery temperature of the core flow, is used as this is assumed to be the temperature driving heat transfer at the surface of interest. If a non-uniform inlet temperature profile exists then inlet streamlines at a range of temperatures, within the thermal boundary layer, are convected into the downstream flow and may interact with the surface of interest. The temperature of a streamline, or the average of multiple streamlines at various initial temperatures which mix together, and interact with the surface at a given location is unknown. It depends on the temperature of the core flow, the temperature of the endwall flow and the temperature profile. As the thermal boundary layer develops over time, due to upstream surfaces heating up, the streamline unknown temperature is a function of time and is generally lower than the measured core flow temperature.

In this situation using the core flow temperature to calculate effectiveness will result in inferred values greater than the value that would be obtained with an uniform inlet temperature profile. For comparative studies (performed in the same facility and similar mainstream temperatures and runtimes) where the difference in effectiveness values, rather than the absolute values, are of primary interest this may not significantly affect conclusions. If experimental results are being scaled to engine conditions to assist in design of new components however then correcting for these erroneously high effectiveness values is critically important as otherwise components may end up under-cooled and fail at engine conditions.

A new processing method was developed, detailed in the paper submitted to journal and included as Chapter 6, to correct adiabatic film effectiveness measurements in the presence of a time-varying non-uniform inlet temperature profile, such as the thermal boundary layer present in this facility. Implications when determining metal effectiveness values for the results presented in this thesis are detailed below.

5.1.2 Heater mesh characteristic

The mainstream temperature rise can be varied as required by setting the voltage on the DC power supply connected to the heater mesh. The voltage can be continuously varied up to 110 V giving theoretical mainstream-to-coolant temperature ratios in the range 1–1.2. A series of test runs were conducted with 5 V set voltage increments to determine the power characteristics of the installed heater mesh. Figure 5.3 shows results of the power dissipated by the mesh and the mainstream temperature rise achieved for the nominal operating point with a mainstream mass flow rate of 2.25 kg s^{-1}. Also indicated on Figure 5.3b are approximate mainstream-to-coolant temperature ratios achieved when using ambient temperature air as coolant. This series of tests was performed to assist in accurately setting desired temperature ratios, by determining the required set voltage from the characteristic, and also to test the capability of the heater design.

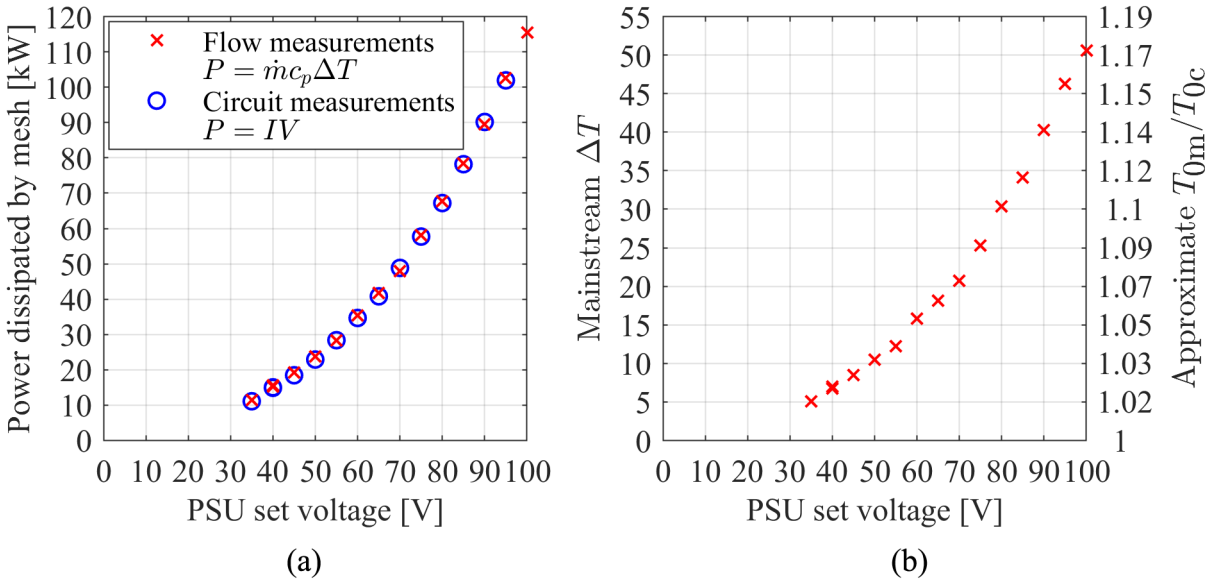


Fig. 5.3 (a) Power dissipated by the heater mesh (b) mainstream temperature rise, and approximate mainstream-to-coolant temperature ratio, achieved in the MERC facility.

Power dissipated by the mesh was calculated by two methods, given by equation 5.2. The circuit current and voltage drop across the mesh were recorded for each run. Voltage across the mesh is measured by wires directly attached to the brass bus bar terminals which connect to a readout on the main control room panel. The circuit current was recorded manually from a readout dial on the DC power supply. Flow measurement of the mainstream temperature rise and mass flow rate were also used to calculate dissipated power. Values calculated from both methods were consistent (within 4.5%) for all runs.

$$P = IV = \dot{m} c_p \Delta T \tag{5.2}$$

The maximum power dissipated by the mesh was 117 kW for a set voltage of 100 V and had a corresponding mainstream temperature rise of 52 K. The DC power supply is nominally able to output 150 kW and can have a maximum set voltage of 110 V however the circuit current reading was unreadable as it was beyond the maximum 1500 A on the dial. It was therefore decided to limit the set voltage to 100 V which may serendipitously have improved the lifespan of the heater mesh.

5.1.3 Coolant capacity calibration

Coolant is supplied to the three plenums of the blade cassette from a common chamber in the blade cassette charged with air from the main facility regulator at 15 bara. The total mass flow rate of coolant is measured via an orifice plate assembly installed along the coolant supply line but is set by the number and diameter of holes in the coolant setting plate (the total choked flow area). A number of plates were manufactured with differing numbers of holes to quickly change coolant supply conditions. Each hole is identical, with a drilled diameter of 1.5 mm. On the downstream side each small hole opens into a larger 6.0 mm diameter hole, angled at the opening to try and assist in flow discharge stability. All of the holes are aligned along slots in the blade cassette so that they vent into one of the three SAS seal plenums. Pressures within the test section are significantly lower than the 15 bara coolant supply pressure (less than 3 bara) and therefore all of the holes remain choked and at a constant mass flow rate throughout each run.

The coolant flow through each, geometrically identical, hole is expected to be evenly distributed. The high pressure supply plenum has a relatively large total volume and a flow area far greater than that of the holes. Each row of holes is sufficiently far away from endwalls that no adverse affects on flow distribution within the supply plenum are expected.

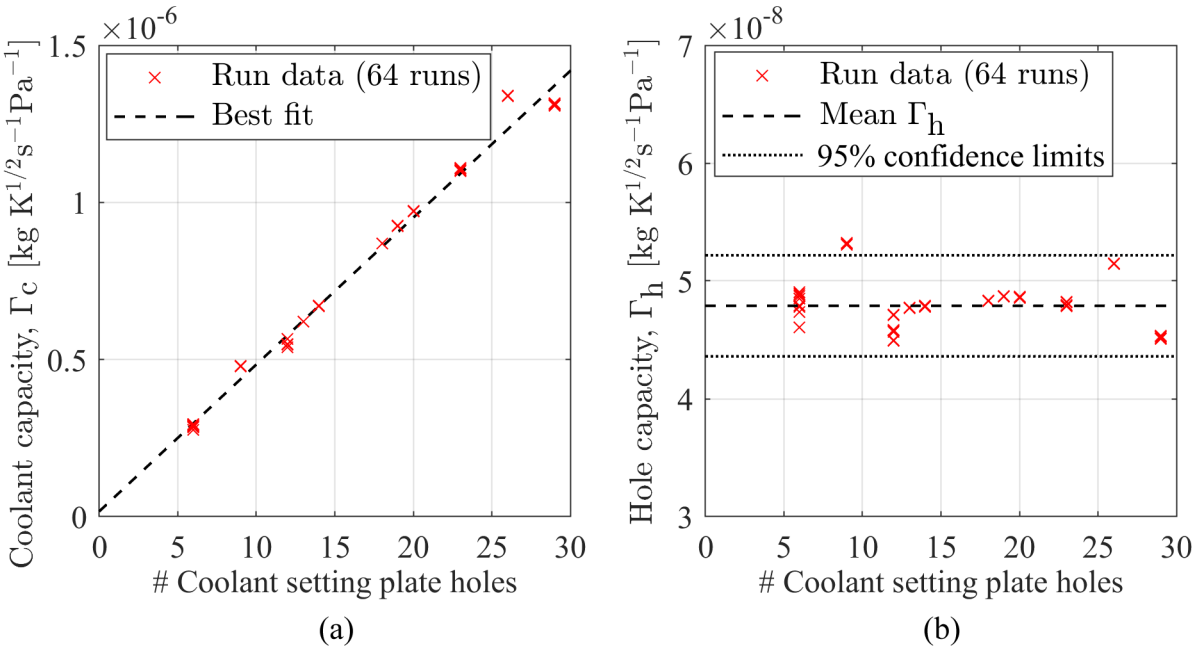


Fig. 5.4 (a) Total coolant capacity of plates, measured across 64 runs in the MERC facility with numbers of coolant setting plate holes in the range 6–29 (b) per-hole capacity calculated from results of the same 64 runs showing mean value and confidence band.

The amount of coolant through each hole is the same and therefore by measuring the total mass flow rate into the blade cassette, and with knowledge of the number of holes, it is simple to accurately calculate how much coolant is supplied to each plenum. During experiments the coolant mass flow rates can be set as a boundary condition by selecting an appropriate number and size of holes. This requires knowledge of the flow capacity (equation 5.1) of this particular hole configuration. The total coolant capacity was calculated from 64 runs in the facility using coolant setting plates of varying numbers of holes in the range 6–29. Figure 5.4a shows the calculated total coolant capacity against the number of holes for these runs. The total coolant capacity increases linearly with hole number which suggests that the initial assumption, that coolant is split evenly between each hole, is correct. The per-hole capacity calculated for each run is plotted in Figure 5.4b against the total number of holes for that run. The mean value across all runs, and

therefore the expected capacity of a single hole, is calculated as 4.8×10^{-8} kg $K^{1/2}$ s^{-1} Pa^{-1} with a $\pm 10\%$ uncertainty to 95% confidence.

5.2 Downstream area traverse

The downstream traverse system is capable of covering almost the entire five-passage exit plane, located 1 mm downstream of the blade trailing edge. The probe is aligned with the exit angle of the blades and points directly into the exit flow. The traverse area is limited by the thickness of the probe arm meaning approximately 70% of the total span can be surveyed. Typically only the central 2.5 blade passages were covered to reduce run times and increase measurement resolution. The location of the traverse plane and the typical measurement area is shown in Figure 5.5 on a schematic of the test section and rendered into 3D CAD geometry. Results shown in the 3D CAD image are real traverse measurements presented in full below.

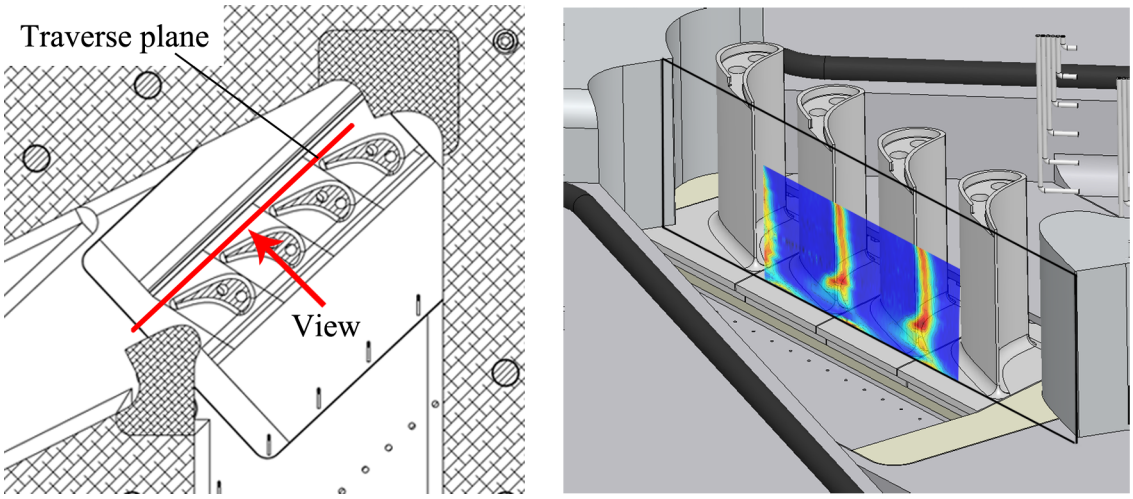


Fig. 5.5 Schematic cross section and 3D CAD render of the MERC facility test section with downstream traverse plane and general measurement area indicated. Arrow indicating view relates to full plots of traverse results included in this chapter.

5.2.1 Transonic probes

A number of bespoke transonic probes were designed for use with the traverse system which share a common mounting bracket to the traverse carriages and are interchangeable between runs. A pair of multi-hole probes were made for aerodynamic measurements: a five-hole probe

for mainstream flow measurements; and a three-hole yaw probe for near-wall measurements. A temperature probe with two aspirated K-type thermocouples was also made. Each of the probes were manufactured via DMLS in stainless steel. Figure 5.6 shows photographs of the three-hole and temperature probes, along with schematics showing the passages passing through the printed steel bodies.

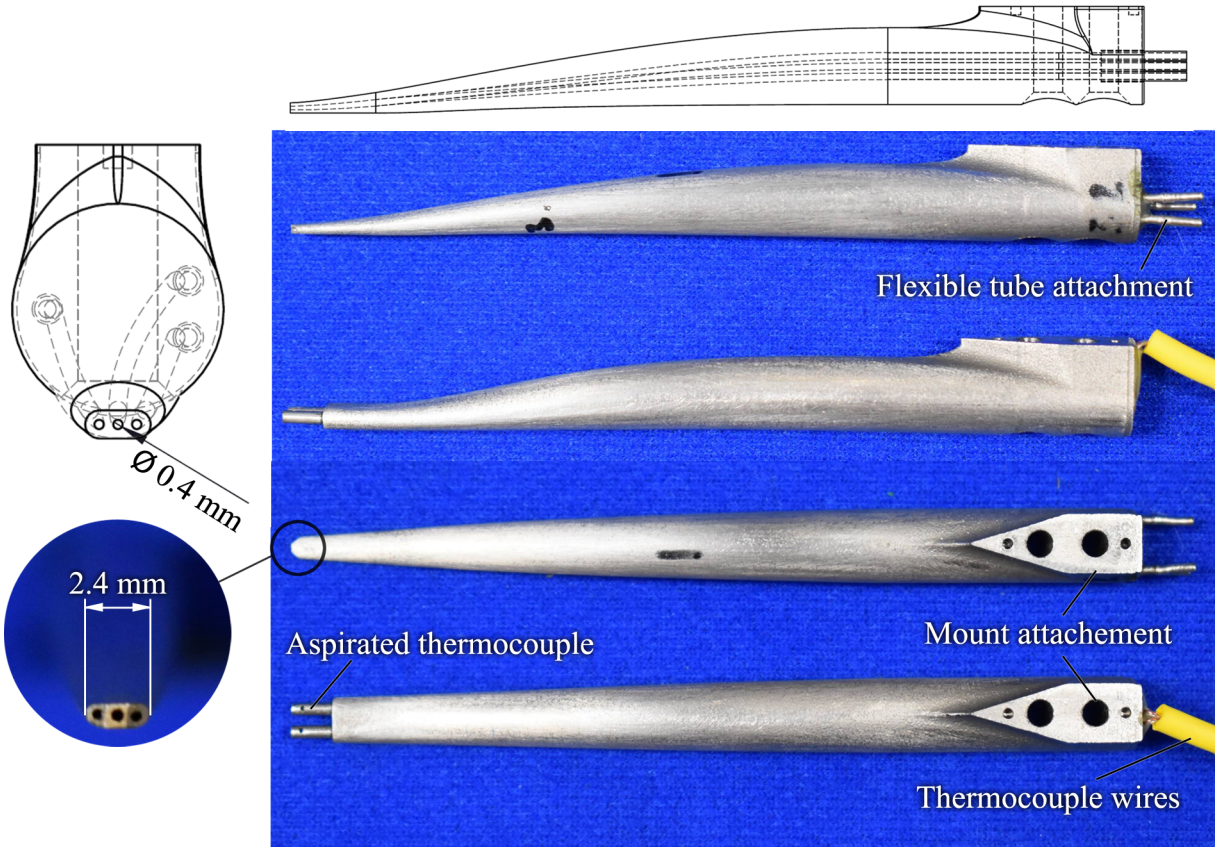


Fig. 5.6 Schematic diagrams and photographs of bespoke transonic probes used within the MERC facility. Two designs are shown of a three-hole probe for aerodynamic measurements and an aspirated thermocouple probe for temperature measurements.

Each probe is 10 cm long, which was the minimum length which could be achieved with the possible location for the traverse slot. The exit flow is transonic and therefore the probes experience large aerodynamic loads during facility operation. They were designed to be stiff enough to prevent flexing and excessive tip vibration, which resulted in the main probe bodies being 9 mm in diameter. The probe tip sizes were made as small as manufacturing tolerances allowed and the bodies aerodynamically contoured, to smoothly transition between the small tips and main body, to reduce loading and flow blockage. The probes are clamped to a mounting arm,

which is in turn connected to the traverse carriages and descends into the test section through a slot, via a pair of bolts. When tightened, the bolts sit flush with the bottom of the probe body. The mounting arm is machined out of steel and has an aerodynamic wedge profile. This profile can be seen in Figure 5.6 as the printed probe body transitions smoothly to match it. A pair of 1 mm dowel pins can also be seen flanking the bolt through-holes and are used for situating the probe during assembly. The probe length and mounted position are fixed meaning only a single downstream plane can be measured with these probe designs. Plans were made for additional probes of differing length, and for a variable distance mounting system which would allow for multiple downstream planes to be surveyed however these have not yet been implemented.

The temperature probe aspirated thermocouples are beaded K-type thermocouples which are held in ceramic inserts and shielded from the transonic flow by steel tubes with a number of perforations in the walls. The purpose of this assembly is to slow down incoming flow to prevent damage to the delicate thermocouple beads whilst still registering an accurate flow temperature. This assembly was based on designs previously successfully implemented in the Oxford Turbine Research Facility. Unfortunately, only a small number of commissioning runs were completed with the temperature probe before both thermocouples were damaged and could no longer be used. The cause of these failures was not identified and due to time limitations a second iteration of the design to try and make the assemblies more robust could not be completed.

The multi-hole probe tips are 2.4 mm in diameter, with each hole having a diameter of 0.4 mm and minimum wall thickness between them of 0.3 mm which were (at the time of manufacture) the minimum recommended dimensions by the manufacturer. Tips are printed flat and then finished off in the laboratory workshop by machining faces at 45° angles. Passages run through the main probe body which connect the tip holes and route around the mounting bolt holes to bulbed stainless steel connectors at the back of the body. These in turn connect to flexible urethane tubing and the *kulite* pressure transducers mounted on the traverse cassette, allowing for tip pressures to be recorded. The flexible tubing, and wires of the thermocouple probe, are taped into a slot on the back of the mounting arm using aluminium tape to protect them from the flow.

Aerodynamic multi-hole probes were calibrated in the laboratory Transonic Calibration Tunnel (TCT). Transonic multi-hole probe calibration techniques are complex and are only briefly discussed here. The basic principle is to create calibration lookup tables of non-dimensional coefficients based on each of the hole pressures measured by the probe in known flow conditions. Pressures recorded in-situ in the MERC facility are used to calculate an equivalent set of coefficients and the lookup tables are then interpolated for these values to infer the pressures, flow angles, and Mach number of the flow. Probes were mounted in the calibration tunnel using the same fixed configuration as in the MERC facility so that probe alignment to incoming flow was identical between facilities. Probe calibration lookup tables were made by putting the probe into a jet, of known Mach number and pressure, and sweeping through a specified range of angles ($\pm 30^{\circ}$) in two axis. The process was repeated for jet Mach numbers in the range 0.1–1.5 to create a complete lookup table over the likely ranges found in the MERC facility.

Significant problems arose during the calibration process, which were not unexpected as calibrating probes over the transonic range of Mach numbers is notoriously difficult due to complex shock systems forming in and around probe tips. Extensive testing and attempts to fix issues were undertaken but the underlying cause, or likely multiple causes, of the calibration problems encountered were not determined and are the subject of ongoing work. This meant that reliable calibrations over the full range of Mach numbers could not be obtained for either of the probe designs, with particular difficulties in the Mach number range 0.9–1.2. Three sets of the multi-hole probes were manufactured and calibrated, all of which experienced similar calibration problems. Similarly, probes of different design, not additively manufactured, which were calibrated in the TCT also had issues with calibrations over this same transonic range.

A number of potential problems with the probe design and manufacture were ruled out. Continuity of the passages through the probe were verified by submerging the probe tip in water, blowing air through each passage in turn, and observing the air bubbles emerging from the tip holes under a magnifying glass. Leaks between passages, which might occur due to small cracks or gaps formed during the printing process, were tested by building a test jig which allowed the probe tip to be pressurised. Pressures from each hole were measured and were disconnected in turn, causing air to flow through one passage at a time. Leaks would be indicated by pressure

drops detected by the transducers still attached to the other holes. The three-hole probe showed no signs of leakage however the five-hole probe tests were inconclusive. Probe response times, which are affected by the tip hole size, fill volume, and passage area changes between the tip and transducers, were estimated by conducting balloon tests where a sudden pressure change (a balloon bursting) is measured at high frequency. Measured probe responses varied between 10 and 50 Hz which was deemed acceptable for the relatively slow traverse speeds conducted.

The five-hole probe calibration was significantly more unstable than the three-hole probe and therefore the three-hole probe was used for all area traverses reported in this thesis. Despite the problems encountered during the calibration process, a calibration for the three-hole probe that was stable and reliable (i.e. the data points used to create the interpolated lookup tables can be very accurately reproduced from the calibration) for Mach numbers in the subsonic and low transonic range of 0.1–0.9 was achieved. Interestingly, the stability of the calibration also improves for fully supersonic flows with Mach numbers greater than 1.2. At these higher Mach numbers the tip flow state is relatively simple as a detached bow shock forms upstream of the probe and pressures measured by the probe are stable, allowing for accurate calibration. Mach numbers encountered at the exit plane of the MERC facility are nominally below 0.9 and therefore this calibration of the three-hole probe was used for all traverses hereafter.

5.2.2 Midspan traverse demonstrating passage periodicity

Periodicity of the cascade was checked by conducting a midspan traverse at the exit plane of the central 2.5 blade passages using the three-hole probe. Results of the measured total pressure, static pressure, and Mach number are presented in Figure 5.7. The total pressure profile is near-perfectly periodic with a central passage total pressure of 2.45 bara. The wakes of each blade can be seen by the drops in pressure to 2.20 bara at 0.5, 1.5, and 2.5 blade passages. Air within the wakes is highly turbulent and has a lower velocity than the central passages, leading to an increase in static pressure in these wake regions. The static pressure distribution varies across each passage, with the lowest static pressure occurring along the suction surface, where the flow velocity is highest. The pressure distribution across blade passages, from 0.5 to 1.5

and 1.5 to 2.5 is approximately periodic. A slight pressure gradient can however be seen, with decreasing pressure when moving from 0 to 2.5 blade passages. This pressure gradient exists across the entire cascade blade row due to the limited number of blades and high flow turning angle. A consequence of this pressure gradient is that the Mach number distribution, Figure 5.7b, is higher in the 1.5 to 2.5 blade passage by approximately 4%.

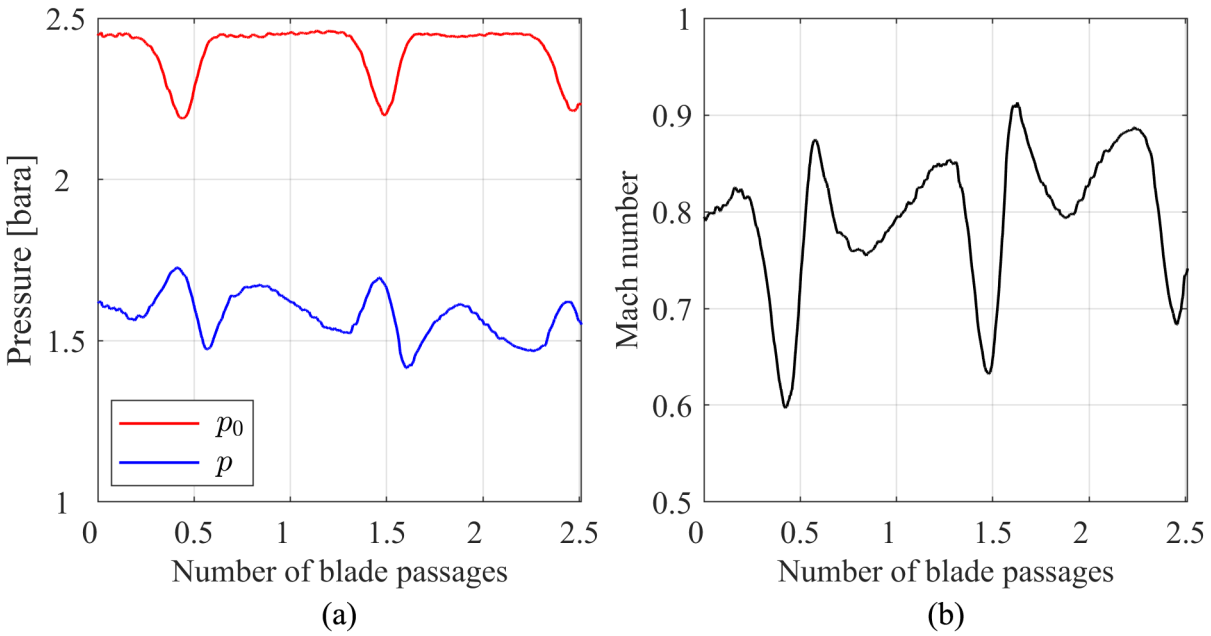


Fig. 5.7 Traverse measurements of midspan (a) pressure and (b) Mach number distributions at the bladerow exit plane. Measured using a calibrated transonic 3-hole probe.

5.2.3 Area traverses with pressure and temperature probes

Area surveys are conducted by setting a predefined traverse path along the exit plane of the bladerow, with the probe aligned with the blade exit flow angle, at a number of discrete spanwise locations. The traverse path can be customised to cover any required area at any resolution, limited only by the runtime of the facility, as it is programmed into the automatic *LabVIEW* controller as a set of coordinate points. Typically, the longer tangential direction (the unrolled circumferential direction) is set for continuous strokes, and spanwise increments carried out at the endpoints of these strokes. Data is measured continuously at 500 Hz throughout each stroke. The probe position is monitored and recorded throughout the traverse by the two LVDTs in the traverse system. Figure 5.8a shows time traces of probe position for an area traverse covering

the large blue traverse path shown in Figure 5.8b. The traverse path starts just above the hub endwall, and makes 25 strokes in total at incremented spanwise locations towards the case. The control system uses the measured probe position to determine move commands to be sent to the motor control units. Slight overshoots occur at the end of each stroke motion, shown by the small bumps at the peaks in tangential position. These overshoots are accounted for in the safe spatial limits when programming traverse points. More detailed surveys were conducted over a smaller area, denoted by the dashed red lines, as this area was found to contain the majority of flow features of interest, namely the blade wake and passage vortex which cause the majority of aerodynamic loss in the system.

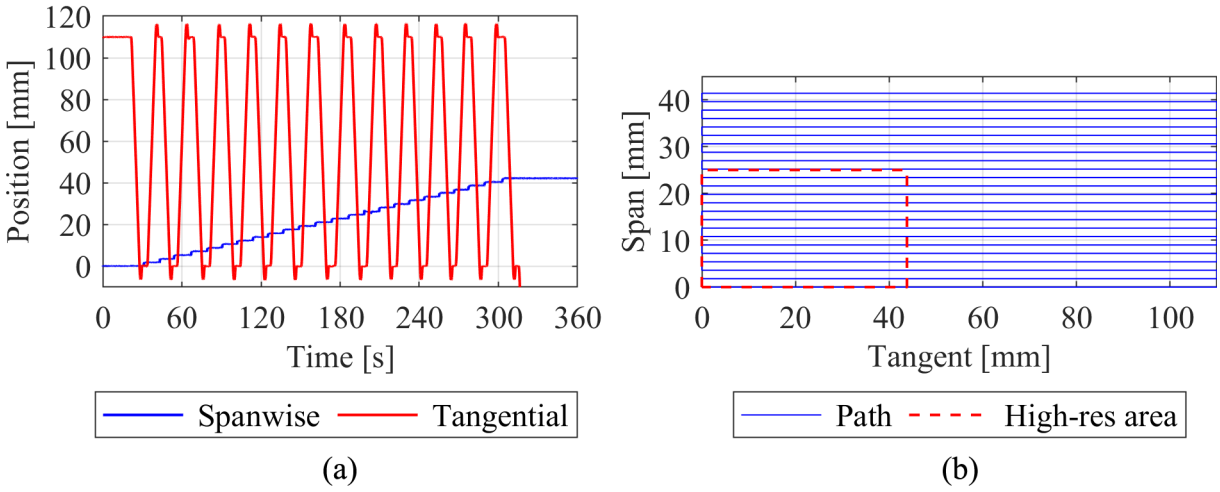


Fig. 5.8 (a) Traces of probe position during a typical area traverse and (b) spatial plot of the traverse path and surveyed areas.

Figure 5.9 shows an area traverse result of total pressure loss coefficient, Y_p, for a survey of the central 2.5 blade passages by 0.7 span fraction, and Figure 5.10 shows a temperature traverse result for the smaller 1 blade passage by 0.4 span fraction. Both of these results are from single experiments, of approximately similar run times. Plots of the two surveyed areas are to scale with each other. The wakes of each blade can clearly be seen by bands of high loss coefficient across the entire span. Towards the hub, between 0.1 and 0.3 span fraction a large loss core can be seen which is caused by the passage vortex. The passage vortex entrains coolant which leaks out of the front seal and flows along the suction surface of the blade. At the trailing edge the vortex sheds and interacts with the wake, pulling it towards the middle of the passage. The peak value of the loss coefficient core, at 0.25 span fraction where the wake and pasage vortex

interaction is strongest, is 0.49. A wider region of high loss coefficient is present along the hub aligned with each blade wake and is caused by the fillet between blade and platform, which is wider than the rest of the blade trailing edge.

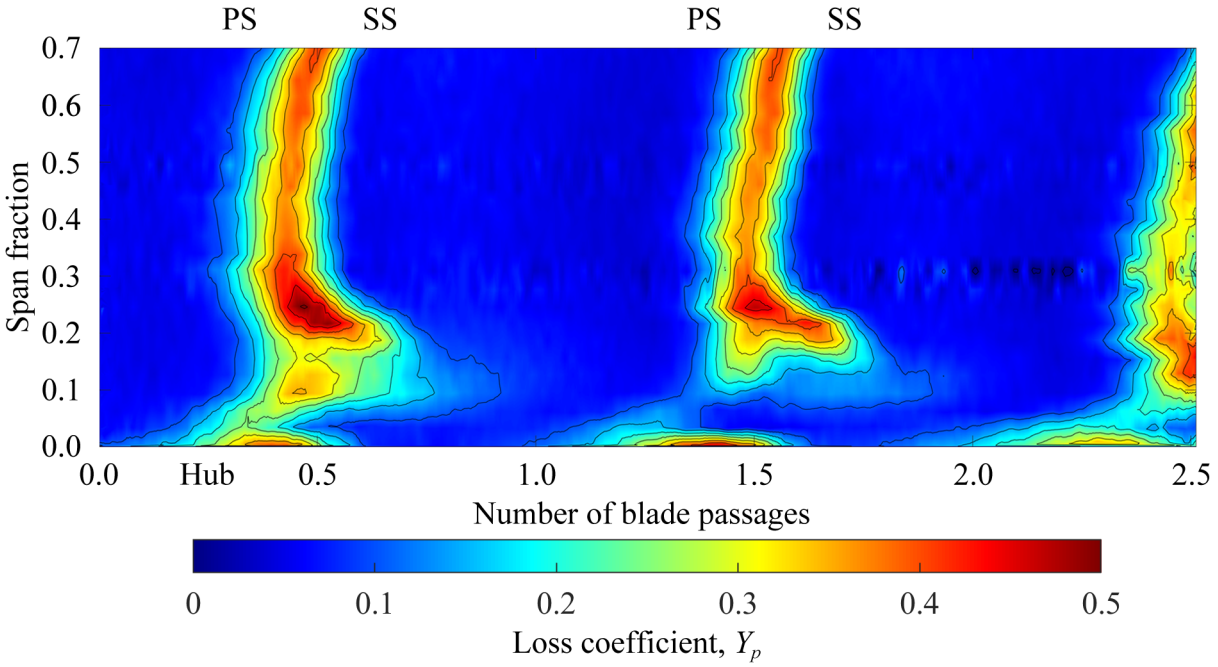


Fig. 5.9 Downstream traverse result of an area of approximately 2.5 blade passages by 0.7 span fraction of total pressure loss coefficient. Measured using a calibrated transonic 3-hole probe. PS and SS indicate approximate locations of blade pressure and suction surfaces repectively.

The distributions of loss coefficient across each passage are very similar and the majority of losses caused by the blade wakes and their interaction with the passage vortex are encapsulated within the lower 0.4 span fraction. Aerodynamic surveys were conducted in the facility to compare the aerodynamic impact of purge flows at a range of test conditions. For these comparative studies, detailed traverses were conducted over the smaller 1 passage by 0.4 span fraction area as shown in the temperature plot. The wake and passage vortex can also be seen in the temperature survey by the regions of relatively cool fluid. Values plotted here are the measured exit plane temperatures, $T_{02,m}$, normalised by the inlet mainstream temperature, $T_{01,\mathrm{m}}$.

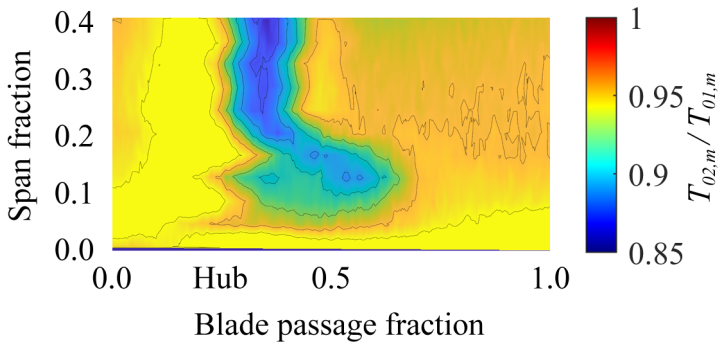


Fig. 5.10 Downstream traverse result of an area of 1 blade passage by 0.4 span fraction of exit temperature normalised by inlet mainstream temperature.

5.3 IR calibration and determination of wall temperature

Area surface temperature distributions are measured with an IR camera mounted externally to the facility with optical access via a ZnSe IR window. The IR camera detects incident thermal radiation using a sensor comprised of an array of microbolometers. These microbolometers have an electrical resistance which varies with temperature, and thus the amount of incident IR radiation absorbed. The camera hardware converts analogue voltage signals from the detector into a digital raw count value which is approximately linear with incident IR radiance. A calibration is applied in order to infer the temperature of the target surface from the detected radiance. This requires two stages: first to calculate the blackbody parameters of the camera hardware; and secondly to account for greybody parameters arising from the target object and surroundings.

The radiance detected by the camera sensor (q_{cam} given in equation 5.3) is the combined total of: emission from the target surface, q_{obj}; reflected emission from surroundings, q_{ref}; and emission along the optical pathway, q_{obj}. An equivalent black body temperature, T_{bb}, is related to the actual target surface wall temperature using Stefan-Bolzmann's law specified in equation 5.4. Here, T_{bb} represents the temperature of an ideal black body that would emit a radiance equal to that detected by the camera. The additional terms in equation 5.4 are: the Stefan-Boltzmann constant, σ_{bb}; the target surface temperature, T_{w}; the effective temperature of the surrounding inner walls of the cascade, T_{s}; the temperature of the optical pathway, T_{opt}; the emissivity of the target surface, ε; and the transmissivity of the optical pathway, τ.

$$q_{\text{cam}} = q_{\text{obj}} + q_{\text{ref}} + q_{\text{opt}} \tag{5.3}$$

$$\sigma_{\text{bb}} T_{\text{bb}}^4 = \sigma_{\text{bb}} \varepsilon \tau T_{\text{w}}^4 + \sigma_{\text{bb}} (1-\varepsilon) \tau T_{\text{s}}^4 + \sigma_{\text{bb}} (1-\tau) T_{\text{opt}}^4 \tag{5.4}$$

A blackbody calibration was performed in the laboratory and used to correct drift in the supplied factory calibration. Transmissivity of the ZnSe window was determined by conducting a calibration in-situ in the MERC facility. The surrounding radiance varies continuously throughout runs due to changes in temperature of the surrounding walls within the cascade. It is measured and corrected by observing surface-mounted thermocouple assemblies embedded in the cascade endwall on a per-run basis. The following sections describe the methods and results of both stages of calibration which were conducted and applied to all subsequent surface temperature results in this thesis.

5.3.1 Blackbody calibration

The *TIM640* IR camera was supplied with a factory blackbody calibration which is automatically applied by the capture software to a stated accuracy of ± 2°C or $\pm 2\%$ of the measurement (output in celsius). The calibration converts the raw count of the camera sensor, C_{r}, to a temperature reading based on a linear regression of the form given in 5.5 where the parameters a and b are blackbody parameters specific to the camera. These parameters can drift over time and so a blackbody calibration was performed in the laboratory. The raw count of the camera cannot be directly read-out from the the manufacturer software meaning the parameters a and b could not be directly determined. Instead, a linear correction was found and applied to the equivalent radiance of the factory-calibrated blackbody temperature.

$$\sigma_{\text{bb}} T_{\text{bb}}^4 = a C_{\text{r}} + b \tag{5.5}$$

The laboratory blackbody calibration was conducted using a portable *CI Systems SR-800N-4A-ET* high-accuracy blackbody emitter and control system. The emitter is an enclosed system

containing a heater and refrigerator attached to a target blackbody emission surface, which is heated or cooled to a uniform temperature as required. The temperature range of the emitter is 0 to 175°C and can therefore cover the full temperature range observed in the MERC facility. The emission surface consists of a 4" square area made up of an array of pyramid structures and has been coated with a very high-emissivity black paint. The stated emissivity of this surface is 0.99 ±0.01 meaning it is an almost-perfect blackbody. A dedicated controller is used to set a target temperature for the surface and a feedback loop allows for achieved surface temperatures to be within ±0.010°C across the central 80% of the surface. The blackbody emitter system is shown in Figure 5.11. To prevent damage to the delicate emission surface, when not in use, a perspex cover is attached to the front of the emitter. This cover is removed when performing calibrations with the system.

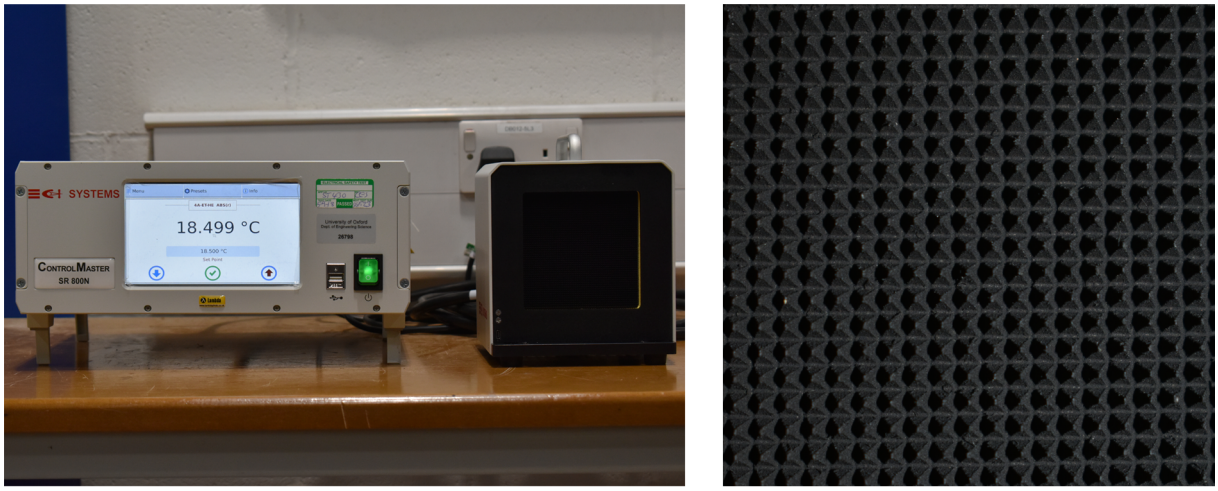
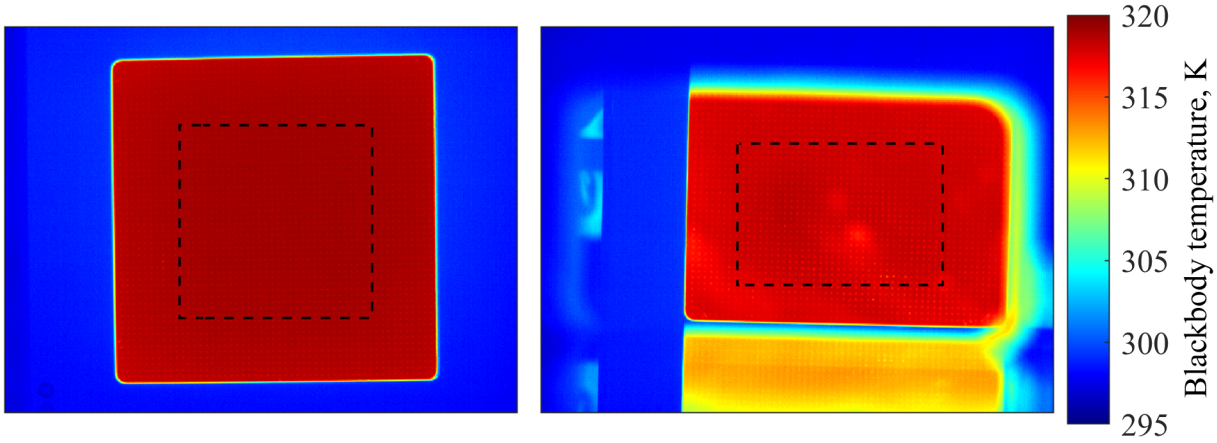


Fig. 5.11 Blackbody radiation emitter system and surface.

The system was set up with the IR camera mounted on a tripod looking directly at the blackbody emitter such that the surface was normal to the camera, the view can be seen in Figure 5.12a. The controller was then manually set to temperatures, at 5°C intervals, from 0 to 100°C. The system was allowed to reach a steady temperature at each interval which generally took a couple of minutes. The controller system notifies the user when a steady average temperature within 0.01°C of the target set point is reached. Snapshot images were taken of the emitter surface for each temperature. Ambient temperature was recorded for each snapshot from

the camera internal sensor, displayed in the camera software interface, and varied by less than 0.1°C over the approximate one hour that testing took place.



(a) (b)

Fig. 5.12 IR camera captured images of blackbody temperature of the blackbody emitter system (a) on a desk with direct view for blackbody calibration (b) in-situ in the cascade behind the ZnSe IR window to determine window properties.

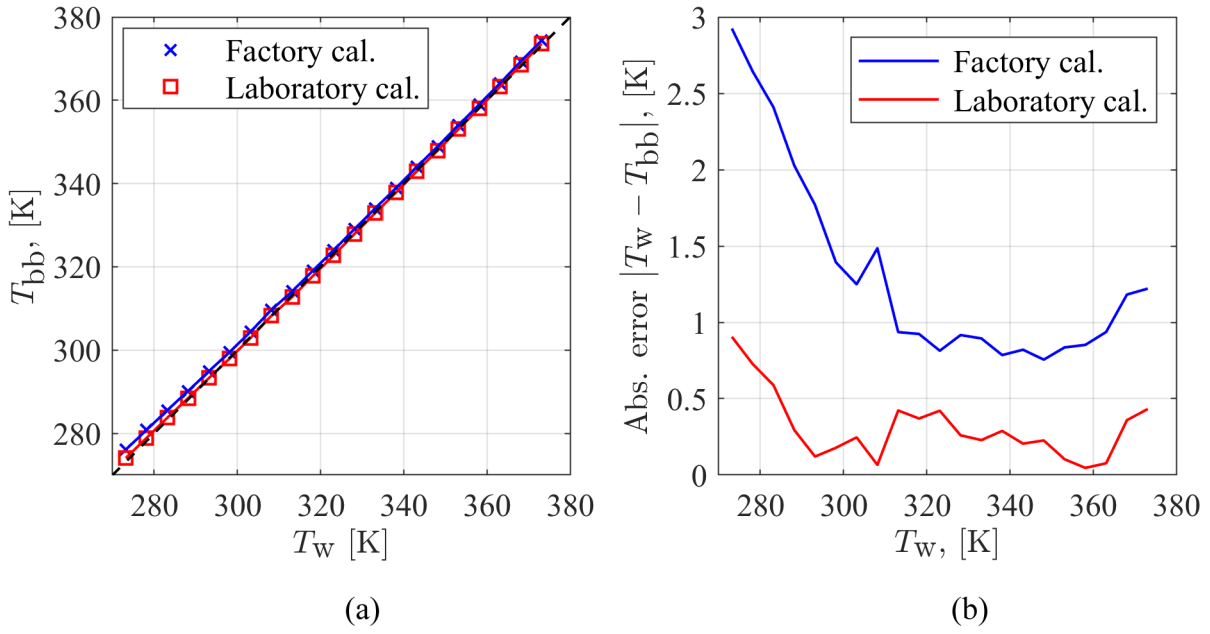
The calibration procedure was conducted in an area without any other external heat sources so that the only reflected radiation was from ambient temperature surroundings. The camera was set up at the minimum focal length of 0.5 m and had an unobstructed view of the emitter surface. Over this small distance the transmittance of air is approximately unity [22]. This means that equation 5.4 can be reduced to equation 5.6 allowing for corrections, from the known surface emissivity of the emitter and ambient temperature, to be made. This can be done automatically by the camera software however this makes correcting values during post-processing more difficult. Values of emissivity and transmissivity are therefore set to unity within the camera software and the output factory blackbody temperature is corrected during post-processing.

$$\sigma_{bb} T_{bb}^4 = \sigma_{bb} \varepsilon T_w^4 + \sigma_{bb} (1-\varepsilon) T_s^4 = aC_r + b \tag{5.6}$$

Images were processed in *MATLAB* and the area enclosed by the dashed lines in Figure 5.12 was averaged to give the factory-calibrated temperature of the camera. The averaged area, in the centre of the target surface, was chosen as it was the most stable temperature region of the emitter surface. The standard deviation of temperatures in the averaged area ranged from a

minimum of 0.13 K at 303 K to maximum 0.43 K at 373 K with values typically increasing further away from ambient conditions. A best-fit linear correction, equation 5.7, was found and applied to the factory-calibrated values. Factory and laboratory calibration values are plotted in Figure 5.13a and whilst both follow the ideal trendline, $T_w = T_{bb}$, it can be seen that the laboratory calibration is consistently more accurate. This is further demonstrated by plotting the absolute error, defined as the absolute temperature difference between calibrated (T_{bb}) and actual blackbody temperature (T_w), in Figure 5.13b. The factory calibration error is between 0.75 and 2.9 K, with most deviation at sub-ambient temperatures. The laboratory calibration is significantly better with values less than 0.50 K for all temperatures greater than 283 K. The greatest error in the laboratory calibration is 0.91 K at 273 K, a temperature unlikely to occur in the MERC facility.

$$T^4_{bb,\,lab} = 0.9902\, T^4_{bb,\,factory} - \frac{5.215}{\sigma_{bb}} \tag{5.7}$$



(a) (b)

Fig. 5.13 (a) Blackbody emitter set surface temperature against calibrated blackbody temperature for factory and laboratory calibrations. (b) Blackbody emitter set surface temperature against absolute error($|T_w - T_{bb}|$) of calibrated blackbody temperatures.

5.3.2 Greybody calibration

Test components within the cascade are coated with a layer of high-emissivity matt black paint. Despite this they cannot be considered blackbodies as the surface emissivity is less than unity. Furthermore, there is a ZnSe IR window in the optical path between IR camera and target surface which has a transmissivity less than unity. A secondary calibration to account for greybody parameters in equation 5.6 is therefore required and can be split into two parts concerning: direct emission from the target surface; and radiance from external sources, namely the reflected radiance emitted from the surroundings and the radiance emitted by the IR window. As in the blackbody calibration, transmittance of air over small distances is approximately unity and so terms related to emission and absorption of the atmosphere and working gas are ignored.

The first part of the greybody calibration requires accurately determining the surface emissivity of test components, ε, and transmissivity of the optical path, τ. This can be achieved by performing an isothermal calibration in-situ. Under isothermal conditions $T_\mathrm{s} = T_\mathrm{opt}$ and so equation 5.4 can be rearranged as in equation 5.8 to find the product $\tau\varepsilon$ from a known surface temperature and calibrated blackbody temperature measured by the IR camera.

$$\tau\varepsilon = \frac{T_\mathrm{bb}^4 - T_\mathrm{s}^4}{T_\mathrm{w}^4 - T_\mathrm{s}^4} \tag{5.8}$$

The second part of the greybody calibration is performed on a per-run basis by grouping external sources of radiance, q_{ref} and q_{opt}, into a surrounding radiance term, q^*, defined in equation 5.9. This surrounding radiance term is normalised by $(1-\tau\varepsilon)$ in order to make it independent of greybody parameters under isothermal conditions. This term can be determined, and surface temperature measurements corrected to give true T_w values, from known values of $\tau\varepsilon$ and an accurate surface measurement obtained during experiments from surface-mounted thermocouples. Details of how each of these methods were employed in the present study are now described.

$$q^* = \frac{\sigma_\mathrm{bb}[(1-\varepsilon)\tau T_\mathrm{s}^4 + (1-\tau)T_\mathrm{opt}^4]}{(1-\tau\varepsilon)} \tag{5.9}$$

Determination of transmissivity-emissivity product

The emissivity of test components depends highly on the type of surface, type of coating being applied and application method. Likewise, the transmissivity of the IR window is highly dependent on material, thickness, and anti-reflective coating. It is therefore advantageous to be able to directly measure these quantities in-situ for a specific application to minimise uncertainty in the greybody correction. Rather than measuring both quantities individually, it is usual to instead determine the product $\tau\varepsilon$ to reduce the overall uncertainty that would be introduced by using multiple calibration methods. It will be shown later that when using a combined surrounding radiance term only the product $\tau\varepsilon$, instead of individual parameter values is required to fully correct for greybody parameters.

An in-situ isothermal calibration can be carried out to determine $\tau\varepsilon$ from equation 5.8. This involves using a test component, or surrogate high-conductivity plate, that has been coated with the same matt black paint and instrumented with thermocouples to give accurate surface temperature measurements. The component is then heated outside of the facility to a uniform temperature, using a heated water bath or otherwise, before being quickly installed into the facility. The IR camera, mounted as in the experimental case, and thermocouples are then used to continuously record T_{bb} and T_{w} respectively as the component cools down. As the surroundings and IR window remain isothermal at ambient temperature this allows for $\tau\varepsilon$ to be measured at a range of surface temperatures.

Whilst this method could, in theory, be applied within the MERC facility by using the instrumented front seal components, the relative complexity of assembling the blade cassette and delicate nature of exposed thermocouples (within the front seal) made it impractical. Also due to time and resource constraints of the project it was not possible to manufacture, instrument, and test a surrogate flat plate which could be fitted into the cascade instead of the blade cassette. It was however possible to use the same blackbody emitter as in the previous blackbody calibration to provide an in-situ known surface temperature whilst keeping the surroundings and window isothermal.

The main drawback of using the emitter is that the surface emissivity is higher than the value expected of the matt black spray paint used on components. The calculated value of $\tau\varepsilon$ will therefore be higher than that of the coated test components. The emissivity of the emitter is accurately known however so τ can be deduced. It may be possible to have a custom emission surface made for the blackbody emitter and control system which is coated with the same paint as used on test components. This would have significant advantages over observing a cooling plate as it would allow for easily repeatable, fast, and highly accurate calibrations to be performed over a wider range of surface temperatures.

An assumed value of ε for the matt black paint can then be used with τ in the greybody calibrations. Relying on tabulated values of emissivity is non-ideal due to the uncertainty in the given values. The matt black paint, which was used on all components throughout the work conducted in this thesis, has been used extensively in prior work conducted in the laboratory and so using an assumed value of $\varepsilon = 0.96$ was deemed acceptable. This value is also consistent with the range indicated on the paint data sheet supplied by the manufacturer (0.92–0.98). It is also worth noting that ε and τ are assumed uniform so that a single value of $\tau\varepsilon$ is used for all surface positions. This assumption is reasonable as a uniform coating of paint is applied and platform surfaces are flat and normal to the camera meaning that viewing angle effects on emissivity should be negligible.

The in-situ calibration was performed in a similar manner as the previously described blackbody calibration. The cascade blade cassette was removed and the blackbody emitter installed such that the emission surface was approximately level with the blade platform endwall. Blackbody temperature was varied from 273 K to 373 K in 5 K increments as before. The IR camera was mounted as during experiments and recorded the emitter through the IR window. The view of the camera is shown in Figure 5.12b. The laboratory blackbody calibration was applied and an average temperature was calculated for each temperature increment from the area within the dashed rectangle in Figure 5.12b (the two patches of relatively low temperature, visible in the rectangle, were not included in the average as they were caused by grease on the window, they line up with the coolant exhaust holes in the blade tip). The maximum standard deviation of values in the averaged region was 0.53 K at 373 K.

Values of τ calculated for each temperature increment are plotted in 5.14. At temperatures close to ambient (296 K) the solution of equation 5.8 breaks down. For temperatures greater than 300 K the calculated τ is approximately constant. The mean value of $\tau\varepsilon$ from measurements greater than 300 K was 0.9234 with a standard deviation of 0.0050. Using an assumed matt black paint emissivity of 0.96, the product $\tau\varepsilon$ used throughout this thesis was taken as 0.8944 which is comparable with past values of similar experimental set ups [47].

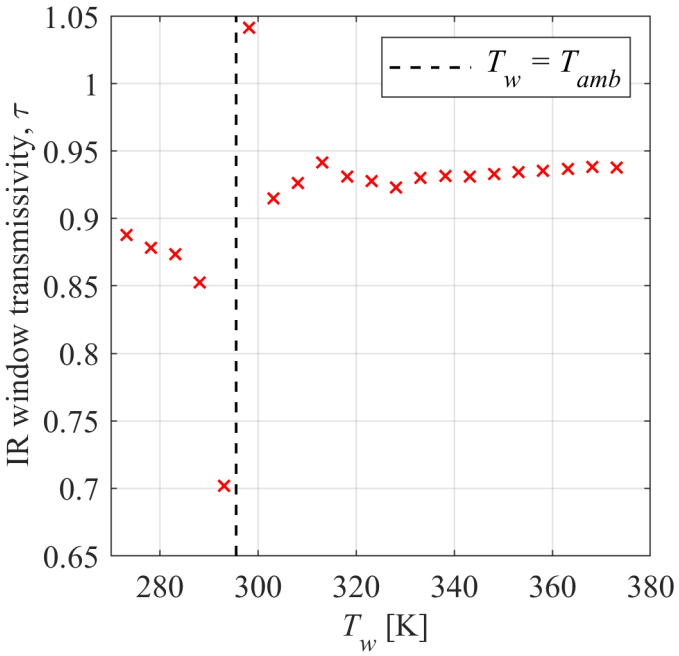


Fig. 5.14 IR window transmissivity, τ, calculated from in-situ blackbody emitter surface temperature and laboratory calibrated IR camera measurements.

Surrounding radiance correction

The surrounding radiance, q^*, within the cascade varies continuously throughout a run, as the temperature of surfaces surrounding the test components change, and is dependent on the initial temperature state of the facility. The test section and all test components are coated with a high-emissivity matt black paint to reduce internal reflections, particularly that of radiation emitted by the hot heater mesh. The surrounding radiation reflected from the black paint coated surface of interest is expected to be small compared to the overall radiation incident on the camera and therefore the majority of the q^* radiance is expected to be due to the IR window heating up during the run. A typical value of q^* within the cascade is 500 W sr^{-1}m^{-2} which is

equivalent to approximately 8% of the total radiation incident on the camera, or a temperature difference of up to 2.5 K. In order to accurately extract surface temperature measurements from the camera blackbody temperature it is therefore necessary to measure the surrounding radiance throughout each individual run and correct for it. Defining q^* as in equation 5.9 we can rearrange equation 5.4 into the form given in 5.10.

$$q^* = \frac{\sigma_{bb}(T_{bb,TC}^4 - \tau\varepsilon T_{w,TC}^4)}{(1-\tau\varepsilon)} \tag{5.10}$$

Whilst variation in surrounding radiance is significant over time, spatial variation across the flat platform surfaces of interest is expected to be negligible. We can therefore apply a surrounding radiance correction determined at any single location, at a given time, to the entirety of the visible surface. Two thermocouple assemblies, embedded in the front seal overhang, are visible in the camera field of view and these are used to monitor surface temperature throughout an experiment. The surrounding radiance term in the cascade is calculated, using equation 5.10, by comparing the thermocouple temperature measurement to the calibrated camera blackbody measurement at these locations ($T_{w,TC}$ and $T_{bb,TC}$ respectively).

Using the methods described above, $\tau\varepsilon$ and q^* can be determined for any surface location at all times during a run. Surface temperature, at any location and at any time, can thus be determined by substituting equation 5.10 into equation 5.4. The final greybody correction to determine surface temperature is given as equation 5.11. This was implemented for all IR measurements throughout this thesis. The corrected surface temperature is then used to calculate values of adiabatic film effectiveness, covered within the journal paper in Chapter 6, and metal effectiveness described below.

$$T_w(x,t) = \sqrt[4]{\frac{\sigma_{bb}T_{bb}^4(x,t) - (1-\tau\varepsilon)q^*(t)}{\sigma_{bb}\tau\varepsilon}} \tag{5.11}$$

5.4 Determination of metal effectiveness

Metal effectiveness distributions across the two visible rotor platforms are calculated from corrected surface temperature distributions and measurements of mainstream and coolant temperature via equation 5.12. As the cascade flow is compressible, the recovery temperature, T_R, is used as this is the relevant driving gas temperature at the platform surface. The recovery temperature is a function of the mainstream temperature and is determined simultaneously with metal effectiveness by using measurements from multiple experiments performed at different temperature ratios but approximately the same aerodynamic conditions.

$$\theta = \frac{T_{\mathrm{R}} - T_{\mathrm{w}}}{T_{0\mathrm{m}} - T_{0\mathrm{c}}} \tag{5.12}$$

5.4.1 Recovery temperature

In compressible flows, a temperature variation exists within the boundary layer, in addition to the existing thermal gradient, due to viscous dissipation. This alters the conductive and convective heat transfer within the boundary layer and results in a gas temperature at the surface, driving heat transfer, that can be either higher or lower than the gas temperature in an equivalent incompressible flow. The difference depends on the gas properties and whether the surface is being heated or cooled. For experiments with air ($\mathrm{Pr} < 1$) heating a cooled surface, the recovery temperature is lower than the mainstream total temperature. The recovery temperature is a function of local Mach number:

$$T_{\mathrm{R}} = T_{0\mathrm{m}} \frac{1 + r\frac{\gamma-1}{2}\mathrm{M}^2}{1 + \frac{\gamma-1}{2}\mathrm{M}^2} = R_{\mathrm{c}} T_{0\mathrm{m}} \tag{5.13}$$

where r is the compressible recovery factor, related to the Prandtl number of the gas, commonly taken as $r = \mathrm{Pr}^{1/2}$ and $r = \mathrm{Pr}^{1/3}$ for laminar and turbulent boundary layers respectively [45]. For a particular environment with fixed Mach number distribution and defined boundary layer state, the recovery temperature is effectively a function of mainstream temperature and surface position. We can therefore simply state the recovery temperature as $T_{\mathrm{R}} = R_{\mathrm{c}} T_{0\mathrm{m}}$ where

R_c is the compressible mainstream recovery ratio. The compressible mainstream recovery ratio is approximately independent of temperature and is primarily a function of the Mach number distribution and state of the boundary layer. It can therefore be matched between experiment and engine and is affected by cooling flows, which alter the boundary layer behaviour.

5.4.2 Determination of compressible mainstream recovery ratio

Direct calculation of R_c, and hence θ, requires knowledge of the local Mach number and state of the boundary layer, which cannot be accurately measured during an experiment. Instead, a series of tests are conducted within the MERC facility at identical flow conditions but with varying mainstream-to-coolant temperature ratio. From these, values of both compressible mainstream recovery ratio and metal effectiveness are determined. Typically, four runs were conducted for each test condition, with T_{0m}/T_{0c} varying between 1.1 and 1.17. Changes to the aerodynamic field over this small range of temperature ratios is expected to be negligible. Both R_c and θ are approximately independent of temperature and therefore distributions should be identical between each of the four runs.

An algorithm, described mathematically by equation 5.15, determines the compressible mainstream recovery ratio distribution on a per-pixel basis by finding the value which minimises the standard deviation in metal effectiveness across each set of four runs. That is, a value of R_c is initially guessed for a given pixel and equation 5.14 used to calculate θ for that pixel from each of the four runs. The standard deviation and mean metal effectiveness is then calculated and stored and R_c incrementally varied. The minimum standard deviation calculated is then used to determine the correct value for R_c and the corresponding θ from the mean value of the 4 runs. This process is repeated for every pixel and results composited to give a complete surface distribution.

$$\theta = \frac{R_c T_{0m} - T_w}{T_{0m} - T_{0c}} \tag{5.14}$$

$$\underset{R_c \in [0,1]}{\arg\min} \sqrt{\frac{1}{N} \sum_{i=1}^{N} (\theta_i - \overline{\theta})^2} \tag{5.15}$$

An advantage of this process is that any irregularities in the mainstream temperature measurements will be mitigated, as the determined value of R_c will implicitly correct for them when calculating θ. For example, the inlet thermal boundary layer present in the cascade means that there is a temperature gradient in the inlet flow. Streamlines from different parts of the inlet thermal boundary layer, with different temperatures, will be convected downstream and interact with positions on the platform surface. This means that the local driving temperature will vary across the platform, independently and in addition to the compressible recovery effect. Essentially, the measured mainstream temperature (taken as the average of the inlet temperature rakes) will be hotter than the actual temperature of the flow that interacts at the platform surface and this will lead to higher inferred metal effectiveness values. However, due to the process used to determine R_c this will be corrected for as the value of R_c will simply be found to be lower than for an identical case with a perfectly flat inlet temperature profile. The calculated R_c value in this situation is less physically meaningful as it then becomes a single term correcting for multiple phenomena.

A *recovery and redistribution* parameter which performs a similar function to R_c was proposed by Luque et al. [54][55]. They corrected for the compressible recovery effect, as well as heat being transferred and redistributed between mainstream and coolant through metal walls in NGV film cooling systems. For the metal effectiveness studies detailed in this thesis, the absolute values of R_c were of secondary interest to corrected θ distributions. It was therefore deemed acceptable by the author to allow the process of determining R_c to automatically correct, to a large degree, for the inlet thermal boundary layer in these cases.

An example of the standard deviation distribution in θ values is given in Figure 5.15 alongside a histogram of pixel values. These are the minimum standard deviations in θ, which were used to select local values of R_c. Values are below 0.012 across the entire image, and lower values closer to 0.002 are achieved across the majority of the platform. The largest deviations occur along the calibration strip upstream of the blade (top of the image) as this region is uncooled and is characterised by very low metal effectiveness values which, as shown below, have increased measurement uncertainties. The regions of relatively high deviation on the platform correspond

with very high (greater than 0.8) metal effectiveness values which similarly have increased measurement uncertainty.

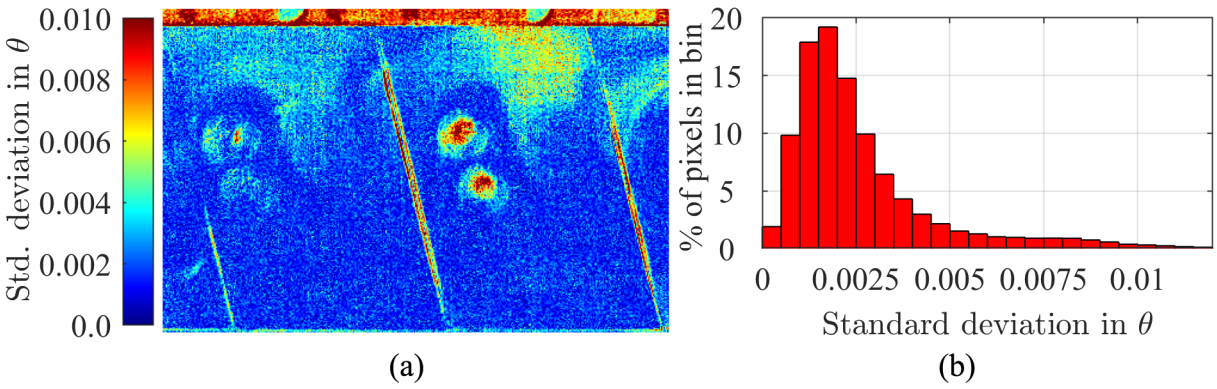


Fig. 5.15 (a) Example distribution of per-pixel minimum standard deviation in metal effectiveness, calculated from 4 experimental runs with T_{0m}/T_{0c} between 1.1 and 1.17. (b) Histogram showing percentage of pixels corresponding to standard deviation in θ values.

Determining the distribution of R_c is important in order to obtain a correct θ distribution. If the mainstream recovery temperature is not used when calculating metal effectiveness, *ie* R_c is assumed equal to unity, then values in regions with elevated local Mach numbers will be incorrect. Calculated metal effectiveness in these regions will be larger than the actual values because the reference temperature of the hot mainstream gas is larger than the actual driving temperature, T_R, and hence the cooling performance on the surface appears better than it actually is. This could be disastrous when used to assess safe operating limits at engine conditions.

An example is given in Figure 5.16 which shows the calculated θ distribution (a) assuming incompressible flow ($R_c = 1, T_R = T_{0m}$), and (b) when R_c is determined as described and T_R correctly used. A histogram comparison of θ values is given in 5.16c and the R_c distribution in 5.16d. The Mach number of the flow increases through the turbine blade passage and hence there are lower values of R_c towards the rear of the platform (bottom of the images). In this region the uncorrected metal effectiveness values can be seen to be up to 0.1 higher than the corrected values, which corresponds approximately to a 50 K difference in engine metal temperatures.

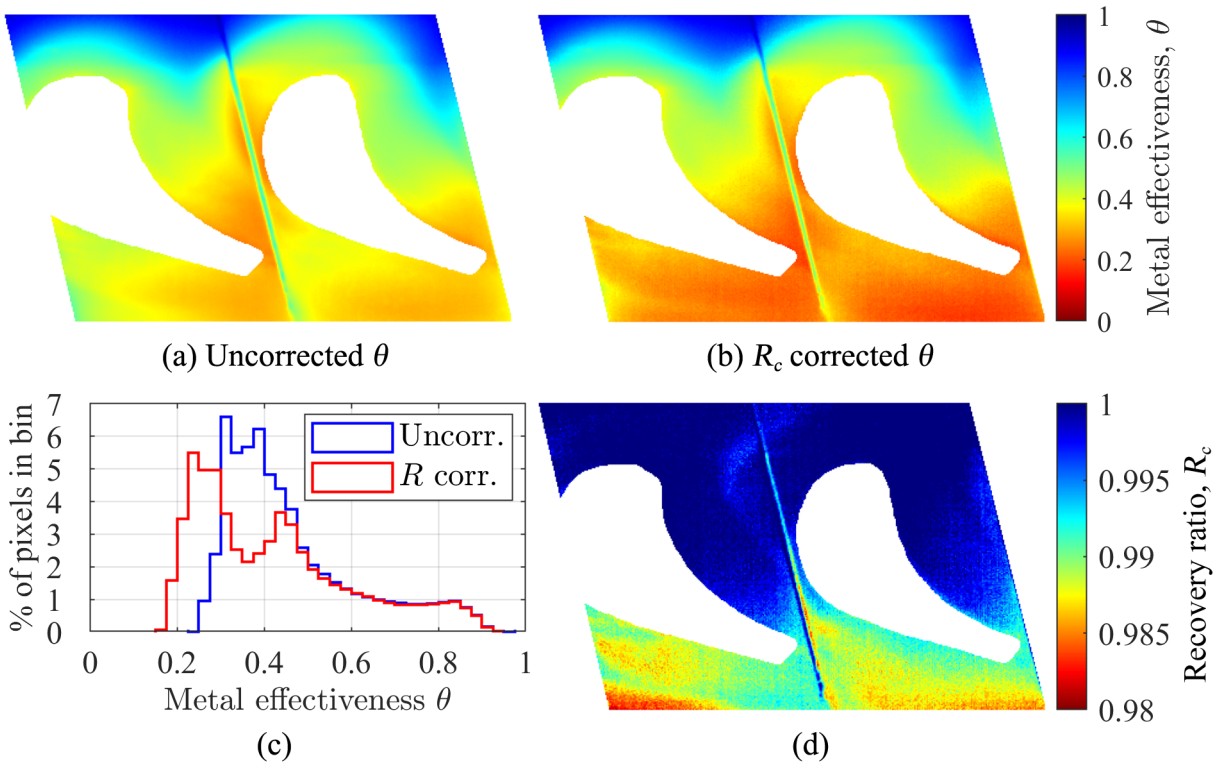


Fig. 5.16 Example surface distributions on two blade platforms of (a) θ calculated assuming incompressible mainstream flow, (b) θ calculated using compressible mainstream recovery ratio, R_c, (c) histogram of pixel θ values for two previous cases, and (d) corresponding R_c distribution.

5.5 Summary of IR image processing

The primary outputs of the MERC facility are the processed infra-red images of the two rotor platforms which allow for surface temperature and effectiveness distributions to be determined. A large number of corrections are required to obtain accurate values for these derived quantities. A summary of the complete image processing method to obtain metal effectiveness distributions is now described and is illustrated in Figure 5.17. The processing method used to obtain adiabatic film effectiveness measurements from corrected surface temperatures differs and is covered in Chapter 6.

The camera output is recorded as a video—of 640×480 pixels at 32 Hz—for the full run duration using the manufacturer *TIM Connect* software which automatically converts the incident radiation on each camera pixel to an equivalent blackbody temperature. A slight delay exists between the start of the experiment and the camera recording (typically in the range of 0.2–

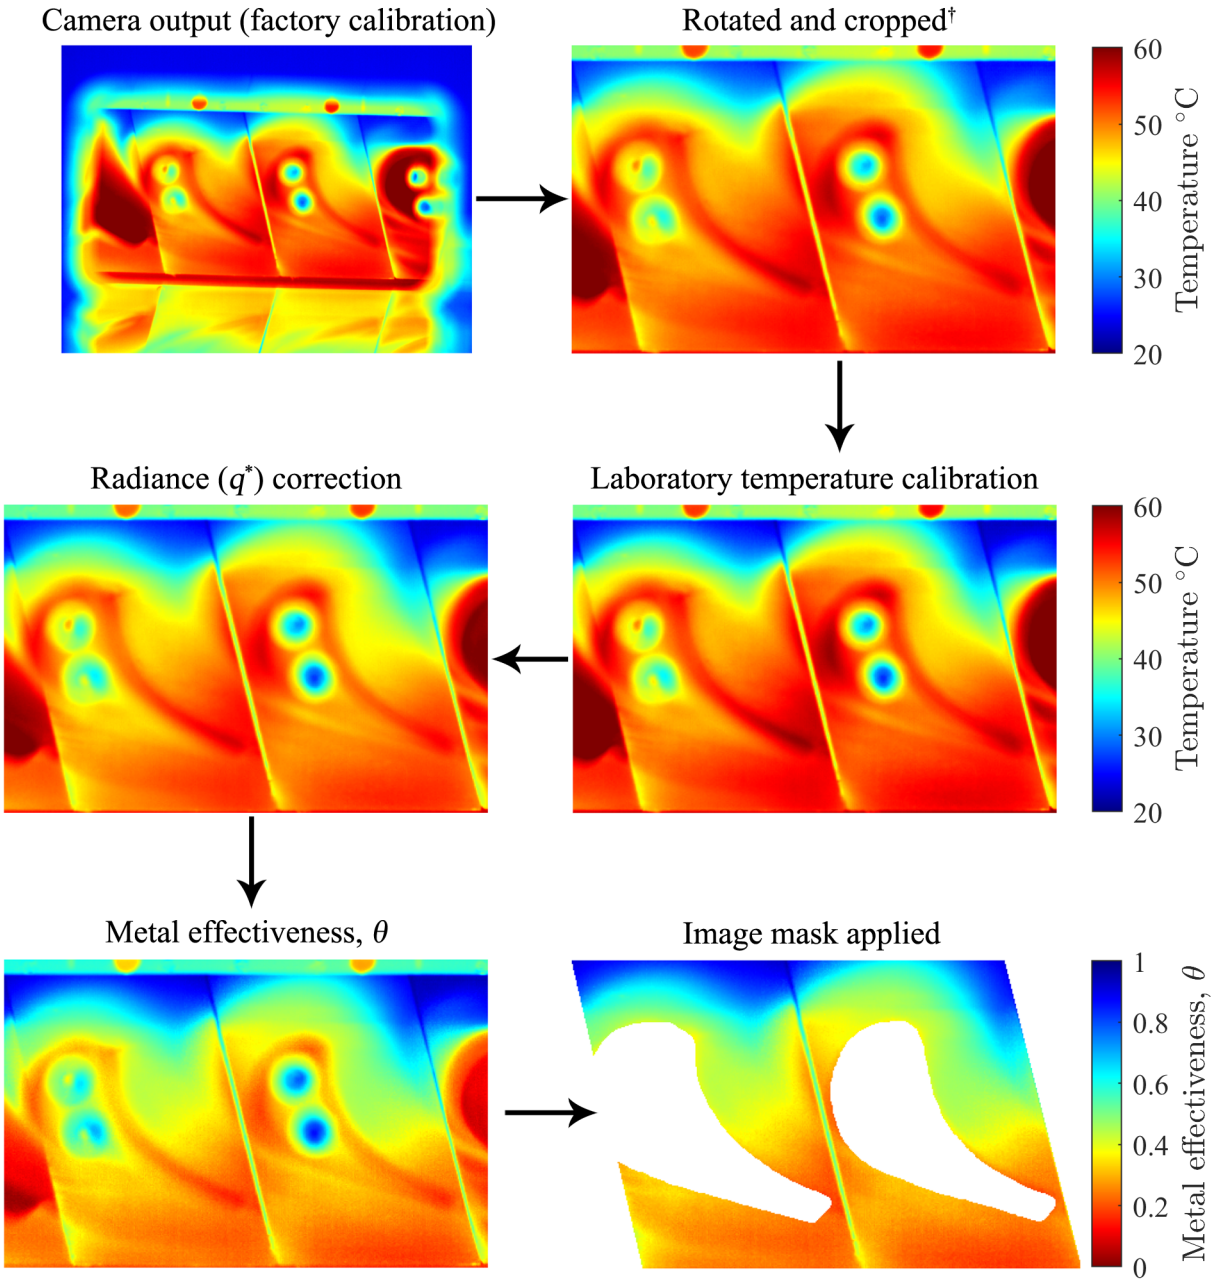


Fig. 5.17 Example of an image, of two blade platforms taken in the MERC facility, being processed through stages of the IR thermography processing routine including application of laboratory blackbody calibration, surrounding radiance correction, and determination of metal effectiveness. †Rotate and crop done first here for presentation clarity, but typically done as final stage to prevent pixel distortion affecting results.

0.8 s). Temperature histories of averaged 5×5 pixel areas are used to sync the video time to pressures and temperatures recorded by the main facility data acquisition system by looking at key observable phenomena such as the initial pressurisation spike and start of the heated period. Snapshot images are then taken from the IR video at known timestamps, and for each image several frames (typically 6) are time-averaged to reduce noise.

Images are then cropped to remove the surrounding aluminium casing and reduce the size, and number of pixels that need processing. As the blackbody calibration drifts over time a laboratory calibration is applied (equation 5.7) over the top of the factory calibration to correct measured blackbody temperatures. The surrounding radiance is then calculated and corrected for (equation 5.11), by examining the difference in surface temperature measured by thermocouples and IR camera at the calibration strip copper patches, to give an accurate surface temperature distribution.

Metal effectiveness and compressible mainstream recovery ratio are then simultaneously determined using processed surface temperature distributions from four runs at different mainstream-to-coolant temperature ratios. An algorithm is used (equation 5.15) to determine the value of R_c by the value for each pixel which minimises the standard deviation in θ. The complete metal effectiveness distribution is then composited together by using the mean value from the 4 runs and R_c distribution.

Finally, image masks are applied to the image to remove the visible portions of the blade lifting surfaces and tip, and surrounding components. The masked images are useful for easily calculating area-averages of metal effectiveness and when comparing cases side-by-side.

5.6 Uncertainty of wall temperature and metal effectiveness

The uncertainty in derived quantities (such as T_w and θ) depends on the uncertainty of the individual measurements which are used to determine them. The expected absolute uncertainty (defined as U_Φ such that for any given function or quantity, Φ, the uncertainty can be expressed as $\Phi \pm U_\Phi$) is estimated here by the root-sum-squares (RSS) method. This method uses a summation of functions of the partial derivatives with respect to each component variable, and

corresponding measurement uncertainty. As an example, the corrected wall temperature can be expressed as in equation 5.16 and the uncertainty can be calculated from equation 5.17.

$$T_{\mathrm{w}}(x,t) = \sqrt[4]{\frac{T_{\mathrm{bb}}^4(x,t) - \left[T_{\mathrm{bb,TC}}^4(t) - \tau\varepsilon T_{\mathrm{w,TC}}^4(t)\right]}{\tau\varepsilon}} \tag{5.16}$$

$$U_{T_w} = \sqrt{\left(U_{T_{\mathrm{bb}}}\frac{\partial T_{\mathrm{w}}}{\partial T_{\mathrm{bb}}}\right)^2 + \left(U_{T_{\mathrm{bb,TC}}}\frac{\partial T_{\mathrm{w}}}{\partial T_{\mathrm{bb,TC}}}\right)^2 + \left(U_{T_{\mathrm{w,TC}}}\frac{\partial T_{\mathrm{w}}}{\partial T_{\mathrm{w,TC}}}\right)^2 + \left(U_{\tau\varepsilon}\frac{\partial T_{\mathrm{w}}}{\partial \tau\varepsilon}\right)^2} \tag{5.17}$$

Due to the method used to determine compressible mainstream recovery ratio, described above, it is difficult to determine the uncertainty of R_{c} and hence the uncertainty of recovery temperature. For this analysis, when calculating metal effectiveness via equation 5.14, it has therefore been assumed that $T_{\mathrm{R}} = T_{0\mathrm{m}}$ and $U_{T_{\mathrm{R}}} \approx U_{T_{0\mathrm{m}}}$. The dominant measurement uncertainty in θ is that of the corrected wall temperature T_{w}, due to the various IR measurement corrections. The additional uncertainty contribution of R_{c} is therefore expected to be relatively small and so estimated uncertainties here should reasonably be representative of the entire blade platform area. All uncertainties are estimated to 95% confidence level.

Measurements of the mainstream and coolant temperatures are taken by K-type thermocouples with an off-the-shelf listed uncertainty of ± 2.2 K.This manufacturer uncertainty is predominantly a bias uncertainty and is typically considered as a maximum (to 99.7% confidence) limit. It is however only relevant in the case of using the thermocouples in an uncalibrated off-the-shelf state and does not account for the thermocouple mounting setup or data acquisition hardware used to record results, each of which introduces additional uncertainty. To offset these uncertainties, all thermocouple measurements are normalised at an isothermal condition (ambient temperature after at least 12 hours of settling time) at the beginning of any experimental test sequence. Temperature ranges during experiments are relatively small (up to 60 K) and test durations are short which minimises the risk of calibration drift during a run. Therefore, by regularly calibrating each thermocouple at this isothermal state the majority of the bias uncertainty is removed. Measurements are also recorded at a high sampling rate of 500 Hz, filtered to remove

high frequency noise, and downsampled to 40 Hz to remove precision uncertainty in the recorded temperature data.

The residual uncertainty is therefore due to differences in individual thermocouple calibrations not accounted for, cold junction compensation drift, and non-linearity and hysteresis in the particular thermocouple calibration taken. Some of these terms are expected to have a fixed uncertainty distribution whilst others should properly be considered to have a distribution proportional to the temperature offset from the normalisation point. A complete end-to-end analysis and calibration of the specific hardware and setup for the MERC facility was not undertaken during the course of this project. Due to the complexity in determining and implementing such contributions to uncertainty, an estimate of ± 0.8 K is taken for all K-type thermocouple measurements ($T_{0\mathrm{m}}, T_{0\mathrm{c}}$, and $T_{\mathrm{w,TC}}$). This value is not considered sufficiently independent to perform a $1/\sqrt{n}$ (where n is the number of thermocouples) reduction. This uncertainty value is taken from Nakos [61] who performed a robust uncertainty analysis on a data acquisition setup that is similar to that used in the MERC facility. They tested two NI terminal blocks to find per-channel uncertainties when measuring a controlled set temperature. They suggest that an uncertainty value of 0.8 K covers the expected bias uncertainty of all channels to a 95 % confidence limit.

Uncertainty in measured wall temperatures (T_{w}), from IR thermography, requires consideration of each of the calibration and radiance correction stages described above. The corrected wall temperature is determined by equation 5.11 or equivalently 5.16 and so the uncertainty can be calculated by equation 5.17. Temperature measurements of the wall surface, T_{bb}, at the point of interest (*i.e.* each pixel on the platform surface) and at the copper patch locations, $T_{\mathrm{bb,TC}}$ are the lab-calibrated blackbody temperatures from the IR camera. The standard error of the lab-calibration was 0.379 K and so the uncertainty to 95% confidence is ± 0.758 K.

The transmissivity-emissivity product ($\tau\varepsilon$) from the greybody calibration was 0.9234 ± 0.010, determined as twice the standard deviation of values. The surface emissivity, and associated uncertainty, of the blackbody emitter ($\varepsilon_{\mathrm{emit}} = 0.99 \pm 0.01$) is different to that of the paint used to coat test components ($\varepsilon_{\mathrm{paint}} = 0.96 \pm 0.04$). The expected overall uncertainty of $\tau\varepsilon$ for experiments is therefore calculated as 0.8954 ± 0.0396.

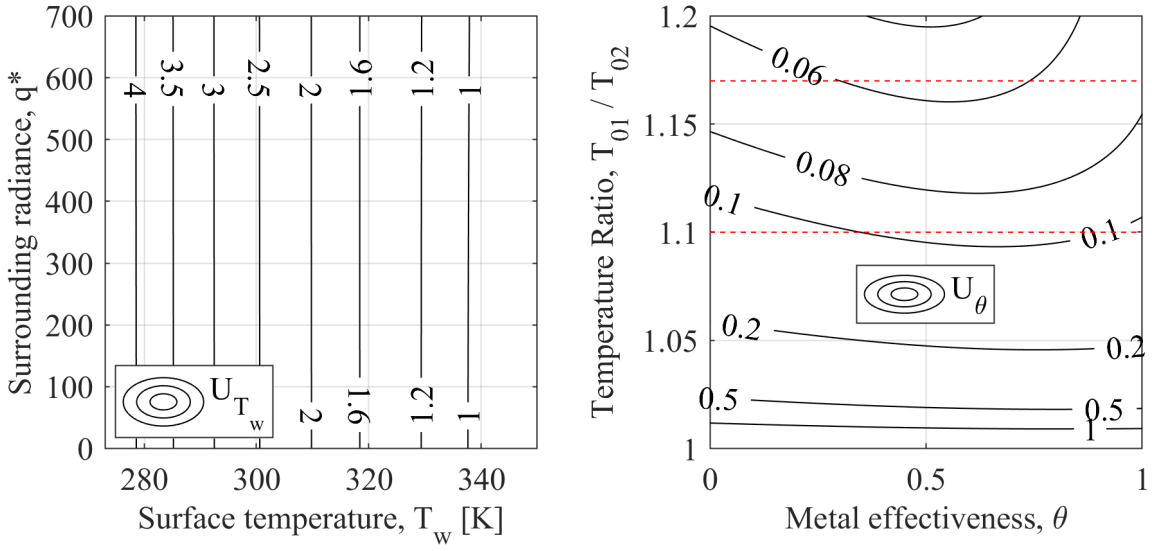
Contours of U_{T_w} as a function of T_w and q^* are plotted in Figure 5.18a. It can be seen from these contours that the uncertainty in wall temperature is insensitive to the absolute value of q^* (vertical contour lines.) The uncertainty decreases as T_w increases, from 3 K at room temperature (293 K) to 1.3 K at 340 K. The dominating contributors to U_{T_w} are the uncertainties of the IR blackbody calibration and emissivity of the matt black paint used to coat test components.

Uncertainty in θ, for given values of T_{0m} and T_{0c}, is dominated by the uncertainty in the corrected wall temperature measurement. Contours of U_θ as a function of θ and temperature ratio T_{0m}/T_{0c} are plotted in Figure 5.18b for a typical cascade surrounding radiance value $q^* = 500$. There is a strong dependence on temperature ratio, with an approximately exponential increase in uncertainty as T_{0m}/T_{0c} approaches unity. At higher temperature ratios (greater than 1.1) the uncertainty towards either limit of θ (0 and 1) is increased. This is because at these limits the individual contributions of T_{0m} and T_{0c} and their associated uncertainty become larger influences on the calculated θ. The uncertainty is also greater at higher values of θ due to the associated lower wall temperatures and therefore increased uncertainty in T_w.

Upper and lower bounds for temperature ratios of the sets of four experiments conducted for each test condition are marked with red dashed lines in Figure 5.18b. For typical minimum and maximum values for θ of 0.2 and 0.8 the uncertainties at the lower temperature ratio bound are 0.11 and 0.095 respectively. At the upper temperature ratio bounds the uncertainties are both 0.062. These values are reasonable as they encapsulate the majority of uncertainty contributions from IR image processing and are not dissimilar from other published works. It should be noted that these uncertainties are highly sensitive to the uncertainty of the K-type thermocouples used to take the majority of temperature measurements. This reiterates two key points: that it is critical to calibrate and remove as much uncertainty from the individual thermocouple measurements as possible; and that increasing the temperature ratio is highly beneficial in reducing uncertainty in calculated metal effectiveness values.

A significant proportion of the estimated uncertainty is likely to be bias uncertainty ‘baked-in’ to the processing method and common to all experimental measurements. For comparative studies, such as those described in this thesis, where the difference in values—rather than the

absolute values themselves—are of interest then the relative uncertainty is expected to be lower than those values calculated here.



(a) (b)

Fig. 5.18 Absolute uncertainty contours for (a) surface temperature, T_w, against surrounding radiance correction, q^* and (b) metal effectiveness, θ, against mainstream-to-coolant temperature ratio.

5.7 Recommendations for improvements to the MERC facility

The MERC facility in its current configuration is a versatile, user-friendly experimental facility that is capable of generating very high quality data for platform cooling investigations. There are however improvements which could be made and a number of brief suggestions of improvements, or checks of the facility which were not completed during the timeframe of the present work, are now outlined.

- Hydraulic pressure testing of the facility demonstrated leaks at higher-than-nominal pressures. At lower pressures there may be air leaks which are undetected and which would alter the actual mainstream and coolant mass flows through the facility. An air leak test could be performed by blanking the facility from the main supply and exhaust line, con-

necting to a low pressure tank of known fill volume, and observing any pressure drop over a large period of time which would indicate slow leaks in the facility. The leak capacity could be determined from this kind of experiment and mass flow rates measured in the facility corrected.

- The IR window has an anti-reflective coating however the effectiveness of this may deteriorate over time. A blackout enclosure could be made around the IR camera and window to eliminate external heat sources from being reflected and registered by the camera.

- The IR camera software is basic and lacking in quality-of-life features. The IR camera view cannot easily be remotely monitored, as the software outputs only a blackscreen when connecting remotely. Recording must also be triggered in an awkward manner using a custom script. The raw count of camera radiance cannot be directly output and so blackbody calibrations must be done on top of the existing factory calibration. These issues may be solvable at a software-level or a new camera model could be used with more advanced features.

- The data acquisition system is limited to 64 pressure transducer measurements during any single run. An upgrade to the hardware to allow for more measurement points would allow for more flexibility in monitoring pressure measurements throughout the facility.

- The current traverse system is, due to the authors lack of *LabVIEW* and electronics expertise, rudimentary. It is functional and performs adequately, however it would benefit from a re-write of the control program and the addition of hard-wired power limit switches to prevent the system from crashing probes into any walls. An overhaul of the *LabView* data logging script is also overdue and these could be combined.

- Improvements could be made to the traverse mounting system to allow for traverses of multiple downstream planes and of the entire span. A second iteration of probes with smaller lengths could be designed for this or the mounting arm modified to allow the

downstream position to be set. In particular, a plane downstream of the rear seal gap would be beneficial.

- The boundary layer profile at the inlet to the test section should be measured. This could be done by designing a secondary traverse system, utilising a flat-headed pitot probe capable of taking near-wall measurements, and replacing one of the upstream probe cassettes with it. If the boundary layer is found to be excessively thick then a new blade cassette could be designed which integrates a boundary layer bleed immediately upstream of the front seal block.

- A second iteration of the blade cassette SAS insert with more instrumentation could be manufactured, allowing for pressure distributions in each seal cavity to be measured. Secondary distribution plates could also be implemented to more evenly distribute coolant within the cavities. Additional specialised blade cassette inserts, which could be interchangeable, to allow for IR camera calibration, blade surface pressure distributions *etc.* to be conducted would also be beneficial to future research programs.

- Finer control over mass flow rates into each blade cassette plenum could be achieved by using an upstream total pressure lower than the current 15 bara. This would reduce the mass flow through each choked hole and would allow for a larger number of holes to be machined on a single plate. Mass flows could then be more accurately set by covering holes as needed. Alternatively, holes of smaller diameter could be used however these may be prone to blockage from pipework particulates.

- The heater mesh DC power supply cuts power after approximately 1 minute at its highest power output. The cause of this was not determined however, assuming the heater mesh design can survive longer heated periods, it would be beneficial to run the facility for slightly longer to improve the steady-state assumption. Efforts to reduce the heat drawn by upstream surfaces, through introduction of internal insulation, would also be beneficial in creating a more uniform inlet temperature profile.

- Due to the strong sensitivity on effectiveness uncertainty to the uncertainty of the K-type thermocouples used to record temperatures, a comprehensive end-to-end uncertainty analysis should be completed for the facility data acquisition system and thermocouple mounting setups, including any extension leads. This could potentially be done by measuring a series of known temperatures from, for example a stirred mineral oil bath, using the same thermocouples as used in experiments and the exact data acquisition system. The results should be processed in an identical way to allow for the residual uncertainty to be quantified for each channel on the NI terminal blocks.

Chapter 6

Method for determining adiabatic film effectiveness in presence of thermal boundary layer

This chapter consists of a preprint journal paper which details a new processing technique for adiabatic film effectiveness measurements. The MERC facility inlet temperature profile is non-uniform and characterised by a time-varying thermal boundary layer, caused by metal component surfaces upstream of the test section heating up during the short heated period of experiments. This non-uniform temperature profile causes traditional processing techniques for adiabatic film effectiveness to overestimate values due to the reference mainstream temperature being higher than the actual gas temperature interacting with the surface of interest. The techniques detailed in this paper correct values of adiabatic film effectiveness to the true values, which would be expected for an uniform inlet temperature profile. Additional corrections are detailed to account for compressible recovery temperature and non-ideal adiabatic behaviour of the target surface.

This paper was written jointly by the author of this thesis and supervisor, Tom Povey. As of the time of writing it has been submitted to journal and is pending review. Nomenclature and reference sections from this paper have been integrated with those of the rest of the thesis.

6.1 Abstract

In this paper we present a new method for determining adiabatic film effectiveness in film-cooling experiments with non-uniform inlet temperature distributions, in particular the situation of an inlet thermal boundary layer. This might arise in a quasi-steady experiment due to loss of heat

from the mainstream flow to the inlet contraction walls, for example. In this situation the thermal boundary layer would be time varying.

Adiabatic film effectiveness is generally normalised by the difference between mainstream and coolant gas temperatures. Most importantly these temperatures are generally assumed to be spatially—and, possibly temporally—uniform at the system inlet. In experiments with non-uniform inlet temperature, the relevant hot-gas temperature for a particular point of interest on a surface is not easily determined, being a complex function of both the inlet temperature profile and the flow-field between the inlet and the point of interest. In this situation, adiabatic film effectiveness cannot be uniquely defined using conventional processing techniques.

We solve this problem by introducing the concept of *equivalent mainstream effectiveness*, a non-dimensional temperature for the mainstream that can be used to represent the thermal boundary layer profile at the inlet plane, or the effective temperature of the mainstream gas—which we refer to as the *equivalent mainstream temperature*—entrained into the mixing layer affecting the wall temperature at a particular point of interest.

By using data from two or more time instants during an experiment, we simultaneously solve for equivalent mainstream effectiveness and true adiabatic film effectiveness, that is, the adiabatic film effectiveness that we would measure in an experiment with both steady and uniform inlet temperature. This is an important transformation because the true adiabatic effectiveness has a clear physical interpretation and is a more *transportable* quantity between systems (comparisons between different experiments, between experiment and CFD *etc.*)

The proposed method is experimentally demonstrated using full-surface infrared thermography measurements of a cooled rotor blade platform, operated in a transonic linear cascade facility at matched engine conditions. Film effectiveness measurements processed in the conventional way suffer from inlet thermal boundary layer effects, rendering them both time-varying (as heat is released to the walls of the tunnel) and, in certain regions, non-physical in magnitude. The proposed technique renders the same data insensitive to time, and everywhere within physically reasonable limits. By demonstrating independence to the particular (time-varying) inlet temperature profile, we demonstrate the advantages of the proposed technique.

6.2 Introduction

Film cooling performance is often characterised using the adiabatic film effectiveness, η, a non-dimensional temperature typically defined by

$$\eta = \frac{T_{01} - T_{\mathrm{w}}}{T_{01} - T_{02}} \tag{6.1}$$

where T_{w} is the surface temperature at a point of interest, T_{01} is the (typically hot) mainstream gas temperature, and T_{02} is the (typically cold) coolant gas supply temperature.

In an adiabatic system the wall temperature can be thought of as being due to mixing of particular mass fractions of hot and cold gas at the surface. In this sense, the adiabatic film effectiveness can be regarded as being directly related to the local mass fractions. More specifically, rearranging 6.1, we write

$$T_{\mathrm{w}} = \eta T_{02} + (1 - \eta) T_{01} \tag{6.2}$$

where η can now be thought of as the mass fraction of cold gas, and $(1 - \eta)$ the mass fraction of hot gas. Strictly speaking, this formulation is only valid in the case of the mainstream and coolant stream having the same specific heat capacity. This is approximately true in the vast majority of experimental situations, and we adopt this simplification in the remainder of this paper. An extension of the arguments to situations where the heat capacities are not the same is possible with some additional complexity in the formulation.

It is common practice to use laboratory experiments at reduced temperature to optimise film cooling systems for engines. One part of this optimisation may be the adiabatic film effectiveness distribution. Provided the temperature ratio is matched, to first order the laboratory measurements will be the same as the expected engine condition and require no further scaling. In cases where the temperature ratio is not matched, and where air is used for both coolant flow and mainstream flow, the scaling can be quite complex (see, for example, [63]). The need to scale arises primarily from a mismatch in the specific heat capacity flux ratio. This problem can be mitigated by using

so-called foreign gas for one stream, typically the coolant stream, to bring the specific heat capacity flux ratio closer to the engine condition. Even if this approach is taken, second-order scaling issues would still need to be addressed [63].

However the experiments are run, and whatever the complexity of the scaling process required, all scaling processes fundamentally rely on the assumption of flow field—including mixing rate—similarity between the two environments (laboratory and engine). The primary driver for this in the case of adiabatic film cooling experiments is the momentum flux ratio between coolant and mainstream. More complex, second-order effects that lead to dissimilarity of the flow-field are discussed in [63]. It is largely true to say that the second-order effects cannot be corrected for analytically, because they change the structure of the flow in a way that tends to be geometry specific. Thus, it is impossible to determine an analytical correction, or experimental correlation for the same.

In a well-designed experiment (that in which sufficient flow-field similarity is achieved), to meaningfully use laboratory-measured adiabatic film effectiveness distributions to infer the equivalent engine property relies on η representing the mass fraction of cold gas. This is equally true regardless of whether significant scaling (mismatched temperature ratio, and no foreign gas) or second-order scaling (matched temperature ratio, or well-designed experiment employing foreign-gas) is required. If this assumption becomes unreasonable the basis for equivalence in the two situations breaks down.

A problem arises when calculating values of adiabatic film effectiveness from equation 6.1 if the inlet temperature profile is non-uniform. An example of this is an inlet temperature profile caused by an inlet thermal boundary layer. In our example we will refer to a wall which is cold relative to the hot mainstream gas. In this situation, the gas temperatures within the thermal boundary layer will be lower than that measured at mid-span. Streamlines at a range of temperatures within the boundary layer are convected into the downstream flow and entrained within the partially mixed-out coolant film that develops on the surface. The average temperature of the flow represented by these streamlines is unknown. It depends on the inlet gas temperature in the core-flow (that which is unaffected by the thermal boundary layer) and the wall temperature. In general, it is also a function of time. The adiabatic film effectiveness no longer represents

the mass fraction of cold gas, and the meaning of the measured film effectiveness becomes ambiguous.

Achieving completely uniform (both spatially and temporally) inlet temperature profiles in experimental facilities is difficult and may reasonably—by inspection of the apparatus—be assumed to affect a significant proportion of heat transfer experiments to some degree. In particular, short duration, blowdown facilities, not dissimilar from the one used to collect data for this paper [64], are likely to have a time-dependent inlet thermal boundary layer, because the upstream surfaces do not reach steady state conditions during the short test period (see also, for example [2], [73], [21]).

The problem of arbitrariness in selecting a hot gas temperature in adiabatic film effectiveness measurements is seldom explicitly discussed in literature. The most common approaches are either using a single temperature measurement from the mainstream flow (normally the core-flow temperature), or taking an average of the inlet temperature profile by using a thermocouple rake. For the case of a cold thermal boundary layer, using the mainstream core-flow temperature would result in inferred effectiveness values greater than the value that would be obtained with an uniform inlet field; using the mean value could result in inferred effectiveness values either higher or lower than the values that would be obtained with an uniform inlet field. Thus, neither approach solves the central problem, and the resulting adiabatic film effectiveness values have ambiguous meaning and cannot easily be transported between environments (for comparison) or used to represent expected engine conditions (with or without scaling).

In this paper we proposed a practical experimental processing technique which restores the direct link between adiabatic film effectiveness and mass fraction in the experimental situation complicated by an inlet thermal boundary layer. We do this using the entire temporal response of the surface to simultaneously solve for local adiabatic film effectiveness, η, and a new local quantity which we refer to as equivalent mainstream effectiveness, η^*. The equivalent mainstream effectiveness can be thought of as a non-dimensional expression for the mainstream flow temperature associated with a particular point of interest. This conceptual temperature could arise as the result of mixing of flow from regions of different temperature (streamlines emanating from different heights within the thermal boundary layer mixing before being entrained into the

partially mixed coolant film), or convection of flow at a single temperature into the film cooling layer.

We demonstrate that adiabatic film effectiveness distributions from a single experiment in which a time-dependent thermal boundary layer is present can be reduced to equivalent distributions independent of the thermal boundary layer, and that these new distributions are therefore transportable to different environments (e.g. to predict engine performance). Complications arising from compressible effects are also briefly considered.

6.3 Related literature

There are a number of well-established experimental techniques for the determination of adiabatic film effectiveness distributions from steady state surface temperature measurements alone. These typically use IR thermography or liquid crystal methods. More advanced techniques, typically using one or more transient experiments, have been used to simultaneously determine both adiabatic film effectiveness and heat transfer coefficient distributions. In almost all these methods, it is assumed that the mainstream temperature is uniform, and the surface is perfectly adiabatic. In more recent work, authors have proposed methods for dealing with non-adiabatic surfaces, and the problem of determining the film effectiveness in the presence of a constant non-uniform inlet temperature field (e.g. that arising from a combustor simulator). We review all these techniques below. It is notable, that—in all the techniques we review—it is a requirement to assume that any inlet temperature non-uniformity is non-varying with time. That is, it would not include the case in which there is a time varying inlet thermal boundary layer. It is the case that almost all facilities suffer from a time varying inlet thermal boundary layer—to some degree—and the purpose of this paper is to present a new method to account for this. In a later section we show that it is possible to do this using a single transient run. The new method is applicable to many well-established experimental techniques, and can be thought of as an additional processing step rather than a fundamentally new method.

Vedula and Metzger [85] proposed a two-test approach using the colour transition of a thermochromic liquid crystal (TLC) coating. In the first test both mainstream and coolant flows

were supplied at the same temperature, with the surface of interest conditioned to start at a uniform temperature different from the temperature of the flows. The transition time of the liquid crystal was used to determine the heat transfer coefficient at each surface location. In the second test, the coolant was heated with respect to the mainstream, and the surface heat transfer coefficient distribution used to calculate the driving surface temperature. From this the distribution of η was calculated. In this method it is necessary to assume that the mainstream temperature is uniform, with no thermal boundary layer, and that the surface is perfectly adiabatic.

A related method was proposed by Ekkad et al. [18]. In a single transient experiment (IR thermography) with heated mainstream, surface images from two timestamps during the transient surface temperature rise were used to set up two simultaneous equations. These were solved at every location on the surface to yield distributions of η and h. Reducing the method to a single experiment was reported to be advantageous in terms of both overall experimental time and uncertainty. The same assumptions about the mainstream temperature and surface must be made as in the Vedula and Metzger [85] method.

More recently, to relax the requirement for perfectly adiabatic surface behaviour, Holgate et al. [41] proposed a technique in which η and h were determined in a single experiment. The proposed method does not require the surface material to be considered semi-infinite nor to have a uniform (in depth) initial temperature distribution. Experiments were characterised by an initial transient phase with increasing surface temperature followed by a period of approximately steady state surface temperature. The test surface was made of a low thermal conductivity material with external film cooling, and unknown external h, and rear side cooling, at an assumed uniform internal h. They calculated η from time-averaged temperatures from the steady state period of the run. The aerodynamic conditions and η were assumed constant throughout the experiment, including the transient phase. A finite-element model and the calculated distribution of η were then used in an iterative procedure to find a converged solution for the external distribution of h (each pixel from an IR camera image).

A number of techniques have been proposed to deal with the problem of spatially varying inlet temperature. To decouple the effects of film cooling on a HP NGV from additional cooling effects due to dilution gases in the inlet profile of a combustor simulator, Holgate et al. [40]

proposed using a series of experiments, both with and without film cooling. The adiabatic film effectiveness when both dilution port flow and film cooling is present was assumed to consist of an apparent effectiveness from the dilution port flow, and an additive film effectiveness due to the film coolant itself. An experiment with dilution port flow and no film cooling was used to estimate the effectiveness due to dilution only and then remove this profile from the film effectiveness in the film and dilution experiment to yield an estimate of the additive film effectiveness.

Other techniques have been used to attempt to account for spatial variation in hot gas reference temperature in compressible flow environments. In the compressible situation, there is a superposed thermal gradient within the boundary layer due to a rise in static temperature as fluid decelerates in the near-wall region. This is often characterised by the local recovery factor, r [45], from which a recovery temperature $T_{01,\mathrm{R}}$ can be defined, which depends on the local total temperature at the outer edge of the viscous boundary layer, and the local Mach number and Prandtl number. Since $T_{01,\mathrm{R}}$ is a function of local Mach number it can vary across a surface and therefore needs to be known at each location in order to determine η. Xue et al. [89] proposed a two-test technique, which allowed for the three unknown parameters η, h and $T_{01,\mathrm{R}}$ to be simultaneously determined. In the method, two tests are run using using a hot mainstream temperature and two different coolant temperatures. Similar aerodynamic conditions must be assumed. They demonstrated the technique in a short-duration transonic blowdown facility with transient inlet Mach number (initial ramp-up to design condition of 0.88) and mainstream temperature (initial peak followed by a gradual drop off). Test surfaces were made of a low conductivity material so that heat conduction was approximately one-dimensional. Measured surface temperature data (IR thermography) from each of the two tests were used as boundary conditions for a one-dimensional semi-infinite conduction model to reconstruct surface heat flux histories. It was assumed that the two tests (at different coolant temperatures) are conducted at similar aerodynamic conditions, so that η and h could be assumed to be unchanged between tests. At each surface location (pixel) a dual linear regression technique was then be used, in which the recovery temperature was incrementally varied to find a minimum difference between the two experiments (at each surface location) in the calculated values of η and h.

For the scaling of metal temperatures (or overall cooling effectiveness, i.e. the conjugate component) in a compressible environment with film cooling, Luque et al [54], [55] demonstrate that conventional definitions of metal effectiveness fail to account for redistribution of heat between regions differently affected by compressibility effects, even when the external recovery factor is accounted for. This is because there is heat redistribution within the thermal network of the part, which means a purely local definition of recovery factor poorly defined. The authors show that both the Mach number distribution of the part and the thermal network of the part combine to give an effective local recovery temperature, which is referred to as the recovery and redistribution temperature. The authors show that for results to be independent of mainstream-to-coolant temperature ratio, it is necessary to define metal effectiveness in terms of a recovery and redistribution temperature, $T_{\mathfrak{R}}$, and associated recovery and redistribution parameter $\mathfrak{R} = T_{\mathfrak{R}}/T_{01}$. Although the analogy to the conventional theory is obvious, these parameters are entirely distinct from parameters used in previous literature. The authors show that in an experimental situation, the local value of $T_{\mathfrak{R}}$ can only be determined from a series of experiments (at least two) in which the flow field is kept approximately constant, but the mainstream-to-coolant temperature ratio is varied.

6.4 Theory

In this section we introduce the concept of equivalent mainstream effectiveness. The schematic shown in Figure 6.1 shows a developing—both spatially and temporally—thermal boundary layer upstream of an ejecting film. In this example we assume a core-flow temperature (outside the thermal boundary layer, which we take as our reference for the mainstream temperature) at a temperature $T_{01}(t)$, and an upstream wall reference temperature, at a particular location $x = x_1$, $T_{\mathrm{w1}}(t)$. Moving away from the wall there is a gradient of temperature between $T_{\mathrm{w1}}(t)$ and $T_{01}(t)$. Flow from within the thermal boundary layer is convected downstream. In this process, streamlines at various initial temperatures mix with each other. Flow from these mixing streamlines is then entrained within the partially mixed coolant film, giving rise to a temperature

distribution on the surface, $T_w(x,t)$. Here, for simplicity, we refer to incompressible conditions in the mainstream flow.

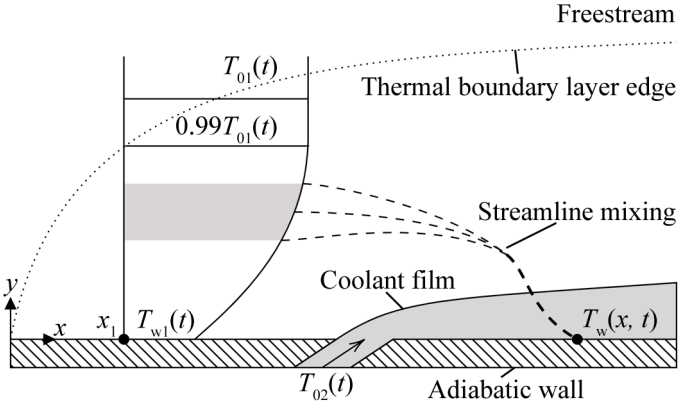


Fig. 6.1 Schematic of a partially mixed coolant film entraining gas from a (potentially time-varying) thermal boundary layer.

The resulting wall temperature $T_w(x,t)$ depends on the local mass fraction of coolant gas, η, the temperature of the coolant gas, $T_{02}(t)$, and the mass fractions and associated temperatures of the entrained regions of hot gas that affect the point of interest on the wall. In our framework, we represent the result of the hot gas streamlines interacting at a point of interest by an equivalent mainstream temperature $T^*(x,t)$ at every point on the surface. In analogy to the adiabatic film effectiveness, we define this temperature in terms of an equivalent mainstream effectiveness by

$$T^*(x,t) = \eta^*(x,t)T_{w1}(t) + (1-\eta^*(x,t))T_{01}(t) \tag{6.3}$$

where, rearranging for $\eta^*(x,t)$, we get

$$\eta^*(x,t) = \frac{T_{01}(t) - T^*(x,t)}{T_{01}(t) - T_{w1}(t)} \tag{6.4}$$

Here $\eta^*(x,t)$ and $1-\eta^*(x,t)$ can be thought of as the equivalent mass fractions of gas (at the point of interest) at the instantaneous temperature of the upstream wall (at the particular location $x = x_1$), $T_{w1}(t)$, and the instantaneous temperature of the core-flow, $T_{01}(t)$, respectively, that would have had to have mixed to give rise to the equivalent mainstream temperature $T^*(x,t)$ at the point of interest. For completeness we note that although this is a simple way of thinking

about it, the analogy is imperfect because the temperature $T^*(x,t)$ arises from the mixing of flow at a range of temperatures between $T_{w1}(t)$ and $T_{01}(t)$. Additionally, we assume that the range of temperature ratio—$T_{01}(t)/T_{w1}(t)$—over which we are operating is sufficiently small that errors introduced by resulting variation (during the test) of specific heat capacity flux ratio are relatively small and can be ignored. Likewise, we assume that there is little change—across the run—in second order correction effects associated with, for example, corrections for compressibility.

Having now defined the equivalent mainstream temperature relevant to each point on a surface, $T^*(x,t)$, we are in a position to re-write equation 6.1 replacing the core-flow temperature, $T_{01}(t)$, with the equivalent mainstream temperature $T^*(x,t)$. We defined the corrected adiabatic film effectiveness

$$\eta'(x,t) = \frac{T^*(x,t) - T_w(x,t)}{T^*(x,t) - T_{02}(t)} \tag{6.5}$$

We effectively have two equations—6.5 and either 6.3 or 6.4—in two unknowns, $\eta'(x,t)$, and—optionally—$\eta^*(x,t)$ or $T^*(x,t)$. The challenge lies in solving simultaneously for $\eta'(x,t)$ and, say, $\eta^*(x,t)$. This requires at least two sets of data.

6.5 Solution of the equations

Consider an experiment in which the mainstream temperature, $T_{01}(t)$, coolant gas temperature, $T_{02}(t)$, and the upstream wall reference temperature at a particular location $T_{w1}(t)$ are all measured. For simplicity we will assume a set of boundary conditions common in short-duration heat transfer testing, namely: constant coolant supply temperature; and mainstream temperature subject to a step change (in time), from the coolant supply temperature to a higher temperature. Thus, in our notation, setting $t = 0$ to be the time at which the step change in mainstream temperature occurs, and considering only the period following this step, we can regard both T_{01} and T_{02} as constant and therefore independent of time. It should be noted that this is not a necessary condition for the technique: it is only necessary that the time constant of the upstream surfaces (i.e. the time constant associated with the temperature deficit in the thermal boundary layer) be much greater than the time constant associated with the temperature rise of the mainstream flow.

The idealised experiment is illustrated in Figure 6.2 for $\eta^* = 0.5$ and $\eta' = 0.5$. In an experiment T_{01} and T_{02} would normally be measured with thermocouple probes, $T_{w1}(t)$ with a wall-mounted thermocouple, and $T_w(x,t)$ with—for example—an infrared camera.

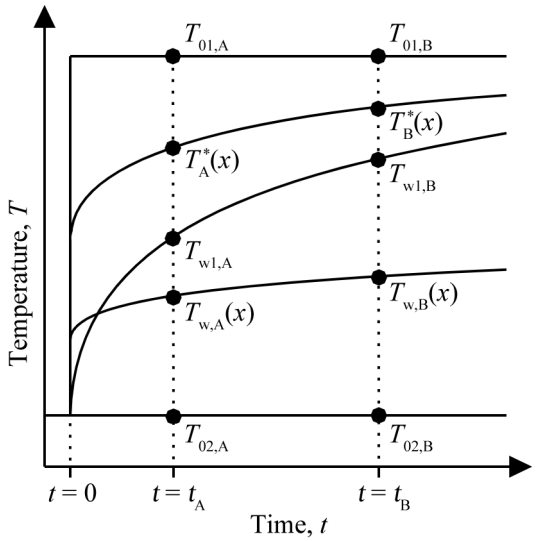


Fig. 6.2 Example of temperature signals from a single transient run, for $\eta^* = 0.5$ and $\eta' = 0.5$, with data selected at two evaluation times at $t = t_A$ and $t = t_B$.

To solve simultaneously for $\eta'(x,t)$ and $\eta^*(x,t)$ (with corresponding $T^*(x,t)$) we require at least two data points. Timestamps $t = t_A$ and $t = t_B$ are marked in the figure; corresponding measured variables have subscripts A and B for compactness. The characteristics shown are typical for transient heat transfer testing, as demonstrated in a later section.

For an unchanged aerodynamic flow field, and accepting the assumptions stated earlier, we can assume that the following hold true

$$\eta^*(x,t) = \eta^*_A(x) = \eta^*_B(x) \tag{6.6}$$

$$\eta'(x,t) = \eta'_A(x) = \eta'_B(x) \tag{6.7}$$

That is, we drop the time dependence in the equations, and equate values of η^* and η' at timestamps A and B. This is equivalent to assuming that the aerodynamic field is unchanged with time between timestamps A and B (same pathlines and mixing) and that the mass fractions (or, more strictly, the specific heat capacity flux fractions) of mixing gas are also constant in time (both $\eta^*(x,t)$ and $\eta'(x,t)$ can be thought of as representing mixing fractions of fluid).

Taking these assumptions (i.e. equations 6.6 and 6.7) allows us to solve for η^* and η' as follows. Substituting 6.5 into 6.7, we get

$$\frac{T^*_{\mathrm{A}}(x) - T_{\mathrm{w,A}}(x)}{T^*_{\mathrm{A}}(x) - T_{02,\mathrm{A}}(x)} = \frac{T^*_{\mathrm{B}}(x) - T_{\mathrm{w,B}}(x)}{T^*_{\mathrm{B}}(x) - T_{02,\mathrm{B}}(x)} \tag{6.8}$$

Then substituting for $T^*_{\mathrm{A}}(x)$ and $T^*_{\mathrm{B}}(x)$ in 6.8 using 6.3, and simplifying using 6.6, gives equation 6.9, which can be rearranged to give an expression for $\eta^*(x)$, equation 6.10. To explicitly evaluate $T^*(x,t)$ at any time instant we can substitute $\eta^*(x)$ into 6.3, and thus, using equation 6.5 we can evaluate $\eta'(x)$.

$$\frac{\eta^*(x)T_{\mathrm{w1,A}} + (1-\eta^*(x))T_{01,\mathrm{A}}(x) - T_{\mathrm{w,A}}(x)}{\eta^*(x)T_{\mathrm{w1,A}} + (1-\eta^*(x))T_{01,\mathrm{A}}(x) - T_{02,\mathrm{A}}(x)} = \frac{\eta^*(x)T_{\mathrm{w1,B}} + (1-\eta^*(x))T_{01,\mathrm{B}}(x) - T_{\mathrm{w,B}}(x)}{\eta^*(x)T_{\mathrm{w1,B}} + (1-\eta^*(x))T_{01,\mathrm{B}}(x) - T_{02,\mathrm{B}}(x)} \tag{6.9}$$

$$\eta^* = \frac{T_{\mathrm{w,A}}(x)\,(T_{01,\mathrm{B}} - T_{02,\mathrm{B}}) - T_{\mathrm{w,B}}(x)\,(T_{01,\mathrm{A}} - T_{02,\mathrm{A}}) + T_{01,\mathrm{A}}T_{02,\mathrm{B}} - T_{01,\mathrm{B}}T_{02,\mathrm{A}}}{(T_{01,\mathrm{B}} - T_{\mathrm{w1,B}})\,(T_{\mathrm{w,A}}(x) - T_{02,\mathrm{A}}) - (T_{01,\mathrm{A}} - T_{\mathrm{w1,A}})\,(T_{\mathrm{w,B}}(x) - T_{02,\mathrm{B}})} \tag{6.10}$$

6.6 Experimental demonstration of technique

We now demonstrate the proposed technique using experimental data from the Metal Effectiveness Rotor Cooling (MERC) facility [64] at the University of Oxford. The facility is a high-pressure transonic blowdown linear cascade designed to investigate turbine rotor platform cooling. A modular working section allows rapid interchangeability of blades and the study of both metal and plastic parts. For the data presented in this paper, the facility was configured to measure the adiabatic film effectiveness arising from leakage flows from hub seals. To do this very low conductivity parts were used.

6.6.1 Facility overview and instrumentation

For the tests described, the modular working section was built with a linear cascade of rotor blades manufactured from glass-filled nylon. The blades included platform, shank and fir-tree root, to ensure representative flow in the cavity beneath the blades. The platforms are cooled

with leakage flows both at the rotor-stator interface, and the so-called slashface between pairs of rotor blades. This is shown schematically in Figure 6.3.

During experiments the mainstream flow is heated using a fast-response heater mesh ($\sim$150 kW power) giving a near step-change (in time) in mainstream temperature of 53 K and a corresponding mainstream-to-coolant temperature ratio of approximately 1.2 which is stable throughout the heated period. A typical core-flow temperature trace $T_{01}(t)$ is shown in Figure 6.4. In this run, the heater is turned on at approximately 40 s, and turned off at approximately 70 s. The temperature spike at approximately 10 seconds is associated with the initiation of flow and pressurisation of the facility. The core-flow temperature was measured using 5 thermocouple rakes, each with 6 K-type bare-bead thermocouples (wire diameter 0.13 mm; approximate time constant, $\tau = 0.15\ s$). These are shown in Figure 6.3. The core-flow temperature rise has a measured time constant of approximately 0.17 s; this is very close to the thermocouple time constant and the true time constant is likely lower (a theoretical estimate gives 0.01 s).

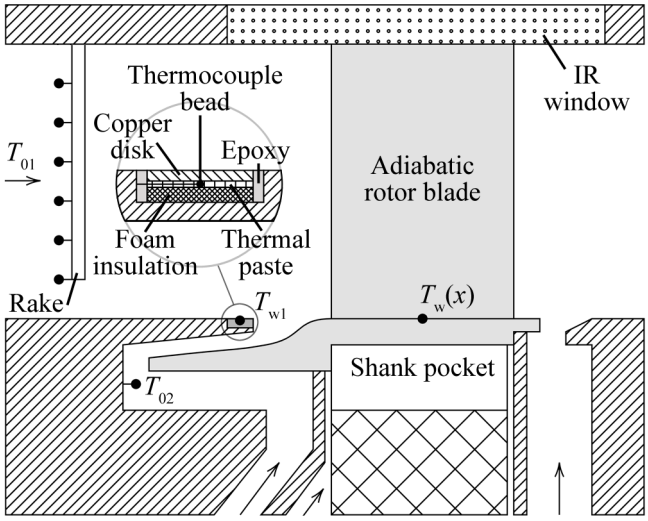


Fig. 6.3 Schematic cross-section of the linear cascade working section showing the locations of measurements: $T_{01}(t)$, the core-flow; $T_{w1}(t)$, the upstream wall reference temperature; $T_{02}(t)$, the coolant flow; $T_w(x,t)$, the platform surface.

The upstream wall reference temperature, $T_{w1}(t)$, was measured using an assembly comprising a 0.5 mm thick copper disk 5 mm in diameter, an underside K-type thermocouple, and a 1 mm thick insulating layer of rigid foam (Rohacell; $k = 0.03$ W m^{-1} K^{-1}). The purpose of the assembly is to register an accurate wall-gas-temperature measurement independent of

the heat flux into the surrounding conducting surface: the assembly is essentially adiabatic. The time constant of the assembly is estimated to be 0.15 s. The location and construction of the measurement system is shown in Figure 6.3. Further details of the technique, and the corresponding uncertainty analysis, are reported in [59].

A typical response of the upstream wall reference temperature, $T_{w1}(t)$, is shown in Figure 6.4. The time constant of this trend is approximately 2.5 s. The time constants of the inlet components are those for approximately 50 mm thick aluminium surfaces: in the range 150–300 seconds for heat transfer coefficients in the range 500–1000 W m^{-2} K^{-1}. The time constant of the upstream wall reference temperature is related both to the time constant of the inlet components and the boundary layer flow (related to the transit time of gas between the heater and the inlet section). The downstream adiabatic surfaces—$T_w(x,t)$ in Figure 6.4—respond to the core-flow temperature, having similar time constants. This makes evaluating the adiabatic effectiveness very challenging by means of conventional processing techniques. The problem is overcome with the method proposed in this paper.

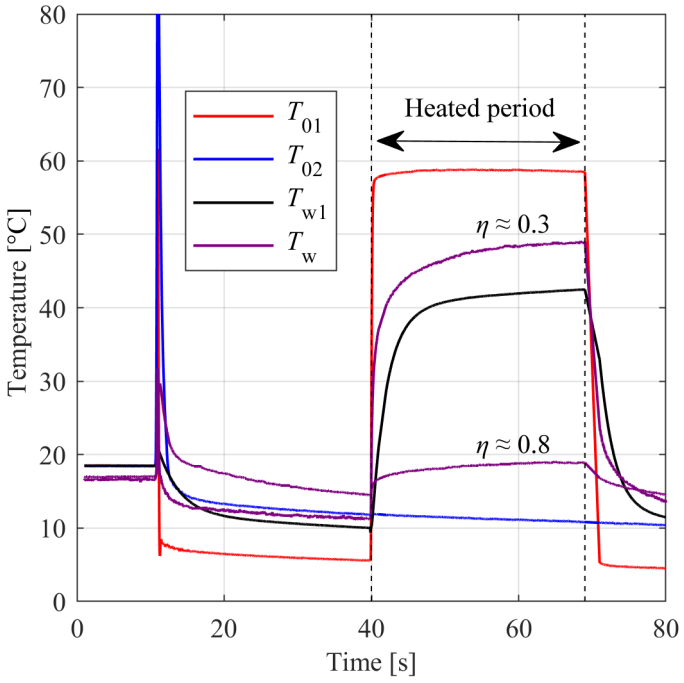


Fig. 6.4 Typical experimental temperature traces for: the core-flow temperature, $T_{01}(t)$; the upstream wall reference temperature, $T_{w1}(t)$; the coolant temperature, $T_{02}(t)$; and the platform surface temperature, $T_w(x,t)$ at two locations with $\eta \approx 0.8$ and $\eta \approx 0.3$.

Coolant outlets are supplied by a common plenum, in which the gas temperatures are measured using K-type thermocouples close to the cavity exit. These are shown in Figure 6.3. The coolant feed temperature and initial temperature of the facility walls are both approximately at ambient conditions, so the individual coolant outlet temperatures are very similar to each other, and vary little with time. A typical coolant temperature characteristic, $T_{02}(t)$, is shown in Figure 6.4.

The platform surface temperature distribution, $T_{\mathrm{w}}(x,t)$, was measured using a calibrated IR camera system. Typical traces from locations with low ($\eta' \approx 0.3$) and high ($\eta' \approx 0.8$) effectiveness regions are shown in Figure 6.4.

Measurements were averaged over 5×5 pixels regions. The blades are manufactured from glass-filled nylon, which has low thermal conductivity (0.25–0.30 W m^{-1} K^{-1}). The platforms were approximately 5 mm thick, rendering them effectively thermally semi-infinite. For an estimated average heat transfer coefficient of 1200 W m^{-2} K^{-1}, the time constant of the surface of interest was estimated to be 0.58 s. This value is in good agreement with the values of 0.41 s and 0.76 s for the sample regions reported in Figure 6.4. Small differences in time constant arise due to local variations in heat transfer coefficient across the surface.

To further illustrate the severity of the processing problem to be overcome, we plot in Figure 6.5 the ratio of the core-flow temperature, $T_{01}(t)$, and upstream wall reference temperature, $T_{\mathrm{w1}}(t)$, as a function of time. Just after the heater is initiated ($\sim 40\ s$) this ratio is approximately 1.15. In the first 5 s this decreases to 1.07. Towards the end of the heated period ($\sim 69\ s$) the ratio reaches 1.05, decaying only very slowly. It is clear that to reach a value of unity would take an impractically long time, beyond the scope of most semi-transient test facilities (common for transonic testing).

In Figure 6.6 we show a typical spanwise temperature profile at cascade inlet (for location, see Figure 6.3) evaluated for the run shown in Figure 6.4, at a time instant just before the heater is turned off ($\sim 69\ s$), using a single thermocouple rake. The hub wall measurement (span fraction of zero) was performed with the adiabatic wall measurement system. Although the spanwise resolution is relatively coarse, the thermal boundary layer can be estimated to extend to a span fraction of approximately 0.2 from either endwall. The central four thermocouples (on each of

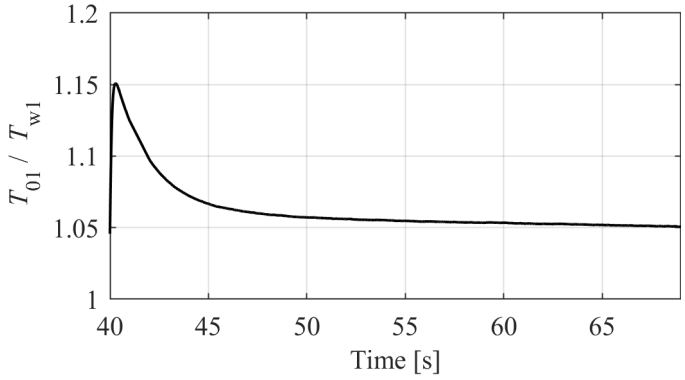


Fig. 6.5 Ratio of the core-flow temperature, $T_{01}(t)$, and the upstream wall reference temperature, $T_{w1}(t)$, as a function of time.

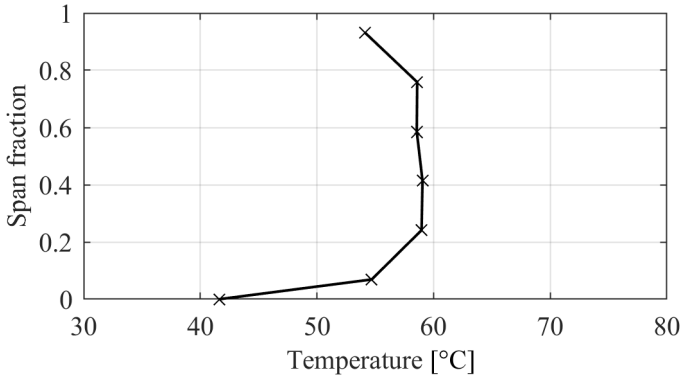


Fig. 6.6 Spanwise inlet temperature evaluated at $\sim$ 69 s for the run shown in Figure 6.4.

the 5 probes) read very similar temperatures allowing the core-flow temperature to be accurately established.

6.6.2 Application of the proposed method

In the presence of an inlet thermal boundary layer, and with a complex aerodynamic field between the inlet measurement plane and the surface of interest, adiabatic film effectiveness distributions evaluated using conventional methods are difficult to interpret on account of arbitrariness in the selection of the hot gas temperature. This makes the film effectiveness data impossible to transport between environments; e.g. to infer an engine condition, or to compare to other experimental data. The technique proposed in this paper is designed to correct the film effectiveness data to the values one would have expected in the absence of an inlet thermal boundary layer.

In this section we take full-surface film effectiveness measurements for a rotor blade platform, and process the data using the conventional method, and the proposed method. We will see

that the conventional method delivers unrealistic—and in certain regions non-physical—results, which are not time-independent. That is, the results depend on the time-varying state of the inlet thermal boundary layer. We then see that the proposed method renders the results independent of time, and corrects non-physical values.

Conventional technique

For comparison with the proposed technique, we first process the data using the conventional method. We arbitrarily define the hot gas temperature, $T_{01}(t)$, as the mean of all the inlet thermocouples. The resulting adiabatic film effectiveness distribution for the run shown in Figure 6.4 is shown in Figure 6.7. The evaluation time was 69 s, at the end of the heated period immediately prior to the heater being switched off. The distribution shows good detail. The low effectiveness strip along the bottom of the image is an aluminium overhang which forms the stator-side of the front rim seal structure and overlaps the front of the blade—this is relatively cool on account of it being highly non-adiabatic. On the same strip we can see two half-circles (C and D) which mark the copper disks for measurement of $T_{w1}(t)$. On these adiabatic patches the measured value of η is approximately 0.30. This is non-physical because there is no upstream coolant ejection (a value of $\eta = 0$ is expected). The dominant features on the main body of the platform are as follows: a well-cooled leading region of the platform; the sweeping of coolant towards the suction surface, caused by the passage crossflow [26]; the introduction of a secondary cooling film towards the rear of the platform as coolant leaks through the blade slashface from the internal cavity; a wedge-shaped region of low effectiveness caused by the trailing edge wake.

To illustrate the difficulty of interpreting the data processed in this way, and also the time dependent nature of the problem, we take the average value of η in regions C, D, E and F in Figure 6.7, and plot the trends $\eta(t)$ during the heated period of the run. These trends are shown in Figure 6.8. E and F mark arbitrary regions of relatively high and relatively low effectiveness, respectively, on the main body of the platform.

We now examine the trends of $\eta(t)$ from regions C, D, E and F, shown in Figure 6.8. The data are plotted as a function of the fraction of the time representing the heated period of the run. Regions C and D mark—approximately adiabatic—surface measurement patches, which are

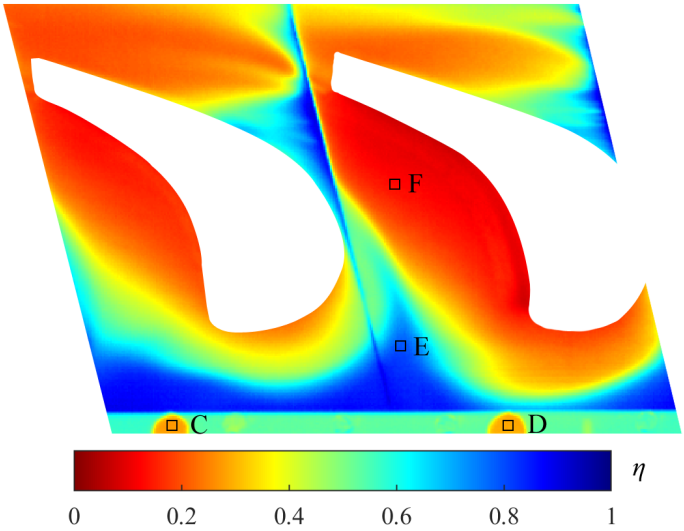


Fig. 6.7 Adiabatic film effectiveness distribution processed using the conventional technique, with $T_{01}(t)$ defined as the mean inlet temperature. C and D mark approximately adiabatic copper disks associated with the wall temperature measurement system. E and F mark arbitrary regions of relatively high and relatively low effectiveness, respectively, on the main body of the platform.

upstream of the cooling flow, and should, therefore have a cooling effectiveness of zero. Although there is good agreement between the two readings, using the conventional processing technique, the reported values of $\eta(t)$ are both non-physical (between 0.79 and 0.33, as opposed to zero) and strongly varying in time. E and F mark arbitrary regions of relatively high and relatively low effectiveness on the main body of the platform. At location E, there is less variation of $\eta(t)$ with time because $T_{\mathrm{w}}(t)$ is dominated by the coolant flow temperature, $T_{02}(t)$, and is therefore insensitive to both the variation in and ambiguity in the definition of $T_{01}(t)$. At F the reverse is true. In this low-effectiveness region, there is significant dependence on $T_{01}(t)$, and the strongly time-varying trend (despite the surface being near-adiabatic) clearly demonstrates the problem in definition of $\eta(t)$. Even the settled value at F, of $\eta \approx 0.3$ is physically unlikely, because this is behind the separation line marked by the horseshoe vortex and is typically completely starved of cooling flow (and effectiveness of zero is expected).

Proposed technique

We now process the same data using the proposed technique in which we simultaneously solve for η^* and η' by using data from a number of timestamps during the run. In the first instance we use two timestamps to show the simplest implementation of the technique.

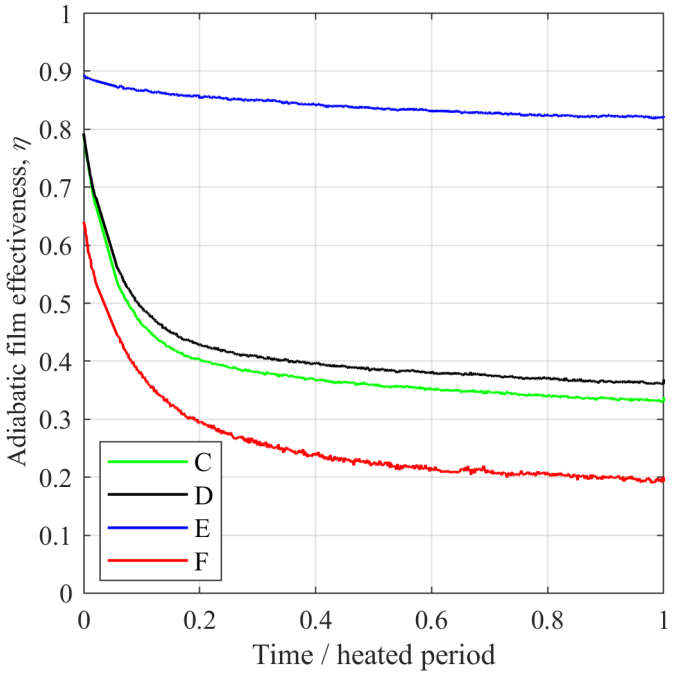


Fig. 6.8 Trends of $\eta(t)$ from regions C, D, E and F in Figure 6.7, using a conventional processing technique with $T_{01}(t)$ defined as the mean inlet temperature.

The data of Figure 6.4 are reproduced in Figure 6.9 for the time period 35 s to 75 s. As an example of the local temperature trends the average temperatures in regions E and F are also shown. The heated period is $t_1 = 39.9$ s to $t_2 = 69.0$ s, a duration of 29.1 s. For comparison with other timestamp combinations, we define a non-dimensional time by

$$t' = \frac{t - t_1}{t_2 - t_1} \tag{6.11}$$

In the current processing we use timestamps at $t_A = 41.0$ s and $t_B = 68.0$ s, or $t'_A = 0.038$ and $t'_B = 0.97$. These are somewhat arbitrary timestamps chosen to give a substantial difference between the two data. We examine the sensitivity to choice of timestep in a later section. The times t_1, t_2, t_A and t_B are marked in Figure 6.4.

Solving simultaneously for η^* and η' across the full surface of the vane, using data at each pixel location of the type given in Figure 6.9, at timestamps $t'_A = 0.038$ and $t'_B = 0.97$, gives the distributions of η^* and η' presented in Figure 6.10. The method used for solution is that presented in the section above entitled solution of the equations.

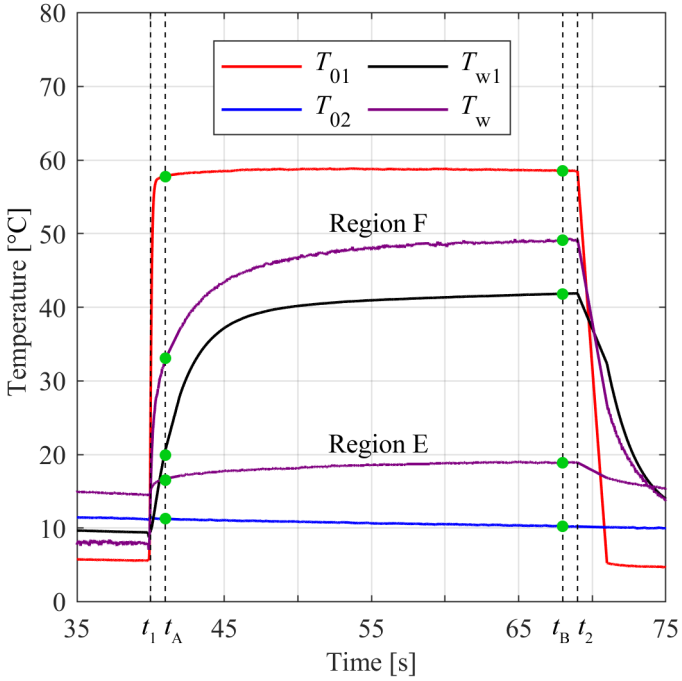


Fig. 6.9 Data from the heated period of a run ($t_1 \leq t \leq t_2$) showing particular timestamps $t_A = 41.0$ s and $t_B = 68.0$ s, or $t'_A = 0.038$ and $t'_B = 0.97$. Wall temperatures, $T_w(t)$, in regions E and F (see Figure 6.7) are also marked.

We first examine the resulting corrected adiabatic film effectiveness, η' (Figure 6.10a). Using the new processing technique, the approximately adiabatic copper patches have an effectiveness close to zero, as expected; we recall that in the conventional processing technique values between 0.79 and 0.33 were returned, which were non-physical. Likewise, in the region behind the horseshoe vortex, close to the suction side of the vane, a wide region of close to zero effectiveness is observed. This is in-line with expectation because the region behind the horseshoe vortex is typically entirely starved of coolant flow. The leading region of the platform, directly downstream of the coolant ejection, and the region downstream of the slashface deeper into the passage have high values of corrected adiabatic film effectiveness ($\eta' > 0.8$) consistent with expectation.

Overall, the distribution of η' is similar to the distribution of η achieved using a mean inlet temperature for $T_{01}(t)$, but values have been corrected downwards. This is to be expected if the choice of mean inlet temperature for $T_{01}(t)$ in the conventional definition of η was unrepresentatively high compared to the actual mainstream fluid temperatures entrained within the mixing cooling layer. As the temperature deficit in the inlet thermal boundary layer is—by definition—concentrated towards the endwalls, it is natural that choice of mean inlet temperature

for $T_{01}(t)$ would generally be an overestimate of the temperature of the entrained gas, and a correction to lower values ($\eta' < \eta$) is expected.

We now briefly examine the distribution of equivalent mainstream effectiveness η^*, presented in Figure 6.10b: our purpose here is not to explain the flow structure in detail, but simply to demonstrate the reasonableness of the profile.

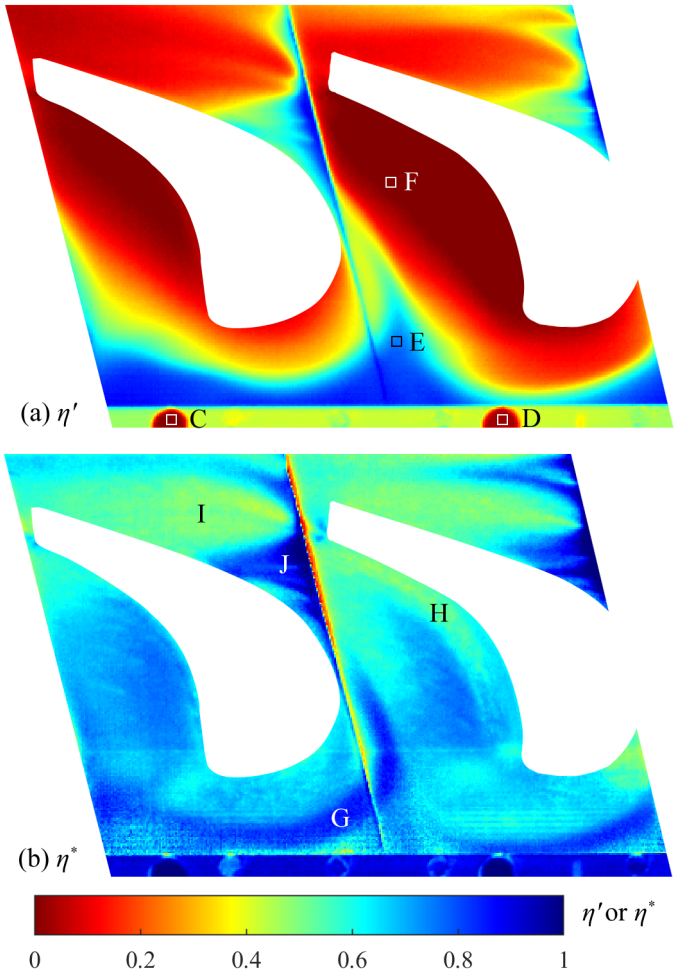


Fig. 6.10 (a) Adiabatic film effectiveness distribution of two turbine rotor blade platforms, calculated using the proposed methodology; (b) corresponding equivalent mainstream effectiveness distributions.

On the approximately adiabatic copper disks (region C, D) the value of η^* is by definition unity. A crescent of high equivalent mainstream effectiveness (region G) can be seen upstream of and wrapping around the leading edge of the blade. Here we have $\eta^* \sim 0.9$. Upstream of this crescent, the lower inlet boundary layer fluid is forced over the coolant flow (leading region of high η' in Figure 6.10a), leading to stratification of the flows. Where the horseshoe vortex

forms, at the position of the crescent, this lower boundary layer fluid is rolled down into the coolant flow, mixing for the first time. Thus, high values of η^* (representing fluid close to the temperature $T_{w1}(t)$) are expected in this region.

Region H, along the pressure side (PS) fillet of the blade have low values of η^* (in the range $0.5 < \eta^* < 0.6$). This is to be expected due to the downwash from the PS of the vane, as low-momentum boundary layer fluid is drawn across the passage (from PS to SS) by the cross-passage pressure gradient. As flow is drawn down the blade it is replaced by fluid from regions of higher span, corresponding to higher temperature and lower values of η^*. Likewise the late SS region, and the region downstream of the trailing edge (marked I) there is a large region of low η^* (in the range $0.4 < \eta^* < 0.6$). Here the flow is highly mixed in the spanwise direction, due both to large scale vortices rolling into the SS corner and off the trailing edge of the blade, and due to turbulence generated by trailing edge shedding.

Region J has higher values of η^* (in the range $0.7 < \eta^* < 1.0$) than might be physically expected. Although this points to instability in the technique, the resulting values of η' are relatively insensitive to η^*, η' taking values close to unity in this region because the near wall flow is dominated by coolant flow leaking through the slashface gap.

To look quantitatively both at the correction η to η', and to look at the physical reasonableness of the values of η^*, we plot histograms of percentage of pixels falling in buckets of values for both the conventional and proposed technique. The same data as presented in Figure 6.7 and Figure 6.10 was used for these images, but the strip extending between and beyond C and D was removed, because this surface is non-adiabatic. The histograms are shown in Figure 6.11. The correction η to η' acts to lower the mean value of effectiveness, with the correction being greater at low values of η than at high values. Lowering the values is expected on account of the average temperature of entrained mainstream gas being overestimated in our implementation of the conventional method. A greater correction at low values of η is expected because in this range the value of η is more sensitive to the temperature of entrained gas.

The uncorrected distribution of η has no values below 0.18, which—alone—suggests an offset error, the cause of which has been discussed. The corrected values η' extend all the way

to zero, giving physically reasonable values in regions where no coolant can reach the surface (as predicted by CFD).

The distribution η^* has most of its weight between $0.5 < \eta^* < 0.7$, with only small tails at high and low values of η^*. Examining Figure 6.6, which shows the thermal boundary layer extending to approximately 20% span, and noting the typical secondary flow patterns and spanwise length scale of flow structures for high-turning blades, this range of η^* is thought to be reasonable. The absence of significant weight in the range $0.7 < \eta^* < 1.0$ is explained by the fact that the thermal boundary layer is forced over the coolant flow in the leading part of the passage, and first interacts with the wall only as it rolls into the horseshow vortex, a process that necessitates moderate mixing. The absence of significant weight in the range $0.0 < \eta^* < 0.5$ is explained by the depth of the thermal boundary layer, and the fact that we are studying only the endwall: that no core-flow, un-mixed with fluid from within the thermal boundary layer, interacts with the endwall is unsurprising as it would require spanwise migration of more than 20%.

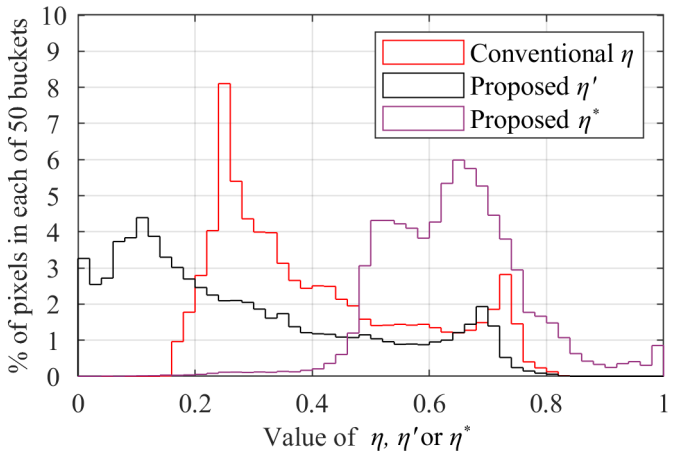


Fig. 6.11 Histograms of percentage pixel counts in each of 50 equal buckets for: conventional adiabatic film effectiveness, η; corrected adiabatic film effectiveness, η'; and equivalent mainstream effectiveness, η^*.

6.6.3 Stability of the process

The aim of the proposed technique is to avoid the sensitivity of the conventional adiabatic film effectiveness, η, to variation with time of the non-dimensional temperature in the inlet thermal boundary layer. In the proposed method, corrected adiabatic film effectiveness, η', should be

independent of time. More precisely perhaps, because the technique relies on the use of at least two timestamps, the value of η' returned should be independent of the choice of timestamps.

In the foregoing analysis, the timestamps A and B were chosen somewhat arbitrarily, but respecting the rather obvious condition that greater sensitivity is achieved by choosing times that allow for significant change in the absolute temperature of the boundary layer fluid. This is a stability requirement for solution of equations 6.9 and 6.10.

To understand the sensitivity to timestamp selection we calculate η' for every possible pair of timestamps (t'_{A} and t'_{B}) taken from the heated period of the run. Results are shown in Figure 6.12 for region E in Figure 6.7.

In the range $0.020 < t'_{\mathrm{A}} < 0.10$ (early start time) and $0.4 < t'_{\mathrm{B}}$ (sufficiently long separation in time to ensure good thermal rise) the value of η' returned is relatively insensitive to timestamp selection, varying between 0.75 and 0.81. For a method using only two timestamps, it is recommended to choose values within this region. For reference, we note that the range $0.020 < t'_{\mathrm{A}} < 0.10$ (or $0.58 < t_A < 2.9$ s) is equivalent to between 0.23 and 1.2 time constants for the inlet thermal boundary layer ($\tau = 2.5$ s) and between 1.0 and 5.0 time constants of the surface of interest ($\tau = 0.58$ s).

As expected, close to the line $t'_{\mathrm{B}} = t'_{\mathrm{A}}$ the solution becomes unstable, because the temperatures $T_{\mathrm{w1,A}}$ and $T_{\mathrm{w1,B}}$ are close in value. Likewise, as t'_{A} increases in region $t'_{\mathrm{A}} > 0.20$ there is significant sensitivity to the timestamp pair due to reducing temperature difference between the two timestamps. At the very start of the experiment, $0 < t'_{\mathrm{A}} < 0.020$, the solution is unstable as $T_{\mathrm{w1,A}}$ is very close in value to that of the coolant, T_{02}: a separation in values between $T_{\mathrm{w1,A}}$ and T_{02} is required in order to distinguish coolant gas from low temperature gas in the mainstream.

6.7 Uncertainty with two timestamps

In this section we perform an uncertainty analysis for the proposed method applied with two timestamps within the heated period. Timestamps are taken from near the end of a run, and within the first two seconds of the heated period. To make the results as generically applicable as possible, we take nominal ranges for parameters that are representative of typical well-

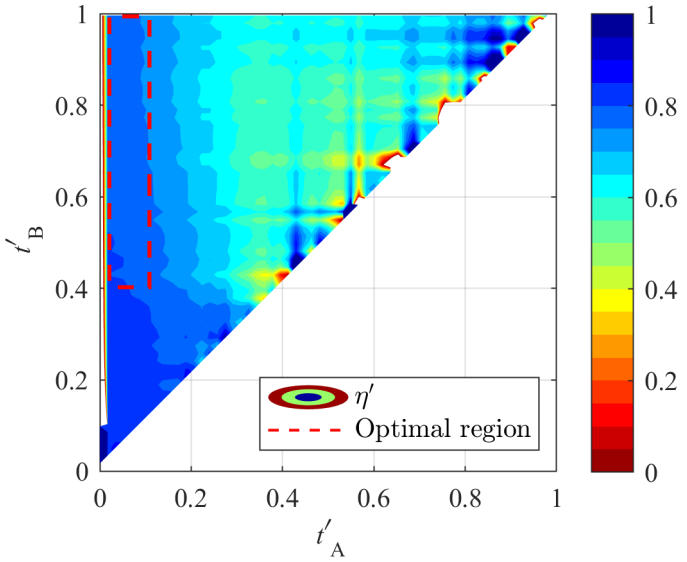


Fig. 6.12 Value of η' returned by the proposed method for region E for every combination of timestamps t'_A and t'_B.

conditioned laboratory experiments, not unlike the experiment of this study. The purpose is to calculate confidence limits in η' and to guide best-practice in application of the proposed technique.

We consider the impact on uncertainty in η' of: mainstream-to-coolant temperature ratio, T_{01}/T_{02}; ratio of upstream wall reference temperature between timestamps, $T_{w1,A}/T_{w1,B}$; and equivalent mainstream effectiveness, η^*. Ranges for each parameter and corresponding estimated uncertainties are presented in Table 6.1. For this analysis, coolant is assumed to be air at room temperature. Core-flow temperature is defined by the ratio T_{01}/T_{02} and is taken as a fixed value for both timestamps. The upstream wall reference temperature, T_{w1}, varies between T_{02} and T_{01}. Surface wall temperature, T_w, is derived from values of η' and η^* covering the range $0-1$.

Measurements taken with thermocouples are assumed to have been normalised at an isothermal condition (normally settling to ambient temperature) and recorded at sufficiently high sampling rate that precision uncertainty in the running (down-sampled) mean becomes negligible. The uncertainty is therefore dominated by bias caused by cold junction compensation drift, differences in individual thermocouple calibrations that are not accounted for, and by nonlinearity and hysteresis in the particular thermocouple calibration taken. Some of these terms can be taken as having a fixed uncertainty (strictly, uncertainty distribution) and others should properly be considered to have a distribution that is proportional to the temperature offset from the normalisation

point. Because of the relative complexity of implementing such a distribution we instead take an estimated upper bound on the uncertainty (i.e. that at the maximum temperature condition) as applicable to the entire temperature range. This upper bound is taken to be ± 0.5 K to 95% confidence. For measurements taken with two (or more) thermocouples we do not consider the measurements sufficiently independent to perform a $1/\sqrt{2}$—or greater—reduction, and retain the ± 0.5 K uncertainty limits.

Table 6.1 Assumed temperature ranges and corresponding estimated absolute uncertainties (to 95% confidence).

Parameter	Units	Value	U(K) 95%	Source
T_{01}	K	293–440	± 0.5	K-type thermocouple
T_{02}	K	293	± 0.5	K-type thermocouple
T_{01}/T_{02}	–	1.0–1.5	–	Derived
T_{w}	K	293–440	± 1.0	IR camera
T_{w1}	K	293–440	± 0.5	K-type thermocouple
$T_{\mathrm{w1,A}}/T_{\mathrm{w1,B}}$	–	0.7–1.0	–	Derived

The uncertainty in T_{w} was taken as ± 1.0 K to 95% confidence, representative of a laboratory-calibrated IR camera setup [46]. This is the combined overall uncertainty including contributions arising from uncertainty in: target surface emissivity; window transmissivity; reflected surrounding radiance; and black body calibration uncertainty.

Expected uncertainties in η' have been calculated by the root-sum-squared (RSS) method. This uses the partial derivatives of the functions used to calculate η^* and η' (equations 6.5 and 6.10) with respect to each input temperature and the associated measurement uncertainty.

6.7.1 Effect of mainstream-to-coolant temperature ratio

We first consider the effect of mainstream-to-coolant temperature ratio on the absolute overall uncertainty in η' (i.e. expressed as $\eta' \pm \mathrm{U}_{\eta'}$). We define an example case, using values from Table 6.1, in which T_{02} is fixed at 293 K and T_{01}/T_{02} is varied between 1.0 and 1.5. We define a surface temperature near the start of the heated period as $T_{\mathrm{w1,A}} = 0.67(T_{01} - T_{02}) + T_{02}$, and a surface temperature near the end of the heated period as $T_{\mathrm{w1,B}} = 0.90(T_{01} - T_{02}) + T_{02}$. These expressions are selected to give optimal timestamps for a two-timestamp method, as discussed in

the next section. For a system in which the mainstream, coolant, and upstream wall reference all start at the same temperature, and in which the coolant temperature is constant during the experiment, the upstream wall reference temperature at any subsequent moment in time will have a temperature offset from the coolant proportional to the temperature difference $(T_{01} - T_{02})$. If the time constants of the system are independent of T_{01}/T_{02} (same aerodynamic field), the expressions for $T_{w1,A}$ and $T_{w1,B}$ above reference the same points in time (i.e. the upstream wall reference temperature trends are scaled versions of each other).

Using these values, we calculate the overall uncertainty in corrected adiabatic film effectiveness, η', as a function of η' and T_{01}/T_{02}. The result is plotted in Figure 6.13 for $\eta^* = 0.90$ (the impact of η^* on $U_{\eta'}$ is discussed in a later section).

The uncertainty in η' is low when the temperature ratio is high (between 0.050 and 0.052 for a typical laboratory temperature ratio of $T_{01}/T_{02} = 1.2$—red line on Figure 6.13—like the current experiment), but increases somewhat exponentially as the temperature ratio approaches unity ($T_{01}/T_{02} = 1.0$). The uncertainty in η' is above 0.10 for $T_{01}/T_{02} < 1.1$, marking a practical lower limit of temperature ratio that would need to be achieved for acceptable uncertainty. This would correspond to mainstream heating of approximately 28 K in typical laboratory experiments ($T_{02} = 290$ K). The uncertainty in η' is only very weakly dependent on the absolute value of η'.

6.7.2 Effect of timestamp selection

We now consider the effect of timestamp selection, by considering the absolute overall uncertainty in η' as a function of the ratio of the upstream wall reference temperatures evaluated at timestamps A and B, $T_{w1,A}/T_{w1,B}$. For a given characteristic response, this ratio can be considered a proxy for the times at which the two data are taken. Somewhat arbitrarily, we take $T_{01}/T_{02} = 1.5$ and $T_{w1,B} = 0.90(T_{01} - T_{02}) + T_{02}$. This represents an experiment with significant mainstream heating (e.g. by approximately 150 K from ambient to an absolute temperature of 440 K) and an—assumed—near-end-of-heated-period condition with moderate inlet boundary layer temperature non-uniformity (approximately 15 K across the boundary layer). We then vary $T_{w1,A}/T_{w1,B}$ in the range $0.68 < T_{w1,A}/T_{w1,B} < 1.00$ (where the lower limit represents $T_{w1,A} =$

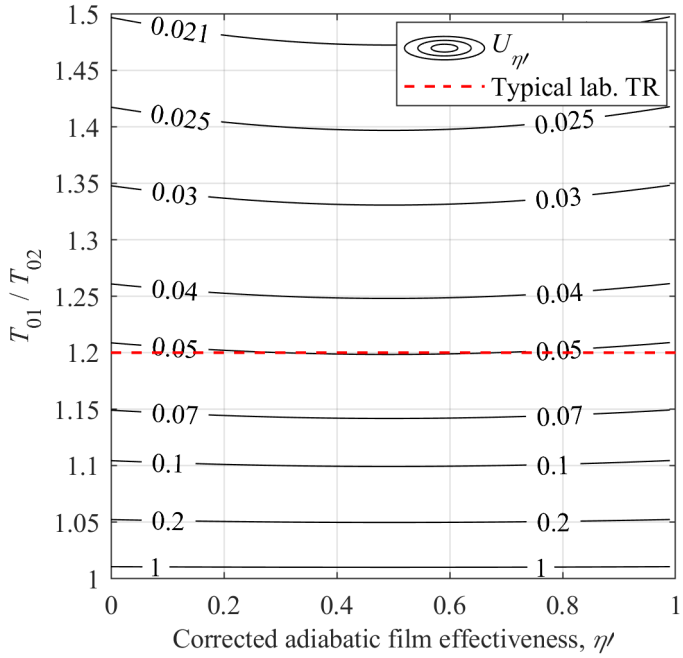


Fig. 6.13 Contours of absolute overall uncertainty in η' as a function of η' and T_{01}/T_{02} for $T_{w1,B} = 0.90(T_{01} - T_{02}) + T_{02}$, $T_{w1,A} = 0.67(T_{01} - T_{02}) + T_{02}$ and $\eta^* = 0.90$.

T_{02}, and the upper limit represents timestamps very close together in time) and examine the impact on the absolute overall uncertainty in η', $U_{\eta'}$. Results are shown in Figure 6.14 for $\eta^* = 0.90$ (flow from lower regions of the thermal boundary layer, i.e. close to the upstream wall reference temperature, T_{w1}).

As expected, in the approach to the upper and lower limits ($T_{w1,A}/T_{w1,B} = 1.00$ and $T_{w1,A}/T_{w1,B} = 0.68$ respectively), the uncertainty in η' increases somewhat exponentially. At the upper limit there is—by definition—no temperature difference between the timestamps (taken at the same time) and therefore the solution of equations 6.5 and 6.10 is extremely unstable. At the lower limit, there is no temperature difference between the measured upstream wall reference temperature at the first timestamp, $T_{w1,A}$, and the coolant temperature, T_{02}, making it difficult to separate contributions from lower boundary layer fluid and coolant.

For $\eta^* > 0.9$ (not shown; equivalent mainstream temperature closer to the upstream wall reference temperature), both the uncertainty and the gradient of uncertainty in η' are higher in the approach to the lower boundary ($T_{w1,A} = T_{02}$) than in the example represented by Figure 6.14. In the limit of the equivalent mainstream effectiveness approaching unity, $\eta^* \to 1$, corresponding to the entrained fluid coming from the bottom of the inlet thermal boundary layer, the equivalent

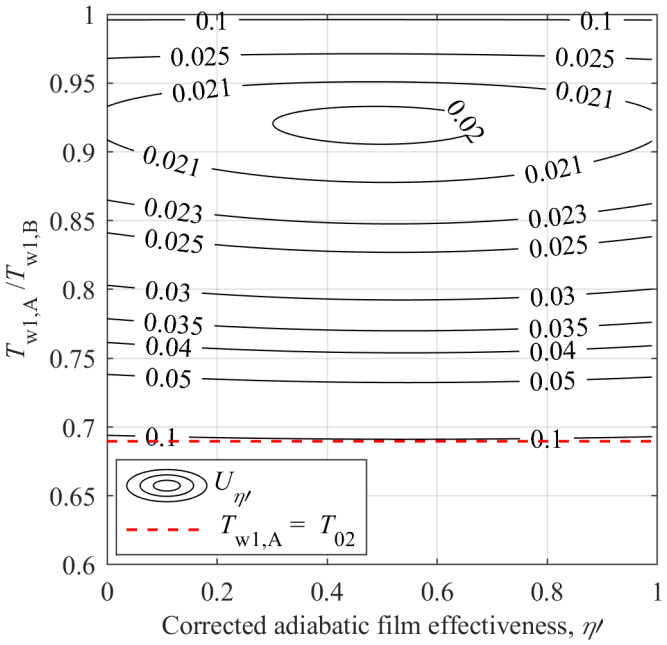


Fig. 6.14 Contours of absolute overall uncertainty in η' as a function of η' and $T_{w1,A}/T_{w1,B}$ for $T_{w1,B} = 0.90(T_{01} - T_{02}) + T_{02}$; $T_{01}/T_{02} = 1.5$; and $\eta^* = 0.90$.

mainstream temperature approaches the upstream wall reference temperature, $T^* \to T_{w1}$. At the start of the heated period, where the upstream wall reference temperature is close to that of the coolant, $T_{w1} \approx T_{02}$, this would result in the equivalent mainstream temperature also being close to that of the coolant ($T^* \approx T_{02}$). In this limit, the temperature difference $T^* - T_{02}$ is small, and hence equation 6.5 becomes unstable as the denominator approaches zero.

For our particular reference values, there is a fairly broad region of low uncertainty in η' (below 0.025) in the range $0.85 < T_{w1,A}/T_{w1,B} < 0.96$, with an optimum at $T_{w1,A}/T_{w1,B} \approx 0.92$ (the condition chosen for preceding section: "Effect of mainstream-to-coolant temperature ratio."). This result is relatively independent of the value of η' but with a shallow minimum in the range $0.30 < \eta' < 0.70$.

Perhaps expectedly, the conclusion is that timestamps should be selected such that there is a large difference between both $T_{w1,A}$ and $T_{w1,B}$ and between T_{w1} and T_{02}. One way of thinking about this is the following: in the proposed technique whilst it is still necessary to maintain a temperature difference between the mainstream gas temperatures and coolant (the minimum difference being dictated by the upstream wall reference temperature) to allow film effectiveness to be determined with acceptably low uncertainty, it is additionally the change in upstream wall

reference temperature ($T_{w1,A}$ to $T_{w1,B}$) that allows the equivalent mainstream effectiveness to be simultaneously determined.

6.7.3 Effect of value of equivalent mainstream effectiveness

The equivalent mainstream effectiveness, η^*, is a function of the inlet thermal boundary layer profile and the flow pattern. Whilst it is hard to imagine a situation in which it could be preferentially controlled, it is nonetheless interesting to consider the impact of the value of η^* on uncertainty in η'. This result is plotted in Figure 6.15 for $T_{01}/T_{02} = 1.50$, $T_{w1,B} = 0.90(T_{01} - T_{02}) + T_{02}$, and $T_{w1,A}/T_{w1,B} = 0.92$ (an optimum value for this condition, as shown in the previous section). The result is that the uncertainty in η' has a relatively weak sensitivity to the value of η^*, varying by a maximum of approximately 22% (for $\eta' = 0.50$) in the range $0 < \eta^* < 1.00$.

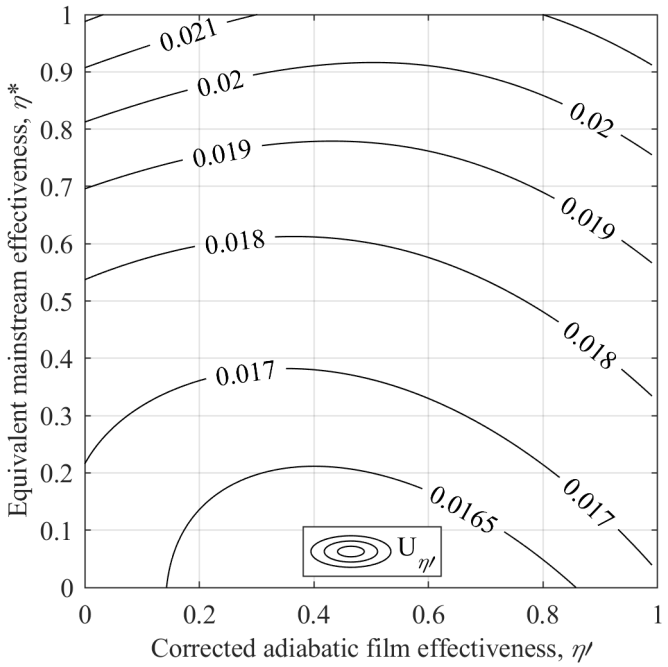


Fig. 6.15 Contours of absolute overall uncertainty in η' as a function of η' and η^* for $T_{01}/T_{02} = 1.50$, $T_{w1,B} = 0.90(T_{01} - T_{02}) + T_{02}$, and $T_{w1,A}/T_{w1,B} = 0.92$.

This example also shows us that in a well-conditioned and optimised experiment (i.e. for the nominal values above), the proposed technique allows very low uncertainty (below 0.022) in corrected film effectiveness, *even for the case of significant inlet thermal boundary layers.*

6.8 Method to reduce uncertainty by using multiple timestamps

In the previous section we calculated the expected overall uncertainty in η' when two timestamps were used. In this section we demonstrate that it is possible to reduce uncertainty in η' by taking a larger number of timestamps ($N > 2$) during the heated period of the run, $0 < t' < 1$. For simplicity we limit ourselves to timestamps evenly distributed between the start and end of the heated period.

The principle is to take every possible pair of timestamps (expressed as a binomial coefficient $\binom{N}{2}$), from the set N, and for each calculate values of η^* and η', and associated uncertainties (using the RSS method, as above). A weighted average value of η' is then calculated by taking a weighted sum according to equation 6.12, where we weight by the square of the inverse of the uncertainty value. This gives low weight to high uncertainty pairs, and high weight to low uncertainty pairs.

$$\overline{\eta'} = \frac{\Sigma_i^{\binom{N}{2}} \eta'_i \left(\frac{1}{U_{\eta'_i}} \right)^2}{\Sigma_i^{\binom{N}{2}} \left(\frac{1}{U_{\eta'_i}} \right)^2} \tag{6.12}$$

A Monte Carlo analysis was performed using 1000 virtual experiments based on an idealised virtual experiment with similar physical conditions to the real data presented in Figure 6.9. The idealised virtual experiment (the initial experiment) is shown in Figure 6.16 and has duration of 30 s, similar to the heated period of the experiment of Figure 6.9. This is similar to the physical conditions of the real experiment but with T_{01}/T_{02} set to a constant value of 1.20, and T_{w1} derived analytically from the response of a semi-infinite surface with physical properties typical of aluminium and a surface heat transfer coefficient contrived to match the measured thermal boundary layer time constant of 2.5 s and measured temperature value $T_{w1} = 316$ K at $t' = 1$. The constant value $\eta^* = 0.5$ was used, with T^* calculated from equation 6.3. Three surface temperature traces were calculated using equation 6.5, taking values of corrected film effectiveness equal to 0.3, 0.5, and 0.8. We refer to a particular true (input) value of corrected

film effectiveness as η'_{TRUE}, and distinguish this from both the corresponding output value of the regression process (for a given pair of timestamps), η'_i, and the weighted-average of the output values $\overline{\eta'}$. For completeness, we make a similar distinction between true (input) values of equivalent mainstream effectiveness, η^*_{TRUE}, and the corresponding weighted-average output values, $\overline{\eta^*}$.

In this method we use a particular constant value of η'_{TRUE}, and a definition of η' (equation 6.5) that is applicable to an adiabatic surface. Thus, output wall temperatures respond instantly to the local driving gas temperature ($\tau = 0$). In systems with a non-adiabatic surface of interest, the surface temperature trace will lag behind (in time) the idealised adiabatic response with a time constant which depends on the material properties and local heat transfer coefficient. The delayed surface temperature response causes the calculated values of $\overline{\eta'}$ to be offset from the true value, by an amount that depends on the time constant of the delayed response. A further correction for this effect is therefore required if the surface of interest is non-adiabatic. This is discussed in Appendix B. For the present analysis we consider the case of a perfectly adiabatic surface. This condition can be approximately met in practice by using very low conductivity materials such as Rohacell.

The 1000 virtual experiments were created by superimposing bias and precision uncertainties on the idealised virtual experiment. To generalise the problem, and to highlight the importance of sampling at sufficiently high rates to drive down the precision uncertainty in the mean, we take values of absolute overall uncertainty as defined in Table 6.1, but ascribe equal contributions to bias and precision uncertainty. This would be representative of processing single data points from an experiment: that is, processing a single IR camera image, and taking a single sample of thermocouple measurements (no averaging of electronic noise). The resulting input uncertainty bands are given in Table 6.2, as are the test conditions for the idealised virtual experiment.

To create the virtual experiments, bias uncertainties (offset) were selected with appropriate weight from the normal distribution defined in Table 6.2, leading to an offset of the entire time history of the idealised virtual experiment. Precision uncertainty (noise) was added in a similar manner but with offsets applied to individual timestamps. This process was performed on every temperature measurement in the idealised virtual experiment, then repeated to create 1000 virtual

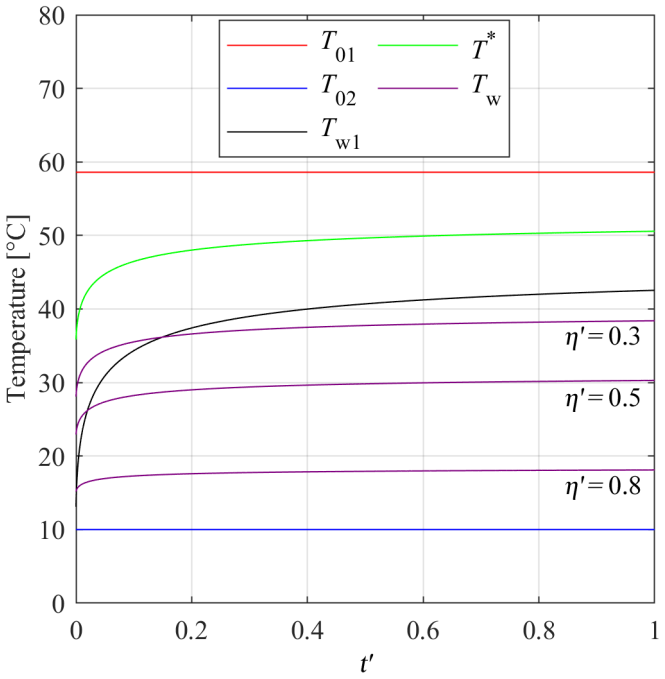


Fig. 6.16 Simulated temperature traces for the idealised virtual experiment (using $\eta^*_{\mathrm{TRUE}} = 0.5$ and $T_{01}/T_{02} = 1.20$) used as the base case for the Monte Carlo analysis.

experiments. The virtual experiments were then processed with varying timestamp number, as we will now describe.

Taking the 1000 virtual experiments as the base data set, the proposed methodology was applied to every virtual experiment for N in the range $2 < N < 50$. Timestamps at $t' = 1$ (end of the run) and $t' = 0.03$ (equivalent to 1 s after the mainstream temperature rise) were used for all runs, with timestamps being evenly distributed between these values for $N > 2$.

We now give an example of the analysis for the surface temperature trace corresponding to $\eta'_{\mathrm{TRUE}} = 0.5$ in Figure 6.16, for 1000 virtual experiments, and for the particular value $N = 50$. A histogram of the 1000 output weighted-average values $\overline{\eta'}$ is given in Figure 6.17. The results are presented in equal buckets of η' (increments of $\eta' = 0.0016$) and show an approximately normal distribution centred on the input value $\eta'_{\mathrm{TRUE}} = 0.5$, with standard deviation 0.0097, or overall uncertainty (to 95% confidence) of 0.0016.

Performing a similar process for $N = 2$ to $N = 50$, we determine values of absolute overall uncertainty (95% confidence) in $\overline{\eta'}$ as a function of N, for input values $\eta'_{\mathrm{TRUE}} = 0.3$, 0.5 and 0.8, and a sample of 1000 virtual experiments.

Table 6.2 Assumed values/ranges of temperatures and corresponding assumed absolute overall, bias, and precision uncertainties to 95% confidence for simulated virtual experiments.

Parameter	Units	Value	U(K) 95%		
			Overall	Bias	Precision
T_{01}	K	340	± 0.5	$\pm\frac{0.5}{\sqrt{2}}$	$\pm\frac{0.5}{\sqrt{2}}$
T_{02}	K	283	± 0.5	$\pm\frac{0.5}{\sqrt{2}}$	$\pm\frac{0.5}{\sqrt{2}}$
T_{01}/T_{02}	–	1.2	–	–	–
T_{w}	K	283–340	± 1.0	$\pm\frac{1.0}{\sqrt{2}}$	$\pm\frac{1.0}{\sqrt{2}}$
T_{w1}	K	283–316	± 0.5	$\pm\frac{0.5}{\sqrt{2}}$	$\pm\frac{0.5}{\sqrt{2}}$

The uncertainty—for a particular value N—is determined by calculating the limits that include 95% of the samples in a particular distribution of the general form of Figure 6.17 (example for $N = 50$). Results are plotted in Figure 6.18 for $N = 2$ to $N = 50$. The general trend in overall uncertainty in $\overline{\eta'}$ is very similar (both qualitatively and quantitatively) for all three values of η'_{TRUE}, starting at a value of approximately 0.048, for $N = 2$, and decaying somewhat asymptotically to a value of 0.019, for $N = 50$. The impact of increasing the sampling number from $N = 2$ to $N = 50$ reduces the overall uncertainty in $\overline{\eta'}$ by approximately 3/5.

Ascribing all of the overall uncertainty to either bias or precision (instead of the equal weighting of Table 6.2) gives rise to the limiting trends shown in Figure 6.18. This shows, at least for the particular sampling strategy chosen (in which optimal timestamps are used), that there is no benefit from increased sampling number if the bias uncertainty is the only contribution to overall uncertainty. This result may be considered surprising, but can be explained as follows. If a perfectly adiabatic surface is assumed, the surface temperature is instantaneously quasi-steady ($\tau = 0$) with respect to the fluid temperatures, and it can be shown that the resulting surface temperature response for particular η' and η^* is still a perfect solution for the system (albeit for values η' and η^* different from the input values), even when offset by a fixed temperature error (a bias error). A bias error (offset) in any measured temperature therefore results in the returned (output) value $\overline{\eta'}$ being offset from input (actual measured, or simulated) η'_{TRUE}. The offset does

not depend on the sampling strategy for the reasons given. In the example 1000 cases presented here, the absolute overall uncertainty, to 95% confidence, in $\overline{\eta'}$ when only bias errors are present is 0.0119.

If we relax the assumption that the surface is perfectly adiabatic then T_w will have a time constant different to that of T^* leading to a lag of the surface temperature response with respect to the driving gas temperature response. In this case a wall temperature with a bias error offset will no longer be a perfect solution for alternate pairings of η' and η^* and we would expect the absolute overall uncertainty to depend on the sampling strategy (particular timestamp pairs).

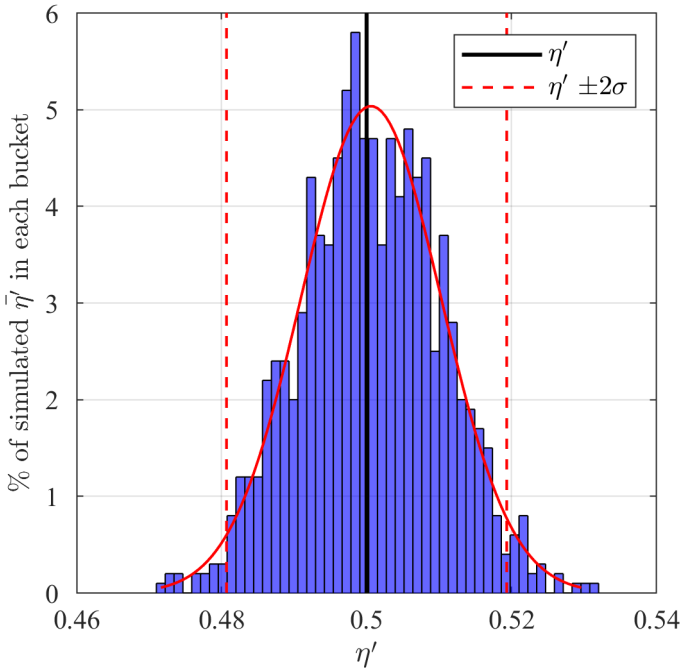


Fig. 6.17 Histogram of values of $\overline{\eta'}$ calculated for 1000 virtual experiments for $N = 50$, presented in equal buckets of η' (increments of $\eta' = 0.0016$).

When precision uncertainty is the only contribution to overall uncertainty, increased timestamp number significantly reduces overall uncertainty. The trend is asymptotic, theoretically approaching 0 as $N \to \infty$, but limited in practice by the sampling frequency of individual variables, and by computational time for high N. For the idealised experiment presented here, and in the case when precision uncertainty is the only contribution to overall uncertainty, the absolute overall uncertainty reduces from 0.065 when $N = 2$ to 0.023 when $N = 50$. For $N = 300$ (not shown), the absolute overall uncertainty is halved again (to approximately 0.011), but at an approximately 50-fold increase in computational cost. Even in experiments with high sampling

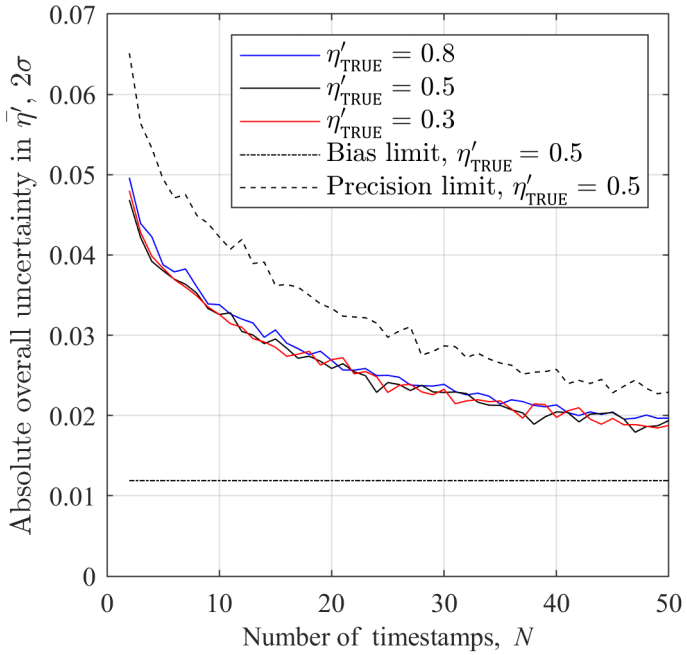


Fig. 6.18 Absolute overall (bias and precision) uncertainty in $\overline{\eta'}$ (95% confidence) for $\eta'_{\mathrm{TRUE}} = 0.3$, 0.5 and 0.8 based on a population of 1000 virtual experiments, and $2 < N < 50$.

rates, it is impossible to entirely eliminate precision error, and therefore reasonably high N would always be a sound choice. We now demonstrate this using simulated data realistic of an actual experiment.

To demonstrate the reduction in uncertainty in $\overline{\eta'}$ with increasing timestamp number, example results from two particular virtual experiments (input values in Table 6.2) are plotted in Figure 6.19 for $\eta'_{\mathrm{TRUE}} = 0.3, 0.5$ and 0.8. The two particular virtual experiment examples were selected to have $\overline{\eta'}$ significantly different from the true value η'_{TRUE} for $N = 2$. We see that as N increases, although there is some random-walk element in the decay of the trends (caused by precision uncertainty), the data collapse so that for $N = 50$, $\overline{\eta'} \approx \eta'_{\mathrm{TRUE}}$. Similar results would be expected in all real experiments.

6.9 Conclusions

In this paper we present a new method to calculate adiabatic film effectiveness in the presence of an inlet thermal boundary layer. The general problem is that in systems with distorted inlet temperature (perhaps due to a cool thermal boundary layer) the convection and mixing processes

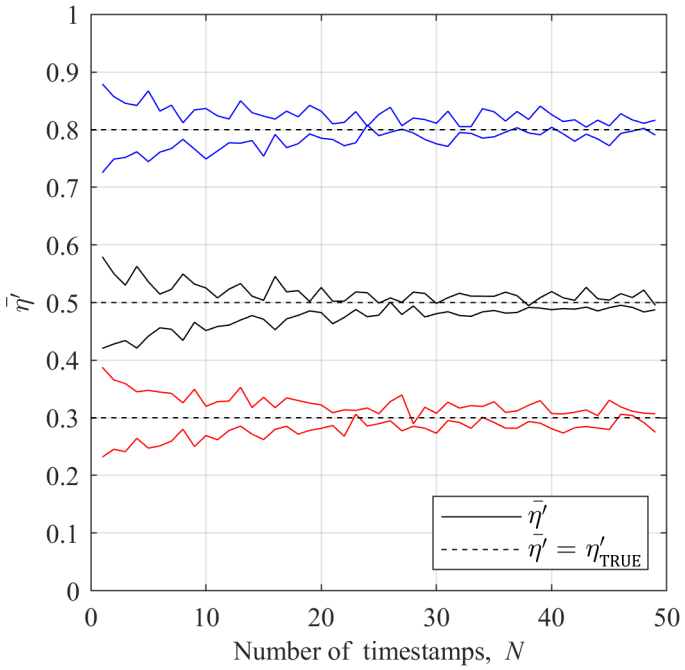


Fig. 6.19 Values of $\overline{\eta'}$ from 2 of 1000 virtual experiments for input values of $\eta'_{\mathrm{TRUE}} = 0.3$, 0.5, 0.8 and $2 < N < 50$.

are complex, and it is impossible to measure the hot gas temperature relevant to each point on the surface. In this situation, conventional processing methods (which use a single reference hot gas temperature) can over-estimate or under-estimate the true value of adiabatic film effectiveness. The arising error generally has both a temporal and a spatial element. This makes interpretation of the data problematic.

In the proposed method, the problem is overcome by calculating an equivalent mainstream temperature, T^*, for every point on the surface as a function of time. This is expressed using a new parameter which we refer to as equivalent mainstream effectiveness, η^*. This can be thought of as a non-dimensional expression for T^* in terms of the core gas temperature and the wall temperature. In our method, the equivalent mainstream effectiveness and the corrected adiabatic film effectiveness η' are solved simultaneously using data from two or more timestamps during the unsteady process of the wall accommodating to the local mixed temperature. The method works even in the case of time-varying inlet distortion, but is more accurate when the time constant of the thermal boundary layer is significantly greater than the time constant of the surface.

Using real experimental data from a high-pressure transonic blowdown linear cascade, we show that the proposed method produces adiabatic film effectiveness distributions which are more physically meaningful than conventional methods.

Using a numerical experiment that is a low-order simulation of a typical heat-transfer experiment, we show that uncertainty in η' can be minimised by: maximising mainstream-to-coolant temperature ratio T_{01}/T_{02}; selecting data timestamps which have significant temperature differences between upstream wall reference temperature measurements T_{w1}; and selecting data timestamps in which the upstream wall reference temperature is significantly different to the coolant supply temperature. The first condition is limited by experimental facility capability. In general, the second and third conditions are achieved by selecting one early timestamp, between 0.23 and 1.20 time constants of the inlet thermal boundary layer, and one timestamp near the end of the heated period of the run. We show that in a well-conditioned experimental setup the overall absolute uncertainty in η' can be below 0.022.

We demonstrate a further reduction in uncertainty in the proposed method by taking more than two timestamps and calculating a weighted-average of results, based on their RSS uncertainties. Through analysis of 1000 virtual experiments with normal uncertainty distributions superimposed on idealised temperature traces, we show that the absolute overall uncertainty of the weighted-average, $\overline{\eta'}$, can be significantly reduced when a large number of timestamps are used. For 50 timestamps, a reduction in uncertainty of 3/5 was achieved over the case with two timestamps.

There are many examples in the literature of adiabatic film effectiveness distributions measured in the presence of a significant inlet boundary layer. Conventional processing methods return measurements with potentially significant errors arising from this effect. The proposed method should significantly increase accuracy of the adiabatic film effectiveness returned in any experiment with significant distortion of the inlet temperature profile.

6.10 Appendix A: Extension to compressible flow

In compressible flows there is a static temperature variation within the boundary layer caused by the so-called recovery effect: that is, in the frame of reference of the wall, there is bulk

kinetic energy reduction and an associated increase in the mean molecular speed or static temperature. This static temperature variation through the boundary layer can be thought of as being superimposed on the existing thermal gradient, and changes the conductive and convective heat transfer within the boundary layer. We refer to the recovery temperature distribution, $T_{01,\mathrm{R}}$, of the surface, which is a function of Prandtl number of the fluid and the local Mach number

$$T_{01,\mathrm{R}} = T_{01}\left[\frac{1+r\left(\frac{\gamma-1}{2}\right)\mathrm{M}^2}{1+\left(\frac{\gamma-1}{2}\right)\mathrm{M}^2}\right] \tag{6.13}$$

where r is the compressible recovery factor which is commonly taken as $r = \mathrm{Pr}^{(1/2)}$ and $r = \mathrm{Pr}^{(1/3)}$ for laminar and turbulent boundary layers, respectively [45]. The Mach number is a function of position and therefore the recovery temperature should properly be expressed as $T_{01,\mathrm{R}}(x,t) = f(\mathrm{M}(x),\mathrm{Pr}(x))T_{01}(t)$, where x represents the surface position. For a particular environment with fixed Mach number distribution and defined boundary layer state, we can simplify this to $T_{01,\mathrm{R}}(x,t) = f(x)T_{01}(t)$ *i.e.* the local recovery temperature is a function of surface position and the local mainstream total temperature.

In an environment with a mixing coolant film, in which the Mach number and Prandtl distributions have been determined in the presence of the mixing film, it would be natural to define the local total temperature at the wall surface in terms of the recovery temperatures of both the mainstream fluid and the cooling fluid, that is

$$T_{\mathrm{w,c}}(x,t) = \eta_{\mathrm{c}}(x)f(x)T_{02}(t) + (1-\eta_{\mathrm{c}}(x))\,f(x)T_{01}(t) \tag{6.14}$$

or, equivalently

$$T_{\mathrm{w,c}}(x,t) = \eta_{\mathrm{c}}(x)T_{02,\mathrm{R}}(x,t) + (1-\eta_{\mathrm{c}}(x))\,T_{01,\mathrm{R}}(x,t) \tag{6.15}$$

where η_{c} is the mass fraction of coolant in the near wall fluid and $(1-\eta_{\mathrm{c}})$ is the corresponding mass fraction of mainstream fluid. Whilst η_{c} should strictly also be specified as a function of time—*i.e.* $\eta_{\mathrm{c}}(x,t)$—in the most general case, in an experiment in which the aerodynamic field is steady, we can simplify to $\eta_{\mathrm{c}}(x)$ as in equations 6.14 and 6.15. Here we also make the

simplification of the same specific heat for both coolant and mainstream, but note that equation 6.14 could easily be extended to the situation where there is a difference in the values of specific heats. Rearranging equation 6.14, we can show that η_c can be interpreted as a compressible adiabatic film effectiveness

$$\eta_c(x) = \frac{f(x)T_{01}(t) - T_{w,c}(x,t)}{f(x)T_{01}(t) - f(x)T_{02}(t)} \tag{6.16}$$

or, restating equation 6.16 directly in terms of the recovery temperatures, we have

$$\eta_c(x) = \frac{T_{01,R}(x,t) - T_{w,c}(x,t)}{T_{01,R}(x,t) - T_{02,R}(x,t)} \tag{6.17}$$

The form 6.17 is discussed in a small number of papers (for example [34], [16]), but it is far more common to assume that the mixing coolant flow is not subject to recovery effects, and to use T_{02} in place of $T_{02,R}$ in equation 6.17 (see, for example, [89], [27]). Nonetheless, the form in equation 6.16 preserves the link between the adiabatic film effectiveness and the mass fraction of coolant and mainstream fluids at a surface location—a physically tangible concept—and is therefore compatible with mass-fraction based methods of the type used in this paper.

We now take this preferred general definition of compressible adiabatic film effectiveness, and develop expressions for corrected compressible adiabatic film effectiveness necessary in the situation which is both compressible and has a time-varying inlet thermal boundary layer.

First, we define a *compressible equivalent mainstream temperature*, T_c^*, in analogy to equation 6.3, by

$$T_c^*(x,t) = \eta_c^*(x,t)T_{w1,R}(x,t) + (1 - \eta_c^*(x,t))T_{01,R}(x,t) \tag{6.18}$$

where, $T_{w1,R}(x,t) = f(x)T_{w1}(t)$, *i.e.* the recovery temperature at a compressible location x, of a streamline originating from the location of the upstream wall reference temperature. Here $T_c^*(x,t)$ is the result at a particular surface location of the mixing of streamlines from within the thermal boundary layer. In equation 6.18, η_c^* and $1 - \eta_c^*$ can be thought of as mass fractions of gas originating from locations instantaneous at the temperature of the upstream wall at the

reference location, $T_{w1}(t)$, and at the core-flow temperature, $T_{01}(t)$, respectively. Rearranging for *compressible equivalent mainstream effectiveness* $\eta_c^*(x,t)$ we get

$$\eta_c^*(x,t) = \frac{T_{01,R}(x,t) - T_c^*(x,t)}{T_{01,R}(x,t) - T_{w1,R}(x,t)} \tag{6.19}$$

In analogy to equation 6.5, we can define a compressible corrected adiabatic film effectiveness

$$\eta_c'(x,t) = \frac{T_c^*(x,t) - T_{w,c}(x,t)}{T_c^*(x,t) - T_{02,R}(x,t)} \tag{6.20}$$

where parameters in the incompressible environment $T^*(x,t), T_w(x,t)$ and $T_{02}(t)$ become equivalent parameters in the compressible environment $T_c^*(x,t), T_{w,c}(x,t)$ and $T_{02,R}(x,t)$. For a steady aerodynamic flow field, η_c' and η_c^* can be assumed to be approximately constant in time allowing us to equate values at any two timestamps A and B, *i.e.*

$$\eta_c^*(x) = \eta_{c,A}^*(x) = \eta_{c,B}^*(x) \tag{6.21}$$

$$\eta_c'(x) = \eta_{c,A}'(x) = \eta_{c,B}'(x) \tag{6.22}$$

Using equations 6.16–6.22 we can derive an expression for $\eta_c^*(x)$—compressible parallel of equation 6.10—in terms of measured temperatures and $f(x)$, where $f(x)$ is calculated from, for example, CFD or supplementary experiment (Mach number distribution required). We can then use equation 6.18 to obtain $T_c^*(x,t)$, and equation 6.20 to obtain $\eta_c'(x)$.

If the surface Mach number distribution is not known from CFD or a supplementary experiment (*i.e.* $f(x)$ not known) then provided additional timestamps are used (three or more), it is possible to extract $f(x)$ at the same time as $\eta_c^*(x)$ and $\eta_c'(x)$. That is, we solve for three variables simultaneously. This introduces additional sensitivities which would need to be assessed.

6.11 Appendix B: Correction for non-ideal adiabatic surface behaviour

In this section we quantify the error in the value of corrected adiabatic film effectiveness determined by our method in the case of non-adiabatic surface behaviour. That is, the error in the result if non-adiabatic behaviour is not accounted for. We refer to the extracted values in the case of non-adiabatic behaviour as inferred (by the method) values.

Materials used in adiabatic film effectiveness studies are typically selected for their low values of thermal conductivity, k, and thermal inertia, $(\rho ck)^{1/2}$, which result in approximately adiabatic behaviour. Nonetheless, no material is perfectly adiabatic, and this leads to heat transfer into the surface and a time lag between the surface temperature response and the driving gas temperature response. The surface time constant is a function of k and $(\rho ck)^{1/2}$, and the local heat transfer coefficient, h. The surface time constant is lower at lower values of k and $(\rho ck)^{1/2}$ and at higher h. True adiabatic behaviour, with $\tau = 0$, occurs for $(\rho ck)^{1/2} = 0$ or $h = \infty$.

The idealised virtual experiment, discussed above (see Figure 6.16 and context), is replotted in Figure 6.20 for an adiabatic surface, and for a surface with material properties of glass-filled nylon ($k = 0.28$, $(\rho ck)^{1/2} = 730$ J m^{-2} K^{-1} s$^{-1/2}$), with $h = 600, 1200$ and 1800 W m^{-2} K^{-1}. For all cases we specify input $\eta'_{\mathrm{TRUE}} = 0.5$, $\eta^*_{\mathrm{TRUE}} = 0.5$, and the time constant for T_{01} is zero. The driving gas temperature was calculated with the specified input values $T_{01}(t)$, $T_{02}(t)$, η'_{TRUE} and η^*_{TRUE}. This driving temperature is equal to the adiabatic surface response $T_{\mathrm{w}}(t)$, shown by the dashed line in Figure 6.20. The surface temperature responses for the non-adiabatic surface were calculated using a one-dimensional semi-infinite transient conduction model implemented in *MATLAB* (purple lines). As expected, there is a time-lag with respect to the driving temperature which increases as h decreases.

The method of this paper assumes an adiabatic surface response. Processing the simulated non-adiabatic surface responses of Figure 6.20 using this same method, and failing to account for the non-adiabatic behaviour, results in inferred output values (from virtual experiment) $\overline{\eta'}_{\mathrm{INFERRED}}$ which are different from the true (input to virtual experiment) values, η'_{TRUE}. In

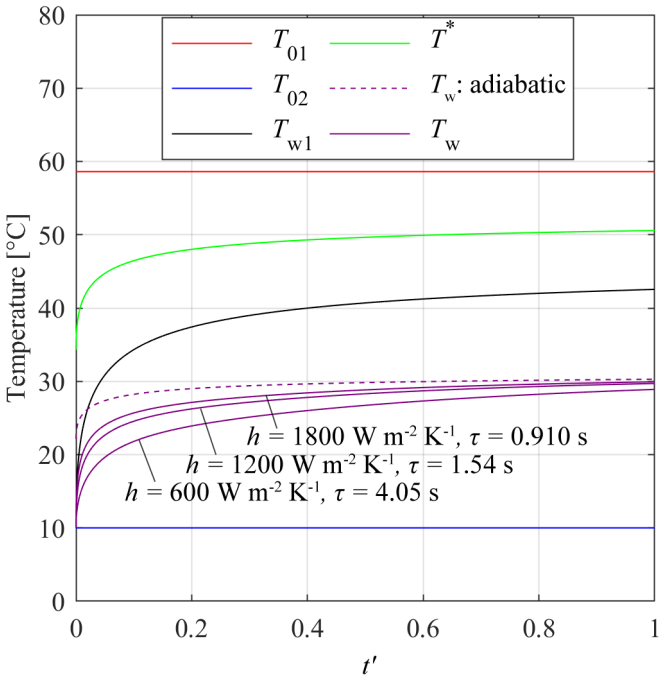


Fig. 6.20 Simulated surface temperature for adiabatic and non-adiabatic surfaces using thermal properties of glass-filled nylon: $(\rho ck)^{1/2} = 730$ J m^{-2} K^{-1} s$^{-1/2}$; $\eta'_{\text{TRUE}} = 0.5$; $\eta^*_{\text{TRUE}} = 0.5$; and $T_{01}/T_{02} = 1.20$.

the same way, inferred output values of $\overline{\eta^*}_{\text{INFERRED}}$ are different from the true (input) values of η^*_{TRUE}.

For target surface materials with large values of k and $(\rho ck)^{1/2}$, or very low h, the time constant of the target surface can exceed that of T_{w1} and best-fit solutions to the equations become non-physical. In this situation, solutions with $\overline{\eta^*}_{\text{INFERRED}} > 1$ are possible, for example, representing an apparent driving temperature, T^*_{INFERRED}, lower than the upstream wall reference temperature. These non-physical solutions can be stable, in that using different timestamps, or a greater number of timestamps, will still result in the same (incorrect) inferred pair $\overline{\eta^*}_{\text{INFERRED}}$ and $\overline{\eta'}_{\text{INFERRED}}$ for a given surface temperature response.

We define the offset error as the difference between the output from the proposed method and input values of corrected adiabatic effectiveness, *i.e.* $\overline{\eta'}_{\text{INFERRED}} - \eta'_{\text{TRUE}}$. This offset error can be estimated by using a series of virtual experiments. As an example, we do this for every value of η'_{TRUE} for four commonly used surface materials (Macor; glass-filled (GF) nylon; acrylic; and Rohacell) of differing k and $(\rho ck)^{1/2}$, but for fixed values $\eta^*_{\text{TRUE}} = 0.5$, and

$h_{\mathrm{TRUE}} = 1200$ W m^{-2} K^{-1}. In our method we choose 50 timestamps, evenly distributed between optimal first and last timestamps. Results are shown in Figure 6.21.

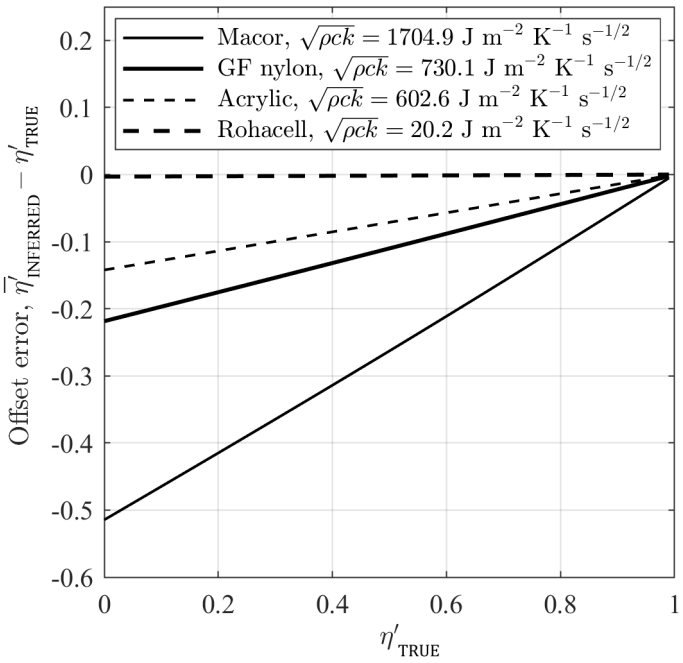


Fig. 6.21 Offset error $\overline{\eta'}_{\mathrm{INFERRED}} - \eta'_{\mathrm{TRUE}}$ plotted as a function of η'_{TRUE} for four materials commonly used in adiabatic film effectiveness experiments, for $\eta^*_{\mathrm{TRUE}} = 0.5$, and $h_{\mathrm{TRUE}} = 1200$ W m^{-2} K^{-1}.

The offset values are everywhere negative, with larger offsets for larger values of $(\rho ck)^{1/2}$ (larger non-adiabatic tendency) and for smaller η'_{TRUE}. We also note that inferred values of equivalent mainstream effectiveness, $\overline{\eta^*}_{\mathrm{INFERRED}}$, are greater than η^*_{TRUE} (not shown) and hence the inferred equivalent mainstream temperature, T^*_{INFERRED}, is lower than the true value at any given timestamp. The offset error reduces in magnitude as η'_{TRUE} increases, because when the wall temperature is closer to the coolant temperature errors in the driving hot gas temperature have reduced effect. For larger values of $(\rho ck)^{1/2}$ the surface time constant is greater, for a given value of h, resulting in greater errors in T^*_{INFERRED}. Although materials like Rohacell, are almost perfectly adiabatic (offset values less than 0.003 for all η'_{TRUE}), we see that for Acrylic the maximum offset is approximately 0.14, and worsens considerably for glass-filled Nylon and Macor. Without further correction (next section) many common materials would give unacceptable error in this method.

For an experiment with known material properties and distribution of external surface heat transfer coefficient, by modelling the surface response we can perform a first order correction to the inferred values $\overline{\eta'}_{\text{INFERRED}}$, to remove the principal component of the offset error ($\overline{\eta'}_{\text{INFERRED}} - \eta'_{\text{TRUE}}$). We now attempt such a correction, and refer to the resulting best estimate of adiabatic film effectiveness as $\overline{\eta'}_{\text{BE}}$. For semi-infinite surfaces, in the absence of lateral conduction, and for known h, the correction could be performed without error ($\overline{\eta'}_{\text{BE}} = \eta'_{\text{TRUE}}$). We characterize the error in the best estimate result, $\overline{\eta'}_{\text{BE}} - \eta'_{\text{TRUE}}$, for semi-infinite surfaces, in the absence of lateral conduction, for surfaces where the actual value of h (h_{TRUE}) is different from the assumed value, h_{ASSUMED} used to model the surface response. We call the quantity $\overline{\eta'}_{\text{BE}} - \eta'_{\text{TRUE}}$ the residual offset error. Results for four materials are plotted in Figure 6.22 for input values $\eta'_{\text{TRUE}} = 0.5$, $\eta^*_{\text{TRUE}} = 0.5$, and $h_{\text{ASSUMED}} = 1200$ W m^{-2} K^{-1}, for the range of values $600 < h_{\text{TRUE}} < 1800$ W m^{-2} K^{-1}.

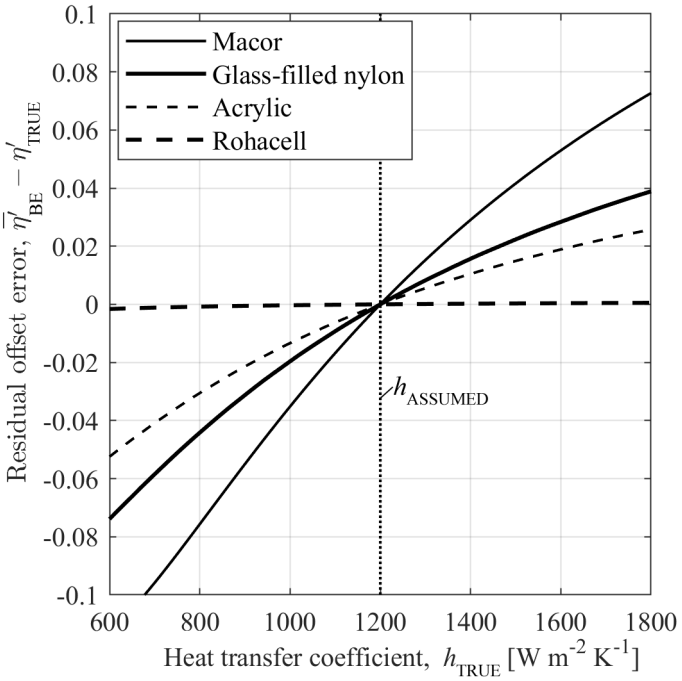


Fig. 6.22 Residual offset error, $\overline{\eta'}_{\text{BE}} - \eta'_{\text{TRUE}}$, for four materials for $\eta'_{\text{TRUE}} = 0.5$, $\eta^*_{\text{TRUE}} = 0.5$, $h_{\text{ASSUMED}} = 1200$ W m^{-2} K^{-1} plotted over the range $600 < h_{\text{TRUE}} < 1800$ W m^{-2} K^{-1}.

At the condition $h_{\text{TRUE}} = h_{\text{ASSUMED}}$ the residual offset error is zero. That is, our best estimate value $\overline{\eta'}_{\text{BE}}$, is equal to the true value η'_{TRUE}. This is trivially the case because, in the situation in which the assumed and true values of h are the same, our characterization of the offset error $\overline{\eta'}_{\text{INFERRED}} - \eta'_{\text{TRUE}}$ is perfect. The residual offset error increases both as $|h_{\text{TRUE}} - h_{\text{ASSUMED}}|$

increases (departure of true surface response from assumed surface response) and as $(\rho ck)^{1/2}$ and k increase. We see that for values h_{TRUE} within 10% of the assumed value, the maximum residual offset errors are 0.020, 0.011, 0.008 and 0.000 for the four materials discussed, in contrast to the absolute offset errors (Figure 6.21) 0.263, 0.110, 0.071 and 0.002. That is, by using a surface response model to estimate the correction term, the absolute errors can be reduced by a factor of approximately 1/10.

Chapter 7

Front seal purge flow angle and platform cooling

Coolant leaks onto the rotor platform from the front seal as a consequence of the mainstream non-uniform pressure distribution present along the seal interface. This purge flow acts as a cooling film along the rotor platform. However, the area effectively cooled is limited as coolant is rapidly swept towards the suction surface by passage crossflow. In a rotating environment the area cooled is further reduced as coolant exits the seal with a tangential velocity component, resulting in a flow angle biased toward the suction surface. The tangential velocity component of the purge flow coolant is a function of, and is lower than, the rotor speed due to shearing between stationary and rotating surfaces within the seal structure and this results in worse cooling coverage at higher rotor speeds. The film cooling performance of front seal purge can be enhanced by altering the inlet flow angle, as it flows onto the platform, to be away from the blade suction surface. This increases the area cooled and the average adiabatic film effectiveness without incurring significant aerodynamic losses.

This chapter details a parametric experimental investigation into the adiabatic film cooling effectiveness of front seal purge flow. The effects of changing the inlet purge flow angle and purge mass ratio were investigated by using sets of guide vanes integrated into the front seal flow path. Full-area adiabatic film effectiveness measurements of one blade passage are presented along with area traverse measurements of downstream total pressure loss coefficient. A practical method is proposed to alter the rotor purge flow angle in real engine configurations by using exhaust coolant from nozzle guide vane platform cooling.

7.1 Front seal purge flow angle in a rotating frame

High pressure air is continuously supplied to the front seal structure, between upstream NGV and rotor disk, via a flowpath that passes through the central core of the engine. This sealing air flows radially up through the seal structure and ultimately exhausts into the main gas path as a purge flow across the platform surface, through an annular front seal slot. As the air moves through the seal it is subjected to a constant shearing between the rotating rotor disk and the static walls of the engine core, seal, and NGV which imparts a tangential swirl component to the flow in the direction of rotation. At the exit of the seal, the purge flow coolant flows onto the platform at an angle, in the absolute frame, that is in the direction of rotor travel and hence towards the pressure surface of the blade. Due to the shearing effect, the tangential velocity component of the purge flow is lower than that of the blade. The result of this is that the purge flow angle in the rotor relative frame is opposite to the absolute angle and towards the suction surface of the blade.

In the following analysis the situation is reduced to its simplest form, of air flowing between a stationary flat surface and a moving flat surface with constant velocity. This allows the complicated flow structures within the seal, dependent on specific seal design, to be approximated in a generalised case. A linear velocity profile is assumed between the two surfaces, with a velocity of zero at the static wall increasing to the rotor tangential velocity at the moving surface. This is shown schematically in Figure 7.1. The average velocity of this profile, equal to half the rotor tangential velocity, is taken as the bulk velocity of the purge flow at the exit plane as it flows onto the platform. More complex velocity profiles could be assumed (either by analogy to boundary layer velocity profiles or to account for more complex spatial variations due to NGV wakes *etc.*) however this complicates the model and simply results in a change to the value of bulk velocity, which is nonetheless always proportional to the rotor speed.

A schematic of the front seal flow path, simplified model geometry and velocity distribution, and velocity triangles of the purge exit flow are given in Figure 7.1. Here α is the flow angle in the absolute (static) frame, β is the flow angle in the relative rotor frame, v_{abs} is the absolute purge flow velocity, v_{rel} is the relative purge flow velocity, v_{x} is the axial velocity component, v_{T}

is the relative tangential velocity component, and U is the tangential rotor speed. The coordinate system used throughout this chapter is that tangential velocity and angles in the direction of rotor motion are positive, such that U and α in Figure 7.1 are positive and v_T and β are negative.

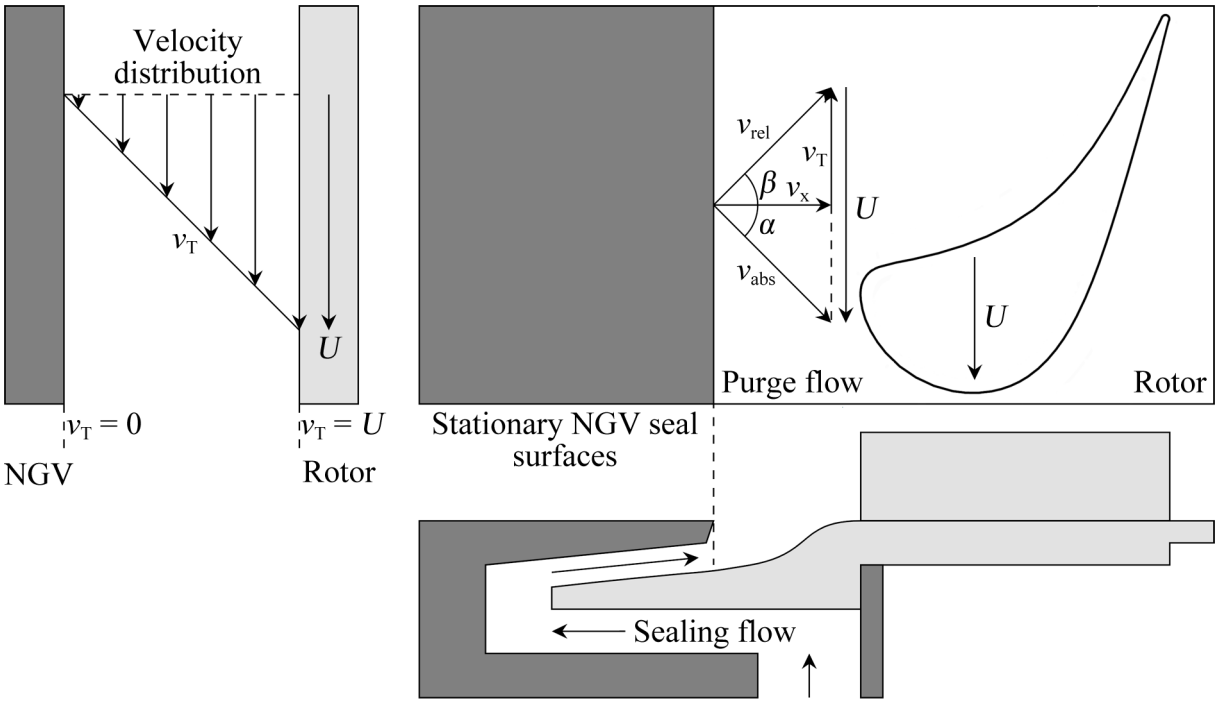


Fig. 7.1 Schematic diagrams showing the velocity distribution between stationary and rotating surfaces within the front seal, the flow path of front seal coolant and resultant purge flow velocity components and angles at the exit plane onto the rotor platform.

The purge flow angles are dependent on the ratios between axial and tangential velocity components, and are therefore dependent on the rotor speed. All of the purge flow passes through an axial plane at the end of the front seal (the front seal slot) and, assuming the purge flow Mach number is low so that the flow is incompressible, the axial velocity can be calculated from equation 7.1. The purge flow pressure and temperature are specified by the sealing requirement and available air supply. The flow area of the annular slot is determined by the required minimum rotor clearance for motion and thermal expansion. The purge flow axial velocity component can therefore only be altered by increasing, or reducing, the mass flow rate of coolant through the seal. An increase in purge mass flow rate causes an increase in axial velocity and, for a given rotor speed, a reduction in purge flow angles.

$$v_x = \frac{\dot{m}_{pf}}{\rho_{pf} A_{pf}} \tag{7.1}$$

Purge flow angles for a range of coolant mass flow rates (specified as mass ratios, MR, in the range 0.5–2.5% between mainstream and coolant) at engine speeds up to 14,000 RPM were calculated for a hypothetical engine with operating conditions, at the maximum take-off condition (MTO), specified in Table 7.1. Tangential velocity was taken as half the rotor speed and axial velocity calculated from equation 7.1. The purge flow area was taken to be half the geometric purge slot area. The non-uniform pressure distribution present along the front seal/mainstream interface blocks flow from emerging from the purge slot in regions of relatively higher pressure. This was therefore approximated as an effective blockage of half the purge slot flow area. Results are given in Figure 7.2 which show the ratio of tangential-to-axial velocity components and the resulting relative purge flow angles.

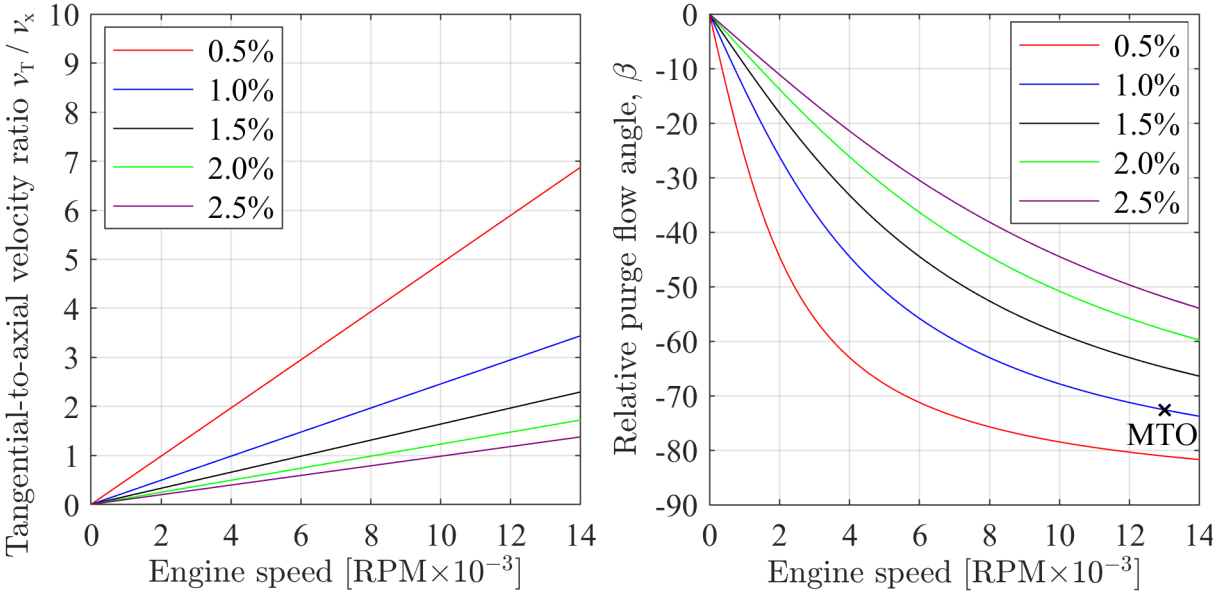


Fig. 7.2 (a) Resultant tangential-to-axial purge flow velocity ratio and (b) relative purge flow angle for purge mass ratios in the range 0.5–2.5% at engine speeds in the range 0–14,000 RPM.

The ratio of tangential-to-axial velocity components (Figure 7.2a) increases linearly with rotor speed since, for a given purge mass ratio, the axial velocity is constant and we assume a linear relationship between rotor speed and tangential velocity. The gradient of this ratio is lower for larger purge mass ratios as the axial velocity is relatively larger compared to the tangential velocity components. The relative purge flow angles (Figure 7.2b) are negative, indicating that flow is biased towards the suction surface in the relative rotor frame, as the tangential velocity

components are lower than the rotor speed. If the purge flow did not experience shearing in the front seal cavity and was accelerated to match the rotor speed then the resulting relative purge angle would be zero, resulting in axial purge flow in the rotor relative frame.

Table 7.1 Representative engine specification and operating conditions for a hypothetical HP rotor at maximum take-off condition.

Parameter	Symbol	Value	Unit
Purge flow pressure	p_{pf}	310000	Pa
Purge flow total temperature	$T_{0,\mathrm{pf}}$	1000	K
Mean platform radius	r_{p}	0.33	m
Purge slot radial height	h_{pf}	2×10^{-3}	m
Purge slot flow area	A_{pf}	4.1×10^{-3}	m^2
Rotor angular velocity	ω	13000	RPM
Rotor tangential velocity	$U = r_{\mathrm{p}}\omega$	450	$\mathrm{m\ s}^{-1}$
Nominal purge mass ratio	$\mathrm{MR} = \dfrac{\dot{m}_{\mathrm{pf}}\sqrt{T_{0,\mathrm{pf}}}}{\dot{m}_{\mathrm{m}}\sqrt{T_{0\mathrm{m}}}}$	1.0	%

At lower engine speeds, the purge flow axial velocity is greater than or equal to that of the tangential velocity ($v_{\mathrm{T}}/v_{\mathrm{x}} \leq 1$) and the magnitude of the purge flow angles are relatively small. As the engine speed increases, the tangential velocity becomes much greater than the axial velocity and the relative flow angle approaches $-90°$. At the nominal MTO engine speed of 13,000 RPM and 1% purge mass ratio, the tangential velocity is approximately three times larger than the axial velocity, resulting in purge flow strongly biased towards the suction surface with a relative flow angle of $-72.1°$. This means that the area cooled by the purge flow will be severely limited as it flows directly towards the suction surface. This effect is independent of the pressure-gradient-induced passage crossflow which also acts to sweep purge flow towards the suction surface.

Increasing the purge flow mass ratio increases the axial velocity and reduces the magnitude of the relative flow angle, for a given engine speed. However, the effectiveness of this is limited with a large increase in purge mass ratio (from 1% to 2.5%) at MTO resulting in only a 36% reduction in flow angle magnitude (from $-72.1°$ to $-46.0°$.) This reduction in purge relative flow angle would be expected to result in wider film cooling coverage across the platform. The increase in mass flow would also be expected to increase cooling effectiveness as the purge flow temperature

would rise more gradually due to having more mass. An increase in purge flow would also result in higher cavity pressure and better sealing performance. The aerodynamic penalty incurred due to bleeding more coolant from the compressor and flooding it back into the main gas path would be expected to be significant and would likely not be acceptable for performance. At high engine speeds, with fast rotor velocities, it is the tangential velocity component which primarily determines the purge flow angle. It would therefore be more beneficial to alter the tangential velocity component to either reduce its magnitude towards the suction surface, or flip it from negative to positive and bias the flow towards the pressure surface. This would increase film cooling coverage without the need to increase the purge mass ratio.

7.1.1 Practical method to alter purge flow angle

One method currently employed in some engine architectures to alter the tangential velocity component of flow within the front seal cavity is to use pre-swirl nozzles to supply coolant. These nozzles are inclined at an angle to swirl coolant into the cavity in the direction of the rotor motion. The primary function of these pre-swirl nozzles is to assist in coolant entry into the already rotating cavity flow and are typically at a relatively low radial position. The shearing effect therefore occurs over a large portion of the seal, slowing the flow, as the coolant flows radially towards the purge flow slot. Conventional pre-swirl systems therefore have a limited possible effect on the purge flow exit angle.

Proposed systems exist in patents which use guide vanes integrated into the front seal geometry, located near to the purge slot exit, which swirl flow immediately before ejection onto the mainstream platform surface. This concept was used successfully to set purge flow angles in the static cascade experiments, described below, but are highly impractical in actual engines. The gap between rotor and stator surfaces exists as a necessity to allow for rotor motion and thermal expansion without collision between moving and stationary parts. The gap represents a significant point of failure, due to the risk of hot gas ingestion into the wheel space, and is purposefully kept to a minimal size. This minimises the amount of coolant required to seal the gap and therefore the associated cycle efficiency penalty. Clearance would still be required

if guide vanes were added to the slot which would mean the gap would have to be increased to accommodate the vanes and would require additional coolant to properly seal. Efforts to minimise coolant usage with guide vanes would result in low-aspect ratio designs, to minimise the additional cavity size, which would be largely ineffective due to the relatively large tip gap of the clearance area.

A more practical method of imparting swirl to the purge flow close to the seal exit would be to implement pre-swirl nozzles close to the exit plane. The proposed method here is to supply a fraction of the sealing coolant from the upstream NGV platform cooling system. Coolant would be exhausted through holes on the internal seal surfaces inclined at a large angle to the axial direction, towards the direction of rotor motion. Nozzle guide vane platforms require cooling similarly to the rotor. Internal cavities, below the NGV platform, with a constant supply of coolant are commonly implemented. Coolant from the internal NGV cooling system is typically also used to supply external film cooling on the NGV surface. The temperature rise in coolant within the internal NGV cooling system is not significant enough that the coolant cannot be reused as film cooling for the NGV and is therefore also a suitable temperature to be used in the rotor front seal cavity.

The internal NGV cooling cavity typically follows the geometry of the vane passage and therefore near to the trailing edge the coolant flows at a large flow angle, in the direction of the rotor pressure surface. Exhausting the internal NGV platform coolant through holes immediately upstream of the platform trailing edge, directly into the rotor front seal cavity, would form high velocity jets at a large turning angle, which would subsequently mix with the seal flow and exhaust onto the rotor platform as purge flow. The resultant flow angle of the mixed purge flow would depend on the flow velocity and angle of the jets exhausted into the seal, and the split of coolant mass flow through the jets and main front seal flowpath. The coolant supplied to the main front seal structure could be appropriately reduced so that the overall amount of coolant used would remain the same and the sealing performance maintained. A schematic of the proposed system showing a number of jets being exhausted into the front seal from the NGV, subsequently mixing with sealing flow, and being ejected onto the rotor platform is given in

Figure 7.3. Velocity triangles of the NGV jet flow (1), seal flow (2), and resultant mixed-out purge flow (3) are also included.

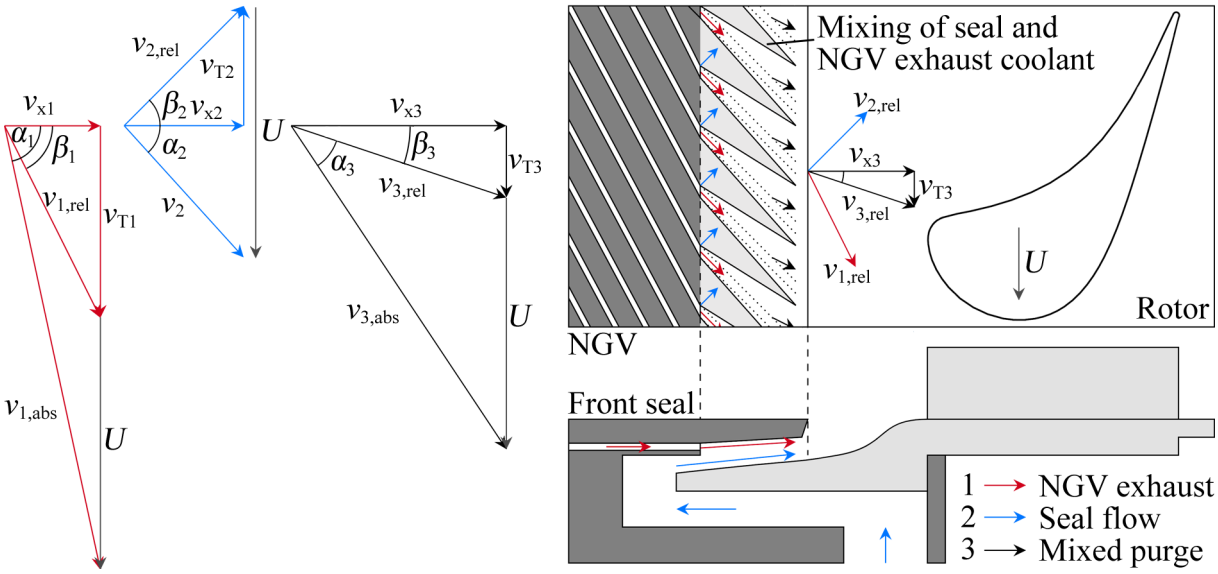


Fig. 7.3 Velocity vector triangles of coolant flows and schematic diagram of a front seal system using jets exhausting from the NGV platform cooling system to supply a fraction of the sealing flow.

A low-order model to demonstrate the viability of such a system was written as a *MATLAB* script. The total amount of coolant supplied to the front seal is specified for the same range of purge mass ratios as before (in the range 0.5–2.5%). The fraction of coolant exhausted from the NGV, f_{pf}, varied in the range 0–1, where 0 indicates all of the coolant is supplied to, and flows through, the main seal structure and unity indicates that all of the flow is exhausted via holes from the NGV platform cooling system. Flow through the main seal structure is assumed to behave as previously described, with an axial velocity determined by the flow area and mass flow rate and tangential velocity equal to half the rotor speed. The holes which form jets from the NGV cooling system are assumed to be choked, due to the large pressure drop across the NGV and subsequent very high coolant supply pressure to the coolant system, and are oriented to emit jets with a large absolute flow angle ($\alpha_1 = 70°$). The flow from these jets is assumed to mix uniformly with the main seal flow such that conservation of momentum (equations 7.2–7.4) can be used to determine the resultant mixed flow velocities and angles.

$$\dot{m}_{pf} v_{x3} = f_{pf} \dot{m}_{pf} v_{x1} + (1 - f_{pf}) \dot{m}_{pf} v_{x2} \tag{7.2}$$

$$\dot{m}_{\mathrm{pf}} v_{\mathrm{T3}} = f_{\mathrm{pf}} \dot{m}_{\mathrm{pf}} v_{\mathrm{T1}} + (1 - f_{\mathrm{pf}}) \dot{m}_{\mathrm{pf}} v_{\mathrm{T2}} \tag{7.3}$$

$$\beta_3 = \arctan\left(\frac{v_{\mathrm{T3}}}{v_{\mathrm{x3}}}\right) \tag{7.4}$$

Results of the mixed purge flow absolute and relative flow angles, as the purge exhausts onto the platform surface, are presented in Figure 7.4 for fractions of flow exhausted from the NGV between 0–1, mass ratios in the range 0.5–2.5%, and at the MTO engine speed of 13,000 RPM. When $f_{\mathrm{pf}} = 0$ all of the purge flow comes through the main seal structure and the resultant relative flow angles are therefore identical to those in Figure 7.2. As f_{pf} increases from zero towards unity there is a corresponding decrease in the flow through the seal structure and the absolute flow angle, for all values of purge mass ratio, tends towards the metal angle of the jet-forming holes in the NGV (70°). As the amount of coolant through the main seal structure decreases, the axial velocity (v_{x2}) reduces, leading to an increase in flow angle (α_2). This results in the characteristic curves of absolute flow angle of the mixed-out purge flow (α_3). For a purge MR = 1.5%, and greater, α_3 initially increases towards the jet metal angle as a larger fraction of the coolant is being supplied at that angle and the axial velocity decreases causing the through-seal flow angle to increase. For f_{pf} in the range 0.6–0.9 the resultant absolute flow angle becomes greater than the blade metal angle because the axial velocity of the through-seal flow is very small, and the corresponding angle of this flow is therefore larger than the jet metal angle.

In the relative frame, the flow angles similarly tend towards a single value of 34°, determined by the rotor speed and NGV jet metal angles. The relative flow angle of the mixed-out purge is negative, corresponding to flow towards the rotor suction surface, for lower f_{pf} values (below approximately 0.6). For the particular modelled case, with fixed engine speed of 13,000 RPM and NGV jet metal angle of 70°, axial flow is established in the relative frame when $f_{\mathrm{pf}} = 0.61$. The mass ratio of the overall purge flow does not affect the value of f_{pf} at which axial flow occurs and all of the lines pass through the same point. This is because we assume that all flow through the main seal structure experiences the same kind of shearing in the seal, with the same resultant linear tangential velocity profile set by the rotor speed. Increasing the overall purge mass ratio therefore increases the axial velocity, but the value of f_{pf} which results in the tangential velocity components of the two flows cancelling-out to give axial flow remains constant.

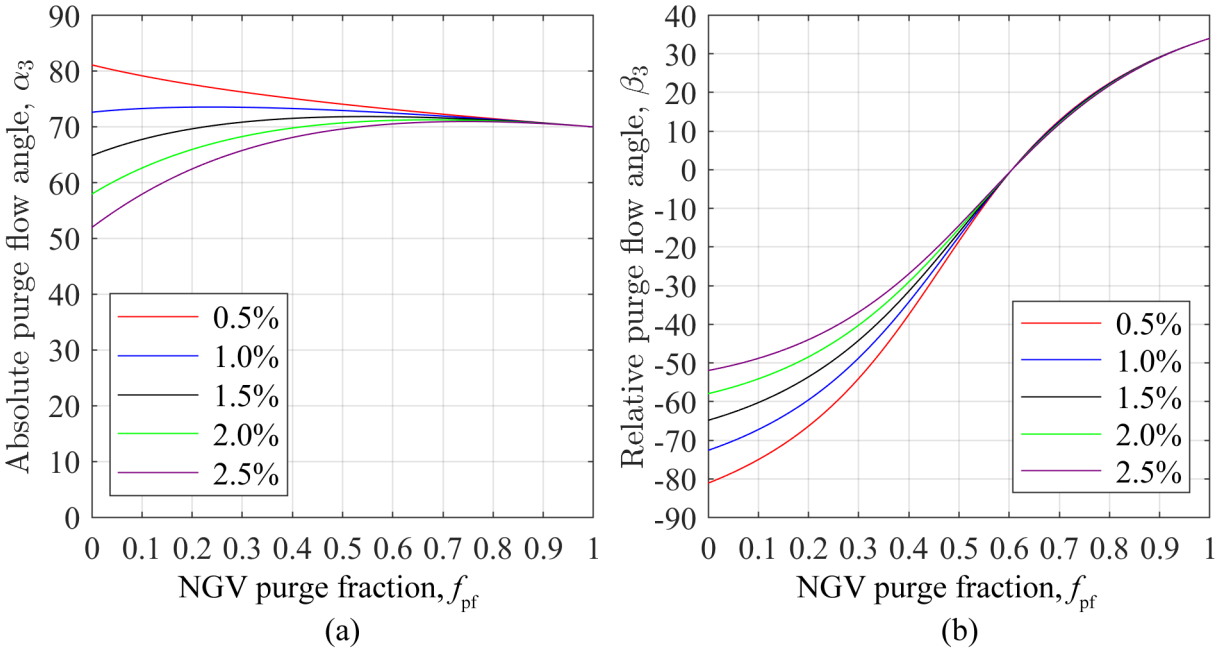


Fig. 7.4 Resultant flow angles of mixed-out purge flow at the exit plane onto the rotor platform, for purge mass ratios in the range 0.5–2.5%, when a fraction of the coolant is supplied from exhaust jets from the NGV platform cooling system. (a) flow angle in the absolute (stationary) frame and (b) flow angle in the relative rotor frame.

This modelling shows that partially supplying front seal coolant from the upstream NGV platform cooling system could be used to control the flow angle of the purge flow as it exits onto the rotor platform. The purpose of this would be to direct flow away from the natural bias, towards the suction surface of the blade, and towards the pressure surface of the blade in order to increase the area of the platform affected by film cooling. For a nominal purge flow mass ratio of 1.0% a realistically achievable fraction of coolant supplied from the NGV is estimated to be f_{pf} in the range 0.2–0.8. From Figure 7.4b this can be seen to correspond to resulting mixed-out purge flow angles in the range $-60°$ to $20°$. Angles within this range, in increments of $20°$, were selected for experimental testing.

7.2 Experimental setup and testing procedure

A parametric study was conducted in the MERC facility to investigate the effect of varying purge flow angle and mass ratio on the film cooling performance across the rotor platform. Front seal blocks, described in detail in Chapter 3, with integrated guide vanes were used to swirl flow

immediately upstream of the front seal exit slot. These guide vanes allowed the exit purge flow angle to be set at exact angles and are analogous to using jets of coolant from the NGV platform cooling system to partially supply the front seal purge flow coolant. Five purge flow angles of $-60°$, $-40°$, $-20°$, $0°$ and $20°$ were each tested at four purge mass ratios of 0.5%, 1.0%, 1.5%, and 2.0% for a total of 20 unique test conditions. To isolate the film cooling performance of the purge flow, baseline test blades with no integrated cooling features were used and coolant was only supplied to the front seal cavity and rear seal within the blade cassette. Liquid silicone sealant was used between adjacent test blades to prevent coolant or mainstream air from entering the shank cavity. The test blades used were made of low thermal conductivity glass-filled nylon so that adiabatic film effectiveness distributions could be determined from IR images. Adiabatic film effectiveness distributions across two blade platforms and downstream total pressure loss coefficient distributions, across an area covering one blade passage up to midspan, were obtained for each test condition from a single run in the MERC facility. This allowed for both the film cooling performance of the purge flow and the aerodynamic penalty incurred to be quantified and compared between test conditions.

All experiments were conducted with the heater mesh power supply set to 100 V to achieve mainstream-to-coolant temperature ratios of approximately 1.15. Each run was 210 s in duration to allow for both area traverse and IR measurements to be taken. As the heater mesh could not be operated for longer than 60 s at a time the area traverse measurements were taken with unheated mainstream flow.

Each of the five front seal blocks had guide vanes machined onto the internal seal surface at fixed angles to direct the purge flow, with the exception of the axial flow '0-degree' block which just has an open flow area. When assembled in the blade cassette the guide vanes are in contact with the extended blade front lip, with no clearance gap. All of the coolant is therefore forced to flow through the guide vane passages. The guide vane profiles are 2 mm wide flat plates and span the full 2.4 mm span between the front seal block and test blade lip. The leading edge profile is an ellipse with an aspect ratio of 2:1. Each vane diffuses to a sharp 0.2 mm trailing edge which terminates 2 mm upstream of the exit plane. Each of the front seal designs therefore has the same exit flow area onto the platform so that mass ratio and flow angle are the only variables affecting

the purge between test cases. The spacing between guide vanes is 8 mm which allows for a reasonable amount of overlap, so that complete passages are formed between adjacent vanes, whilst allowing for a relatively large flow area to mitigate the area affected by boundary layers on the main seal surface. This vane geometry was largely determined for ease of manufacture to reduce cost and production time. The location and design of the guide vanes used in the MERC facility are shown in Figure 7.5.

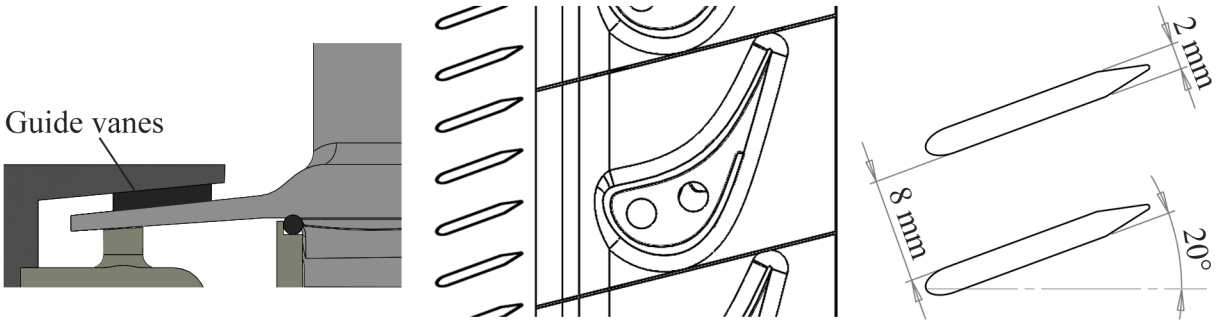


Fig. 7.5 Schematic diagrams of the guide vane design and location within the front seal, immediately upstream of the purge exit slot.

Purge mass ratios were set by changing the coolant mass setting plates which had different numbers of holes. The nominal design points for the setting plates were for mass ratios of 0.5%, 1.0%, 1.5%, and 2.0%. The nominal engine operating point is 1.0% and values were tested both above this, to examine the impact of simply flooding the platform with additional coolant, and below this to allow for characteristic behaviour to be determined. At the time this study was conducted the orifice plate assembly had not been fitted to the coolant line and per-run mass ratios were therefore not monitored. These were back-calculated based on each runs measured pressures and temperatures, and the calibrated hole capacity from later experiments. Estimated values from this method were found to be within 10% of the nominal design points for each experiment.

7.3 Adiabatic film effectiveness results

Adiabatic film effectiveness distributions have been determined from IR video recorded during each experiment using the processing techniques described in Chapter 6. This involves a number of corrections to the data to account for the non-uniform and time-varying inlet temperature

profile within the MERC facility, and compressible recovery temperature. The additional correction to account for non-ideal adiabatic surface behaviour has not been applied as the processing technique is still being refined to work correctly for complete surface images without introducing spatial distortion. This means that offset errors are present in the reported values of adiabatic film effectiveness and therefore these are lower than the expected true values. A consequence of this is that negative effectiveness values occur, which are non-physical, in regions of little or no film cooling. All negative values have been considered equal to zero in analysis. Whilst the presence of these offset errors is non-ideal, it is a consistent problem (quantifiable and reproducible) which predominantly affects low values of adiabatic film effectiveness and affects results from all experiments equally. It is therefore expected to have a negligible impact on the comparative performance of the test cases examined.

Values of adiabatic film effectiveness reported within this chapter are therefore the equivalent of $\overline{\eta'_c}$ which is the compressible weighted average corrected adiabatic film effectiveness from multiple timestamp method. For brevity, this is simply referred to as adiabatic film effectiveness and denoted with η hereafter.

In the following sections a typical result showing the adiabatic film effectiveness distribution along the central blade passage is presented and discussed to illustrate the characteristic features of purge flow film cooling. The results from each of the 20 test conditions are then compared side-by-side with discussion. Surface distributions were obtained for two complete blade platforms however for display clarity and ease of analysis only the central blade passage data is used and presented for each case.

7.3.1 Typical film cooling characteristics of front seal purge flow

The behaviour of front seal purge flow once it has exited the purge slot onto the blade platform is governed primarily by the secondary flows within the blade passage. The interactions between the purge flow and secondary flows directly affects the adiabatic film cooling performance of the purge flow. We can therefore infer these main features and interactions of the flows from processed adiabatic film effectiveness distributions on the blade platform. A typical blade

passage distribution of η is given in Figure 7.6 from an experimental run in the MERC facility with purge MR = 1.0 and a front seal angle of 0°. These conditions are at the nominal engine purge mass ratio and with no front seal guide vanes.

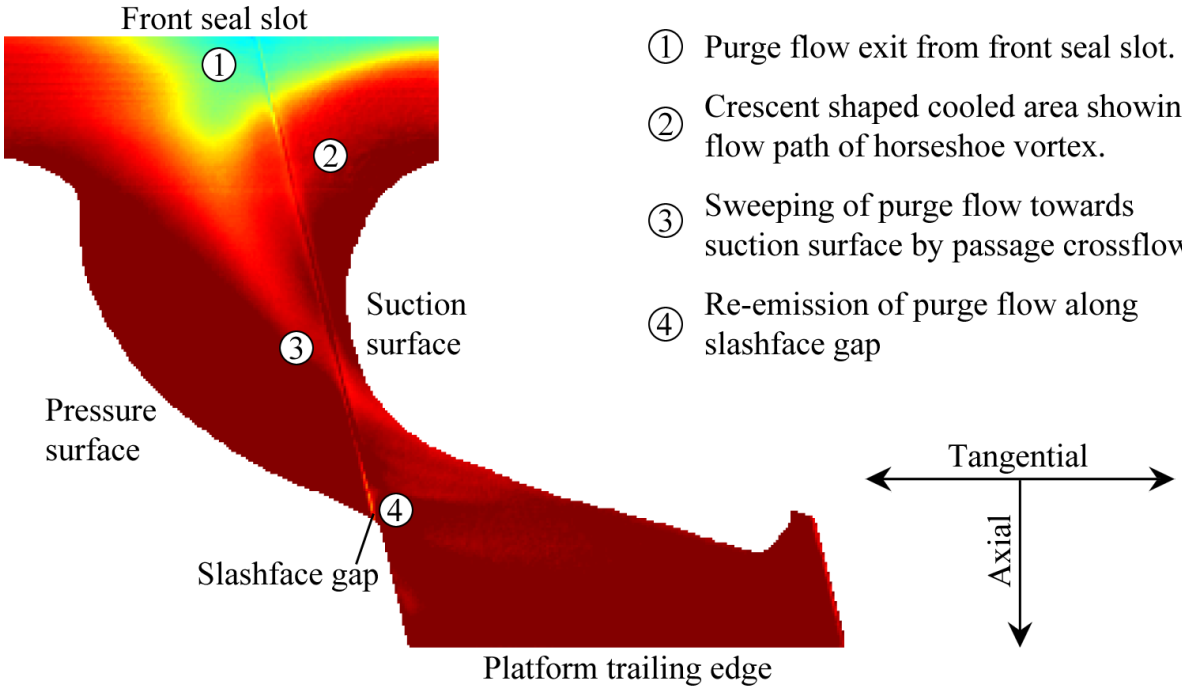


Fig. 7.6 Example front seal purge flow adiabatic film effectiveness distribution on the rotor platform down one blade passage. With a seal angle of 0° and purge mass ratio of 1.0%.

In Figure 7.6, purge flow exits the front seal slot along the top of the image, which is cropped along the edge of the front seal block surface, and then flows downwards through the blade passage. The purge flow is initially at the coolant supply temperature and cools the platform as it flows across the surface, leading to a wedge-shaped region of high adiabatic film effectiveness. The adiabatic film effectiveness is greatest towards the centre of the cooled region, with a peak value of 0.65, and gradually reduces to a value of 0 towards either of the blade surfaces. This wedge shaped cooled region occurs as a result of the purge flow concentrating towards the middle of the blade passage, being pushed away from the high static pressure of the leading edge stagnation zone and blade pressure surface. The adiabatic film effectiveness decreases towards the tip of the wedge, further into the blade passage for two reasons. Firstly, the coolant begins to mix with hot mainstream gas which increases the temperature and reduces the cooling potential

of the flow. Secondly, the purge flow is entrained within the passage vortex which convects the coolant further up into the blade span, away from the platform surface.

A crescent-shaped area of low adiabatic film effectiveness, of approximately 0.14, surrounds the leading edge stagnation region. This shows the location and path of the leading edge horseshoe vortex which entrains the purge flow near the leading edge and subsequently forms into the passage vortex. The tip of the wedge-shaped cooled region also forms the tip of the crescent and this arcs towards the suction surface of the blade, with contact at approximately 50% axial chord. Purge flow not immediately entrained by the horseshoe and passage vortex is swept towards the suction surface by passage crossflow, from the high pressure surface to the relatively low pressure of the suction surface. Ultimately, the purge flow caught in the passage crossflow is also entrained by the passage vortex near the suction surface of the blade and is convected downstream along the blade surface.

Streaks of low adiabatic film effectiveness, of approximately 0.05, can be seen near to the blade suction surface marked (4). One cause of these streaks is the downwash of purge flow coolant within the passage vortex. Another cause for these cooled streaks is purge flow which travels along the blade slashface surfaces between adjacent blade platforms. Liquid silicone sealant was used to seal the clearance gap between test blades and prevent ingestion of mainstream gas into the shank cavity. The slashface gap can be seen in Figure 7.6 by the straight line through the blade passage. The silicone was not applied to the entire slashface however in order to prevent it from leaking onto the mainstream surface during assembly and being visible to the IR camera. The silicone sealant could not be painted and has a different emissivity to that of the painted blade surfaces and this would have corrupted the IR data. The consequence of this is that a narrow flow path exists along the blade platform slashfaces. Purge flow from the front seal flows into this slashface gap and along the surfaces towards the relatively low pressure towards the rear of the platform. At a point where the blade passage pressure is sufficiently low (just downstream of the blade throat where the flow is fastest), the purge flow within the slashface gap re-emerges into the mainstream and cools a small area of the platform surface. The amount of coolant flowing along the slashface is small but remains relatively cool as it does not mix with mainstream flow until it re-emerges into the main gaspath and causes small cooled streaks.

The front seal purge flow only effectively cools a small region of the rotor platform, predominantly the region upstream of the blade leading edge. Additional film cooling occurs in a limited wedge-shaped area towards the middle of the blade passage, towards the suction surface, as coolant is swept across the passage. Areas of the platform near the pressure surface, and in the rear 50% of the axial chord are almost entirely uncooled. The vast majority of the front seal purge is entrained by the passage vortex which pulls the coolant towards the blade suction surface, and lifts it up away from the platform.

Now that the characteristic features of the adiabatic film effectiveness distributions have been discussed, a comparison of each of the 20 test conditions is given so that the effects of changing mass ratio and seal flow angle can be observed.

7.3.2 Comparison of blade passage adiabatic film effectiveness distributions

Distributions of adiabatic film effectiveness similar to Figure 7.6 were determined for each test condition. A composite of these distributions side-by-side is given in Figure 7.7. Results are arranged such that purge flow MR increases left–right (from 0.5% to 2.0%) and seal angle increments from towards the pressure surface (positive flow angle) to towards the suction surface (negative flow angle) top-to-bottom. Arrows have been included to indicate the relative direction of purge flow as it emerges from the front seal slot at the top of each respective distribution image.

It is immediately apparent that the general characteristics of the purge film cooling, described above, are similar across all the tested conditions. The purge flow emerges from the front seal slot and cools a wedge shaped region of the rotor platform. The purge flow is then swept rapidly towards the suction surface, leaving the platform area towards the pressure surface and trailing edge almost entirely uncooled. The general path of the passage vortex is consistent across all of the tested conditions as its formation and flow path is determined by the mainstream flow and is largely unaffected by the relatively small amount of purge coolant flow. Key differences between

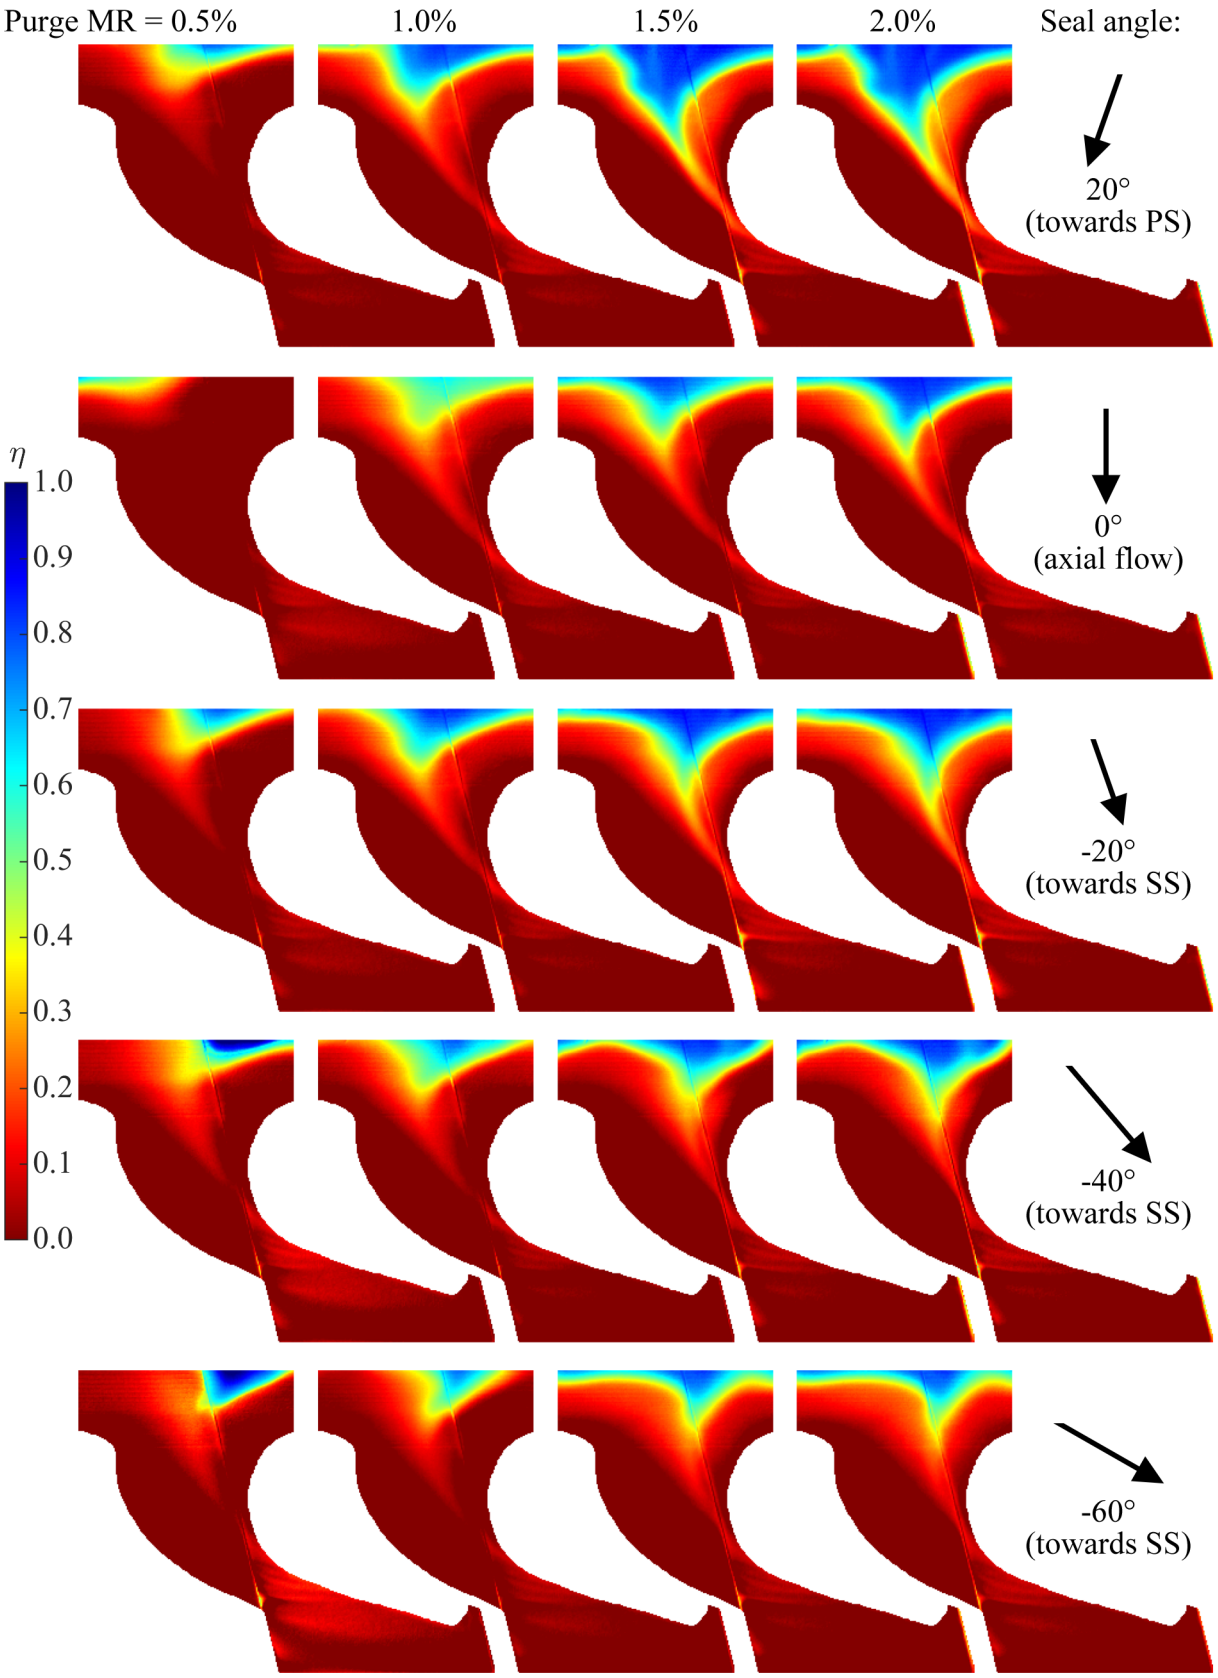


Fig. 7.7 Corrected adiabatic film effectiveness distributions for one blade passage at 20 total test conditions with varying purge mass ratio and flow angle, set by using front seal guide vanes.

the various test conditions are limited to the size and shape of the wedge shaped cooled region as well as the location of the area with highest adiabatic film effectiveness within the cooled region.

Effect of varying purge mass ratio

Increasing the purge flow mass ratio increases the size of the cooled area and the peak η achieved. This is consistent for each of the tested seal angles. The larger amount of coolant at higher purge mass ratios increases in temperature at a slower rate as more heat is required to increase the bulk coolant fluid temperature. This results in better film cooling performance as the temperature of the purge flow remains relatively low as it flows further into the blade passage. In addition, the axial velocity of the purge flow as it emerges onto the platform is higher for larger mass ratios and the increased momentum of the fluid allows it to flow further across the platform surface before being entrained by passage vortex or being swept across the passage by crossflow. This leads to better film cooling area coverage. This suggests that one way of enhancing the film cooling performance of purge flow is simply to increase the amount of coolant supplied to the front seal to flood the platform with coolant. The area effectively cooled by the purge flow is still limited primarily to within the platform leading 50% axial chord. The platform area towards the pressure surface remains entirely uncooled. The area towards the suction surface leading edge has elevated adiabatic film effectiveness at higher purge mass ratios. The best example of this is the 20° seal angle in the region near to the leading edge (towards the right of the image in Figure 7.7) in which η increases from 0–0.1 for a purge MR of 1.0% to 0.1–0.3 for a purge MR of 2.0%.

For purge mass ratios of 0.5% and 1.0% there are portions to the left of the distributions, towards the pressure surface, along the top edge which have very low adiabatic film effectiveness, below 0.1. This indicates regions where the front seal coolant is prevented from flowing by the high pressure of the leading edge stagnation region. This allows some hot mainstream gas to be ingested into the front seal. Whilst this ingestion may only be partial, with hot gas being confined close to the exit slot and diluted by purge coolant, this nonetheless introduces the risk of overheating internal components. At higher mass ratios a band of adiabatic cooling effectiveness greater than 0.7 is present along the entire length of the front seal slot indicating that

any mainstream gas which enters the front seal is heavily diluted by purge coolant. Increasing the purge mass ratio enhances the film cooling effectiveness of the purge flow in addition to increasing the sealing performance of the front seal. The only negative impact of increasing purge mass flow is therefore related to the cycle efficiency penalty incurred by bleeding a larger amount of coolant from the compressor.

The ingestion of mainstream gas into the seal is most prominent in the cases with purge mass ratio of 0.5%. In each of these cases the purge flow is concentrated towards the suction surface and cools only a very small area. An outlier is the case which has a seal angle of 0° and purge mass ratio of 0.5%, where the purge flow ejects close to the pressure side and seems to be convected into a neighbouring blade passage. This large-scale redistribution of purge flow coolant is due to significant ingestion of mainstream gas into the seal. Also of note is that values of adiabatic film effectiveness are typically lower at similar locations at each purge mass ratio for the 0° seal compared to equal purge MR for the other seal angles. The guide vanes, present in all of the front seals other than the 0° seal, inhibit large-scale flow redistribution within the seal. This may be advantageous to increase sealing effectiveness and a similar effect may occur if sealing flow were supplied from jets from the NGV platform cooling system. The jets would not be expected to as effectively block lateral flows within the seal as solid vanes but would assist in redirecting and mixing-out ingested mainstream flow. Sealing performance was not assessed in this study but is a good candidate for future work.

Effect of varying purge seal flow angle

Purge flow exits the front seal slot at an angle determined by the seal angle of internal guide vanes. The effect of seal angle is to alter the shape of the cooled region and the location of the area with highest adiabatic film effectiveness. In Figure 7.7 the 0° seal distributions have wedge shaped cooled regions approximately centred in the blade passage, with a very slight skew towards the suction surface. For cases with negative seal angles this cooled region is shifted towards the suction surface. Similarly, the cases with a seal angle of 20° the cooled region is shifted closer to the pressure side. This suggests that the guide vanes are indeed altering the

purge flow angle as expected and that the flow angle does have a significant impact on the purge film cooling performance along the rotor platform.

Before discussing the differences caused by seal angle in detail, it is important to appreciate the aerodynamics of the purge flow and how the streamlines are affected by the presence of different seal geometries. In order to better understand these aerodynamics, which cause the observed differences between adiabatic film effectiveness distributions for the different seal angles, a set of computational fluid simulations were run of cascade configurations of the 20°, 0°, and $-20°$ seals. The purpose of these CFD simulations was to qualitatively understand the way in which the guide vanes affected the purge flow streamlines, rather than to perfectly match the cascade experiments for comprehensive thermal comparisons.

Simulations were run in *Ansys CFX* of a periodic blade passage domain which included the mainstream gas path and the topmost portion of the front seal. Meshes of the domain were unstructured, with approximately 2.4 million cells and inflation layers added to all surfaces to allow for boundary layers to be resolved. Boundary conditions were set based on the cascade operating conditions and are summarised in Table 7.2. Total pressure was specified at the inlet and mass flow rates specified for both the purge flow coolant inlet and overall outlet of the fluid domain. Each simulation had the same purge mass ratio of 1.0% and the only difference between cases was therefore the seal angle. The turbulence model used was k-ε and solutions were iterated until root-mean-squared residuals, monitoring changes in conservation equations, were below 1×10^{-6} and stable which typically took under 1000 iterations.

Figure 7.8 shows a comparison of streamlines from the three CFD cases showing the path that purge flow takes as it passes through the internal guide vanes, exhausts from the front seal slot and becomes entrained within the passage vortex. In each case a tangential vortex can be seen which flows through the space between the front seal exit plane and the blade leading edge. This vortex forms as a result of the pressure gradient across the blade preventing the purge flow from flowing towards the pressure side. The vortex therefore flows from the pressure surface towards the suction surface and feeds into the main passage vortex as it emerges at approximately mid-passage.

Table 7.2 Boundary conditions for CFD simulations of a periodic cascade blade passage and front seal, including internal guide vanes.

Parameter	Value	Unit
Inlet total pressure	275000	Pa
Inlet total temperature	328	K
Inlet velocity x-component	0.731	–
Inlet velocity y-component	-0.682	–
Purge flow mass flow rate	0.045	kg s^{-1}
Purge flow total temperature	295	K
Purge flow inlet velocity x-component	1.0	–
Purge flow inlet velocity y-component	0.0	–
Outlet total mass flow rate	0.4545	kg s^{-1}

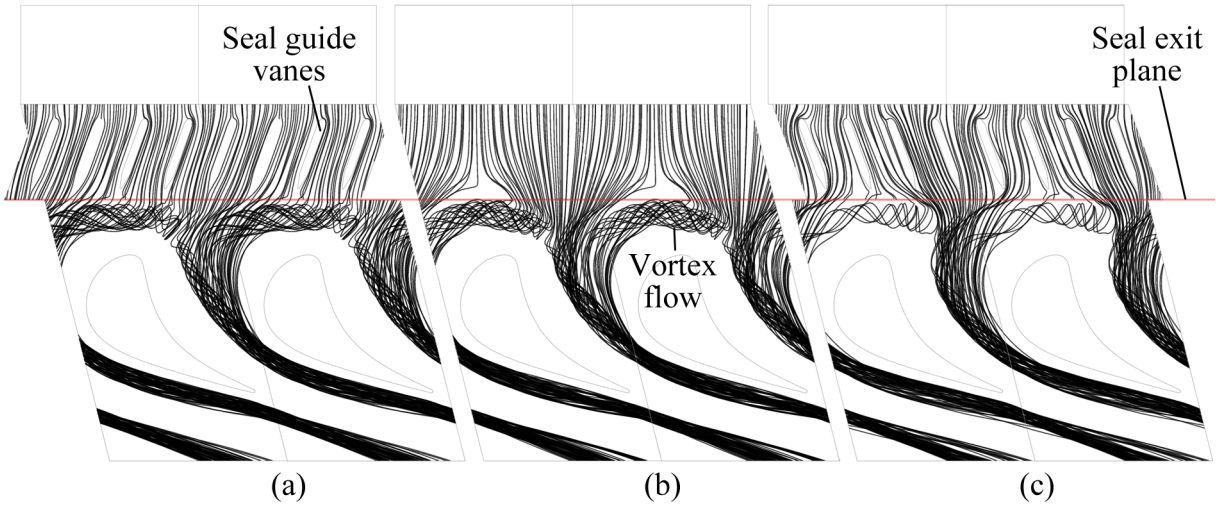


Fig. 7.8 Purge flow streamlines from CFD simulations of three cascade front seal configurations with seal angles of (a) 20°, (b) 0°, (c) −20° showing flow within the front seal and the path of purge flow once exhausted onto the rotor platform.

The amount of purge flow entrained by the tangential seal vortex varies between the three cases. Mainstream pressure gradients propagate upstream into the seal structure causing the purge flow streamlines to be significantly altered within the seal. In the 0° case the purge flow is blocked from emerging from the front seal slot near the leading edge which results in a no-flow area (the white area absent of streamlines upstream of the seal exit plane) and the streamlines coalesce into a tighter flow area approximately centred in the blade passage. The purge splits into two distinct flows, with part being entrained in the passage vortex and the other part forming the tangential vortex. The tangential vortex flows into the adjacent blade passage where it joins with the passage vortex flow. The streamlines of flow within the tangential vortex all join the

vortex at the same point, on the pressure surface side of the leading edge. In contrast to this is the 20° case which has an even distribution of purge flow across the seal. The purge angle is in the same general direction as the tangential vortex flow and so streamlines merge into this vortex at all points along its path before ejection into the mainstream.

The −20° case shows comparatively fewer streamlines within the tangential vortex. The purge flow is directed in opposition to the tangential vortex direction and the vortex therefore acts as a blockage. The purge flow along the tangential vortex path is forced to significantly change direction when being entrained. This results in a larger proportion of the purge flow directly joining the passage vortex, rather than flowing along the front seal gap. The negative seal angle distributions in Figure 7.7 have narrower cooled areas, located closer to the suction surface where the passage vortex initially forms. This is a result of less coolant flowing within the tangential vortex reducing the cooling towards the pressure surface in the blade passage. More coolant immediately joins the passage vortex in the −20° seal angle case which results in a longer extension of the cooled area down the passage, compared to the 0°. For larger negative seal angles the purge flow is directed both more aggressively against the tangential vortex flow direction and towards the leading edge stagnation region. This results in the purge flow stagnating and being swept away from the platform, into the passage vortex and results in further reductions to the size of the cooled area and lower adiabatic film effectiveness values.

The 20° seal angle cases have larger cooled areas and elevated adiabatic film effectiveness values since the purge flow more easily joins the tangential vortex flow. Purge flow is evenly distributed out of the front seal slot and more of its momentum is preserved allowing it to flow across more of the platform surface and enhancing its film cooling performance. In each seal angle case, once the purge flow joins the main passage vortex and is lifted away from the platform the film cooling performance drops off dramatically. The general flow path of the passage vortex is unaffected by the purge flow angle which can be seen by the similarity in streamlines within the blade passages in Figure 7.8 and the view of the axial-span plane in Figure 7.9. It can be seen from these figures that the purge flow ultimately all joins the passage vortex and that this flow then lifts up away from the platform and is convected downstream towards and along the suction surface.

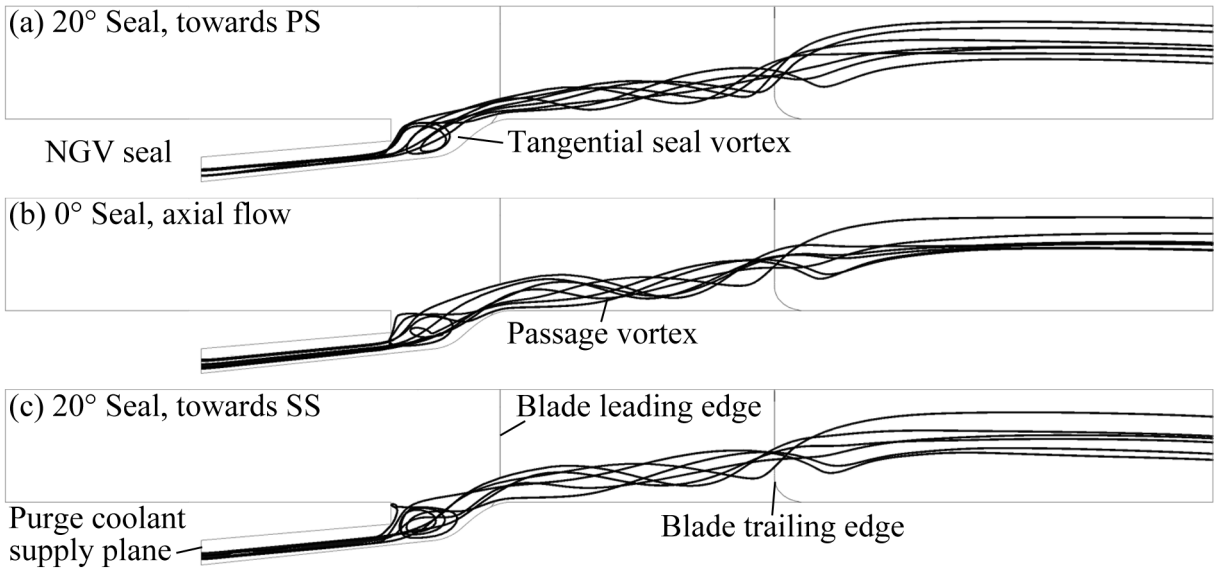


Fig. 7.9 Purge flow streamlines from CFD simulations of three cascade front seal configurations with seal angles of (a) 20°, (b) 0°, (c) −20° showing the entrainment of purge flow into the passage vortex.

7.3.3 Comparison of axial platform adiabatic film effectiveness profiles

In order to more easily quantify the differences in adiabatic film effectiveness distributions discussed above, the axial profiles were determined by averaging values across tangential lines at each axial location through the blade passage. Comparisons of these profiles are plotted in Figure 7.10. The profiles are plotted in groups to demonstrate the effect of mass ratio and seal angle respectively, with (a)–(e) comparing different purge mass ratios for each of the seal angles, and (f)–(i) comparing the same results of the different seal angles at each of the four tested purge mass ratios.

It can be seen that the majority of adiabatic cooling effectiveness is concentrated in the leading 20% of the platform axial length in all cases. Each profile starts with a peak at the exit of the front seal slot (normalised axial position of zero) which sharply decreases towards zero approximately 50% of the way through the blade passage. A short plateau occurs in each profile between normalised axial positions of 0.20 and 0.25 as a result of the tip of the wedge shaped cooled regions extending down into the platform as the purge flow joins the passage vortex. A small secondary peak, less than 0.1, occurs between normalised axial positions of 0.7–0.8 due to

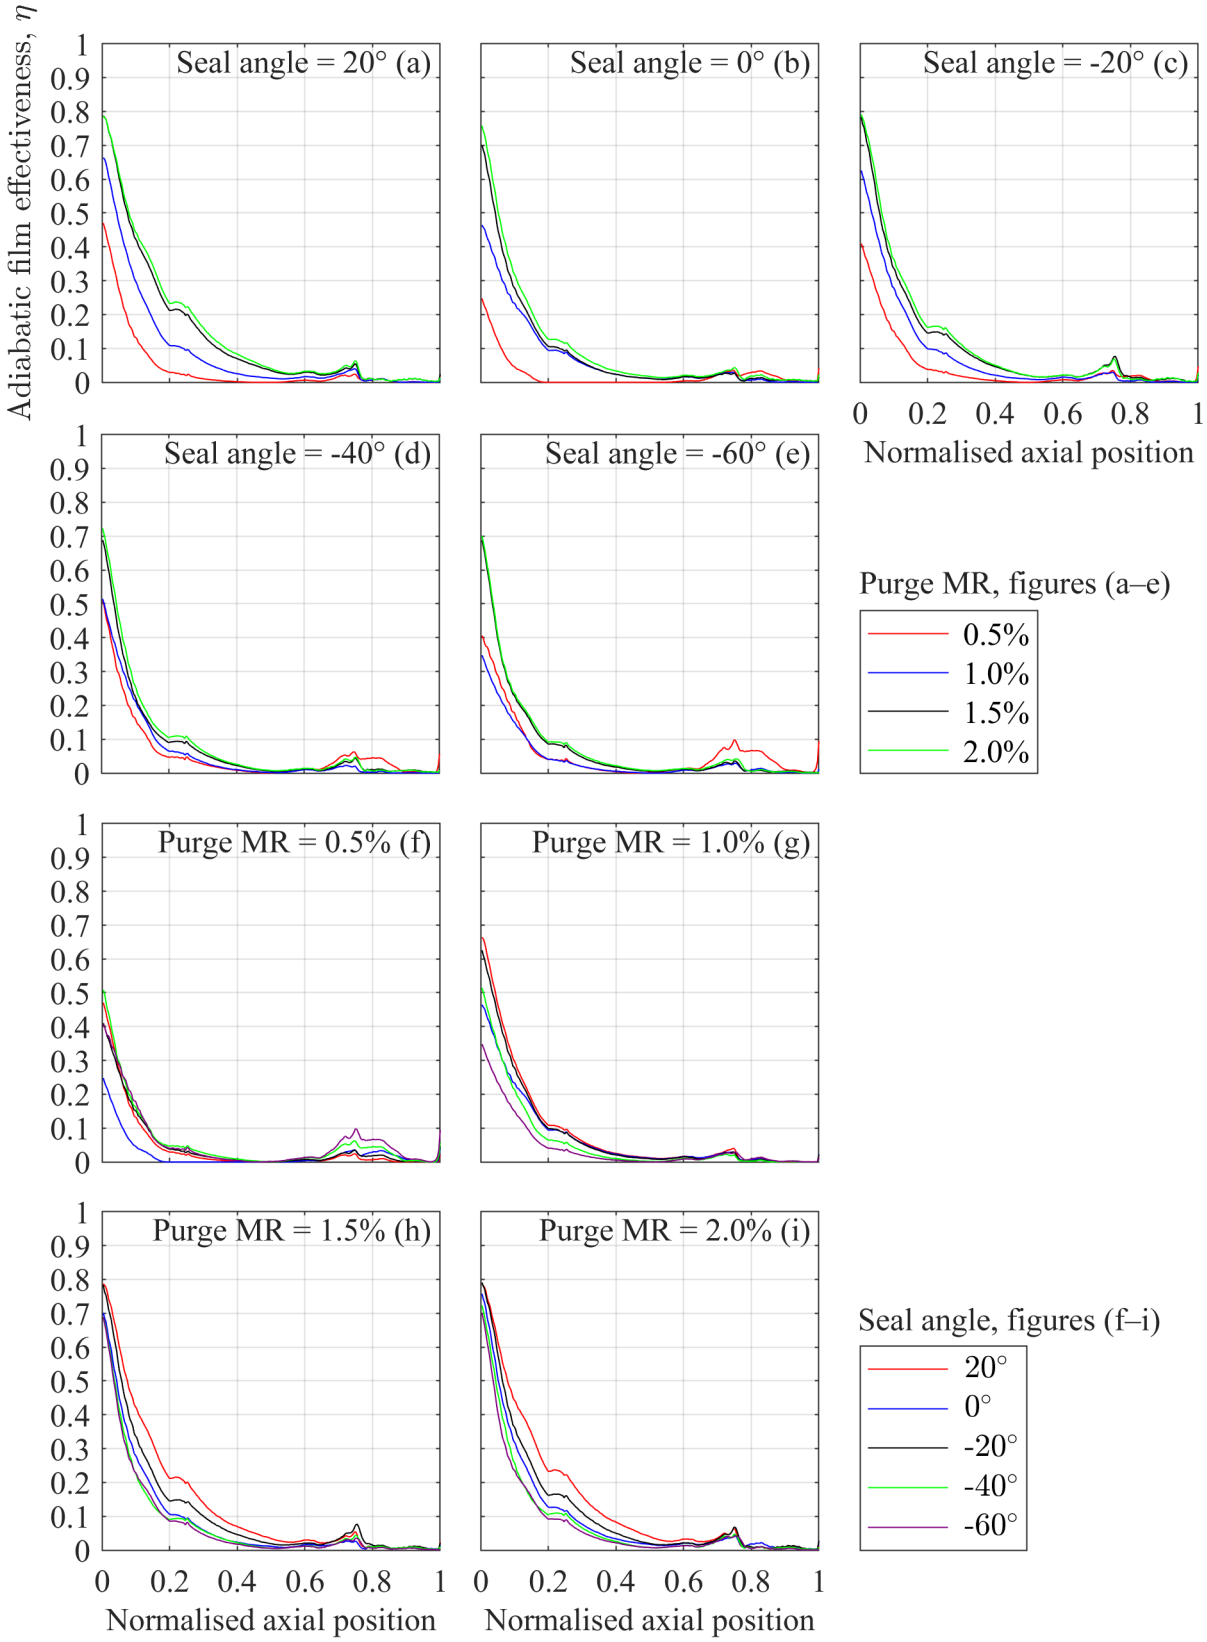


Fig. 7.10 Axial profiles of corrected adiabatic film effectiveness, averaged in the tangential direction down one blade passage. (a)–(e) compare results for each seal angle at four purge mass ratios and (f)–(i) compare results for 5 seal angles at each of the four purge mass ratios.

leakage from the slashface slot and downwash of the partially mixed coolant within the passage vortex.

In plots (a)–(e) in Figure 7.10 it can be seen that higher purge mass ratios result in better cooling at all axial locations for each seal angle. The profiles are all similar but absolute values of adiabatic film effectiveness are larger at larger mass ratios. Significant improvements in cooling performance occur when increasing the purge mass ratio from 0.5% to 1.0% and from 1.0% to 1.5%. This is demonstrated most clearly by the 20° seal angle in which the initial peak adiabatic film effectiveness increases from 0.47 to 0.66 and then 0.79 for purge mass ratios of 0.5%, 1.0% and 1.5% respectively. The lower, horizontal plateau values in these three cases increase from 0.040 to 0.10 and then 0.22 for purge mass ratios of 0.5%, 1.0% and 1.5% respectively. There is however very little increase in cooling performance when increasing the purge mass ratio from 1.5% to 2.0% which may be due to reaching a critical purge flow velocity at which the purge begins to simply feed more directly into the passage vortex rather than covering the platform surface.

Comparing the axial profiles of the seal angles at each purge mass ratio shows that the 20° seal provides comparatively better cooling, as the values of adiabatic film effectiveness are the highest at all axial locations. The increased cooling performance is marginal at the nominal purge mass ratio of 1.0%, with the 20° and −20° seal angle profiles being very close together. However the 20° performs significantly better at purge mass ratios of 1.5% and 2.0%. The profiles can be seen to be arranged in an approximate ranked order, with the larger negative angles having lower adiabatic film effectiveness values. This demonstrates a clear cooling advantage to turning the purge flow away from the suction surface and towards the pressure surface. The 0° seal angle cases do not follow the ranked order, with generally lower values than the −20° seal angle cases. This suggests that the presence of the guide vanes themselves enhances the film cooling performance of the purge flow. This is likely due to the vanes disrupting large-scale flow migration within the front seal.

7.3.4 Comparison of area-averaged platform adiabatic film effectiveness values

To further display the differences between the purge flow configurations the overall adiabatic film effectiveness distributions have been reduced to singular representative value for each case. These are the area-averaged values, which represent the overall film cooling performance in each case. Values are plotted in Figure 7.11 for each of the purge flow angles and as a function of purge mass ratio.

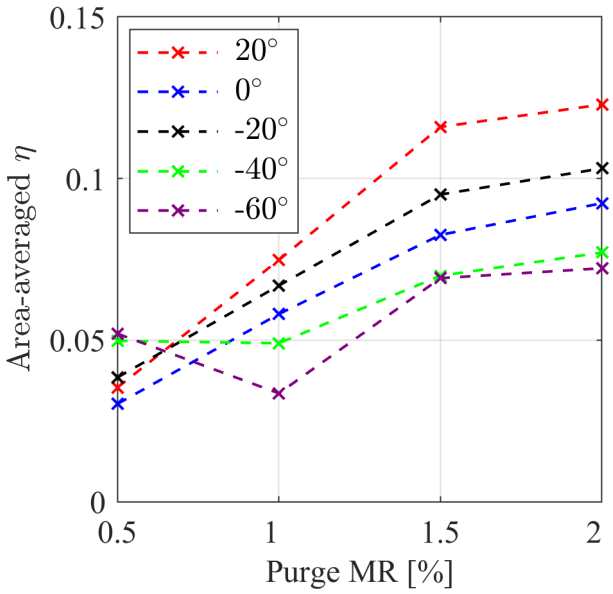


Fig. 7.11 Overall comparison of area-averaged adiabatic film effectiveness for 20 test conditions covering 4 mass ratios and 5 seal angles.

The two main trends previously discussed are clear in Figure 7.11a which shows an increase in area averaged adiabatic film effectiveness at higher purge mass ratios, and a ranked order of seal angles. The rates of increase (the gradients of the dashed lines) are greatest between purge mass ratios of 0.5% and 1.5%. Between 1.5% and 2.0% these gradients level off, showing diminishing returns on film cooling performance from using greater and greater amounts of coolant. At the nominal design purge mass ratio, the 20° seal has an area-averaged adiabatic film effectiveness of 0.075 which is 220% larger than that of the −60° seal with a value of 0.034. The −60° seal is the closest representation to the expected flow angle within an engine of the seal angles tested. Whilst the overall fraction of the blade platform cooled by the purge flow is small,

leading to these small values, the increase in cooling performance at the same purge mass ratio due to altering the purge flow angle is significant.

7.4 Downstream total pressure loss measurements

The enhanced film cooling performance of the front seal purge flow, achieved by directing flow towards the pressure surface and increasing the purge mass ratio, has now been established. Any penalties to cycle efficiency, caused by using a larger amount of coolant or increasing aerodynamic losses through the bladerow, are of equal importance to the benefit of better film cooling performance. Downstream traverses covering an area of one blade passage up to midspan were conducted using a 3-hole pressure probe. Distributions of total pressure loss coefficient, Y_p, were determined from pressure measurements in each case to allow for any differences caused by varying purge mass ratio or flow angle to be observed.

7.4.1 Comparison of total pressure loss area distributions

Downstream loss coefficient distributions for each of the 20 test conditions are given in Figure 7.12. Black contour lines have been overlaid onto interpolated data maps to highlight the profile shape. These contour bands are loss coefficients of 0.2, 0.3 and 0.4. The general profiles are very similar, with differences between test conditions being subtle. In each case the dominant feature is an 'L', or 'boot', shaped region of high loss, greater than 0.3, centred on the 0.4 passage fraction (tangential direction). This is caused by the trailing edge wake. At an approximate span fraction of 0.2 the loss region extends further into the passage along the tangential direction, forming the characteristic L-shape. This feature is caused by the interaction of the swirling passage vortex and the trailing edge wake. The swirling passage vortex flow entrains the blade wake pulling it up and into the passage. The blade wake would otherwise be expected to be a straight band from the hub up to the casing following the blade trailing edge profile. A loss core, the peak loss coefficient value, occurs at the corner of this interaction as highly turbulent mixing occurs. An additional region of loss, of an approximate value of 0.2, occurs in the bottom-left of

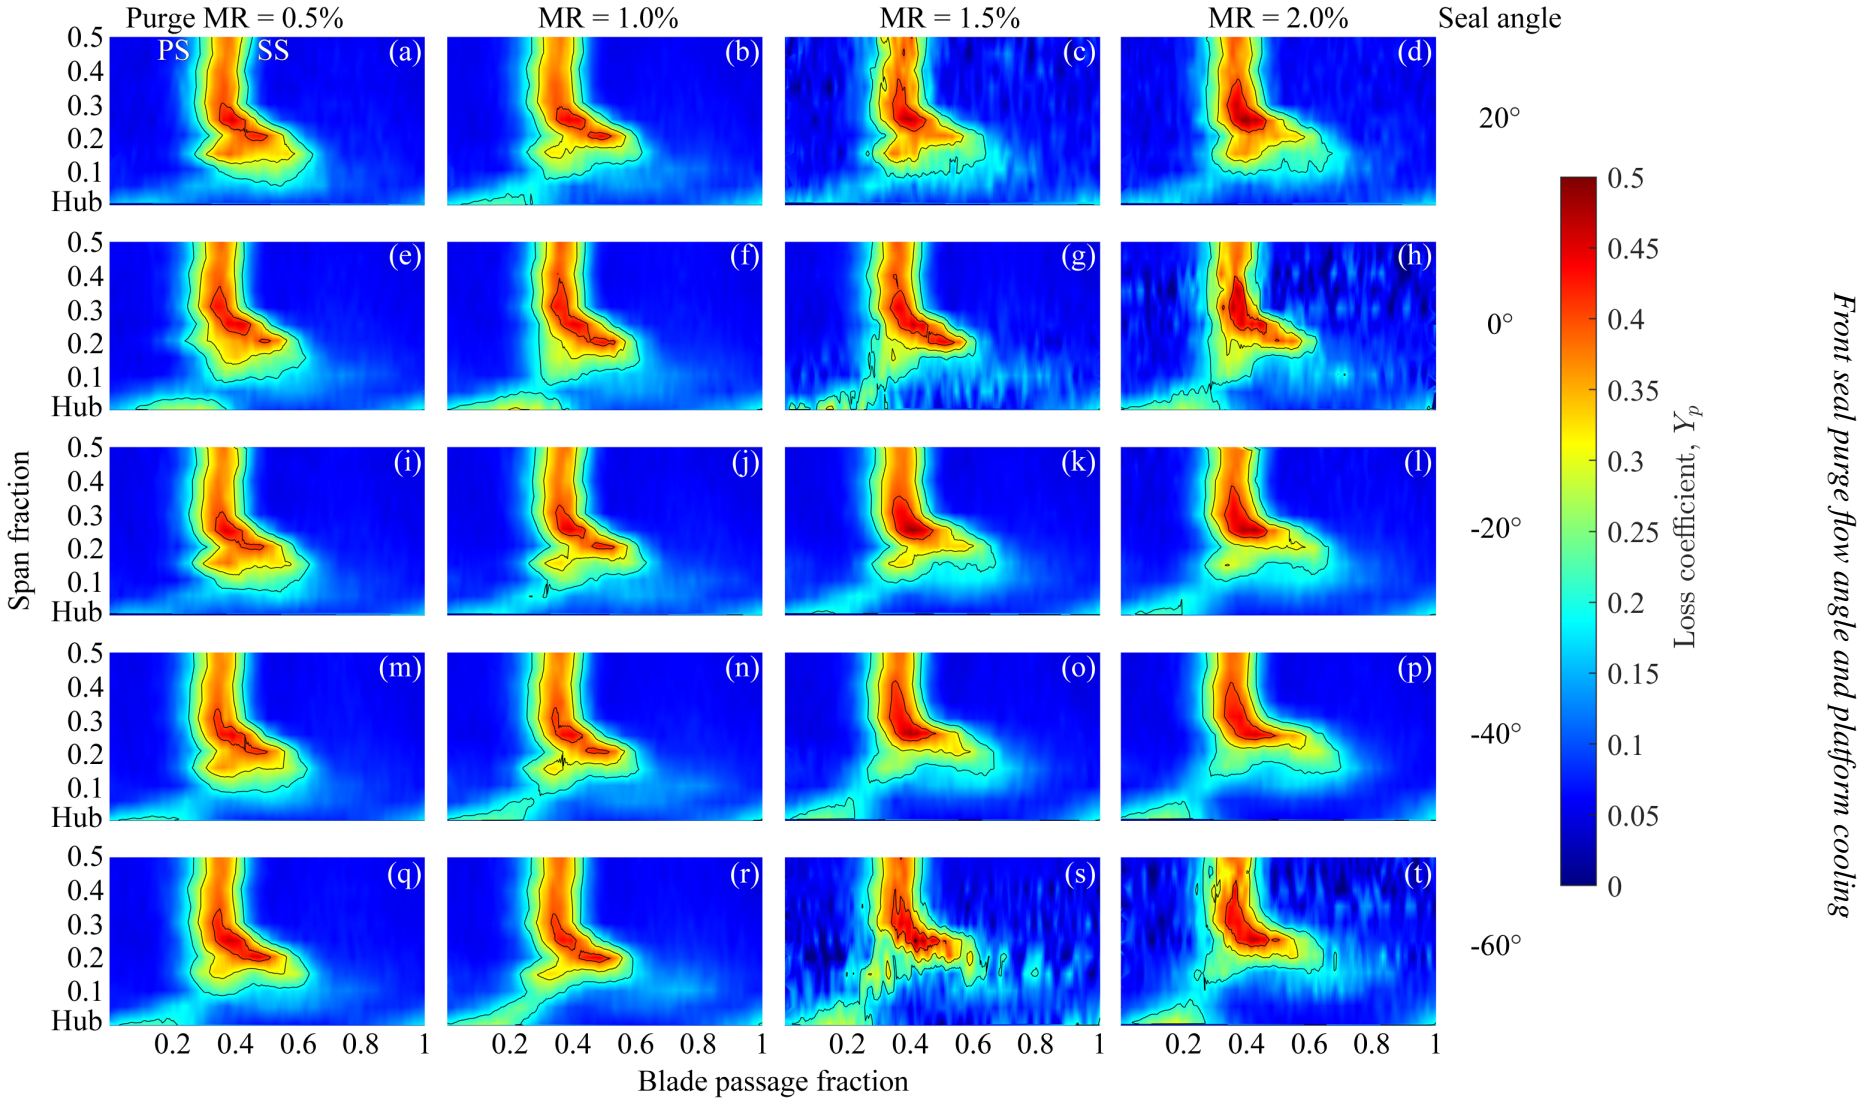


Fig. 7.12 Total pressure loss coefficient distributions of 20 test conditions with varying seal angle and purge mass ratio. Distributions measured across an area covering one blade passage up to midspan, in a plane 0.06 axial chord fraction (2 mm) downstream of the blade trailing edge.

each distribution. This is the part of the blade wake associated with the fillet between the blade and platform and so is concentrated near the hub.

Some of the results in Figure 7.12—(c), (d), (g), (h), (s), and (t)—have increased measurement noise in the mainstream low-loss measurement areas. This is likely due to vibration of the probe tip in the mainstream transonic flow in these cases. The fluctuations are relatively small compared to the primary loss core within the blade wake and passage vortex wake interaction.

The overall shape of the high loss coefficient areas can be seen to change only subtly between cases. For negative seal angles, as the purge mass ratio is increased, the loss region extends further into the blade passage and is pulled higher up into the span. More coolant having been entrained by the passage vortex may have strengthened it so that the trailing edge wake is affected more than in lower purge mass ratio cases. In general, the more the loss region extends up into the passage and span, the lower the loss value at the 'heel' of the boot shaped region which suggests that in these cases it is the passage vortex more strongly interacting with the blade wake and pulling the fluid away from its usual flow path.

7.4.2 Comparison of spanwise and tangentially averaged loss coefficient profiles

The differences in total pressure loss coefficient distributions in Figure 7.12 are subtle and difficult to meaningfully parse from the complete images. In order to examine the differences in a more detailed way, the distributions were averaged in the spanwise and tangential directions to give two spatial loss profiles for each test case. These are plotted in Figures 7.13 and 7.14 respectively. Similarly to the axial adiabatic film effectiveness profiles in Figure 7.10 these plots show comparison groups of both seal angle and purge mass ratio.

The spanwise loss coefficient profile, in Figure 7.13, can be seen to be largely constant at span fractions greater than 0.35 with a tangentially-averaged loss, due to the blade wake, of 0.11. This is the loss associated with just the blade trailing edge wake. Higher losses occur between span fractions of 0.10 and 0.30, where the blade wake interacts with the passage vortex. The peak loss in each case is similar, in the range 0.16–0.17. In Figure 7.13 (a), (c), (d), and (e) two

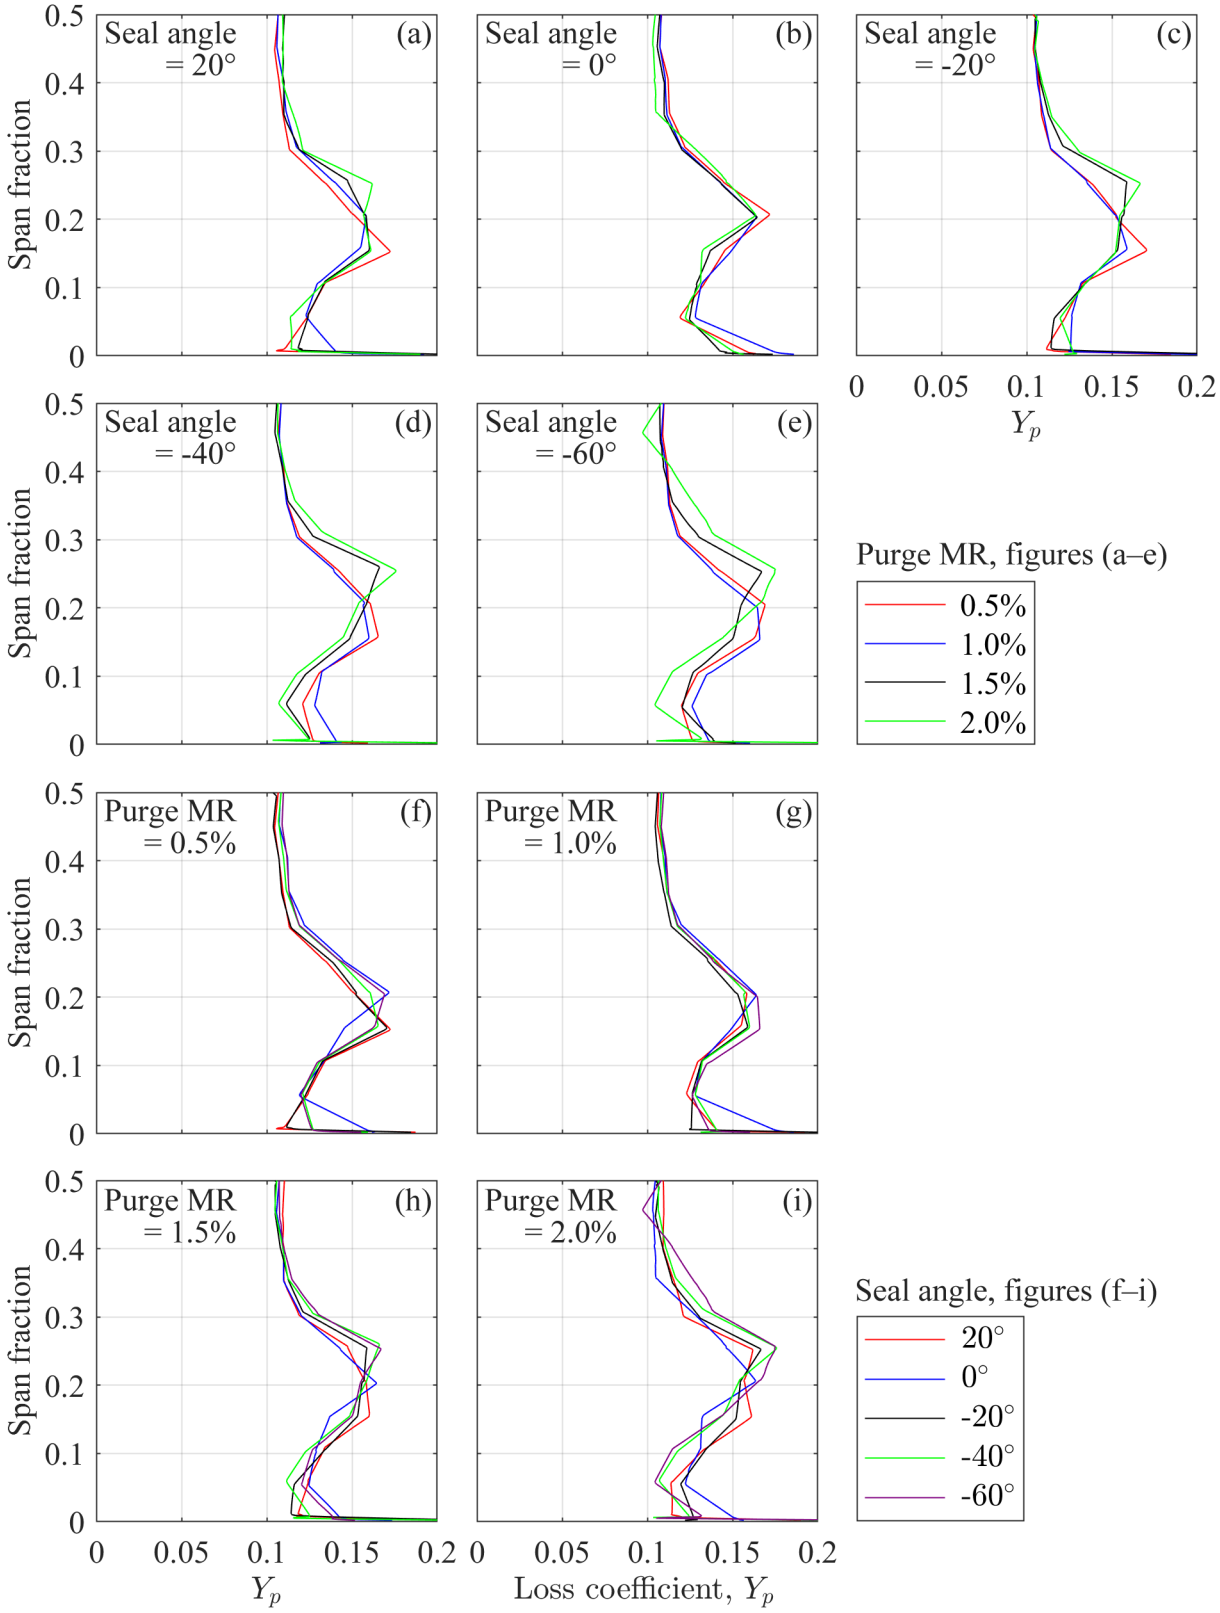


Fig. 7.13 Spanwise profiles of total pressure loss coefficient, averaged in the tangential direction across an area covering one blade passage. (a)–(e) compare results for each seal angle at four purge mass ratios and (f)–(i) compare results for 5 seal angles at each of the four purge mass ratios.

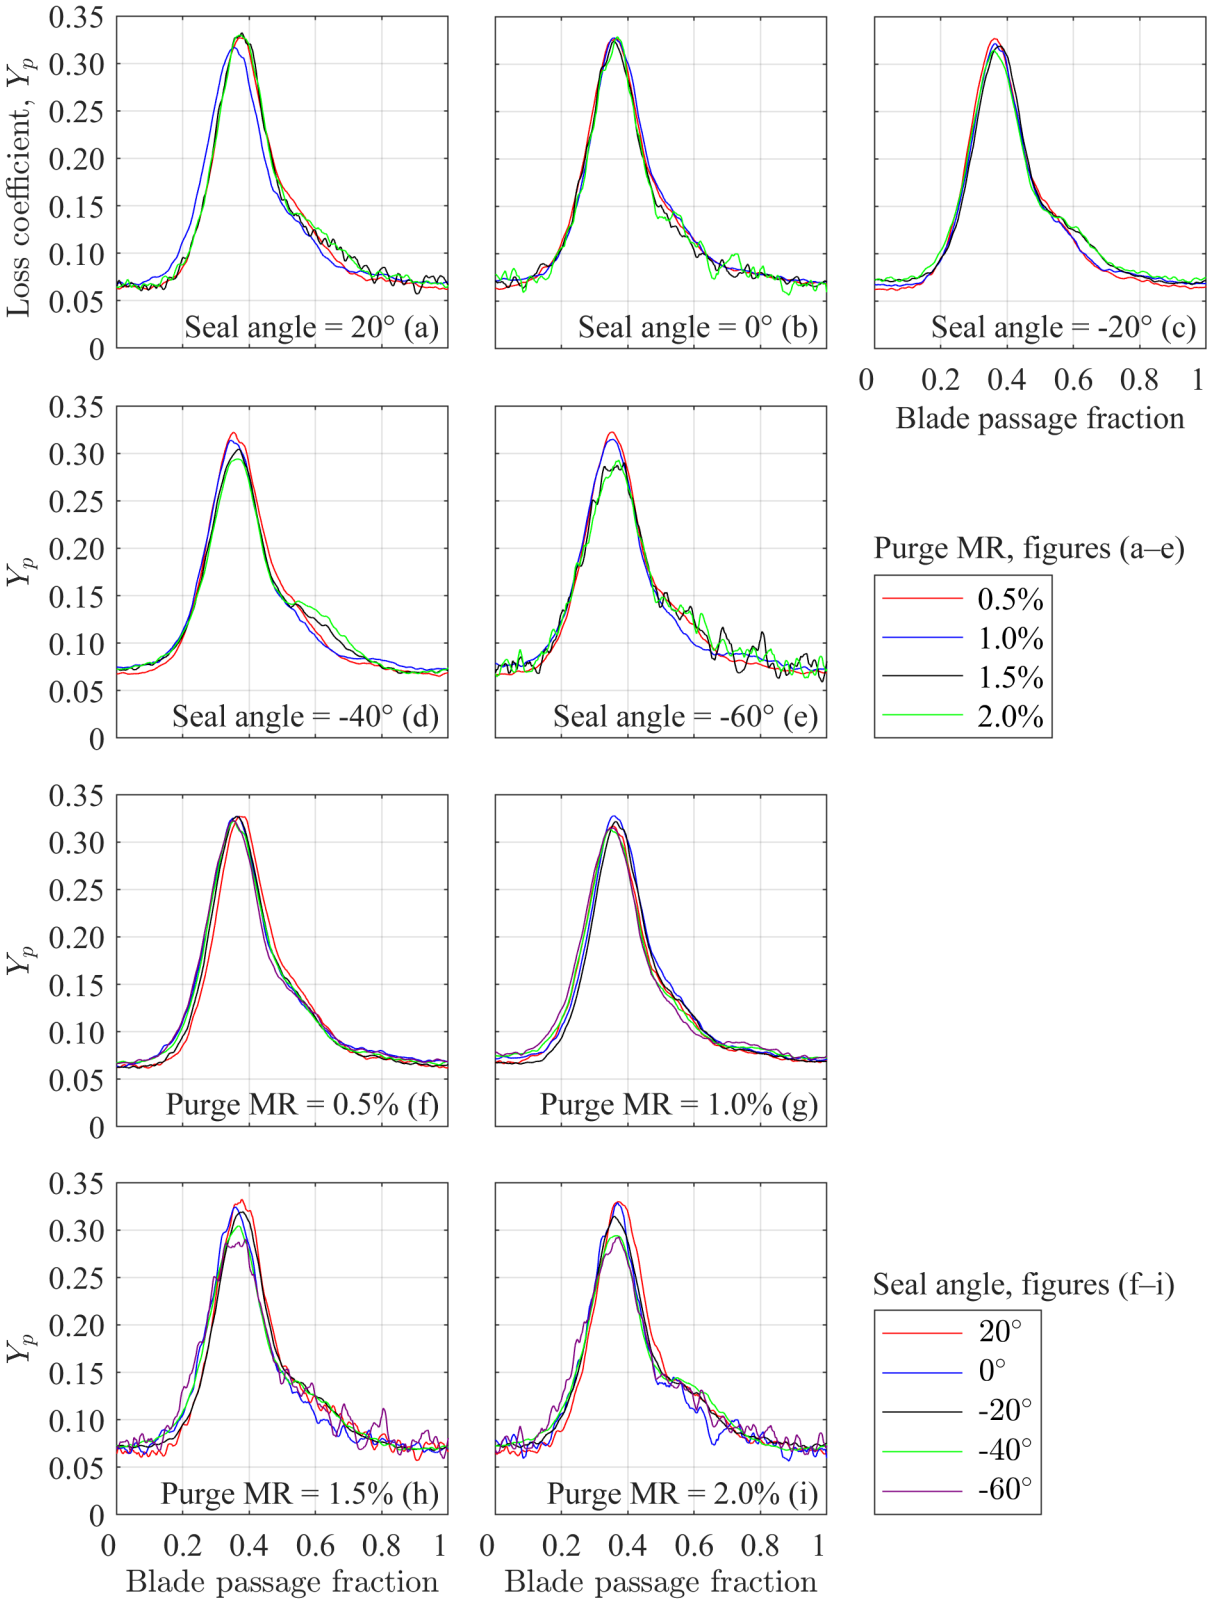


Fig. 7.14 Tangential profiles of total pressure loss coefficient, averaged in the spanwise direction across an area covering one blade passage. (a)–(e) compare results for each seal angle at four purge mass ratios and (f)–(i) compare results for 5 seal angles at each of the four purge mass ratios.

distinct peaks at span fractions of 0.18 and 0.25 can be seen in the groups of loss profiles for each seal angle. The peaks higher up in the span are associated with test cases run with larger purge mass ratios of 1.5% and 2.0%. In contrast, the 0° seal profiles are remarkably uniform for each of the tested purge mass flows. This suggests that the differences present in the other seal geometries are due to the presence of the internal seal guide vanes.

Comparisons of span loss profiles between seal angles, shown in plots (f)–(i) in Figure 7.13 show no clear differences caused by seal angle. A broad difference is that the peak loss generally occurs over a larger span fraction for the seal angles in which guide vanes are present than in the 0° seal cases which have sharp peaks.

The spanwise averaged tangential loss profiles in Figure 7.14 are all similar with no major differences between any of the cases. A large central peak associated with the blade wake occurs at approx 0.35 blade passage fraction. This peak decays to a base mainstream loss coefficient of 0.06 and is typical of a blade wake loss profile. Within the mainstream flow a loss coefficient closer to 0 may be expected and the value of 0.06 determined here may therefore be associated with interactions between the transonic mainstream flow and probe. One small difference can be seen, most clearly in plots (c)–(e) in Figure 7.14, in which negative seal angles have increased loss coefficients in blade passage fractions in the range 0.6–0.75 at higher purge mass ratios of 1.5% and 2.0%. These elevated loss values are due to stronger passage vortex interactions with the trailing edge due to a strengthening of the passage vortex from having ingested a larger amount of coolant from the front seal purge.

7.4.3 Comparisons of maximum and mass-weighted average total pressure loss coefficient

A singular value which characterises the overall aerodynamic loss through the blade passage associated with each is the mass-weighted average total pressure loss coefficient. Mass-weighting the average attempts to account for the relative amount of flow through the high loss regions of the blade wake. Mass flux distributions across the traverse area were derived from probe-measured pressures. Mach number distributions are obtained from the probe calibration and

the total temperature of both coolant and mainstream are assumed to be the same and equal to that measured at the inlet of the cascade since the traverses were conducted 'cold' without mainstream heating. This allows for the mass flux distribution to be calculated and used to mass-weight an average of the total pressure loss coefficient across the surveyed area. Results are plotted in Figure 7.15a.

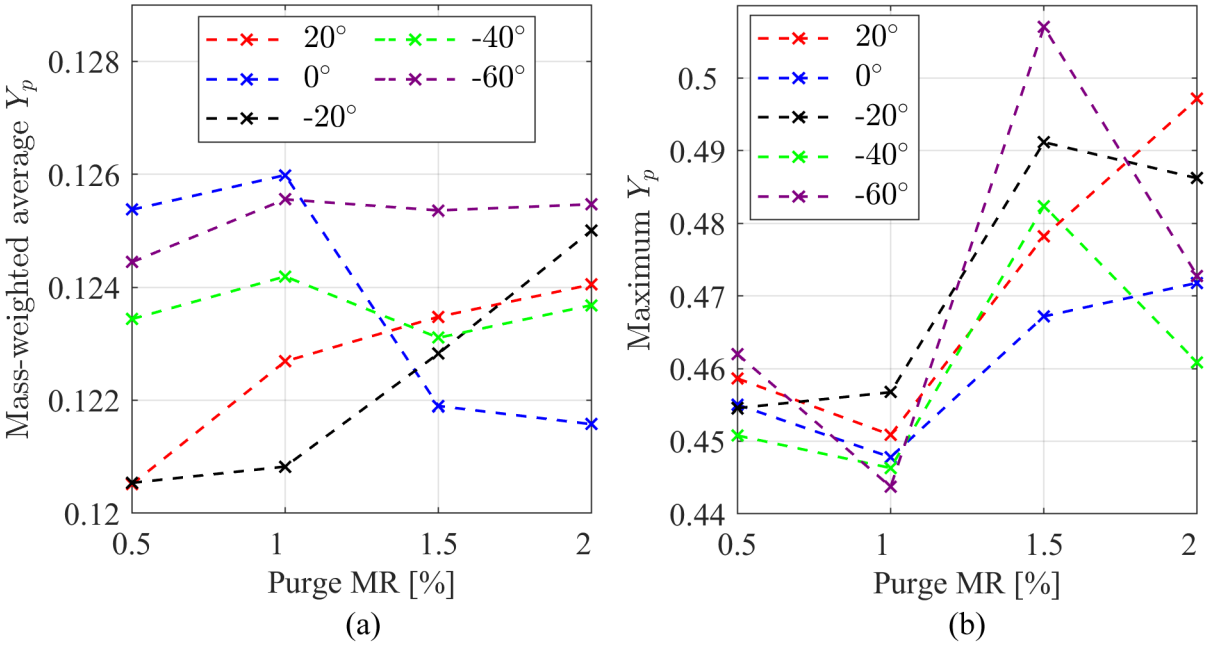


Fig. 7.15 Comparison of all 20 test conditions for (a) mass-weighted average total pressure loss coefficient, and (b) maximum total pressure loss coefficient, against purge flow mass ratio.

Perhaps surprisingly, although subtle differences in downstream loss coefficient distributions were observed there are very little differences in the overall mass-weighted average of loss coefficient across the surveyed areas. No clear trends exist between cases for either changing seal angle or purge mass ratio, with the maximum difference in values being less than 0.006 equivalent to less than 5%. If we instead look at the peak loss coefficient values (averaged over small areas to avoid anomalously high measurements skewing results) which occur in each case, Figure 7.15b, then a general trend of higher peak loss occurring at higher purge mass ratios can be seen which may be expected as more coolant is being introduced into the passage vortex. Also of interest is that the maximum loss coefficients for each seal at a purge mass ratio of 0.5% are higher than at the nominal design point of 1.0%. This may be caused by additional losses generated by the ingestion and subsequent re-emission and mixing of mainstream fluid into the

front seal. It should be noted that the dominating feature of loss within the passage is the passage vortex, inherently linked to the roll-up of the inlet boundary layer which is expected to be thicker than that in an engine environment. The thick boundary layer in the cascade may cause the passage vortex to be relatively stronger, containing more fluid, which may be dominating any subtleties caused by subsequent ingestion of front seal purge flow into the vortex. Additional CFD studies may be able to quantify the effect of the inlet boundary layer on the strength of the passage vortex and hence whether this result, of little-to-no impact of front seal purge on downstream loss holds in an engine environment.

These results suggest that the overall aerodynamic losses associated with the flow through the blade passage are dominated by the blade wake and its interaction with the passage vortex. All of the purge flow ultimately ends up entrained by the passage vortex, however the amount of coolant is relatively small and largely has no overall impact on the overall total pressure loss coefficient. Subtle changes to the total pressure loss coefficient distribution seem to have little overall effect on the total value. This is important as it means that the flow angle of the purge flow can be altered, to improve film cooling performance towards the leading edge of the platform, without incurring significant additional aerodynamic losses within the blade passage. Methods which alter the purge flow angle are therefore purely beneficial when used with the same amount of sealing flow required to seal the front seal. Increasing the mass ratio of purge flow also seems to have a negligible impact on blade passage losses, however this will have an impact on cycle efficiency due to more coolant being bled from the compressor.

7.5 Conclusions

A summary of the key findings from the work described in this chapter are given below:

- The angle at which purge flows out of the front seal onto the rotor platform depends on the axial and tangential components of velocity. The axial velocity is set by the available front seal slot flow area and the purge mass ratio. The tangential velocity component is a function of rotor speed, is in the direction of rotor motion but is lower than the tangential velocity of the rotor due to shearing between stationary and rotating surfaces within the

front seal. This results in the purge flow being directed towards the blade suction surface in the rotor relative frame.

- The purge flow acts as a cooling film across the blade platform. A series of experiments were conducted to determine surface adiabatic film effectiveness distributions across the blade platform with purge flow angles of $20°$, $0°$, $-20°$, $-40°$, and $-60°$ set by guide vanes within the front seal. Purge flow directed towards the pressure surface of the blade had the best cooling performance, covering a larger area of the platform surface and with higher adiabatic film effectiveness values. Negative seal angles had progressively worse cooling performance as the angle increased towards the suction surface, covering a small platform area and with lower adiabatic effectiveness values.

- Increasing the purge mass ratio, the amount of coolant through the front seal, leads to an increase in both the area cooled and absolute values of adiabatic film effectiveness distribution. Purge mass ratios of 0.5%, 1.0%, 1.5% and 2.0% were tested.

- Film cooling provided by the purge flow was limited only to the leading 20% axial chord of the blade platform for all tested conditions as it is swept across the passage towards the suction surface and lifts off from the platform due to being entrained by the passage vortex.

- Purge flow is ultimately all entrained by the passage vortex, either joining the vortex flow as the leading horseshoe vortex rolls up into the passage or being swept across the passage by crossflow to join the vortex nearer to the blade suction surface.

- A secondary tangentially-flowing vortex forms within the gap between the front seal slot and blade leading edge, which entrains part of the purge flow and emerges into the mainstream close to the mid-passage. The seal angle affected the amount of purge flow entrained within this vortex. For a positive seal angle (towards the pressure surface) the purge flow was evenly distributed along the seal width and joined the tangential vortex gradually. Negative seal angles direct purge flow against the direction of the tangential

vortex, causing a blockage effect, which reduces purge coolant ejection near the blade leading edge and causing more to exit midspan and directly join the passage vortex.

- Downstream area traverses allowed for detailed total pressure loss coefficients to be determined. No significant effects of altering either seal angle or purge mass ratio were observed. Mass-weighted average loss coefficient values of all tested cases varied by less than 5%, with an average value of 0.12. This suggests that the film cooling performance of the front seal purge can therefore be increased by directing the purge flow towards the pressure surface without incurring any additional aerodynamic losses.

- Blade passage losses associated with increasing purge mass ratio are insignificant compared to the dominant loss associated with the blade wakes and interaction of the wake and passage vortex. Cycle efficiency penalties are associated with increasing purge mass ratio, which are unquantified here but which make altering seal angle a better option for enhancing purge flow film cooling performance as this can be done without the addition of more coolant.

- A potential method of altering the purge flow angle is to supply part, or all, of the sealing coolant to the front seal from the upstream NGV platform cooling system. This fraction of coolant could be supplied via angled holes within the seal structure at fixed metal angles so that the mixed-out purge flow is directed away from the suction surface. This method could be used to alter the purge flow angle without using more coolant than current designs, and may actually allow for a reduction in total engine coolant usage as coolant used in the NGV platform cooling system could be reused in the blade front seal.

Chapter 8

Metal effectiveness of platform cooling concepts

In this chapter, distributions of metal effectiveness are presented for three titanium rotor blade designs: a solid baseline; integrated convective passages; and shank-fed impingement array. A total of 16 conditions, spread across the three blade designs, were run. Test conditions varied by the presence of leakage flows, the mass flow of leakage flows, and the mass flow of coolant through each integrated cooling system. Full area metal effectiveness distributions across two blade platforms for each test case are presented and discussed. Comparisons between test cases are assessed using representative quality parameters: area-averaged metal effectiveness to quantify overall cooling performance; minimum metal effectiveness, to quantify the maximum expected metal temperature; non-dimensional mass ratio, MR, which quantifies the aerodynamic cost of supplying coolant.

8.1 Experimental setup and boundary conditions

Experiments were conducted between July and August 2020 and due to very limited time and resource constraints only a small number of test cases were run. Efforts were made to set boundary conditions and perform series of runs which demonstrated the cooling performance, and maximised understanding, of each system within the available time frame. Test configurations varied in which leakage flows were present, and how much mass flow was supplied to the three seal plenums. The front seal block used for all test cases was the '0-degree' seal, with no guide vanes present. The total mass ratio supplied varied in the range 1.1–4.8% between cases. The total mass ratio gives an indication of the cycle efficiency penalty of each system caused by

drawing coolant from the compressor. It is therefore useful as a means to gauge the relative 'costs' of each system when comparing performance. The additional aerodynamic losses due to reintroduction of coolant into the mainstream will vary between test configurations however this was not assessed in these experiments. The current traverse system installed in the MERC facility covers the exit plane of the blade row which is upstream of the rear seal. As the integrated cooling systems both exhaust into the rear seal, a traverse plane downstream would be required to properly assess the aerodynamic impact as the rear seal flow then exhausts as a downstream purge flow. The aerodynamic traverse system was therefore not used for these tests, which reduced runtimes and made it possible to run more configurations. A summary of each of the sixteen test cases (denoted A-P) is given in Table 8.1.

Table 8.1 Summary of sixteen test case conditions for metal effectiveness experiments

Case	Blade	$\mathrm{MR} = \frac{\dot{m}_c\sqrt{T_{0c}}}{\dot{m}_m\sqrt{T_{0m}}}[\%]$			
		Front seal	Shank pocket	Rear seal	Total
(A)	Baseline	1.1	0*	0	1.1
(B)		2.0	0*	0	2.0
(C)		1.1	1.8	1.3	4.2
(D)		2.0	1.8	1.3	5.1
(E)	Impingement array	0	0.88*	1.3	2.2
(F)		0	1.1*	1.3	2.4
(G)		0	1.3*	1.3	2.6
(H)		1.1	0.88*	1.3	2.2
(I)		1.1	1.1*	1.3	3.5
(J)		1.1	1.3*	1.3	3.7
(K)	Convective passages	1.1 (0.16†)	0*	0	1.1
(L)		1.7 (0.18†)	0*	0	1.7
(M)		2.0 (0.19†)	0*	0	2.0
(N)		1.1 (0.16†)	1.8	1.3	4.2
(O)		1.7 (0.18†)	1.8	1.3	4.8
(P)		2.0 (0.19†)	1.8	1.3	5.1

* Silicon-sealing between blades

† Amount through convective passages as opposed to front seal purge.

Baseline blades (cases A–D) were cooled only by leakage flows and were tested at four conditions. Front seal purge flows with mass ratios of 1.1% and 2.0% were tested. Each front

seal mass ratio was tested with and without silicone-sealed shanks to allow leakage flow between blade slashfaces. In cases with no sealant between blade slashfaces the total coolant mass ratio into the shank cavity was 1.8%. Coolant flow into the shank cavity is split between the blade pseudo multi-pass system (exhausting at the blade tip) and leakage between blade slashfaces into the mainstream. The split of coolant was not measured. As heat conducts through the platform (especially near to the blade fillet) to the blade, the coolant flowing through the blade itself is therefore expected to have some cooling effect on the platform. All of the coolant supplied to the shank cavity is therefore considered, when comparing coolant usages between cases, as doing cooling duty for the platform.

The impingement array cooled blades (cases E–J) draw coolant from the shank cavity and exhaust to the rear seal. The coolant mass ratio through these systems is therefore set by the pressure ratio between the shank cavity and rear seal. Since flow tests—like those conducted for the convective passage blades which allowed for capacity curves to be determined—were unable to be completed for these blades, the measured pressure ratio could not be used to determine the mass ratio through the system. Instead, all tests were conducted with silicone-sealed shanks and the holes at the top of the blades were completely sealed to prevent tip leakage. All coolant supplied to the shank cavity was therefore constrained to flow through the impingement system. The shank was supplied with three coolant mass ratios: the impingement system design point of 1.1%, 0.88% and 1.3% which are approximately $\pm 20\%$ of the design point. These values were selected to give an indication of the impingement systems cooling performance at off-design conditions and were the minimum changes in mass ratio possible by covering single holes in the existing coolant setting plates.

The convective passage blades (cases K–P) draw coolant from the front seal and exhaust to the rear seal. The mass flow rate through each passage is dependent on the pressure ratio between seals. A total of six conditions were tested with the convective passage blades. Half of the cases were conducted with silicone-sealed shanks such that the only leakage onto the mainstream side of the platform was from the front seal. The other 3 cases were tested with open shank slashfaces with the same total shank cavity mass ratio as in the baseline cases of 1.8%. The front seal cavity was supplied with purge flow mass ratios of 1.1, 1.7, and 2.0%. For each of the six cases the

mass flow ratio through the platform passages was estimated from the pressure ratio between front and rear seals and the capacity curves determined in Chapter 4. The front seal purge flow mass ratio was calculated as the difference between the supplied mass ratio and the flow through the platform convective passages.

8.2 Baseline blades: results and discussion

Surface metal effectiveness distributions for two blade platforms for baseline blade test cases A–D are presented in Figure 8.1. The images are oriented such that the inlet flow to the blade row is at the top of the image and axial flow is towards the bottom of the image. The baseline cases are cooled solely by coolant flowing inside the hub seal structures and leakage flows from these seals into the mainstream gas path.

The effect of leakage from the front seal can be seen in all four cases by the regions of high metal effectiveness (greater than 0.7) along the leading portion of the platforms, upstream of the blade leading edge. This part of the platform is characterised by a curved surface geometry which forms the rotor-side of the rotor-stator seal interface. High values of θ in this area are therefore expected as fresh coolant is continuously supplied to the seal. A repeated wave-pattern can be seen in these regions of high θ, caused by the high pressure of stagnated mainstream flow at the blade leading edge. The high pressure in this region prevents coolant from flowing near to the blade leading edge and instead forces it to travel laterally along the seal before ejecting onto the platform surface. This occurs in the middle of the passage between pressure and suction surfaces, where the pressure distribution is relatively low compared to the leading edge stagnation region and along the blade pressure surface. This can be seen by the peaks of the wave-pattern, marked (1). The front seal leakage onto the platform is quickly convected towards the blade suction surface due to the cross-passage pressure gradient. The flow is entrained by the passage vortex and as such contributes very little to platform cooling from approximately halfway through the blade passage.

The effect of doubling leakage mass flow (from Case A to B) is to slightly increase values of θ in the previously described regions, most dramatically along the leading edge of the platform

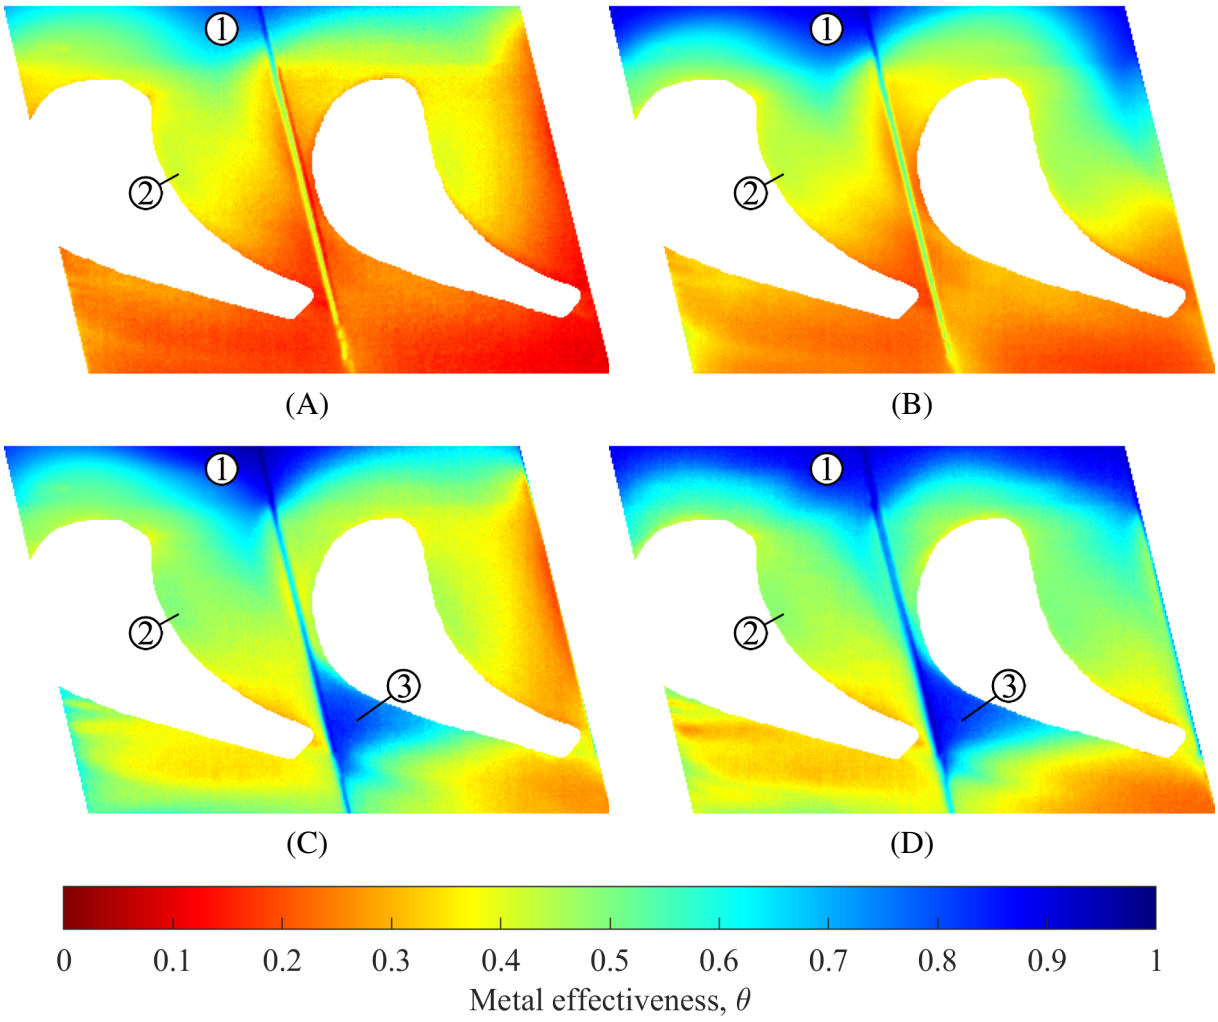


Fig. 8.1 Metal effectiveness distributions of two platforms for baseline blades. Cases (A) and (B) have front seal leakage and silicone-sealed shanks. Cases (C) and (D) have front seal and shank slot leakage.

within the seal. The additional cooling and higher values of θ between cases A and B decays along the blade passage. The distributions of both cases become more similar towards the trailing edge since the secondary flows sweeping coolant into the passage vortex dominate and lift all of the purge flow away from the platform. Towards the trailing edge of the platform, very low values of θ (below 0.2) are observed in both cases. The limited increase in cooling effectiveness achieved by doubling leakage mass flow is unlikely to be worth the additional penalty to engine cycle efficiency from using more coolant.

Towards the pressure side of the blade, a region of low metal effectiveness might be expected due to no purge flow coolant reaching this area. There is however a region in the range $0.35 < \theta < 0.6$ present along the pressure surface, most clearly seen in cases A and B and

marked (2). Figure 8.2 shows a comparison of the metal effectiveness and adiabatic film effectiveness distributions (discussed in Chapter 7) for Case A. The purge flow provides no film cooling towards the pressure surface of the blade and therefore cannot be responsible for this region of high metal effectiveness. This region does however correspond with the geometric profile of the fir tree root and loft which attaches to the rear side of the platform. These profiles are superimposed on the right platform in Figure 8.2a.

It is speculated that this unexpectedly-high θ region is therefore caused by heat conducting from the platform into the underlying fir tree root, which constitutes a large portion of the blade volume. In cases A and B the shank is sealed and the air within it is therefore stagnant and gradually heats up. This high θ region is likely to be a residual transient effect in these experimental cases. This is because the fir tree root is held in a solid steel block, representative of the blade disk, and this is unlikely to achieve a thermal steady-state during the relatively short runtime of the experiments. However, in cases C, D, and in real engine configurations, the fir tree root is surrounded by continuously replenishing coolant (due to topping up seal flows). The actual steady-state temperature of the fir tree in these situations is expected to be close to that of the coolant supply temperature and therefore the region of higher θ observed on the platform towards the pressure surface would be expected to be present within an engine environment.

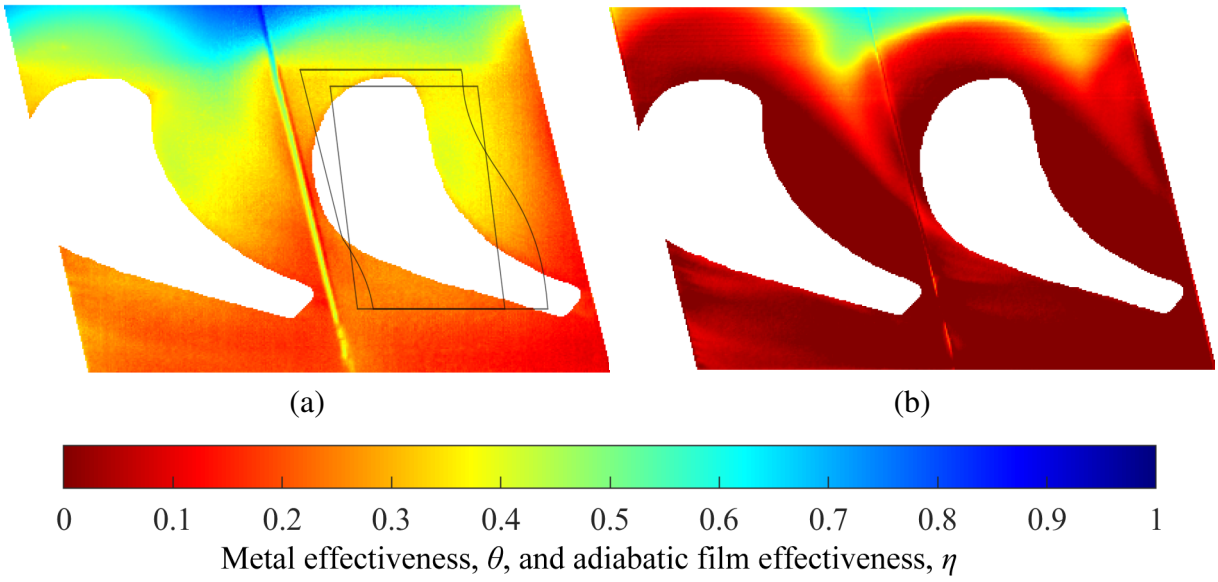


Fig. 8.2 (a) Metal effectiveness distribution across two platforms for Case A, with fir-tree profile superimposed on rightmost platform. (b) Adiabatic film effectiveness distribution from similar test conditions to Case A.

The line of low metal effectiveness between blades, which spans the full height of each image in Figure 8.1, is the clearance gap between neighbouring blade slashfaces. In cases A and B this slashface gap is sealed with liquid silicone sealant, which was not painted and which has different thermal properties to the titanium blades. The values along this line in these cases are therefore not pertinent. In cases C and D the slashface gap does not contain silicone sealant and so high values, close to unity, are achieved since coolant from the shank pocket flows along this gap. As flow accelerates through the blade passage the corresponding drop in static pressure causes coolant from the shank pocket to leak onto the platform surface through the slashface gap. This can be observed in cases C and D where a wedge-shaped region of high θ is present towards the rear of the platform, marked (3), emanating from the slashface gap. This leakage flow is swept toward the suction surface and is ultimately entrained by the passage vortex. The same wedge region is not the same across both blade platforms due to differences in gap tolerances and therefore the amount of coolant leakage. The leftmost blade neighbours a glass-filled nylon blade (not shown) which was slightly oversized (due to manufacturing tolerance variations) and the interference between components has almost entirely sealed the gap. Engine leakage is expected to fall within these two examples and will vary with damper-seal effectiveness and part tolerances.

The results of the baseline blade tests show that leakage flows partially cool some regions of the platform. The film cooling provided by the front seal purge flow results in high metal effectiveness values towards the leading edge of the blade platform and the distributions in this area are similar to those of adiabatic film effectiveness reported in Chapter 7. Additional film cooling is provided by leakage flow through the gap between adjacent blades and this cools a region of the platform between the slashface gap and the blade suction surface. An unexpected high region of metal effectiveness occurs towards the pressure surface, caused by heat being conducted from the platform into the fir tree root and then dissipated into seal coolant. Low metal effectiveness regions which will require additional cooling are found: towards the trailing edge of the platform, starting approximately 50% through the passage; near to the blade leading edge; and, in low slashface-leakage conditions, along the slashfaces.

8.3 Impinging jet array: results and discussion

Surface metal effectiveness distributions for two blade platforms for the impingement array design, cases E–J, are presented in Figure 8.3. A wire-frame of the impingement system has been superimposed over the metal effectiveness distributions to aid in visualisation. Cases E–G have no leakage flows into the mainstream gas path and hence the platform cooling and resultant metal effectiveness values are due entirely to flow through the impingement system and within the rear seal. In these cases the leading edge of the platform, along the seal surface up to the blade leading edge marked (1), have low effectiveness values (below 0.25) which is in stark contrast to the baseline cases with front seal purge flow.

Regions of increased effectiveness (in the range $0.20 < \theta < 0.50$) lie along each of the impingement system ducts. The platform areas which lie above the first two rows of impingement holes towards the leading edge of the platform, marked (2), have lower values of θ. There are several potential contributing factors. The heat load towards the leading edge is relatively high, compared to elsewhere along the platform, due to the mainstream gas being low velocity and not having been diluted by any coolant from the front seal. Additionally, the amount of coolant mass at this location within the ducts is very low and therefore a relatively small amount of heat can be picked up. The impingement cooling of each individual jet is one part of the overall cooling effect of the system. Additional convective cooling along the ducts occurs as spent coolant, which has already impinged on the platform surface, flows along the duct towards the rear seal exhaust. Another explanation for the lower performance of the first impinging jets may be a non-uniform inlet flow to each jet. The top end of the impingement cavity is approximately in line with an internal wall (representing the sealing faceplates between blades and front seal cavity) and so the mass flow into jets closest to this internal wall may be reduced.

The interaction between convecting cross-flow and impinging jets is complex and can either enhance or diminish cooling performance. The results shown here suggest that the crossflow enhances cooling within the system, leading to increased values of θ along the ducts from leading-to-trailing edges of the platform. Past the final row of impinging jets the cross-flow coolant continues to convectively cool the platform, with diminishing effect as heat is transferred

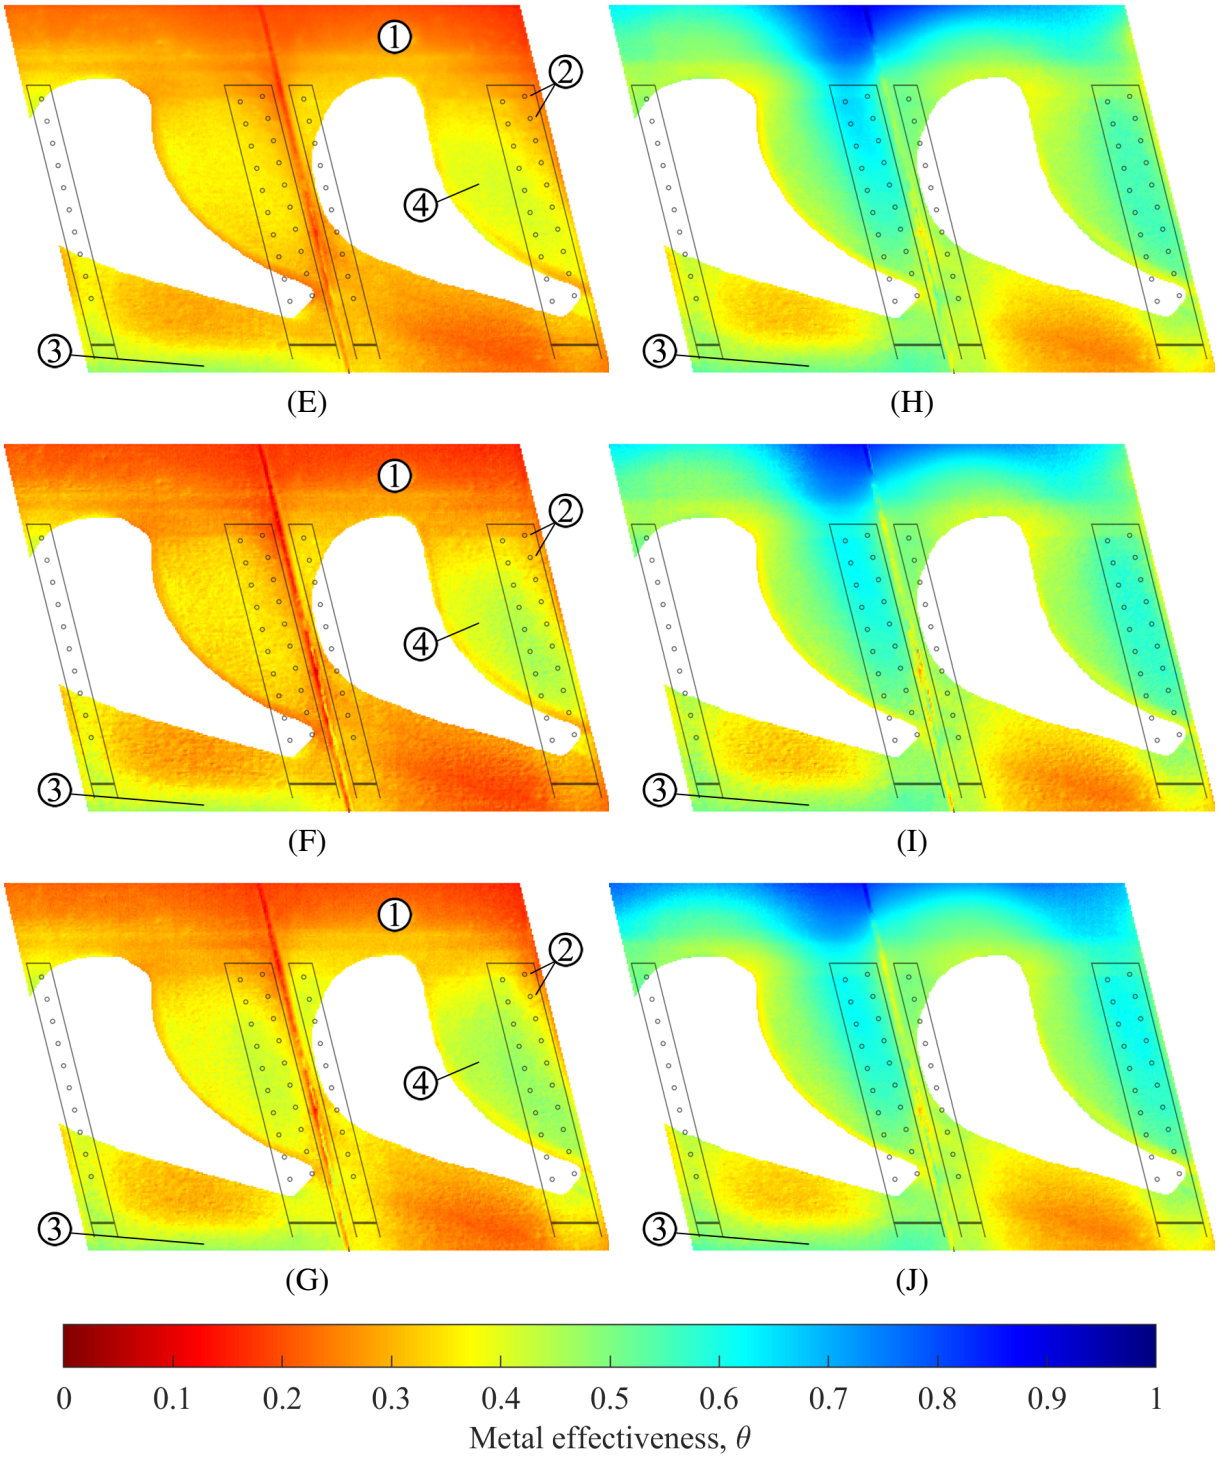


Fig. 8.3 Metal effectiveness distributions of two platforms for impingement cooled blades. All displayed cases have silicon-sealed shanks. Cases (E), (F), and (G) have no front seal leakage. Cases (H), (I), and (J) have front seal leakage.

into the flow. This can be seen by the regions which align with the ducts path near the platform trailing edge with $0.30 < \theta < 0.45$.

Notable differences in metal effectiveness distributions can be seen between left and right blade platforms in cases E–G. A band of relatively high θ (approximately 0.5) can be seen along the trailing edge of the left platform, marked (3). The width of this band corresponds with the size of the overhang of the platform over the rear seal internal structure. Values of metal effectiveness are greater in the region of the platform near to the blade pressure surface of the rightmost platform, marked (4). The cause of this discrepancy is non-uniform distribution of coolant within the rear seal. This is likely due to the pressure gradient across the five blade passages and the geometry of the rear seal cavity. A redesign of the blade cassette insert which includes a secondary distribution plate within the rear seal would be able to produce more even coolant distribution in future studies.

Cases H–J have the same mass flow of coolant through the impingement system as the corresponding cases E–G but have an additional leakage flow from the front seal. This leakage flow effectively cools the leading portion of the blade platforms and can be seen to enhance the cooling performance of the impingement system. The characteristic features of θ distributions remain and the effect of the additional film cooling provided by the leakage flow is to increase values in all locations, and to smooth-out gradients. Although the leakage flow was shown previously to have little impact on cooling past approximately 50% of the passage, in this case it has also led to increased effectiveness values towards the platform trailing edge. This is because the leakage film reduces the heat transferred into the impingement system coolant in the leading portion of the platform. The impingement duct cross-flow therefore has a lower temperature towards the trailing edge resulting in enhanced metal effectiveness values compared to cases E–G with no front seal leakage.

In the case of front seal purge flow and impingement system, values of $\theta > 0.45$ are achieved across the majority of the blade platform. A small region with $0.2 < \theta < 0.35$ exists toward the trailing edge of the platform, near to the suction surface of the blade and in between the slots where the impingement cavities exhaust into the rear seal. If slashface leakage flows were present

then they would partially cool this region and more uniform coverage of the entire platform would be expected to be achieved.

8.4 Convective passage: results and discussion

The surface metal effectiveness distributions for two blade platforms for the convective passage design, cases K–P, are presented in Figure 8.4. All cases have front seal leakage and coolant passing through the integrated convective passages, the amount of which is proportional to the pressure ratio between front and rear seals. Cases N–P also have leakage from the shank cavity.

The leading portion of the blade platforms can be seen to be very well cooled, with metal effectiveness values of $\theta > 0.55$ across the leading region (up to the blade leading edge). The internal convective passages follow the curved profile of the front seal portion of the platform in this region and are unaffected by the mainstream pressure distribution caused by the blade leading edge. This allows the entire leading portion of the platform to be cooled internally, in addition to the film cooling provided by front seal purge flow. In cases K–M small regions of lower effectiveness, marked (1), (approximate minimum $\theta = 0.31$) can be seen along the edges of the platform (along the slashface gap). Due to the pressure field of the blade leading edge, very little leakage coolant flows in this region and so it is cooled solely by the internal convective passages.

The metal effectiveness distribution can be seen, in all cases, to gradually decrease along the blade passages. This is due in part to the diminishing film effectiveness of the leakage flow as it is swept across the passage and lifted away from the platform. As coolant flows along the convective passages it picks up heat from the surrounding metal and increases in temperature. The temperature difference between the coolant and platform is therefore reduced and the cooling effectiveness reduces, leading to lower values of metal effectiveness further along the convective passage lengths. In cases K–M, towards the rear of the platform and past the blade throat, effectiveness values are reasonably uniform in the tangential direction and are similar in all three cases, reducing to a minimum of 0.35 immediately prior to exhausting into the rear seal. In this rear portion of the platform the front seal purge flow has a very limited effect on the

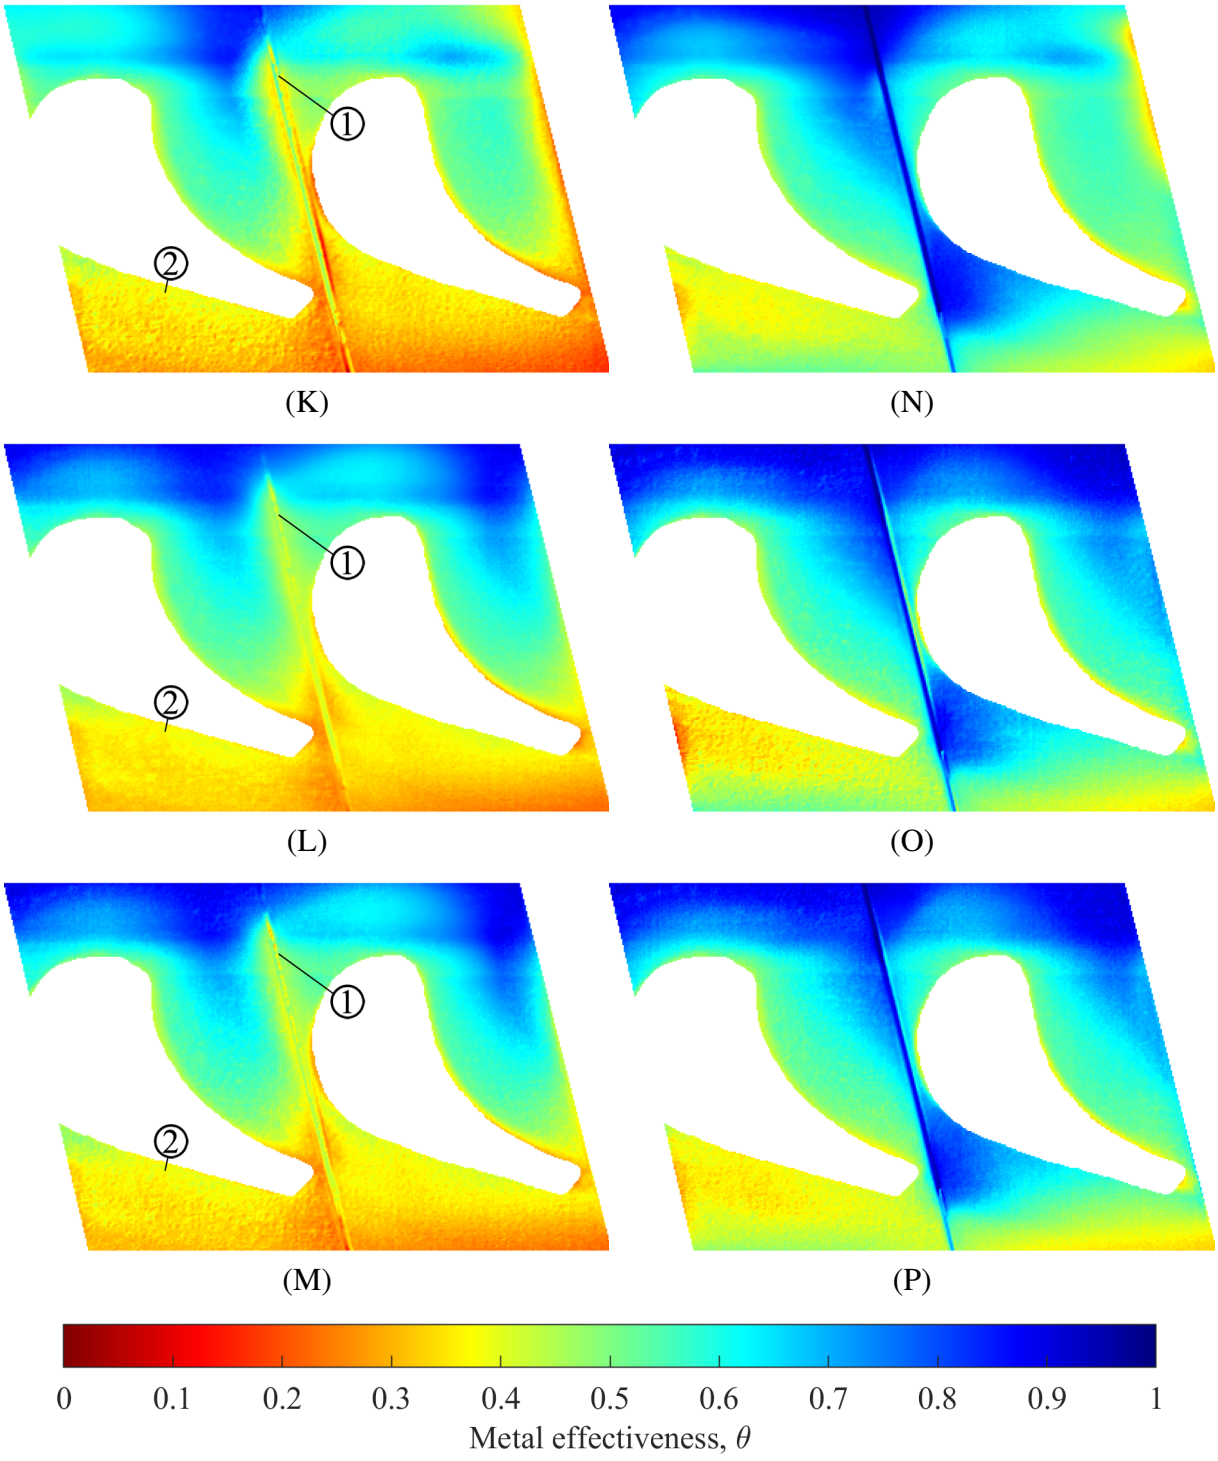


Fig. 8.4 Metal effectiveness distributions of two platforms for convective cooling passage blades. All displayed cases have front seal leakage. Cases (K), (L), and (M) have silicon-sealed shanks. Cases (N), (O), and (P) have shank slot leakage.

metal effectiveness distribution and only slightly increases values along the blade suction surface, marked (2). The rear seal overhang in these cases has low effectiveness values of approximately 0.20 as it is only partially cooled by coolant exhausting from the convective passages.

The addition of slashface leakage in cases N–P increases metal effectiveness values in this trailing edge region to greater than 0.50. The leakage between slashfaces in these cases is, similarly to the baseline cases, seen to be different between left and right platforms due to tolerance differences between parts. The leakage can be seen on the right platform to enhance cooling effectiveness values downstream of the blade throat. The band of low effectiveness along the slashface edges, seen in cases K–M, does not occur when slashface leakage is present as coolant flows along these surfaces towards the rear portion of the platform where the drop in static pressure allows it to leak onto the surface.

In cases K–M no coolant was supplied to the rear seal due to a previously discovered non-uniform flow within the rear seal cavity. This means that the pressure ratio between front and rear seal will be slightly higher than at realistic operating conditions. It does however mean that we expect each convective passage to have approximately the same exit static pressure. In the previous discussion of the impingement design blades there were notable differences in the values at corresponding locations between left and right platform distributions. This is not true for the convective passage distributions which, beyond the initial leading region at the front seal, are very similar in magnitude along the blade passage. In cases N–P, which had coolant supplied to all 3 seal plenums, a band of higher effectiveness (greater than 0.4) can be seen along the trailing edge of the leftmost platform, similar to the impingement cases and also likely caused by non-uniform flow within the rear seal. In an evenly sealed rear seal system it would therefore be expected for the metal effectiveness along the trailing edge to be closer to the values of the leftmost blade in cases N–P.

8.5 Comparison of all cooling configurations

The key features of the metal effectiveness distributions for the three pairs of tests blades have now been discussed in some detail. In this section the three blade designs are compared to examine

how well each of the two integrated cooling systems perform compared to each other, and what additional cooling benefit they provide over the baseline case. The key performance criteria which are used are the area-averaged metal effectiveness, the minimum metal effectiveness, and the total coolant mass ratio. The area-averaged metal effectiveness is taken across the entire two platforms and gives an indication of the overall cooling performance in each case. The minimum metal effectiveness value corresponds to the highest metal temperature on the platform. The total coolant mass ratio gives an indication of the cycle efficiency penalty of each case.

8.5.1 Comparison of area-averaged metal effectiveness

The area-averaged metal effectiveness of each test case A–P has been computed as the mean value of every pixel in the masked images shown in Figures 8.1, 8.3, and 8.4. A band, approximately 4 pixels wide, along the slashface gap has been excluded from the average. The resulting area-averaged metal effectiveness is plotted in Figure 8.5 as a function of total coolant mass ratio. The total massflow is defined as coolant which contributes to cooling and does not include the rear seal as the majority of coolant in the rear seal does not contribute to cooling the platform and is non-uniformly distributed within the rear seal system. Symbols denote the blade design used in each case and colour denotes the type of leakage flows present. Trend lines of the performance for each of the designs have also been plotted.

The baseline blade, cooled only by leakage flows, shows an approximately linear increase in area-averaged effectiveness as coolant mass flow is increased. A threefold increase in total coolant mass ratio causes the average θ to increase from 0.36 to 0.53, an increase of 47.2%. The additional aerodynamic penalty of the leakage flows mixing with the mainstream is expected to be high compared to the other systems, at similar total coolant mass ratios, which keep more of the used coolant internal. For cases C and D, where slashface leakage is a significant proportion of the coolant, the aerodynamic losses due to mixing would be expected to be significantly higher due to the shank leakage exiting into a higher Mach number flow.

The impingement system performs fairly consistently across the $\pm$20% variation in total coolant mass ratio tested. The maximum difference in average θ for cases E–G was 0.033, or

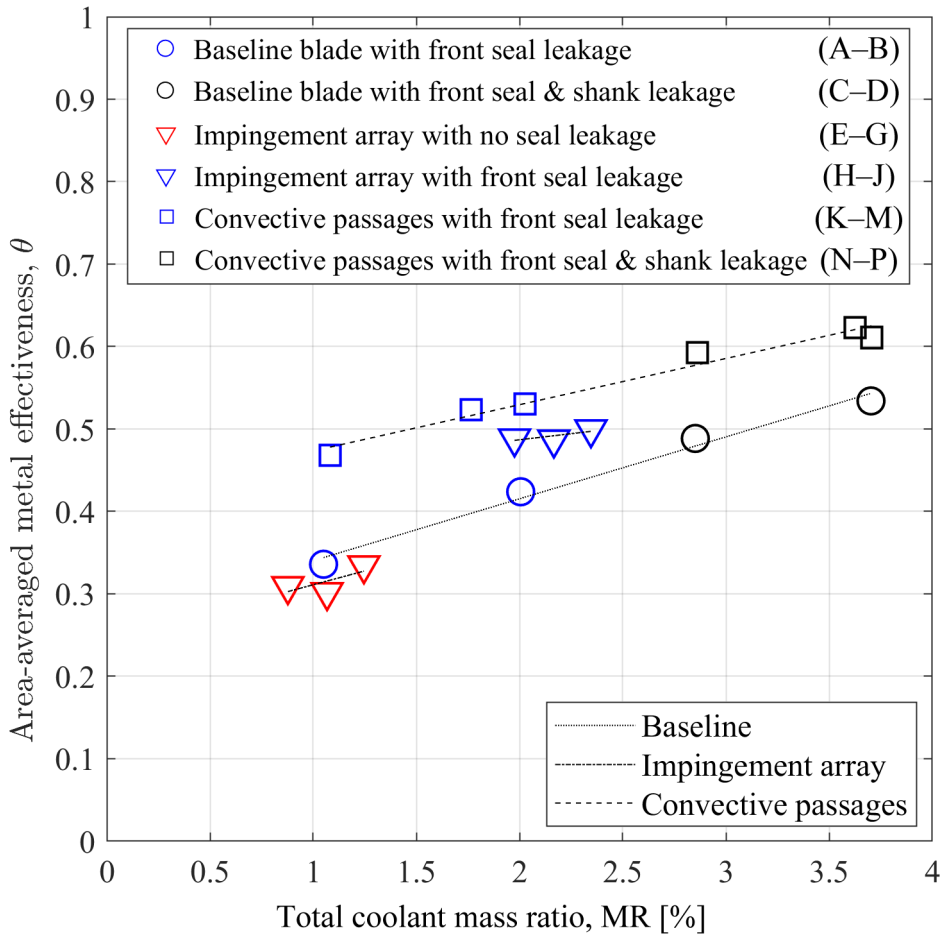


Fig. 8.5 Area-averaged metal effectiveness across two rotor platforms plotted as a function of total coolant mass ratio for 16 test case conditions (A–P).

11%. For the case where total coolant mass ratio is approximately unity, the impingement system performed slightly lower, $\Delta\theta = 0.039$, than the baseline case. This is due to no leakage coolant being present leading to low values of effectiveness along the leading portion of the platform. The aerodynamic penalty incurred by mixing of coolant and mainstream would however be expected to be significantly lower at this condition since all of the impingement flow is exhausted into the rear seal. For case H, with MR $= 2\%$ the impingement system performed better than the baseline case, as one would expect for a design with additional cooling features, and had an average effectiveness higher than the baseline by 0.065.

The impingement cases which were cooled only by the internal impingement system and which had no external leakage flows (E–G) are the most similar to the situation modelled by the low-order model solver in Chapter 4. Although no tuning of the model was conducted based

on these limited test cases, it is worth noting that the area-averaged metal effectiveness is very similar to values predicted by the model, $\theta \approx 0.3$. This lends confidence to the models ability to predict behaviour and can potentially be expanded upon to design more optimal iterations of the design. The convective cooling model is not compared to the results here as it did not model external cooling films and it was impossible to allow for flow through the convective cooling passages whilst not having front seal purge flow in the current cascade configuration.

The convective passage design outperformed the impingement array design at all test conditions. An approximately linear increase in average effectiveness as mass flow is increased is observed. Average effectiveness values were greater than the baseline case by 0.13 and 0.077 for total coolant MR values of 1.1% and 3.7% respectively. Although the convective passage design outperformed the impingement system, the aerodynamic losses would be expected to be greater due to the leakage flows being introduced to the mainstream passage, especially the shank leakage into the high Mach number region.

These results show a clear advantage in cooling performance, for similar leakage flow conditions, for both of the integrated cooling systems over the baseline case. Of the two integrated blade designs, the convective cooling passage design performed the best, due to its ability to directly cool the entire platform surface (whereas the area targetting by the impingement system is limited) and enhancement by additional external film cooling. It would be interesting to conduct experiments on the convective passage design with no external cooling films, so that the cooling effect of just the convective passages could be assessed. This could be achieved by having a new front seal block manufactured that fully seals the front seal slot to allow the cavity to be pressurised without having a front seal purge flow. This would allow the convective passage and impingement designs to be more fairly compared at a total coolant mass ratio of unity. Although only a small number of test cases were completed, the demonstrable cooling performance of each system warrants further investigation and optimisation of both designs.

8.5.2 Comparison of minimum metal effectiveness

The minimum value of metal effectiveness on the platform corresponds with the highest metal temperature on the platform which must be below acceptable operating limits for the metal and desired component life. This value is useful when comparing these different blade designs as it gives an indication of problematic hot spots; a system which has an higher area-averaged metal effectiveness may be less useful than a competing design if that is being achieved by overcooling some regions and inadequately cooling others. A good cooling system is therefore indicated by having a high minimum metal effectiveness value.

The minimum 'stable' value of θ was determined from area-averaging groups of 3×3 pixels and finding the lowest value on the platform. This prevents erroneous single pixel values from interfering with comparisons. The minimum metal effectiveness value is plotted as a function of total coolant mass ratio in Figure 8.6 for each of the 16 test cases. Trend lines of the performance for each of the designs have also been plotted.

The baseline blade, cooled only by leakage flows, has the lowest minimum metal effectiveness values of the three designs at all comparable test conditions. The lowest of these was case A with a value of 0.094. In an engine setting, values below 0.2 would typically be a cause for concern. The two cooling systems both have higher minimum metal effectiveness values at this total mass ratio of unity, with 0.157 for the impingement design and 0.164 for the convective passage design. Both values are therefore approximately 50% higher than the baseline case at the same total coolant mass ratio. When the total coolant mass ratio is increased to 2% both of the cooling designs have minimum metal effectiveness values greater than 0.2. The impingement system here performs slightly better than the convective passage design, having values of approximately 0.24 compared to 0.22. This is because the convective cooling passage design has proportionally more of the coolant exhausted as the front seal purge as less passes through the internal system. As previously discussed, the front seal purge cooling performance is impeded by the passage secondary flows. This therefore shows a clear advantage in keeping cooling flows internal as more evenly distributed cooling can be achieved.

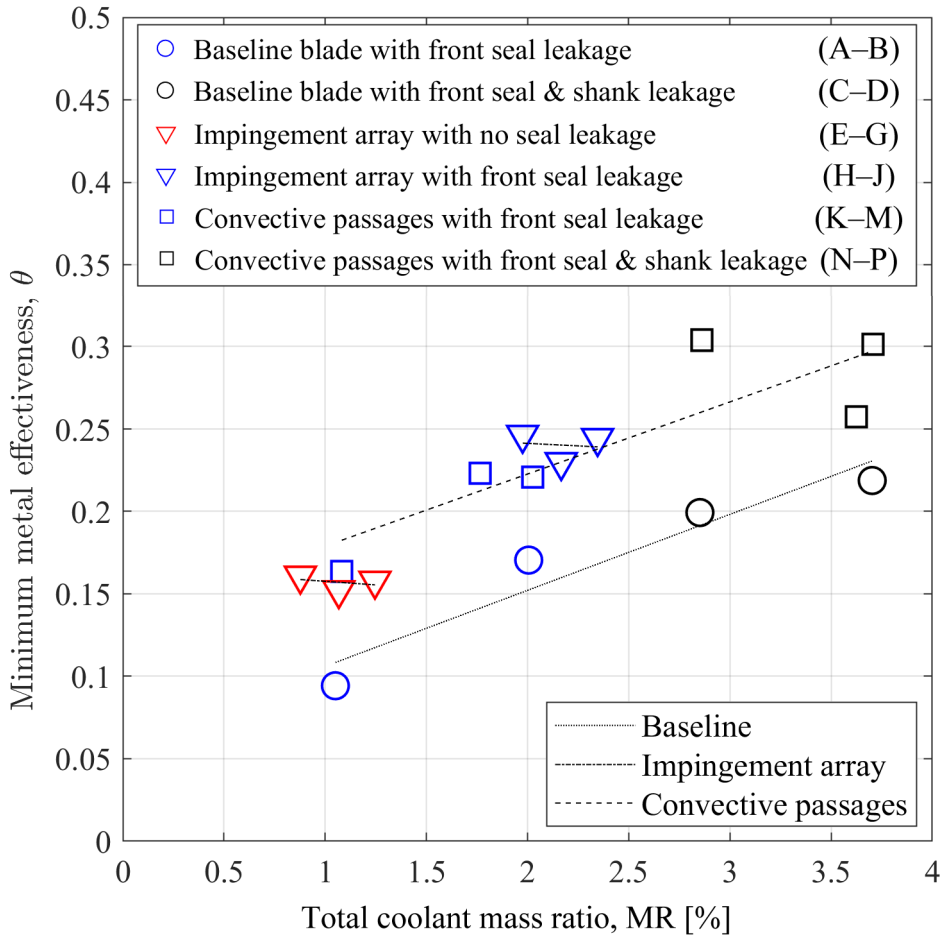


Fig. 8.6 Minimum stable value (defined as a 3x3 pixel area) of metal effectiveness across two rotor platforms plotted as a function of total coolant mass ratio for 16 test case conditions (A–P).

8.6 Conclusions

A summary of the key findings from the work described in this chapter are given below:

- Three pairs of titanium test blade have been tested in the MERC facility and metal effectiveness distributions obtained across two platforms at a variety of flow conditions. The 3 blade designs tested were: a baseline solid platform cooled by hub seal leakage; an impingement array integrated into the platform and supplied with coolant from the shank pocket; and a number of convective passages running from the front seal to the rear seal. Test conditions which were varied: the presence and amount of leakage flow from the front seal and shank pocket; the amount of coolant passing through integrated systems. A total of 16 test cases, spread across the three designs, at various conditions were examined.

- Metal effectiveness results from the baseline blade cases showed that leakage flows partially cool the platform surface. The extent of this cooling is limited by the strong blade passage secondary flows which sweep leaked coolant towards the suction surface of the blade where it is entrained by the passage vortex. A 300% increase in mass ratio resulted in an increase in average effectiveness from 0.36 to 0.53 and increased the minimum effectiveness value from 0.094 to 0.22. In addition to the direct cooling effect of the seal leakage flow, a region of higher-than-expected metal effectiveness was present near to the pressure side of the blade, likely caused by the fir tree root acting as a heat sink, wicking heat away from the platform and being internally cooled by seal coolant.

- The impingement array design effectively cools the platform along the integrated ducts and the surrounding areas. The design performed better than the baseline in the presence of similar leakage flows, increasing the average metal effectiveness by 0.065 and minimum effectiveness by 0.076. The platform region with lowest effectiveness was found to lie along the trailing edge, in between the two impingement ducts and up to the suction surface of the blade. Front seal leakage flows enhance the cooling performance of the impingement system, increasing effectiveness values and reducing gradients, with a front seal purge MR = 1.1% increasing the average effectiveness by 0.185.

- The convective passage design outperformed the impingement system at all conditions with average effectiveness values between 0.47 and 0.62. As the convective passages are spread over a larger proportion of the platform surface the effectiveness distribution is more even than in the impingement design. Cooling effectiveness reduces gradually from the leading to trailing edge as heat is transferred into the coolant. The lowest effectiveness values for the convective passage design (between 0.16 and 0.30) are found along the trailing edge overhang into the rear seal, where only exhausted coolant from the passages is able to partially cool the metal. The convective passage design is expected to incur a higher aerodynamic penalty for similar values of mass ratio to the impingement system due to a larger proportion of the cooling coming from leakage flows which mix with the mainstream.

- Both integrated cooling designs exhaust coolant into the rear seal where it could potentially be reused to help seal the downstream rotor-stator interface. It may be possible to reduce the coolant supplied to the rear seal to account for this and partially offset the aerodynamic cost of the coolant used to cool the rotor platforms. Systems such as these which effectively reuse coolant for the dual purpose of cooling the rotor platform and sealing the hub seals have the potential to adequately cool rotor platforms whilst minimising overall coolant usage and mixing losses.

Chapter 9

Summary and recommendations for future work

In this chapter, the key findings and conclusions from the work described throughout this thesis are summarised and recommendations for future work outlined. More detailed conclusions can be found at the end of relevant chapters. Recommendations for improvements to the MERC facility can be found at the end of Chapter 5.

9.1 Summary of facility development and cooling system design work

An experimental facility was designed, commissioned, and utilised in a number of experimental campaigns to investigate novel platform cooling systems. The facility is a highly modular, high pressure, high temperature, transonic, linear cascade which operates at scaled engine conditions in the rotor relative frame. Similarity to engine conditions is achieved by matching Mach number, Reynolds number, inlet turbulence intensity and Biot number (by using titanium alloy test components). The facility is highly versatile and allows for a wide variety of platform cooling systems to be tested. Metal effectiveness and adiabatic film effectiveness distributions across two blade platforms are determined from measurements of surface temperature via IR thermography. Aerodynamic loss studies are conducted by using a downstream area traverse system and measurement probes.

Two novel platform cooling systems were developed: a convective passage system where coolant flows from the front seal to the rear seal through passages within the platform; and an impingement array system which uses impinging jets fed from the shank cavity. An advantage

of both these designs is that cooling is achieved by coolant flowing within the platforms, rather than being ejected into the mainstream gas path and incurring additional mixing losses. Coolant is exhausted into the rear hub seal and can be re-used as sealing flow such that it may be possible to reduce the required coolant supplied to the rear seal. The main advantage of this would be to increase cycle efficiency by reducing the total coolant mass flow rate used to cool the turbine.

These cooling systems were developed using low-order numerical models. These models solve the aerodynamics of each system and estimate cooling performance by additionally modelling heat transfer in discretised flow domains. They were used to estimate the performance of a large number of systems and hence to select optimal design parameter values (passage diameter, number of jets *etc.*) Each cooling system was subsequently integrated into cascade test blade geometry, manufactured via DMLS, and tested in the new MERC facility.

9.2 Summary of adiabatic film effectiveness experiments

A new processing method was developed to allow for accurate adiabatic film effectiveness measurements to be obtained from a single heated experiment when a non-uniform, time-varying inlet temperature profile (such as that caused by a thermal boundary layer developing over time) is present. This inlet temperature profile results in the temperature of streamlines interacting with a surface, and thus driving heat transfer into it, being different to the mainstream core-flow temperature typically measured in experiments and used when calculating effectiveness. The result of this was shown to be that adiabatic film effectiveness values calculated in the conventional manner are higher than the true values. Images from two, or more, timestamps during the experiment are used to simultaneously determine the true adiabatic film effectiveness and a newly proposed equivalent mainstream effectiveness. Additional corrections were detailed which account for: compressible recovery temperature, by using a known Mach number distribution to calculate compressible recovery factor; and non-ideal adiabatic surface behaviour, by using virtual experiments and a transient conduction model to estimate surface behaviour.

A parametric study of 20 test conditions was conducted to assess the impact of varying front seal purge flow angle and mass ratio on the adiabatic film cooling performance. Five purge

flow angles (20° towards the pressure surface, 0°, −20° towards the suction surface,−40°,−60°) and four mass ratios (0.5%, 1.0%, 1.5%, and 2.0%) were tested. Adiabatic film effectiveness distributions along one blade passage were presented for each case, as well as total pressure loss distributions on a plane immediately downstream of the trailing edge. Front seal purge flow was found to effectively cool a part of the platform towards the leading edge, though the area was limited due to the strong interaction with passage secondary flows. All of the purge flow coolant is ultimately ingested by the passage vortex. Purge flow angles towards the pressure surface had enhanced film cooling performance, with higher adiabatic film effectiveness values and greater cooling coverage. Increasing the purge mass ratio led to an increase in cooling performance. Somewhat surprisingly, varying both the purge flow angle and mass ratio did not have a significant impact on the total aerodynamic loss through the blade passage.

A system was proposed to bias the purge flow towards the pressure surface in an engine environment and thus take advantage of the increased cooling performance. This system imposes a strong flow angle bias towards the pressure surface by supplying part, or all, of the purge flow coolant from the NGV platform cooling system. A secondary benefit of this system is that it reuses coolant from the NGV cooling system and may thus allow for the total coolant used to cool the turbine to be reduced, increasing cycle efficiency.

9.3 Summary of metal effectiveness experiments

Three pairs of titanium test blades were tested in the new MERC facility and detailed metal effectiveness distributions were obtained across both platforms. The blade pairs tested were: a solid baseline design with no integrated cooling features which was cooled solely by leakage flows; a convective passage design with convective cooling passages integrated directly within the platform, passing coolant from front to rear seal; and an impingement jet array design feeding jets from the shank cavity. Experiments were conducted for a total of 16 test configurations. Test configurations differed by the presence of leakage flows and the mass ratios of coolant to these leakage flows and cooling systems. Both integrated cooling systems were shown to have increased cooling performance, at similar total coolant mass ratios, compared to the baseline

blades. The convective passage system outperformed the impingement system, in terms of cooling performance, at all total coolant mass ratios tested.

9.4 Recommendations for future work

Some suggestions for future work are now given. The work of this thesis forms a strong foundation and framework for future studies, providing many opportunities for more detailed experimental campaigns and for optimisation of each cooling system design.

- The two integrated platform cooling systems both showed clear cooling benefits in the metal effectiveness tests. The designs each have a large scope for optimisation and further development into more complex systems. The convective cooling passages could be developed to have less evenly distributed coverage in order to target hotter regions of the platform more directly, by having closer spaced passages in those areas, with the aim to minimise platform thermal gradients. The major limitation of the impingement system is the area which can be targeted from the shank cavity. A double-skin system could be developed so that jets could be used to cool any part of the platform.

- The numerical models used to design cooling systems could be tuned using available experimental data, and with more data from additional experiments, and used to further optimise the designs. Extensions to the numerical models could be made to account for lateral conduction of heat, to give more accurate prediction of cooling performance, and to additionally include more cooling features such as film cooling.

- The system proposed to control the angle of front seal purge flow could be tested in the MERC facility by having custom blade cassette inserts and front seal blocks manufactured. Additionally, many other potential platform cooling systems have been proposed in open literature. The MERC facility is versatile enough that it would be feasible to test a wide variety of these such as: platform contouring; discrete film cooling holes; reverse-pass convective cooling. The facility could also be fairly easily adapted for tip-cooling experiments.

- There remains a lot of potential for further development of the IR thermography methods used to obtain data. Spatial variation in surrounding radiance and window transmissivity are typically regarded as negligible. Experiments to verify this, and extensions to processing to account for variations if not, would be beneficial. A comprehensive set of experiments and uncertainty analysis on K-type thermocouples to definitively characterise their typical bias and precision uncertainties, when using the data acquisition system and typical lab set up, would be highly useful.

- The extension to the adiabatic film effectiveness processing technique, to account for non-ideal adiabatic surface behaviour, still needs work to be reliably applied to complete images. If this could be fully implemented it would be highly beneficial both to further adiabatic film effectiveness work in the MERC facility and in other similar facilities worldwide. Implementing transient conduction models into the IR processing could also lead to metal effectiveness distributions being obtained with fewer required runs at different temperature ratios. It may also allow for heat transfer coefficient distributions to be determined simultaneously with effectiveness distributions.

- Detailed CFD studies of the current MERC facility test section and blade cassette would be useful to understand the non-uniformities in inlet pressure and temperature profiles. CFD is ideal to study and understand how differences between the cascade and engine environments affect the behaviour of leakage flows and the impact this has on cooling. In particular, it would be beneficial to examine the differences in the inlet hub boundary layer and the impact of the addition of unsteady NGV wakes. These could be used to develop upgrades to the facility and in the design of updated blade cassette inserts. Similarly, CFD studies could also be conducted to assess the impact of rotation on the investigated internal cooling systems and to characterise differences between the MERC facility experiments and rotating environments.

- Parallel CFD studies, especially conjugate heat transfer studies, to support the development and optimisation of the investigated cooling designs would also be highly useful for developing more complex cooling systems. In particular, a CFD study of the front seal

jet mixing concept, to control purge flow angle, may be more appropriate and yield better understanding of the system than experiments in the MERC facility could provide.

References

[1] Baldauf, S., Scheurlen, M., Schulz, A., and Wittig, S. (2002). Correlation of film-cooling effectiveness from thermographic measurements at enginelike conditions. *Journal of Turbomachinery*, 124(4).

[2] Baldauf, S., Schulz, A., and Wittig, S. (1999). High-resolution measurements of local effectiveness from discrete hole film cooling. *Journal of Turbomachinery*, 123(4):758–765.

[3] Barigozzi, G., Franchini, G., Perdichizzi, A., Maritano, M., and Abram, R. (2013). Purge flow and interface gap geometry influence on the aero-thermal performance of a rotor blade cascade. *International Journal of Heat and Fluid Flow*, 44:563–575.

[4] Bergman, T. L., Lavine, A. S., Incropera, F. P., and DeWitt, D. P. (2002). *Fundamentals of Heat and Mass Transfer*. Wiley, 5th edition.

[5] Blair, M. F. (1974). An experimental study of heat transfer and film cooling on large-scale turbine endwalls. *Journal of Heat Transfer*, 96(4):524–529.

[6] Bogard, D. G. (2006). Airfoil film cooling. In *The Gas Turbine Handbook*, pages 309–320. Department of Energy, Office of Fossil Energy, National Energy Technology Laboratory.

[7] British Standards Institution (1981). *BS 1042: Measurement of fluid flow in closed conduits*.

[8] Bru Revert, A., Beard, P., Chew, J., and Bottenheim, S. (2020). Performance of a turbine rim seal subject to rotationally-driven and pressure-driven ingestion. In *ASME Turbo Expo 2020: Turbomachinery Technical Conference and Exposition*.

[9] Bunker, R. S. and Bailey, J. C. (2005). Turbine blade platform impingement cooling with effects of film extraction and roughness. In *ASME Turbo Expo 2005: Turbomachinery Technical Conference and Exposition*, volume 3, pages 419–429.

[10] Cattafesta, L., Bahr, C., and Mathew, J. (2010). *Fundamentals of Wind-Tunnel Design*.

[11] Chen, A. F., Shiau, C.-C., and Han, J.-C. (2016). Turbine blade platform film cooling with simulated swirl purge flow and slashface leakage conditions. *Journal of Turbomachinery*, 139(3).

[12] Chew, J. W., Gao, F., and Palermo, D. M. (2018). Flow mechanisms in axial turbine rim sealing. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 233(23-24):7637–7657.

[13] Crafton, J., Lachendro, N., Guille, M., Sullivan, J., and Jordan, J. (1999). Application of temperature and pressure sensitive paint to an obliquely impinging jet. In *AIAA 37th Aerospace Sciences Meeting and Exhibit*.

[14] Crystran Ltd. (2014). *The Design of Pressure Windows*.

[15] Cumpsty, N. (2003). *Jet Propulsion: A Simple Guide to the Aerodynamic and Thermodynamic Design and Performance of Jet Engines*. Cambridge University Press, 2nd edition.

[16] Dellimore, K., Marshall, A., and Cadou, C. (2010). Influence of compressibility on film-cooling performance. *Journal of Thermophysics and Heat Transfer*, 24(3):506–515.

[17] Dixon, S. and Hall, C. (2014). *Fluid Mechanics and Thermodynamics of Turbomachinery*. Butterworth-Heinemann, Boston, 7th edition.

[18] Ekkad, S. V., Ou, S., and Rivir, R. B. (2004). A transient infrared thermography method for simultaneous film cooling effectiveness and heat transfer coefficient measurements from a single test. *Journal of Turbomachinery*, 126(4):597–603.

[19] Ellis, S. E. and Hynum, D. A. (2012). Apparatus and methods for cooling platform regions of turbine rotor blades. US Patent 2012/0082566.

[20] Erickson, R. D., Simon, T. W., Zhang, L., and Moon, H. (2011). Experimental investigation of disc cavity leakage flow and hub endwall contouring in a linear rotor cascade. In *ASME Turbo Expo 2011: Turbine Technical Conference and Exposition*, volume 5: Heat Transfer, Parts A and B, pages 1769–1780.

[21] Facchini, B., Tarchi, L., Toni, L., and Ceccherini, A. (2010). Adiabatic and overall effectiveness measurements of an effusion cooling array for turbine endwall application. *Journal of Turbomachinery*, 132(4).

[22] FLIR (2012). *The Ultimate Infrared Handbook for R&D Professionals*.

[23] Florschuetz, L. W., Berry, R. A., and Metzger, D. E. (1980). Periodic streamwise variations of heat transfer coefficients for inline and staggered arrays of circular jets with crossflow of spent air. *Journal of Heat Transfer*, 102(1):132–137.

[24] Florschuetz, L. W., Truman, C. R., and Metzger, D. E. (1981). Streamwise flow and heat transfer distributions for jet array impingement with crossflow. *Journal of Heat Transfer*, 103:337–342.

[25] Friedrichs, S. (1997). *Endwall Film-Cooling in Axial Flow Turbines*. PhD thesis, University of Cambridge, Whittle Laboratory.

[26] Friedrichs, S., Hodson, H. P., and Dawes, W. N. (1997). Aerodynamic aspects of endwall film-cooling. *Journal of Turbomachinery*, 119(4):786–793.

[27] Furukawa, T. and Ligrani, P. (2002). Transonic film cooling effectiveness from shaped holes on a simulated turbine airfoil. *Journal of Thermophysics and Heat Transfer*, 16(2):228–237.

[28] Gao, Z., Narzary, D., and Han, J. C. (2009). Turbine blade platform film cooling with typical stator-rotor purge flow and discrete-hole film cooling. *Journal of Turbomachinery*, 131(4).

[29] Gao, Z., Narzary, D., Mhetras, S., and Han, J. C. (2012). Upstream vortex effect on turbine platform film cooling with typical purge flow. *Journal of Thermophysics and Heat Transfer*, 26:75–84.

[30] Gilboy, M. T. and Gan, X. P. (2003). Validation of a new correlation parameter against film cooling data from a detailed literature survey. QinetiQ ltd., Cody Technology Park, Farnborough, Hampshire GU14 0LX.

[31] Gillespie, D. (1996). *Intricate internal cooling systems for gas turbine blading*. PhD thesis, University of Oxford, Department of Engineering Science.

[32] Girgis, S. and Marini, R. (2007). Deflectors for controlling entry of fluid leakage into the working fluid flowpath of a gas turbine engine. US Patent US7244104B2.

[33] Gladden, H. J. and Simoneau, R. (1988). Review and assessment of the database and numerical modeling for turbine heat transfer.

[34] Goldstein, R. J. (1971). Film cooling. volume 7 of *Advances in Heat Transfer*, pages 321 – 379. Elsevier.

[35] Guimbard, J. M., Kueny, O., and Pintat, L. (2015). Blade with 3d platform comprising an inter-blade bulb. US Patent US9518467B2.

[36] Han, J. C., Dutta, S., and Ekkad, S. (2013). *Gas Turbine Heat Transfer and Cooling Technology*. CRC Press, 2 edition.

[37] Han, J. C. and Rallabandi, A. (2010). Turbine blade film cooling using psp technique. *Frontiers in Heat and Mass Transfer*, 1.

[38] Han, S. and Goldstein, R. J. (2005). Influence of blade leading edge geometry on turbine endwall heat (mass) transfer. *Journal of Turbomachinery*, 128(4):798–813.

[39] Hirsch, C. (1993). *Advanced Methods for Cascade Testing*. AGARD/AG-328. Advisory Group for Aerospace Research & Development.

[40] Holgate, N. E., Ireland, P. T., and Romero, E. (2018). The effects of combustor cooling features on nozzle guide vane film cooling experiments. *Journal of Turbomachinery*, 141(1).

[41] Holgate, N. E., Ireland, P. T., and Romero, E. (2019). An experimental-numerical method for transient infrared measurement of film cooling effectiveness and heat transfer coefficient in a single test. *The Aeronautical Journal*, 123(1270):1982–1998.

[42] Howatson, A. M., Lund, P. G., and Todd, J. D. (2009). *Engineering tables and data*.

[43] Hultgren, K. G., Zagar, T. W., North, W. E., and S., H. (2000). Gas turbine blade platform cooling concept. US Patent 6120249.

[44] Ireland, P., Neely, A., Gillespie, D., and Robertson, A. (1999). Turbulent heat transfer measurements using liquid crystals. *International Journal of Heat and Fluid Flow*, 20(4):355–367.

[45] Kays, W., Crawford, M., and Weigand, B. (2004). *Convective heat and mass transfer*. McGraw-Hill Higher Education.

[46] Kirollos, B. and Povey, T. (2017). Laboratory infrared thermal assessment of laser-sintered high-pressure nozzle guide vanes to derisk engine design programs. *Journal of Turbomachinery*, 139(4).

[47] Kirollos, B. W. (2015). *Aerothermal Optimisation of Novel Cooling Schemes for High Pressure Components using Combined Theoretical, Numerical and Experimental Techniques*. PhD thesis, University of Oxford, Department of Engineering Science.

[48] Kurzke, J. and Halliwell, I. (2018). *Propulsion and Power*. Springer, Cham.

[49] Lakshminarayana, B. (1996). *Fluid Dynamics and Heat Transfer of Turbomachinery*. John Wiley & Sons, Inc.

[50] Lampart, P. (2006). Tip leakage flows in turbines. *TASK Quarterly*, 10:139–175.

[51] Langston, L. (1980). Crossflows in a turbine cascade passage. *Journal of Engineering for Power*, 102:866.

[52] Langston, L. (2001). Secondary flows in axial turbines—a review. *Annals of the New York Academy of Sciences*, 934(1):11–26.

[53] Logan, E. (2003). *Handbook of Turbomachinery*. CRC Press, 2 edition.

[54] Luque, S., Jones, T., and Povey, T. (2016). Theory for the scaling of metal temperatures in cooled compressible flows. *International Journal of Heat and Mass Transfer*, 102:331–340.

[55] Luque, S., Jones, T., and Povey, T. (2017). Scaling of turbine metal temperatures in cooled compressible flows—experimental demonstration of a new theory. *Journal of Turbomachinery*, 139(8).

[56] Luque Martinez, S. (2011). *A Fully Integrated Approach to Gas Turbine Cooling System Research*. PhD thesis, University of Oxford, Department of Engineering Science.

[57] Lynch, S. P., Sundaram, N., Thole, K. A., Kohli, A., and Lehane, C. (2011). Heat transfer for a turbine blade with nonaxisymmetric endwall contouring. *Journal of Turbomachinery*, 133(1).

[58] Mahmood, G. I., Gustafson, R., and Acharya, S. (2005). Experimental investigation of flow structure and nusselt number in a low-speed linear blade passage with and without leading-edge fillets. *Journal of Heat Transfer*, 127(5):499–512.

[59] Michaud, M., Ornano, F., Chowdhury, N. H. K., and Povey, T. (2020). Methodology for high-accuracy infrared calibration in environments with through-wall heat flux. *Journal of the Global Power and Propulsion Society*, 4:1–13.

[60] Morel, T. (1975). Comprehensive Design of Axisymmetric Wind Tunnel Contractions. *Journal of Fluids Engineering*, 97(2):225–233.

[61] Nakos, J. T. (2004). Uncertainty analysis of thermocouple measurements used in normal and abnormal thermal environment experiments at sandia's radiant heat facility and lurance canyon burn site. *Sandia National Laboratories Report*.

[62] Narzary, D. P. (2009). *Experimental Study Of Gas Turbine Blade Film Cooling And Heat Transfer*. PhD thesis, Texas A&M University.

[63] Ornano, F. and Povey, T. (2020). Theory of non-dimensional groups in film effectiveness studies. *Journal of Turbomachinery*, 142(4).

[64] Parker, J. A., Romero, E., and Povey, T. (2019). A modular transonic turbine cascade for cooled rotor metal effectiveness investigations. In *ASME Turbo Expo 2019: Turbomachinery Technical Conference and Exposition*, volume 5A: Heat Transfer.

[65] Pau, M., de la Loma, A., Paniagua, G., and Delhaye, D. (2007). Film cooling and hub disk leakage flow experiments in a fully rotating hp turbine stage. In *Proceedings of the 5th IASME/WSEAS International Conference on Fluid Mechanics and Aerodynamics*.

[66] Pau, M., Paniagua, G., Delhaye, D., de la Loma, A., and Ginibre, P. (2010). Aerothermal impact of stator-rim purge flow and rotor-platform film cooling on a transonic turbine stage. *Journal of Turbomachinery*, 132(2).

[67] Piggush, J. and Simon, T. (2007a). Measurements of net change in heat flux as a result of leakage and steps on the contoured endwall of a gas turbine first stage nozzle. *Applied Thermal Engineering*, 27(4):722–730. Energy: Production, Distribution and Conservation.

[68] Piggush, J. and Simon, T. W. (2007b). Heat transfer measurements in a first-stage nozzle cascade having endwall contouring: Misalignment and leakage studies. *Journal of Turbomachinery*, 129:782–790.

[69] Popović, I. and Hodson, H. P. (2013). Aerothermal impact of the interaction between hub leakage and mainstream flows in highly-loaded high pressure turbine blades. *Journal of Turbomachinery*, 135(6).

[70] Povey, T., Jones, T. V., and Oldfield, M. L. G. (2004). On a novel annular sector cascade technique. *Journal of Turbomachinery*, 129(1):175–183.

[71] Povey, T., Oldfield, M. L. G., and Haselbach, F. (2008). Transonic turbine vane tests in a new miniature cascade facility. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 222(5):529–539.

[72] Rahim, A. and He, L. (2015). Rotor Blade Heat Transfer of High Pressure Turbine Stage Under Inlet Hot-Streak and Swirl. *Journal of Engineering for Gas Turbines and Power*, 137(6). 062601.

[73] Ramesh, S., Ekkad, S., Straub, D., Lawson, S., and Alvin, M. A. (2014). Experimental and computational analysis of film cooling hole performance on a high temperature test rig. In *ASME International Mechanical Engineering Congress and Exposition*, volume 8B: Heat Transfer and Thermal Engineering.

[74] Roach, P. (1987). The generation of nearly isotropic turbulence by means of grids. *International Journal of Heat and Fluid Flow*, 8(2):82–92.

[75] Rolls-Royce plc. (1996). *The jet engine*. 5th edition.

[76] Scobie, J. A., Sangan, C. M., Michael Owen, J., and Lock, G. D. (2016). Review of ingress in gas turbines. *Journal of Engineering for Gas Turbines and Power*, 138(12).

[77] Seitzman, J. M. (2001). Lecture notes for course ae 3450: Thermodynamics and compressible flow. Georgia Tech College of Engineering.

[78] Sharma, O. P. and Butler, T. L. (1987). Predictions of Endwall Losses and Secondary Flows in Axial Flow Turbine Cascades. *Journal of Turbomachinery*, 109(2):229–236.

[79] Sieverding, C. H., Van Hove, W., and Boletis, E. (1984). Experimental study of the three-dimensional flow field in an annular turbine nozzle guidevane. *Journal of Engineering for Gas Turbines and Power*, 106(2):437 – 444.

[80] Singh, R., Ameyugo, G., and Noppel, F. (2012). 4 - jet engine design drivers: past, present and future. In Young, T. M. and Hirst, M., editors, *Innovation in Aeronautics*, pages 56–82. Woodhead Publishing.

[81] Suryanarayanan, A. (2009). *Film Cooling, Heat Transfer And Aerodynamic Measurements In A Three Stage Research Gas Turbine*. PhD thesis, Texas A&M University.

[82] Thole, K. (2006). Airfoil endwall heat transfer. In *The Gas Turbine Handbook*, pages 353–362. Department of Energy, Office of Fossil Energy, National Energy Technology Laboratory.

[83] Thrift, A. A., Thole, K., and Hada, S. (2011). Effects of an axisymmetric contoured endwall on a nozzle guide vane: Adiabatic effectiveness measurements. *Journal of Turbomachinery*, 133.

[84] Thrift, A. A. and Thole, K. A. (2012). Influence of flow injection angle on a leading-edge horseshoe vortex. *International Journal of Heat and Mass Transfer*, 55(17):4651 – 4664.

[85] Vedula, R. and Metzger, D. (1991). A method for the simultaneous determination of local effectiveness and heat transfer distributions in three-temperature convection situations. ASME Paper No. 91-GT-345.

[86] Watmuff, J. H. (1987). Wind tunnel contraction design. In *9th Australasian Fluid Mechanics Conference*, pages 472–475.

[87] Wright, L. M., Gao, Z., Varvel, T. A., and Han, J.-C. (2005). Assessment of steady state PSP, TSP, and IR measurement techniques for flat plate film cooling. In *ASME Heat Transfer Summer Conference*, pages 37–46.

[88] Wright, L. M., Malak, M. F., Crites, D. C., Morris, M. C., Yelavkar, V., and Bilwani, R. (2014). Review of platform cooling technology for high pressure turbine blades. In *ASME Turbo Expo 2014: Turbine Technical Conference and Exposition*, volume 5B: Heat Transfer.

[89] Xue, S., Roy, A., Ng, W., and Ekkad, S. (2015). A novel transient technique to determine recovery temperature, heat transfer coefficient, and film cooling effectiveness simultaneously in a transonic turbine cascade. *Journal of Thermal Science and Engineering Applications*, 7.

[90] Zuckerman, N. and Lior, N. (2006). Jet impingement heat transfer: Physics, correlations, and numerical modeling. volume 39 of *Advances in Heat Transfer*, pages 565 – 631. Elsevier.