

HYPERSONIC MODELLING AND TESTING OF SPACE DEBRIS DURING PLANETARY ENTRY

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Engineering Science

---

Nathan Donaldson
St. Catherine's College

Supervised by Prof. Peter Ireland

Submitted: Trinity 2019

ABSTRACT

HYPERSONIC MODELLING AND TESTING OF SPACE DEBRIS DURING PLANETARY ENTRY

Nathan Donaldson, St. Catherine's College, University of Oxford

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Engineering Science

In the study of atmospheric entry, some of the most challenging conditions to model numerically manifest in the slip and transition regimes. Denser than free molecular gases and sparser than continuum flows, these regimes are typically characterised by Knudsen numbers in the range $0.001 \leqslant Kn \leqslant 10$. The complex interaction between gas particles in these sparse flows make the determination of aerothermal phenomena especially difficult, as traditional Navier-Stokes based CFD solutions cease to be accurate, and the Newtonian-style flow which characterises the free molecular regime does not manifest until $Kn \gtrapprox 10$.

This thesis investigates this complex regime in the context of destructive satellite re-entry. Modern space debris demise codes often utilise an "object-oriented" approach to modelling intricate spacecraft during atmospheric entry, whereby simple geometric shapes such as cuboids, spheres, and cylinders are used to represent electronics, tanks, reaction wheels, etc. This approach is widespread in the debris demise community, allowing many current codes to utilise tumble-averaged aerodynamic and blunt-body heating correlations to assess the aerothermal loads acting upon the aforementioned debris analogues. However, while these methods are extremely computationally efficient, they tend to suffer from inaccuracies inherent in simplifying such complex phenomena into compact expressions. This in turn increases the likelihood of errors in estimations of the survivability of debris during atmospheric entry.

There are several alternate approaches which may be used to calculate the aerothermal loads in debris re-entry calculations, many of which are capable of providing greater insight into these loads than the aforementioned tumble-averaged and blunt body correlations. These methods differ in both complexity and accuracy, ranging from panel methods (which are simple enough to be evaluated using hand calculations) all the way to large scale CFD and DSMC simulations (which typically require computation on powerful workstations or large clusters). Additionally, it must also be noted that there is currently very little experimental data regarding the aerothermal behaviour of sharp-edged geometries in these regimes. This prevents experimental validation of numerical methods specific to debris re-entry problems, hence further increasing uncertainty.

The present work addresses this uncertainty in multiple ways. Firstly, panel methods are used to generate high resolution aerodynamic profiles for common debris shapes (cylinders and cuboids) with varying freestream Mach and Knudsen numbers, aspect ratios, and angles of incidence relative to the flow. These are then correlated using various techniques in order to create highly detailed aerodynamic models which may be queried in an extremely computationally efficient manner. Hence, the first source of inaccuracy mentioned – specifically, the use of tumble-averaged aerodynamic correlations – is addressed and mitigated.

Secondly, high fidelity simulations of the flow around cylindrical debris analogues are performed using DSMC codes, allowing the thermal distributions on these shapes to be studied. These data are then compared with experiments performed at the same conditions in Oxford's Low Density Wind Tunnel (LDWT), whereby heat transfer coefficients are compared between simulation and experiment. Various correlation techniques are also utilised here in order to normalise and reproduce surface distributions of heat transfer coefficient. As such, the second and third sources of uncertainty, i.e. the used of blunt body correlations and a lack of experimental validation, are similarly mitigated.

"I love deadlines. I like the whooshing sound they make as they fly by."

DOUGLAS ADAMS

ACKNOWLEDGEMENTS

I would, first of all, like to thank my family for their unending love and support, without which I could never have gotten as far in life as I have.

Many thanks are also due to my friends at the Osney lab (both past and present), who have gifted me with some of my fondest memories during the course of my DPhil.

The technical and workshop staff at Osney certainly deserve my gratitude, as without their expertise, my work would likely have taken twice as long to complete. A special thank you goes to Hal and Bill, whose work on the LDWT helped it live again.

Many thanks are due to Colin Wilson and Matt McGilvray, whose continual support and vast technical expertise have allowed me to not only refurbish the LDWT, but to also complete my experiments despite a great many setbacks.

Finally, I would like to thank my supervisor, Peter Ireland, without whom I would not be the engineer I am today.

CONTENTS

Nomenclature . . . 7

1. Introduction . . . 12
1.1 Motivation and scope . . . 14
1.2 Thesis content . . . 17

2. Literature review . . . 19
2.1 Satellite demise simulations . . . 19
2.2 Hypersonic fundamentals . . . 28
2.2.1 Rarefied gas dynamics . . . 28
2.2.2 The slip and transition regimes . . . 31
2.2.3 Mach number independence . . . 31
2.2.4 Surface heat flux distribution . . . 32
2.2.5 Thermal soak and peak heating . . . 36
2.3 Numerical analysis techniques . . . 38
2.3.1 Panel methods . . . 38
2.3.2 Bridging functions . . . 40
2.3.3 Stagnation point heat flux . . . 41
2.3.4 The DSMC method . . . 44
2.3.5 DSMC implementations . . . 46
2.3.6 DSMC analyses of re-entry vehicles . . . 48
2.4 Hypersonic wind tunnels . . . 49
2.4.1 Existing hypersonic rarefied facilities . . . 50
2.5 Experimental techniques . . . 51
2.5.1 Infrared thermography . . . 51
2.5.2 Thin film gauges . . . 52
2.5.3 Transient liquid crystal thermography . . . 53
2.5.4 Transient thermal experiments . . . 55
2.6 Statistical analysis . . . 58
2.7 Surrogate modelling . . . 60
2.7.1 Curve and surface fitting . . . 61
2.7.2 Gaussian process regression . . . 61
2.8 Chapter summary . . . 64

3. Numerical modelling . . . 66
3.1 Atmospheric modelling . . . 67
3.2 Trajectory simulations . . . 69
3.3 Panel method simulations . . . 73
3.3.1 Freestream conditions . . . 74
3.3.2 Geometry preprocessing . . . 75
3.3.3 Cross sectional area calculation . . . 76
3.3.4 Shielding analysis . . . 77
3.3.5 Aerodynamics . . . 79
3.3.5.1 Continuum aerodynamics . . . 79
3.3.5.2 Free molecular aerodynamics . . . 82
3.3.5.3 Transition aerodynamics . . . 83
3.3.6 Forces, moments, and aerodynamic coefficients . . . 84

3.3.7 Heat transfer . . . 86
3.3.8 Code validation . . . 87
3.3.8.1 Sphere . . . 87
3.3.8.2 Cylinder . . . 88
3.3.9 Correlation of results in terms of Knudsen number, $\boldsymbol{Kn}$ . . . 90
3.3.10 New aerodynamic models . . . 92
3.3.10.1 Cylinder . . . 92
3.3.10.2 Cuboid . . . 98
3.4 DSMC modelling . . . 102
3.4.1 Sampling locations . . . 105
3.4.2 DSMC results scaling . . . 107
3.4.3 *DS2V* simulations . . . 108
3.4.4 *dsmcFoam+* simulations . . . 115
3.4.4.1 Stagnation line . . . 117
3.4.4.2 Correlation of results in terms of angular coordinate, θ . . . 127
3.4.4.3 Front edge . . . 129
3.4.4.4 Side edge . . . 134
3.5 Summary and discussion of results . . . 138
3.6 Chapter summary . . . 143

4. Experimental facility . . . 145
4.1 Present experimental work . . . 145
4.2 The Oxford Low Density Wind Tunnel . . . 147
4.2.1 System description . . . 147
4.2.2 Refurbishment . . . 148
4.2.3 Test chamber layout . . . 149
4.2.4 Traverse mechanisms . . . 150
4.2.5 Gas conditioning . . . 151
4.2.6 Flow initialisation . . . 152
4.2.7 Instrumentation . . . 153
4.2.8 Leak rate determination . . . 156
4.3 Spatial variation of test flow . . . 157
4.3.1 Pitot measurements . . . 157
4.3.2 Stagnation chamber measurements . . . 161
4.3.3 Total temperature measurement . . . 161
4.3.4 Plume simulation . . . 163
4.3.4.1 Definition of CFD domain . . . 165
4.3.4.2 Slip versus no-slip CFD modelling . . . 166
4.3.4.3 Mesh independence . . . 169
4.3.4.4 No-slip CFD simulations . . . 172
4.3.5 Mach and Knudsen number calculation . . . 176
4.3.6 Surveys of regions of interest in the plume . . . 178
4.3.7 Temporal variability of test conditions . . . 182
4.4 Chapter summary . . . 186

5. Experimental methods . . . 188
5.1 Operating conditions . . . 188
5.2 Liquid crystal probes . . . 192
5.2.1 Calibration procedure . . . 193
5.3 Data analysis . . . 195
5.3.1 Video processing . . . 195
5.3.2 One dimensional (1D) heat transfer analysis . . . 197
5.3.3 Iterative FEA solution . . . 198
5.4 Chapter summary . . . 208

6. Experimental results . . . 210
6.1 Hemisphere . . . 210
6.2 Cylinder . . . 215
6.2.1 Performance of the 3D HTC solver code . . . 217
6.2.2 Flow condition 1 ($\varnothing = 15\,mm$, $Kn = 0.02$[16]) . . . 222
6.2.3 Flow condition 2 ($\varnothing = 10\,mm,\ Kn = 0.10$[12]) . . . 236
6.2.4 Comparison of both flow conditions . . . 240
6.3 Chapter summary . . . 245

7. Summary and conclusions . . . 249
7.1 Experimental facility . . . 249
7.2 Numerical studies . . . 250
7.2.1 Re-entry aerothermodynamics . . . 251
7.2.2 DSMC modelling of re-entry heating . . . 252
7.3 Experimental studies . . . 252
7.4 Future work . . . 254
7.5 Conclusion . . . 257

Bibliography . . . 258

A. Appendix A . . . 276
Thermochemical properties of air at low temperatures . . . 276

B. Appendix B . . . 278
Analytic LDWT performance model . . . 278

C. Appendix C . . . 283
Semi-empirical LDWT performance model . . . 283

D. Appendix D . . . 287
Nozzle throat diameter measurement . . . 287

E. Appendix E . . . 289
Original LDWT specifications . . . 289

NOMENCLATURE

PHYSICAL QUANTITIES

A	Area (m^2)
a	Speed of sound $(m \cdot s^{-1})$
c	Molecular speed $(m \cdot s^{-1})$, Specific heat of a solid $(J \cdot kg^{-1} \cdot K^{-1})$
C_P	Specific heat of a fluid at constant pressure $(J \cdot kg^{-1} \cdot K^{-1})$
C_V	Specific heat of a fluid at constant volume $(J \cdot kg^{-1} \cdot K^{-1})$
$\varnothing$	Diameter (m)
D	Mass diffusivity $(m^2 \cdot s^{-1})$, Number of dimensions
d	Mean molecular diameter (m), Index of a dimension
E_k	Kinetic energy (J)
f	Velocity distribution function
G	Gravitational constant $(m^3 \cdot kg^{-1} \cdot s^{-2})$
g	Gravitational acceleration $(m \cdot s^{-2})$
g_0	Gravitational acceleration at Earth's surface $(m \cdot s^{-2})$
h	Altitude (m), Specific enthalpy $(J \cdot kg^{-1} \cdot K^{-1})$, Heat transfer coefficient $(W \cdot m^{-2} \cdot K^{-1})$
I	Pixel intensity
K	Kernel function
k	Conductive heat transfer coefficient $(W \cdot m^{-1} \cdot K^{-1})$
k_B	Boltzmann constant $(m^2 \cdot kg \cdot s^{-2} \cdot K^{-1})$
ℓ	Lengthscale for GPR kernels
L	Characteristic length (m)
M	Molar mass $(kg \cdot mol^{-1})$
m	Mass (kg)
M_E	Mass of Earth (kg)
mcs	Mean collision separation (m)
N	Normal direction, Total number of particles, Total number of elements in a set
n	Number density (m^{-3}), Index of an item in a set
N_A	Avogadro constant (mol^{-1})
p	Pressure (Pa)
$\dot{q}$	Heat flux $(W \cdot m^{-2})$
Q	Volume flow rate $(l \cdot s^{-1})$
R	Perfect gas constant $(J \cdot kg^{-1} \cdot K^{-1})$, Range of a data set
R_E	Radius of Earth (m)
R_n	Equivalent nose radius (m)
s	Molecular speed ratio
$\bar{t}$	Mean free time (s)
T	Temperature (K)
t	Time (s)
U, v	Velocity $(m \cdot s^{-1})$
V	Volume (m^3)
w_{DSMC}	DSMC particle weighting factor
x	Observation coordinate
x'	Prediction coordinate
x, y, z	Principal Cartesian directions
x_{tp}	Thermal pulse penetration depth (m)
y	Observed value
y'	Predicted value

GREEK

α	Thermal diffusivity ($W \cdot m^{-1} \cdot K^{-1}$), Pitch angle(*rad*), Additional factor in GPR-RQ kernels
α_T	Thermal accommodation coefficient
β	Yaw angle (*rad*), Non-dimensional temperature variable
δ	Angle between freestream and panel normal vectors (*rad*), Precentage error (%)
ϵ	Emissivity, Absolute error
η	Relative error
γ	Ratio of specific heats
κ	Density ratio between freestream and the region downstream of a normal shock
λ	Mean free path (m)
μ	Dynamic viscosity ($Pa \cdot s$), Mean of a data set
ν	Kinematic viscosity ($m^2 \cdot s^{-1}$)
Ω	Solid angle (rad)
ψ	Roll angle (*rad*)
ρ	Density ($kg \cdot m^{-3}$)
σ	Effective particle diameter (m), Surface accommodation coefficient, Stefan–Boltzmann constant ($W \cdot m^{-2} \cdot K^{-4}$)
σ_f	Variance from the mean
τ	Shear force (N)
θ	Surface or shock inclination angle (dependant on subscript) (rad), Non-dimensional temperature, Geometry angular coordinate (rad)
φ	Flight path angle (rad)

DIMENSIONLESS NUMBERS

Kn	Knudsen number
Le	Lewis number
Ma	Mach number
Nu	Nusselt number
Pr	Prandtl number
Re	Reynolds number
Sc	Schmidt number
St	Stanton number

COEFFICIENTS

C	General coefficient
C_A	Coefficient of axial force
C_A, C_M, C_N	Aerodynamic moment coefficients (object reference frame)
C_D, C_L, C_S	Aerodynamic force coefficients (object reference frame)
C_f	Coefficient of friction
C_p	Coefficient of pressure
C_x, C_y, C_z	Aerodynamic force coefficients (Cartesian reference frame)
C_τ	Coefficient of shear
C_{mx}, C_{my}, C_{mz}	Aerodynamic moment coefficients (Cartesian reference frame)
R^2	Coefficient of determination

VECTORS

$\boldsymbol{\beta}$	Vector of coefficients
$\hat{N}$	Unit surface normal vector
$\hat{S}$	Unit surface shear vector
$\hat{V}$	Unit freestream direction vector
$\mathbf{x}'$	Vector of prediction coordinates
$\mathbf{x}$	Vector of observation coordinates
$\mathbf{y}'$	Vector of predicted values
$\mathbf{y}$	Vector of observed values
$\overrightarrow{F}$	Force vector (N)
$\overrightarrow{L}$	Moment arm (vector from COG to panel centre)
$\overrightarrow{M}$	Moment vector
$\mathbf{c}$	Molecular velocity vector ($m \cdot s^{-1}$)
$\mathbf{F}$	External force vector (Boltzmann eq.)
$\mathbf{r}$	Molecular position vector (m)

SUBCRIPTS

0	Stagnation quantity/coordinate
ARD	Automatic relevance determination
A	Axial
cat	Catalytic
$cond$	Conduction
$cont$	Continuum condition
$conv$	Convection
c	Conic surface quantity
CFE	Cylinder front edge
cs	Cross section
CSE	Cylinder side edge
CSF	Cylinder front stagnation line
CSS	Cylinder side stagnation line
D	Drag
EQ	Exponential quadratic
eq	Equilibrium
$exp.$	Exerimental result
E	Equivalent
e	Boundary layer edge coordinate
ex	Condition at nozzle exit
FE	Finite element simualtion result
fm	Free molecular condition
fr	Frozen
∞	Freestream quantity
I	Pixel intensity
i	Initial value
i, j, n	Indices
L	Lift
new	Current solution in an iterative process
$noncat$	Non-catalytic
old	Previous solution in an iterative process
$pumps$	LDWT pumping system
rad	Radiation
RQ	Rational quadratic
r	Relative value, Value at a specific radial coordinate, Recovery value
ref	Reference quantity
SE	Squared exponential

s	Post-shock quantity, Surface quantity
tp	Thermal pulse
$trans$	Transition condition
T	Transverse direction
th	Condition at nozzle throat
w	Wall quantity

SUPERSCRIPTS

(i)	Index value from a range of integers
$*$	Post collision value

ACRONYMS / CONVENTIONS

Abs.	Absolute measurement
AFRL	Air Force Research Lab (Wright-Patterson Air Force Base)
ATV	Automated Transfer Vehicle
BNC	Bayonet Neill–Concelman
BRL	Belstead Research Limited
CBAERO	Configuration Based Aerodynamics
CDG	Capacitive Diaphragm Gauge
CEV	Crew Exploration Vehicle
CFD	Computational Fluid Dynamics
CFL	Courant–Friedrichs–Lewy
CNES	Centre National d'Etudes Spatiales
COG	Centre Of Gravity
CPU	Central processing unit
DAS	Debris Assessment Software
Diff.	Differential measurement
DOF	Degree(s) of freedom
DRAMA	Debris Risk Assessment and Mitigation Analysis
EDL	Entry, Descent, and Landing
EDL	Entry, descent, and landing
EM	Electromagnetic
EPS	Encapsulated Post Script (file format)
ESA	European Space Agency
FEA	Finite Element Analysis
FEM	Finite Element Method
FGE	Fluid Gravity Engineering
FOSTRAD	Free Open Source Tool for Re-entry of Asteroids and Debris
FPS	Frames Per Second
GPR	Gaussian Process Regression
GUI	Graphical user interface
HVV	High vacuum valve
IR	Infrared
ITAR	International Traffic in Arms Regulations
LC	Liquid Crystal
LDWT/LDT	Low Density Wind Tunnel
LEO	Low Earth Orbit
LHS	Latin Hypercube Sampling
MD	Molecular Dynamics
MDG	Metallic Diaphragm Gauge
NASA	National Aeronautics and Space Administration
ONERA	Office National d'Etudes et de Recherches Aérospatiales
OpenFOAM	Open source Field Operation and Manipulation Toolkit

ORSAT Object Re-entry Survival Analysis Tool

OTI Oxford Thermofluids Institute

PDF Probability Density Function

PID Potential-Integral-Derivative

PIL Python Imaging Library

PMMA Polymethyl Methacrylate (Perspex)

PSI Pounds per Square Inch

RAC Re-entry Aerothermal Calculator

RANS Reynolds Averaged Navier-Stokes

RMS Root-Mean-Squared

RSD Relative Standard Deviation

RSS Root of Sum of Squares

S/HABP Supersonic/Hypersonic Arbitrary Body Program

SAM Spacecraft Aerothermal Model

SARA Survival And Risk Analysis module

SCARAB Spacecraft Atmospheric Re-entry and Aero-thermal Breakup

SESAM Spacecraft Entry Survival Analysis Module

SLPM Standard Litres Per Minute

STL Stereolithography (file format)

STS Space Transport System (Space shuttle)

TFHTG Thin Film Heat Transfer Gauge

THTAC Transient Heat Transfer Analysis Code

VKI The von Karman Institute for Fluid Dynamics

INTRODUCTION 1

Space debris is fast becoming one of the great environmental challenges of the 21st century. With improvements in the miniaturisation of electronics and the reduction of space launch costs, it is easier than ever to place satellites in Earth's orbit. However, since the first artificial satellite launched in 1957, there has been relatively little concern given to the end of a spacecraft's life, or indeed, how to remove them from orbit once they are no longer of use.

Now, over six decades later, there are approximately 4000 defunct satellites in orbit around the Earth, which have contributed to over 500 collisions, explosions, and break-ups. These events, amongst others, often result in the fragmentation of this so-called "space junk", releasing swarms of dangerous debris – much of which cannot be tracked using current technology – into the paths of other spacecraft. As a result, there are now an estimated 34000 objects which are less than 10 cm in diameter, and a projected 128 million objects with diameters between 1 mm and 1 cm [1].

This debris poses a significant threat to vital orbital infrastructure such as telecommunications and weather monitoring satellites, not to mention manned missions such as the International Space Station (ISS). At orbital velocities, even a chip of paint can cause severe damage to a spacecraft, as is evidenced by Figure 1.1, which shows a 7 mm gouge in a window of the International Space Station's Cupola module. The cause of this damage was determined to be a metal fragment of the order of micrometers in diameter.

An additional – and perhaps even more dire – risk that arises from space debris is that of an object successfully re-entering the Earth's atmosphere and impacting the ground. Events like this have already occurred several times in recent decades, one of the most well known being that of a Delta 2 upper stage propellant tank crashing near Georgetown, Texas in 1997 (see Figure 1.2). Although no harm to the populace or damage to property occurred, it is a reminder of the threat that these pieces of debris pose.

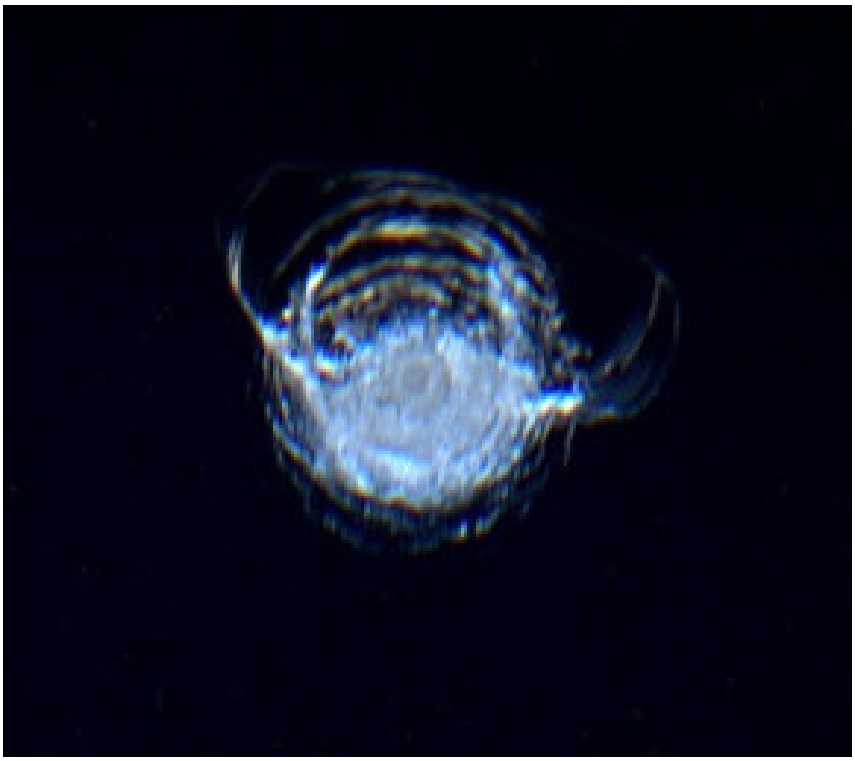


Figure 1.1: A 7 mm chip in a window of the Cupola module aboard the International Space Station caused by a piece of space debris of the order of micrometers in diameter [1]



Figure 1.2: A propellant tank from a Delta 2 rocket which re-entered the atmosphere and impacted near Georgetown, Texas, USA in 1997 [2]

[1] `https://m.esa.int/spaceinimages/Images/2016/05/Impact_chip`
[2] `https://m.esa.int/spaceinimages/Images/2015/12/Debris_landed_in_Texas`

In order to mitigate the ever growing population of space debris, a number of initiatives have been launched by global space entities, such as the European Space Agency's (ESA) Clean Space programme. This endeavour seeks to prevent pollution of the near-Earth space environment by reducing orbital impacts, safely disposing of existing debris, and mitigating future growth of the debris population. Figure 1.3 shows an infographic produced by ESA which outlines the goals of the Clean Space initiative.

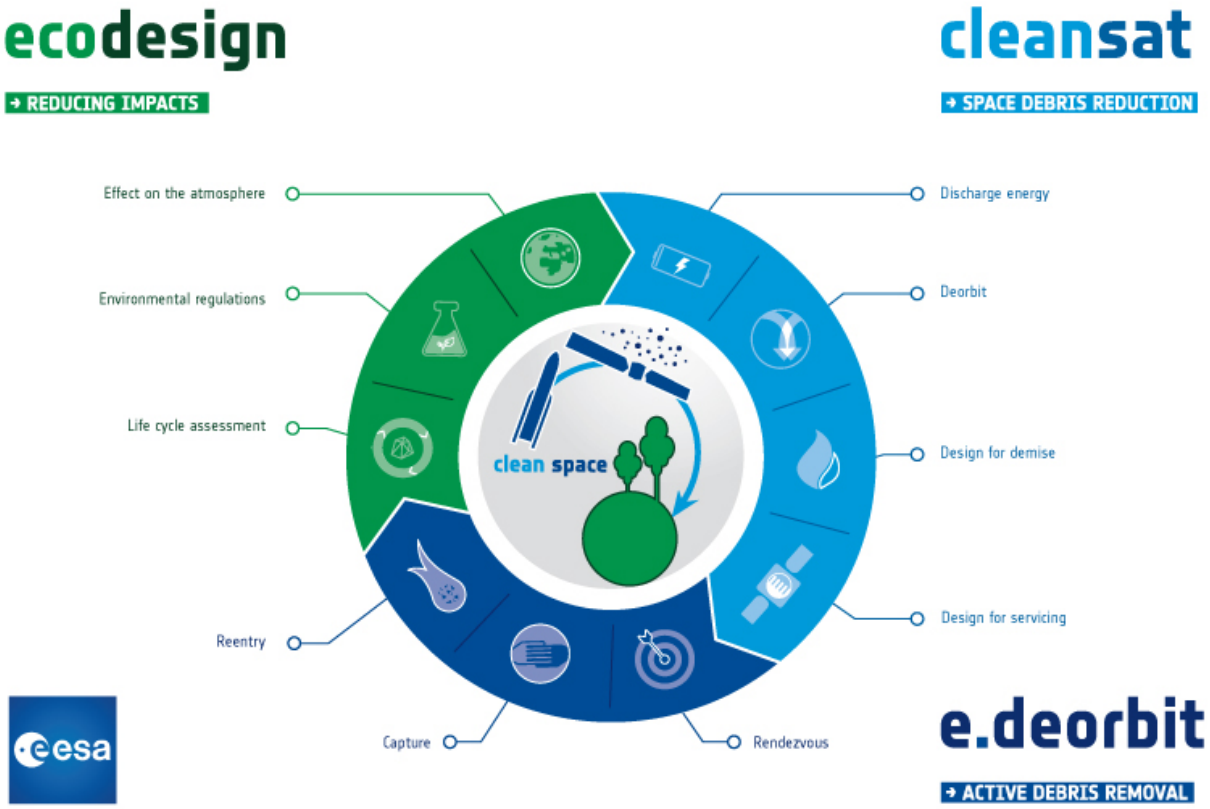


Figure 1.3: An infographic showing an overview of ESA Clean Space initiate activities [2]

1.1 MOTIVATION AND SCOPE

This thesis deals with those issues raised by the "cleansat" section of the Clean Space initiative, wherein improvements to the accuracy and reliability of destructive re-entry calculations for defunct satellites are sought. The goals of this initiative, however, are not exclusively of interest to ESA. Due to the importance of keeping Earth's orbital environment free of harmful debris, space agencies, research institutions, and private companies around the world have developed their own tools for predicting the survivability of space debris during atmospheric entry. One of the most popular methods for performing such calculations is known as the "object-oriented"

approach, whereby complex satellite geometries are built from simple geometric shapes such as cylinders, cuboids, spheres, etc. These shapes serve as geometric analogues of common spacecraft components like propellant tanks, electronics enclosures, structural elements, reaction wheels, etc.

These tools, however, often suffer from inaccuracies, borne of the need for computational efficiency. When calculating re-entry trajectories, it is necessary to iteratively model the passage of an object through an atmosphere. This is due to variances in an object's altitude (which affects the freestream conditions), velocity, orientation (attitude), etc., which require recalculation at every step of a trajectory. As such, models which average the aerothermodynamic characteristics of these debris analogues over a range of attitudes are often utilised. Owing to their simplicity, models like these are extremely computationally efficient, but can lead to large uncertainties in the overall survivability and potential final landing sites of re-entering debris.

This is especially important in the slip and transition regimes, those types of fluid flow which are denser than the upper atmosphere yet sparser than the altitudes at which commercial aircraft fly. These regimes are extremely important for demise simulations as they can have significant effects on the trajectories and heating distributions of re-entering spacecraft and debris. Despite this, however, these regimes are still not well understood, with many so-called "bridging laws" being used to interpolate between aerodynamic and heating correlations developed for the more well understood continuum and free molecular regimes. There is also a distinct lack of experimental data pertaining specifically to space debris re-entry, as most studies in the slip and transition regimes tend to focus on aerodynamically optimal shapes such as spheres, cones, lifting bodies, capsules, etc., and not the sharp-edged geometries which are so often used as debris analogues.

Figure 1.4 shows an example of a tumble-averaged model used in ESA's DRAMA debris demise toolkit [3]. The model, which returns the drag coefficient for cylindrical debris items, only has two variables which it can accept as inputs: Knudsen number (a freestream quantity linked to the level of rarefaction in a gas – see Section 2.2.1), and the cylinder's aspect ratio (D/L, or $L/\varnothing$, depending on the formulation, where $\varnothing$ or D is the cylidner's diameter, and L is its length). Note that the model assumes constant magnitudes for the drag coefficients when

outside of the transition regime, and bridges these values with a simple sine-squared law.

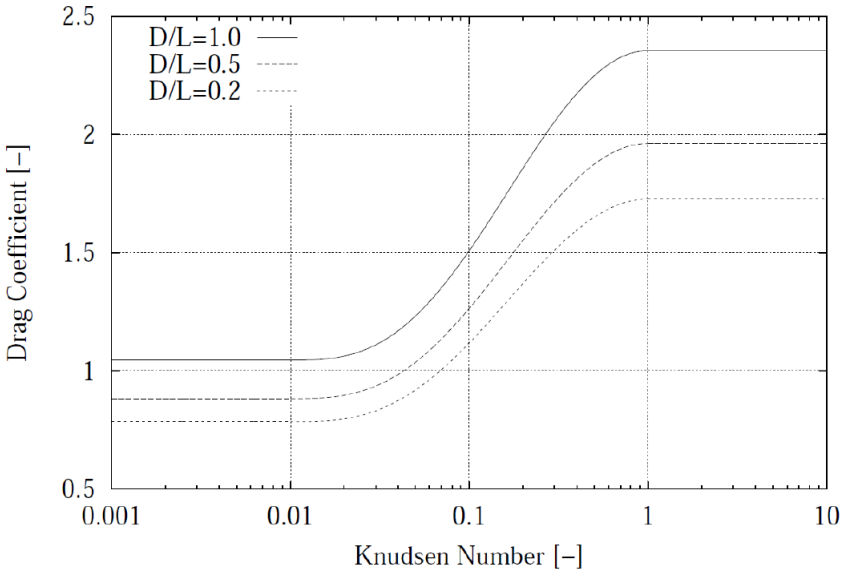


Figure 1.4: An example of tumble-averaged aerodynamic drag coefficient models for a series of cylinders with differing aspect ratios. These curves are utilised in the aerodynamic module of the ESA DRAMA debris demise simulation toolkit [3].

In order to mitigate these uncertainties for future debris demise predictions, two goals are set for the present work.

The first is the development of a method for quickly calculating the aerodynamic coefficients of re-entering debris analogues (the aforementioned cylinders, cuboids, sphere, etc.). This method must be more accurate than the tumble-average correlations currently in use by many debris demise codes, but must also be computationally efficient enough to ensure that iterative trajectory simulations can be performed in a timely fashion.

The second goal of the present work deals with the examination of the thermal behaviour of these debris analogues. At present, blunt body correlations are often used to assess the thermal surface distributions of debris items during re-entry simulations. Like the tumble-averaged aerodynamic correlations, these equations are usually heavily simplified representations of the actual physical laws governing these complex heat transfer phenomena. As such, they tend to have large inaccuracies associated with their predictions, due in part to the aforementioned presence of sharp edges on the debris analogues, features which these correlations cannot easily take into account. Additionally, without experimental data to validate the data generated by new codes and correlations, the uncertainties in these aerothermodynamic predictions cannot truly be

mitigated.

As such, the second goal of the present work is to study these debris analogues in the relevant thermal environment, and hence, obtain experimental data pertaining to the their behaviour in the slip and transition regimes. A campaign in an appropriate experimental facility is necessary, whereby the heat transfer to debris analogues may be measured and characterised. These data may then be used to validate current tools and correlations, gauging their suitability for the study of atmospheric re-entry of space debris items in the slip and transition regimes.

1.2 THESIS CONTENT

At the beginning of the present work, a literature review of the current state of the art in rarefied gas dynamics simulation and transient thermal experiments is presented, as are descriptions of many current debris demise codes in use around the world. This comprises chapter 2.

Chapter 3 consists of the development and validation of a reduced-order panel method hypersonic aerothermodynamics code called *RAC* (Re-entry Aerothermal Calculator). This code is used in conjunction with the Gaussian Process Regression (GPR) modelling package *GPy* and a series of correlative curves newly developed by the present author to characterise the aerodynamic behaviour of debris analogues (cylinders and cuboids of varying size and aspect ratio). These models are found to be highly accurate (relative to their respective input datasets), presenting a clear improvement over tumble-averaged correlations, whilst still being computationally efficient enough to be an order of magnitude faster than *RAC* when queried.

Simulations performed using state of the art DSMC (Direct Simulation Monte Carlo) codes (Bird's *DS2V*, and the University of Strathclyde's *dsmcFoam+*) are then presented, wherein GPR modelling, correlative curves, and a novel normalisation technique for non-optimal DSMC simulations are incorporated. These tools and techniques are subsequently shown to be suitable for studying the complex thermal behaviour of satellite analogues in the slip and lower transition regime.

Chapter 4 describes the experimental facility used during the present work. The Oxford Low Density Wind Tunnel (LDWT) is a continuous flow, vacuum wind tunnel intended for hypersonic rarefied gas studies in the slip and transition regimes. The refurbishment of the

facility is elaborated upon, wherein new data acquisition and experimental support equipment is described. This section then culminates in a characterisation study of the facility's hypersonic plume, utilising pitot pressure and total temperature surveys alongside traditional no-slip CFD simulations.

Chapter 5 enumerates the methods used to setup, perform, and subsequently post-process the present work's experimental studies. The process used to determine and select operating conditions in the LDWT is described, followed by the selection of regions of interest in the experimental plume produced by the LDWT nozzle. The design and construction of the liquid crystal probes is then described, as is the calibration procedure for the LC coatings. Finally, the algorithm used to analyse the experimental videos and hence determine surface heat transfer coefficient distributions is described. An in-house code called THTAC (which is based upon the semi-infinite substrate assumption) is used to analyse some of these cases, while others are analysed with a newly developed post-processing code. This program, named the "3D HTC solver" (and abbreviated "3D-HTCS") is based upon work by Ryley [4], formerly of the Osney Thermofluids Institute (OTI), and is used to map a video of a liquid crystal experiment on an FEA mesh. Thermal finite element simulations are then run automatically by the code, wherein simulated liquid crystal responses are generated after each iteration of the FEA solver. These are then compared to the experimental video, and corrections to the heat transfer coefficient (HTC) profile made. In this fashion, the code converges on the "true" HTC distribution experienced by an experimental probe.

Experiments performed for the current work are then described in Chapter 6, wherein the results from various experiments on hemispherical and cylindrical probes are described. It is found that the results from the 3D HTC solver and DSMC simulations agree extremely well in many cases, with most errors arising from poor video quality or shallow viewing angle during experiments.

Finally, in Chapter 7, a summary of the present work is given, overall conclusions are drawn, and suggestions for future development of the research described herein are presented.

LITERATURE REVIEW 2

This chapter deals with the current state of the art in those research areas related to the present work. Firstly, a summary of current debris demise software is presented, followed by a review of the fundamental physical phenomena relevant to the study of atmospheric entry, and numerical methods used for the simulation of hypersonic aerothermodynamics. Experimental facilities and techniques are then elaborated upon, followed by statistical analysis and surrogate modelling techniques for analysing the experimental and simulated data produced during the course of the present work.

2.1 SATELLITE DEMISE SIMULATIONS

An important capability in modern space engineering is the ability to quickly and accurately assess the danger posed by spacial debris to equipment and personnel. It is currently mandated by international agreements that defunct space equipment must be removed from orbit following the conclusion of its mission lifetime [5], and with the current lack of craft like the STS (NASA Space Transportation System, i.e. the Space Shuttle) which could facilitate a recovery mission for orbiting debris, there has been an increase in the interest of space agencies in the simulation and rapid assessment of destructive re-entry trajectories. As such, several advanced software programs have been developed by NASA and ESA (amongst others) in order to perform these complex calculations, and hence assess the probability of dangerous ground impacts.

There are currently multiple programs in use in this field, but in a general sense there are only two operative levels of fidelity: low resolution solvers for simple, bulk analyses of objects, and high resolution solvers with improved models for more precisely calculating the survivability of re-entering spacecraft via heat fluxes and aerodynamic forces. These higher fidelity models are only employed when the low fidelity solvers indicate that an object will be non-compliant with a safe re-entry criterion e.g. incomplete destruction resulting in a significant risk to populated areas. The codes either feature object-oriented geometry modelling (wherein spacecraft components are represented by primitives such as cuboids, spheres, cones, cylinders, and partial segments thereof), or spacecraft-oriented geometry modelling (wherein the entire spacecraft is represented

by generalised mesh elements).

SESAM

The ESA low fidelity program, SESAM, is a component of the DRAMA suite's SARA module. SESAM is an object-oriented code, and assesses only the aerothermodynamic properties of a spacecraft. Its heat transfer model uses a lumped mass approach derived from NASA's ORSAT code, which applies a tumble-averaged, uniform heat flux across the surface of each component [6]. This heat flux is derived from the empirical stagnation point convective heat flux expressions of Detra, Kemp and Riddell (Eqs. (2.37) and (2.38)). Using these measures, and its built-in temperature independent material database, SESAM is able to estimate the survivability of a spacecraft (and of its individual components after a fragmentation event) during re-entry, and hence plot the likely impact locations of surviving fragments [7] [8] [9].

SCARAB

The ESA high fidelity code, SCARAB (developed by HTG GmbH), is a modular spacecraft-oriented analysis code which assesses the survivability of components using a series of coupled equations and a high resolution geometric model. The program simulates the evolution of aerothermal loads using a two-dimensional heat conduction model and assesses structural loads and spacecraft motion using a 6-DOF system of dynamic equations [10] [11]. SCARAB is unique among the tools described here in that it analyses a spacecraft geometry as a whole, decomposing it into a surface mesh composed of triangulated elements. The geometry may still be generated using the placement of primitive three dimensional shapes, but analyses are conducted using local panel methods, which are generalised over the entire model as opposed to being specially formulated for a certain shape. SCARAB is also able to simulate partial melting of elements using the temperature dependant models included in its materials database. In this fashion, SCARAB may assess the survivability of a spacecraft both before and after a fragmentation event, tracking dynamically any generated debris fragments as their trajectories evolve independently of one another [7] [8] [12].

DAS

The NASA low fidelity program, called DAS is an integrated analysis tool (unlike SESAM, which is split into separate modules) which utilises a lumped parameter, object-oriented analysis to assess the survivability of objects. DAS subjects spacecraft models to a generalised analysis using simplified equations (such as the US 1976 Standard Atmosphere for atmospheric density calculations, and a basic, overall heat accumulation model) to iteratively assess the thermal loads acting on the components. The thermal analyses of all components begin at the same, pre-defined altitude and, as partial melting of components and fragmentation events are not modelled, the spacecraft is only considered demised when the overall measure of the thermal energy it has absorbed rises above the heat of ablation for the given materials. Using these data, and a library of temperature independent material properties, DAS is able to arrive at highly conservative estimates as to whether a given object will comply with the relevant safety regulations [7] [8].

ORSAT

The high fidelity NASA tool, ORSAT, is used to model objects which are deemed non-compliant with safe re-entry regulations following an analysis run using DAS. ORSAT utilises the same geometry modelling procedure as DAS, but instead of simplified expressions features a one dimensional heat conduction model with the ability to model partial component melting, and a 3-DOF equation of motion routine. This system of equations allows the program to iteratively assess the evolution of thermal loads across the craft and to also determine its tumbling attitudes and rates, ultimately culminating in a high resolution determination of the object's survivability, the probability of its fragmentation, and the landing location of any components following a successful, i.e. not entirely destructive, re-entry [7] [8] [12].

SAM

SAM is an engineering level tool developed by Fluid Gravity Engineering Ltd. It is designed to marry the computational efficiency of the object oriented model with the physical accuracy of a breakup methodology. The code features numerically defined joints

between components which are included in the shielding solver's calculations, i.e. the determination of whether the freestream flow will interact with parts of a spacecraft due to them being obscured by other parts of the geometry [13]. Once these joints reach their respective melting temperatures (based on their assigned materials), they are deleted from the simulation, and the component(s) that they coupled are separated from the main object.

Initially, the re-entering object is modelled using a full 6-DOF trajectory solver, such that the heating of the aforementioned joints is modelled as accurately as possibly. Following breakup, the motions of individual components are modelled using a reduced order 3-DOF trajectory solver, and utilise a combination of analytical and empirical calculations in order to assess their aerothermal behaviour [14]. Hence, the probability of their re-entering successfully, and their likely landing coordinates (if applicable) are calculated.

PAMPERO

Developed by CNES, PAMPERO is a spacecraft-oriented code [15]. Its main innovation over other engineering level re-entry codes is its calculation of local geometric curvature during simulations. These data are then used in heating calculations along with a full 3D heat transfer model in order to better predict the thermal loads on objects. This is accomplished by considering the effective surface curvature of panels around an element whose thermal loads are being evaluated, thereby better informing what is otherwise an entirely local calculation.

FOSTRAD

FOSTRAD is an implementation of the object oriented method developed by the Aerospace Centre of Excellence at the University of Strathclyde [16]. The code is written in Matlab, and is meant to be entirely open source, allowing it to be utilised by academics and industry analysts alike.

The code operates in a similar fashion to the SAM FGE code, but makes use of a local curvature calculation much like that in PAMPERO. It also features a suite of surrogate models computed using experimental data gathered from literature and DSMC simulations. These models are applied to objects in the code during a simulation as corrective factors that

bring the predictions of the code's panel methods into line with higher fidelity experimental and numerical results. This has been shown in several test cases to produce superior agreement with validation data [17] [18].

MUSIC/FAST

MUSIC/FAST is a spacecraft-oriented code developed by ONERA [19], which features both an aerothermal calculation component (FAST) and a trajectory propagation tool (MUSIC). Like PAMPERO, the code's FAST module features local curvature based heat flux methods which allow thermal concentrations near sharp edges to be more accurately modelled. An efficient panel shielding algorithm is also included in order to assess which faces of a mesh are obscured by upstream portions of the geometry.

DEBRISK

DEBRISK is a object-oriented code developed by CNES [20] which operates in a similar fashion to other such codes described above. It features a series of different Earth atmosphere models, a lumped-mass thermal model for individual components, and an ablation subroutine which allows reductions in a component's mass due to melting of surface layers to be simulated. An interface to Orekit (ORbits Extrapolation KIT), a Java-based trajectory propagation code is also included in the code.

An extended summary of the above codes is presented in Tables 2.1 through 2.4. Therein, question marks (?) indicate that no information could be found at the time of writing in open literature regarding the functionality of a specific code. Also, note that the term "surrogate model" is used in a general sense, and refers to any form of correlative model informed by external data, from simple curve fits to multi-dimensional Gaussian process regression (GPR) or radial basis function (RBF) models.

Code name	Developer	Geometry model	Aerodynamic model	Heat transfer model (internal)
SESAM	ESA	Object	Tumble-averaged, $(Ma = \infty)$	Lumped mass, (w. material melting)
SCARAB	HTG	Spacecraft	Panel method	2D finite element
DAS	NASA	Object	Fixed aero coefficients	Lumped mass
ORSAT	NASA/Lockheed-Martin	Object	Tumble-averaged, $(Ma = \infty)$	1D finite element
SAM	BRL/FGE	Object (joint-based)	Panel method	Heat Balance Integral, Lumped mass, 1D FEA
PAMPERO	CNES	Spacecraft	Panel method	3D finite element
FOSTRAD	Uni. of Strathclyde	Spacecraft	Panel method, w. surrogate model corrections	1D finite element
MUSIC/FAST	ONERA	Spacecraft	Panel method, w. surrogate model corrections	N/A (planned development)
DEBRISK	CNES	Object	Tumble-averaged, $(Ma = \infty)$	Lumped mass, (w. material ablation)
RAC	Present author	Spacecraft	Panel method	N/A (planned development)

Table 2.1: Summary of developers and modelling strategies of current debris demise codes.

Code name	Continuum regime	Transition regime	Free molecular regime
SESAM	Geometry-specific sigmoid	Global bridging, Geometry-specific sigmoid	Geometry-specific sigmoid
SCARAB	Modified Newtonian	Global bridging, Empirical law	Schaaf-Chambre
DAS	Fixed ($C_D = 2.2$)	Fixed ($C_D = 2.2$)	Fixed ($C_D = 2.2$)
ORSAT	Geometry-specific sigmoid	Global bridging, Geometry-specific sigmoid	Geometry-specific sigmoid
SAM	Modified Newtonian	Global bridging, Geometry-specific sigmoid	Schaaf-Chambre
PAMPERO	Modified Newtonian	Global bridging, Empirical law	?
FOSTRAD	Modified Newtonian	Local bridging, Radius-based surrogate model	Schaaf-Chambre
MUSIC/FAST	Modified Newtonian	Local bridging, Blanchard-derived sine-squared law	Bird
DEBRISK	Geometry-specific sigmoid	Global bridging, Geometry-specific sigmoid	Geometry-specific sigmoid
RAC	Modified Newtonian, Tangent wedge, Tangent cone	Global bridging, Wilmoth sine-squared law, Local bridging, Potter-Peterson (planned development)	Schaaf-Chambre

Table 2.2: Summary of aerodynamic models used in current debris demise codes.

Code name	Continuum regime	Transition regime	Free molecular regime
SESAM	Detra-Kemp-Riddell	Global bridging, Stanton number (spheres, cylinders), Exponential (flat plates, boxes)	$\dot{q}_{0_{cont}} = \frac{\alpha_T \rho U_\infty^3}{2}$; $\alpha_t = 0.9$ Geometry-specific shape factors
SCARAB	Modified Lees	Local bridging, Exponential heat flux (Legge)	Schaaf-Chambre (assume $St \sim 1$)
DAS	?	?	?
ORSAT	Detra-Kemp-Riddell (stagnation) Geometric shape factor	Global bridging, Stanton number (spheres, cylinders), Exponential (flat plates), Empirical (boxes)	$\dot{q}_{0_{cont}} = \frac{\alpha_T \rho U_\infty^3}{2}$; $\alpha_t = 0.9$ $St = f(s)$ (flat plates) Empirical (boxes)
SAM	Detra-Hidalgo (spheres), Poll (cylinder surface), Eckert (plates, boxes), Klett (boxes, cylinder ends)	Local bridging, Stanton number bridging (Legge)	Schaaf-Chambre
PAMPERO	Experiment/CFD derived correlations (curvature-based, proprietary)	Local bridging, Empirical bridging law	?
FOSTRAD	Fay-Riddell, Van Driest, Detra-Kemp-Ridell	Local bridging, Exponential heat flux (Legge)	Schaaf-Chambre
MUSIC/FAST	Detra, Verant-Lepage, Verant-Sagnier	Local bridging, Blanchard-derived sine-squared law	Bird
DEBRISK	Tumble-averaged (Cropp/Klett)	Tumble-averaged (Cropp/Klett)	Tumble-averaged (Cropp/Klett)
RAC	Brandis-Johnston, Fay-Riddell, Kemp-Rose-Detra, Lees, Detra-Hidalgo-Newtonian	Local bridging, Stanton number bridging (Legge)	Schaaf-Chambre

Table 2.3: Summary of external heat transfer models used in current debris demise codes.

Code name	Atmospheric model	Trajectory model	Breakup model	References
SESAM	US76	3 DOF	Pre-defined altitude	[3] [21]
SCARAB	MSISe-00	6 DOF	Material melting/overstress (mesh element demise)	[22] [23] [24]
DAS	Jacchia77, US76	3 DOF	Pre-defined altitude	[25]
ORSAT	MSISe-90	3 DOF	Pre-defined altitude	[23] [24] [26]
SAM	US76, CIRA-JB-2008, Jacchia-Roberts, NRLMSISE00	3/6 DOF	Material melting/overstress (simulated joint demise)	[13] [27]
PAMPERO	?	6 DOF	Material melting/overstress (mesh element demise)	[15] [28] [29]
FOSTRAD	US76	3 DOF	N/A	[17] [18] [16] [30]
MUSIC/FAST	US76-Barlier	3/6 DOF	N/A	[19] [28] [29]
DEBRISK	CIRA88-MSISE-90, US76, NASA Glenn, Exponential	3 DOF	Pre-defined altitude	[20]
RAC	NRLMSISE00	3 DOF, 6 DOF (planned development)	N/A	[31] [32]

Table 2.4: Summary of atmospheric, trajectory, and breakup models used in current debris demise codes.

2.2 FUNDAMENTALS OF HYPERSONIC FLUID PHYSICS

Before further examining space debris re-entry, it is important to first describe those elements of fluid mechanics and thermodynamics that act as prerequisites for the current work. In this section, the equations and non-dimensional quantities used to describe the fluid flows encountered during re-entry are explained, followed by the "regimes" into which such flows are grouped. This grouping allows treatment with different numerical and experimental methods which are often only valid in a single regime.

Further to the above, principles relevant to the present work (and specific to rarefied hypersonic flows and atmospheric re-entry) are then described. These include the Mach independence principle (the effects of which are an important consideration when performing high speed flow experiments), variations in surface heat flux distribution with Knudsen number (i.e. the level of rarefaction at a given altitude), and the "thermal soak" phenomenon encountered during a re-entry trajectory (which affects the overall heat flux and maximum temperature of a re-entering object).

2.2.1 Rarefied gas dynamics

Rarefied gas dynamics is the name given to the branch of fluid mechanics that deals with the behaviour of low density flows, and is described in great detail by both Bird [33] and Anderson [34]. Assessments of heat loads and aerodynamic forces during simulations of atmospheric entry rely on the specialised formulations that underpin this branch of physics.

There are many dimensionless parameters which allow the categorisation, comparison, and scaling of such flows, four of the most important being the Mach (Eq. (2.1)), Reynolds (Eq. (2.2)), Knudsen (Eq. (2.3)), and Stanton (Eq. (2.4)) numbers. Note that first three of these quantities may also be related to one another as is also noted in Eq. (2.3)). These measures are useful in the comparison and scaling of flows because they represent ratios rather than absolute quantities, and as such are able to describe relative differences or similarities between flows, regardless of otherwise vast differences in single quantities like pressure or velocity.

$$Ma = \frac{U}{a} = \frac{U}{\sqrt{\gamma RT}} \tag{2.1}$$

$$Re = \frac{UL}{\nu} = \frac{\rho UL}{\mu} \tag{2.2}$$

$$Kn = \frac{\lambda}{L} = \frac{k_B T}{\sqrt{2}\pi\sigma^2 pL} = \frac{Ma}{Re}\sqrt{\frac{\gamma\pi}{2}} \tag{2.3}$$

$$St = \frac{h_{conv}}{\rho U C_P} \tag{2.4}$$

$$\lambda = \frac{k_B T}{\sqrt{2}\pi\sigma^2 p} \tag{2.5}$$

$$\gamma = \frac{C_P}{C_V} \tag{2.6}$$

The Mach number (Eq. (2.1)) represents a dimensionless velocity, and is expressed as the ratio of the flow's velocity vector and the local speed of sound i.e. the scalar speed of a pressure wave in the fluid being described. The Reynolds number (Eq. (2.2)) represents a force comparison, and is expressed as the ratio between the inviscid (inertial) and viscous forces acting upon the flow as a result of its motion and fluid properties. The Knudsen number (Eq. (2.3)) represents a comparison of scale and density, and is expressed as the ratio of a flow's mean free path, λ (Eq. (2.5)), (the mean distance that a particle in the flow must travel before interacting with another particle) and the characteristic length scale, L, of the analysis as a whole (note that, in the present work, the value of σ was set to the mean molecular diameter of air, $4 \times 10^{-10}\, m$ [35]). Finally, the Stanton number (Eq. (2.4)) represents an energy comparison, and is expressed as the ratio between the heat transferred into a fluid and its thermal capacity.

The quantity γ (Eq. (2.6)) is the ratio of the specific heats of a fluid (otherwise known as the Poisson constant or heat capacity ratio). Bird [33] notes that, for maximum accuracy when calculating Kn across an entire flowfield, or when performing a DSMC calculation, the *precise* value of the local length scale, L, may be determined by comparing macroscopic gradients (Eq. (2.7)). The viscosity, μ, may be calculated using a variety of equations and models. These include Sutherland's law [36] (a semi-empirical relationship which correlates temperature to viscosity) and Chapman-Enskog theory (a derivation of hydrodynamics transport properties based upon kinetic theory) [37]. If the viscosity of a gas mixture's components are known (or

indeed have been calculated using one of the aforementioned methods), then the overall viscosity of the mixture may be calculated using the Wilkie mixing rule [38]. This method is utilised by the Cantera chemical kinetics package used extensively in the present work [39] [40].

$$L = \frac{\rho}{\left(\frac{\mathrm{d}\rho}{\mathrm{d}x}\right)} \tag{2.7}$$

The four non-dimensional quantities listed above are all useful in determining the type of analysis techniques that one must utilise when assessing a flow, but none more so in the context of atmospheric entry than the Knudsen number. Depending on the value of this measure, an analyst may determine whether a flow can truly be classed as rarefied. If a flow has $Kn < 0.001$, then it is classed as a *continuum* flow, and traditional computational fluid dynamics (CFD) calculations using the Navier-Stokes equations may be performed with an acceptable error as the Knudsen number rises towards this threshold [41]. If a flow has $Kn > 10$, it is classed as a *free molecular* flow, where the probability of intermolecular collisions is negligible, and Newtonian collision behaviour manifests.

Flows with high Knudsen numbers require that a different simulation strategy be utilised for analysis, as one of the underlying assumptions of the Navier-Stokes equations (that of the fluid being a continuum, i.e. a continuous substance, as opposed to discreet particles) ceases to be true for sufficiently small mesh sizes or sparse flows. Two such simulation strategies are molecular dynamics (MD), and direct simulation Monte Carlo (DSMC). The latter of these is used extensively in the present work, and is described in depth in Section 2.3.4. The MD method, which is not utilised in the present work, was first proposed by Alder and Wainwright in 1959 [42], and models the fluid in motion as a continuous flow of individual particles, calculating quantities such as pressure and temperature based upon intermolecular and wall collisions.

DSMC operates in much the same fashion, but instead utilises a stochastic solution based upon the averaging of randomly seeded particle positions and velocities over large timescales. A weighting factor is utilised in order for a single DSMC particle to represent an arbitrary number of real particles, thereby greatly increasing the method's range of applicability. To summarise, MD methods simulate every particle in a sparse fluid flow, whereas DSMC simulates only a small fraction of these, instead relying on ensemble averaging and numerical weighting factors

to produce physically accurate results. Considering these ranges of applicability, it may be stated that CFD deals with the *macroscopic* scale, MD with the *microscopic* scale, and DSMC with the intermediate *mesoscopic* scale.

MD methods are able to yield extremely accurate numerical results and often feature impressively robust solvers, due mostly to their simplicity, but are hindered by computational limitations. Since each particle in a flow must be modelled both independently and simultaneously, flows that exist below the free molecular limit represent an extremely large computational hurdle, with particle counts rising quickly beyond the capabilities of modern hardware. DSMC simulations, as mentioned above, utilise a particle weighting factor, and as such may be utilised in the analysis of significantly denser flows. The primary disadvantage of DSMC is the requirement for long averaging times in order to mitigate stochastic noise in the results.

2.2.2 The slip and transition regimes

The *slip* and *transition* regimes are the names given to the classes of fluid flows whose Knudsen numbers (Eq. (2.3)) lie in the range $0.001 \leqslant Kn \leqslant 10$, that is, between the continuum and free molecular regimes.

The *slip* regime refers to the departure from continuum behaviour as Knudsen number becomes non-negligible ($0.001 \leqslant Kn < 0.1$), and where the so-called "no slip" boundary condition (i.e. zero fluid velocity at a wall) often used in continuum flow descriptions, ceases to hold true [43].

The *transition* regime refers to flows which have $0.1 \leqslant Kn < 10$ in this thesis (although it is used in some literature to refer to the range $0.001 \leqslant Kn < 10$). This regime is characterised by the large number of intermolecular collisions i.e. there are too many to allow for a reasonable MD solution, and too few for an accurate Navier-Stokes solution.

2.2.3 Mach number independence

The Mach independence principle is well described by Anderson [44], and states that aerothermal quantities converge to a steady, constant value as $Ma \to \infty$. This principle has been shown

to apply to all regimes of flow as described in Section 2.2.2 of the present work, and was demonstrated experimentally in the continuum regime by Grabow et al. [45] and Masson, Morris, and Bloxsom [46]; and numerically by Schaaf and Chambre [47], and Parham and Barba [48] in the free molecular regime. (Note that all four of the aforementioned studies were performed on spheres.)

In practice, this principle means that beyond a certain freestream Mach number (typically taken to be $Ma_\infty \approx 4 - 6$, depending on geometry [49]), non-dimensional quantities like aerodynamic coefficients and thermal flux trends change very little. Although this assumption is not universally true, it is often used by ground test facilities to design hypersonic experiments.

Generating a high Mach number flow generally requires a very large pressure differential between the stagnation reservoir and the facility's test section, which in turn either requires extensive vacuum equipment or high pressure gas supplies (and indeed sometimes both). As such, it is advantageous from an engineering standpoint to keep the target Mach number for a hypersonic facility as low as the planned experiments will allow.

It should be noted, however, that Mach independence does not apply to absolute quantities like heat flux, temperature, pressure, velocity, etc. If such quantities are to be studied directly, an experiment with the appropriately large stagnation pressure, temperature, and flow enthalpy would have to be designed.

2.2.4 Surface heat flux distribution

The dependence of heat flux distribution on mean free path (and, hence, Knudsen number) may be inferred from the kinetic theory of gases. In this representation of a fluid, molecules are modelled as spherical particles moving through a vacuum, and the overall thermodynamic properties of the fluid (such as temperature, Eq. (2.9), and pressure, Eq. (2.8)) are expressed as functions of the particles' mass, m, and mean molecular velocity, $\bar{c}$. The kinetic energy of the flow is expressed as the sum over all particles, and may be related to temperature and pressure via Eq. (2.10). Note that for simplicity, all particles in the following equations are assumed to have the same mass, m.

$$p = \frac{Nm\overline{c^2}}{3V} \tag{2.8}$$

$$T = \frac{pV}{Nk_B} = \frac{m\overline{c^2}}{3k_B} \tag{2.9}$$

$$E_k = \sum_{i=1}^{N} \frac{m_i c_i^2}{2} = \frac{Nm\overline{c^2}}{2} = \frac{N3k_BT}{2} = \frac{3pV}{2} \tag{2.10}$$

$$n = \frac{N}{V} = \frac{\rho N_A}{M} \tag{2.11}$$

A popular derivation of the mean free path, λ, posits a stream of moving particles impacting a number of stationary ones. In this fashion, the value of λ may be written as Eq. (2.12), where A_{cs} is the effective cross sectional area available for collisions, and n is the number density of particles (as expressed in Eq. (2.11)).

$$\lambda = \frac{1}{A_{cs}n} \tag{2.12}$$

If the particles in a system exist in an established equilibrium, however, their mean relative velocity in vector form may be expressed as in Eq. (2.13) (where **c** is the molecular velocity vector):

$$\overline{\mathbf{c}_r{}^2} = \overline{(\mathbf{c}_1{}^2 + \mathbf{c}_2{}^2)} = \overline{\mathbf{c}_1{}^2 + \mathbf{c}_2{}^2 - 2(\mathbf{c}_1 \cdot \mathbf{c}_2)} \tag{2.13}$$

Assuming that $\mathbf{c}_1$ and $\mathbf{c}_2$ have random values, and are unrelated to one another (also known as being uncorrelated), the mean dot product over all particles will be $\overline{\mathbf{c}_1 \cdot \mathbf{c}_2} = 0$. Also, $\mathbf{c}_1 \approx \mathbf{c}_2$, and thus the relative speed, c_r, becomes:

$$c_r = \sqrt{\overline{\mathbf{c}_r{}^2}} = \sqrt{\overline{\mathbf{c}_1{}^2 + \mathbf{c}_2{}^2}} = \sqrt{2}c \tag{2.14}$$

So the number of collisions in the gas is $\sqrt{2}$ times larger for particles in equilibrium as opposed to the aforementioned case (a moving stream with stationary target particles). Thus, the expression for mean free path, λ, in this case is;

$$\lambda = \frac{1}{\sqrt{2}A_{cs}n} \tag{2.15}$$

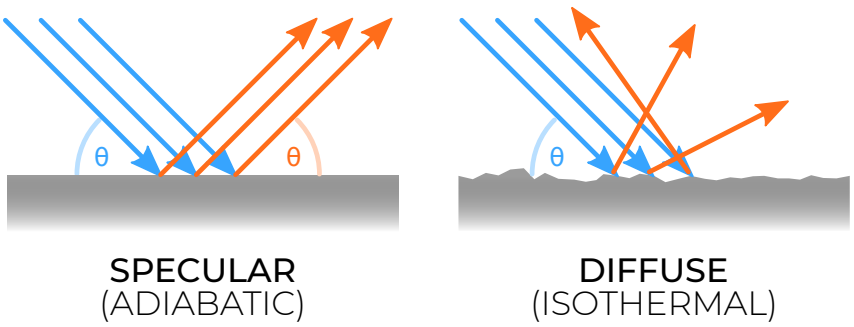


Figure 2.1: Diagram demonstrating the difference between specular and diffuse particulate reflections.

Substituting in the ideal gas law ($p = nk_BT$) and setting the cross sectional area available for collisions equal to that of two spheres of radius, r, and diameter, σ, ($A_{cs} = \pi 2r^2 = \pi 4/4\sigma^2 = \pi\sigma^2$) then yields the kinetic gas theory expression for mean free path (note that this expression is identical to that shown in Eq. (2.5)):

$$\lambda = \frac{k_BT}{\sqrt{2}\pi\sigma^2 p} \tag{2.16}$$

According to the aforementioned kinetic theory of gases, heat transfer occurs as the result of collisions between particles and solid surfaces present in the flow. The number density of the gas can be seen to be proportional to the number of particles per unit volume, and is also related to mass density (Eq. (2.11)). It follows from this that the probability of particles colliding with one another will decrease with n (resulting in a higher mean free path, λ).

Additionally, it must be noted that different types of particle-surface collisions result in different physical effects. These collision types are part of the Maxwellian gas-surface interaction model. The two distinct types of collision in this model are termed “specular” and “diffuse”. Specular reflections behave as light does when impacting a mirrored surface, i.e. reflecting at an angle equal to that of the initial particle impact relative to the surface. The physical effect of this is called an *adiabatic* boundary condition, wherein the net heat transfer to a surface is 0, but its temperature is affected by particulate collisions. Diffuse reflections behave in the opposite fashion: random scattering occurs, usually as the result of an uneven or rough surface. The physical effect of diffuse reflection is what is known as an *isothermal* boundary condition, wherein thermal energy may be transferred to the surface by incident particles, without altering the overall temperature of the surface.

Variables known as "accommodation coefficients" are used to express the ratio of specular reflections to diffuse reflections. The first of these is the normal accommodation coefficient, σ_N, which deals with pressure forces acting on a surface, i.e. the component of a collision force perpendicular to the surface. The second is the tangential accommodation coefficient, σ_T, which deals with the shear forces acting on a surface, i.e. the component of a collision force parallel to the surface. Both coefficients are expressed as numbers in the range $0 \leqslant \sigma \leqslant 1$, such that $\sigma_N = 1 - \sigma_T$ and $\sigma_T = 1 - \sigma_N$.

With these assertions, it may be concluded that a surface situated in a flow will encounter a number of collisions proportional to the number density of the particles, with the amount of heat being transferred per collision being related to each particle's velocity. Note that this is one of the fundamental reasons for heat transfer correlations being expressed as functions of freestream density and velocity (see Section 2.3.3). A similar determination could also be arrived at by using dimensional analysis via the Buckingham Pi theorem [50].

The heat flux trend, i.e. the shape of the distribution of heat flux on a surface, is heavily dependent on the number density of particles in the incident flow. The probability of a particle colliding with another in a given time step scales with n, meaning that sparser flows have fewer inter-particle collisions. As a result of this, the mean free paths of particles rebounding after collisions with a solid surface become higher as n decreases, and their momentum is far less likely to be affected by other particles in the flow. Hence, flow structures (which form as a result of momentum being distributed through a fluid flow) are typically more diffuse. Similarly, heat transfer trends tend towards Newtonian-style uniformity over flat surfaces i.e. they develop a strong dependence on surface inclination angle (this has been observed by many researchers experimentally, including Masson [46]).

The distributions of heat flux that come about as a result of this phenomenon scale not just with the mean free path length of the flow, but also with the relative size of the object on which the flow is impinging. For example, a flat plate in a high Mach number flow would encounter a rise in heat flux magnitude towards its edges, and a region of near constant heat flux towards it centre (thus resembling a power law). If a smaller plate were inserted into the same flow (thereby increasing the Knudsen number by reducing characteristic length), the trend of heat flux across

its surface would retain the same power law shape, but would have a far smaller region of near constant heat flux in the centre, and higher magnitude peaks towards its edges (i.e a different distribution and peak magnitude). In order to non-dimensionalise this relationship between the flow and an object on which it is incident, both the mean free path of the freestream and a characteristic dimension of the object must be considered. This expression takes the simple form λ/L, which is the equation for Knudsen number shown in Eq. (2.3).

These effects have been observed in the works of Klett [51], and Schaaf and Talbot [47] wherein analytical studies of the aerodynamic forces and heating on simple geometries were performed.

2.2.5 Thermal soak and peak heating

The concept of thermal soak is simply that of preheating through an object via conduction. During atmospheric reentry, this phenomenon is extremely important, as one of the driving factors of heat transfer is the difference in temperature between the an object and its surroundings [52]. This relationship is evident from the basic expressions for the three modes of heat transfer: conduction (Eq. (2.17)), convection (Eq. (2.18)), and radiation (Eq. (2.19)), respectively.

$$\dot{q}_{cond} = -k\frac{dT}{dx} \tag{2.17}$$

$$\dot{q}_{conv} = h_{conv}\left(T_w - T_\infty\right) \tag{2.18}$$

$$\dot{q}_{rad} = \epsilon\sigma\left({T_w}^4 - {T_\infty}^4\right) \tag{2.19}$$

As with any engineering level calculation, the peak loads encountered by the subject of an analysis are of paramount importance. In the case of atmospheric re-entry, these loads are the peak deceleration and heating encountered during an object's transit through a planetary atmosphere (regardless of whether one is designing a spacecraft, or analysing a piece of debris). Additionally, the total heat load encountered over an entire entry trajectory is also of interest to designers, particularly in the case of debris demise calculations (such as those preformed by the NASA *DAS* code).

It is well established that the region of peak heating occurs in the lower parts of a re-entry

trajectory. Depending on an object's area-to-mass-ratio, $\frac{A}{m}$, this may occur in the continuum regime ($Kn < 0.001$) for small values of $\frac{A}{m}$, or the lower slip regime ($0.001 \leqslant Kn < 0.1$) for high values of $\frac{A}{m}$. As a result of this relative proximity to the continuum regime, many numerical studies utilise continuum heating formulations for their thermal calculations during peak heating.

Analyses demonstrating continuum peak heating on crew capsule geometries have been performed by Viviani and Pezzella [53], and are also included as part of standard simulation procedures presented during a course on capsule aerothermodynamics at VKI, the von Karman Institute for Fluid Dynamics (distributed as AGARD-R-808 [54]).

Lu, Kim, and McNay [55] have performed lower fidelity calculations for a small payload recovery concept which concisely demonstrate the dependence of peak heating regime on $\frac{A}{m}$. Baumann, Pagan, and Herdrich [56] have also performed low fidelity calculations for a small sphere-cone re-entry demonstrator whose geometry is similar to that of capsules, but with a far larger $\frac{A}{m}$. As such, their calculations showed that its peak heating region would occur in the lower slip regime.

In the case of time-variant thermal calculations, a full analytic solution is only possible with analytic models of the thermal environment and the thermal behaviour of the object being analysed e.g. steady cooling between layers of an insulated solid. Since a typical atmospheric re-entry calculation must take variations in thermal environment with altitude ($p_\infty, T_\infty = f(h)$), and object attitude into account, analytic solutions may only rarely be formulated for this particular class of calculation. Such solutions have previously been formulated by Allen and Eggers [57] for ballistic trajectories, with enhancements made by Putnam and Braun [58].

In general, heating which manifests prior to the peak heat load during an atmospheric entry is termed the 'thermal pre-soak', and affects the wall temperature of an object from very early on in its transit through an atmosphere. Taking the dependence of heat transfer on temperature difference (demonstrated in Eqs. (2.17), (2.18), and (2.19)), and the late onset of the peak heating region into account, it becomes apparent that knowledge of the thermal environment at every step of a re-entry trajectory is required for peak heating calculations to be accurate.

As such, in order to take this phenomenon fully into account, analysts are required to model the evolution of temperature from the initial entry of an object into the upper atmosphere to

touchdown (or, indeed, to the point that it is destroyed by aerothermal stresses).

2.3 NUMERICAL ANALYSIS OF HYPERSONIC AEROTHERMODYNAMICS

In the study of rarefied hypersonic flows, many numerical techniques have been developed in order to calculate the aerothermodynamics of re-entering objects. These vary widely in both complexity and accuracy, and are often grouped by the latter of these metrics into what are termed "low fidelity" and "high fidelity" methods. Here, "low fidelity" refers to numerical methods which do not attempt to fully model the physics of a fluid flow, instead relying on empirical correlations or reduced order governing equations in order to provide useful data. Conversely, "high fidelity" refers to those numerical methods which attempt model fluid flows as completely (and hence, as accurately) as is feasible. For the present work, panel methods (low fidelity) and Direct Simulation Monte Carlo simulations (high fidelity) are both used to assess the aerothermodynamics of re-entering space debris.

In this section, hypersonic panel methods are introduced (these are elaborated upon in Chapter 3), followed by a discussion of bridging functions (which allow aerothermodynamic quantities in the transition regime to be easily estimated), and stagnation point heat flux correlations (which facilitate the calculation of thermal fluxes incident at a re-entering object's stagnation point). These comprise the low-fidelity methods used in the present work.

Finally, the high fidelity DSMC (Direct Simulation Monte Carlo) method is described, followed by those software implementations which were evaluated for use during the present work. This section then concludes with descriptions of previous simulations of re-entry flows that have been performed using the DSMC method.

2.3.1 Panel methods

Panel methods have been used for low fidelity aerodynamic analyses since the late 1960s, and were originally developed before computer technology was advanced enough to support solutions to large finite element or finite volume problems, such as those presented by the Navier-Stokes and Boltzmann equations, to name but two.

The equations and algorithms used in these methods have changed very little since their

initial implementation by the Douglas Aircraft company in 1964. Development of this program, which was originally known as the "Douglas Arbitrary-Body Aerodynamic Computer Program", was continued in earnest by the US Air Force Research Lab (AFRL) at Wright-Patterson Air Force Base, where the code was renamed to *S/HABP* (Supersonic/Hypersonic Arbitrary Body Program). This effort culminated in the Mark IV version in 1973 [59] with the Mark V version following some years later.

Although panel methods have been superseded in terms of accuracy by high fidelity finite element codes (which were themselves made possible by the advent of high performance digital computers), their popularity amongst engineers and academics has far from diminished. With the aforementioned technological improvements, panel methods have become extremely computationally efficient, allowing them to be used extensively in the early design stages of missiles [60], rockets [61], high speed aircraft [59], re-entry capsules [62], single-stage-to-orbit vehicles [63] [64], and even in the prediction of the aerothermal behaviour of space debris [22] [14] [16] and meteorites [65].

Typically, panel methods rely on the discretisation of a geometry (a surface model of the object or vehicle being analysed) into discrete facets, known as panels. The angle between these panels and the freestream flow vector is then calculated, and any shielding effects between individual panels are taken into account e.g. if a solar panel or canard obscures a section of a craft's fuselage. The calculated inclination angle of the panels are then used as inputs to a series of simple, highly computationally efficient equations which assess the local pressure, shear, and heat transfer coefficients. This dependence on panel inclination leads us to the alternate name for panel methods: "local surface inclination methods".

Finally, these coefficient distributions are processed using additional calculations in order to characterise the aerothermodynamic behaviour of the vehicle or object. Thermal distributions are usually calculated using a freestream heat flux correlation (see Section 2.3.3) by which the heat transfer coefficient distributions are multiplied, while pressure and shear forces are summed over the entire object, and the six aerodynamic coefficients derived.

2.3.2 Bridging functions

Flows in the slip and transition regimes have historically proven very challenging to analyse or categorise numerically, with empirical functions based on experimental data often being employed to bridge the gap between those Knudsen numbers that are considered "high" (free molecular) and "low" (continuum). This difficulty arises as a result of the fluid densities often encountered in these regimes i.e. they are often too sparse for CFD to yield suitably accurate results, while MD methods require particle counts so high as to make the simulations too computationally expensive to run satisfactorily.

These empirical curves (known as "bridging functions") are often used in order to interpolate values between analyses which are performed on either side of the regime gap. There are two types of these functions: "local", wherein surface skin friction and pressure coefficients (C_f and C_p, respectively) are calculated and then integrated to compute the total aerodynamic coefficients, and "global" wherein the aforementioned total aerodynamic coefficients are computed first with reference values calculated in the continuum and free molecular regimes (either side of the slip regime bridge) (Eq. (2.21)) [66]. A bridging function is then used to interpolate between these bounding values with respect to the freestream Knudsen number.

The curves are most commonly generated using experimental or flight data. Owen [67] has previously collated a number of useful global type functions for drag coefficient, while suggesting his own correlation based upon experimental measurements of cone drag (Eq. (2.20)). Similarly, Morsa et al. [66] have also assembled a number of separate correlations (although in this case both global and local functions are presented) and have sought to validate and augment them (where necessary) using DSMC results computed with Bird's *DS2V* code. Included in these equations is Eq. (2.21) the sine-squared law developed by Wilmoth, Blanchard, and Moss [68], which is based upon reconstructed flight data and DSMC simulations of a number of vehicles with spherically blunted heat shields and conic afterbodies (these include Viking, Pathfinder, Microprobe, Mars '01, and Stardust). The variables a_1 and a_2 in this expression allow the upper and lower bounds of Knudsen number (where the continuum and free molecular regimes begin) to be altered, although the most common values are $a_1 = 0.375$ and $a_2 = 0.125$, which correspond to Kn limits of 0.001 and 10, respectively.

In addition to aerodynamic coefficient correlations, there are also bridging functions for such quantities as pressure coefficient, C_p, and skin friction coefficient, C_f, which Morsa et al. [66] have also collated from various authors in order to validate them using Bird's *DS2V* code. Additionally, Gelhaus utilised Stanton number, St, and heat flux bridging functions in his work on the ESA DRAMA suite, which are formulated for primitive three dimensional geometries and are based upon the best fit solutions of Detra, Kemp and Riddell [6]. Of these functions, that which deals with Stanton number bridging may be applied to spheres and cylinders (Eq. (2.22)), while the heat flux bridging function may be applied to cuboids and flat plates (Eq. (2.23)). The bridging formulae of Legge (as presented in [22]) are also frequently utilised, and may be expressed in terms of Stanton number, St, as in Eq. (2.24) or heat flux, $\dot{q}$, as in Eq. (2.25).

$$\frac{(C_D - C_{D_{cont}})}{(C_{D_{fm}} - C_{D_{cont}})} = \text{erf}\left(1.5\sqrt{Kn}\right) = 1 - e^{-3\,Kn^{0.6}} \tag{2.20}$$

$$C_{trans} = C_{cont} + \left\{\left(C_{fm} - C_{cont}\right)\sin^2\left(\pi\left[a_1 + a_2\log_{10}(Kn)\right]\right)\right\} \tag{2.21}$$

$$St_{trans} = \frac{St_{cont}}{1 + \frac{St_{cont}}{St_{fm}}} \tag{2.22}$$

$$\dot{q}_{trans} = \dot{q}_{cont}\left[1 - e^{-\frac{\dot{q}_{fm}}{\dot{q}_{cont}}}\right] \tag{2.23}$$

$$St_{trans} = \frac{St_{cont}}{\sqrt{1 + \left(\frac{St_{cont}}{St_{fm}}\right)^2}} \tag{2.24}$$

$$\dot{q}_{trans} = \frac{\dot{q}_{cont}}{\sqrt{1 + \left(\frac{\dot{q}_{cont}}{\dot{q}_{fm}}\right)^2}} \tag{2.25}$$

2.3.3 Stagnation point heat flux

With regard to heat flux, a common strategy in the past has been to calculate the heat transfer rate at a geometry's stagnation point. This value may then be utilised as a worst case assessment or, in the case of a somewhat higher fidelity analysis, a reference value by which other key quantities or temperature distributions may be estimated. This approach was pioneered by Fay and Riddell [69] who studied the flowfield around a sphere by formulating from first principles a

streamline whose origin was coincident with said sphere's stagnation point. Their expression, known commonly as the Fay-Riddell equation, is presented as Eq. (2.26). Note that the variables h_D, h_0, and h_w in this expression are, respectively, the enthalpy of formation, stagnation enthalpy, and flow enthalpy at wall conditions.

By adjusting the input parameters of this equation, Fay and Riddell were able to estimate the convective heat transfer at the sphere's stagnation point. This value could be adjusted for a cylinder through the modification of an empirical factor, A (Eq. (2.27)), and for differing boundary layer and wall conditions via the expression B (Eq. (2.28)). These conditions comprise an equilibrium boundary layer, a fully catalytic wall with a chemically frozen boundary layer, and a non-catalytic wall with a chemically frozen boundary layer, respectively. Note that in these expressions, the variable α is the thermal diffusivity of the gas, the equation for which is given as Eq. (2.32). Additionally, the non-dimensional numbers, Pr, Sc, and Le are the Prandtl, Schmidt, and Lewis numbers, respectively. The Prandtl number represents the ratio of momentum diffusivity to thermal diffusivity, the Schmidt number is the ratio of viscous diffusivity to mass diffusivity, and the Lewis number is the ratio of thermal diffusivity to mass diffusivity. Equation (2.33) represents the velocity gradient between the shock and wall. Note that the variable R_n is the radius of the sphere or cylinder being modelled in the calculation. In the event of a non-spherical, non-cylindrical geometry, R_n represents the radius of a sphere with the equivalent surface area.

$$\dot{q}_{0_{conv}} = A\, Pr^{-0.6}(\rho_e \mu_e)^{0.4}(\rho_w \mu_w)^{0.1}\sqrt{\left(\frac{dU_e}{ds}\right)_0}(h_0 - h_w)B \tag{2.26}$$

$$A = \begin{cases} 0.76 & \text{sphere} \\ 0.53 & \text{cylinder in cross-flow} \end{cases} \tag{2.27}$$

$$B = \begin{cases} 1 + (Le^{0.52} - 1)\left(\frac{h_D}{h_0}\right) & \text{equilibrium boundary layer} \\ 1 + (Le^{0.63} - 1)\left(\frac{h_D}{h_0}\right) & \text{frozen boundary layer, fully catalytic wall} \\ 1 - \left(\frac{h_D}{h_0}\right) & \text{frozen boundary layer, non-catalytic wall} \end{cases} \tag{2.28}$$

$$Pr = \frac{\nu}{\alpha} = \frac{C_P\, \mu}{k} \tag{2.29}$$

$$Sc = \frac{\nu}{D} = \frac{\mu}{\rho D} \tag{2.30}$$

$$Le = \frac{\alpha}{D} = \frac{Sc}{Pr} \tag{2.31}$$

$$\alpha = \frac{k}{\rho C_P} \tag{2.32}$$

$$\frac{dU_e}{ds} = \frac{1}{R_n}\sqrt{\frac{2(p_0 - p_\infty)}{\rho_0}} \tag{2.33}$$

Further correlations were computed in later studies, with an emphasis on the empirical determination of stagnation heat flux for re-entry vehicles. Such expressions, like those of Detra and Hidalgo (Eq. (2.34)) [70], Brandis and Johnston (Eq. (2.35), Eq. (2.36)) [71], or Detra, Kemp and Riddell [6] are far more compact and simple to solve than the Fay-Riddell expression. Additionally, several similar expressions for the heat flux contributions of radiation have been computed, the most recent and accurate of which was also devised by Brandis and Johnston [71].

Detra and Hidalgo state that, for continuum conditions:

$$\dot{q}_{0_{cont}} = 5.16 \times 10^{-5}\sqrt{\frac{\rho}{R_n}U^{3.15}} \tag{2.34}$$

For $3km/s \leqslant U \leqslant 9.5km/s$ Brandis and Johnston yields:

$$\dot{q}_{0_{conv}} = 4.502 \times 10^{-9}\rho^{0.4704}U^{3.147}{R_n}^{-0.5038} \tag{2.35}$$

For $9.5km/s \leqslant U \leqslant 17km/s$ Brandis and Johnston yields:

$$\dot{q}_{0_{conv}} = 1.270 \times 10^{-6}\rho^{0.4678}U^{2.524}{R_n}^{-0.5397} \tag{2.36}$$

The solution of Detra, Kemp and Riddell yields two solutions. One for continuum flow (Eq. (2.37)), and another for the free molecular regime (Eq. (2.38)). In the latter of these expressions, the variable α_T (which varies in the interval $0 \leqslant \alpha \leqslant 1$) is the thermal accommodation coefficient, which allows the behaviour of gas molecules during collisions with a solid body to be modelled (i.e. such that energy transfers between molecules and surfaces are not modelled as being perfectly efficient).

$$\dot{q}_{0_{conv}} = \frac{11028.5}{\sqrt{R_n}}\left(\frac{\rho_\infty}{1.225}\right)^{0.5}\left(\frac{U_\infty}{7802.88}\right)^{3.15} \tag{2.37}$$

$$\dot{q}_{0_{fm}} = \frac{\alpha_T \rho_\infty U_\infty^3}{2} \tag{2.38}$$

2.3.4 The DSMC method

Direct Simulation Monte Carlo, or DSMC, is a numerical analysis methodology developed by Prof. Graeme Bird, Emeritus Professor of Aeronautics at the University of Sydney [41]. It was originally created in order to bridge the gap in analysis techniques between the continuum and free molecular limits by addressing the computational limitations of pure MD methods, and has since found great favour with space agencies and research institutions as an assessment tool for rarefied aerodynamic studies and microfluids simulations.

DSMC operates in a similar fashion to pure MD solvers, but with only a fraction of the actual number of particles being simulated. A weighting factor is incorporated, allowing a single DSMC particle to represent an arbitrary number of actual fluid particles in the simulation, and averaging of stochastic variables over time is used to arrive at a final solution. This allows the calculation of intermolecular collisions to be simplified, at the same time as retaining the magnitudes and vital features of the flow, such as standing shock structures, body force distributions and heat flux patterns.

Simulations may be initialised using a Boltzmann distribution of pseudorandom velocity vectors, or a unified velocity vector (which represents a freestream). Collision and sampling cells are also defined such that collision calculations may be further streamlined and poor computational scaling may be avoided. (Without such discretistion of the computational domain, a DSMC simulation would scale in a similar fashion to the notorious "Travelling Salesman" problem in computer science [72] i.e. *every* particle would have to be considered for collision with every *other* particle for *every* timestep, an enormous and entirely unnecessary number of calculations.)

For each time step, the simulated particles are translated in space according to their velocity vectors, whereby new positions, and whatever collisions occur as a result, are calculated. In this fashion, the Boltzmann equation (Eq. (2.39), as presented in the work of White et al. [73]) is solved numerically, and energy transfers between elements of the computational domain are

found. This equation describes the state of a thermodynamic system in terms of probability, i.e. how likely it is that a given particle would be at a certain position in space, travelling at a certain finite velocity

$$\frac{\partial\left(n_i f_i\right)}{\partial t}+\mathbf{c}_i \cdot \frac{\partial\left(n_i f_i\right)}{\partial \mathbf{r}_i}+\mathbf{F} \cdot \frac{\partial\left(n_i f_i\right)}{\partial \mathbf{c}_i}=\sum_{j=1}^{N} \int_{-\infty}^{\infty} \int_{0}^{4\pi} n_i n_j\left(f_i^* f_j^*-f_i f_j\right)\left|c_{r_{ij}}\right| \sigma_{ij}\, d\Omega\, d\mathbf{c}_j \quad (2.39)$$

Here, n is the number density of the flow, f is the normalised velocity distribution function, **r** is the particle position vector, **c** is the particle velocity vector, Ω is the solid collision angle, σ is the molecular cross-section, and **F** is any external force which acts on the particles (such as gravity, or Lorentz force).

Quantities such as pressure, temperature, and velocity, as well as surface integrals like force distribution and heat flux are averaged over all time steps. The simulation is deemed "complete" when a steady state has been achieved (wherein the rate of change of variables such as mean kinetic energy or particle count approach zero), following which statistical time averaging of the stochastic solution is performed (except where unstable flow is expected at large values of t, in which case ensemble sampling is utilised [43]).

Bird [43] has put forward several rules of thumb by which a DSMC analysis may be judged in order to assess its accuracy. These are:

Time step suitability

The time step of the simulation should be smaller than the smallest value of mean free time, $\overline{t}$, calculated in the simulated flowfield, which is computed using Eq. (2.40) (where $\overline{c}$ is the mean particle velocity). Ideally, the time step should be smaller than $\overline{t}/3$.

$$\overline{t}=\frac{\overline{c}}{\lambda} \quad (2.40)$$

Collision cell scale

The largest individual collision cell dimension in a particular section of the computational domain should be less than the fluid mean free path length, λ, at that point (ideally, the largest collision cell should be smaller than $\lambda/3$).

Mean collision separation

The ratio of the mean collision separation between DSMC particles, mcs, to the physical mean free path of the flow being simulated, λ, should be less than unity (ideally, the mean collision separation should be less than $\lambda/5$).

Particle count

The number of DSMC particles present in the simulation must be sufficient such that there is a mean of at least 7 particles per collision cell in all areas of interest. Note that this advice is valid only for Bird's codes *DS2V* and *DS3V* – the ideal value of this parameter will depend upon the formulation of a given code. For example, the optimal particle count per collision cell for *dsmcFoam+* is between 20 and 30 [74]). This parameter is often adjusted by altering a code's particle weighting factor, which determines how many real gas particles are simulated by a single DSMC particle. The relationship between the real and simulated number of particles is described by Eq. (2.41), where w_{DSMC} is the particle weighting factor, N is the number of real gas particles in the flowfield, and N_{DSMC} is the number of simulated gas particles.

$$w_{DSMC} = \frac{N}{N_{DSMC}} \tag{2.41}$$

It should be noted that the criteria listed above are the theoretical minimum requirements for a "good", i.e. physically accurate, DSMC simulation. In performing actual simulations, the specific code an analyst uses must be sufficiently understood in order to properly set the simulation parameters and hence gain accurate results.

2.3.5 DSMC implementations

There are several mature implementations of the DSMC method which are available for use as analysis tools. Of these, only a small number have been utilised for the current body of work (owing mostly to the cost of licences and ITAR regulations which restrict the availability of some codes). The programs selected for the analyses pertaining to this body of work are: *dsmcFoam+* [75], *DS2V* [76], and *DS3V* [76]. These codes were selected for the present work owing to their

extensive use in literature, and the relative the ease of obtaining them.

DS2V / DS3V

Windows executables which were written, and are actively maintained by Professor Bird himself, and together represent a long period of continual development. As is alluded to in their naming convention, the former is designed to analyse 2D flows, while the latter is configured for 3D scenarios. *DS2V* is an extremely useful code in that it is self contained and entirely automatic i.e. it features a GUI, a geometry definition tool, an automatic meshing algorithm, a built-in post processor, and automatic calculation and adaption of sample cell size, collision cell size, time step, and particle count. *DS3V* also features many of these functions, but is hampered by a somewhat convoluted geometry definition system wherein an operator must specify the properties of each individual surface element of a geometry prior to executing an analysis. Additionally, neither code is programmed to take advantage of multi-core CPUs, and their computational times suffer significantly as a result. As such, *DS2V* is regarded by the author as an indispensable tool for 'initial' analyses which may be used as baselines for more complex simulations, providing the relevant data for mesh refinement, time step selection and particle count.

dsmcFoam+

An open source DSMC code developed by the James Weir Fluids Laboratory at the University of Strathclyde [73], which is built using the *OpenFOAM* C++ library. This code, which was developed 'from scratch', as it were, was written with parallel computation and ease of use in mind, and as such allows for scalable parallel decomposition and processing. It also makes use of the many utilities already built into *OpenFOAM* for functions such as mesh generation, data management, and post-processing. Data is input into the application via C++ formatted text files (known as "dicts", or dictionaries) which are actively monitored during the simulation, allowing for the adjustment of control options during runtime. Recent updates to the original release (which was simply titled *dsmcFoam* [77]) have included the implementation of the quantum-kinetic (QK) chemistry model [78], which, under certain circumstances wherein high energy chemical reactions become significant, can significantly improve the accuracy of simulation results. Another important

improvement is the inclusion of calculations for simulation quality metrics, such as those described in Section 2.3.4. A notable limitation of current versions of *dsmcFoam+* is the lack of an automatic mesh refinement feature. As such, it is left up to the user to judge the quality of a simulation domain with the supplied tools, and hence make the appropriate refinements themselves.

2.3.6 DSMC analyses of re-entry vehicles

There are numerous examples of previously published DSMC analyses that are useful in the construction of new simulations. With regard to the analysis of planetary entry scenarios, the works of Shang and Chen [79], and Moss et al. [80] [81] are of particular interest. Shang and Chen's work involves a basic simulation of the external flow around a CEV (Crew Exploration Vehicle) geometry and includes a good example of what can be accomplished with the *dsmcFoam* solver. The works of Moss et al. involve the comparison of actual flight data and DSMC simulations for the Mars Pathfinder, in which the generation of the aerodynamics database for the EDL (entry, descent, and landing) phase of the mission is described [80]. In the second of Moss et al.'s works which is cited here, the simulation of Mars entry geometries is outlined in a similar fashion. The works involve extensive examinations of the in-flight force coefficients experienced by the spacecraft (and an aerothermal assessment of the geometry at zero incidence).

The simulation of basic shapes has also received a great deal of attention, but mostly with respect to in-flight force coefficients (most recently in reference to the generation of generalised dynamics models for re-entry simulations). One of the earliest of these is the work of Vogenitz et al. [82] wherein spheres and cones are analysed using an early implementation of the DSMC method. Somewhat more relevant to the current work are the analyses performed by Walker [83] and Cheng et al. [84]. Walker's work involves highly detailed force coefficient correlations across many different varieties of three dimensional primitives, and was conducted in order to improve the prediction of satellite re-entry trajectories. The work of Cheng et al. [84] deals with a cuboid shaped geometry, analysed over a range of Knudsen numbers. Flowfield surveys are performed and mean heat flux coefficients are calculated across separate sides of the model.

More recently, Falchi et al. [17] [18] [16] have performed a number of DSMC simulations

on simplistic debris analogues (mainly cuboids and cylinders) in order to develop surrogate corrective models for their FOSTRAD code (see Section 2.1). These simulations were designed to provide determinations of the skin friction and heat transfer coefficients at transition regime Knudsen numbers, and were shown to provide excellent agreement to experimental and flight data. Particular attention was paid in these studies to the objects' angles of incidence, and to variations near to sharp corners. As such, these studies are of particular interest to the present work.

2.4 INTRODUCTION TO HYPERSONIC WIND TUNNELS

In this section, the type of experimental facilities used for both the present work and similar research are elaborated upon. Additionally, examples of this type of facility are listed and described.

Hypersonic wind tunnels are fluid mechanics test facilities that produce Mach numbers beyond the hypersonic limit (see Section 2.2.3). This class of flows are utilised for the study of high speed aerothermodynamics, often associated with traversal of a planetary atmosphere during re-entry. Other applications include the study of radiation heat transfer induced by the formation of plasmas in shock structures, high temperature chemical reactions, scramjet design, and shock-shock interaction, to name but a few.

Hypersonic facilities may be split into two distinct categories: high enthalpy, and low enthalpy. The former of these are capable of reproducing a hypersonic flow as it would occur naturally, both in terms of thermodyanmic and aerodynamic properties. The latter, however, are only capable of reproducing the aerodynamic characteristics of the flow being simulated.

A number of types of experimental facility are used to model this type of high enthalpy flow, a sub-category of which are typically referred to as "shock tunnels". These operate in different fashions, depending on their designs, but generally involve releasing a mass of heated, pressurised gas into an evacuated test section, thereby generating extremely high enthalpies and Mach numbers. Due to the engineering challenges of containing such energies, the actual experimental duration is usually measured in milliseconds [85] [86]. These short test times also allow thermal fatigue of experimental equipment to be avoided, since the high temperatures

associated with hypersonic flows only manifest for a short amount of time in such facilities.

Low enthalpy facilities (which often operate continuously in a steady state) also exist which allow tests of long durations to be run. These generally focus on aerodynamic studies, although thermal tests utilising appropriate scaling and normalisation factors may also be performed in them [87].

Regardless of which category they fall into, facilities of this type operate using the same basic procedure, i.e. introducing a gas into a sealed test chamber with a total-to-static pressure ratio greater than ~ 530 (calculated using the isentropic nozzle equation with $Ma = 5$ and $\gamma = 1.4$).

2.4.1 Existing hypersonic rarefied facilities

There are many hypersonic wind tunnels currently operating around the world. Table 2.5 shows a summary of low enthalpy facilities which are capable of reproducing the Knudsen numbers required for the studies described herein.

Facility	**Location**	$\max(Ma)$	$\max(Kn)$	$Re\ \left[m^{-1}\right]$	$\max(T_0)\ [K]$
LDWT	OTI, Uni. of Oxford, UK	10	0.4	$200 - 500$	523
HRWT	JAXA, Tokyo, Japan	13	0.3	< 200	800
MARHy	ICARE, Orléans, France	20	0.36	$0.263 - 7522$	1206
H2K	DLR, Cologne, Germany	11.2	3.1×10^{-6}	$[2 - 21] \times 10^6$	1100

Table 2.5: Details of assorted low enthalpy hypersonic wind tunnels

The MARHy (Mach Adaptable Rarefied Hypersonic) [88] facility features a test chamber of comparable size to that of the LDWT ($\varnothing 2\,m \times 3.5\,m$), and an extensive pumping system comprising 14 Rootes type pumps connected in parallel, backed by a pair of rotary vacuum pumps. This arrangement allows extremely high mass flow rates and Mach numbers for this type of facility.

The HRWT (Hypersonic Rarefied Wind Tunnel) [89] facility, in contrast, is a relatively compact system, featuring a test chamber measuring $\varnothing 1.0\,m \times 1\,m$, and a pumping system comprising a dry vacuum backing pump ($Q = 100\ l/s$) and three turbomolecular pumps ($Q = 3300\ l/s$).

The H2K tunnel [90], unlike those previously mentioned, is a "blow-down" facility which operates by evacuating the test section using a vacuum sphere. A pressurised air supply is then vented through a set of $5\,MW$ electric heaters (which raise the gas temperature to $1100\,K$) before arriving at the test chamber via a nozzle, wherein hypersonic test conditions are produced. The tunnel's operators estimate that between 8 and 10 tests may be performed each day.

The LDWT – described in detail in Chapter 4 – fulfils a role between these three facilities, providing a large test chamber capable of accommodating bulky experimental equipment, whilst featuring a compact pumping system (compared to MARHy) capable of maintaining a high mass flow rate during operation (compared to HRWT). It is also capable of achieving a significantly longer test duration than H2K. This "hybrid" capability is achieved by the unique inclusion of a vapour booster pump in the LDWT's backing system, making it the only civilian operated facility of its kind in the world capable of operating in the aforementioned niche.

2.5 EXPERIMENTAL HEAT FLUX MEASUREMENT

The experimental measurement of heat fluxes across models of re-entry bodies is often one of the most important issues in hypersonic tests, and there are several methods by which this data may be gathered. Those that are most relevant to the current work are infrared thermography, visual thermography, and thin film gauge measurements. These methods are described below.

2.5.1 Infrared thermography

This method involves the use of cameras sensitive to the infrared portion of the electromagnetic (EM) spectrum. This type of EM radiation is characterised by wavelengths far larger than those of visible light ($7 - 14\ \mu m$ as opposed to the $380 - 700\ nm$ of the visible spectrum). Infrared cameras have the advantage of being able to measure the temperature of an object directly without any modification to said object, making them an ideal choice for non-intrusive measurements.

Infrared thermography is described by de Luca et al. [91] and Le Sant et al. [92] wherein it is shown to be a viable alternative to thermochromatic techniques for large wind tunnel facilities. The hypersonic experiments of Henckels and Maurer [93] and Asma and Van der Haegen [94] highlight the usefulness of the technique in the context of high speed facilities. These two works

also draw attention to how the method has matured over the two decades between the two studies: from an evaluation of its potential usefulness in 1989 to full confidence in the results it produces for a heavily funded ESA mission by 2010.

2.5.2 Thin film gauges

Thin film gauge measurements involve the use of surface-mounted temperature measurement devices whose electrical resistance changes as a function of their temperature. Gauges of this type have previously been used extensively at OTI (Oxford Thermofluids Institute), having been described in detail by Schultz and Jones [95] in 1973. The most recent method for their manufacture was described by Collins et al. [96], whereby high precision photo-resist layers are used to etch flexible circuitry onto sheets of copper-coated polyamide film. These are then combined with a second photo-resist layer and placed in a sputter deposition chamber where platinum is used to form the gauges themselves.

By applying a constant current to such devices, their temperature history may be derived by measuring the change in voltage over time, and hence deriving the variation in resistance via Ohm's law. Heat fluxes at the gauge location may then be derived via knowledge of this transient, as well as the material properties of the underlying substrate. A similar method involves the use of "double-sided" gauges, wherein two standard gauges are sputtered onto opposing sides of a thin insulating layer (usually a polyamide film, such as Kapton, with a layer of copper affixed to either side). In this case, the film is regarded as the substrate, and the temperatures of the two gauges are measured simultaneously during an experiment. With knowledge of the film properties, it is then possible to derive the thermal fluxes at each gauge location, without any knowledge of the thermal behaviour of the underlying substrate (i.e. the material onto which the film is fixed for a given experiment). Oldfield [97] has shown that determinations of heat flux from such gauges (when properly manufactured and calibrated) are extremely accurate, as well as being computationally efficient.

2.5.3 Transient liquid crystal thermography

This technique involves the application of a thermally thin layer of thermochromic liquid crystals (LC) to a substrate with known thermal properties. The chemical composition of the liquid crystals allows them to remain transparent both below and above a certain temperature, which is determined during manufacture. As such, temperature changes may be tracked visually by recording the colour play (the variation in colour of the liquid crystals) over the surface of a model.

This method has been used extensively in previous studies at the Osney lab, including those performed by Ryley [4], and Owen [87]. The methodology of utilising liquid crystals as a thermographic technique is described by Schultz and Jones [95], and later in much greater detail by Ireland and Jones [98], wherein their applicability to thermal flux experiments, and strategies for post processing the resultant data, are described. Additional studies, conducted using temperature sensitive paints by Liu [99], [100] reinforce the conclusion that high fidelity data may be gathered using these thermochromatic coating techniques over a wide range of flow conditions.

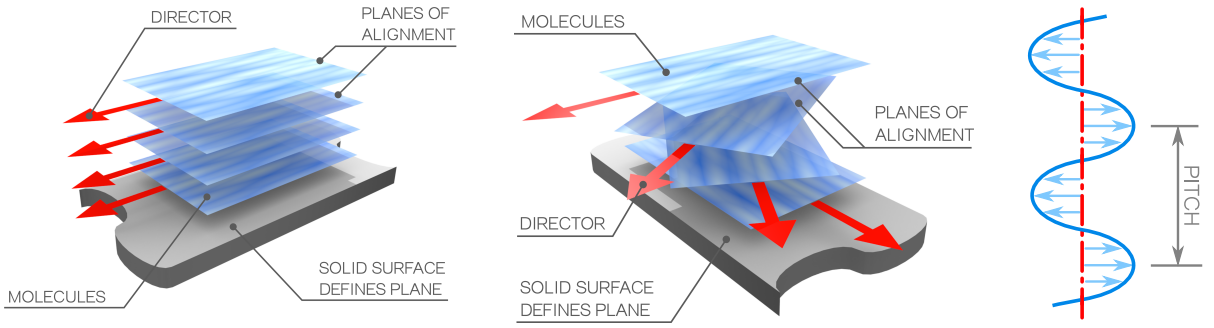
Schultz and Jones [95] noted that early LC formulations were sensitive to pressure, which interferes with the temperatures at which colour transitions are observed. Later work by Ireland and Jones [98] indicated that encapsulated liquid crystals (wherein the LC molecules are sealed inside polymer spheres of $40\mu m \leqslant \varnothing \leqslant 50\mu m$) were insensitive to pressure variations, allowing their use in a wide variety of test conditions.

At temperatures lower than their transition temperature, liquid crystals exist in a what is known as the "nematic" phase, wherein they align in stacked planes with the same orientation (referred to as their "director"). This phase is illustrated in Figure 2.2(a). As crystals approach their transition temperature, the aforementioned planes begin to rotate relative to one another, forming a helical structure. This is referred to as the "chiral nematic" phase, and is illustrated in Figure 2.2(b), while the resulting helical structure is shown in Figure 2.2(c). The pitch of a liquid crystal in the chiral nematic phase is the distance required for a given crystal to undergo a full 360^o twist.

In the nematic phase, liquid crystals appear to be translucent. As they enter the chiral nematic

phase, however, the wavelength of light that they reflect shifts with the rotation of the molecular planes. Provided that the pitch of the helical structure is within the visible part of the EM spectrum, this transition will cause circularly polarised light to be reflected, and hence a visible colour change will occur. Typically, high wavelengths are reflected first during a transition such that the colour play on a liquid crystal goes from red, through green, and finally arrives at blue before returning to transparency as the crystals enter their isotropic phase.

Liquid crystal thermography is easily one of the simpler methods of those described above to implement, as it only requires that the crystals be mixed with a stabilising binder solution and applied to a substrate via spraying. A layer of black ink is usually applied first in order to improve the visual contrast of transitions (this is particularly important when dealing with transparent substrates such as Perspex). A liquid crystal solution may comprise several types of crystals, each with a different transition temperature, thereby allowing multiple temperature transitions to be observed during a single experiment. An example of one such experiment performed during the present study is shown in Figure 2.3, wherein a cylindrical geometry with a hemispherical end was exposed to a heated, high Mach number flow. The three transitions may be observed as rings (labelled) on the hemispherical nose of the geometry.



(a) Nematic phase (b) Chiral nematic phase (c) Director helix

Figure 2.2: Schematics of nematic and chiral nematic phases of encapsulated thermochromatic liquid crystals.

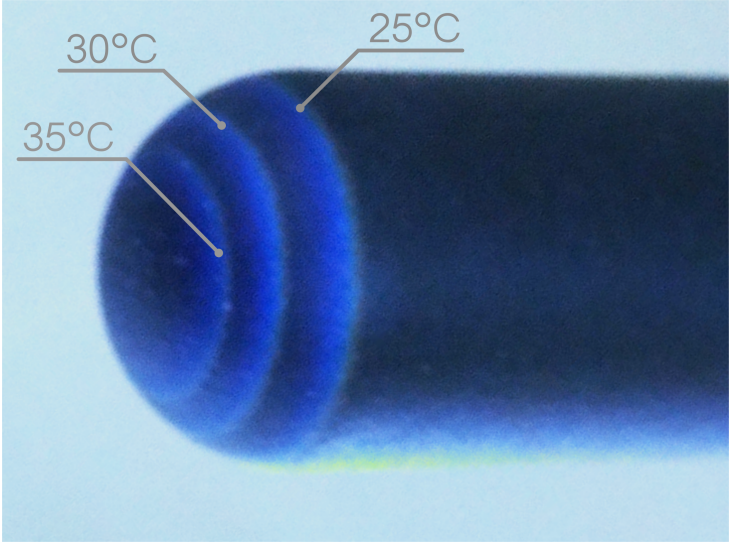


Figure 2.3: Photograph of the colour play on a test piece coated with a slurry comprising three encapsulated liquid crystal solutions with transition temperatures 25^oC, 30^oC, and 35^oC (labelled).

Ultimately, liquid crystals were chosen for the thermal flux experiments described in a later section of this thesis. This choice was motivated by the wide range of flow conditions that liquid crystal coatings may be applied to, coupled with the ease with which the relevant probes may be manufactured. Although additional effort is required to post process data from certain LC studies, this was thought to be an acceptable compromise given the advantages that LC experiments offer.

2.5.4 Transient thermal experiments

When studying thermal fluxes in a physical experiment, an ideal setup would include a step change in heat flux, $\dot{q}$, at the start of a test, where said heat flux then remains constant throughout. Changes in temperature would then be easily manipulated during post-processing in order to find the thermal response of the test piece. In high Mach number flow experiments, however, it is often challenging to generate a perfect step of $\dot{q}$, and virtually impossible to ensure a constant value throughout a test.

Instead, a step change in flow temperature is often sought, whereby the model under test is heated asymptotically towards the recovery temperature of the gas, T_r. The expression for T_r in a supersonic flow is presented in Eq. (2.42), where r is the flow's recovery factor. This strategy naturally causes $\dot{q}$ to decrease from an initially high value as the temperature of a test piece's substrate material increases from its initial temperature, T_i, towards T_r. The convective heat

transfer coefficient, h_{conv}, then becomes the desired output of this kind of study, as it does not vary in time during an experimental run as long as the flow conditions remain consistent. The relationship between convective heat transfer coefficient, heat flux, and temperature difference is presented in Eq. (2.18).

$$T_r = T_\infty \left(1 + r\frac{\gamma - 1}{2} {Ma_\infty}^2\right) \tag{2.42}$$

Schultz and Jones [95] have described this technique in detail, describing the analytic rise in temperature at the surface of a model with Eq. (2.44) where the variable β is given by Eq. (2.43). In analysing the data from a transient heat flux experiment, θ allows one to model the progression of temperature against time, and hence calculate the value(s) of h_{conv} at discrete points.

$$\beta = \frac{h_{conv}\sqrt{t}}{\sqrt{\rho c k}} \tag{2.43}$$

$$\theta = \frac{T_s - T_i}{T_r - T_i} = 1 - e^{\beta^2}\,\mathrm{erfc}\,(\beta) \tag{2.44}$$

A notable disadvantage of this method, however, lies in its initial assumption of one-dimensional conduction within the substrate material. This model also assumes that the thermal fluxes within the model substrate behave such that it may be considered semi-infinite in all directions except depth, i.e. the vast majority of heat transfer occurs normal to the probe surface. In the case of a blunt body probe in a hypersonic plume, this assumption has only a limited time during which it may be considered valid. This time is characterised by Eq. (2.45), where x_{tp} is the depth to which a thermal pulse penetrates a semi-infinite substrate in a time, t_{tp}, and α is the thermal diffusivity of the substrate material (Eq. (2.46)). Essentially, thermal pulse penetration depth is a thermal analogue of mean free time, $\bar{t}$ (see Eq. (2.40)), i.e. the time before a thermal pulse from one face of a geometry interacts with — and thereby interferes with — a pulse from another.

$$x_{tp} = 4\sqrt{\alpha t_{tp}} \tag{2.45}$$

$$\alpha = \frac{k}{\rho c} \tag{2.46}$$

Setting x_{tp} in Eq. (2.45) to half the diameter of a blunt body probe, $\varnothing/2$, and rearranging for t_{tp}, yields the maximum time for which the semi-infinite assumption is valid. In the event that a test's duration exceeds t_{tp}, any data gathered after this point will not yield accurate results if analysed using Eq. (2.44). While this method may still yield usable data up to t_{tp}, this will only apply for blunt bodies where there are no "sharp" corners or other discontinuities which may present regions with low thermal mass to the flow.

Consider the case of a cylinder of finite length with one of its flat ends facing into a heated hypersonic flow. The flow will surround the cylinder, impinging upon both the forwardmost circular face and the curved sides simultaneously. Thermal fluxes transfer heat into the substrate normal to its surface, so at the corner of a cylinder, the depth to which a thermal pulse may travel without encountering another will tend towards zero ($x_{tp} \rightarrow 0$). From Eq. (2.45), it is clear that t_{tp} will also tend towards zero ($t_{tp} \rightarrow 0$). As such, it may be determined that the semi-infinite substrate assumption will not function at all at the sharp corners of non-blunt test pieces. This is an important issue for the present work, in that many of the geometries which are used as analogues for satellite debris feature such areas where the semi-infinite assumption will not be valid.

Both Ling [101] and Ryley [4] have previously developed a method to combat this limitation which utilises experimental data in concert with FEM analyses (up to 2D in Ling's work, and 3D in Ryley's) in order to model the internal thermal behaviour of a test piece. This strategy works by initially processing the thermal history of an experiment using Eq. (2.44) (utilising the aforementioned semi-infinite substrate assumption), and extracting estimates of convective heat transfer coefficient, h_{conv}. These data, along with the thermal properties of the test gas, are then set as boundary conditions for a finite-element transient thermal simulation of the test piece (the duration of this simulation is set to be equal to that of the initial experiment). Once the simulation has been completed, the temperature rise across the surface of the test piece is reconstructed and compared to the experimental data. Errors are calculated, and new estimates of h_{conv} are estimated using Eq. (2.47) (where the variable ϕ is a damping term used to prevent divergence in estimations of h_{conv} between FEM simulations). This process is then repeated until the calculated errors are within a given tolerance. Ryley reports that temperature histories

generated using this method converge to within 5% of experimental values within 10 iterations.

$$h_{FE_{new}} = \phi \, h_{FE_{old}} \left(\frac{\sqrt{t_{FE}}}{\sqrt{t_{exp.}}} \right) + h_{FE_{old}} \left(1 - \phi\right) \tag{2.47}$$

Although the cases that Ryley considered featured subsonic, heated gas flowing over models of turbomachinery components, the implementation of this method for supersonic and hypersonic experiments is not expected to present significant issues. (LC techniques have previously been used by Owen [87] among others in hypersonic flow generated in the LDWT, and were observed to produce good results. Additionally, the method described above is mathematically robust, and is not physically affected by any of the factors which differ between subsonic and hypersonic studies.)

2.6 STATISTICAL ANALYSIS

When comparing the results of experiments and simulations, it is important to consider how errors between data sets are expressed. In many cases, the absolute (ϵ), relative (η), and relative percentage (δ) errors — presented in Eqs. (2.48), (2.49), and (2.50), respectively — are sufficient to characterise the differences between data points. Here, y is taken to be a "true" value (such as those derived from experimental data), while y' is taken to be a predicted value (such as those generated by simulations).

$$\epsilon = \left| y - y' \right| \tag{2.48}$$

$$\eta = \frac{\epsilon}{|y|} = \left| \frac{y - y'}{y} \right| = \left| 1 - \frac{y'}{y} \right| \tag{2.49}$$

$$\delta = 100 \, \eta = 100 \, \frac{\epsilon}{|y|} = 100 \left| \frac{y - y'}{y} \right| = 100 \, \left| 1 - \frac{y'}{y} \right| \tag{2.50}$$

In other cases, however, it may be advantageous to consider errors in such a way that they are compared to the full range of a given data set. The relative range percentage error, δ_R, is presented in Eq (2.51), and allows one to calculate a percentage error based upon the range of the data set being considered i.e. the difference between the maximum and minimum, as opposed to a specific value. This is useful in cases where there are large differences (usually greater than

one order of magnitude) between values in a data set.

For example, when considering lift coefficient, C_L, predictions, it is theoretically possible to obtain a value of zero. In reality, a calculation or measurement of such a value would likely be extremely small, but would almost certainly not be exactly zero. As such, the relative error between two disparate predictions (calculated using either Eq. (2.49) or Eq. (2.50)) could potentially be enormous, giving the impression of severe inaccuracy when in fact the absolute difference between the two values is very small. The use of relative range error (as in Eq. (2.51)), however, would essentially provide context to this comparison relative to the rest of the data set.

$$\delta_R = 100\ \eta_R = 100\ \left| \frac{y - y'}{\max(y) - \min(y)} \right| \tag{2.51}$$

When considering large data sets, it is sometimes useful to characterise errors using a single metric, as opposed to a point-by-point calculation (which must usually be graphed or tabulated for presentation). The expressions for root-mean-squared (RMS) and normalised root-mean-squared (NRMS(μ) and NRMS(R)) errors are shown in Eqs. (2.52), (2.53), and (2.54)).

The RMS error of a data set is similar to the absolute error, in that the resulting quantity will have the same units as the data set. The two equations for NRMS error return non-dimensional values, normalised to different quantities from the data set being considered. The first, NRMS(μ), is normalised to the mean of the data, while NRMS(R) is normalised to the data range (in a similar fashion to δ_R).

$$\delta_{RMS} = 100\ \eta_{RMS} = 100\ \sqrt{\frac{\sum_{n=1}^{N} |y_n{}' - y_n|^2}{N}} \tag{2.52}$$

$$\delta_{NRMS(\mu)} = 100\ \eta_{NRMS(\mu)} = 100\ \frac{\eta_{RMS}}{\overline{y}} \tag{2.53}$$

$$\delta_{NRMS(R)} = 100\ \eta_{NRMS(R)} = 100\ \frac{\eta_{RMS}}{|\max(y) - \min(y)|} \tag{2.54}$$

The mean values of δ and δ_R may be calculated trivially as in Eq. (2.54), but are included for completeness in Eqs. (2.55) and (2.56).

$$\overline{\delta} = 100\ \overline{\eta} \tag{2.55}$$

$$\overline{\delta_R} = 100\ \overline{\eta_R} \tag{2.56}$$

Finally, when considering the quality of a model fitted to a set of data, the coefficient of determination, R^2, may be used to express the quality of the model's predictions on a scale from 0 to 1 (where 0 denotes no correlation, and 1 denotes a perfect fit to the data). The expression for this metric is shown in Eq. (2.57).

$$R^2 = 1 - \frac{\sum_{n=1}^{N}(y_n - y_n{}')^2}{\sum_{n=1}^{N}(y_n - \overline{y})^2} \tag{2.57}$$

More advanced uncertainty quantification techniques exist, such as the "high dimensional derivative" method which Mehta et al. applied to the modelling of space debris re-entries [102]. These more advanced methods, however, are outwith the scope of the present work, and they are as such not applied to the data presented herein.

2.7 SURROGATE MODELLING

Surrogate modelling is a general term for any technique whereby correlations between a system's input parameters and its outputs are determined. As such, they are intended to infer continuous relationships between discrete points of data in an incomplete set; typically where data has been obtained via "expensive" means, such as large scale simulations or complex experiments. Surrogate models may be used to predict the response of a system based upon prior observations of its behaviour, and can be produced with little or no knowledge of the system's underlying dynamics.

Models of this type are used widely throughout all branches of science to interpolate between (or, less commonly, extrapolate beyond) data points, with some being optimised for as few as two variables (in the classic case of fitting a linear trend). Owing to the "expensive" nature of the simulations and experiments described herein, surrogate modelling is utilised extensively to correlate results. The two most frequently used methods in the present work are described below.

2.7.1 Curve and surface fitting

Curve and surface fitting are the most computationally straightforward surrogate modelling techniques, representing correlation of one and two independent variables, respectively. Typically, functions are fit to data by minimising the sum of squared residuals between the predicted function and the known values of the dependent variable (also known as least-squares fitting). This error takes the form shown in Eq. (2.58) for a simple case with two variables (curve fitting), where $\boldsymbol{\beta}$ is a vector of the coefficients which condition the response, y_i is an individual data point, and x_i is the coordinate at which y_i is measured. This optimisation technique is ubiquitous in modern scientific software, and as such will not be fully described here.

$$\min(\epsilon) = \min\left(\sum_{i=1}^{N}\left[y_i - f(x_i, \boldsymbol{\beta})\right]^2\right) \tag{2.58}$$

2.7.2 Gaussian process regression

Gaussian process regression (also known as GPR modelling, Kriging, or Wiener-Kolmogarov prediction) is a type of surrogate modelling that has gained a great deal of attention in machine learning research and as an optimisation technique for engineering problems, owing primarily to its generality and scalability. Rasmussen and Williams [103] have exhaustively described the process of formulating and exploiting GPR models as a general machine learning tool, while their use in engineering design was pioneered by Sacks et al. [104], who used them to design and correlate computer simulations. More recently, GPR models have been used by Forrester, Sóbester, and Keane [105] as a major component in engineering optimisation problems. In the field of geostatistics, GPR models are known as Kriging models; named after Danie Krige, a mining engineer who first used the procedure to estimate the density of gold deposits from a small scattering of exploratory boreholes.

Initially, a GPR model is defined solely by its mean and covariance (kernel) functions; this is termed the "prior distribution" (Duvenaud [106] observed that a GPR model's mean is often assumed to be zero, as fluctuations in the magnitude of the mean may be compensated for by adding additional terms to the kernel). Then, as with any surrogate model or regression method,

the coordinates of an observed variable ($\mathbf{x}$) are specified, as well as the value of said variable ($\mathbf{y}$) at that point. The GPR model is then "trained" using these data by combining the prior with a Gaussian likelihood function for each of the observed values; this is termed the "posterior". As such, the parameters of the covariance function at each observed point are optimised for the given data set, generating what is known as the "covariance matrix". Once this process is complete, a prediction (y') may be performed by calculating the weighted average of the model at some new point (denoted by the vector $\mathbf{x}'$). This provides the best linear unbiased estimator, as well as the value of the covariance of the GPR model at $\mathbf{x}'$.

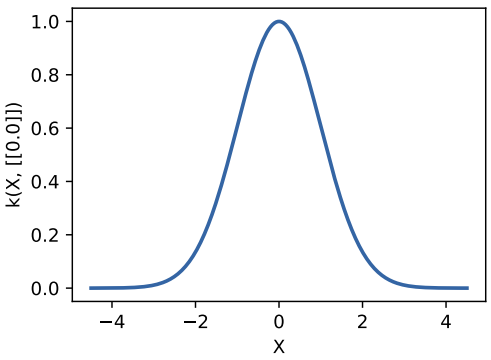
Kernel functions are typically defined by a lengthscale, ℓ, and their variance from the mean, σ_f. Since σ_f is simply a scaling factor, it is the value of ℓ which determines the "relevance" of the function to the value of a predicted variable. As it is unlikely that the response of a system with multiple input parameters will depend on them all to the same degree, a technique known as "automatic relevance determination" (ARD) is used, whereby separate covariance functions are utilised for different input parameters (termed "dimensions" here). These are then added or multiplied together across all dimensions of the GPR model for each observation.

$$K_{EQ,\,ARD}(\mathbf{x},\mathbf{x}') = \prod_{d=1}^{D} \sigma_{f_d}^2 \exp\left(-\frac{(\mathbf{x_d}-\mathbf{x'_d})^2}{2\ell_d^2}\right) \tag{2.59}$$

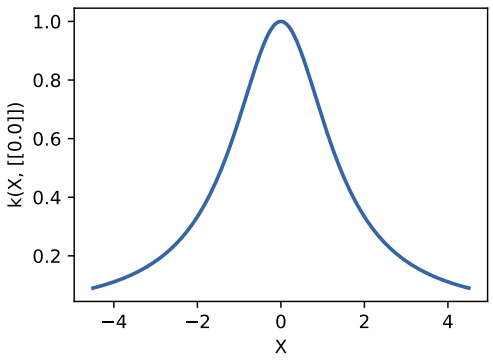
$$K_{RQ,\,ARD}(\mathbf{x},\mathbf{x}') = \prod_{d=1}^{D} \sigma_{f_d}^2 \left(1+\frac{(\mathbf{x_d}-\mathbf{x'_d})^2}{2\alpha_d\ell_d^2}\right)^{-\alpha_d} \tag{2.60}$$

$$K_{EQ\times RQ,\,ARD}(\mathbf{x},\mathbf{x}') = \sigma_{f_1}^2\sigma_{f_2}^2 \exp\left(-\frac{(\mathbf{x_1}-\mathbf{x'_1})^2}{2\ell_1^2}\right)\left(1+\frac{(\mathbf{x_2}-\mathbf{x'_2})^2}{2\alpha_2\ell_2^2}\right)^{-\alpha_2} \tag{2.61}$$

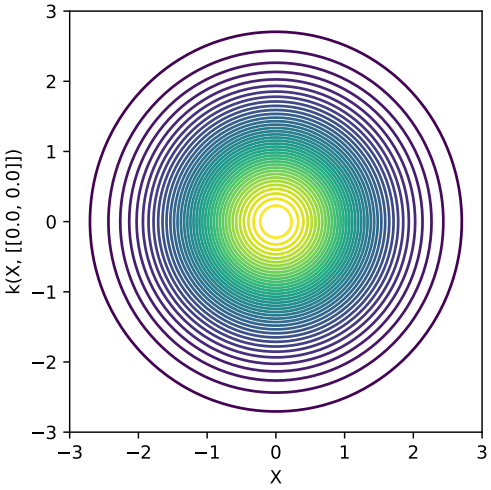
Three ARD kernels are featured in the present work, depending on which one produced the best fit for a given data set. Their suitability in each case was determined via statistical analysis of the GPR model's predictions versus observations that were not included in the training data. The first, presented in Eq. (2.59), is known as an exponential quadratic (K_{EQ}) or squared exponential (K_{SE}) kernel, and is a common starting point for many different GPR models, mostly owing to its wide applicability (as observed, once again, by Duvenaud [106]).



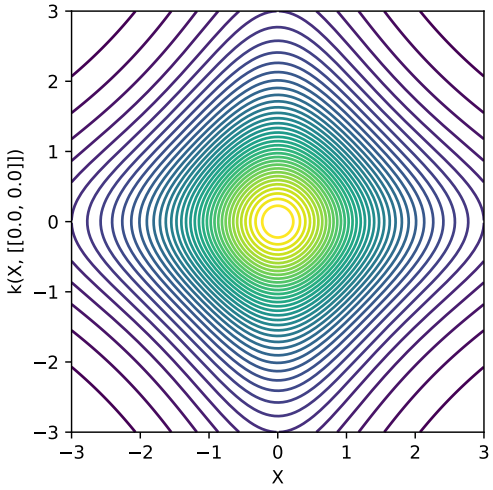
(a) K_{EQ}, $\ell = 1.0$, $\sigma_f = 1.0$, $D = 1$



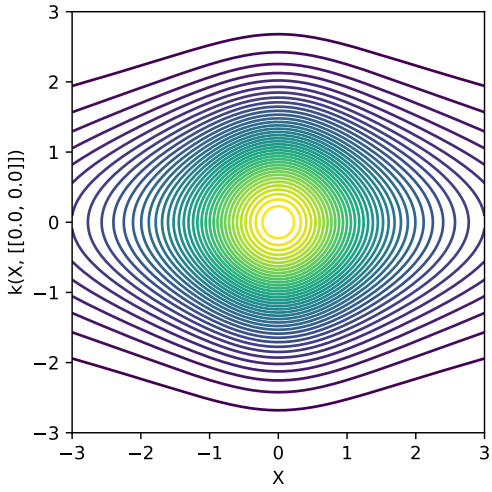
(b) K_{RQ}, $\alpha = 1$, $\ell = 1.0$, $\sigma_f = 1.0$, $D = 1$



(c) K_{EQ}, $\ell = 1.0$, $\sigma_f = 1.0$, $D = 2$



(d) K_{RQ}, $\ell = 1.0$, $\sigma_f = 1.0$, $\alpha = 1$, $D = 2$



(e) $K_{EQ \times RQ}$, $\ell = 1.0$, $\sigma_f = 1.0$, $\alpha = 1$, $D = 2$

Figure 2.4: Plots of various GPR kernel functions using normalised coordinates

The second is a slight variation, known as the rational quadratic (K_{RQ}) kernel, which features an additional parameter (α) that is used to alter the conditioning of the response. Note that in each of these cases, the final kernel function, K, is a product of the original function at each

observed point across D dimensions, denoting that these are ARD type kernels. Finally, the third kernel (presented in Eq. (2.61)) is a combination of one-dimensional versions of the previous two. Hence, this kernel is limited to two input dimensions (denoted here by the subscripts 1 and 2), but serves to demonstrate the ease with which individual covariance functions may be combined to create GPR models with complex responses.

Plots of the kernel functions with $D = 1$ are presented in Figures 2.4(a) and 2.4(b), while the ARD kernel functions for $D = 2$ are presented in Figures 2.4(c), 2.4(d), and 2.4(e). The first four of these plots demonstrate both the similarity of the two selected kernel functions to one another, as well as highlight the minor differences between them. The final figure shows an example of a hybrid kernel, demonstrating the variations in topology which may be achieved by arithmetically combining other kernel functions. Note that each of these plots show the kernel responses for a single input point in each case ($x = 0$ in the $D = 1$ cases, and $x = (0, 0)$ in the $D = 2$ cases).

The software utilised for all of the GPR modelling in the present work is *GPy*, which was developed by the Sheffield Machine Learning Group [107].

2.8 CHAPTER SUMMARY

In this chapter, the current state of the art in a number of subject areas are described. These include debris demise simulation technology, rarefied gas modelling, surrogate modelling, heat flux prediction, low and high fidelity aerothermal modelling, and the experimental evaluation of thermal surface distributions and loads experienced during re-entry.

It is apparent from this review that the debris demise simulation tools which are currently in use often use a great variety of different methods to perform their calculations. Comparative studies of such tools like that performed by Lips and Fritsche [108] serve to confirm that these codes suffer from large uncertainties in aerothermal loads when running calculations, especially in the slip and transition regimes. This is often due to the need for computational efficiency, which is generally accomplished via the implementation of heavily simplified aerodynamic and thermodynamic correlations. The use of such equations — as well as various other differences in calculation strategies between codes — can potentially result in large inaccuracies and uncertainties in the predicted survivability of certain space debris items. It is therefore feasible

that a re-entering piece of debris may survive transit through the Earth's atmosphere and impact a populated area, thereby causing damage to property, injury, and — in extreme circumstances — loss of life.

In order to improve these predictions, and hence decrease this uncertainty, improvements to both the aerodynamic and thermodynamic modelling of these debris items is necessary. However, there are currently very few experimental or numerical studies which focus on the non-aerodynamic blunt bodies that are used by debris demise tools as analogues for defunct satellite components.

As such, the relevant areas for improvement (and hence, the direction for the present work) are clear. Firstly, it is necessary to develop a new method for the rapid assessment of aerodynamic properties for integration into future debris demise tools. This method must be both more accurate than the simplistic tumble-averaged equations currently in use in some codes, and more computationally efficient than the advanced engineering-level panel methods used in others. (Note that the intended application of this method is that of an accurate aerodynamic database that may be rapidly queried by an iterative trajectory solver.)

Secondly, a study of the heating to these debris analogues must be conducted such that the thermal trends can be characterised and better understood. Both numerical and experimental evaluation of these thermal phenomena are necessary, in order to validate the capabilities of current high fidelity modelling tools for use with debris analogue geometries in the slip and transition regimes.

NUMERICAL MODELLING 3

A wide array of numerical models were used during the present work to characterise the aerothermal behaviour of geometric debris analogues. These were split into two distinct categories: those performed using the high fidelity DSMC method, and those performed using low fidelity hypersonic panel methods. DSMC (Direct simulation Monte Carlo) simulations were run with freely available or open source software (*DS2V* and *dsmcFoam+*, respectively - see Section 3.4), while panel method calculations were performed using a code developed by the present author (*RAC*, Re-entry Aerothermal Calculator - see Section 3.3).

In this chapter, the atmospheric modelling strategy used during the present work is examined, followed by 3 DOF (Degree of freedom) trajectory simulations on spheres of various masses and diameters. The development of the aforementioned panel method code *RAC* is then described, and the results of the trajectory calculations used as the upper and lower bounds of a sampling plan of input conditions for simulations. These panel method calculations are then used to train a series of GPR (Gaussian process regression) models which are intended to replace other aerodynamic calculation methods in future debris demise codes.

Following these low-fidelity calculations, a series of DSMC simulations are presented wherein the edge heating effects on cylinders are examined at various freestream conditions and angles of incidence. These data are then also correlated using GPR models such that results could be effectively interpolated between cases.

The goal of these numerical studies was twofold. Firstly, the aerodynamic behaviour of the debris analogues during atmospheric entry was to be characterised such that the accuracy and computational efficiency of future trajectory calculations (which currently use either simplified tumble-averaged aerodynamic models or potentially expensive engineering level panel methods - see Section 2.1) may be improved. This is vital to the correct calculation of thermal loads, as variations in aerodynamic properties can greatly affect the duration of a re-entry trajectory, potentially introducing large inaccuracies in estimations of heat transfer (and, hence, whether or not debris will survive re-entry intact). This work is also intended to increase the computational efficiency of such calculations, such that future debris demise codes may more quickly run full

trajectory calculations without sacrificing accuracy (a useful improvement when one considers the need to run uncertainty quantification on a candidate re-entry trajectory/breakup).

The second of these goals involves understanding the thermal behaviour of the debris analogues during re-entry. This is important as it provides reference data which may be compared to experimental results, as well as identifying peak loads on areas of interest to the present work, such as sharp corners and edges. Since these features are neither blunted nor aerodynamically optimised, they are likely to heat extremely quickly during re-entry. As such, the phenomenon of heat transfer at sharp edges and corners, i.e. regions where the orientation of a surface changes suddenly, must also be characterised in order to accurately predict whether or not similar debris items will survive their transit through the atmosphere.

The DSMC simulations were performed using input conditions that could be matched by the experimental facilities available, as well as those which could not. This allowed comparisons to be made between the DSMC method and experimental results, as well as for correlations in different flow regimes (see Section 2.2.2) to be generated.

Hypersonic panel methods were used to produce large quantities of low-fidelity data in order to generate and iteratively improve correlative aerodynamic models without the computational expense of the DSMC method. GPR modelling (see Section 2.7) was used extensively in conjunction with both the high and low fidelity numerical methods in lieu of curve or surface fitting.

Additionally, trajectory simulations were performed in order to better inform the choice of input conditions for the other models.

3.1 ATMOSPHERIC MODELLING

During the course of the numerical simulations, it was necessary to accurately model the properties of the working gas (atmospheric air in this case). While simulations performed using freestream conditions from the LDWT could be characterised using measurements and thermochemical databases (see Appendix A), this was not true of those simulations that were to be performed using full scale quantities like those experienced during an actual atmospheric entry. As such, the NRLMSISE00 [109] Earth atmospheric model was utilised in the form of a Python

Variable	**NRL abbreviation**	**Unit**	**Value used in *RAC***
Day of year	`doy`	days	172
Year	`year`	years	0
Time of day	`sec`	seconds	29000
Geodetic latitude	`g_lat`	degrees	60
Geodetic longitude	`g_long`	degrees	120
Local apparent solar time	`lst`	hours	16
Average of F10.7 flux over 81 days	`f107A`	sfu	150
Daily F10.7 flux	`f107`	sfu	150
Daily magnetic index (K-index)	`ap`	N/A	4

Table 3.1: Input variables used in the NRL atmospheric model for the *RAC* linear interpolation model. Note that those variables which mention the "F10.7 flux" refer the $10.7\,cm$ $(2800\,MHz)$ radio flux, which is used as an indicator of solar activity.

implementation [110]. This model was selected based upon its versatility, its accuracy, and the fact that it was originally intended for use in satellite drag calculations. It features a spherical harmonic representation of the Earth's atmosphere from ground level upwards, and is capable of producing profiles of static pressure and temperature with respect to altitude, longitude, latitude, solar activity level, and time of year. During actual calculations, fixed values were used for all of these input parameters with the exception of altitude (the location chosen was $13^o19'00.0"N$, $169^o09'00.0"W$, the splashdown location of the Apollo 11 capsule in the North Pacific Ocean). Plots of the generated static pressure and temperature profiles are presented in Figure 3.1.

The NRL model is also able to generate profiles of molecular species concentration. In the present work, this capability is coupled with the Cantera [39] thermodynamics toolkit in order to allow the calculation of quantities not normally assessed by the NRL model such as specific heats, viscosity, etc. Cantera performs its calculations using the 9-term NASA thermodynamic polynomials for air which, when coupled with the NRL model, produce profiles such as those presented in Figure 3.2.

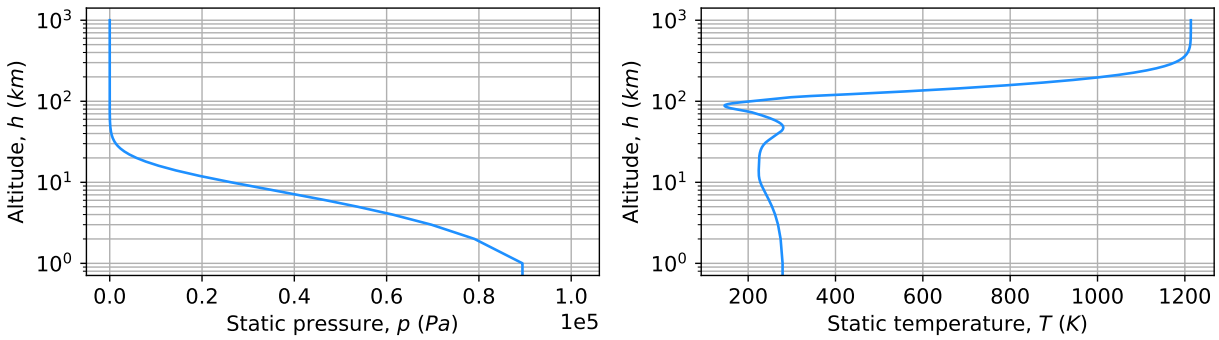


Figure 3.1: Pressure and temperature profiles generated by the *NRLMSISE-00* atmospheric model.

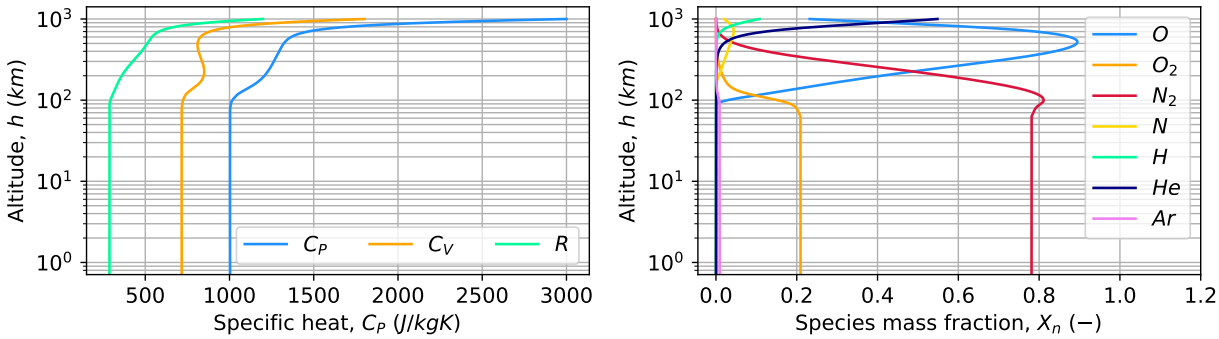


Figure 3.2: Specific heat and species mass fraction variation with altitude generated using the *NRLMSISE-00* atmospheric model and the Cantera chemical kinetics package.

3.2 TRAJECTORY SIMULATIONS

As with the atmospheric modelling described above, it was necessary to perform calculations in order to derive the freestream properties that would be achieved during "real" atmospheric entry trajectories. These calculations were performed using the coordinate system presented in Figure 3.3, which features three degrees of freedom. These are: altitude (h), ground range (r), and flight path angle (φ). A simple 3DOF trajectory calculator code called *traj_calc* was developed by the present author for these calculations.

The governing equations for this type of trajectory simulation are presented in Eqs. (3.1) through (3.5), in a form similar to that utilised by Tetzman [111]. These equations are capable of modelling a lifting re-entry trajectory (i.e. both lift and drag forces are considered) occurring in the atmosphere of a spherical, non-rotating planet. This type of model is well understood and is frequently utilised in debris re-entry codes owing to its simplicity and robustness (see Section 2.1).

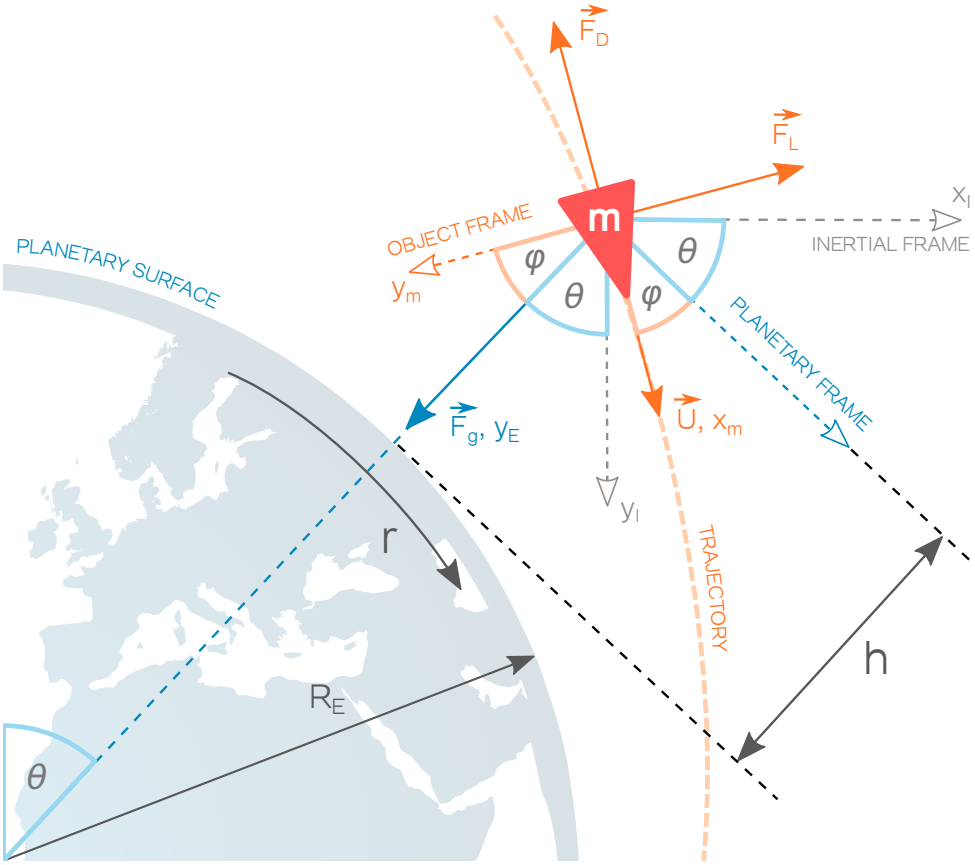


Figure 3.3: Coordinate system used to derive the governing dynamic equations for 3 DOF trajectory modelling. The object (body-centred) reference frame (subscript m) is shown in orange, the planet-centred frame (subscript E) in blue, and the inertial frame (subscript I) in grey. Velocity, force and direction vectors are also included.

The first of these expressions, Eq. (3.1), models the re-entering object's rate of change of velocity in its own reference frame (termed the "object" frame in Figure 3.3, and also referred to as the "body-centred" frame), while the second, Eq. (3.2), models the rate of change of the flight path angle, φ (which describes the rotation of the body frame relative to the planetary frame). The next two expressions operate in the "planet-centred" frame, calculating the rates of change of altitude (Eq. (3.3)) and ground range (Eq. (3.4)), respectively. Finally, Eq. (3.5), models the variation in gravitational acceleration with altitude on a spherical analogue of the Earth.

$$\frac{dU}{dt} = g\sin(\varphi) - \frac{\rho A U^2 C_D}{2m} \tag{3.1}$$

$$\frac{d\varphi}{dt} = \frac{g\cos(\varphi)}{U} - \frac{\rho A U C_L}{2m} - \frac{U\cos(\varphi)}{h + R_E} \tag{3.2}$$

$$\frac{dh}{dt} = -U\sin(\varphi) \tag{3.3}$$

$$\frac{dr}{dt} = \frac{U R_E \cos(\varphi)}{R_E + h} \tag{3.4}$$

$$g = g_0 \left(\frac{R_E}{R_E + h} \right)^2 \tag{3.5}$$

Re-entry trajectory simulations were performed for a series of spheres with varying diameters and masses, in the ranges $0.1 \leqslant \varnothing \leqslant 10.0\,m$ and $10 \leqslant m \leqslant 1000\,kg$, respectively. Each simulation was run using the same initial altitude ($200\,km$) with each sphere having its orbital velocity calculated using the Eq. (3.6), where G is the gravitational constant ($6.67384 \times 10^{-11}\,m^3/(kg \cdot s^2)$), and M_E and R_E are the mass ($5.9742 \times 10^{24}\,kg$) and radius ($6378.136\,km$) of the Earth, respectively. This initial altitude was selected as it is generally accepted to be the upper limit of strong atmospheric drag effects on orbiting objects (with freestream densities, ρ_∞, of the order of $1 \times 10^{-10}\,kg/m^3$).

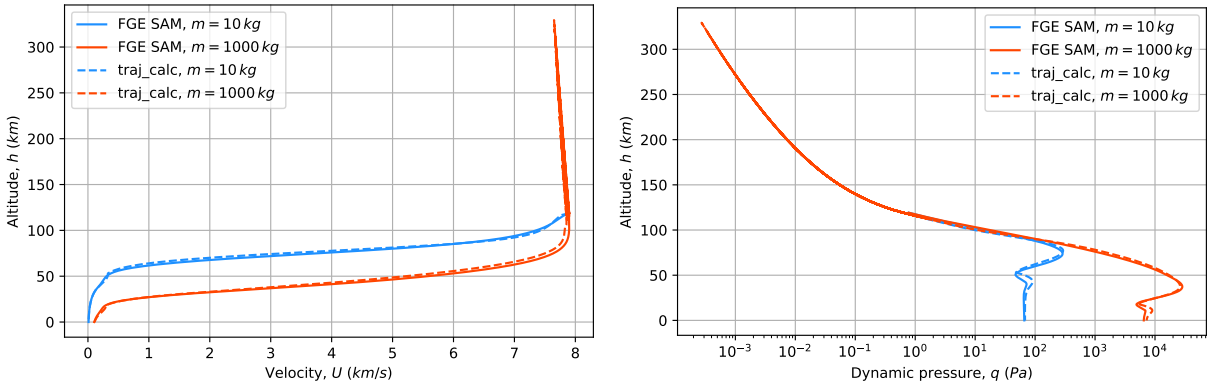
$$v = \sqrt{\frac{G(M_E + m)}{R_E + h}} \tag{3.6}$$

The aerodynamic properties of the sphere in hypersonic flight are well understood in both the continuum and free molecular regimes, having been investigated in the former by Masson, Morris, and Bloxsom [46], Grabow et al. [45], and Hodges [112], and in the latter by Schaaf and Chambre [113]. Experimental measurements were performed using spheres with rough surfaces (compared to polished metal) at Mach numbers in the range $4 \leqslant Ma_\infty \leqslant 58.7$ and Reynolds numbers in the range $2 \times 10^1 \leqslant Re_\infty \leqslant 1 \times 10^6$. Analytic calculations in the free molecular regime utilised a smooth sphere assumption.

As such, spheres are ideal candidates for exploratory simulations such as those presented herein. The drag coefficient, C_D, of the spheres was set to 2.0 in the free molecular regime ($Kn \geqslant 10$), and 0.92 in the continuum regime ($Kn \leqslant 1 \times 10^{-3}$), with Wilmoth's sine-squared bridging law [68] being utilised in the transition regime. The lift coefficient, C_L, of the spheres was set to 0 in all simulations.

Simulations of sphere re-entries were also used to validate the *traj_calc* code. Figure 3.4 shows the results of trajectory simulations performed using *traj_calc* and a similar 3 DOF code developed by Fluid Gravity Engineering (FGE). Two spheres, with $\varnothing 1\,m$ and masses of $10\,kg$

and $1000\,kg$, were simulated using both codes. As can be seen in the aforementioned figure, both codes agree well for the vast majority of the re-entry trajectories. Small differences between them are visible in the dynamic pressure, q, predictions, at $\sim 50\,km$ and $\sim 10\,km$ for the $10\,kg$ and $1000\,kg$ spheres, respectively. This was found to be a side effect of differing atmospheric models between the codes (*traj_calc* uses the previously described NRLMSISE00 model, while the FGE code uses the US76 model). Other than this minor discrepancy, both codes were found to agree well, hence validating *traj_calc*.



(a) Velocity, U, against altitude, h

(b) Dynamic pressure, q, against altitude, h

Figure 3.4: Plots of dynamic pressure and velocity against altitude for various trajectory simulations showing the similarity between results from different codes. The geometries being modelled are spheres of diameter $\varnothing 1\,m$ and mass $10\,kg$ and $1000\,kg$.

The results of the *traj_calc* trajectory simulations are presented in Figure 3.5 in terms of Mach and Knudsen numbers, allowing the cases to be effectively compared to one another (Knudsen numbers were calculated using each sphere's diameter as the characteristic length, L, and the freestream mean free path, λ_∞). It is clear from these results that the trajectories are extremely similar initially, only diverging well into the transition regime. The Mach numbers of these simulations prior to the onset of the continuum regime lay in the range $8 \leqslant Ma \leqslant 33$, confirming that the slip-transition regime is dominated by hypersonic flight conditions. As such, the assumptions on which subsequent numerical analyses and experiments are based, i.e. that all slip and transition regime cases considered have Mach numbers above the hypersonic limit, is shown to be valid (see Section 2.2.3).

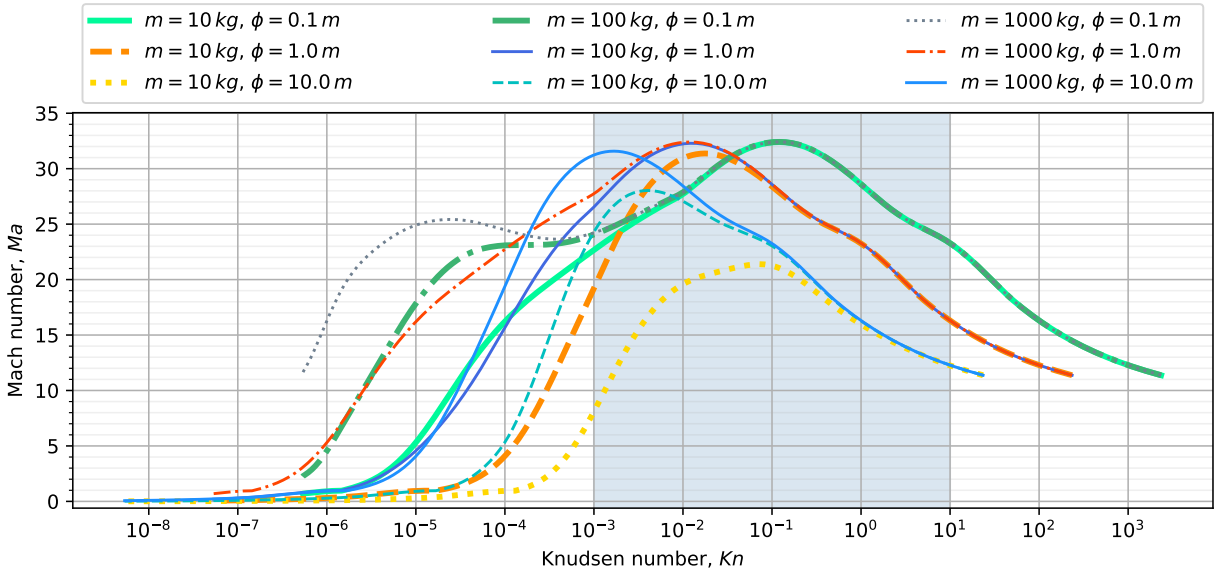


Figure 3.5: Re-entry trajectory profiles of various spheres returning from low Earth orbit (LEO) in Mach-Knudsen space. The slip-transition regime is highlighted in blue.

From these data, the bounds of subsequent aerodynamic simulations were set to $5 \leqslant Ma \leqslant 35$, and $1 \times 10^{-5} \leqslant Kn \leqslant 1 \times 10^{4}$ in order to cover the full range of expected Mach numbers across all examined flow regimes. The lower bound of Knudsen number was selected owing to the low likelihood of debris items having extremely low area-to-mass-ratios, such as that of the case with $m = 1000\,kg$ and $\varnothing = 0.1\,m$.

3.3 PANEL METHOD SIMULATIONS

A panel method code was developed by the author for the current work, such that simulations and the development of correlations could be more effectively automated. The result of this effort was *RAC*, the Re-entry Aerothermal Calculator [32].

RAC was written entirely using the Python programming language, and is open source. Although the code does rely heavily on the compiled C and FORTRAN routines of the *NumPy* and *SciPy* libraries, these are all linked, interfaced, or wrapped using the Python C API, meaning that controlling (or indeed editing) the program is as easy as manipulating any other Python script.

RAC operates in much the same way as other panel method codes, but places an emphasis on interoperability above all else. The code is operated using a simple input and execution script, and is written using an object-oriented class structure (note that this refers to the programming

style used during development, and *not* the debris breakup methodology mentioned in Section 2.1).

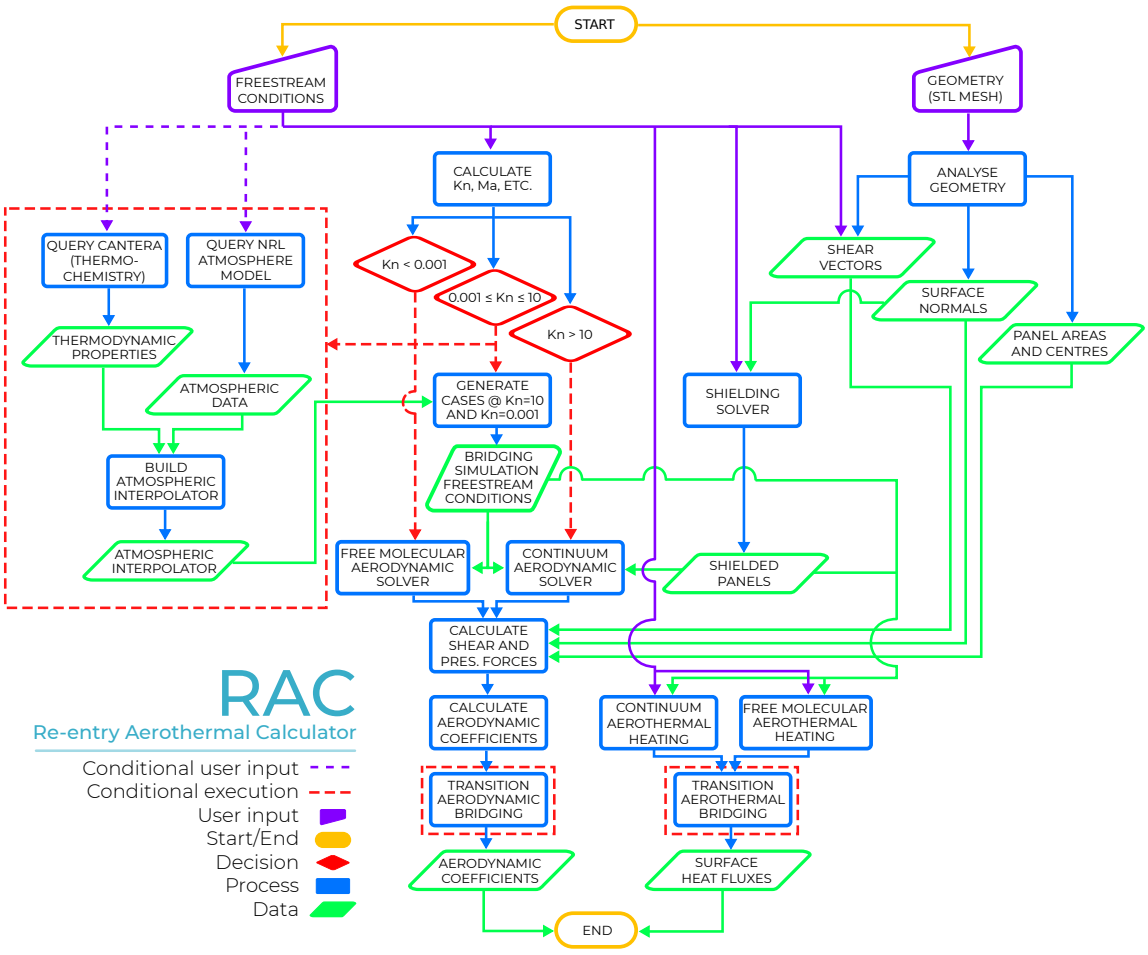


Figure 3.6: RAC program flow diagram

3.3.1 Freestream conditions

Freestream conditions are entered into *RAC* as two separate arrays. The first consists of the attitudes at which the analysis is to be performed, stored as an $n \times 3$ array of pitch, yaw and roll angles, $[\alpha, \beta, \psi]$, (relative to an appropriate Cartesian coordinate system). The second consists of the freestream conditions at which the analyses are to be performed, stored as an $m \times 6$ array. The freestream parameters input using this array are the Mach number, ratio of specific heats, static pressure, velocity, density, and static temperature, $[Ma, \gamma, p_\infty, U_\infty, \rho_\infty, T_\infty]$. From these parameters, all other fluid flow quantities may be calculated, resulting in $m \times n$ simulated cases.

This results in what is known as a full factorial sampling plan; a simulation setup whereby

an analysis is performed for every combination of the input parameters (in this case, attitudes and freestream conditions). In the event that a user wishes to utilise a sampling plan for their simulations which is not full factorial, *RAC* allows the entry of single case definition arrays of the form $[\alpha, \beta, \psi, Ma, \gamma, p_\infty, U_\infty, \rho_\infty, T_\infty]$.

3.3.2 Geometry preprocessing

RAC operates on geometric surface models made up of tessellated triangular faces or "panels". The wide availability of triangular mesh preprocessing tools was a driver for this decision. Meshes are input using the popular *STL* format, and are quickly analysed using a series of simple equations (described below).

When a geometry is imported into *RAC*, each of the mesh panels is processed such that its geometric centre, surface area, and unit normal are calculated. A reference length and area are also required from the user as a basis for later calculations. Additionally, the user must also enter the XYZ coordinates of the object's centre of gravity (COG) so that moment arms may be calculated, as well as the surface temperature of the object (which may either be a constant or a per-face distribution). Finally, the user must enter the surface accommodation coefficients σ_N and σ_T, which are used to define the response of a surface to particulate collisions in the normal and tangential directions, respectively (see Figure 3.7).

These coefficients are important for the calculation of pressure and shear coefficients (C_p and C_τ, respectively) in the free molecular regime, and lie in the range $0 \leqslant \sigma \leqslant 1$. Mechanically, the accommodation coefficients represent some fraction of particles incident upon a surface, σ, undergoing diffuse reflection, i.e. scattering in a random direction. The remaining fraction of incident particles, $(1 - \sigma)$, is reflected specularly, i.e. a "mirror-like" reflection in a single direction. A value of 0.9 was used for both of these coefficients throughout the present work, which is in line with the typical values used for interactions between air and heated metals, and is also commonly used in satellite re-entry simulations [114] [115].

RAC utilises the panel orientation angle relative to the freestream vector, θ for its pressure coefficient calculations, and the complimentary angle (between the surface normal and the freestream vector), δ for its shielding and geometric calculations. These angles are illustrated in

Figure 3.7, and the appropriate relations for calculating each one are presented as Eqs. (3.7) and (3.8) below. Note that since *RAC* operates using unit vectors for the panel normals and freestream direction, a slight simplification of the geometric expressions may be made (as $|\hat{N}||\hat{V}| = 1$).

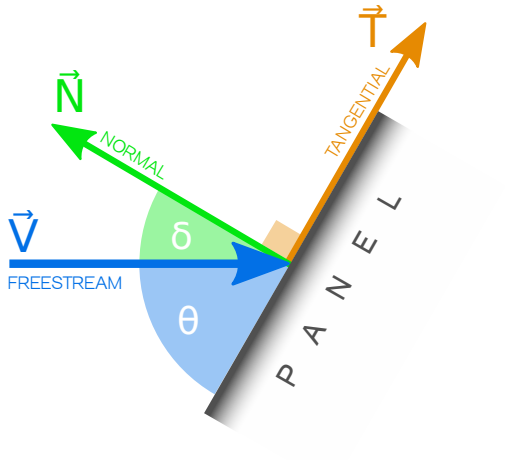


Figure 3.7: Diagram showing the angles defined in Eqs. (3.7) and (3.8) for panel method calculations.

$$\delta = \cos^{-1}\left(\frac{\hat{N}\cdot\hat{V}}{|\hat{N}||\hat{V}|}\right) = \cos^{-1}\left(\hat{N}\cdot\hat{V}\right) \tag{3.7}$$

$$\theta = \sin^{-1}\left(\frac{\hat{N}\cdot\hat{V}}{|\hat{N}||\hat{V}|}\right) = \sin^{-1}\left(\hat{N}\cdot\hat{V}\right) \tag{3.8}$$

Geometry preprocessing is repeated for each separate attitude in the input array. If desired, the user may also input the mesh curvature, although this must be computed by an external tool.

3.3.3 Cross sectional area calculation

The user has the option to define the total reference area with which *RAC* will calculate the final aerodynamic coefficients. Either an override may be enabled, wherein the reference area is used unilaterally, or an automatic calculation may be invoked. This calculation involves rotating the mesh in memory using Euler rotation matrices such that the local X axis coincides with the freestream vector. The mesh is then flattened i.e. the depth (X) dimension is ignored, and one of two algorithms is used to calculate the enclosed area:

1. **Convex hull** This method is useful for large, convex bodies such as re-entry capsules. The algorithm is supplied by the *QHull* [116] library via the Python *SciPy* module and is used to detect the outermost vertices of the point cloud. This polygon then undergoes a Delaunay triangulation, whereafter the projected area may be easily calculated by summing the areas of each individual triangle.

2. **Concave hull** This method is supplied by the *CGAL* [117] library via the Python *cgal-bindings* module. The algorithm numerically calculates the border around a point cloud, conforming the boundary according to the density of the cloud (this also allows the identification of "holes" in the middle of the cloud). Once this border is calculated, a separate algorithm identifies the individual loops in the geometry and utilises the so-called Shoelace formula (Eq. (3.9)), a special case of Green's theorem for 2D polygons, to calculate the enclosed area [118].

$$A = \frac{1}{2}\left\{\sum_{i=1}^{n-1} x_i y_{i+1} + x_n y_1 - \sum_{i=1}^{n-1} x_{i+1} y_i - x_1 y_n\right\} \tag{3.9}$$

An example of both methods being used to calculate the cross section of a complex geometry (the ESA Automated Transfer Vehicle – abbreviated ATV – oriented at $\alpha = 75^o$) is presented in Figure 3.8. The result of the convex hull calculation is shown in green, while the result of the concave hull calculation is shown in red.

As is apparent from the image, the convex hull procedure is suited primarily to convex geometries such as spheroids, aeroshells and crew capsules. It also has the advantage of being totally automatic. The concave hull calculation, by contrast, may be applied to arbitrary geometries, but may also require modification to produce a suitable polygon via the adjustment of a sensitivity parameter. *CGAL* is capable of estimating the optimal value of this parameter, but user entry is possible in *RAC* in the event that fine tuning is required.

The accuracy of the convex and concave hull calculations was evaluated using a simple sphere test case and the aforementioned ATV case. The true cross sectional area of each case was evaluated using a trivial circle area calculation and CAD software (projection onto a plane), respectively. The maximum error between the calculated areas and the "true" values was 1.5%.

3.3.4 Shielding analysis

Shielding is the property of a geometric mesh to obscure parts of itself from the perspective of a "viewer", which in the case of the present work is the freestream vector representing the oncoming flow. In the context of *RAC*, a shielded (or "leeward") panel is any face which will

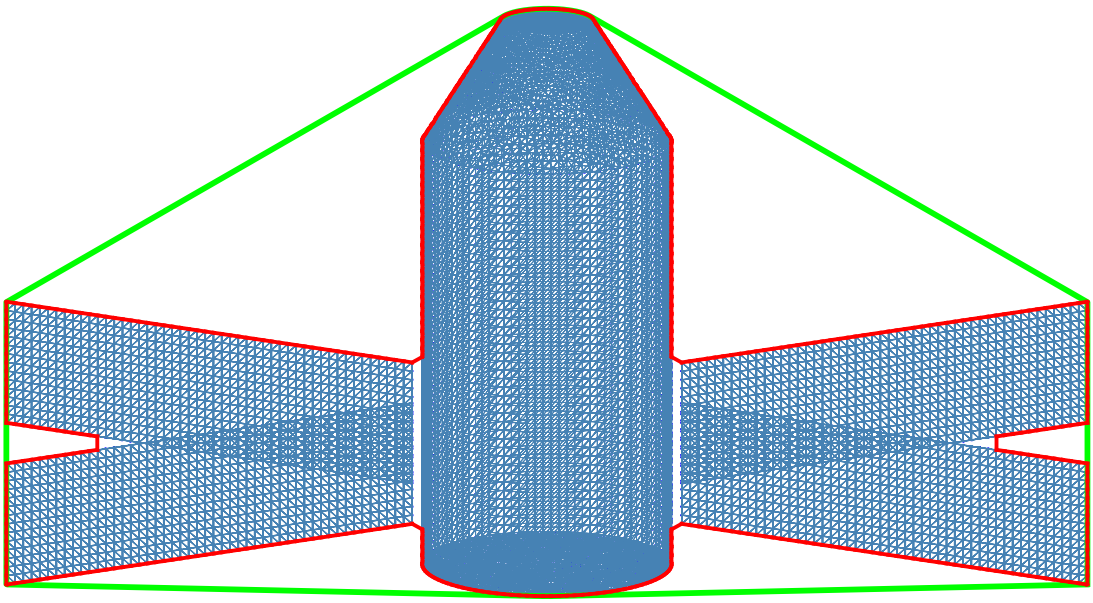


Figure 3.8: An example of convex and concave hull calculations on the ATV surface mesh performed using *RAC* at $\alpha = 75^o$. The convex hull boundary is shown in green, while the concave hull boundary is shown in red.

not be exposed to the flow during flight at a given attitude, while unshielded (or "windward") panels will. There are separate flow calculations for leeward panels which differ significantly from their windward counterparts, necessitating the identification of which panels are in fact shielded during an analysis. *RAC* accomplishes this in one of two ways:

1. **Panel orientation checking** A fast method ideally suited to simple convex shapes such as spheroids. The angle between each panel normal and the freestream vector, δ, is calculated. If $\delta \geqslant 90^o$, the panel is treated as leeward. If $\delta < 90^o$ the panel is treated as windward.

2. **Visual panel identification** This method, which is based upon the Pixelator script [119], requires the use of a plotting package for which *RAC* utilises *MayaVI* [120]. An orthographic vector image of the geometry is plotted in memory, with each panel being assigned a unique colour. This view is then rotated so that the viewpoint attitude coincides with the freestream vector. The resulting image is saved as a vectorised *EPS* (Encapsulated Postscript) and processed by *PIL* (Python Imaging Library), wherein the unique colours present in the image are catalogued. These are then linked to their respective faces by way of a lookup table, and a boolean operation is performed: panels whose colours have been detected in the saved image are tagged as windward, while all "unseen" panels are tagged

as being leeward. This process is described graphically in Figure 3.9.

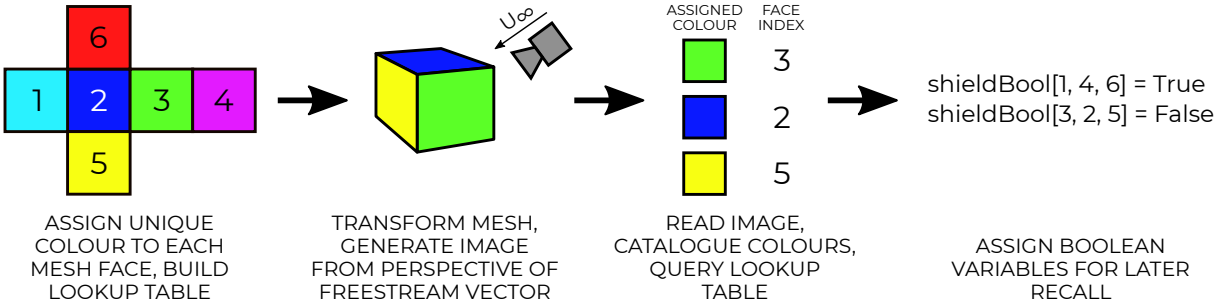


Figure 3.9: Diagram describing the algorithm used for the visual shielding check in *RAC*.

3.3.5 Aerodynamics

Aerodynamic analyses in *RAC* utilise surface inclination methods. These reduced order methods are extremely computationally efficient, and relate the local pressure and shear coefficients, C_p and C_τ, respectively, to the angle of incidence of a particular panel with respect to the freestream vector and conditions.

Several different varieties of calculation method may be used, and these vary automatically based upon which "regime" a particular analysis case lies in. *RAC* considers three separate regimes independently of one another based upon the Knudsen number, Kn, which is calculated using Eq. (2.3), and a reference length, L, which must be provided by the user (this value remains fixed for the sake of consistency throughout the course of a simulation). These regimes, as well as their associated Kn bounds, are listed in Table 3.2.

Regime name	***Kn* bounds**
Continuum	$Kn < 0.001$
Transition	$0.001 \leqslant Kn < 10.0$
Free molecular	$Kn \geqslant 10.0$

Table 3.2: Flow regime Knudsen number bounds used in *RAC*. Note that the reference length, L, used in the calculation of Knudsen number is fixed in *RAC*, and is a required input for any simulation.

3.3.5.1 Continuum aerodynamics

The continuum regime utilises panel methods derived from the *S/HABP* code of the AFRL. This code was originally based upon the Douglas Aircraft Company's Gentry code [121], and has

since inspired the development of multiple other codes such as NASA's *CBAERO* [122].

S/HABP was originally formulated to assist in the design of such vehicles as combat aircraft and missiles, which are generally launched from ground based facilities or low-altitude aircraft. As such, the operational velocities used in analyses of these vehicles are significantly lower than those of objects or spacecraft entering the atmosphere from orbit. This in turn results in a need to evaluate the behaviour of such craft at low Reynolds numbers, and a corresponding necessity to include viscous force calculations in simulations involving the continuum regime.

Since *RAC* is designed to exclusively simulate descents into planetary atmospheres, it was assumed that all objects analysed using the code would have a high residual velocity upon reaching the continuum regime. Owing to the characteristically high Reynolds numbers of high velocity continuum flow, as well as the comparatively short duration of flight in this regime for re-entering objects, it was therefore assumed that effects of viscous forces in the continuum regime would be negligible. This is a common assumption in re-entry studies which dates back to the earliest panel method calculations [59]. Hence, the shear coefficient in the continuum regime is automatically set to zero during a simulation i.e. $C_{\tau_{cont}} = 0$.

Several different options exist for continuum regime analyses, each one having its own specific area of application. The windward and leeward continuum methods currently available in *RAC* are described below.

Windward methods

1. **Modified Newtonian** A widely used method based on Newtonian theory. It is most accurate when applied to blunt bodies in high Mach number, idealised flow ($M = \infty$; $\gamma = 1$). This modified version utilises the stagnation pressure of the freestream as a corrective factor for finite Mach number flows.

$$C_{p_{cont}} = C_{p0_{cont}} \sin^2(\theta) \tag{3.10}$$

2. **Tangent Wedge** An empirical method based on the expansion of an oblique shock beyond

a semi-infinite wedge.

$$C_{p_{cont}} = \frac{4(Ma^2 \sin^2(\theta_s) - 1)}{Ma^2(\gamma + 1)} \tag{3.11}$$

3. **Tangent Cone** An empirical method originally based on the expansion of a conic shock beyond a semi-infinite cone. This method is derived from second order slender body theory and the empirical method of Hammitt and Murthy [123]. It is well suited to analysing bodies of revolution whose surface panels have a small inclination angle relative to the flow. *RAC* also includes the older Edwards, and HABP Mk. III versions of the tangent cone relation, which are slightly more computationally efficient but significantly less accurate at lower Mach numbers.

$$C_{p_{cont}} = \frac{2p_\infty}{\rho_\infty {U_\infty}^2} \left\{ \frac{\left[\frac{2\gamma}{\gamma+1} {Ma_\infty}^2 \sin^2(\theta_s) - \frac{\gamma-1}{\gamma+1}\right]}{\left[1 + \frac{\gamma\, Ma_\infty (\theta_s - \theta_c)^2 \cos^2(\theta_s)}{1 + \frac{\gamma-1}{2} {Ma_\infty}^2 \sin^2(\theta_s)}\right]} \right\} \tag{3.12}$$

Leeward methods

1. **Newtonian** The bounding assumption of Newtonian flow is that any shadowed regions have the same surface pressure as the freestream, hence they exhibit a surface pressure coefficient of zero.

$$C_{p_{cont}} = 0 \tag{3.13}$$

2. **High Mach base pressure (full vacuum)** Another bounding assumption of high speed flow is that a total vacuum may form in the wake region of the spacecraft.

$$C_{p_{cont}} = \frac{-2}{\gamma {Ma_\infty}^2} \tag{3.14}$$

3. **High Mach base pressure (70% vacuum)** Recirculation and expansion into the wake region may drive the surface pressure up above the bounding assumption of a full vacuum, as is briefly discussed by Bonner, Clever, and Dunn [124]. The semi-empirical expression shown in Eq. (3.15) predicts a pressure slightly above vacuum, hence returning a pressure

coefficient between the Newtonian and full vacuum bounds.

$$C_{p_{cont}} = \frac{-1}{{Ma_\infty}^2} \tag{3.15}$$

RAC allows the user to set switching values for panel orientation (relative to the freestream vector) and curvature (as interpolated from surrounding vertices) such that different methods may be automatically selected by the program on a panel-by-panel basis. For example, an Apollo-style blunt body of revolution may be treated with both the modified Newtonian method on high incidence faces, and with the tangent cone method on faces near the curved, low incidence edge of the body. Similarly, for a SHEFEX-style body constructed of flat facets, a combination of modified Newtonian and tangent wedge could be employed. It must be noted, however, that this capability can potentially result in non-physical step changes in pressure and shear coefficients if ill-informed switching values are specified for panel orientation and curvature.

Other methods that are available for continuum aerodynamics in *RAC* include OSU blunt body (windward), Hankey flat surface (windward), Dahlem-Buck empirical (windward), and van Dyke unified (both windward and leeward). These legacy methods are included in the code for the sake of evaluating them against modern simulation and experimental data, and are not actively employed during the analyses presented herein. Note that, in the expressions for Eqs. (3.11) and (3.12), shock angle, θ_s, may be computed using a number of methods. In *RAC*, much like in the *S/HABP* code, the empirical expressions of NACA report TR-1135 [125] are used for shock angle and normal shock component Mach number.

3.3.5.2 Free molecular aerodynamics

In the free molecular regime, *RAC* utilises the analytical methods of Schaaf and Chambre [113] for both C_p (Eq. (3.16)) and C_τ (Eq. (3.17)). Due to the nature of the flow in this regime, both windward and leeward faces may be treated effectively with these equations. The hypersonic speed ratio, s, is computed using Eq. (3.18). The momentum accommodation coefficients σ_N and σ_T are used to describe the interaction of gas particles with a surface, and are further elaborated upon in Section 3.3.2.

$$C_{p_{fm}} = \frac{1}{s^2}\left\{\left[\frac{2-\sigma_N}{\sqrt{2}}s\sin(\theta)+\frac{\sigma_N}{2}\sqrt{\frac{T_w}{T_\infty}}\right]e^{-(s\sin\theta)^2}+ \left[(2-\sigma_N)((s\sin(\theta))^2+0.5)+\frac{\sigma_N}{2}\sqrt{\frac{\pi T_w}{T_\infty}}s\sin(\theta)\right]\left[1+\mathrm{erf}(s\sin(\theta))\right]\right\} \tag{3.16}$$

$$C_{\tau_{fm}} = \frac{\sigma_T\cos(\theta)}{s\sqrt{\pi}}\left\{e^{-(s\sin(\theta))^2}+\sqrt{\pi}s\sin(\theta)\Big[1+\mathrm{erf}(s\sin(\theta))\Big]\right\} \tag{3.17}$$

$$s = Ma\sqrt{\frac{\gamma}{2}} \tag{3.18}$$

3.3.5.3 Transition aerodynamics

In the transition regime, which exists between the continuum and free molecular regimes, *RAC* does not directly perform calculations, but rather interpolates from bounding cases.

When a set of flow conditions is preprocessed by *RAC*, its Knudsen number is calculated using Eq. (2.3). If this value lies in the transition regime as defined in Table 3.2, *RAC* begins the process of bridging. To perform transition bridging analyses, *RAC* calculates the mean free path lengths, λ, which correspond to Knudsen numbers of 0.001 and 10.0 (utilising the characteristic length, L, of the geometry being analysed). The program then queries the NRL atmospheric model to find the altitudes at which these values of λ exist (note that this is possible only due to the fact that λ increases monotonically with altitude).

From these altitudes, two bounding sets of freestream conditions are generated for each unique set of Euler angles (α, β, and ψ) – one at the upper bound of the continuum regime, and one at the lower bound of the free molecular regime. This calculation involves once again querying the NRL model at the aforementioned altitudes, thereby returning freestream static pressures and temperatures, p_∞ and T_∞, respectively. The freestream Mach number, Ma_∞, is then held constant between the two bounding cases, and the remainder of the required input variables calculated (γ, U_∞, and ρ_∞). Note that the ratio of specific heats, γ, is calculated using

the Cantera chemical kinetics package.

Simulations are then run on the full range of geometry attitudes using the appropriate continuum and free molecular solvers, totalling $n \times 2$ additional simulations for each transition regime case (where n is the number of unique sets of Euler angles to be analysed). By holding the Mach number constant, the freestream Reynolds number, Re_∞, is altered instead, and pressure and shear coefficient distributions are generated for the boundary cases.

RAC incorporates a global bridging method which utilises the sine-squared function of Wilmoth [68] (see Eq. (2.21) in Section 2.3.2) to bridge aerodynamic coefficients between the continuum and free molecular boundary cases. This bridging law is utilised in *RAC* owing both to the ease of its implementation, and the broad use (and hence, validation) which it has seen to date in the study of re-entry aerodynamics. The parameters of the function are set as $a_1 = 0.375$ and $a_2 = 0.125$ so that the curve's gradient reduces at the appropriate values of Knudsen number (0.001 and 10).

Once the pressure and shear coefficient distributions of the aforementioned boundary cases have been evaluated, the equations described in Section 3.3.6 are used to calculate the aerodynamic coefficients of the geometry. These data are then used in the bridging formula (Eq. (2.21)), with the transitions case's original Knudsen number being used to generate the final set of transition aerodynamic coefficients.

3.3.6 Forces, moments, and aerodynamic coefficients

Forces are calculated by *RAC* using a summation of pressures and shear stresses over all of the model's panels, as well as each individual panel's surface area. Pressures and shear stresses are calculated using Eq. (3.19) and Eq. (3.20), and the shear vector, $\hat{S}$, is calculated using Eq. (3.21). Surface forces and moments are subsequently calculated using Eq. (3.22) and Eq. (3.23), before the corresponding coefficients are calculated using Eq. (3.24) and Eq. (3.25), respectively. Note that the reference area, A_{ref}, may be treated by *RAC* either as a constant (for the sake of consistency between analyses), or calculated using the methods described in Section 3.3.3. If A_{ref} is set as a constant, it must be entered by the user during the setup of a simulation (along with the reference length, L_{ref}).

The final force coefficients, C_D, C_L, and C_S lie in the reference frame of the geometry being analysed, and are calculated by multiplying an Euler rotation matrix (in zyx order, the inverse of xyz order used for initial orientation alteration) with the appropriate values of pitch (α), yaw (β), and roll (ψ) to C_x, C_y, and C_z, respectively (using Eq. (3.26)). The moment coefficients C_{mx}, C_{my}, and C_{mz}, are kept in the global coordinate system.

$$p = p_\infty + \frac{1}{2} C_p \rho_\infty {U_\infty}^2 \tag{3.19}$$

$$\tau = \frac{1}{2} C_\tau \rho_\infty {U_\infty}^2 \tag{3.20}$$

$$\hat{S} = \left(\hat{N} \times \hat{V}\right) \times \hat{N} \tag{3.21}$$

$$\overrightarrow{F} = pA\hat{N} + \tau\hat{S} = \begin{bmatrix} F_x \\ F_y \\ F_z \end{bmatrix} = pA \begin{bmatrix} \hat{N}_x \\ \hat{N}_y \\ \hat{N}_z \end{bmatrix} + \tau A \begin{bmatrix} \hat{S}_x \\ \hat{S}_y \\ \hat{S}_z \end{bmatrix} \tag{3.22}$$

$$\overrightarrow{M} = \overrightarrow{F} \times \overrightarrow{L} = \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} \left(F_y L_z\right) - \left(F_z L_y\right) \\ -\left(F_x L_z\right) + \left(F_z L_x\right) \\ \left(F_x L_y\right) - \left(F_y L_x\right) \end{bmatrix} \tag{3.23}$$

$$\begin{bmatrix} C_x \\ C_y \\ C_z \end{bmatrix} = \frac{\overrightarrow{F}}{\frac{1}{2}\rho_\infty {U_\infty}^2 A_{ref}} \tag{3.24}$$

$$\begin{bmatrix} C_{mx} \\ C_{my} \\ C_{mz} \end{bmatrix} = \frac{\overrightarrow{M}}{\frac{1}{2}\rho_\infty {U_\infty}^2 A_{ref} L_{ref}} \tag{3.25}$$

$$\begin{bmatrix} C_D \\ C_L \\ C_S \end{bmatrix} = \begin{bmatrix} \cos(\alpha) & \sin(\alpha) & 0 \\ -\sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos(\beta) & 0 & -\sin(\beta) \\ 0 & 1 & 0 \\ \sin(\beta) & 0 & \cos(\beta) \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\psi) & \sin(\psi) \\ 0 & -\sin(\psi) & \cos(\psi) \end{bmatrix} \cdot \begin{bmatrix} C_x \\ C_y \\ C_z \end{bmatrix} \tag{3.26}$$

3.3.7 Heat transfer

Heat transfer in *RAC* is calculated independently on each surface element of a computational mesh. These calculations are performed using the equation developed by Lees [126] in the continuum regime (Eq. (3.27)), and the analytical expression derived by Schaaf and Chambre [113] [47] in the free molecular regime (Eq. (3.28)). Note that continuum regime heat fluxes are normalised to the stagnation heat flux, $\dot{q}_0$, which must be calculated separately, whereas the free molecular heat fluxes are calculated directly. Here, the variable s is the hypersonic speed ratio, which is defined in Eq. (3.18).

$$\left(\frac{\dot{q}}{\dot{q}_0}\right)_{cont} = \frac{2\theta\sin(\theta)\left[1-\frac{\cos^2(\theta)}{\gamma\, Ma^2}+\frac{1}{\gamma\, Ma^2}\right]}{\sqrt{\left(1-\frac{1}{\gamma\, Ma^2}\right)\left(\theta^2-\frac{\theta\sin(4\theta)}{2}+\frac{1-\cos(4\theta)}{8}\right)+\frac{4}{\gamma\, Ma^2}\left(\theta^2-\theta\sin(2\theta)+\frac{1-\cos(2\theta)}{2}\right)}} \tag{3.27}$$

$$\begin{aligned}\dot{q}_{fm} =& \frac{\alpha_T \rho V^3}{4s^3\sqrt{\pi}}\Bigg\{\left[s^2+\frac{\gamma}{\gamma-1}-\frac{\gamma+1}{2(\gamma-1)\frac{T_w}{T_\infty}}\right] \\ & \left[e^{-\left(s\cos(\theta)\right)^2}+\sqrt{\pi}\left(s\cos(\theta)\right)\left[1+\mathrm{erf}\left(s\cos(\theta)\right)\right]\right]-\left[\frac{1}{2}\,e^{-\left(s\cos(\theta)\right)^2}\right]\Bigg\}\end{aligned} \tag{3.28}$$

It should be noted that Lees' equation is formulated based upon a spherical blunt-nosed body, much like the stagnation heat transfer equations described in Section 2.3.3 (which are almost all derived from experiments and calculations performed on similarly spherical geometries). A planned improvement to *RAC* for future heat transfer studies is the inclusion of Eckert's flat plate heating equation [127], which has been used to great effect in the *SAM* debris demise code. This expression allows for the characteristic increase in heat transfer near a sharp edge to be far more effectively predicted, whereas Lees' equation has no dependence on surface length scale, thereby making the consideration of sharp geometric features impossible. Eckert's equation was not implemented during the present work as DSMC simulations were intended to be used for heat transfer studies from the outset.

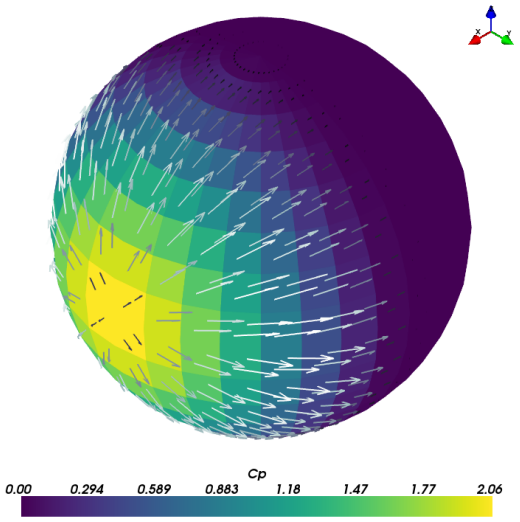
3.3.8 Code validation

A series of simulations were run using *RAC* in order to validate the code against known reference cases. Specifically, these were a sphere and a cylinder. In each case, the modified Newtonian equation (Eq. (3.10)), Wilmoth bridging equation (Eq. (2.21)), and Schaaf and Chambre equations (Eqs. (3.16) and (3.17)) were used to analyse the geomtries' windward faces in the continuum, slip-transition, and free molecular regimes, respectively. Shielded faces were treated using the 70% vacuum high Mach base pressure expression (Eq. (3.15)) in the continuum regime, and with the same equations as windward faces in the slip-transition and free molecular regimes.

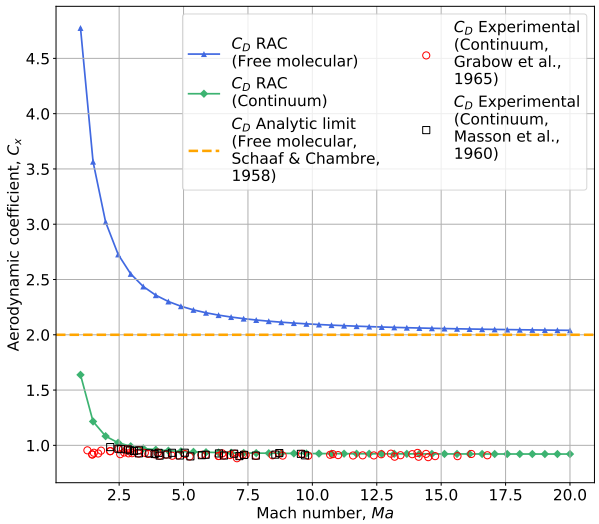
3.3.8.1 Sphere

The results of the sphere validation are shown in Figure 3.10. The calculation of free molecular regime pressure coefficient, C_p, and shear stress vectors, τ, are shown in Figure 3.10(a). As can be seen from aerodynamic coefficient data in Figure 3.10(b), *RAC* accurately predicts that the drag coefficient in the continuum regime, $C_D(C_x)$, will tend towards 0.92 as Mach number tends towards infinity, a result confirmed experimentally by Grabow, Pannabecker, and Plate [45], Masson et al. [46], and Hodges [112] amongst others. The free molecular results presented in the same figure conform well to expected values as calculated by Schaaf and Chambre [113], and approach the correct limiting value of 2.0. The comparatively gentler approach of C_D towards this asymptote is the result of the inherent rarefaction of the gas in the free molecular regime, which also results in such phenomena as diffuse shocks.

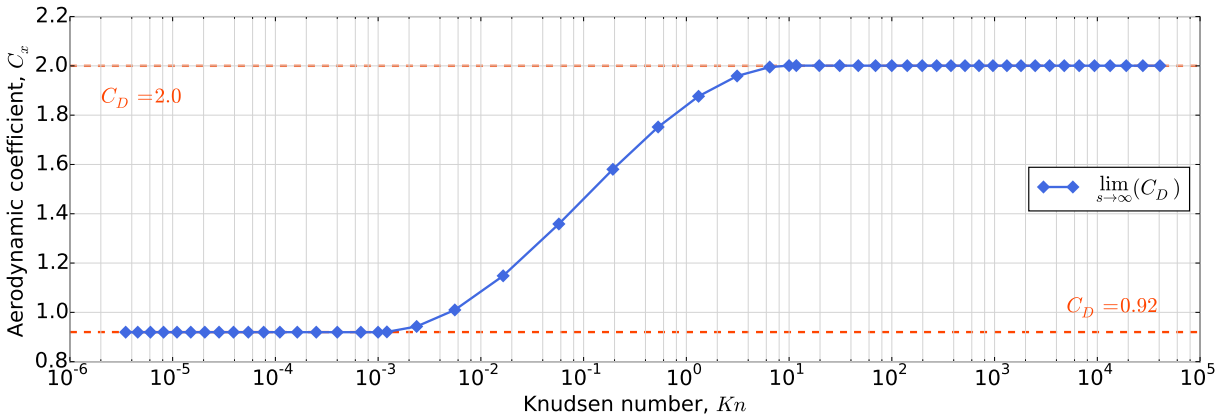
Figure 3.10(c) shows a series of calculations performed over a range of Knudsen numbers with a high Mach number (and, hence, a high molecular speed ratio, s). The resulting data represent the limits of the presented equations as $Ma \to \infty$ and $s \to \infty$ for the given geometry, and conform exactly to established $Kn\,/C_D$ correlations such as those used in object oriented analysis codes [3]. Additionally, the values of lift coefficient, $C_L(C_z)$, and sideslip coefficient, $C_S(C_y)$, remained at zero throughout the simulation, as expected according to the symmetric nature of the mesh.



(a) C_p distribution and vectors of τ in the free molecular regime



(b) Drag coefficient against Mach number



(c) Limit of drag coefficient against Knudsen number as $Ma \to \infty$

Figure 3.10: RAC simulations on a sphere at various Mach and Knudsen numbers.

3.3.8.2 Cylinder

The cylinder validation exercise was performed on a geometry of unit diameter and length. This geometry was chosen as it is a common debris analogue used in object-oriented calculations of destructive re-entry trajectories, representing among other things propellant tanks, reaction wheels, structural members, etc.

A sweep of freestream Knudsen number, Kn, as $Ma \to \infty$ and $s \to \infty$ was once again performed, and the results compared to the correlations of Klett [51] for the continuum and free molecular regimes. Klett performed his calculations analytically, developing expressions for the aerothermal behaviour of cylinders in fixed orientations and randomly tumbling along one or all axes in both the continuum and free molecular regimes. As such, his determinations represent

limiting values for the cases considered, similar to that of the sphere case presented above.

Orientation	**Continuum (laminar) C_D**	**Free molecular C_D**
Side on ($\alpha = 90^o$)	$\frac{2}{3}(2 - \kappa)$	2.0
End on ($\alpha = 0^o$)	$(2 - \kappa)\frac{\pi D}{4L}$	$\frac{\pi D}{2L}$

Table 3.3: Drag coefficient relations for a sphere at fixed attitudes in continuum and free molecular regimes

Klett determined the aerodynamic drag relations for a cylinder based on a set cross-sectional reference area of LD (cylinder diameter multiplied by length). These relations are presented in Table 3.3. It should be noted that Klett utilised an empirically-derived corrective factor of 0.909 for the mean pressure experienced by a cylinder oriented with its circular face facing the freestream direction ($\alpha = 0^o$) in the continuum regime. As *RAC* utilises the modified Newtonian method for its pressure calculations, and does not automatically detect common shapes, it was not possible to apply this correction automatically in the presented analyses. Additionally, the corrective factor is not attitude dependant, and applies only to the $\alpha = 0^o$ case, whereas *RAC* allow for calculations to be performed at an arbitrary attitude. As such, the values presented for comparison from Klett's work have had this corrective factor removed in order to show that the correct generalised limits for the Newtonian method are being approached.

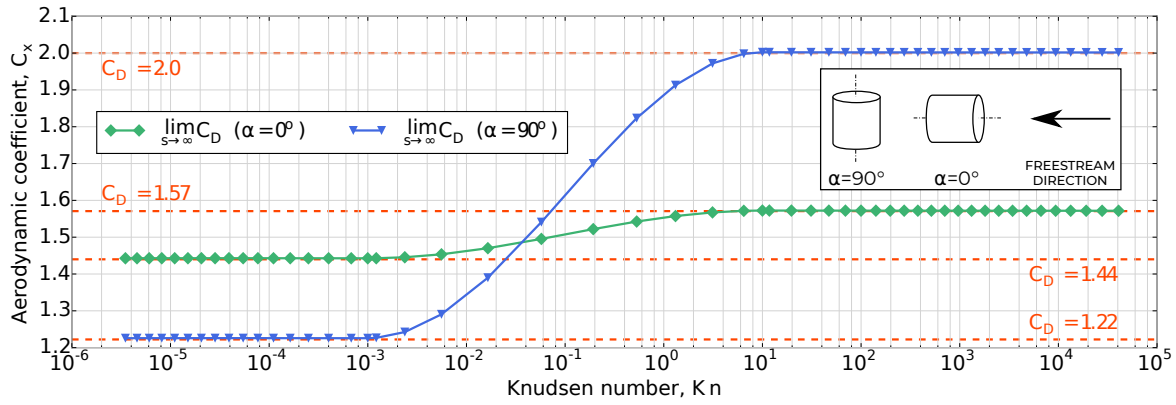


Figure 3.11: Drag coefficient, C_D, of a unit cylinder at fixed attitudes across a wide range of Knudsen numbers, Kn

As $Ma \to \infty$, the value of the pressure ratio across a normal shock, κ tends toward a limit of $\frac{1}{6}$. Hence, for a unit cylinder ($D = 1$, $L = 1$), the limiting values of drag coefficient, C_D,

in the laminar continuum regime may be calculated as 1.22 and 1.44 for "side-on" ($\alpha = 90^o$) and "end-on" ($\alpha = 0^o$) orientations, respectively. Similarly, for the free molecular regime, the calculated values are 2.0 and 1.57, respectively.

Figure 3.11 shows the values of drag coefficient, C_D, as calculated using *RAC* across a wide range of Knudsen numbers, Kn, for the given fixed orientations. The limits as calculated from the work of Klett (including the uncorrected Newtonian limit mentioned above) are superimposed as dotted lines. Excellent agreement may be observed between the limiting simulations and the calculated values as determined using Klett's methods.

3.3.9 Correlation of results in terms of Knudsen number, Kn

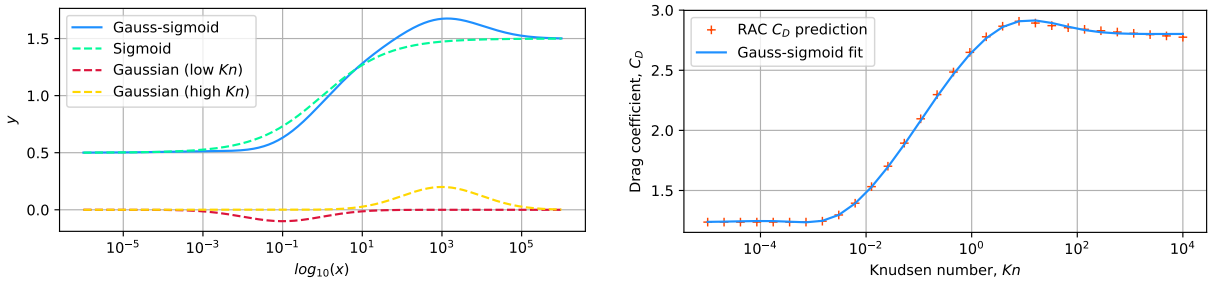
It was found following several simulations that the trends of aerodynamic coefficients against Knudsen number were consistent regardless of geometry size or attitude (this included studies featuring spheres, cylinders, and cuboids). Additionally, small variations in the shape of this trend with Mach number at the fringes of the transition regime were found to be poorly reconstructed by the debris correlations obtained from literature [3].

As such, a curve was developed in order to better facilitate the correlation of these variables. A sigmoid was used as the base, and two Gaussians added in locations close to the lower and upper bounds of the $log_{10}(Kn)$ range. The result of this was termed the "Gauss-sigmoid" curve, the equation for which is presented as Eq. (3.29).

$$f(x) = \left[\frac{a_1}{1+\exp\left(-c_1\left(\log_{10}(x)-b_1\right)\right)} + d_1\right] + \left[a_2\exp\left(-\frac{\left(\log_{10}(x)-b_2\right)^2}{2c_2^2}\right)\right] + \left[a_3\exp\left(-\frac{\left(\log_{10}(x)-b_3\right)^2}{2c_3^2}\right)\right] \tag{3.29}$$

Figure 3.12(a) illustrates the construction of the curve from its constituent functions (which are overlaid as dotted lines for clarity). Figure 3.12(b) shows the suitability of the fit provided by the Gauss-sigmoid from sample simulations performed on a cylinder (following optimisation of the coefficients through use of a least-squares fitting method). Similarly, the Gauss-sigmoid curve was found to be appropriate for the correlation of all aerodynamic coefficients throughout the

present work. As is mentioned above, this suitability persisted regardless of geometry, attitude, or freestream conditions.



(a) Curve construction (b) Least sqaures fit to C_D data

Figure 3.12: Gauss-sigmoid curve fitting to cylinder C_D data in terms of $log_{10}(Kn)$

Note the peak in the data presented in Figure 3.12(b) at $Kn \approx 10^1$. This feature was observed to manifest at low Mach numbers, and reduce in magnitude as $Ma \to \infty$ (note that $s \to \infty$ as $Ma \to \infty$). It is essentially a density-based analogue of the rapid change in aerodynamic coefficients which is apparent around $Ma = 1$, an effect that has been observed experimentally during hypersonic rarefied studies of spheres by Masson, Morris, and Bloxsom [46]. Although this study also featured high Kn experimental studies of sphere aerodynamics, the aforementioned peak was not apparent owing to the high Mach number of the experimental flow (as high as $Ma = 56.7$ in some experiments).

Parameter	Sub-curve affected	Value
b_1	Sigmoid	-1.0487
c_1	Sigmoid	1.5329
b_2	Gaussian (low Kn)	0.8000
c_2	Gaussian (low Kn)	0.7919
b_3	Gaussian (high Kn)	-2.620
c_3	Gaussian (high Kn)	0.5928

Table 3.4: Fixed Gauss-sigmoid parameters and their effects on the overall curve

Over the course of these simulations, several of the parameters in the Gauss-sigmoid fits were found to vary by only small amounts. Accordingly, their values were fixed, both reducing the complexity of the fitting procedure, and significantly improving its robustness. The fixed

parameters and their values are presented in Table 3.4. (It should be noted for the sake of completeness that altering the selected Kn bounds for the transition regime would require that these parameters be recalculated.)

3.3.10 New aerodynamic models

Gaussian process regression was used to correlate the remaining parameters of the Gauss-sigmoid curves, utilising the freestream Mach number, and the aspect ratio and orientation of the modelled geometries as inputs. This resulted in models comprising four separate GPR models (one for each of the non-fixed Gauss-sigmoid parameters a_1, d_1, a_2, and a_3). Two geometries were examined during this part of the present work: a cylinder and a cuboid.

The freestream conditions used for each simulation were generated using the atmospheric model described in Section 3.1, in conjunction with the Mach and Knudsen number bounds presented in Section 3.2. Sets of parameters were generated using an optimal Latin hypercube sampling (LHS) routine (as described by Forrester, Sobester, and Keane [105]), and the appropriate *RAC* simulations run via an automated script.

3.3.10.1 Cylinder

The cylindrical geometry was supplied to *RAC* as a triangulated STL file, wherein the parameters generated by the Latin hypercube sampling plan could be used to rotate and scale it. This allowed the aspect ratio (calculated as length over diameter, $L/\varnothing$) and orientation of the geometry to be adjusted for different simulations, while the freestream conditions for *RAC* were generated using the Mach and Knudsen numbers computed by the sampling plan, in conjunction with the NRL atmospheric model. Table 3.5 shows the bounds used in the LHS routine to produce the input conditions for *RAC*, as well as the scales on which the variables were evaluated.

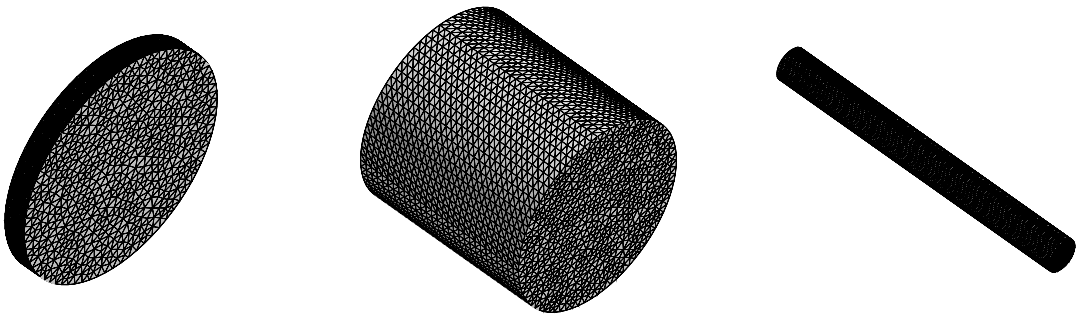
As is shown in Figure 3.14, the centre of gravity (COG) of the geometry was set as the centre of the rearmost circular face as opposed to the geometric centre of the shape as a whole. This facilitated the calculation of moment coefficients about an offset centre of mass, as well as allowing future users of the resulting GPR model to scale these coefficients by adjusting the COG location (and, hence, the lengths of the moment arms) relative to the cylinder's geometric

centre.

Parameter	Symbol	Lower bound	Upper bound	Scale
Mach number	Ma	5	35	Linear
Knudsen number	Kn	1×10^{-5}	1×10^{4}	Logarithmic
Pitch angle	α	0.1^o	89.9^o	Linear
Aspect ratio	$L/\varnothing$	0.1	100	Logarithmic

Table 3.5: Parameter bounds and scales used in generating input conditions for *RAC* simulations utilising the cylinder geometry.

Although the geometry mesh shown in Figure 3.14 is that of a unit cylinder ($L = 1\,m$, $\varnothing = 1\,m$), the ranges of aspect ratio supplied to *RAC* facilitated scaling from a thin disk all the way to a long rod, allowing a single GPR model to be computed for what are essentially different geometries. The results of this type of scaling are shown in Figure 3.13 (note that, for clarity, the upper bound of $L/\varnothing = 100$ is not shown).



(a) $L/\varnothing = 0.1$ (b) $L/\varnothing = 1$ (c) $L/\varnothing = 10$

Figure 3.13: Examples of mesh scaling during *RAC* simulations of the cylinder geometry

As cylinders feature rotational symmetry, any combination of roll, pitch, and yaw may be converted by a simple vector transform to an equivalent pitch angle at zero yaw and roll. As such, only the pitch angle, α, of the geometry in this case was varied between simulations (this rotation is visualised in Figure 3.14). Additionally, this meant that the lateral force coefficient, C_S, and the rolling and yawing moment coefficients (C_{mx} and C_{mz}, respectively) did not have to be included in the GPR models. Accordingly, only the drag, lift, and pitching moment coefficients (C_D, C_L, and C_{my}, respectively) were correlated.

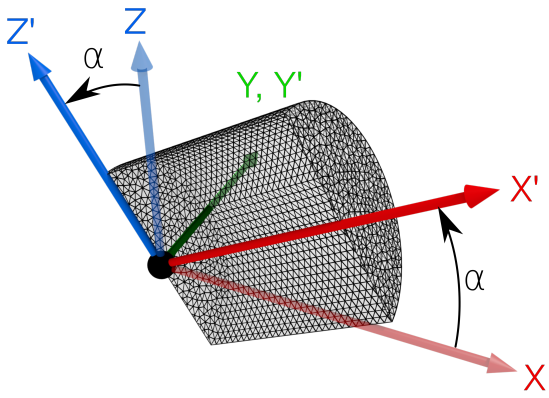


Figure 3.14: Section view of the mesh used for cylinder simulations in *RAC* showing Euler rotation about the global y axis (pitch, α) with an offset centre of mass.

The procedure for converting a known orientation (expressed in terms of Eulerian pitch, α, yaw, β, and roll, ψ, angles) to a single equivalent pitch angle, α_E, for input to the GPR models is outlined below. Firstly, the Euler rotation matrices, $\mathbf{R_{xyz}}$, are reproduced here in Eq. (3.30), as is the equation for calculating the angle between two vectors, shown here in Eq. (3.31).

Note that the term "global" here refers to the Cartesian xyz axes that remain stationary as a geometry is transformed. The term "local" refers to a body-centred set of axes which is transformed along with the geometry. An illustration of these two coordinate frames is shown in Figure 3.14, where those labelled XYZ are the global axes, and those labelled $X'Y'Z'$ are the local axes.

$$\mathbf{R_{xyz}}\left(\alpha,\beta,\psi\right) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\psi) & \sin(\psi) \\ 0 & -\sin(\psi) & \cos(\psi) \end{bmatrix} \cdot 1 \begin{bmatrix} \cos(\beta) & 0 & -\sin(\beta) \\ 0 & 1 & 0 \\ \sin(\beta) & 0 & \cos(\beta) \end{bmatrix} \cdot \begin{bmatrix} \cos(\alpha) & \sin(\alpha) & 0 \\ -\sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{3.30}$$

$$\theta_V\left(\overrightarrow{a},\overrightarrow{b}\right) = \cos^{-1}\left(\frac{\overrightarrow{a}\cdot\overrightarrow{b}}{|\overrightarrow{a}|\cdot|\overrightarrow{b}|}\right) \tag{3.31}$$

To begin, the unit vector in the global x direction ($\hat{x}$) is defined:

$$\hat{x} = [1\ 0\ 0] \tag{3.32}$$

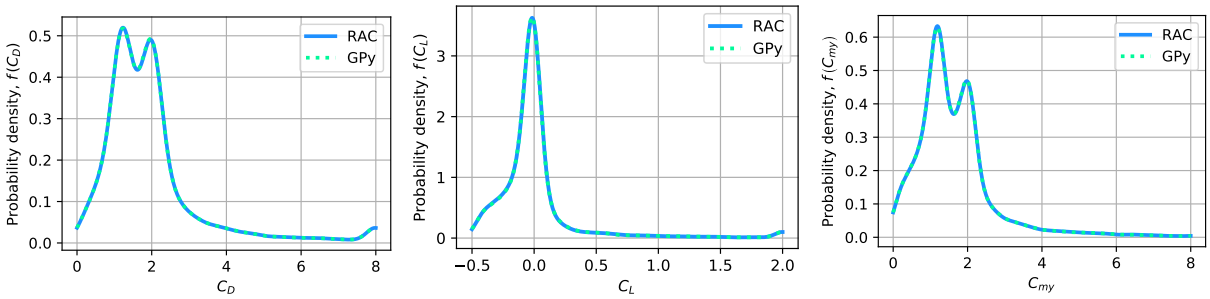
This unit vector, $\hat{x}$, is then transformed using the Euler angles, α, β, and ψ such that it becomes coincident with the geometry's local x axis:

$$\overrightarrow{x}' = \mathbf{R_{xyz}} \cdot \hat{x} \tag{3.33}$$

Finally, the angle between the transformed vector, $\overrightarrow{x}'$, and global x axis is calculated. In the present work, this angle is called the "equivalent local pitch", α_E, and is the sole attitude related variable required for axisymmetric GPR model predictions.

$$\alpha_E = \theta_V \left(\hat{x}',\ \hat{x}\right) \tag{3.34}$$

Overall, four input parameters were utilised by the GPR models (Mach number, Knudsen number, equivalent pitch angle, and aspect ratio - see Table 3.5). A total of 30,000 *RAC* simulations were run for the cylinder case (1000 different combinations of Ma, α, and $L/\varnothing$, with a sweep of 30 values of Kn each), and the predictions tested using a randomly selected set of 20,000 input parameters (which were computed within the correlation bounds, as defined in Section 3.2 and Table 3.5).



(a) C_D predictions (b) C_L predictions (c) C_{my} predictions

Figure 3.15: Plots of probability density comparing aerodynamic coefficient results produced by *RAC* simulations and subsequent Gaussian process regression models. The geometry being analysed was a cylinder of variable orientation and aspect ratio.

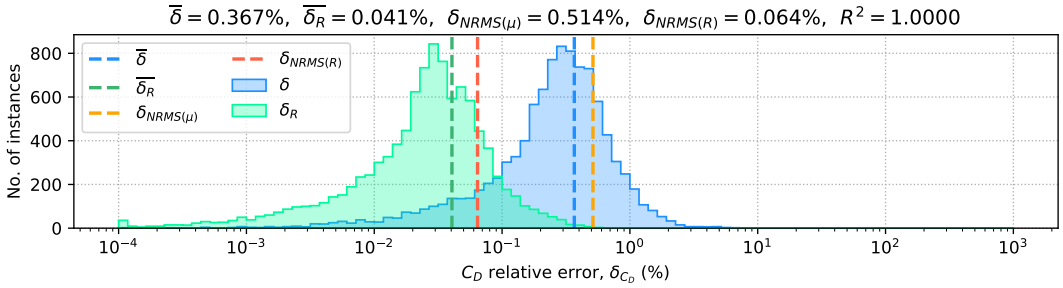
A statistical measure called "probability density" is used to analyse the results. In probability theory, a probability density function (abbreviated PDF) provides the *relative* likelihood that a given variable will equal a given value. While the *absolute* likelihood of some variable equalling a specific value is 0 (since there are an infinite number of possibilities), a PDF provides the probability that a given variable will fall in a certain range. In this fashion, PDFs behave much like histograms.

Figure 3.15 features a series of probability density plots which compare the likelihood of obtaining a certain value from *RAC* versus the resulting GPR models. These were computed based upon a test set of attitudes and freestream conditions which was input to both *RAC* and the GPR models, and which had not been used at all in the initial computation of the GPR models. These data then represent a true test of the GPR model's predictive accuracy. As can be seen, the agreement between the two sets of results is excellent, indicating that an extremely high quality correlation has been achieved.

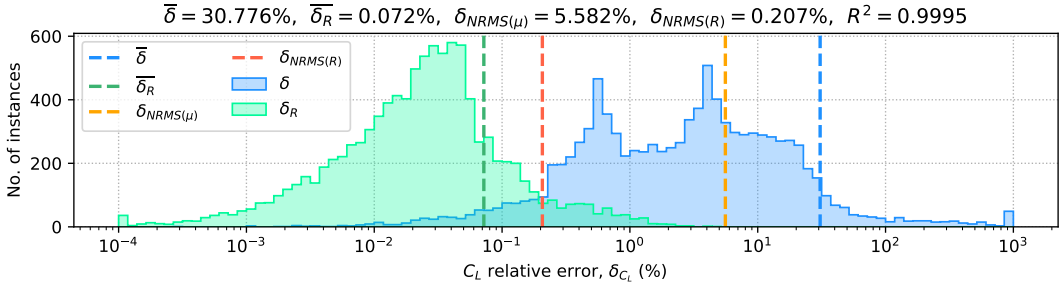
Histograms of the errors between test data and GPR predictions are presented in Figure 3.16. As discussed in Section 2.6, there are several different ways to calculate the errors between datasets, and these are utilised here to better illustrate the suitability of the GPR predictions. As expected from the probability density plots in Figure 3.15, the correlation coefficient, R^2, of the models is uniformly high, with a minimum value of 0.995 for the C_L GPR model.

It can be observed that the relative errors, δ, are lowest for the C_D and C_{my} predictions, with maxima below 10%. Errors in the C_L predictions are far higher, with a maximum of $\sim 1000\%$. Considering the range-normalised errors, δ_R, provides a much clearer perspective of the suitability of the predictions. As these error calculations take into account the magnitude of the predictions relative to the range (the difference between the maximum and minimum values) of the dataset as a whole, context may be applied to the computed errors. From these measures, it is clear that the vast majority of the errors arise from small values of C_L, which is to be expected since $C_L \to 0$ as $\alpha \to 0^o$ or $\alpha \to 90^o$ for a cylinder (or indeed any right prism). Considering the probability density plot of C_L in Figure 3.15(b), this is clearly the case, as the highest probability is centred about $C_L = 0$. These small scale inaccuracies in the predictions of aerodynamic coefficients are unlikely to severely impact a trajectory calculation utilising the GPR models as aerodynamic databases.

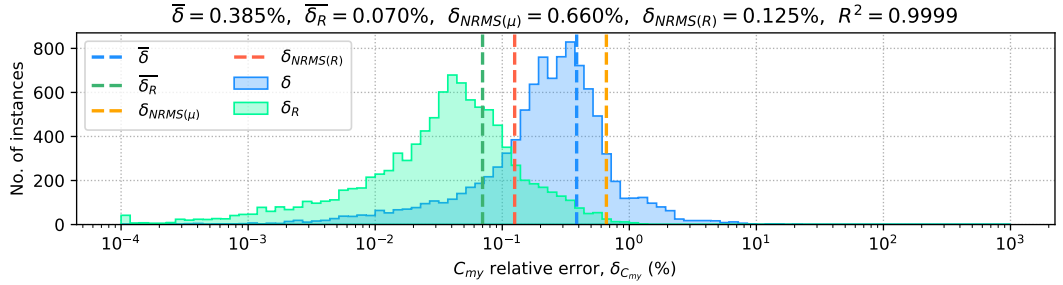
Range-normalised errors are an order of magnitude smaller for the C_D and C_{my} predictions, and four orders of magnitude smaller for the C_L predictions. The conclusion that the models are well fitted to the data is further supported by the mean-normalised and range-normalised root-mean-square errors ($\delta_{NRMS(\mu)}$ and $\delta_{NRMS(R)}$, respectively), which are included at the top of each plot in Figure 3.16. Range normalised RMS errors, $\delta_{NRMS(R)}$, are less than 1% for all



(a) C_D error



(b) C_L error



(c) C_{my} error

Figure 3.16: Histograms showing relative error magnitudes and associated metrics between aerodynamic coefficient data produced by *RAC* simulations and subsequent Gaussian process regression models. The geometry being analysed was a cylinder of variable orientation and aspect ratio.

three models, as are mean normalised RMS errors, $\delta_{NRMS(\mu)}$, for the C_D and C_{my} models. The value of $\delta_{NRMS(\mu)}$ is much higher for the C_L model, which also supports the separate conclusion that most of the errors in this model arise due to the small magnitudes of lift coefficient that are expected at small values of α (as a set comprising a large number of small values will inevitably lead to a similarly low average).

Although one would expect the values of the pitching moment coefficient, C_{my}, to have errors similar to those of the lift coefficient predictions, the aforementioned offset of the COG position resulted in most of these values being non-zero.

3.3.10.2 Cuboid

The aerodynamic study of the cuboid geometry was performed in much the same fashion as that of the cylinder described above. However, as cuboids are not rotationally symmetric, rotations cannot be simply converted to an equivalent pitch angle as is the case with the cylinder. As such, the Latin hypercube sampling plan used to generate input data for the cuboid simulations was computed using five dimensions as opposed to the three in the cylinder case. These were the three Euler angles (pitch, yaw, and roll), the aspect ratio of the cuboid, and the freestream Mach number. Table 3.6 shows a summary of the bounds used in these calculations, as well as the scaling used for each variable.

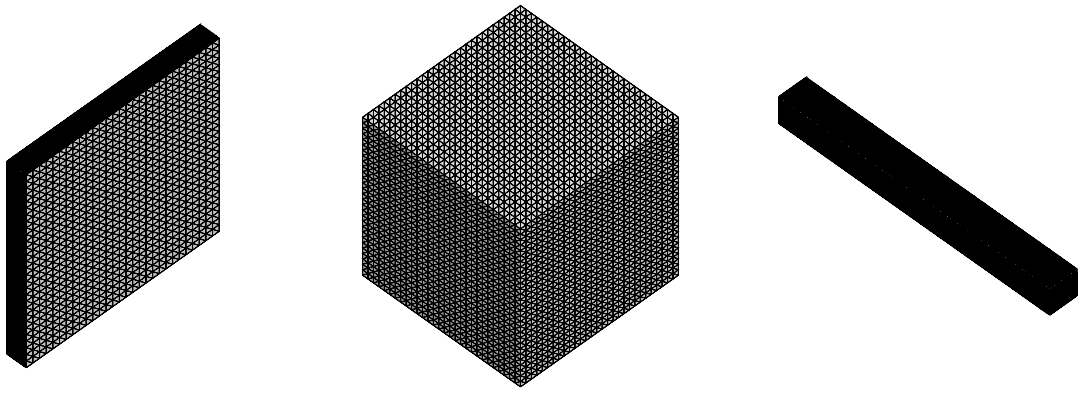
Note that the aspect ratio of a cuboid (computed as length over width), is presented here using the same nomenclature as the cylinder case i.e. $L/\varnothing$, where $\varnothing$ is the side length of the cuboid's square face.

Parameter	**Symbol**	**Lower bound**	**Upper bound**	**Scale**
Mach number	Ma	5	35	Linear
Knudsen number	Kn	1×10^{-5}	1×10^{4}	Logarithmic
Pitch angle	α	0.1^o	89.9^o	Linear
Yaw angle	β	0.1^o	89.9^o	Linear
Roll angle	ψ	0.1^o	89.9^o	Linear
Aspect ratio	$L/\varnothing$	0.1	100	Logarithmic

Table 3.6: Parameters bounds and scales used in generating input conditions for *RAC* simulations utilising the cuboid geometry.

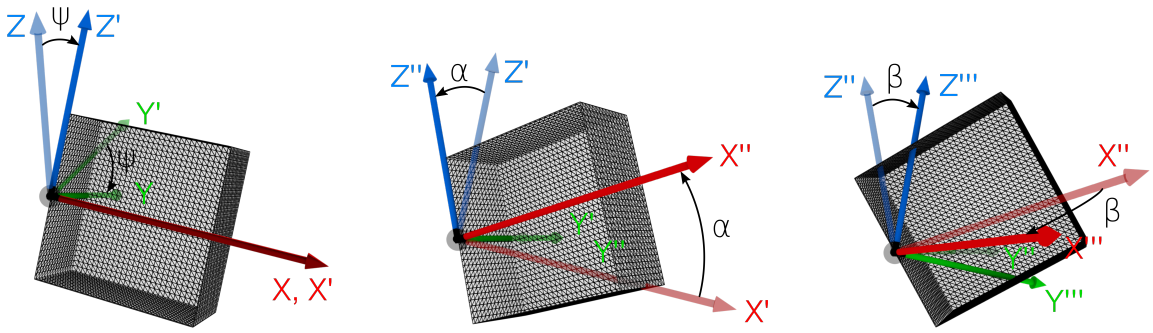
The scaling procedure operates along the global x axis of the geometry prior to any alterations of its orientation, examples of the results of which are presented in Figure 3.17. As a result of this, the ends of the scaled cuboids retain square profiles.

Figure 3.18 shows the process of rotating the cuboid model in order to alter its orientation relative to the freestream. *RAC* utilises xyz order rotation, wherein vertices are rotated using Euler rotation matrices first about the global x axis (see Figure 3.18(a)), then around the global y axis (see Figure 3.18(b)), and finally about the global z axes (see Figure 3.18(c)).



(a) $L/\varnothing = 0.1$ (b) $L/\varnothing = 1$ (c) $L/\varnothing = 10$

Figure 3.17: Examples of mesh scaling during *RAC* simulations of the cuboid geometry



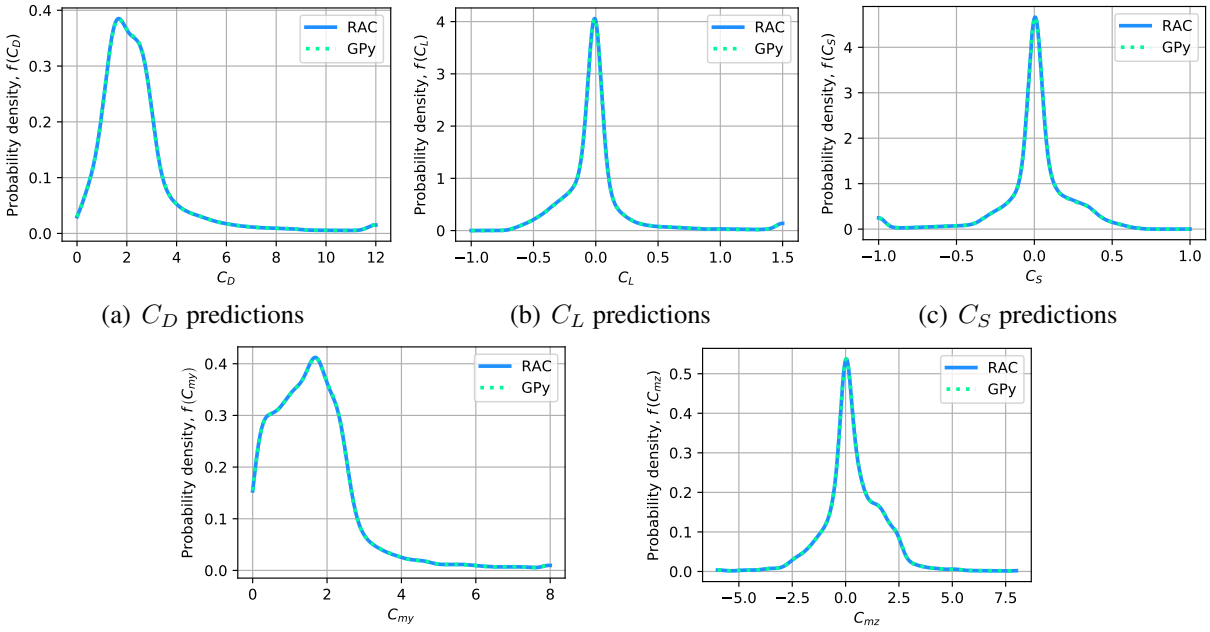
(a) Rotation about global x axis (b) Rotation about global y axis (c) Rotation about global z axis

Figure 3.18: Section views of the mesh used for cuboid simulations in *RAC* showing staged Euler rotation about the global x, y, and z axes (roll, ψ; pitch, α; yaw, β, respectively) with an offset centre of mass.

A total of 90,000 *RAC* simulations were run for the cuboid case (3000 different combinations of Ma, α, β, ψ, and $L/\varnothing$, with a sweep of 30 values of Kn each), and the predictions tested using a randomly selected set of input parameters (which, like those of the cylinder case, were computed within the correlation bounds, as defined in Section 3.2 and Table 3.6). While generating the GPR models based upon the results, it quickly became apparent that the rolling moment coefficient, C_{mx}, did not vary significantly, its value remaining very close to zero throughout virtually every computation. This was an expected side effect of the selected COG position (the middle of one of the geometry's faces, as opposed to its geometric centre see Figure 3.18), as well as the constant square cross-section of the cuboid (see Figure 3.17). Hence, the moments acting about the cuboid's longitudinal axis were minimised.

Accordingly, C_{mx} was omitted from the correlation exercise, and five models in total

produced; one for each remaining aerodynamic coefficient (C_D, C_L, C_S, C_{my}, and C_{mz}). Like the cylinder simulations, each model featured 4 GPR "sub-models", which were used to correlate the parameters of a Gauss-sigmoid curve fitted to the trend of a given aerodynamic coefficient with Knudsen number (see Section 3.3.9).

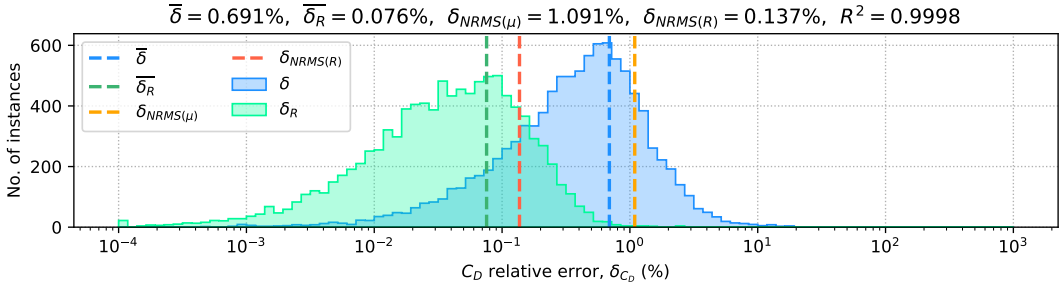


(a) C_D predictions (b) C_L predictions (c) C_S predictions

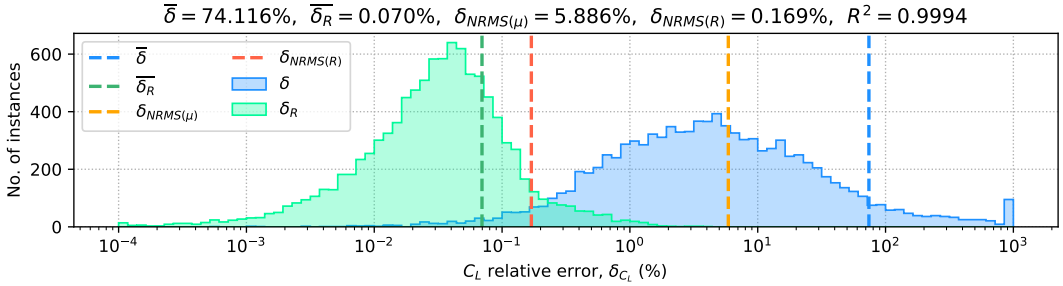
(d) C_{my} predictions (e) C_{mz} predictions

Figure 3.19: Plots of probability density comparing aerodynamic coefficient results produced by *RAC* simulations and subsequent Gaussian process regression models. The geometry being analysed was a cuboid of variable aspect ratio.

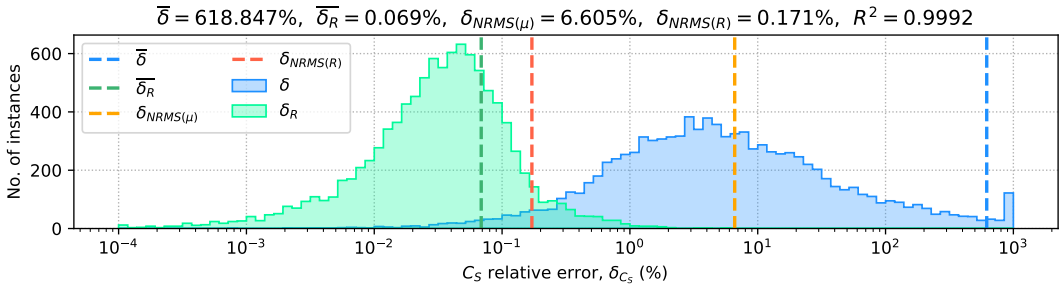
Probability density plots of results from the cuboid GPR models compared to those obtained directly from *RAC* are presented in Figure 3.19. Much like the cylinder case, all of the GPR models show excellent correlation with the results of the test data set (whose values were not included in the training of any of the GPR models). Once again, the most likely C_L value is 0, and this is mirrored in the lateral force coefficient, C_S. The pitching and yawing moment coefficients (C_{my} and C_{mz}, respectively), however, did not show this similarity. This was caused by the coordinate system transform performed on the force coefficients (C_D, C_L, and C_S), which is described in Section 3.3.6. While the force coefficients are transformed into the body frame of the geometry being analysed, the moment coefficients are allowed to remain in the global coordinate system, hence resulting in the aforementioned disparity in Figure 3.19.



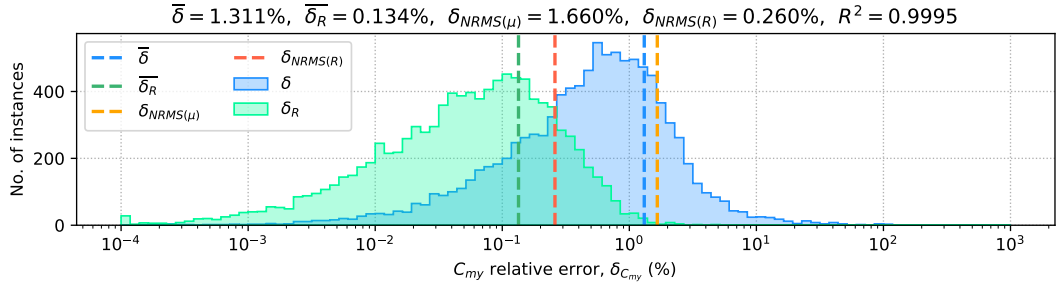
(a) C_D error



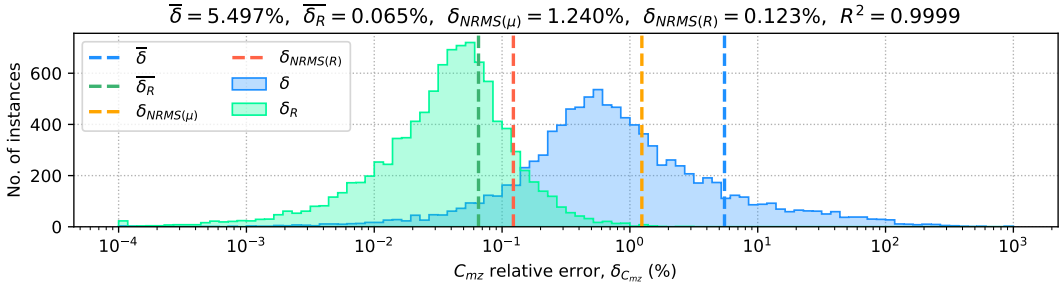
(b) C_L error



(c) C_S error



(d) C_{my} error



(e) C_{mz} error

Figure 3.20: Histograms showing relative error magnitudes and associated metrics between aerodynamic coefficient data produced by *RAC* simulations and subsequent Gaussian process regression models. The geometry being analysed was a cuboid of variable aspect ratio.

Histograms of the errors between test data and GPR predictions for the cuboid geometry are presented in Figure 3.20. The correlation coefficients are once again very high, lying in the range $0.9992 \leqslant R^2 \leqslant 0.999$ for all GPR models. Relative errors are highest for the C_S and C_L predictions, with maxima above 1000% and mean values of 618.85% and 74.12%, respectively. Range normalised errors, δ_R, are far lower, however, as was the case with the cylinder models, with a maximum of 0.134% for the C_{my} model and a minimum of 0.065% for the C_{mz} model. Range normalised RMS errors, $\delta_{NRMS(R)}$ are similarly low, with a maximum of 0.26% for the C_{my} model, and a minimum of 0.123% for the C_{mz} model. Mean normalised RMS errors, $\delta_{NRMS(\mu)}$, were slightly higher than the cylinder case, with a maximum of 6.605% for the C_S model, and a minimum of 1.091% for the C_D model. This is most likely due to the increased complexity of the models, each of which features five correlation dimensions as opposed to the cylinder case's three.

3.4 DSMC MODELLING

Thermal simulations were performed using the *DS2V* and *dsmcFoam+* codes described in Section 2.3.5 on two different cylindrical geometries. The first, a hemispherically capped cylinder, was intended to be analogous to the windward surface of a regular sphere, and will be referred to as "the hemisphere" from here onward. The hemisphere was analysed using *DS2V* at a single angle of incidence and two Knudsen numbers, totalling 2 simulations. The second geometry, a flat cylinder (referred to as "the cylinder" from here onward), was analysed using both *DS2V* (at zero angle of incidence, $\alpha = 0^o$) and *dsmcFoam+* at several angles of incidence and Knudsen numbers, totalling 72 simulations. This gives an overall total of 76 simulations for the present work.

For each different set of freestream conditions and attitudes, one simulation was run using an isothermal wall boundary condition, while another was run using an adiabatic boundary (see Section 2.3.4). The results from these two types of simulations were intended to be used in the calculation of convective heat transfer coefficients, h_{conv}.

Of the cylinder simulations, 32 were performed at conditions achievable in the LDWT (these conditions are elaborated upon in Section 4.3.6). The first 16 utilised freestream conditions

equivalent to $80\,mm$ downstream of the nozzle exit, giving a Knudsen number of 0.02. The next set of 16 utilised the freestream conditions equivalent to $108\,mm$ downstream of the nozzle exit, giving a Knudsen number of 0.1, which is approximately one order of magnitude higher than at the previous location. The selection of these locations in the LDWT plume is further elaborated upon in Section 5.1.

The final 40 simulations were performed using isothermal boundaries at various Knudsen numbers in the range $1 < Kn < 10$, with "real scale" freestream conditions which were computed using a Latin hypercube sampling plan. The conditions in this LHS plan were computed with Cantera and the NRL atmospheric model described in Section 3.1 using the bounds of Mach and Knudsen number computed by the trajectory modelling in Section 3.2. Note that the values $Kn = 1$ and $Kn = 10$ were used as limits for the LHS such that the DSMC simulations in this range featured Knudsen numbers in the range $1 < Kn < 10$. This was done in order to better test the capabilities of the curve-augmented GPR models to both interpolate and extrapolate thermal fluxes using sparse input data.

The angles of incidence, α, for the *dsmcFoam+* simulations performed at LDWT conditions were varied in 15^o increments in the range $0^o \leqslant \alpha \leqslant 90^o$, while the angles of incidence for the "real scale" *dsmcFoam+* simulations were selected by the LHS routine in the range $0^o < \alpha < 90^o$. The *DS2V* simulations were performed at these same freestream conditions but at zero angle of incidence ($\alpha = 0^o$), owing to the code's nature as a 2D solver, i.e. only simulations that are either planar or axially symmetric can be modelled. These simulations were performed in order to gather reference data for the *dsmcFoam+* cases; this procedure is elaborated upon in Section 3.4.4

A summary of the simulations performed, as well as the geometries, freestream angles of incidence and Knudsen numbers utilised during setup are presented in Table 3.7.

Kn	**Hemisphere**		**Cylinder**	
	DS2V	***dsmcFoam+***	***DS2V***	***dsmcFoam+***
0.02	$\alpha = 0^o$ [1×2]	$-$ [0]	$\alpha = 0^o$ [1×2]	$\alpha = \{0^o, 15^o, \ldots, 90^o\}$ [7×2]
0.1	$\alpha = 0^o$ [1×2]	$-$ [0]	$\alpha = 0^o$ [1×2]	$\alpha = \{0^o, 15^o, \ldots, 90^o\}$ [7×2]
$1 - 10$	$-$ [0]	$-$ [0]	$\alpha = 0^o$ [20]	$0^o < \alpha < 90^o$ [20]

Table 3.7: Summary of DSMC simulations performed during the present work. The number of simulations run using a given code, geometry and Knudsen number are presented in square brackets. The angles of incidence (or ranges thereof) at which the simulations were performed are also shown. Note that the nomenclature [1×2] and [7×2] is used to signify the two different wall boundary types (mentioned above) run at the same freestream conditions and attitudes.

In each simulation, the geometry wall was set to a fixed temperature of $293\ K$, and heat fluxes calculated directly by the DSMC codes. Non-reactive chemistry models were utilised for all simulations, including those at high Knudsen and Mach numbers, where chemical reactions are known to play a non-negligible role in aerothermal heating. This type of chemistry was selected due to the lack of a finite-rate catalytic wall boundary condition in the most recent version of *dsmcFoam+*. Without a wall condition of this kind, calculations utilising the code's quantum kinetic (Q-K) chemistry model would produce higher heat fluxes at the wall boundary (as a result of molecular dissociation reactions in the shock layer that inevitably forms in front of an object entering a planetary atmosphere). In reality, (assuming a sufficiently catalytic surface) these thermal fluxes would be partially mitigated by the recombination of monatomic species at the wall, e.g. $N + N \rightarrow N_2,\ O + O \rightarrow O_2,\ N + O \rightarrow NO$, etc. However, with the aforementioned lack of a boundary condition capable of modelling this behaviour, it was concluded that the results obtained from subsequent simulations would be more accurate overall if chemistry induced heating were disabled outright, as opposed to considering only dissociation [128]. Also, since most DSMC simulations performed for the present work were based upon LDWT conditions, no chemical reactions were expected (the LDWT is a low-enthalpy experimental facility - see Chapter 4).

Although *DS2V* includes both of the features described above, thermochemistry was also disabled in all simulations utilising this code for the sake of consistency. This facilitated direct comparisons between the two codes' results for certain cases, and also allowed data from both

codes to be used in the same analyses.

The *dsmcFoam+* simulations were monitored by an automated script for steady state convergence (i.e. the point at which the rate of change of linear kinetic energy in the flow reaches zero) before being switched automatically to their averaging mode. They were then allowed to run for 96 hours each before being reconstructed and post-processed. This averaging period was selected because it allowed stochastic noise in both heat flux and adiabatic wall temperature distributions to be reasonably well mitigated. *DS2V* incorporates a fully automatic version of this monitoring process by default, and as such did not require any external intervention during the convergence phase of simulations.

3.4.1 Sampling locations

Several sampling locations were selected on those geometries analysed during the present study in order to facilitate the extraction of heat flux data from regions of interest.

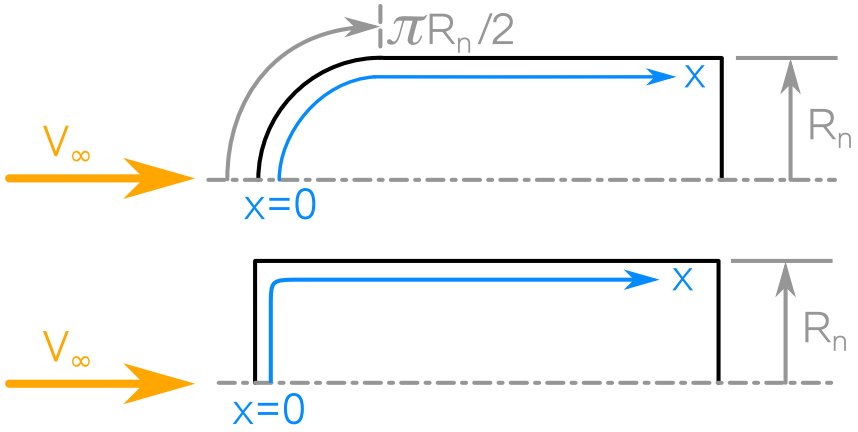


Figure 3.21: Sampling convention for 2D axisymmetric DSMC simulations

For those geometries which were simulated using *DS2V* (using 2D axisymmetric representations), the convention for data sampling is presented in Figure 3.21. This convention utilises a single unified measurement (termed "arc length") to represent the distance from the stagnation point at zero incidence ($\alpha = 0^o$). In cylinder cases, the arc length is normalised to the geometry's radius, R, while in hemisphere cases, it is normalised to the perimeter of the 2D geometry's circular section, $\pi R/2$.

Sampling data on the 3D geometries used in *dsmcFoam+* simulations was accomplished in a similar fashion. As the heat fluxes at the edges of the cylinder geometry are of particular interest

to the present work, several groups of cells on the meshes used were identified as sampling locations, and the data from the *dsmcFoam+* simulations extracted during post-processing. These locations are shown in Figure 3.22, as are the reference points and directions for the angle of incidence, α, and the surface radial coordinate, θ.

The "front edge" is the outermost ring of cells on the front circular face of the cylinder, and the "side edge" is the forwardmost ring of cells on the curved side face. The "stagnation line" begins at the centre of the front-facing circular face and continues around the edge of the cylinder, extending downstream toward the rear of the cylinder (it is split into two sections in order to make plotting of the extracted data more convenient). The "stagnation point" refers to the centre of the front circular face, but only for those simulations which were performed at an angle of incidence, α, of 0^o.

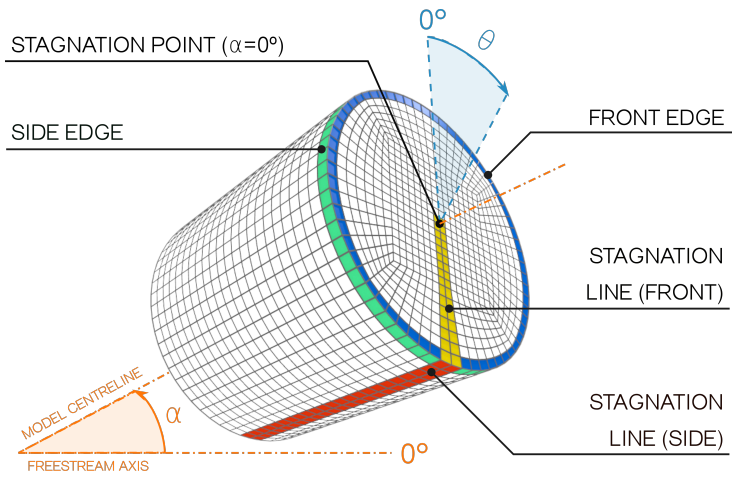


Figure 3.22: Heat flux sampling locations on the cylinder surface mesh used in DSMC simulations. Note that the cell layout shown is representative of meshes used during actual simulations; however the element sizes have been exaggerated for illustrative purposes. The freestream vector in each case is coincident with the "freestream axis" located at $\alpha = 0^o$.

Data sampling from the *DS2V* simulations was accomplished via the use of a built-in utility, while the results of the *dsmcFoam+* simulations were sampled using a series of macros and programmable filters set up using ParaView [129].

References to "arc length" in subsequent figures utilise a similar convention to that used for the *DS2V* results, i.e. a line that begins at the stagnation point and continues around the corner of the cylinder, following the stagnation line shown in Figure 3.22. Similarly, references to "radial

coordinate" refer to an angle measured clockwise from the 12 o'clock position of the cylinder (θ in Figure 3.22).

3.4.2 DSMC results scaling

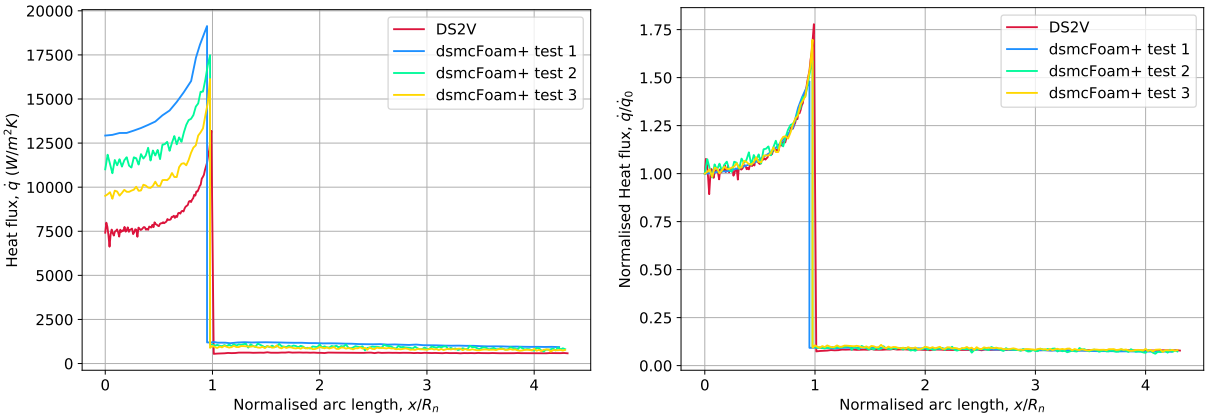
During the initial setup of the DSMC cases necessary for the present study, it was noticed that the normalised trends of surface heat flux did not vary significantly when non-optimal simulation parameters were utilised. The ideal values for these parameters vary depending on the freestream conditions, and are described in more detail on Section 2.3.4.

Simulation	**Mesh cell size**	**Time step**	**DSMC particle weighting factor**
DS2V	Optimal	Optimal	Optimal
dsmcFoam+ test 1	1.0×10^{-4}	5×10^{-7}	1.8×10^{10}
dsmcFoam+ test 2	3.33×10^{-5}	5×10^{-8}	1.0×10^{10}
dsmcFoam+ test 3	3.33×10^{-5}	5×10^{-8}	1.0×10^{9}

Table 3.8: Details of DSMC simulations used in heat flux scaling study

Table 3.8 shows a summary of the parameters which were used in three test simulations on an axially-aligned cylindrical geometry ($\alpha = 0^o$) performed using *dsmcFoam+*. A further test simulation was run using *DS2V* which, as mentioned in Section 2.3.5, is capable of automatically selecting the optimal values for mesh cell sizing, time step, DSMC particle weighting factor, and simulated particle count. The freestream conditions and geometry were consistent across all of the simulations, with the exception of the *DS2V* case, where the geometry was converted to a 2D axisymmetric form (as opposed to the *dsmcFoam+* simulations, which all utilised the same basic 3D geometry).

As can be observed in Figure 3.23(a), the heat fluxes vary in magnitude with different combinations of non-optimal simulation parameters. This behaviour is expected, as is the improvement in accuracy (compared to the *DS2V* results) with finer parameters. For example, *dsmcFoam+* test 3 utilises the smallest values of time step, mesh cell size, and DSMC particle weighting factor, and is hence the closest to the optimal solution produced by *DS2V*. The opposite is true of *dsmcFoam+* test 1, which utilises the largest values of time step, mesh cell size, and



(a) Heat flux, $\dot{q}$

(b) Heat flux normalised to stagnation value, $\dot{q}/\dot{q}_0$

Figure 3.23: Comparison of heat flux results from DSMC simulations using various non-optimal parameters. Results are sampled over the corner of a cylinder aligned with the freestream direction ($\alpha = 0^o$, $\beta = 0^o$). Both full scale results and values normalised to the local stagnation point heat flux are presented.

weighting factor.

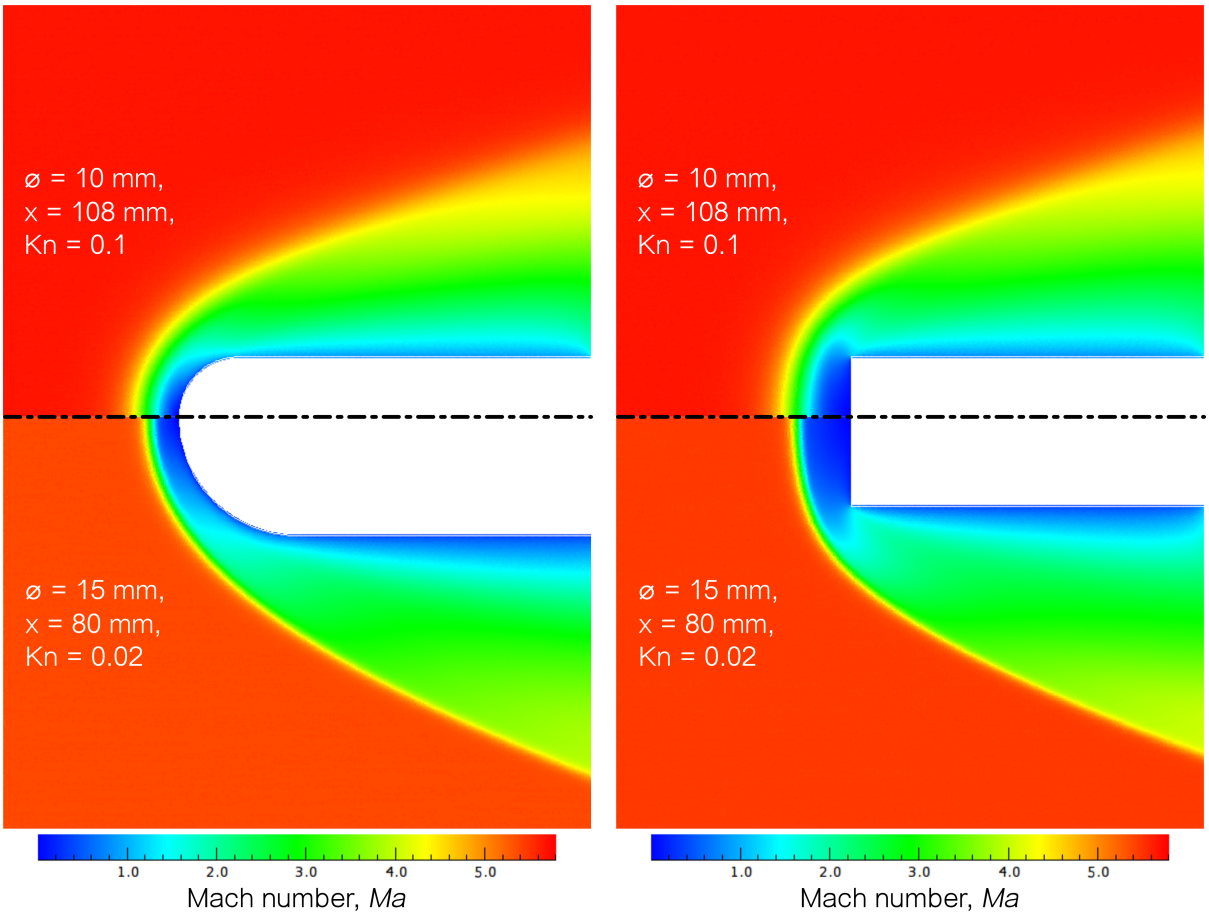
These data show that non-optimal DSMC parameters are unsuitable for determining surface heat fluxes; however, when normalised to their local stagnation values (which occur at $x/R = 0$ for cylinder cases), all four of the test simulations produce similar results. This behaviour, which may be observed in Figure 3.23(b), indicates that the normalised trends of surface heat flux may be determined reliably using DSMC simulation even when optimal parameters for said simulations are not achieved. As running a DSMC simulation is typically very resource intensive, this conclusion is invaluable for low Knudsen number simulations where the required number of simulated particles would ordinarily require the use of many CPU-hours on a large cluster computer. Instead, smaller scale simulations utilising non-optimal parameters could be used to determine normalised thermal trends, which could then be multiplied by a reference stagnation heat flux value (generated by the likes of the Fay-Riddell or Brandis-Johnston expressions - see Section 2.3.3) in order to obtain a full scale heat flux distribution over the entire surface of a geometry.

3.4.3 *DS2V* simulations

As is mentioned in previous sections, axisymmetric *DS2V* simulations were performed for the hemisphere and cylinder geometries at various Knudsen numbers and zero angle of incidence

($\alpha = 0^o$). Contour plots of the the simulated flowfield's Mach number are presented in Figure 3.24 for both geometries at the LDWT-derived conditions, i.e. $Kn = 0.02$ and $Kn = 0.1$.

Since *DS2V* automatically calculates the optimal DSMC parameters for a given simulation, these results are extremely useful to the present work, as they allow both heat flux magnitudes and normalised trends to be assessed at zero angle of incidence, $\alpha = 0^o$ (the reference orientation for normalised results produced by all subsequent *dsmcFoam+* simulations).



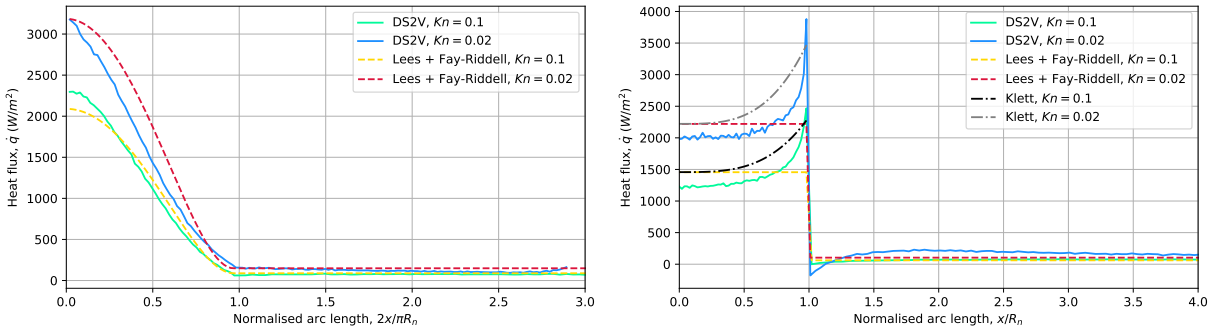
(a) Hemisphere (b) Cylinder

Figure 3.24: Contour plots of Mach number simulated at LDWT conditions using DS2V. Results from simulations featuring cylinders with flat and hemisphere caps are shown. The upper and lower parts of the image show results from different flow conditions (labelled).

In Figure 3.24, the lower half of each subfigure shows the simulated flowfield for a Knudsen number of 0.02, while the upper portions show the flowfield of those simulations with a Knudsen number of 0.1. The departure from continuum regime behaviour with increasing Knudsen number can clearly be seen, with the low Kn cases featuring a well defined bow shock while those at the higher value of Kn feature a far more diffuse shock with a larger standoff distance

(despite the geometry itself being smaller). Additionally, the thickness of the boundary layer at the surface of the examined geometries is noticeably thinner at higher Knudsen numbers.

These effects are consequences of increasing levels of rarefaction in the flow, whereby a lower freestream density results in both a less compact step at the shock boundary, and fewer gas molecules to decelerate at the wall and form a boundary layer.



(a) Hemisphere (b) Cylinder

Figure 3.25: DS2V heat flux results for the hemisphere and cylinder simulations. Predictions of various correlations are overlaid.

Heat fluxes along the simulated surfaces are presented in Figure 3.25, and are plotted according to the convention described in Section 3.4.1 (specifically, Figure 3.21). Predictions of these quantities produced using various correlations are also plotted for comparison. The heat flux distribution predictions are calculated using the equations of Lees (Eq. (3.27)), Schaaf and Chambre (Eq. (3.28)) and Klett (Eq. (3.35)), in the case of the cylinder only), with the relevant stagnation heat fluxes being calculated using the method of Fay and Riddell (Eq. (2.26)).

Although other correlations exist for stagnation point transfer, Fay-Riddell was selected in this case because it was derived from first principles as opposed to being based upon empirical data or simulation results (see Section 2.3.3). As such, it is far more applicable to experimental data derived from experiments run at LDWT conditions than the likes of the Brandis-Johnston or Detra-Hidalgo correlations.

$$\left(\frac{\dot{q}}{\dot{q}_0}\right)_{cont} = 1 + 0.6\left(\frac{x}{R_n}\right)^{3.3} \left\{\frac{x}{R_n} \in \mathbb{R} \,|\, 0 \leqslant \frac{x}{R_n} \leqslant 1\right\} \tag{3.35}$$

It was assumed that, due to the low enthalpy nature of the flow being simulated, the boundary layers would be in equilibrium. In this type of flow, chemical reactions (if present at all) occur

so rapidly that chemical equilibrium exists at every point in the flow, and hence these reactions' influence on overall heat fluxes is negligible. As such, the wall/boundary layer variable, B, in the Fay-Riddell equation was substituted with the appropriate version of Eq. (2.28), and the geometry scaling factor, A, changed accordingly for each case (see Eq. (2.27)). It should be noted that the cylinder version of the Fay-Riddell equation (where A is taken to be 0.53) is formulated for a cross-flow scenario, as opposed to the "head-on" orientation simulated here, hence the higher errors reported below (see Table 3.9).

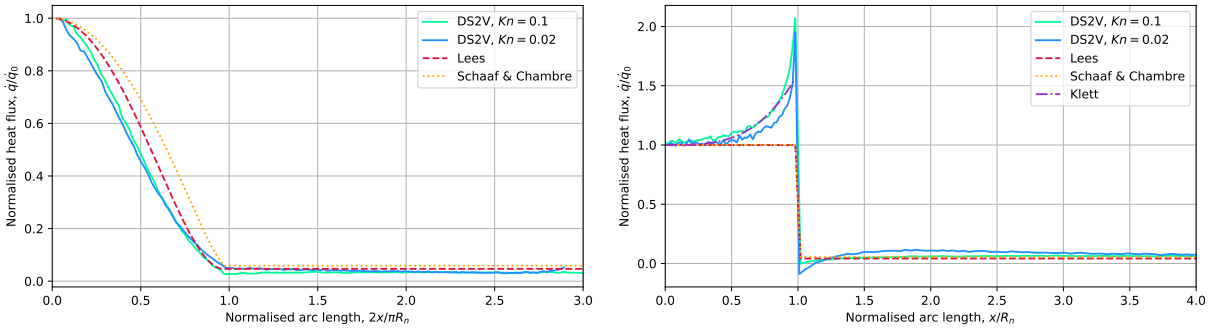
As one would expect, the low Knudsen number cases show lower overall heat transfer rates than those at higher values of Kn. This is due to a higher level of rarefaction in the flow and a smaller probe diameter in high Kn cases, as well as a small difference in freestream total temperature between simulations. Since total temperature magnitudes at the centreline are reduced as the test plume expands downstream of the nozzle exit, the thermal energy of the flow is spread out over a wider area. This results in there being both less thermal energy per unit volume of gas, and less gas per unit volume of space (lower density) in the flow further downstream of the nozzle exit. This then causes there to be less energy in the high Kn which may encounter the surface of the test geometries; hence the difference in heat flux results mentioned above (see Section 2.2.4). It must also be noted that the freestream conditions for these cases are based upon CFD simulations of the flow in the LDWT (see Section 4.3.4.4), which allow the gas conditions to be derived at any point in the flow.

Geometry	***Kn***	$\dot{q}_0$ ***(DS2V)***	$\dot{q}_0$ **(FR)**	$\dot{q}_0$ **(SC)**	δ **(FR)**	δ **(SC)**
Hemisphere	0.1	2296	2148	391	6.4%	224.8%
Hemisphere	0.02	3176	3272	1294	2.9%	573.9%
Cylinder	0.1	1228	1498	6989	18.0%	468.9%
Cylinder	0.02	1980	2282	21985	13.2%	1010.4%

Table 3.9: Stagnation heat flux results for *DS2V* cases. Both simulation results and correlation predictions are shown, along with relative percentage error calculations. The notations "FR" and "SC" refer to the Fay & Riddell and Schaaf & Chambre correlations, respectively. All heat fluxes are presented in W/m^2.

Good agreement between the simulation data and predictions of Lees' correlation can be observed in all four cases, with a maximum stagnation heat flux error of 18.0% for the cylinder

case at $Kn = 0.1$, and a minimum of 2.9% for the hemisphere case at $Kn = 0.02$. A summary of these results is presented in Table 3.9. It is also evident from these data that, although the stagnation heat fluxes agree reasonably well, the distribution of these values over the simulated geometries are not nearly as well predicted.



(a) Hemisphere (b) Cylinder

Figure 3.26: DS2V heat flux results for the hemisphere and cylinder simulations. Predictions of various correlations are overlaid. All values are normalised to the heat flux at each geometry's stagnation points.

The heat flux predictions produced by the correlation of Schaaf and Chambre are not shown in Figure 3.25 as they were found to be highly inaccurate for the present cases (see Table 3.9). This is due to these expressions having been derived based upon the assumptions of the free molecular regime – assumptions which are therefore invalid for flow that is incident upon the windward surface(s) of experimental probes in the LDWT (free molecular flow only begins to manifest when $Kn \geqslant 10$, a level of rarefaction one order of magnitude above the experiments presented herein).

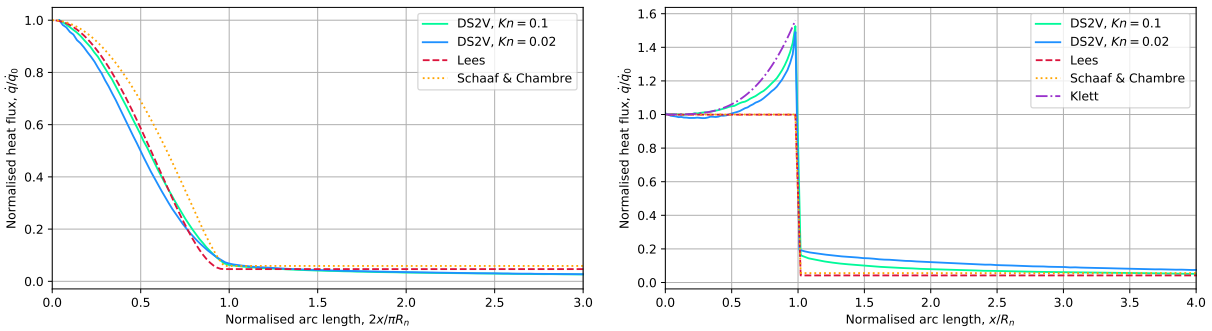
In order to better compare the heat flux distributions produced by DSMC simulations and correlation predictions, the results were normalised to their respective stagnation heat fluxes, $\dot{q}_0$. These data are presented in Figure 3.26, and include the normalised predictions produced by the correlations of Lees, Schaaf & Chambre, and Klett (once again, for the cylinder cases only) for comparison. Note that the Knudsen numbers of these predictions are not stated due to the aforementioned normalisation – the distribution of $\dot{q}$ with respect to surface inclination angle is independent of Kn with such correlations and, as such, only full magnitude heat flux predictions may be influenced by changes in freestream conditions.

Considering the hemisphere cases first (see Figure 3.26(a)), it is clear that the DSMC results

produce a noticeably sharper increase in $\dot{q}$ near the stagnation point when compared to the correlations of both Lees, and Schaaf and Chambre, and a significantly gentler slope near $2x/\pi R_n = 1$ (the point at which the hemispherical cap meets the cylindrical body of the probe). This discrepancy comes about as a result of the assumptions on which the correlations are based. In the case of Lees [126], it is assumed that the ratio between the wall and flow stagnation temperatures will be very small i.e. $T_w/T_0 \ll 1$. In the simulated cases, however, the minimum value of this ratio is ~ 0.78, and as such the limiting case which Lees considered is not fully reproduced. Similarly, the expression developed by Schaaf and Chambre is based upon the analytical examination of free molecular flow wherein the Knudsen number is high enough that the flow itself may be treated as Newtonian. Since the Knudsen numbers of the simulated flows are significantly lower than those of a free molecular flow ($Kn = 0.02$ and $Kn = 0.1$, as opposed to the $Kn > 10$ for free molecular flows), the behaviour of the gas at the hemisphere's surface is not well predicted.

Normalised heat fluxes from the cylinder cases are presented in Figure 3.26(b), and it is immediately apparent that these data feature significantly larger discrepancies between simulation results and correlation predictions than the hemisphere cases. This is due once again to the simplified formulation of the correlations employed, specifically their dependence on a surface's angle relative to the freestream vector, δ. The correlations do not feature any terms which allow the shape of an object being analysed to be taken into account, and as such predict heat fluxes on a point-by-point basis. As a result of this, the cylinder cases — which essentially feature only two unique values of δ — are modelled by the correlations as having constant value of $\dot{q}$ across the entirety of the front circular face, and curved side face.

The correlation of Klett [51], which is specifically formulated for an axially aligned cylinder, shows much better agreement with the DSMC simulation results, especially for the $Kn = 0.1$ case. However, the peak magnitudes of heat flux which manifest toward the corner of the cylinder (where $x/R = 1$), are underpredicted by $\sim 25\%$.



(a) Hemisphere (b) Cylinder

Figure 3.27: DS2V heat flux results for hemisphere and cylinder simulations at $Ma \approxeq 30$. Predictions of various correlations are overlaid. All values are normalised to the heat flux at each geometry's stagnation points.

As the cases considered in this section use freestream conditions derived from the LDWT, the Mach numbers being simulated are only slightly higher than the typical threshold for hypersonic flows ($Ma > 5$). In order to characterise the level of Mach number dependence present in these analyses, a further set of simulations was performed wherein the Mach numbers were increased to ~ 30. These high Mach cases feature the same geometries and freestream conditions as those described above; only the flow velocities have been altered. The heat flux results of these simulations were subsequently normalised to the appropriate stagnation values, and are presented in Figure 3.27.

Agreement between the correlations and simulation results is slightly improved, especially in the case of the $Kn = 0.1$ hemisphere run, which shows good agreement with the predictions of Lees' method up to a normalised arc length of ~ 0.7. Similarly, the cylinder cases show improved agreement with Klett's method, with the peak normalised heat flux at the corner of the cylinder almost exactly matching the upper limit of $1.6\,\dot{q}_0$. This behaviour is to be expected as Klett based his correlations on the limiting case of an arbitrarily high Mach number in the continuum regime. As such, raising the Mach number of the simulations only reproduces part of the condition that Klett considered, and also accounts for the deviation that is still visible between the correlations and simulations in the range $0 < x/R_n < 1$.

From these data, it may be concluded that the correlations of Fay and Riddell, Lees, and Klett agree reasonably well with simulation data, and are applicable in the lower portions of the transition regime (with the notable exception of Lees theory when applied to sharp-cornered

bodies such as the cylinder case). As one would expect, the expression of Schaaf and Chambre is not applicable in this regime, and generally shows poor agreement with simulation results. It may also be concluded that the freestream conditions which the LDWT can replicate have a small degree of Mach number dependence which, for the purposes of the present study, may be considered negligible.

3.4.4 *dsmcFoam+* simulations

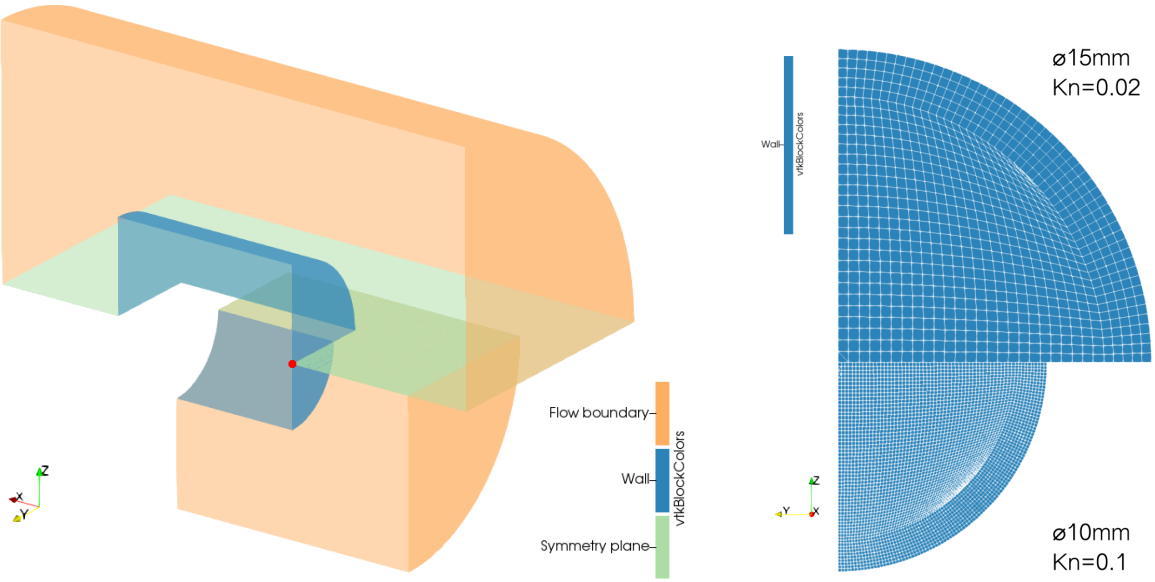
Simulations performed using the *dsmcFoam+* code focussed entirely on the cylinder geometry, and utilised a symmetric 3D mesh such that variations in angle of incidence, α, could be introduced.

Figure 3.28 shows section views of the computational domains used in the 14 aforementioned LDWT derived simulations. The smaller of the two was used for those cases with $Kn = 0.1$, while the larger was used for cases with $Kn = 0.02$. The flow boundaries (shown in orange in Figure 3.28(a)) were sized such that the shocks produced during the course of the simulations would be entirely contained within the computational domain, thereby mitigating numerical anomalies in the DSMC solution. The symmetry planes (shown in green in Figure 3.28(a)) were introduced into the meshes in order to reduce the computational expense of the simulations, since only variations in angle of incidence, α, were to be examined, and so only half of the cylinder needed to be modelled.

Note that the domain used for the $Kn = 0.1$ cases does not include the region behind the cylinder. This portion of the domain was omitted in order to reduce the computational expense required to run the relevant simulations (which utilised optimal DSMC parameters and were, as such, more expensive than the $Kn = 0.02$ cases). This is also why no comparisons were made between results of *RAC* and DSMC simulations; with no modelling of the flow in the wake region of the cylinder for the optimal DSMC cases, any aerodynamic coefficients derived from the results would be inaccurate.

The density and layout of the respective surface meshes is shown in Figure 3.28(b). For those simulations that were performed at $Kn = 0.1$, the mesh density — as well as the DSMC variables described in Section 2.3.4 — were defined such that the simulations would be optimal,

and hence feature accurate predictions of the full-scale heat fluxes. For those cases run at $Kn = 0.02$, however, the simulation variables and mesh were deliberately defined such that they would be non-optimal, relying instead upon the scaling procedure for normalised heat flux trends described in Section 3.4.2. This allowed these simulations to be performed in a much less computationally expensive fashion than with optimised variables, thereby avoiding the use of a cluster computer (which would have been a necessity had optimised simulation variables and an appropriately dense mesh been used). Since these simulations were primarily focussed on determining normalised heat flux distributions, the use of the aforementioned scaling procedure was not expected to impact the accuracy of the results.



(a) Structure of mesh domains showing bondary types.

(b) Cylinder mesh element detail.

Figure 3.28: Section views of the computational domains and surface meshes used in those *dsmcFoam+* simulations based upon LDWT derived conditions. In each subfigure, the upper domain was utilised in $Kn = 0.02$ cases, and the lower in $Kn = 0.1$ cases. The origin of both domains, $x, y, z = (0, 0, 0)$, is marked by the red dot.

For the final set of simulations (those whose freestream Knudsen numbers lie in the range $1 < Kn < 10$), a scaling function was implemented in order to optimise the mesh density for each individual case, and hence ensure that optimal DSMC simulation variables were used for each simulation. Unlike the previous set, which only examined two different Knudsen numbers at a range of angles of incidence, α, these high Kn cases each featured a unique Knudsen number and angle of incidence. As such, the study of normalised heat flux distributions would not have

been possible without reference simulations preformed at $\alpha = 0$. These reference cases were run using *DS2V* at the same freestream conditions, and the relevant stagnation heat fluxes extracted and used to normalise the *dsmcFoam+* heat fluxes. *DS2V* was selected for these simulations as it is a serial code (meaning that it only utilises a single processor core during calculations) that exclusively performs 2D DSMC simulations, allowing the required cases to be run in a short time (compared to those run using *dsmcFoam+*).

Note that the reason for each of these cases having a unique set of freestream conditions is that the resulting data was intended for input into a GPR model. The conditions were selected using an LHS (Latin Hypercube Sampling) plan in order to minimise the number of required simulations while still providing a good representation of the underlying thermal trends. As such, subsequent figures which present the data from these simulations only show the interpolation surfaces produced by the GPR models when queried, and not the actual data points used to train them (as they are scattered throughout the parameter space).

The following sections examine the trends of normalised heat flux, $\dot{q}/\dot{q}_0$, at the sampling locations defined in Section 3.4.1 with respect to angle of incidence, α.

3.4.4.1 Stagnation line

Normalised stagnation line heat fluxes are presented in Figure 3.29 for the $Kn = 0.02$, $Kn = 0.1$, and $1 < Kn < 10$ cases. Those sections of the plotted surfaces which lie in the range $x/R_n > 1$ show GPR model predictions from the side stagnation line described in Figure 3.22, while those data which lie in the range $0 \leqslant x/R_n \leqslant 1$ show predictions over the front stagnation line. The raw DSMC simulation data that the GPR models were initially trained on are also plotted in black for the first two cases.

The covariance function used for the predictions in Figures 3.29(a) and 3.29(b) is the same as that presented as Eq. (2.60) with a dimensionality of 2 ($D = 2$), where the relevant inputs are the angle of incidence, α, and normalised arc length, x/R_n. The cases at higher Knudsen numbers required more complex GPR kernels, since their initial conditions featured variations in both angle of incidence, and Knudsen number (as opposed to the previous two cases, whose Knudsen numbers were fixed). These kernels took the forms shown in Eqs. (3.36) and (3.37) for

the front and side stagnation lines, respectively. In the latter of these cases, the subscripts 1, 2, and 3 refer to the variables Knudsen number (Kn), angle of incidence (α), and normalised arc length (x/R_n), respectively.

In these covariance functions, $\mathbf{x}$ is the vector of freestream conditions on which the model has already been trained, and $\mathbf{x}'$ is a set of new conditions which is to be used to query the model. The variables σ_f, ℓ, and α are the model's variance from the mean of the training data, lengthscale, and an arbitrary scaling factor, respectively. These parameters are not listed here as it is the kernel matrix $K_{RQ,\,ARD}$ which conditions the model's response. This matrix is $n \times n$, where n is the number of input points used to train the model. As such, even small models can end up having extremely large kernel matrices, which are unsuited to printing.

$$K_{RQ,\,ARD}(\mathbf{x},\mathbf{x}') = \sum_{d=1}^{D} \sigma_{f_d}^2 \left(1 + \frac{(\mathbf{x_d} - \mathbf{x'_d})^2}{2\alpha_d \ell_d^2}\right)^{-\alpha_d} \tag{3.36}$$

$$\begin{aligned} K_{RQ,\,ARD}(\mathbf{x},\mathbf{x}') = &\left[\sigma_{f_1}^2 \sigma_{f_2}^2 \left(1 + \frac{(\mathbf{x_1} - \mathbf{x'_1})^2}{2\alpha_1 \ell_1^2}\right)^{-\alpha_1} \left(1 + \frac{(\mathbf{x_2} - \mathbf{x'_2})^2}{2\alpha_2 \ell_2^2}\right)^{-\alpha_2}\right] \\ &+ \sigma_{f_3}^2 \left(1 + \frac{(\mathbf{x_3} - \mathbf{x'_3})^2}{2\alpha_3 \ell_3^2}\right)^{-\alpha_3} \end{aligned} \tag{3.37}$$

These data are separated in this fashion for two specific reasons. Firstly, the discontinuity between the front and side faces is of particular interest in the study of heating on cylindrical geometries, so plotting the two as a combined surface would have likely led to reduced clarity in the presentation of the data. Secondly, the GPR models which were computed based upon these data were similarly split about the geometry's corner such that the selected covariance functions would not have to deal with the discontinuity. This would likely have resulted in reduced accuracy when predicting normalised heat fluxes close to the corner.

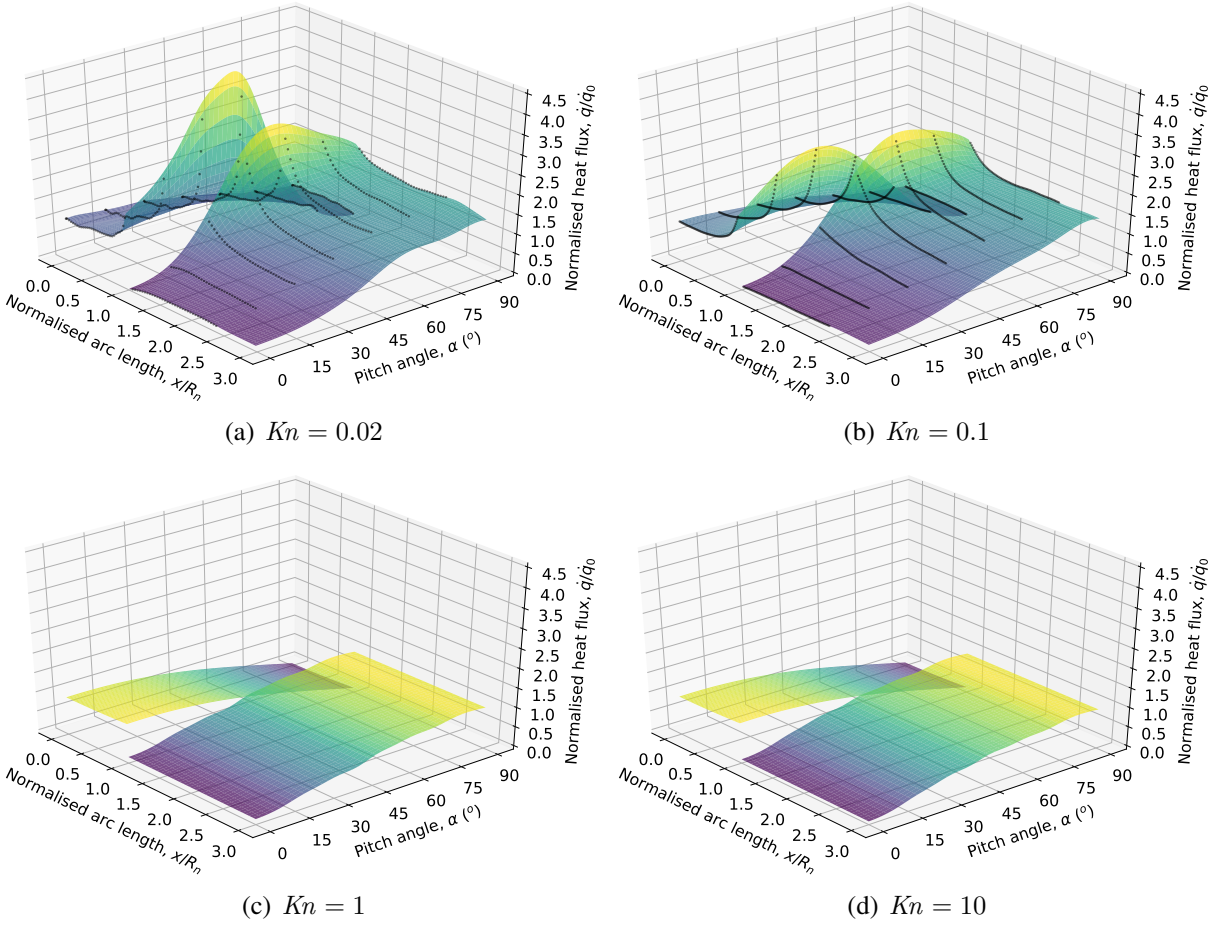
Similarly, separate GPR models were constructed for different Kn ranges such that the accuracy of the model responses when queried would be maximised. A model covering the full range of examined Knudsen numbers ($0.02 \leqslant Kn < 10$) would likely be poorly conditioned compared to the separately trained models described here. During training, the suitability of

the models was evaluated based upon both the smoothness of the response when queried, and how well the training points were matched, i.e. the initial conditions of the training were altered iteratively until a response which matched both the input data points and the expected trends between them was achieved. These initial conditions for the GPR training process are highly influential in the final response of the model and, if poorly chosen, allow the models to produce predictions which reduce to 0 between training points or, in the opposite case, over-smooth the response into a flat or gently curved hyperplane.

The GPR models generated during this part of the present work were designed to function as interpolation functions for the plotting of data. Additionally, it was desirable to study their effectiveness when supplied with "real" training data, i.e. sparse or otherwise non-ideal distributions of test conditions from various expensive experiments and/or simulations (particularly in the case of the $1 < Kn < 10$ model). As such, their suitability for future modelling of this type could be evaluated.

Considering Figure 3.29, the onset of Newtonian-style behaviour with increasing Knudsen number is clear (note that this refers to an increasing overall dependence of normalised heat flux trends on a face's orientation relative to the freestream, and a reduction in edge proximity effects). This occurs due to the lower number of intermolecular collisions in the flow brought on by high levels of rarefaction at high Kn (this is the same reason that the shock structures in Figure 3.24(b) become more diffuse at $Kn = 0.1$ compared to at $Kn = 0.02$).

By comparing Figures 3.29(a) and 3.29(b), a reduction in the magnitudes of local maxima can be observed, whereby peaks in the data toward the cylinder corner (where $x/R_n = 1$) are lower overall in the $Kn = 0.1$ cases than those run at $Kn = 0.02$. Additionally, the heat flux trends surrounding these points show slightly less steep transitions in the higher Kn cases. This behaviour is brought about by the trend, i.e. distribution, of normalised heat flux tending towards dependence solely on the angle between a surface and the local freestream vector (this is the Newtonian thermal behaviour described by Schaaf and Chambre's expression – see Eq. (3.28)). As such, the heat flux trends become "flatter" as $Kn \to \infty$, i.e. there is less variation of $\dot{q}/\dot{q}_0$ over the same geometry at higher Kn.

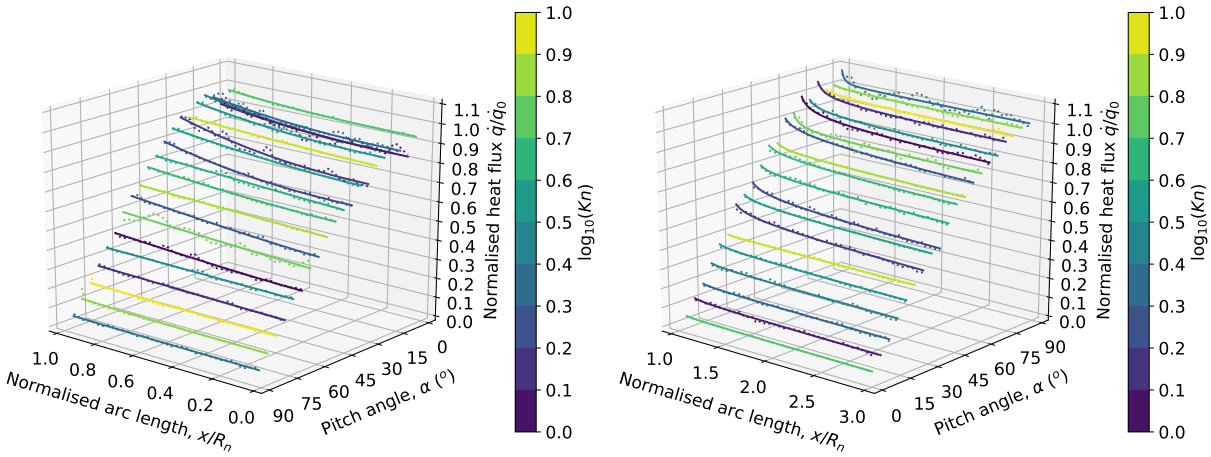


(a) $Kn = 0.02$ (b) $Kn = 0.1$

(c) $Kn = 1$ (d) $Kn = 10$

Figure 3.29: Gaussian process interpolation of heat fluxes on a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. GPR interpolation of normalised heat flux trends along the stagnation line are plotted as surfaces, and simulation data on which these models were trained are plotted in black (for LDWT derived cases only).

Considering Figures 3.29(c) and 3.29(d), it is evident that this transition to Newtonian-style behaviour progresses significantly as Kn increases past 0.1. The peaks in normalised heat flux that previously manifested towards the corner of the cylinder are almost completely absent from these data, with the trends at $Kn = 1$ appearing identical to those at $Kn = 10$. Other Kn predictions in the range $1 \leqslant Kn \leqslant 10$ were not performed owing to this apparent uniformity as $Kn > 1$.



(a) Front stagnation line (b) Side stagnation line

Figure 3.30: Normalised heat fluxes along the stagnation line of a cylinder at various angles of incidence (α) and Knudsen numbers ($1 \leqslant Kn \leqslant 10$). DSMC simulation data (points) and associated power law fits (lines) are presented. Note that the orientation of the pitch angle axis has been adjusted such that both sets of data are oriented similarly for presentation.

During the analysis of these data, it became apparent that the GPR model used to interpolate the high Knudsen number DSMC cases (those with Knudsen numbers in the range $1 < Kn < 10$) was not correctly capturing the small increases in normalised heat flux at the corner. This conclusion was drawn after the raw data from these DSMC simulations was examined visually. For reference, these data, i.e. the stagnation line results of each individual DSMC simulation used in the $1 < Kn < 10$ GPR model, are presented in Figure 3.30. It is clear that non-negligible increases in $\dot{q}/\dot{q}_0$ manifest as $x/R_n \to 1.0$ at lower values of Kn (power law curve fits computed using Eq. (3.38) and least squares regression are also plotted). These data indicate that the transition to Newtonian-style heating behaviour does in fact not occur until $Kn > 1$, as opposed to the indications given by Figures 3.29(c) and 3.29(d).

$$y = a(x - b)^c + f \tag{3.38}$$

Initially, a sub-model correlation strategy similar to that described in Section 3.3.9 was used to correlate these data. The variables of the aforementioned power law fits were correlated using three GPR models, but did not produce adequate predictions when queried (likely due to the inherent sensitivity of power laws as $x \to 0$). As such, the single models described in Eqs. (3.36) and (3.37) were used instead.

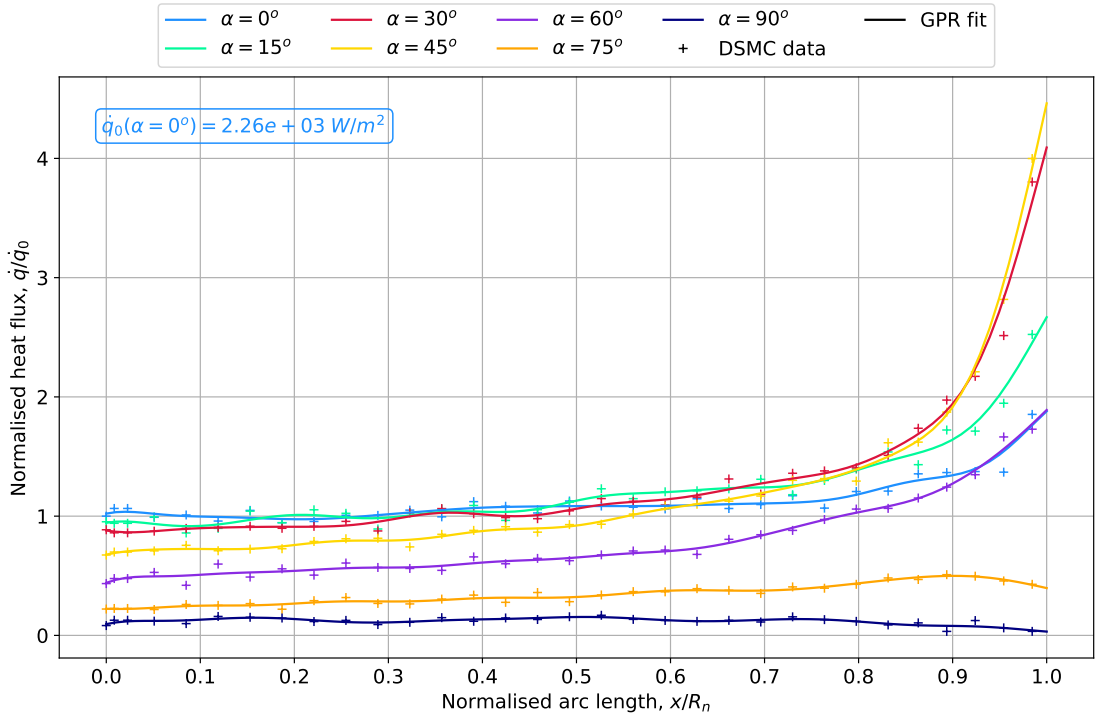


Figure 3.31: Normalised heat fluxes along the front stagnation line of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.02$). DSMC simulation data and accompanying GPR model interpolations are plotted.

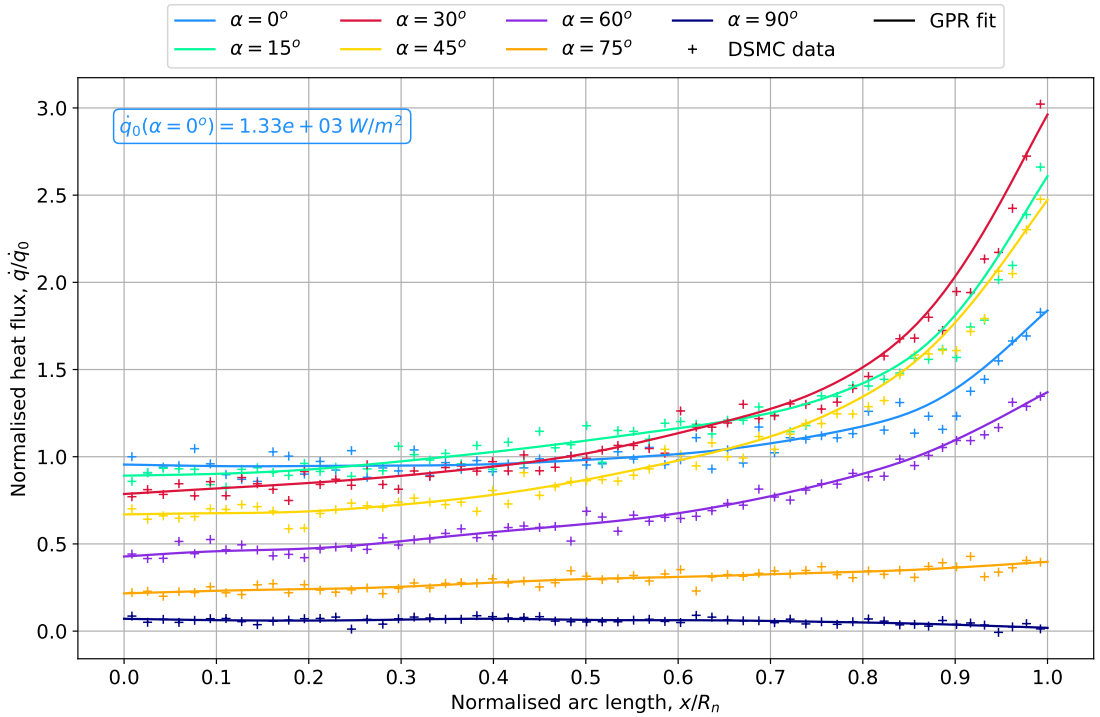


Figure 3.32: Normalised heat fluxes along the front stagnation line of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.1$). DSMC simulation data and accompanying GPR model interpolations are plotted.

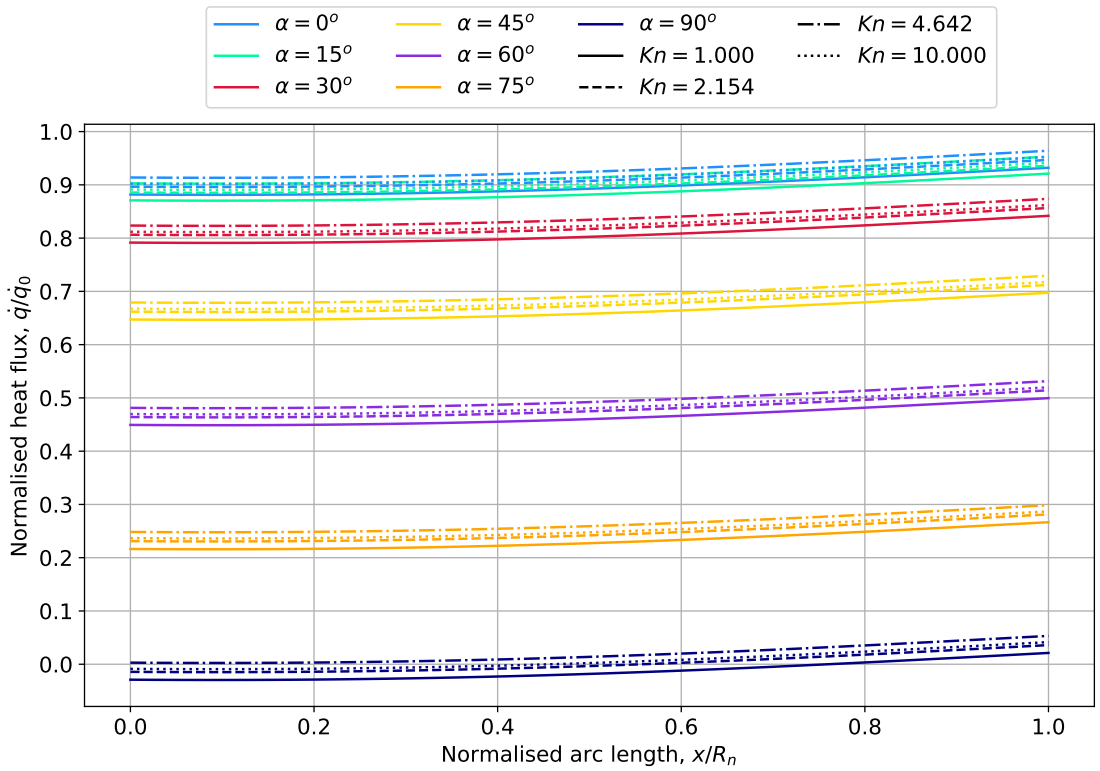


Figure 3.33: Normalised heat fluxes along the front stagnation line of a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. Only predictions from the $1 < Kn < 10$ GPR models are plotted.

Normalised heat fluxes along the front stagnation line are plotted against their respective normalised arc lengths for $Kn = \{0.02,\ 0.1,\ 1,\ 2.154,\ 4.642,\ 10\}$ in Figures 3.31, 3.32, 3.33 (note that these values are best interpreted on a logarithmic scale, i.e. $log_{10}(Kn) = \{-1.699,\ -1,\ 0,\ 0.333,\ 0.667,\ 1\}$). Also, no discrete data points are shown in Figure 3.33 as the GPR models used to produce these predictions were trained on an LHS plan, meaning that the input data points are sparsely distributed throughout the underlying parameter space (which comprises variations in Kn, α, and x/R_n). Finally, the value of $\dot{q}_0$ used to normalise these data was obtained from the accompanying *DS2V* simulations described previously in this section.

Examining the first two of these reinforces those conclusions that have previously been arrived at; specifically, that normalised heat flux trends decrease in severity (i.e. peak magnitudes are reduced) with increasing Knudsen number. The latter figure, which includes data from the GPR models in the range $1 < Kn < 10$, further supports this finding, indicating that heat flux trends flatten significantly as Kn rises above 0.1. Additionally, the simplification of the boundary

layer structure toward Newtonian behaviour with decreasing freestream density (and hence, rising Knudsen number) is also evident, with those cases at high angles of incidence, α, having uniformly lower heat flux trends than those at low angles.

As can be observed from Figures 3.31, 3.32, and 3.33, the highest normalised heat fluxes on the front stagnation line occur at $\alpha = 45^o$ and $\alpha = 30^o$, respectively, for the $Kn = 0.02$ and $Kn = 0.1$ cases. As one would expect from cases with higher Knudsen numbers, the maximum normalised fluxes in this location occur at progressively lower values of α, eventually falling to 0^o as $Kn \rightarrow \infty$. The lowest normalised heat fluxes occur at $\alpha = 90^o$ in all cases. This behaviour is once again due to the onset of Newtonian heating behaviour as $Kn \rightarrow \infty$, i.e. the trends of $\dot{q}/\dot{q}_0$ become increasingly dependant on the surface inclination angle relative to the freestream vector as $Kn \rightarrow \infty$. As such, higher values of surface inclination angle result in higher normalised heat fluxes.

Close scrutiny of these data, however, reveals a series of counter-intuitive results. The most evident of these is that the lowest Knudsen number examined (1.0) also has the lowest normalised heat flux trends with respect to x/R_n; the inverse of what one would typically expect from the other data presented in this section. Additionally, the other Knudsen numbers predicted using this particular GPR model do not produce trends in a logical order, with predictions at $Kn = 4.642$ being greater than those for $Kn = 10$. One would typically expect predictions at $Kn = 10$ to have the lowest normalised heat fluxes, while those performed at $Kn = 1$ should logically have the highest. As this is clearly not the case, it casts doubt on the validity of the GPR model in this instance for determining Knudsen number specific heat flux distributions, despite the fact that the model has evidently captured the other trends (with respect to α and, to a lesser extent, x/R_n) correctly.

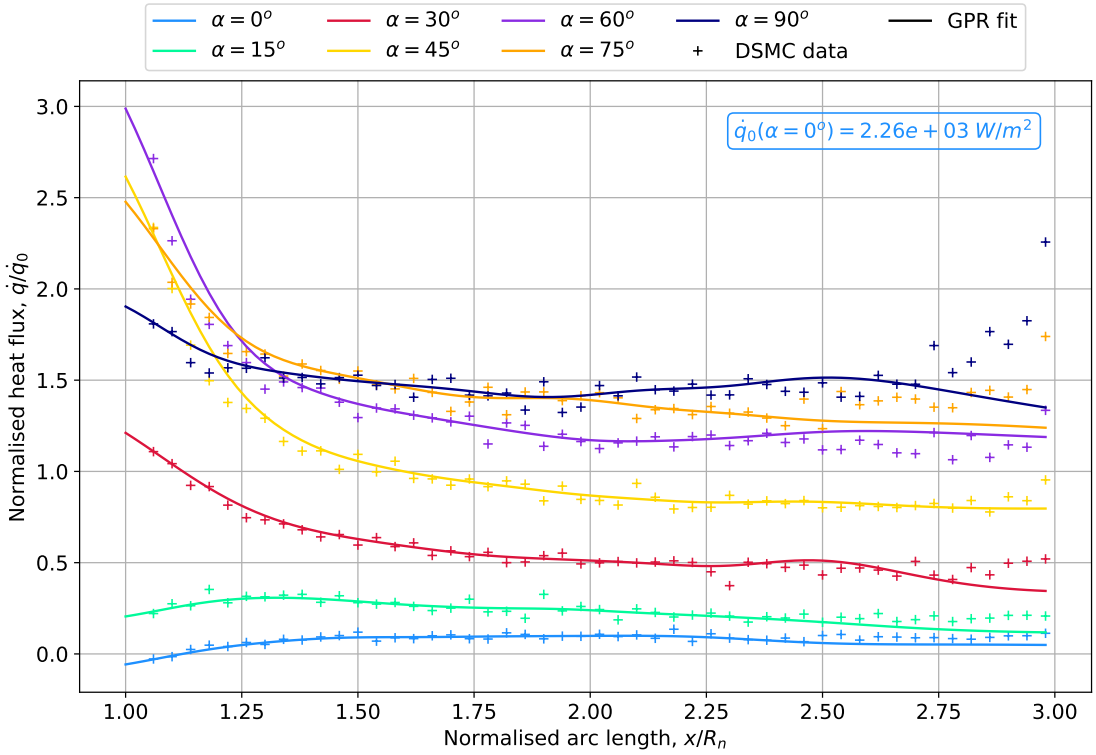


Figure 3.34: Normalised heat fluxes along the side stagnation line of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.02$). DSMC simulation data and accompanying GPR model interpolations are plotted.

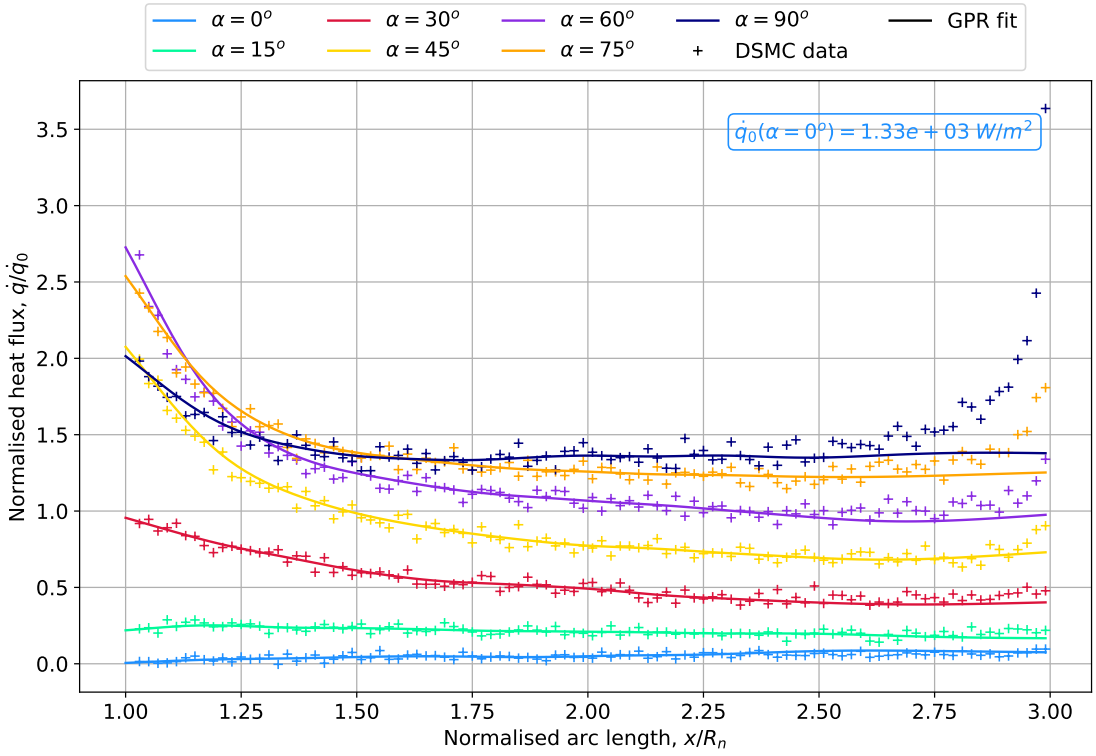


Figure 3.35: Normalised heat fluxes along the side stagnation line of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.1$). DSMC simulation data and accompanying GPR model interpolations are plotted.

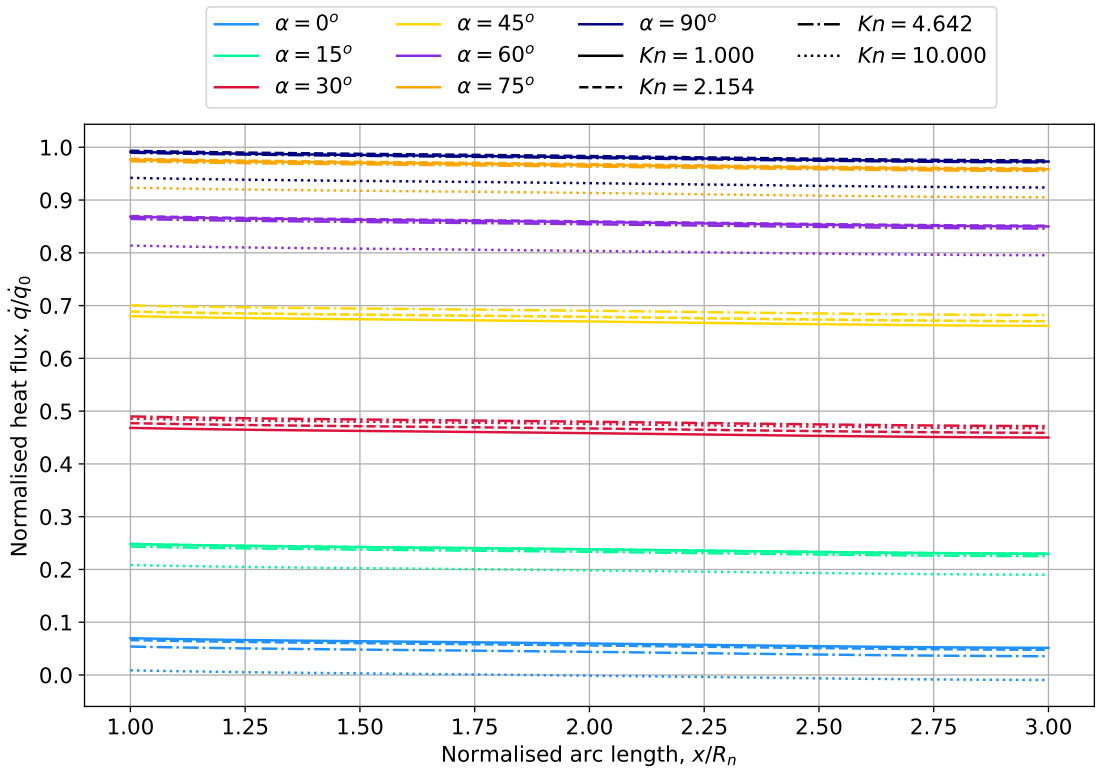


Figure 3.36: Normalised heat fluxes along the side stagnation line of a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. Only predictions from the $1 < Kn < 10$ GPR models are plotted.

Normalised heat fluxes along the side stagnation line are presented in Figures 3.34, 3.35, and 3.36 using the same order and formatting as those on the front stagnation line. The first two of these feature the same trend observed previously, i.e. increasingly Newtonian behaviour as $Kn \to \infty$, but also feature significant increases in $\dot{q}/\dot{q}_0$ as $x/R_n \to 3$. This is a non-physical effect in the simulations with $Kn = \{0.02, 0.1\}$ wherein the cylinder wall meets the flow boundary. As such, these data were excluded from the GPR models such that they would not interfere with the quality of the fits.

Considering Figure 3.36, it may be observed that the GPR model used to interpolate the high Kn heat flux trends on the side stagnation line shows significantly better conditioning than that used for the front stagnation line. This conclusion is supported by the logical spacing of most of the normalised heat fluxes, as well as their relative magnitudes, i.e. those trends at $Kn = 10$ are lower in most cases than those at $Kn = 1$. However, illogical predictions like those returned by the front stagnation line GPR model still manifest at $\alpha = 30^o$ and $\alpha = 45^o$, where low Kn normalised fluxes are higher than those at lower values of Kn. This indicates that the

$1 < Kn < 10$ GPR models are only partially suitable for such predictions, and have not captured small details in the heating trends (such as the slight upwards trends as $x/R_n \to 1$ in the low Kn cases presented in Figure 3.30). This is most likely caused by a combination of inappropriate kernel functions (sudden variations in data sets such as the slight upwards trends mentioned previously are unlikely to be captured well by the chosen kernel functions) and a comparatively small number of data available for training (20 variations of Ma and Kn, as opposed to several hundred x/R_n points).

As can be observed from Figures 3.34, 3.35, 3.36, the highest normalised heat fluxes on the side stagnation line occur at $\alpha = 60^o$ for both the $Kn = 0.02$ and $Kn = 0.1$ cases. As one would expect from cases with higher Knudsen numbers, the maximum normalised fluxes in this location occur at progressively lower values of α, eventually rising to 90^o as $Kn \to \infty$. The lowest normalised heat fluxes occur at $\alpha = 0^o$ in all cases.

Another interesting result of this study may be revealed by examining the preceding six figures at once. Typically, in Newtonian flow, a cylinder would have the same heat fluxes on its windward stagnation line faces at $\alpha = 45^o$ (since the angle between the cylinder surfaces and the freestream vector would be the same along both stagnation lines). The angle of incidence at which this crossover point occurs is dependant on the shape of the shock that forms in front of the cylinder, and shifts from $\alpha \approx 55^o$ for the $Kn = 0.02$ case (see Figures 3.31 and 3.34) to $\alpha \approx 50^o$ in the $Kn = 0.1$ case (see Figures 3.32 and 3.35), before finally settling to $\alpha = 45^o$ as $Kn \to 1$ (see Figures 3.33 and 3.36). This value of α then remains constant as $Kn \to \infty$, and may be observed clearly from Figures 3.33 and 3.36.

3.4.4.2 Correlation of results in terms of angular coordinate, θ

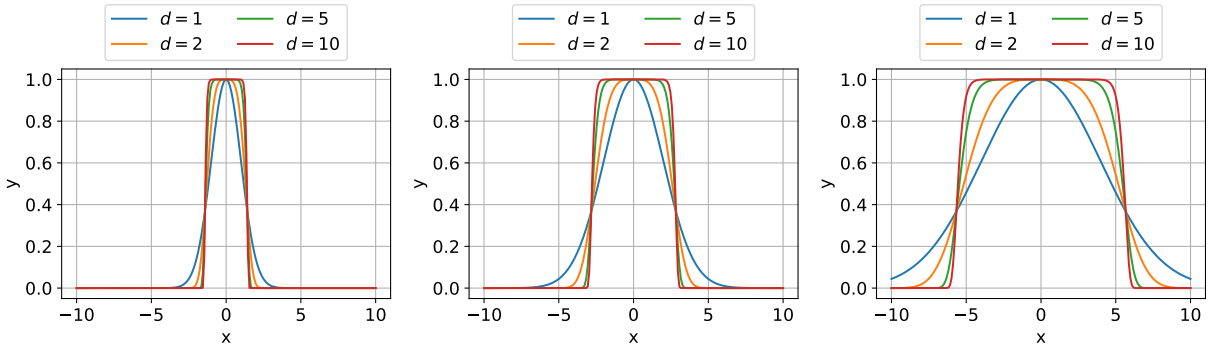
For those sampling locations (see Section 3.4.1) whose dependent geometric variable is angular in nature, i.e. θ, a correlation strategy similar to that described in Section 3.3.9 was implemented. Specifically, a super-Gaussian was employed, the general equation for which is presented as Eq. (3.39). Equation (3.40) is included for completeness, and shows the variables which replace y and x during actual fitting to DSMC results (b is also omitted, as it was set to 0 for simplicity - see below).

This curve is identical to an ordinary 2D Gaussian, except for the inclusion of the additional term d, which allows the peak of the curve to flatten to a plateau, thereby granting greater flexibility to fitting heat flux trends at higher Knudsen numbers.

$$y = a \exp\left(-\left[\frac{(x-b)^2}{2c^2}\right]^d\right) + f \tag{3.39}$$

$$\frac{\dot{q}}{\dot{q}_0} = a \exp\left(-\left[\frac{\theta^2}{2c^2}\right]^d\right) + f \tag{3.40}$$

The effects of varying the super-Gaussian's coefficients c and d are presented in Figure 3.37. Note that, in these plots, the value of the other coefficients are: $a = 1$, $b = 0$, and $f = 0$.



(a) $c = 1$ (b) $c = 2$ (c) $c = 4$

Figure 3.37: "Super Gauss" curves demonstrating the function response with varying input parameters

In subsequent sections, some GPR models used for interpolating the normalised heat flux trends with respect to θ are augmented with this sub-model in order to allow smaller, more simplistic models to be utilised, rather than relying on large single models trained on all of available data (as in the case of the front and side stagnation lines). Much like the aerodynamic models described in Section 3.3.9, sub-model augmented GPR schemes are faster both to train, and to derive predictions from, although they do typically require a larger set of training data. Models of this kind also take up significantly less hard disk space, and generally demonstrate higher accuracy than single models.

As was the case with the aforementioned aerodynamic models, the variable b in Eq. (3.39) was set to 0 prior to fitting such that the resulting curves would be symmetric about $x = 0$ (which translates to $\theta = 180^o$ in the chosen coordinate system). The final predictive models comprised

four separate GPR sub-models; one for each of the remaining coefficients (a, c, d, and f), i.e. a separate GPR model was defined and trained for each parameter, with the input variables for each model being Ma, Kn, and α. In this fashion, a single prediction of $\dot{q}/\dot{q}_0$ involves all four sub-models being queried with the case's freestream conditions, and their predictions of a, c, d, and f being used in conjunction with Eq. (3.40) to generate a super-Gauss curve of $\dot{q}/\dot{q}_0$ against θ.

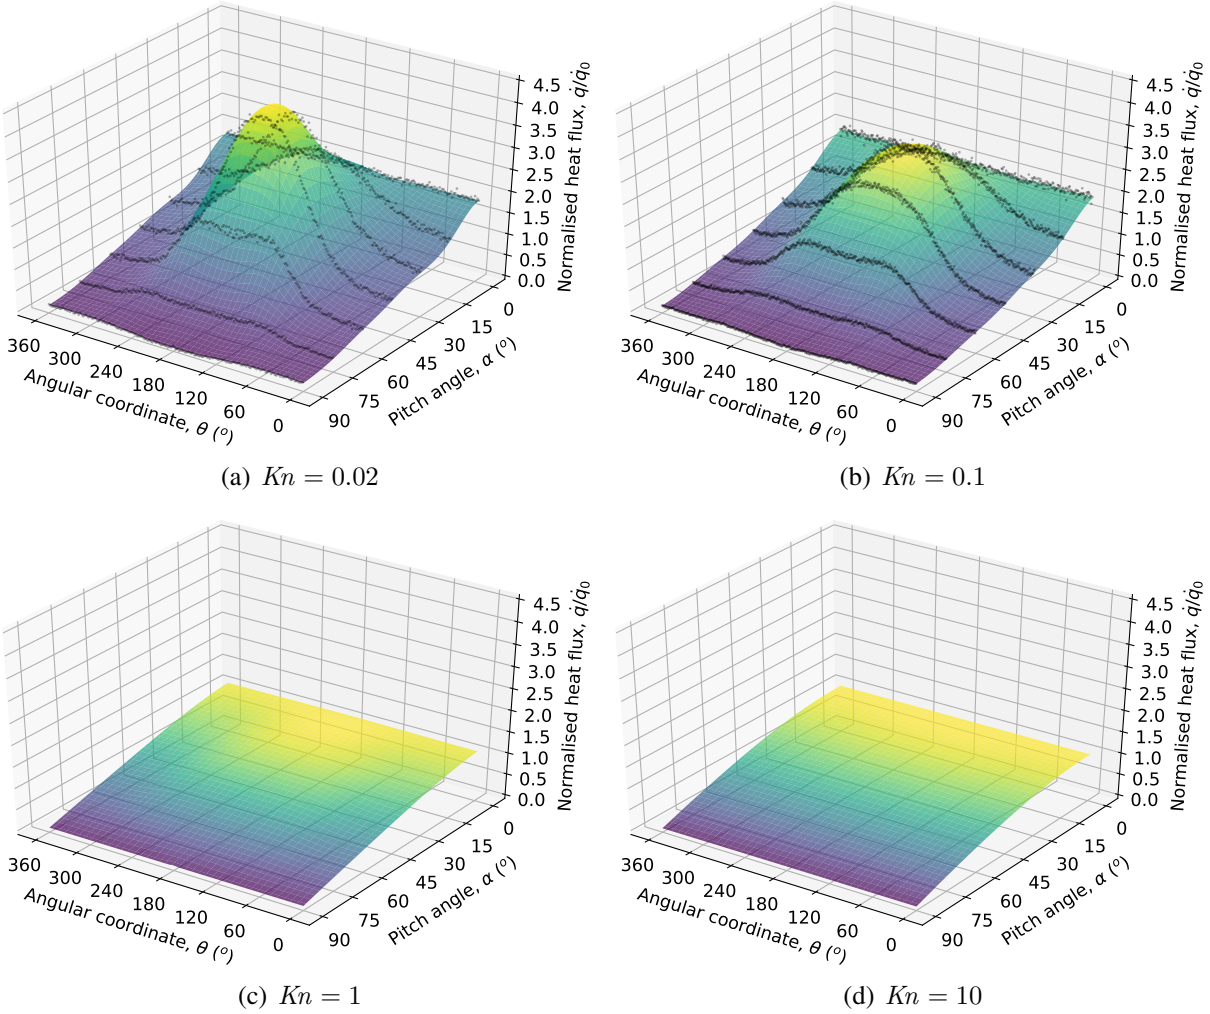
3.4.4.3 Front edge

Normalised heat flux trends on the front edge of the cylinder (as defined in Section 3.4.1) are presented in Figure 3.38 for $Kn = \{0.02,\ 0.1,\ 1,\ 10\}$. For clarity, plots of these trends at $\theta = 180^o$ are presented in Figure 3.39 as the highest heat fluxes were observed to manifest at this angular location. Local maxima are marked on the plots, and GPR predictions of the normalised heat flux trends at $Kn = \{2.154, 4.6342\}$ are also featured. The transition to Newtonian thermal behaviour is even more evident at these locations than along the stagnation line examined previously, with the peak heat fluxes decreasing significantly as $Kn \to \infty$.

As was noted with the stagnation line results (see Figure 3.30), full transition to this behaviour does not occur until $Kn > 1$. This can be observed in Figures 3.38(c) and 3.39, where a small remnant of the central peak that features prominently in Figures 3.38(a) and 3.38(b) is visible.

The maxima of these data agree well with the front stagnation line results (see Figure 3.39), with peaks at $\alpha = 39.86^o$ in the $Kn = 0.02$ case, $\alpha = 24.35^o$ in the $Kn = 0.1$ case, $\alpha = 9.20^o$ at $Kn = 1$, and finally $\alpha = 4.69^o$ at $Kn = 10$. The conclusion previously arrived at from the front stagnation line results holds here as well, i.e. that the angle of incidence at which peak normalised heat flux occurs approaches 0^o as Knudsen number grows arbitrarily large: $\lim\limits_{Kn\to\infty} \alpha \left(\max\left(\dot{q}/\dot{q}_0\right)\right)_{CFE} = 0^o$. Somewhat counter-intuitively, these data also indicate that the normalised heat fluxes at $\alpha = 0^o$ are in fact higher for the $Kn = 0.1$ case than they are at $Kn = 0.02$, indicating the existence of a trend in Kn that does not uniformly decrease as $Kn \to \infty$. This is due to the shape of the shock and boundary layer structure in front of the cylinder, which reduces in complexity towards Newtonian behaviour as $Kn \to \infty$. Figure 3.40 illustrates this expansion region at the edge of a cylinder oriented at $\alpha = 90^o$ for some Kn below

the free molecular (Newtonian) limit.



(a) $Kn = 0.02$ (b) $Kn = 0.1$

(c) $Kn = 1$ (d) $Kn = 10$

Figure 3.38: GPR interpolation of normalised heat fluxes on a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. Interpolated trends along the outermost edge of the circular face are plotted as surfaces, and DSMC simulation data are plotted in black (for LDWT derived cases only).

Plots of normalised heat flux against angular coordinate, θ, on the front circular edge are presented in Figures 3.41, 3.42, and 3.43 for $Kn = \{0.02,\ 0.1,\ 1,\ 2.154,\ 4.642,\ 10\}$.

Similar to the front stagnation line results, these data have a common minimum at $\alpha = 0^o$, but also feature a visible dip in the value of $\dot{q}/\dot{q}_0$ at $\theta = 180^o$. This is likely an effect of the position of the stagnation point at high values of α, which is coincident with the forwardmost point of the cylinder edge. The result of this positioning is that flow accelerates past these sections of the cylinder's front edge, enclosing the region around $\theta = 180^o$ in an expansion fan, and hence reducing the local heat transfer rate. This phenomenon can be seen to manifest at all angles of

incidence above 45^o ($\alpha > 45^o$) in the $Kn = 0.02$ (see Figure 3.41) and $Kn = 0.1$ (see Figure 3.42) cases. High Kn cases in the range $1 < Kn < 10$ (see Figure 3.43) were examined for this trend prior to their correlation with the super-Gaussian curve described in Section 3.4.4.2, as said curve would be unable to capture such a feature. It was subsequently found that the effects of the expansion fan at $\theta = 180^o$ identified above were negligible in all high Kn cases, and the corresponding reduction in $\dot{q}/\dot{q}_0$ did not manifest.

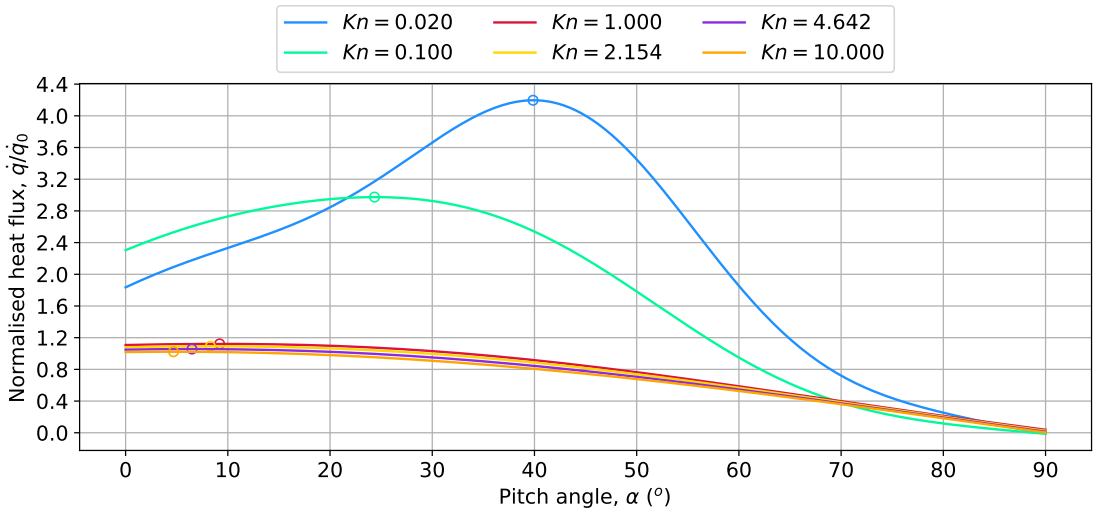


Figure 3.39: GPR model interpolation of normalised heat fluxes at $\theta = 180^o$ along the outermost edge of the circular face of a cylinder. Local maxima are denoted by markers.

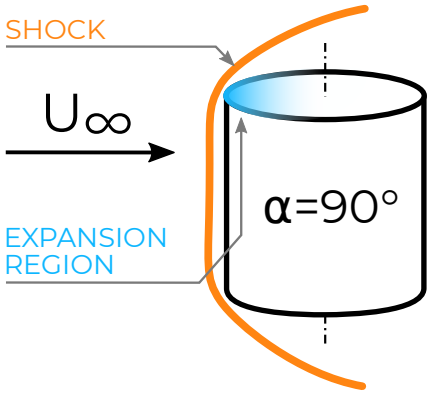


Figure 3.40: Diagram illustrating the expansion region behind the edge shock near $\theta = 180^o$ on cylindrical geometries oriented at high values of α.

This finding also serves to confirm the previous supposition regarding the onset of Newtonian-style heating as $Kn \to \infty$. However, it is evident from Figure 3.43 that fully Newtonian heating (which, in the case of a cylinder at some non-zero angle of incidence, α, satisfies $\dot{q}/\dot{q}_0 \neq f(\theta)$ such that there is no variation in normalised heat flux with angular

coordinate, θ) does not set in until the freestream Knudsen number rises above $Kn = 10$, as there is a clear dependence of $\dot{q}/\dot{q}_0$ on θ for all examined cases.

Unlike the high Kn predictions of $\dot{q}/\dot{q}_0$ along the stagnation line (see Figure 3.33), the GPR models computed for the front edge in the range $1 < Kn < 10$ show the expected behaviour as $Kn \to \infty$. Considering Figure 3.43, it may be observed that the highest normalised heat fluxes occur at low Knudsen numbers, while the smallest occur at high Knudsen numbers (as one would logically expect). Additionally, the spacing between the high Kn normalised heat flux predictions appears to be linear, which, when queried with logarithmically spaced values of Kn, is exactly the expected trend i.e. $\dot{q}/\dot{q}_0 \propto \log_{10}(Kn)$ as $Kn \to \infty$ (as derived analytically by both Klett [51] and Schaaf and Chambre [113]). Note that the small "bump" in the $\alpha = 90^o$ predictions is a non-physical numerical error – a minor side effect of the GPR/super Gauss sub-modelling strategy.

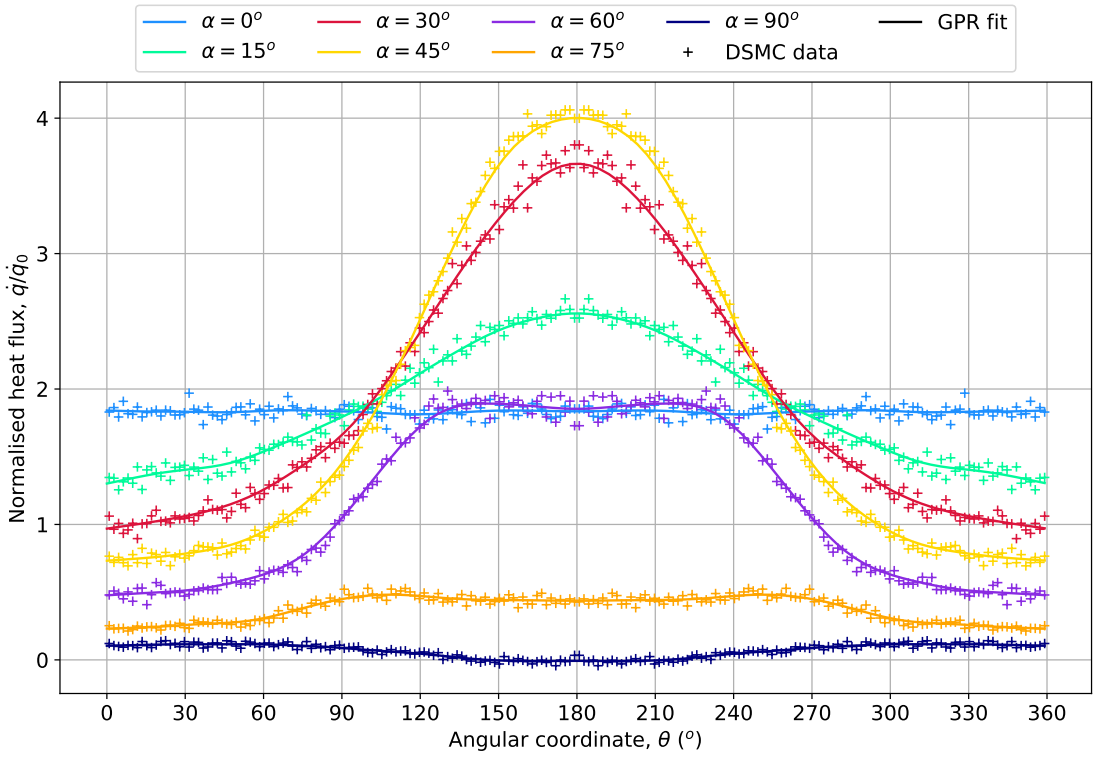


Figure 3.41: Normalised heat fluxes along the front circular edge of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.02$). DSMC simulation data and accompanying GPR model interpolations are plotted.

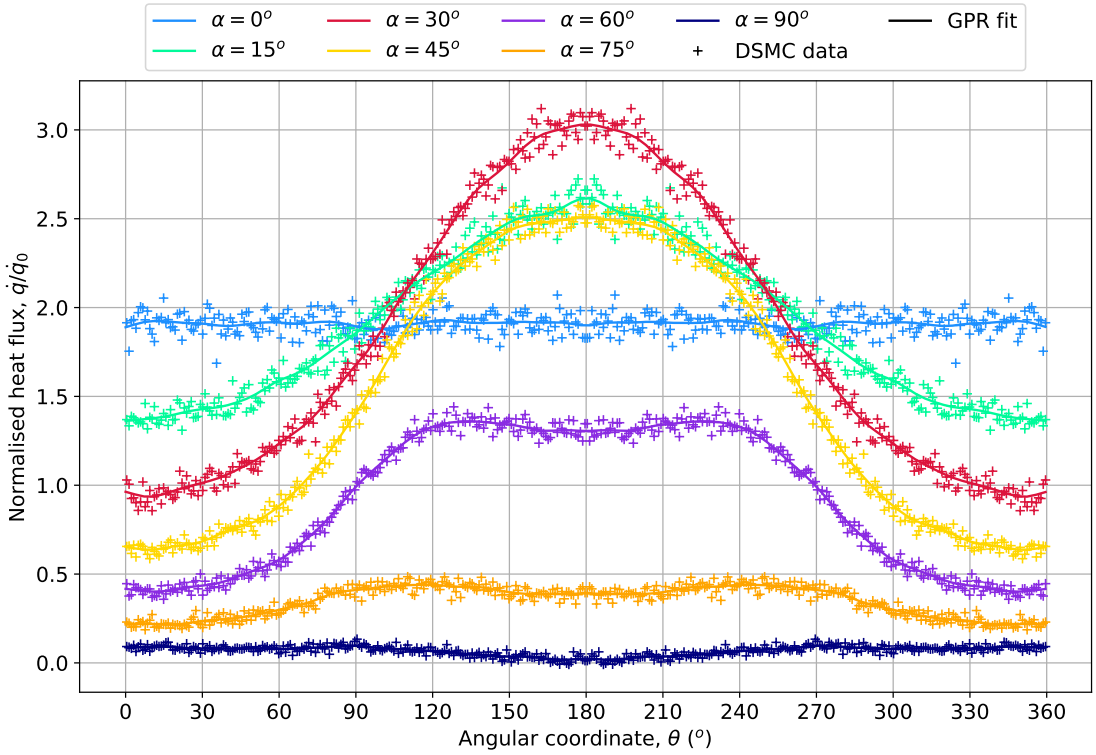


Figure 3.42: Normalised heat fluxes along the front circular edge of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.1$). DSMC simulation data and accompanying GPR model interpolations are plotted.

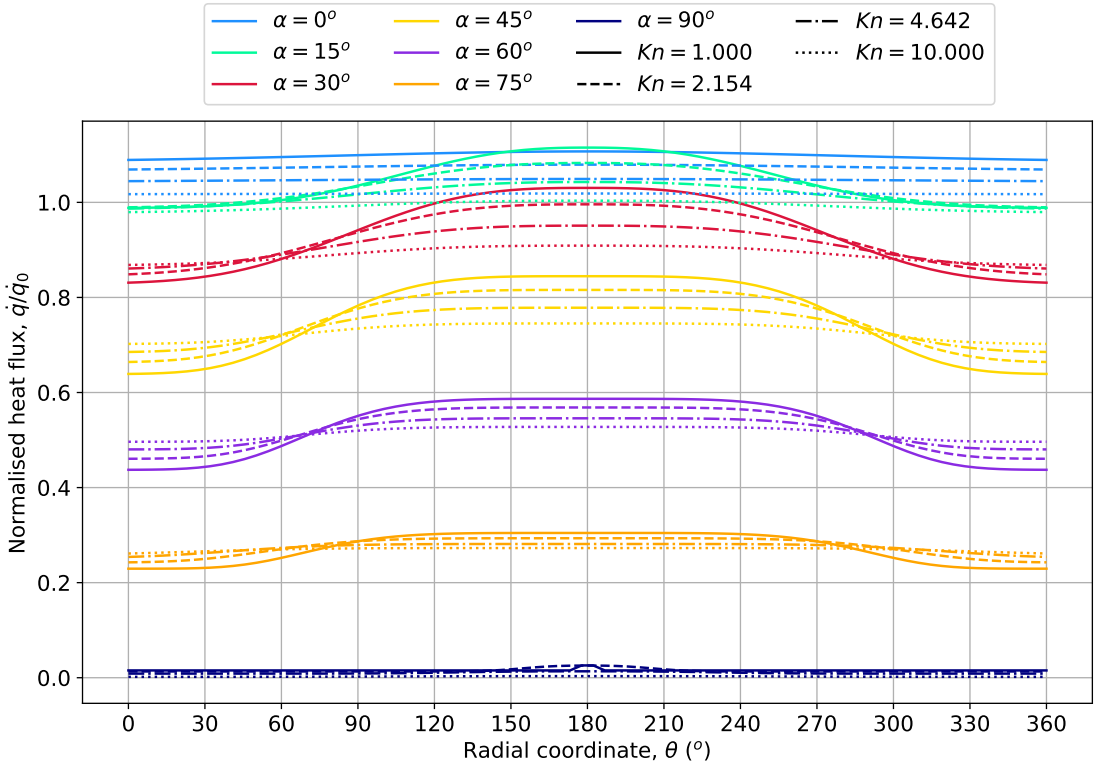


Figure 3.43: Normalised heat fluxes along the front circular edge of a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. Only predictions from the $1 < Kn < 10$ GPR models are plotted.

3.4.4.4 Side edge

Normalised heat flux trends on the side edge of the cylinder (as defined in Section 3.4.1) are presented in Figure 3.45 for $Kn = \{0.02, 0.1, 1, 10\}$. For clarity, plots of these trends at $\theta = 180^o$ are presented in Figure 3.39, with the local maxima marked (GPR predictions of the normalised heat flux trends at $Kn = \{2.154, 4.6342\}$ are also featured). As is the case with the front edge results examined previously, the transition to Newtonian thermal behaviour is also evident at these locations, with the peak heat fluxes decreasing significantly as $Kn \to \infty$.

Unlike the front section of the stagnation line or front circular edge, the side edge's normalised heat flux trend does not flatten to linearity as $Kn \to \infty$. Predictions of these fluxes with respect to the inclination of a surface's normal vector relative to the freestream vector, δ, are presented in Figure 3.44. These data were calculated using the equations of Lees [126] (presented as Eq. (3.27) which is valid for laminar flow in the continuum regime), and Schaaf & Chambre [113] (presented as Eq. (3.28), which is valid in the free molecular regime). An arbitrarily high Mach number was utilised in these calculations such that the upper limits of each equation with respect to δ would be accurately presented.

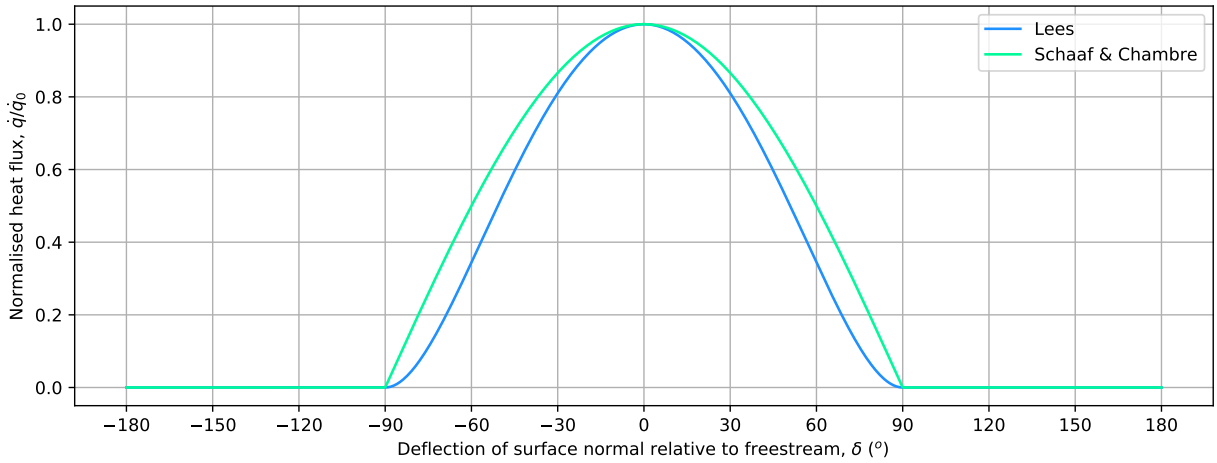
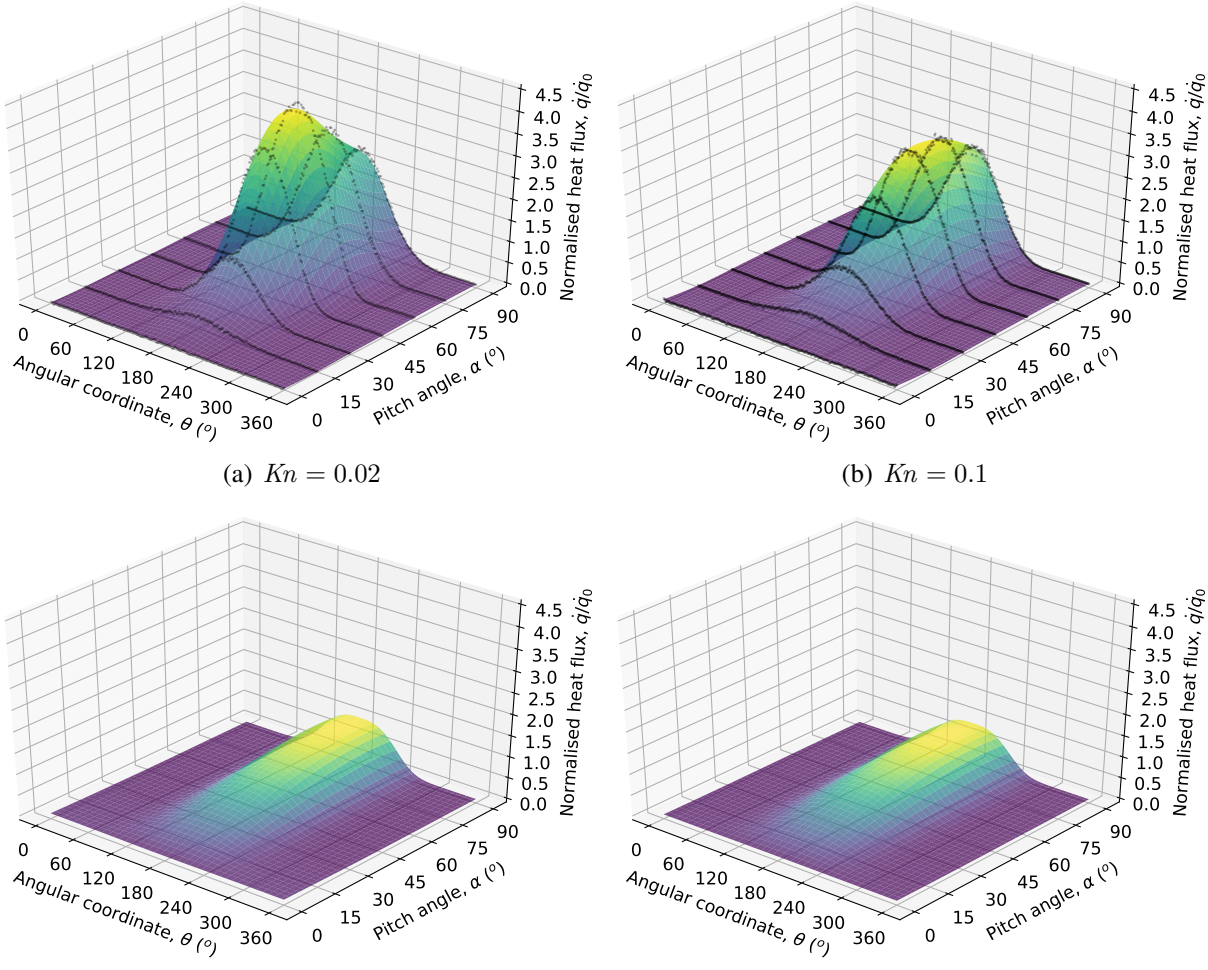


Figure 3.44: Predictions of normalised heat flux with respect to the angle between a surface normal vector and that of the freestream, δ, as $Ma \to \infty$. Calculations were performed using the equations of Lees and Schaaf & Chambre.

It may be observed from Figures 3.45 and 3.46 (note that, for the sake of comparison, here $\theta = \delta + (\pi/2)$), the normalised heat flux trends approach this behaviour as $\alpha \to 90^o$ and $Kn \to \infty$. As is the case with values of $\dot{q}/\dot{q}_0$ on the front edge (see Section 3.4.4.3), the peak

magnitudes of these trends decrease significantly as $Kn \to \infty$.

Plots of the normalised heat fluxes with respect to angular coordinate, θ, are presented in Figures 3.47, 3.48, and 3.49 for $Kn = \{0.02,\ 0.1,\ 1,\ 2.154,\ 4.642,\ 10\}$. These data demonstrate a similar trend to that of the side stagnation line results, and the inverse of that identified in the front circular edge results. Specifically, this is the increase in peak values of $\dot{q}/\dot{q}_0$ as $\alpha \to 90^o$.



(a) $Kn = 0.02$ (b) $Kn = 0.1$

(c) $Kn = 1$ (d) $Kn = 10$

Figure 3.45: GPR interpolation of normalised heat fluxes on a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. Interpolated trends along the forwardmost edge of the curved face are plotted as surfaces, and DSMC simulation data are plotted in black (for LDWT derived cases only).

Considering Figure 3.46, it may be observed that these peaks occur at $\alpha = 56.27^o$ in the $Kn = 0.02$ case, $\alpha = 65.65^o$ in the $Kn = 0.1$ case, and decrease in magnitude as $Kn \to \infty$. The high Kn cases also demonstrate the same improved behaviour seen in the front edge GPR model (see Figure 3.49), with linear spacing between logarithmically spaced Knudsen numbers, and a logical scaling of normalised heat flux as $Kn \to \infty$. However, the ordering of the local maxima

of these data is the inverse of the expected trend, with reducing values of $\alpha \left(\max\left(\dot{q}/\dot{q}_0\right)\right)_{CSE}$ in the range $1 \leqslant Kn \leqslant 10$. This is due to the relative "flattening" of the trend of $\dot{q}/\dot{q}_0$ as $Kn \to \infty$, wherein the GPR models, although well conditioned overall, find the small effects of Knudsen number as $Kn \to \infty$ difficult to incorporate into their predictions.

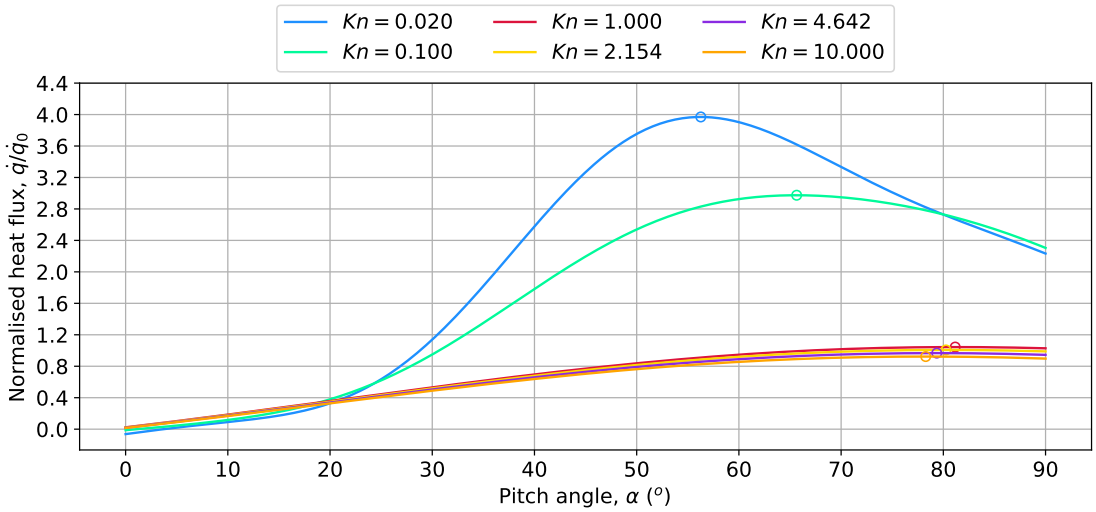


Figure 3.46: GPR model interpolation of normalised heat fluxes at $\theta = 180^o$ on the side curved edge of a cylinder. Local maxima are denoted by markers.

Once again, like the front edge results, these data share common minima, which for these cases occur at $\alpha = 0^o$. Additionally, the counter-intuitive trend of $\dot{q}/\dot{q}_0$ at low Knudsen numbers is also visible here, except that it occurs (as one would expect) at $\alpha = 90^o$. Although the normalised heat fluxes are indeed higher for the $Kn = 0.1$ case than they are at $Kn = 0.02$, the relative difference between the values at these points is significantly less pronounced than in the front edge data.

Comparing Figures 3.47, 3.48, and 3.49 to Figure 3.44, it is clear that the distribution of $\dot{q}/\dot{q}_0$ most closely resembles that predicted by Lees' expression, owing to the gentle transition to $\dot{q}/\dot{q}_0 = 0$ as $\delta \to \pm 90^0$ ($\theta \to 90^o$ and $\theta \to 270^o$ in the context of the former three figures). This transition is present even at the highest Knudsen numbers examined herein, and although it becomes visibly more severe (i.e. occurs over a smaller range of δ) as $Kn \to \infty$, it does not approach the Newtonian thermal behaviour of Schaaf & Chambre's expression (i.e. an instant, sharp transition to $\dot{q}/\dot{q}_0 = 0$ at $\delta = \pm 90^o$). Hence, this finding reinforces the conclusion that fully Newtonian thermal behaviour does not begin to manifest until $Kn > 10$. Note that, although the maximum Knudsen number examined in the present work was $Kn = 10$, clear signs of

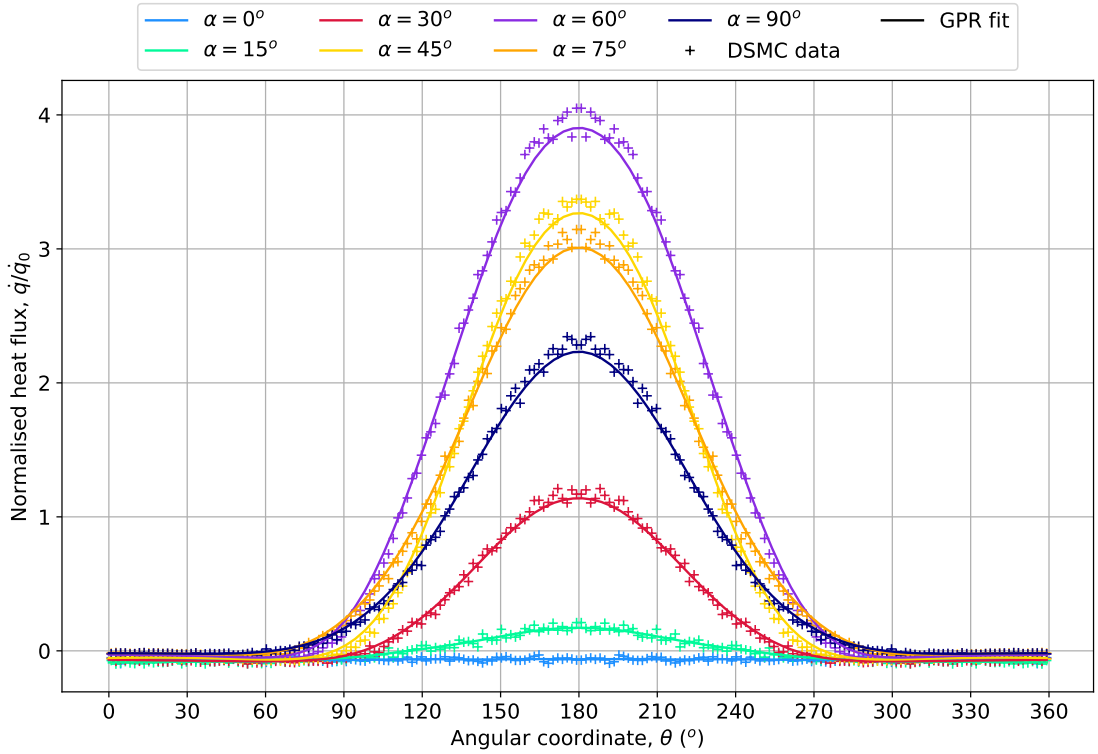


Figure 3.47: Normalised heat fluxes along the side circular edge of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.02$). DSMC simulation data and accompanying GPR model interpolations are plotted.

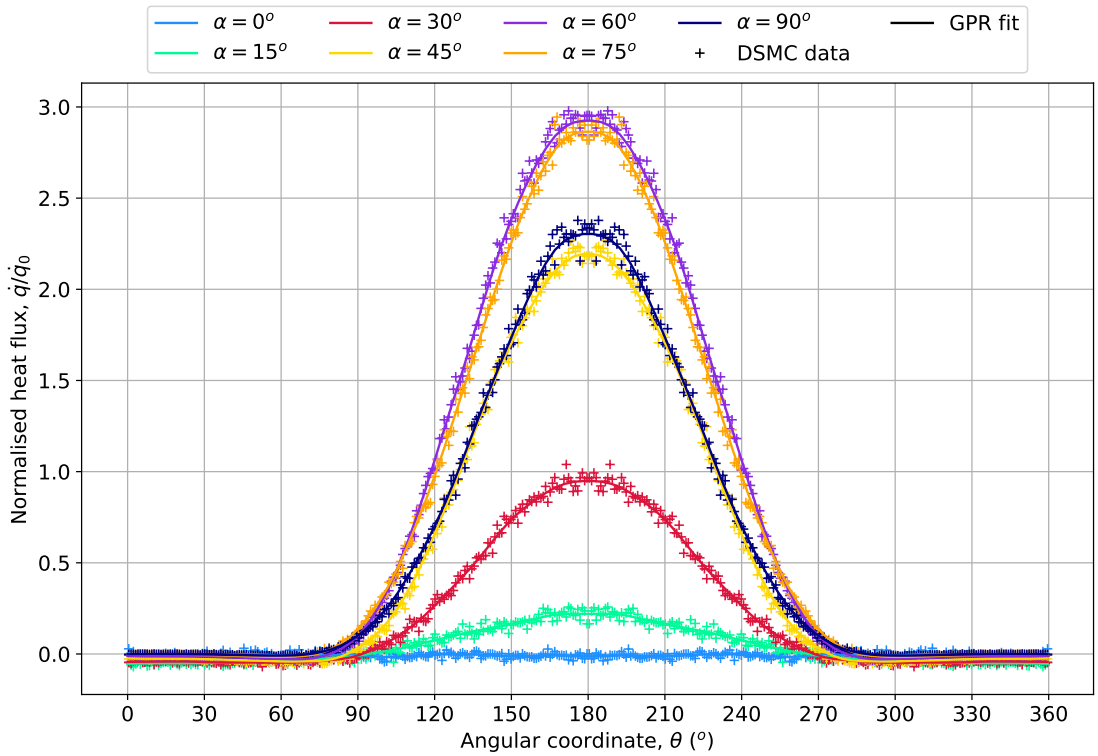


Figure 3.48: Normalised heat fluxes along the side circular edge of a cylinder at various angles of incidence (α) in hypersonic flow (based upon LDWT cases of interest at $Kn = 0.1$). DSMC simulation data and accompanying GPR model interpolations are plotted.

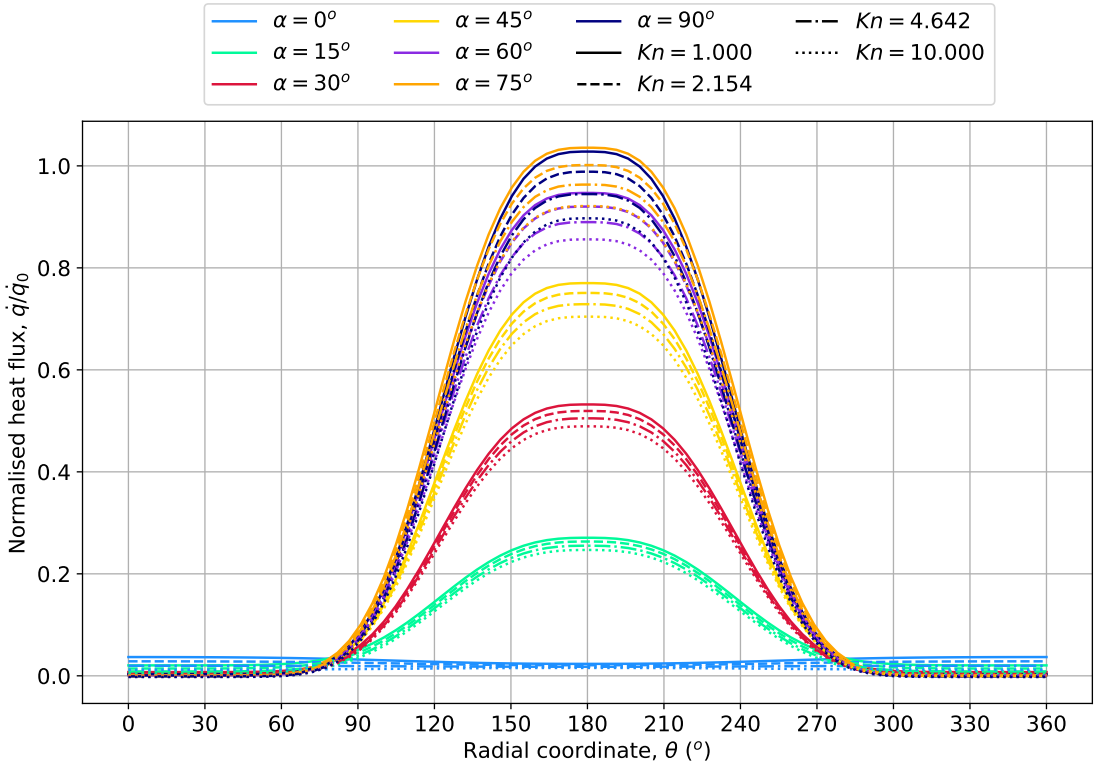


Figure 3.49: Normalised heat fluxes along the side circular edge of a cylinder at various angles of incidence (α) and Knudsen numbers (Kn) in hypersonic flow. Only predictions from the $1 < Kn < 10$ GPR models are plotted.

asymptotic behaviour were apparent at this value of Kn, indicating strongly that little additional change would occur in the observed heat flux trends as $Kn \to \infty$ (this is in line with the analytic determinations of Klett [51] and Schaaf and Chambre [113]).

3.5 SUMMARY AND DISCUSSION OF RESULTS

In this chapter, a number of different aerothermodynamic simulations are described, and the corresponding results discussed. The properties of the Earth's atmosphere (with respect to altitude) are predicted in these simulations via use of the NRLMSISE00 model [109] and the Cantera chemical kinetics package [39]. Simplified 3 degree of freedom trajectory calculations are then described, wherein idealised spheres are simulated as they re-enter the Earth's atmosphere from LEO. Results from these calculations are utilised in the formulation of initial conditions for subsequent simulations.

Following this, hypersonic panel methods are elaborated upon, as is the formulation of *RAC* (Re-Entry Aerothermal Calculator), a code developed for the present work. This code utilises the

aforementioned panel methods to calculate the aerodynamic coefficients and surface thermal flux distributions of arbitrary shapes during atmospheric entry. Validation results for *RAC* are then presented, where the code is shown to agree excellently with limiting expressions for lift and drag coefficients (C_L and C_D, respectively) on spheres and cylinders at high Mach numbers.

The generation of large sets of freestream conditions for input to *RAC* is described, wherein an LHS (Latin Hypercube Sampling) plan is utilised, and simulations are run on cylindrical and cuboidal geometries of varying aspect ratio and attitude. The resulting datasets, which consist of aerodynamic coefficients (C_D, C_L, and C_{my} in the case of the cylinder, and C_D, C_L, C_S, C_{my}, and C_{mz}, in the case of the cuboid), are then correlated with each cases' Mach number, Knudsen number, attitude, and aspect ratio using GPR modelling augmented with a novel Gauss-Sigmoid curve. These models are subsequently queried with randomised test data, and compared to the output from *RAC* when using the same test set as input. Excellent agreement is found between the two methods, with a minimum R^2, i.e. coefficient of correlation, value of 0.9992, and a peak mean range-normalised relative error, $\overline{\delta_R}$, of 0.134%.

It is apparent from these tests that the C_D, C_{my}, and C_{mz} GPR models feature the best conditioning, as they produce results with mean relative errors, $\overline{\delta}$, significantly lower than the C_S and C_L models. It is then found that apparent inaccuracies in the predictions of these last two models are the result of extremely small values of C_L and C_S, wherein even a negligible inaccuracy produces a large relative error. When considering range-normalised mean errors, ($\overline{\delta_R}$), however, it is clear that the C_L and C_S GPR models produce predictions which are suitable in the vast majority of cases. Hence, it is concluded that the GPR models generated as part of the present work are suitable for predicting the aerodynamic properties of arbitrary cylinders and cuboids with aspect ratios in the range $0.1 \leqslant L/\varnothing \leqslant 100$ at Mach numbers in the range $5 \leqslant Ma \leqslant 30$, Knudsen numbers in the range $10^{-5} \leqslant Kn \leqslant 10^{4}$, and attitudes in the range $0.1^o \leqslant \{\alpha, \beta, \phi\} \leqslant 89.9^o$. A summary of the test findings is presented in Table 3.10.

Following this description of aerodynamic models, the results of DSMC simulations are discussed wherein two geometries (a hemisphere and a cylinder) are examined at hypersonic conditions. The former of these is simulated at 2 different LDWT-derived conditions using the 2D *DS2V* code, while the latter is simulated using both *DS2V* and the 3D DSMC code

Geometry	Coefficient	$\overline{\delta}$	$\overline{\delta_R}$	$\delta_{NRMS(\mu)}$	$\delta_{NRMS(R)}$	R^2
Cylinder	C_D	0.367%	0.04%	0.514%	0.064%	1.0000
	C_L	30.776%	0.072%	5.582%	0.207%	0.9995
	C_{my}	0.385%	0.070%	0.660%	0.125%	0.9999
Cuboid	C_D	0.691%	0.076%	1.091%	0.137%	0.9998
	C_L	74.116%	0.070%	5.886%	0.169%	0.9994
	C_S	618.847%	0.069%	6.605%	0.171%	0.9992
	C_{my}	1.311%	0.134%	1.660%	0.260%	0.9995
	C_{mz}	5.497%	0.065%	1.240%	0.123%	0.9999

Table 3.10: Summary of error and correlation statistics for *RAC* GPR models

dsmcFoam+ at multiple conditions and attitudes.

The results from the *DS2V* simulations illustrate several important phenomena in the examined cases. Firstly, it is shown that the accepted conditions for an "optimal" DSMC simulation must be satisfied for accurate estimations of surface heat flux to be obtained. However, it is also shown that simulations utilising non-optimal DSMC variables are still suitable for examining heat fluxes when they are normalised to a suitable local stagnation value, i.e. $\dot{q}/\dot{q}_0$. The accuracy of common hypersonic heat transfer correlations is then studied with reference to these DSMC results. It is found that the Fay-Riddell stagnation heat flux correlation (Eq. (2.26)) is valid for those cases based upon the operating conditions of the LDWT, with maximum prediction errors of 2.9% and 18.0% for the hemisphere and cylinder, respectively. The higher errors observed in the cylinder cases arose as a result of the theoretical basis of the Fay-Riddell equation, which was derived for a sphere in hypersonic flow. This is despite the correlation incorporating an optional factor to account for cylinders, as this correction is based upon a cross-flow scenario, and not the head-on orientation that is the case in the present work.

The heat flux distribution correlations of Lees (Eq. (3.27)), Schaaf & Chambre (Eq. (3.28)), and Klett (Eq. (3.35)) are also examined, and are found to agree reasonably well with the *DS2V* results. Lees' equation is shown to be suitable for use with the hemisphere geometry in the transition regime, with only minor differences between the normalised DS2V results and predictions generated by the correlations. The expression of Schaaf and Chambre, however, is

shown to poorly reproduce these trends, with obvious inaccuracies on the hemisphere geometry manifesting in the range $0 < 2x/\pi R_n < 1$. Heat flux distributions in the cylinder case are shown to be poorly predicted by both correlations, owing primarily to their heavy dependence on surface inclination angle. This prevents the heat flux distributions on large flat surfaces (such as the circular face of a cylinder) from being correctly estimated. Klett's correlation for $\dot{q}/\dot{q}_0$ on the circular face of a cylinder (which was developed in order to mitigate the aforementioned issue with surface orientation dependence) is shown to agree well with high Mach number simulations, and reasonably well with those cases run at LDWT conditions. From these data, it is concluded that a small degree of Mach number dependence is present at LDWT conditions (as one would logically expect given the proximity of these conditions to the hypersonic limit of $Ma \approx 5$) which, for the purposes of the current work, may be considered negligible.

Next, the results of the *dsmcFoam+* simulations are discussed, wherein the cylinder geometry is studied at various angles of incidence in both LDWT derived and "full-scale" (i.e. derived from re-entry trajectories) freestream conditions. These simulations cover Knudsen numbers in the range $0.02 \leqslant Kn < 10$, and angles of incidence, α, in the range $0^o \leqslant \alpha \leqslant 90^o$ (in 15^o increments, i.e. $\alpha = \{0^o, 15^o, 30^o \ldots 90^o\}$). GPR modelling was utilised extensively in this part of the numerical analyses to interpolate values of normalised heat flux, $\dot{q}/\dot{q}_0$, between the selected Knudsen numbers and angles of incidence. Data from various regions of interest on the surface of the cylinder are extracted and analysed using this method; these are the stagnation line, and either side of the cylinder's forwardmost circular edge (see Figure 3.22).

From these data, the transition to Newtonian heating distributions as $Kn \to \infty$ is examined, and it is concluded that the full onset of this behaviour does not occur on the cylinder until $Kn > 10$. Peak values of normalised heat flux are shown to reduce with increasing Knudsen number, and the angles of incidence at which each location reaches its maximum $\dot{q}/\dot{q}_0$ for a given value of Kn are also determined.

It may be concluded from the various presented GPR models that this correlation strategy is extremely useful in the analysis of engineering data, as the majority of the resulting models are both highly accurate, and computationally inexpensive to evaluate. The exceptions to this conclusion are the GPR models computed based upon the front and side portions of the stagnation

line. These models are shown to accurately reproduce the rough trend of $\dot{q}/\dot{q}_0$ with α and x/R_n, but fail to properly infer the variations that occur with differing Kn. This is a side effect of the extremely small differences which were observed at $Kn > 1$. Since the effects of this variable on the values of $\dot{q}/\dot{q}_0$ are extremely small, the GPR models are unable to properly incorporate them into their predictions, despite the use of an ARD type covariance function.

As such, it may also be concluded that GPR modelling is suitable for the present application, but that sub-modelling strategies such as those utilised for the aerodynamic models and thermal distributions on the front and side circular edges must be utilised if there are a small number of sample points. Note that, in the case of the partially successful normalised heat flux models, 20 simulations are used to correlate $\dot{q}/\dot{q}_0$ with respect to 3 variables, Kn, α, and x/R_n, while the sub-model augmented GPR models eliminate the spatial dimension, allowing correlation over Kn and α alone. Additionally, it was noticed during the formulation of these models that the aforementioned augmentation results in significantly smaller file sizes, and shorter training and evaluation times when compared to other models trained using higher dimensionality, i.e. a larger number of correlation variables.

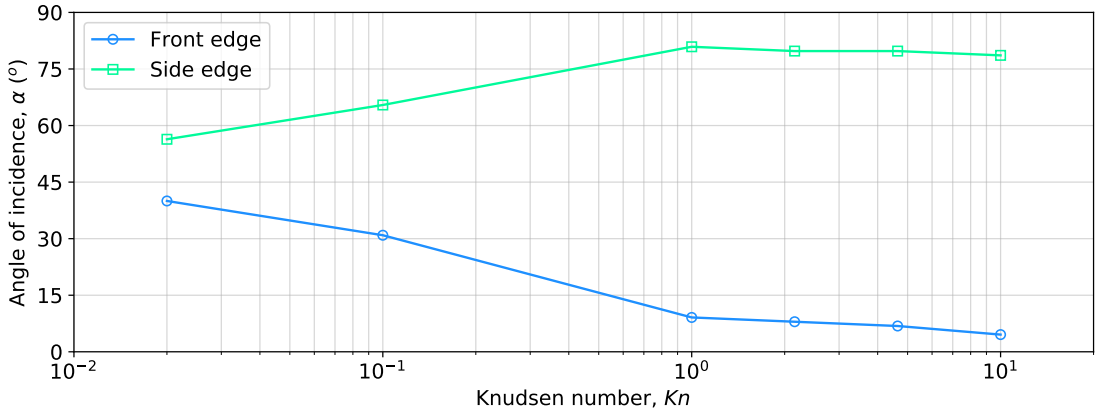


Figure 3.50: Angles of incidence, α, at which peak values of normalised heat flux occur in cylinder GPR models (based upon DSMC simulation data) for a range of Knudsen numbers, Kn.

The various GPR interpolation models may be queried in order to find the true angle of incidence, α, at which the peak normalised heat fluxes occur at a given Knudsen number. The results of such a calculation are presented in Figure 3.50 for the cylinder's front and side circular edges. It may be observed that the trends identified earlier in this section result in a logarithmic trend as $Kn \to \infty$ (which is linear when evaluated over $log_{10}(Kn)$) for both the front and side

edges. However, only the front edge GPR predictions converge to the expected value, i.e. $\alpha = 0^o$. Extrapolation of this trend returns a value of $Kn = 127.6$ at which the front edge's peak heat flux will occur at $\alpha = 0^o$.

The side edge, as mentioned previously, demonstrates a negative trend wherein the angle of incidence at which the peak normalised heat flux occurs actually decreases in the range $1 \leqslant Kn \leqslant 10$. Referring to Figure 3.46, it may be observed that this inaccuracy is a result of the relative "flattening" of the trend of $\dot{q}/\dot{q}_0$ at $\theta = 180^o$ as $\alpha \to 90^o$. As has been previously discussed in reference to the stagnation line GPR models, the effects of Knudsen number in the range $1 \leqslant Kn \leqslant 10$ are extremely slight when compared to the other correlation variables. This fact, coupled with the aforementioned flattening of the trend of $\dot{q}/\dot{q}_0$, makes the interpolation of $\lim_{Kn\to\infty} \alpha \left(\max\left(\dot{q}/\dot{q}_0\right)\right)_{CSE}$ challenging for the GPR model to calculate, even with augmentation. The front edge, on the other hand, has a well defined peak at all examined values of Kn (see Figure 3.39), making it far more suitable for the aforementioned extrapolation. Hence, it may be concluded that the limit of $Kn = 127.6$ which was arrived at for the onset of Newtonian heating distributions is correct, given the available data.

3.6 CHAPTER SUMMARY

Overall, in this chapter, novel models of aerodynamic coefficients for satellite debris analogues have been generated and evaluated, and an in depth study of the thermal flux distributions on hemispheres and cylinders has been performed.

It may be concluded that the aerodynamic models are well conditioned, and demonstrate impressive accurate for a wide range of flight conditions, attitudes, and geometric variations of the examined debris analogues, and provide extremely computationally efficient aerodynamic databases for future re-entry trajectory calculations. The thermal modelling of the hemisphere geometry has also been shown to be extremely useful in evaluating the suitability of various correlations for use in future work.

Finally, the study of thermal flux distributions on the cylinder geometry has shown both the advantages and disadvantages of the selected GPR modelling approaches, and has provided a great deal of insight regarding the onset of Newtonian heating behaviour as $Kn \to \infty$. Future

space debris re-entry simulation codes will be able to take advantage of these data in their estimations of heating to cylindrical components in the slip and transition regimes, hence improving the overall accuracy of satellite debris demise calculations. The data from simulations which were based upon LDWT derived freestream conditions will also be utilised in subsequent analyses of experimental results from the LDWT.

EXPERIMENTAL FACILITY 4

This chapter details the wind tunnel facility used for the present study. A description of the main experimental hardware, supporting equipment, and instrumentation are presented, including the wind tunnel facility itself, its pumping systems, and test chamber. This is followed by descriptions of supporting equipment such as the traversing mechanisms (used for flow characterisation and the adjustment of experimental probe attitudes), the gas conditioning system (which controls the flowrate and temperature of the working gas used to generate the experimental plume), flow initialisation equipment (which allows near instantaneous impingement of the hypersonic plume on experimental probes), and instrumentation (which allows the test conditions to be monitored and recorded during experiments). This is followed by determinations of the leak rates into the pumping system and test chamber, and their impact on the capabilities of the facility.

The spatial variation of the hypersonic plume is then described with measurements taken using the aforementioned traversing mechanisms, a differential Pitot probe (for post-shock total pressure measurements), an aspirated thermocouple (for total temperature measurements), and the sensors which monitor the stagnation conditions in the gas conditioning system.

Simulations of the experimental plume performed using CFD are presented in order to better understand the structure of the flowfield in the experimental facility, and are compared to the aforementioned Pitot measurements. These simulations are used to predict the Mach and Knudsen numbers of the flowfield, and hence select regions of interest in the flow where experimental probes may be situated in order to take advantage of specific conditions in the experimental plume.

Finally, the temporal variability of freestream conditions in the flow is examined in order to characterise the uncertainty that should be attributed to experimental conditions during the analysis of results.

4.1 PRESENT EXPERIMENTAL WORK

As discussed in Section 2.2.5, accurate calculations of the total heat load require simulation not only of the peak heating phase of descent, but also of the earlier stages of an atmospheric entry

where an object is already experiencing heating (known as the the "thermal pre-soak"). The spatial distributions of heat loads in this regime vary with Knudsen number, Kn, (see Section 2.2.4) but are relatively independent of Mach number (see Section 2.2.3).

The focus of the present work is the study of convective heat flux distributions. This class of heat transfer tends to dominate the thermal environment experienced by objects as they pass through the transition regime of entry into planetary atmospheres from low ($h < 200\,km$) orbits. As such, it is not necessary to achieve the high-enthalpy conditions of a 'real' re-entry flow while studying them. Instead, normalised heat fluxes are of interest, whereby the convective heat transfer to a re-entering object would be normalised to a reference stagnation value, and heat transfer coefficients derived.

Many factors affect heat fluxes in high energy regimes such as re-entry flow, including radiation, surface catalycity, ablation, etc. Although these can all have a significant impact on overall heat flux at different stages of a re-entry trajectory, they do not impact the convective heat transfer distribution. Instead, it is the level of rarefaction in the flow that impacts the distribution of convective heat fluxes. Owing to this dependence, (see Section 2.2.4), the freestream Knudsen number, Kn, was the main focus when selecting hardware and designing experiments. It was also recognised that the freestream Mach number, Ma, would not be a significant driver in the development of normalised heat fluxes, as long as experiments were performed close to or beyond the hypersonic limit (see Section 2.2.3).

Additionally, the desired thermal measurements require that the flow be initialised in such a way that a step change in heat transfer i.e. a step change in freestream temperature, is generated. This may be accomplished in a number of ways:

- Rapidly heating an established flow in which a test model is situated.
- Establishing a stable heated flow and physically moving a test model into the plume.
- Establishing a stable heated flow and deflecting it away from a test model, whereby the obstruction may be rapidly removed, allowing the flow to suddenly impinge upon the test piece.

The above methods may also be employed using preheated models and a comparatively cold flow, whereby the direction of heat transfer is reversed. This approach has previously been

utilised in the LDWT by Owen [87], while the inverse, i.e. a cooled model inserted into a heated flow, was previously used by Dahlen [130].

In summary, the conditions required to verify the distribution of heat fluxes in this regime are $1 \times 10^{-3} \leqslant Kn \leqslant 1 \times 10^{-1}$ and $Ma > 5.0$. A step change in freestream temperature must also be achieved.

4.2 THE OXFORD LOW DENSITY WIND TUNNEL

The Oxford Low Density Wind Tunnel (LDWT) is an open circuit, low enthalpy wind tunnel designed to produce hypersonic, slip regime flows. Unlike shock tunnel facilities, the LDWT is capable of operating for many hours uninterrupted, and can generate gas plumes with Mach numbers in the range $3.0 \leqslant Ma \leqslant 10.0$ and Knudsen numbers in the range $0.02 \leqslant Kn \leqslant 0.4$ (based on a characteristic length of 10 mm), depending on the pressure and temperature of the test gas. The current work utilises a single contoured nozzle designed by Peter Rockit to nominally produce a Mach 6 flow (see Section 4.2.5).

4.2.1 System description

The LDWT is split into three main sections, each of which has a distinct function within the system. These are the pumping system, the test chamber, and the gas conditioning system.

As their names suggest, the test chamber is designed to house all of the experimental equipment required for aerothermal experimentation, the pumping system maintains the low pressures required for slip regime studies, and the gas conditioning system preheats the test gas and maintains constant inlet mass flow rates. The test chamber interior measures $\sim \varnothing 1.54\,m$ in diameter by $\sim 1.52\,m$ in depth. Note that this latter measurement does not include the concave features at either end of the chamber, i.e. the rear wall and access hatch, set into each of which are inspection windows. The ultimate pumping speed of the facility is 22000 litres per second, which may be achieved as long as the test chamber pressure is below $\sim 20\,Pa$. These pumping speeds are achieved by the pumping system, which consists of a large vapour booster pump, backed by a Roots type mechanical booster, which in turn is backed by three rotary piston pumps

connected in parallel.

While in operation, the internal pressure in the test chamber is typically in the range 0.7 - 3 Pa when a gas flow of between 5 and 30 SLPM (between ∼0.1 g/s and ∼0.6 g/s) is injected into the wind tunnel via the gas conditioning system. Without an inlet flow, the test chamber pressure reaches a minimum of ∼0.01 Pa.

The system's layout and major components are described in Figure 4.1.

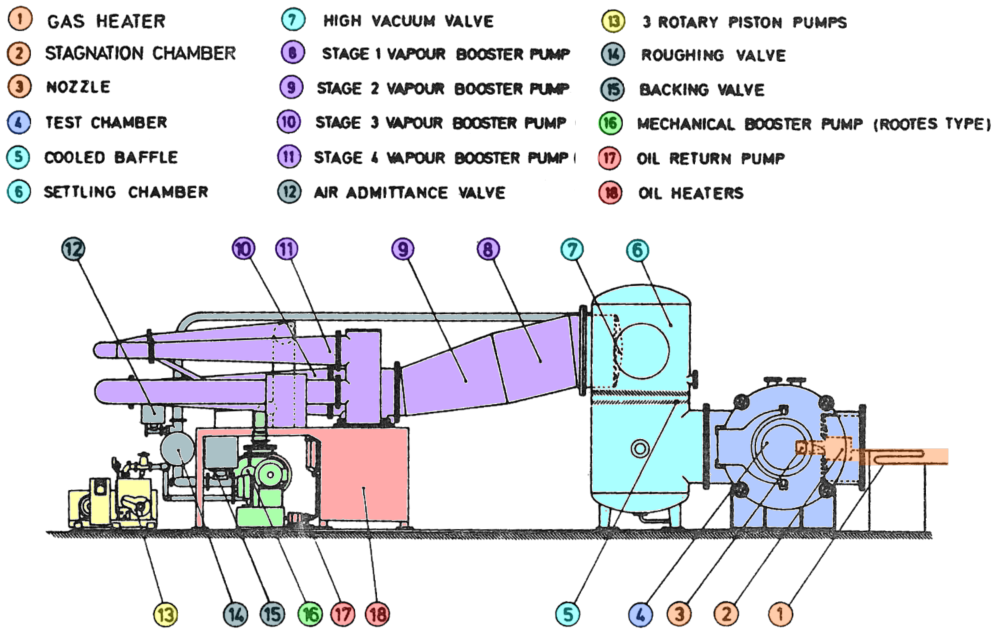


Figure 4.1: Diagram of the LDWT showing its major components [131]. Individual subsystems are colour coded for clarity.

4.2.2 Refurbishment

Aside from improvements to the gas conditioning and DAQ systems (which are described in Section 4.2.7), the LDWT also required extensive refit and recommissioning activities to be performed prior to the experimental campaign.

The following list summarises these refurbishments, which were ongoing for a three year period (in parallel with experimental preparation and numerical studies):

- All manually serviceable O-rings and unused equipment flanges replaced with new nitrile seals and 316 stainless steel blanks (respectively) in order to improve the facility's leak rate.

- Cooling water manifold installed near stagnation chamber.
- Seal kits obtained for future refurbishment of outmoded rotary piston pumps.
- Vapour booster internals (comprising the boiler and vapour ejector assembly) inspected visually and cleaned using industrial solvents.
- Asbestos insulation elements removed and replaced with a modern, non-carcinogenic alternative.
- Vapour booster oil drained and refilled with fresh batch (Apiezon AP201).
- Rotary piston lubricant drained and refilled with fresh batch (Edwards Ultragrade 17).
- Outdated and/or custom components researched and documented. Engineering drawings and CAD models were also produced for remanufacture should the need arise in the future.
- Control system components inspected, replaced, and re-calibrated (including coolant temperature sensor, control relays, vapour booster boiler setpoint, and emergency automatic heater cutoff switch).
- Traverse interface flange designed, manufactured and installed.
- Operating procedure updated and illustrated for novice users.

4.2.3 Test chamber layout

The LDWT test chamber is large enough to allow multiple pieces of supporting equipment and instrumentation to be installed. Past users of the tunnel have used a variety of devices including an analogue motor traverse, a magnetic suspension system, and a liquid nitrogen cooling jacket. A diagram showing the current internal layout of the test chamber is presented in Figure 4.2, below.

The traverse support frame was designed and manufactured for the present study, and allowed for the repositioning of experimental equipment relative to the exit of the nozzle using a set of heavy duty roller bearings. The support frame was designed such that it could be translated back and forth in order to accommodate experimental equipment of a wide range of sizes.

The configuration shown in Figure 4.2 was used to position compact traversing equipment close to the exit of the contoured nozzle.

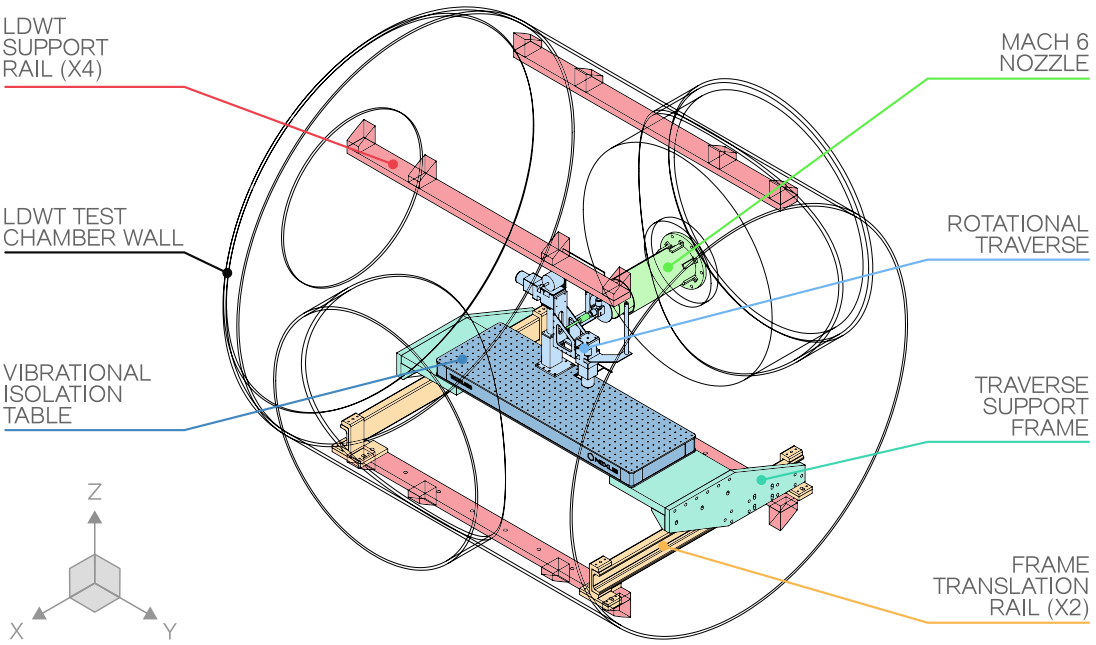


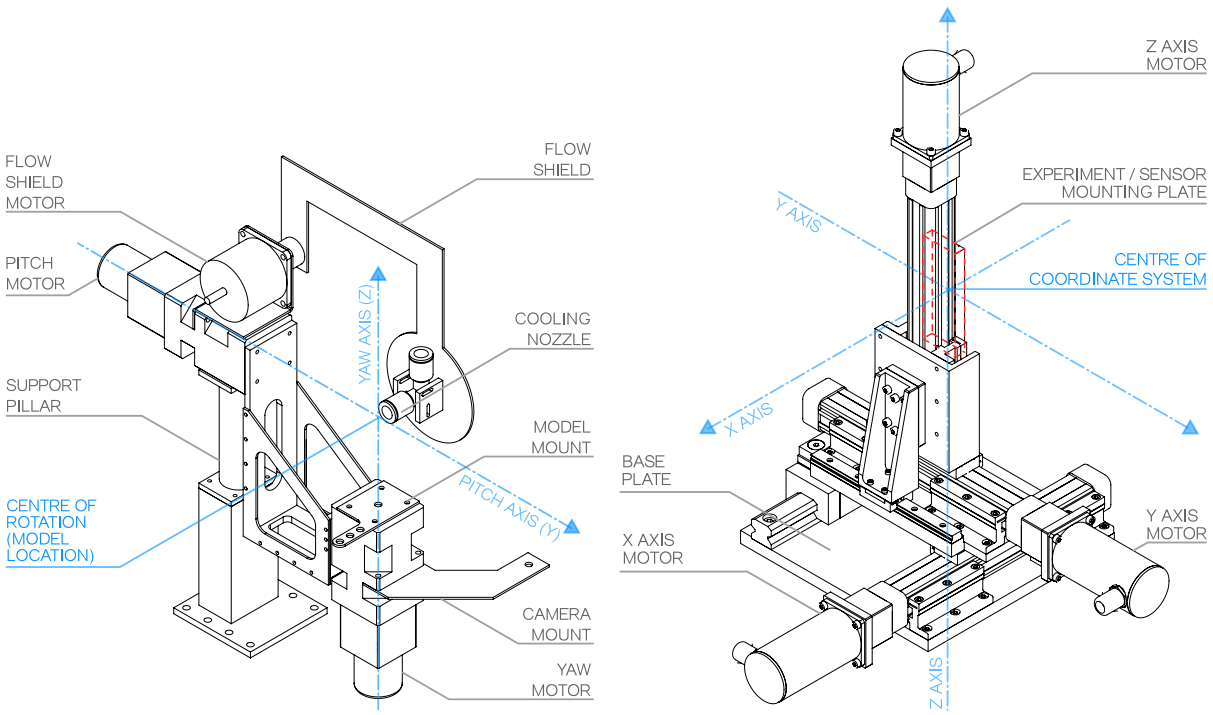
Figure 4.2: Schematic of the LDWT test chamber in its pitch/yaw traversing configuration, showing major components. Note that gas flow is in the positive X direction.

4.2.4 Traverse mechanisms

In order to characterise conditions within the test section, and hence adequately position test pieces, a pair of traversing mechanisms were designed and manufactured by Micromesh Ltd., a company external to the Osney Lab.

The first, shown in Figure 4.3(a), is designed to rotate thermal flux probes about the pitch and yaw axes (labelled), allowing sweeps of 90^o whilst keeping probes in the centre of the wind tunnel's flow field. It is also fitted with a flow deflection device (labelled "flow shield" in Figure 4.3(a) which is further described in Section 4.2.6 (the requirement for this feature is explained in Section 2.4.1).

The second, shown in Figure 4.3(b), is designed to translate sensors along the X, Y, and Z axes (labelled), allowing a 110 mm cube of possible positions to be explored. Both traverses share power and controller equipment.



(a) Rotational traverse (b) Linear traverse

Figure 4.3: Schematics of the new LDWT traversing mechanisms.

4.2.5 Gas conditioning

As is mentioned at the beginning of Section 2.4.1, the study of heat transfer is central to the present work, and as such it is necessary to introduce a temperature differential, ΔT, between the flow and any geometries under test. This was accomplished previously by Owen [87] via heating of test geometries, and the introduction of air at ambient temperature into an unheated stagnation chamber. Dahlen [130] used a similar approach, both heating the flow and cooling the model. This approach has been chosen for the present work as it allows a greater variety of materials and instrumentation types to be considered when constructing test geometries.

Conditioning of the inlet gas is accomplished through use of a sealed stagnation chamber developed by Dahlen [130]. Two heaters (one inline and one tape) are used to heat the gas to a specified inlet temperature, and were operated via a software defined PID controller.

The LDWT Mach 6 nozzle was designed by Peter Rockit, and was previously utilised by both Dahlen [130] and Owen [87]. This nozzle is utilised for the present work such that a Mach number beyond the hypersonic limit may be achieved (see Section 2.2.2) whilst still maintaining a reasonable degree of uniformity within the plume. The nozzle has an exit diameter, $\varnothing_{ex}$, of

$104\,mm$, and a throat diameter, $\varnothing_{th}$ of $9.65\,mm$, giving an overall area ratio, A_{ex}/A_t, of 116.15.

The inlet mass flow rate is governed by a PID driven mass flow controller connected to the OTI's internal air supply. This air is dehumidified by a centralised drying unit in order to mitigate the thermodynamic effects of water vapour content, and is distributed throughout the lab at a pressure of $100\,PSI$.

4.2.6 Flow initialisation

The stagnation chamber used in the present study was also originally designed to initialise the flow entering the LDWT test chamber. During a typical test, the linear valves shown in Figure 4.4 may be manually operated in order to block the main nozzle, and allow flow to enter the test chamber via the bypass nozzle. This theoretically facilitates stable pressures throughout the system until a test is due to begin, at which point, the linear valves are operated in unison, essentially switching the flow from the bypass nozzle to the main nozzle.

When recommissioning the stagnation chamber for the present work, it was noticed that a great deal of trial and error was required in order to seamlessly perform this switching operation. If the nozzles were not switched in unison, a test run would begin with a corresponding spike or drop in stagnation pressure (depending on whether the operator closed one valve before opening the other fully, or vice versa). This deviation from stable conditions was noted to take several seconds to stabilise, by which time the thermal transients of a typical test would have been well underway.

In order to remove this uncertainty in initial flow conditions, the flow shield shown in Figure 4.3(a) was developed and fitted to the rotational traverse. The purpose of this flow shield is to divert the flow exiting the contoured nozzle around any test models affixed to the rotational traverse until the start of a test, at which point the shield is quickly removed from the flow field. This then allows the already stable test plume to impinge upon a test model in a more consistent manner than could be accomplished with the linear valves in the stagnation chamber.

The flow shield is actuated with a stepper motor, and is able to vacate the area in front of the nozzle flow within ~ 0.5 seconds. This is roughly one order of magnitude smaller than the test time used during the experimental studies, thereby minimising any effects of the flow shield's

motion on the results of the heat transfer experiments.

4.2.7 Instrumentation

Control and monitoring of the environment within the LDWT is accomplished via several different pieces of equipment. Figure 4.4, below, shows the major components of the test chamber used for conditioning of the test gas, as well as the instrumentation used to record the test conditions during operation of the tunnel.

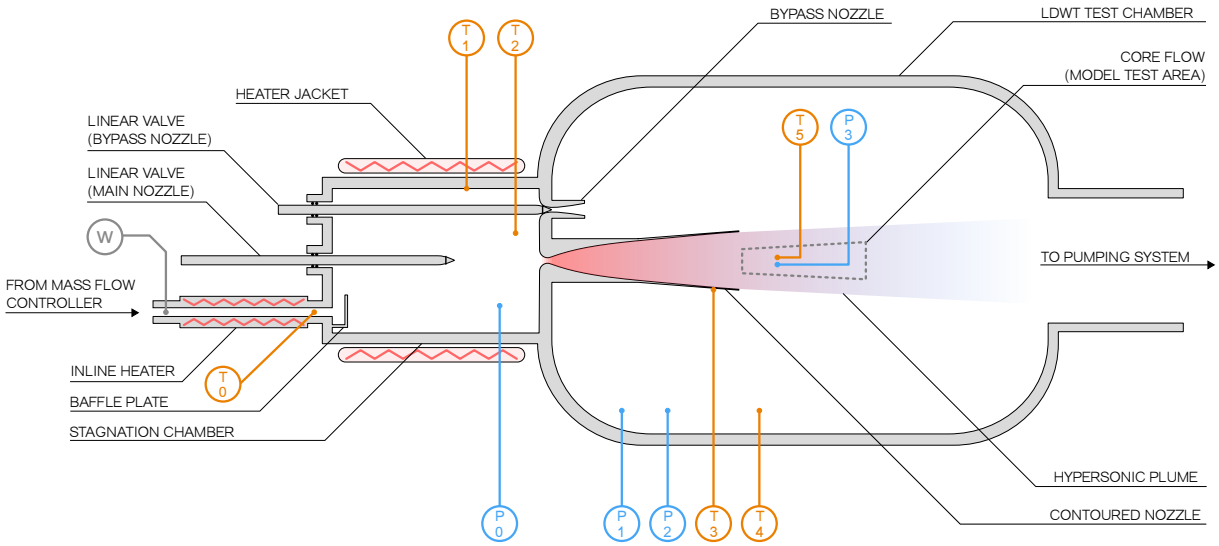


Figure 4.4: Schematic view of instrumentation in the LDWT test chamber and gas conditioning system. Labels p and T refer to numbered entries in Table 4.1.

The instrumentation in the LDWT test chamber was updated during the refurbishment summarised in Section 4.2.2. A series of pressure transducers and thermocouples were installed in order to monitor the system during operation of the tunnel, as well as characterising the testing environment in the nozzle plume. The locations of the installed instrumentation are presented in Figure 4.4. The purpose of these sensors, as well as the variables they measure, their manufacturers, model numbers and standard measurement uncertainties, are presented in Table 4.1.

Pressure transducer P0 and thermocouples T0, T1, and T2 are used to monitor the conditions in the stagnation chamber, where T0 is installed at the exit of the inline heater, T1 is welded to the interior wall of the chamber itself, and T2 is suspended inside the chamber ($3/4$ of a radius from the wall) through a stainless steel tube. Measurements from T0 and T1 are used as inputs

for the software defined PID controllers that control the inline and tape heaters (respectively), while T3 is used to log the total temperature of the gas in the stagnation chamber prior to its expansion through the nozzle.

Conditions in the test chamber are monitored by pressure transducers P1 and P2, and thermocouples T3 and T4. T3 is affixed to the exit of the contoured nozzle in order to monitor the nozzle wall temperature, while T4 is suspended in the chamber far from the test plume, and is used to log the test chamber background temperature.

Pressure transducer P3 and thermocouple T5 are used in conjunction with the linear traverse described in Section 4.2.4 to characterise the flowfield generated by the contoured nozzle. P3 is fixed to the traverse support frame, and is connected to the Pitot probe via a flexible hose. T5 is part of an aspirated thermocouple whose design is based upon equipment developed by Widodo and Buttsworth [132]. This piece of instrumentation is constructed from a hollow tube of Rohacell, and features a thin-wire ($1/1000$ inch diameter) K-type thermocouple whose weld bead is situated at the exit of the tube's opening. As such, the total temperature of the flow may be measured in the subsonic region behind the normal shock generated by the tube's presence in the flow, thereby avoiding difficulties inherent in measuring total temperatures in high Mach number flows.

Label	Measurement	Symbol	Sensor	Location	Range	Accuracy
W	Mass flow rate	$\dot{m}$	Omega FMA-2609A	Stag. chamber	$0 - 50\,SLPM$	$\pm 0.8\% + 0.2\%$ FSR
P0	Stag. pressure	p_0	MKS Baratron 122A	Stag. chamber	$0 - 13300\ Pa$	$\pm 0.5\%$ FSR
P1	BG pressure	p_∞	Edwards Pirani APG100-XLC	Test chamber	$0.01 - 100000\ Pa$	$\pm 15\%$ FSR @ $< 1000\,Pa$
P2	BG pressure	p_∞	Inficon CDG025D	Test chamber	$0 - 133\ Pa$	$\pm 0.2\%$ FSR
P3	Pitot pressure	p_{0_2}	First Sensor 144LP05D-PCB	Nozzle rake	$0 - 500\ Pa$ (Diff.)	$\pm 0.05\%$ FSO
T0	Stag. temperature	T_{inline}	K thermocouple	Stag. chamber	$3 - 987\,K$	$\pm 2.2^oC$ or $\pm 0.75\%$
T1	Stag. temperature	T_{wall}	K thermocouple	Stag. chamber	$3 - 987\,K$	$\pm 2.2^oC$ or $\pm 0.75\%$
T2	Stag. temperature	T_0	K thermocouple	Stag. chamber	$3 - 987\,K$	$\pm 2.2^oC$ or $\pm 0.75\%$
T3	Nozzle temperature	T_{nozzle}	K thermocouple	Test chamber	$3 - 987\,K$	$\pm 2.2^oC$ or $\pm 0.75\%$
T4	BG temperature	T_∞	K thermocouple	Test chamber	$3 - 987\,K$	$\pm 2.2^oC$ or $\pm 0.75\%$
T5	Stag. temperature	T_{0_2}	K thermocouple	Nozzle rake	$3 - 987K$	$\pm 2.2^oC$ or $\pm 0.75\%$

Table 4.1: Summary of LDWT test section instrumentation. Note that the abbreviation "BG" refers to the background pressure of the LDWT test chamber.

4.2.8 Leak rate determination

Ideally, there would be no appreciable leak rates in the LDWT. Due to the facility's size, and the number of separate seals between its components however, some ingress of gas into the system is inevitable. As the LDWT has a finite pumping capacity, leaks in the test chamber and pumping system detract from the ultimate allowable input mass flow rate and test chamber pressure (and, hence, the properties of the test plume). Therefore, it is important to minimise the leak rate, and characterise the effect of any remaining leaks on the performance of the system.

The leak rate in the LDWT test section and pumping system were determined by measuring the pressure rise over time once the pumping system had been shut down. This was performed initially with the high vacuum valve closed, allowing the test section leak rate to be calculated. The measurements were then repeated with the high vacuum valve open in order to find the total leak rate and, by subtracting the former from the latter, the pumping system leak rate.

The internal volumes of the LDWT test section and pumping system were calculated using CAD models, which were generated using a combination of the facility's original schematics and manual measurements. Leak rates then were calculated using the rate of change of pressure over time to arrive at a mass flow rate.

The results of this calculation over a period of several minutes were used in subsequent calculations to include the effects of leaks in performance estimations. A mean leak rate of $0.0024\,g/s$ was calculated for the test chamber. A similar calculation for the pumping section leak rate returned a mean value of $0.1469\,g/s$. As is mentioned in Section 4.2.2, the leak rates in the LDWT were minimised during refurbishment via the replacement of numerous seals in the system. At the beginning of this process, the leak rates in the test chamber and pumping system were 0.25 g/s and 4.36 g/s, respectively. Note that the leak rate is so much higher in the pumping system due to the comparatively large number of couplings and gas flow connections (and hence, a larger number of small potential leaks which must be sealed).

The test chamber leak rate is approximately two orders of magnitude lower than the range of inlet mass flow rates used in the present study, while the pumping system rates are of the same order of magnitude. Although this does not initially sound at all favourable, it was found during characterisation activities that these leaks are present mostly in the backing pump segment of the

pumping system, and not the vapour booster itself. As such, the high volumetric flow rates that facilitate the correct operation of the LDWT are not severely impacted, reducing by a maximum of 15% compared to the original design specifications.

Leaks in the test chamber, however, were found to significantly affect the performance of the facility during characterisation and refit activities. This is due to the fact that these leaks, unlike those mentioned above, are present in parts of the facility upstream of the vapour booster, and as such are capable of impacting the achievable pressures and mass flow rates in the test chamber. In order to characterise these effects, an analytic model of the facility's performance was developed. This model is presented in Appendix B.

4.3 SPATIAL VARIATION OF TEST FLOW

Flow characterisation activities in the LDWT comprised total Pitot surveys and total temperature measurements downstream of the nozzle. These are compared to axisymmetric CFD simulations of flow both upstream and downstream of the nozzle exit. The goal of this work is to determine the spatial dependence of the flow generated in the LDWT test section on the system's input variables ($\dot{m}_1$ and T_0). These properties play a key role in the design of experiments, as well as the post-processing of results.

4.3.1 Pitot measurements

The first activity during the characterisation of the LDWT was a series of Pitot surveys, performed using a differential (with respect to the chamber background pressure, p_∞) pressure transducer (see Table 4.1) mounted on the linear traverse (see Section 4.2.4).

The results of an initial set of Pitot surveys conducted at an inlet mass flow rate, $\dot{m}_1$ of 25 SLPM are shown in Figure 4.5. These surveys were performed along three planes centred on the middle of the nozzle plume. The terms "radial" and "axial" are used in subsequent sections to refer to surveys which were performed along the intersections of the Y-Z and X-Y planes, and X-Z and X-Y planes, respectively.

The conditions for all subsequent characterisation CFD studies were derived from the quantities presented in Table 4.2, which were averaged over the course of the Pitot studies.

Results of the radial and axial Pitot surveys across the conditions in Table 4.2 are presented in Figures 4.6 and 4.7, respectively. Error envelopes were calculated by applying the manufacturer's stated uncertainty for the transducers used to measure the Pitot pressure (P3 in Table 4.1) and the test chamber background pressure (P2 in Table 4.1)

The radial surveys show a region of constant Pitot pressure towards the centre of the plume, the size of which increased with inlet mass flow rate, $\dot{m}_1$. This region — termed the "core flow" — was shown to be stable at the nozzle exit for mass flow rates exceeding 10 SLPM. Total temperature, T_0, was found to have little effect in the spatial distribution of the plume, and only a small effect on the magnitudes of the measured Pitot pressures (whereby p_{0_2} increased with T_0).

The axial surveys showed that the Pitot pressure in the core flow dropped with increasing distance downstream of the nozzle exit. The variation of these data between values measured at the nozzle exit and 110mm downstream was $\sim 50\%$ for the lowest mass flow rate and total temperature, reducing to $\sim 15\%$ for the highest combination of $\dot{m}_1$ and T_0.

$\dot{m}_1$ (SLPM)	$\dot{m}_1$ (g/s)	$\overline{p_0}$ (Pa)		$\overline{p_\infty}$ (Pa)	
		$T_0 = 293$ K	$T_0 = 385$ K	$T_0 = 293$ K	$T_0 = 385$ K
5	0.099	812.2	872.3	0.58	0.70
10	0.198	1308.0	1430.3	1.14	1.17
15	0.296	1788.7	1961.8	1.45	1.52
20	0.395	2239.6	2746.5	1.75	1.79
25	0.494	2675.4	2980.1	2.33	2.32
30	0.592	3134.5	3470.7	2.71	2.94

Table 4.2: LDWT test quantities averaged over characterisation runs. Note that the stated values of T_0 are the set points used for the temperature controller in the stagnation chamber gas conditioning system, and that p_∞ is the chamber background pressure.

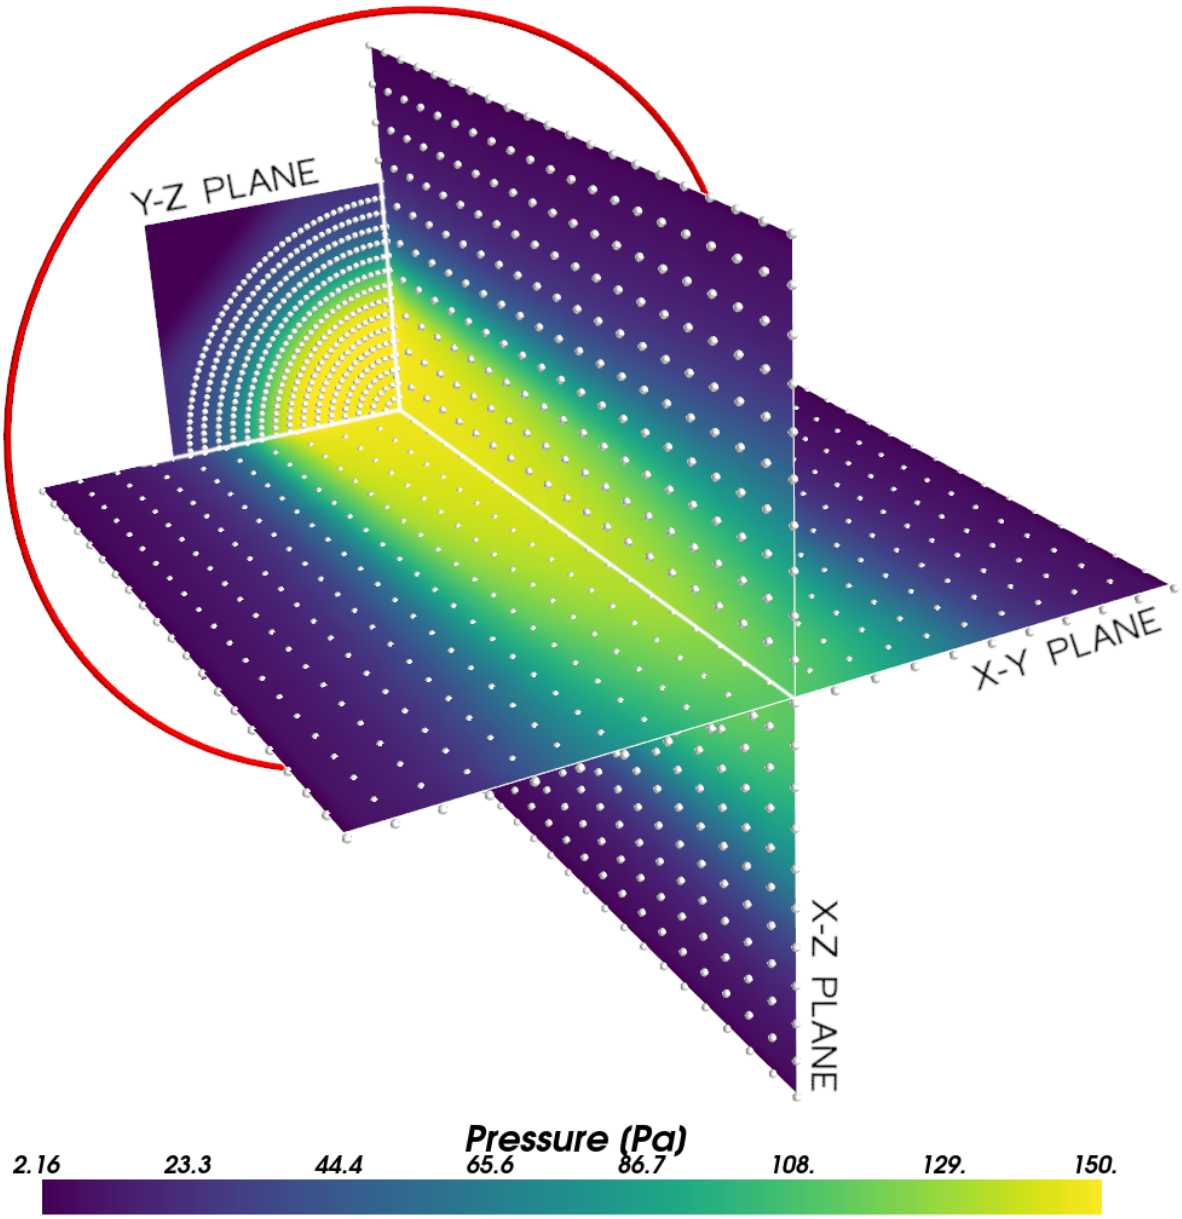


Figure 4.5: Pitot survey of the contoured nozzle plume inside the LDWT test chamber across planes in the x, y, and z directions. Survey sampling points are overlaid as white spheres, and the nozzle exit is represented by a red circle. This survey was conducted using an inlet mass flow rate, $\dot{m}_1$, of 25 SLPM.

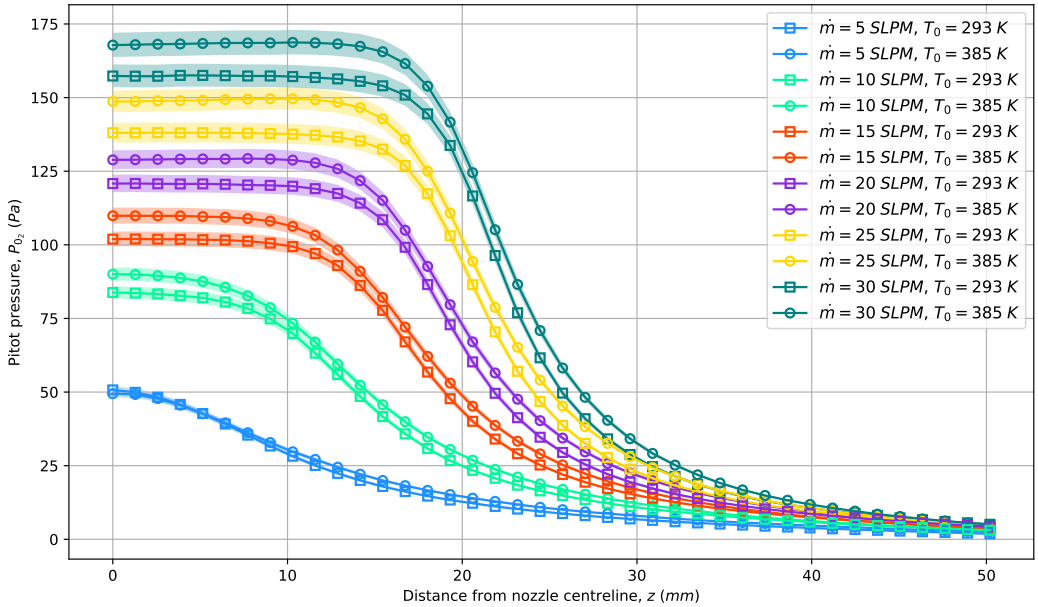


Figure 4.6: Radial Pitot survey of the contoured nozzle plume ($x = 0$). Measurement error is indicated by shaded regions.

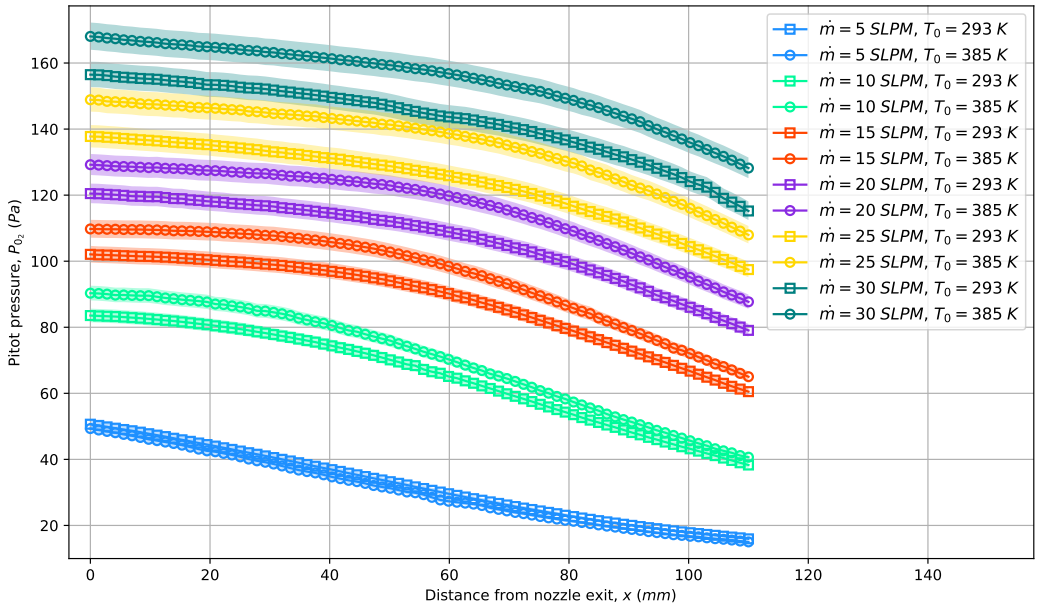


Figure 4.7: Axial Pitot survey of the contoured nozzle plume ($y = 0$, $z = 0$). Measurement error is indicated by shaded regions.

4.3.2 Stagnation chamber measurements

The average stagnation conditions for each of the experiments presented in Chapter 6 are shown in Figure 4.8. Very little scatter is visible for both test conditions, which were conducted over a two week period, with maximum relative standard deviations (RSD) of 0.26% for total temperature, T_0, and 0.22% for stagnation pressure, p_0. As such, it may be concluded that the stagnation conditions of the LDWT are repeatable.

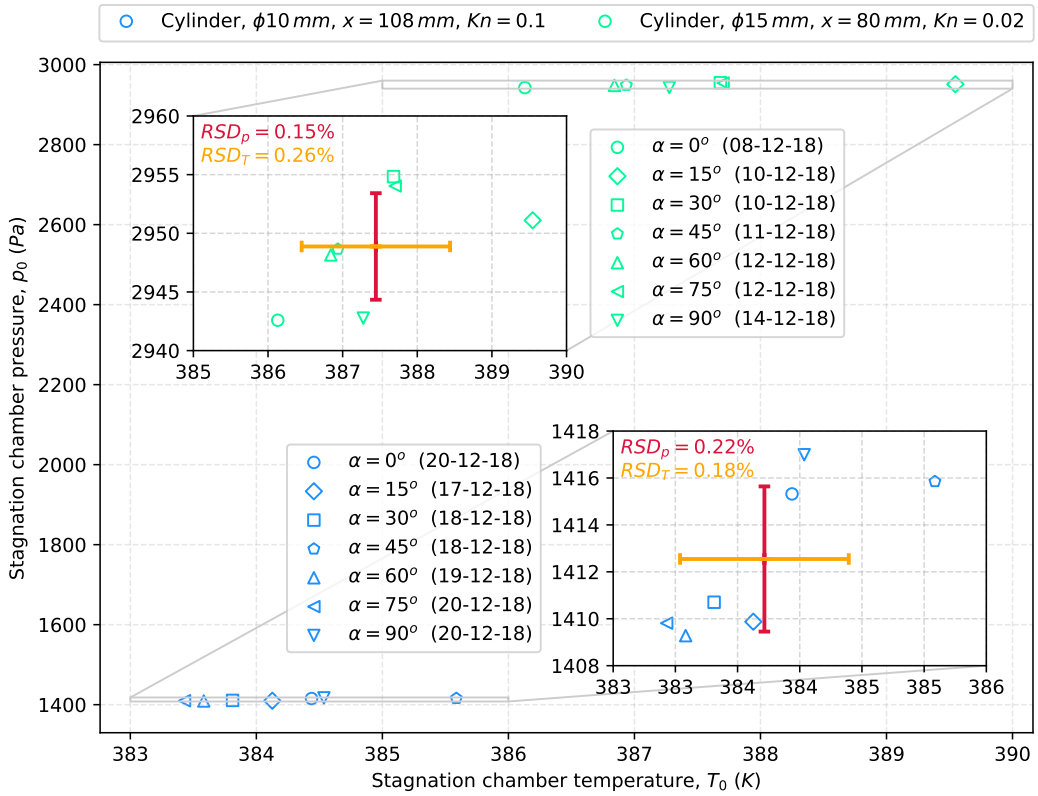


Figure 4.8: Average stagnation conditions in the LDWT stagnation chamber recorded during LC experiments. These experiments are described in Sections 2.5.3 and 2.5.4, while their results are presented in Chapter 6.

4.3.3 Total temperature measurement

Total temperature measurements were taken at the centre of the nozzle exit in order to determine the flow recovery factor. This experiment utilised the aspirated thermocouple and flow shield described in Section 4.2.6. The temperature measurements of the LDWT thermocouples are presented in Figure 4.9 and show the readings from both the aspirated thermocouple and those of its cylindrical body near the thin-wire junction. Also plotted are the data from the gas conditioning system thermocouples, which show the periodic rise and fall of the stagnation

chamber and inline heaters, as well as the gas temperature near the nozzle inlet.

Finally, the corrected thermocouple temperature is plotted, which was computed according to the rod heating model of Carslaw and Jaeger [133] using the junction temperature to adjust the readings of the thin-wire thermocouple. A recovery factor of 0.998 was also applied, as suggested by Widodo and Buttsworth [132]. This correction was introduced in order to correct for the transient response of a homogeneous, constant cross-section wire of high aspect ratio (i.e a wire significantly thinner than it is long) in a convective environment, once again, as suggested by Widodo and Buttsworth. The maximum difference between the measured and actual temperatures was found to be $\sim 1.25\,K$, which is less than the standard uncertainty of a K type thermocouple ($\pm 2.2\,K$).

A sudden rise of temperature in the aspirated thermocouple is evident at the beginning of the plot. This trend was recorded following the removal of the flow shield from in front of the nozzle, constituting the instantaneous initialisation of the flow (note that the rapid cooling recorded immediately following this was a result of the shield being once again lowered into the flow). Following this test of the flow shield's ability to thermally shield a test piece from the plume, it was once again removed and the temperature readings allowed to rise to a steady state (where $t \gtrsim 250\ s$).

In order to further study the behaviour of the aspirated thermocouple, readings were also taken during the cooling of the gas conditioning system. This involved simply turning off power to the inline and stagnation chamber heaters, whilst maintaining the mass flow rate of gas into the stagnation chamber. These data are presented in Figure 4.10, and show the cooling rates of the various components in the gas conditioning system. As expected, the aspirated thermocouple readings gradually come into line with those of the interior of the stagnation chamber, implying that there are no scaling or offset issues in the data.

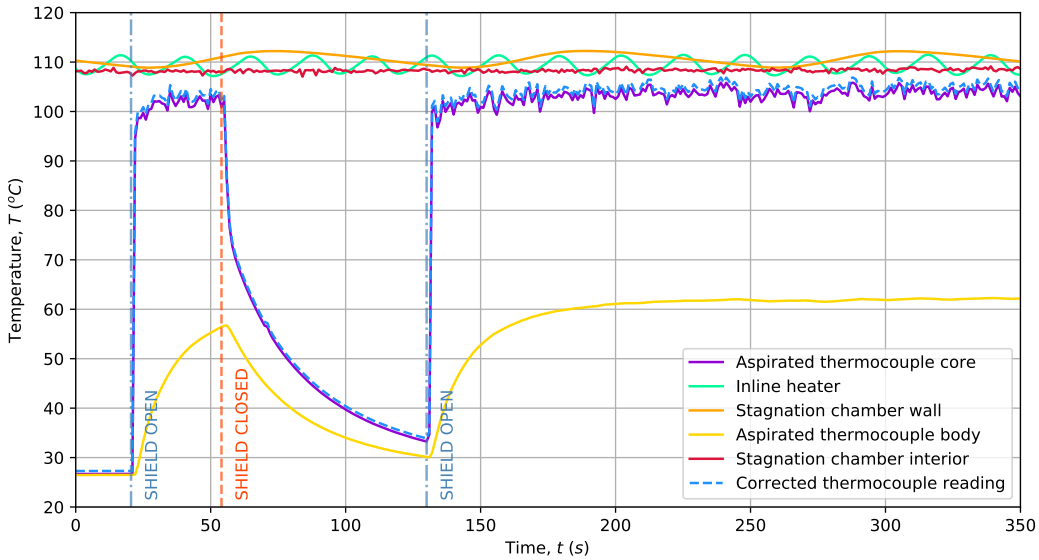


Figure 4.9: Thermocouple readings at the exit of the LDWT nozzle showing variations in total temperature during the removal and replacement of the flow shield in the LDWT test chamber.

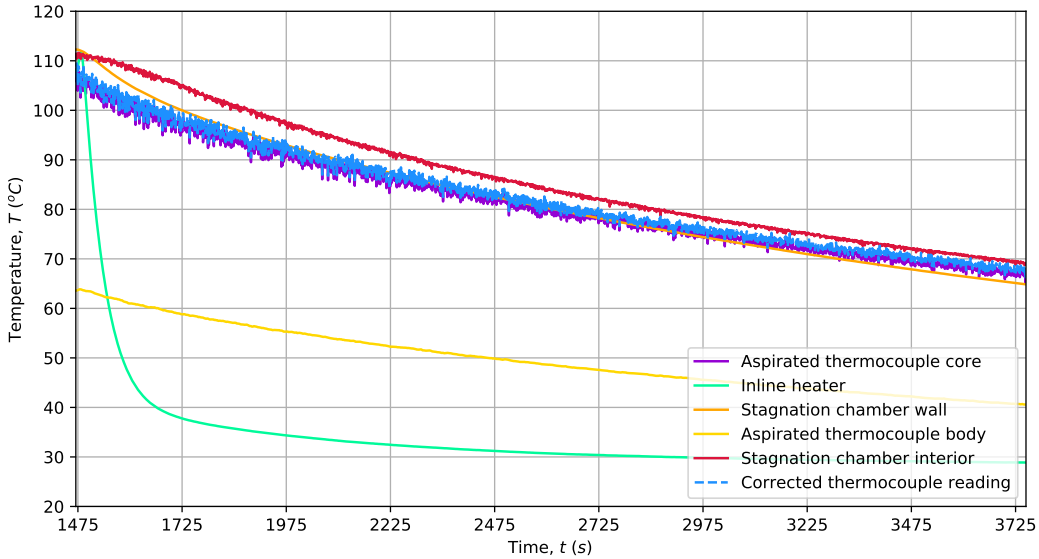


Figure 4.10: Thermocouple readings at the exit of the LDWT nozzle showing variations in total temperature following the shut-down of heaters in the gas conditioning system.

4.3.4 Plume simulation

Simulations of the flow in the contoured nozzle were performed in order to better model the spatial dependence of the flow properties that were not measured directly (such as Mach number, Ma, and static temperature, T) on the input conditions of the LDWT. Although certain assumptions can be made in order to allow the calculation of these variables without resorting to CFD simulations, it was desirable to model them as accurately as possible such that positions of

interest in the flow could be identified for subsequent thermal flux experiments.

The simulations were performed using the RANS solver of ANSYS FLUENT, and were run for the same conditions at which the Pitot surveys were taken. The boundary conditions for the simulations were based on total and static pressures and utilised the automatic CFL (Courant–Friedrichs–Lewy) number adjustment algorithm (referred to as "solution steering" in FLUENT) in order to aid convergence. These boundary condition data are presented in Table 4.2, and the settings used for the simulations are summarised in Table 4.3.

Setting/parameter	**Value**
FLUENT version	19.2
Energy equation	On
Solver formulation	Density
Geometry formulation	Axisymmetric
Gas thermochemical model	Air, NIST derived (see Appendix A)
Numerical formulation	Implicit
Numerical scheme	Roe-FDS/AUSM
Turbulence model	None; laminar
Convergence criterion	1×10^{-6}
Solution steering	On; hypersonic preset

Table 4.3: Parameters and settings of the CFD simulations performed for the present work.

In order to permit a direct comparison between experimental and simulation data, the Pitot pressure, p_{0_2}, was calculated from the CFD results. Total pressure, specific heat ratio, and Mach number data (p_0, γ, and Ma, respectively) were extracted from simulations in the same locations as were traversed during the Pitot surveys. The formula for the total pressure ratio across a normal shock (Eq. (4.1)) was then used to calculate the expected Pitot pressure such that the CFD simulations could be effectively validated using the survey data.

$$\frac{p_{0_2}}{p_{0_1}} = \left[\frac{(\gamma+1)\,Ma^2}{(\gamma-1)\,Ma^2+2}\right]^{\frac{\gamma}{\gamma-1}} \left[\frac{\gamma+1}{2\gamma\,Ma^2-\gamma+1)}\right]^{\frac{1}{\gamma-1}} \tag{4.1}$$

4.3.4.1 *Definition of CFD domain*

Owing to a lack of documentation, the computational domain used for the CFD calculations was designed using measurements of the LDWT Mach 6 contoured nozzle as opposed to the design schematics or engineering drawings. Figure 4.11 shows the structure and boundary condition types used in the axisymmetric domain designed for the nozzle CFD simulations. Table 4.4 summarises the settings used for each of the boundaries.

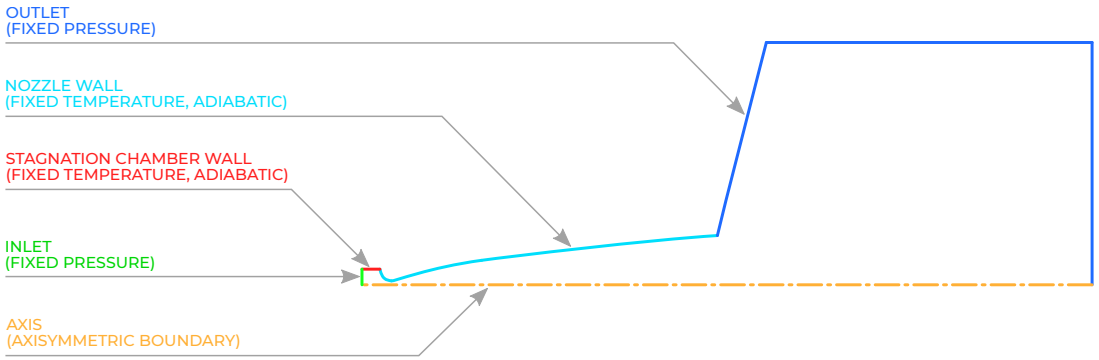


Figure 4.11: Diagram showing the boundary conditions set during CFD simulations of the LDWT nozzle plume.

Boundary name	**Type**	p_0	p_∞	T
Outlet	`pressure-outlet`	N/A	P2	$293\ K$
Axis	`axis`	N/A	N/A	N/A
Inlet	`pressure-inlet`	P0	P2	PID setpoint
Stagnation chamber wall	`wall (isothermal)`	N/A	N/A	PID setpoint
Nozzle wall	`wall (isothermal)`	N/A	N/A	$293\ K$

Table 4.4: Boundary condition settings used in CFD simulations of the LDWT nozzle flow. Values such as "P0" and "P2" refer to the instrumentation listed in Table 4.1. The phrase "PID setpoint" refers to the control system setpoint for the LDWT gas conditioning system, i.e. the controller for the inline and stagnation chamber heaters, either $293\ K$ or $385\ K$.

Pressure inlet and outlet boundaries were found to provide a good degree of stability and convergence in the simulations. As such, these boundary types were selected for the simulations, with the inlet being positioned just upstream of the nozzle throat, and the outlet covering all boundaries outside of the nozzle. The stagnation pressures, p_0, from LDWT measurements were

used as the total pressure for the inlet, and the test chamber background pressure measurements, p_∞, were used as the static pressure setting for both inlet and outlet.

The lowermost boundary was defined as an axis, facilitating the 2D axisymmetric setting in FLUENT to be enabled, thus greatly reducing the computational expense of the simulations. Finally, both the stagnation chamber and nozzle walls were configured as no-slip isothermal boundaries (the rationale for selecting this boundary type is presented in Section 4.3.4.2).

All boundaries which support temperature as a boundary condition were set to $293\,K$, with the notable exception of the inlet in the stagnation chamber, which was set to either $293\,K$ or $385\,K$ depending on which set of boundary conditions from Table 4.2 were being utilised as inputs.

4.3.4.2 Slip versus no-slip CFD modelling

Initially, both slip and no-slip boundary conditions were investigated due to the extremely low density of the expected flow. In modern CFD software, the no-slip boundary condition sets the fluid velocity at a wall to zero, allowing shear stresses to be calculated locally. A slip model inverts this behaviour, setting local shear stress to zero at the wall, and allowing velocity to be calculated instead.

Although ANSYS FLUENT features a partial slip wall formulation, it is only available for use in simulations utilising the code's pressure based solver (it is primarily intended for use in microfluidic simulations). As the flow in the LDWT is both highly rarefied and highly compressible, it was most desirable to utilise FLUENT's density based solver, which only features the traditional slip and no-slip models described above.

Figure 4.5 shows the results of two such simulations, both performed using the same boundary conditions (note that the whole computational domain is presented, whereas subsequent figures are cropped).

As expected, velocities near the nozzle wall in the slip simulation are non-zero, showing a boundary layer whose lowest Mach number is of the order of 50% that of the core flow. Referring to Figure 4.12, it can be observed that despite significantly different boundary layer thicknesses, the degree of plume radial expansion into the freestream is approximately equal between the slip

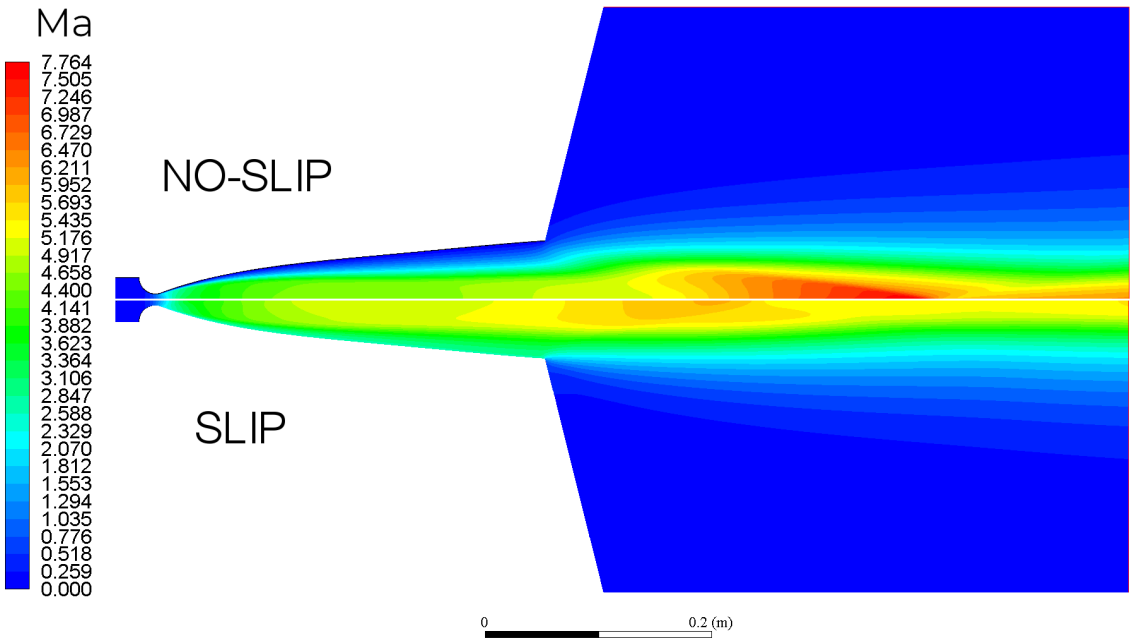


Figure 4.12: Mach number results from both slip and no-slip 2D CFD simulations of the plume from the LDWT contoured nozzle. The boundary conditions of the $T_0 = 385\,K$, $\dot{m}_1 = 30\,SLPM$ case presented in Table 4.2 are shown.

and no-slip simulations. The core region of the flow, however, is markedly smaller when using the no-slip model, and peaks further downstream of the nozzle exit. This is due to the decrease in effective area generated by the presence of a fully developed no-slip boundary layer.

Comparisons between slip wall CFD and experimental data gathered during Pitot surveys in the LDWT are presented in Figures 4.13 and 4.14 (radial and axial surveys respectively).

The radial Pitot survey data show that the boundary layer at the nozzle exit is distinct, with a well defined region of invariant pressure at the core. This is in contrast to the slip CFD simulations, which indicate a much less extreme change in flow properties towards the wall, and significantly lower core Pitot pressures. Similarly, the axial survey data show poor agreement with slip CFD simulations. Trends between different mass flow rates do not match those of the experimental data, and the magnitudes of the core flow Pitot pressures are underpredicted by as much as 45% by the slip model.

Figure 4.15 shows a comparison between experimental Pitot data, and results from a pair of CFD simulations. All data were obtained from the $T_0 = 385\,K$, $\dot{m}_1 = 25\,SLPM$ case. It is immediately apparent that the no-slip model reproduces both the trend and magnitudes of the

experimental data significantly more effectively than the slip model. From these comparisons, it was determined that a full-slip wall model was unsuitable for modelling the flow in the contoured nozzle.

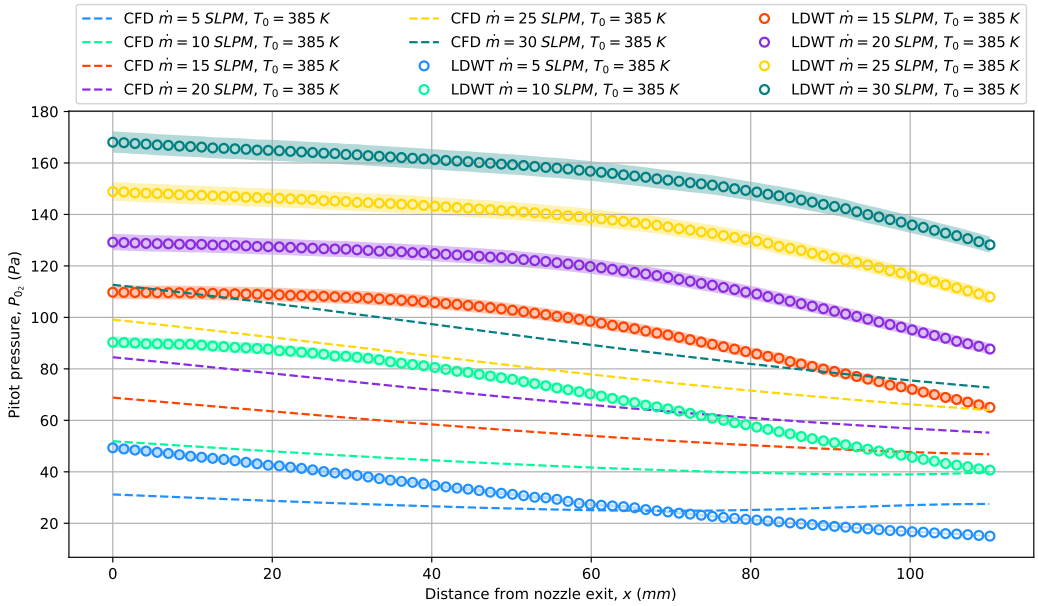


Figure 4.13: Axial Pitot survey of the contoured nozzle plume downstream of the nozzle exit ($y = 0$, $z = 0$). Measurement error is indicated by the shaded area around each series of data. CFD simulation results utilising a slip boundary as the nozzle wall are overlaid.

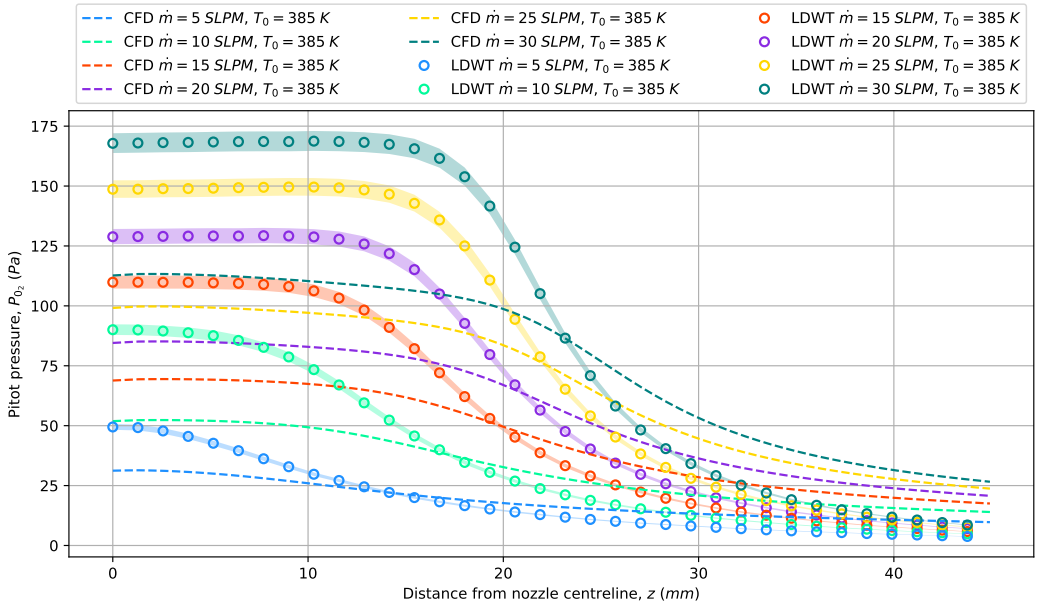


Figure 4.14: Radial Pitot survey of the contoured nozzle plume at the exit of the nozzle ($x = 0$). Measurement error is indicated by the shaded area around each series of data. CFD simulation results utilising a slip boundary as the nozzle wall are overlaid.

Note that, while a partial slip boundary condition is included in FLUENT, it is only accessible

when utilising the pressure-based flow solver, which is primarily intended for the study of microfluidic devices. Due to the highly compressible nature of the flow being modelled, it was necessary to utilise FLUENT's density-based flow solver. As such, a no-slip model was employed for all subsequent CFD simulations.

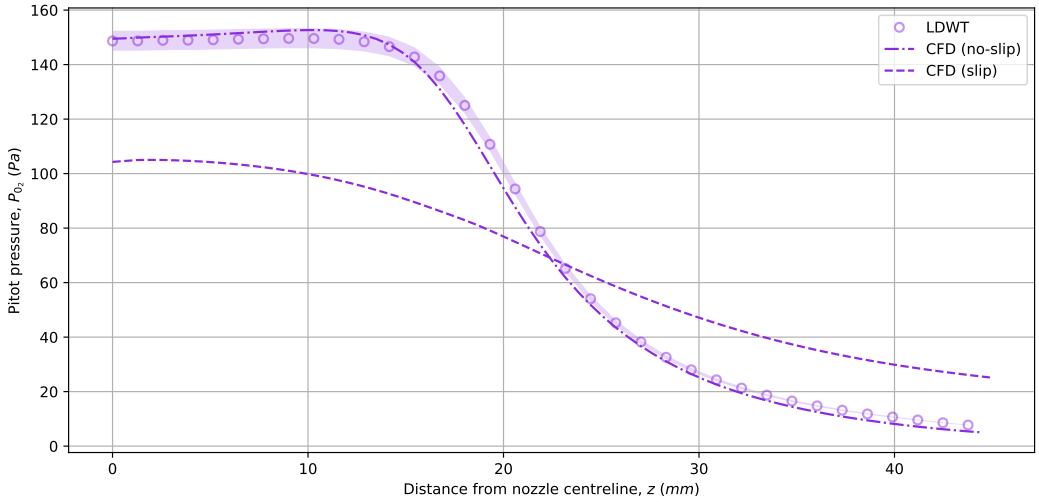


Figure 4.15: Radial Pitot survey of the contoured nozzle plume at the exit of the nozzle ($x = 0$) for the $T_0 = 385\,K$, $\dot{m}_1 = 25\,SLPM$ case. Measurement error is indicated by the shaded area around the data. CFD results from simulations utilising both slip and no-slip boundaries as the nozzle wall are overlaid.

4.3.4.3 Mesh independence

A mesh independence study was performed in order to verify that the CFD solutions were not influenced by the resolution of the computational domain. The same base mesh and boundary conditions were used for all simulations. As with the simulations described in Section 4.3.4.1, the boundary conditions used were those of the $25\ SLPM$, $385\ K$ case shown in Table 4.2. The mesh density was varied via the use of a refinement factor which was applied uniformly to all areas of the domain. Meshes with approximate cell counts of 27×10^3, 110×10^3, 250×10^3, 500×10^3, and 1×10^6 were evaluated. Plots of the mesh density for each domain showing the stagnation chamber and nozzle throat region are presented in Figure 4.16.

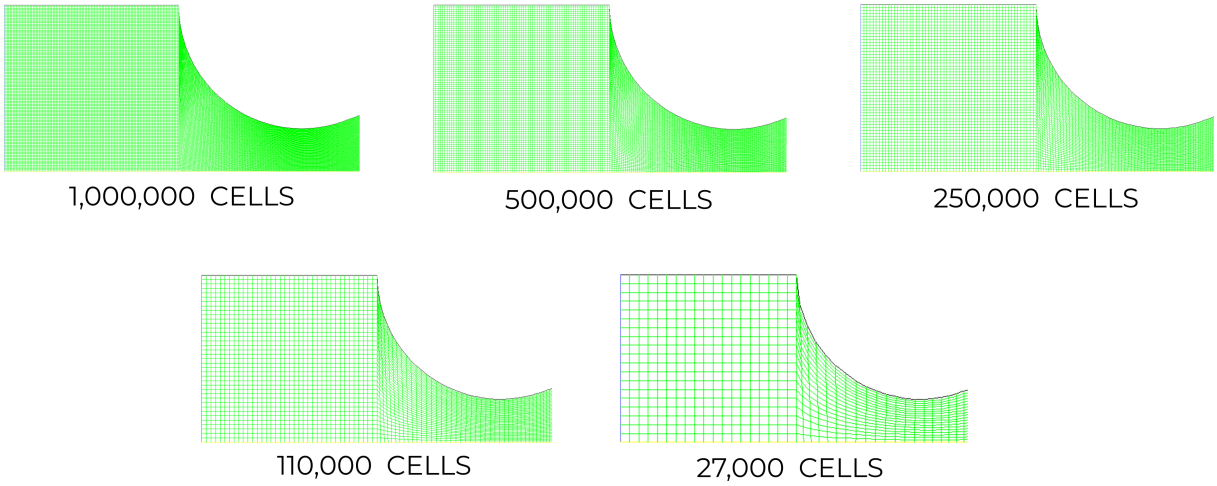
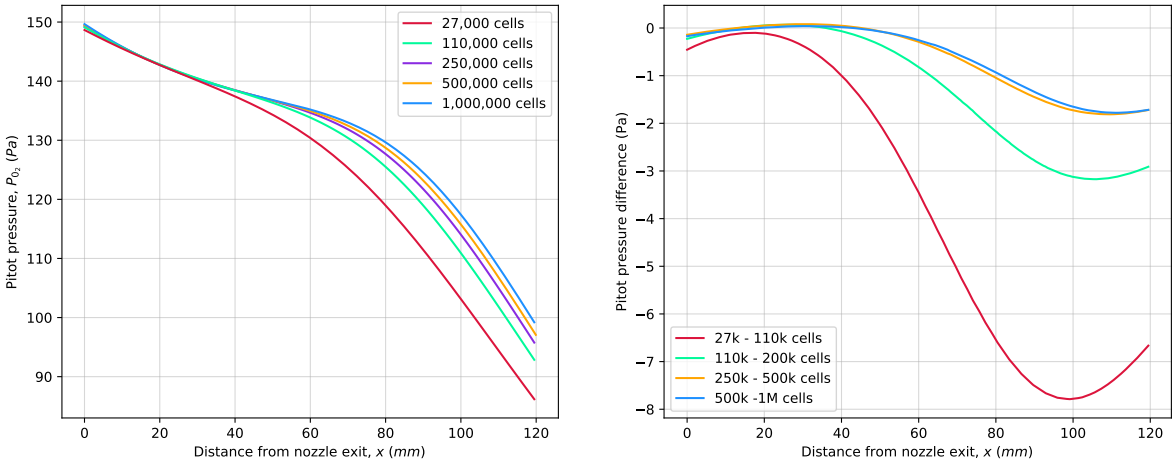


Figure 4.16: Plots of CFD mesh density in stagnation region and nozzle throat for simulations performed during the mesh sensitivity study. The mesh with $\sim 500,000$ cells was selected for all subsequent simulations.

In order to judge the effects of mesh refinement on the quality of the simulations, the Pitot pressure, p_{0_2} was calculated using Eq. (4.1) at locations of interest in the domain. Specifically, these are the plume centreline downstream of the nozzle exit ($y = 0\ mm$, $z = 0\ mm$), and the radial profiles at the nozzle exit ($x = 0\ mm$), and at $100\ mm$ downstream of the nozzle exit ($x = 100\ mm$).



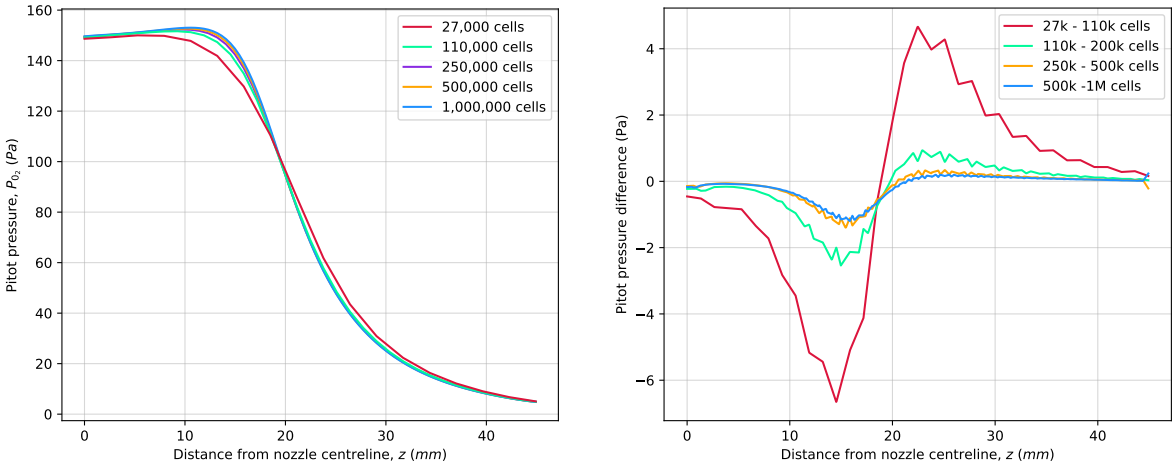
(a) Pitot pressure, p_{0_2} (Pa)

(b) Pitot pressure difference between adjacent mesh densities (error), Δp_{0_2} (Pa)

Figure 4.17: Pitot pressure variation between differing mesh densities downstream of the nozzle exit ($y = 0\ mm$, $z = 0\ mm$).

These data are presented in Figures 4.17, 4.18, and 4.19, respectively, as are the *differences*, Δp_{0_2}, between adjacent cases i.e. errors. It can be clearly seen that simulations using fewer than

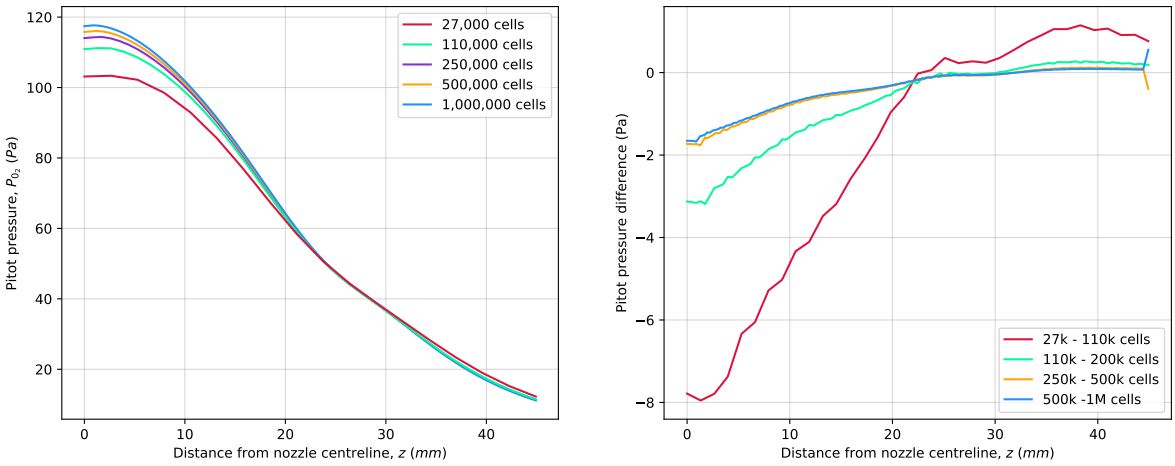
250,000 cells feature far larger discrepancies in their results than those using meshes with a higher cell count.



(a) Pitot pressure, p_{0_2} (Pa)

(b) Pitot pressure difference between adjacent mesh densities (error), Δp_{0_2} (Pa)

Figure 4.18: Pitot pressure variation between differing mesh densities at the nozzle exit ($x = 0$ mm).



(a) Pitot pressure, p_{0_2} (Pa)

(b) Pitot pressure difference between adjacent mesh densities (error), Δp_{0_2} (Pa)

Figure 4.19: Pitot pressure variation between differing mesh densities 100 mm downstream of the nozzle exit ($x = 100$ mm).

The error towards the centre of the plume is low in those cases with more than $27,000$ cells, although these errors become larger further downstream. The difference between the cases with $500,000$ and $1,000,000$ cells is, as expected, the smallest, and appears to reach an asymptote, with the maximum difference between the two cases being of the order of 2 Pa ($\sim 1.4\%$ of the maximum Pitot pressure reported by experiments).

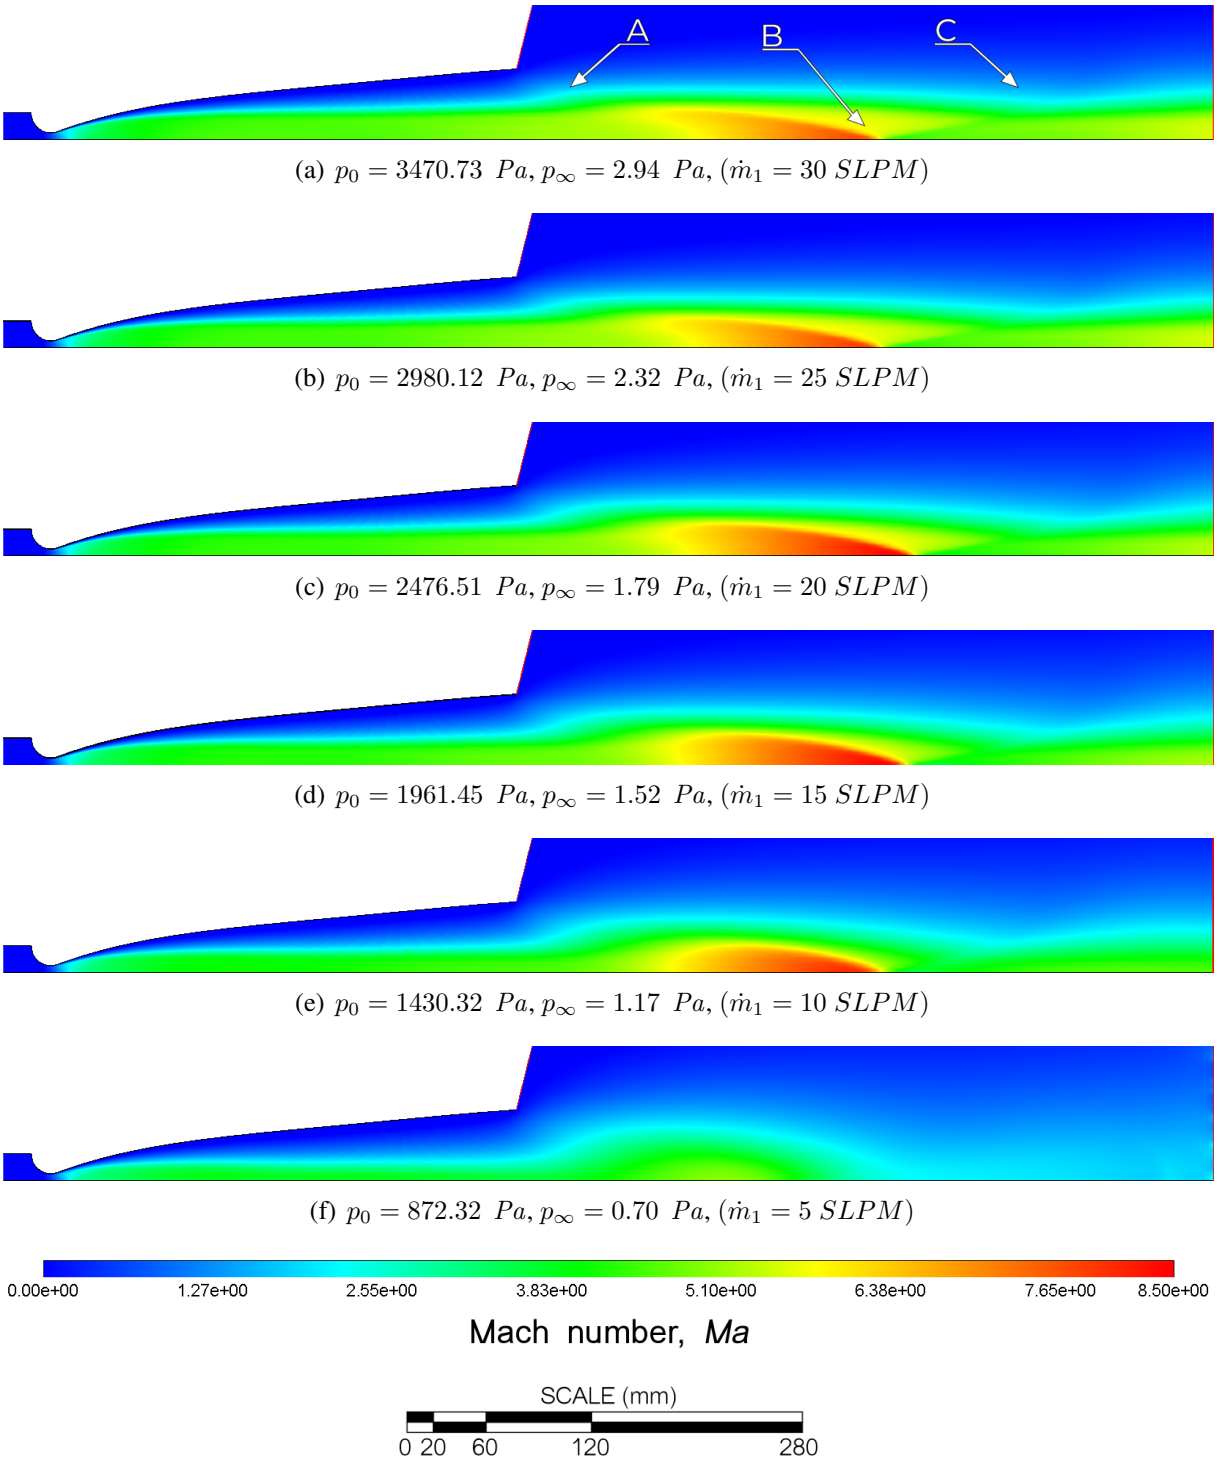
As such, the mesh with $\sim 500,000$ cells was selected for all subsequent simulations utilising the no-slip boundary at the nozzle wall. This choice was also motivated by the computational expense required to run the $\sim 1,000,000$ cell case, which would have increased overall simulation time for all of the required cases. The final cell count of the selected mesh was 441940.

4.3.4.4 No-slip CFD simulations

Simulations utilising a no-slip wall boundary condition were performed at all of the conditions presented in Table 4.2. The results of these are shown in Figures 4.21 and 4.20 for total temperatures, T_0, of $293\ K$ and $385\ K$, respectively.

As expected, the core region increases in size with increasing mass flow rate, $\dot{m}_1$, as does the maximum Mach number in the flow. Underexpansion of the plume is also visible (label "A" in Figure 4.20(a)), resulting in the formation of a Mach disk at the downstream extremity of the core flow region (label "B" in Figure 4.20(a)), and the formation of a shock diamond structure towards the end of the computational domain (label "C" in Figure 4.20(a)). Varying the total temperature appears to have little effect on the flow structure, with Mach number distributions remaining consistent between both sets of simulations.

Simulations performed using the no-slip wall model were found to agree reasonably well with experimental data, both in terms of integrated mass flow rate, and spatial distribution. The largest discrepancy between the data is evident in the radial survey (Figure 4.22), and represents an over-prediction of the boundary layer thickness by the CFD simulations. Considering the results of the slip regime simulations in conjunction with those of the no-slip cases, it is evident that the flow in the plume behaves in line with neither a full slip or no-slip boundary condition. It was suspected that a partial slip wall condition like that described at the beginning of this section would provide the best agreement with experimental data; however, the implementation of this model in the available CFD solvers was beyond the scope of the present characterisation study.



(a) $p_0 = 3470.73\ Pa, p_\infty = 2.94\ Pa, (\dot{m}_1 = 30\ SLPM)$

(b) $p_0 = 2980.12\ Pa, p_\infty = 2.32\ Pa, (\dot{m}_1 = 25\ SLPM)$

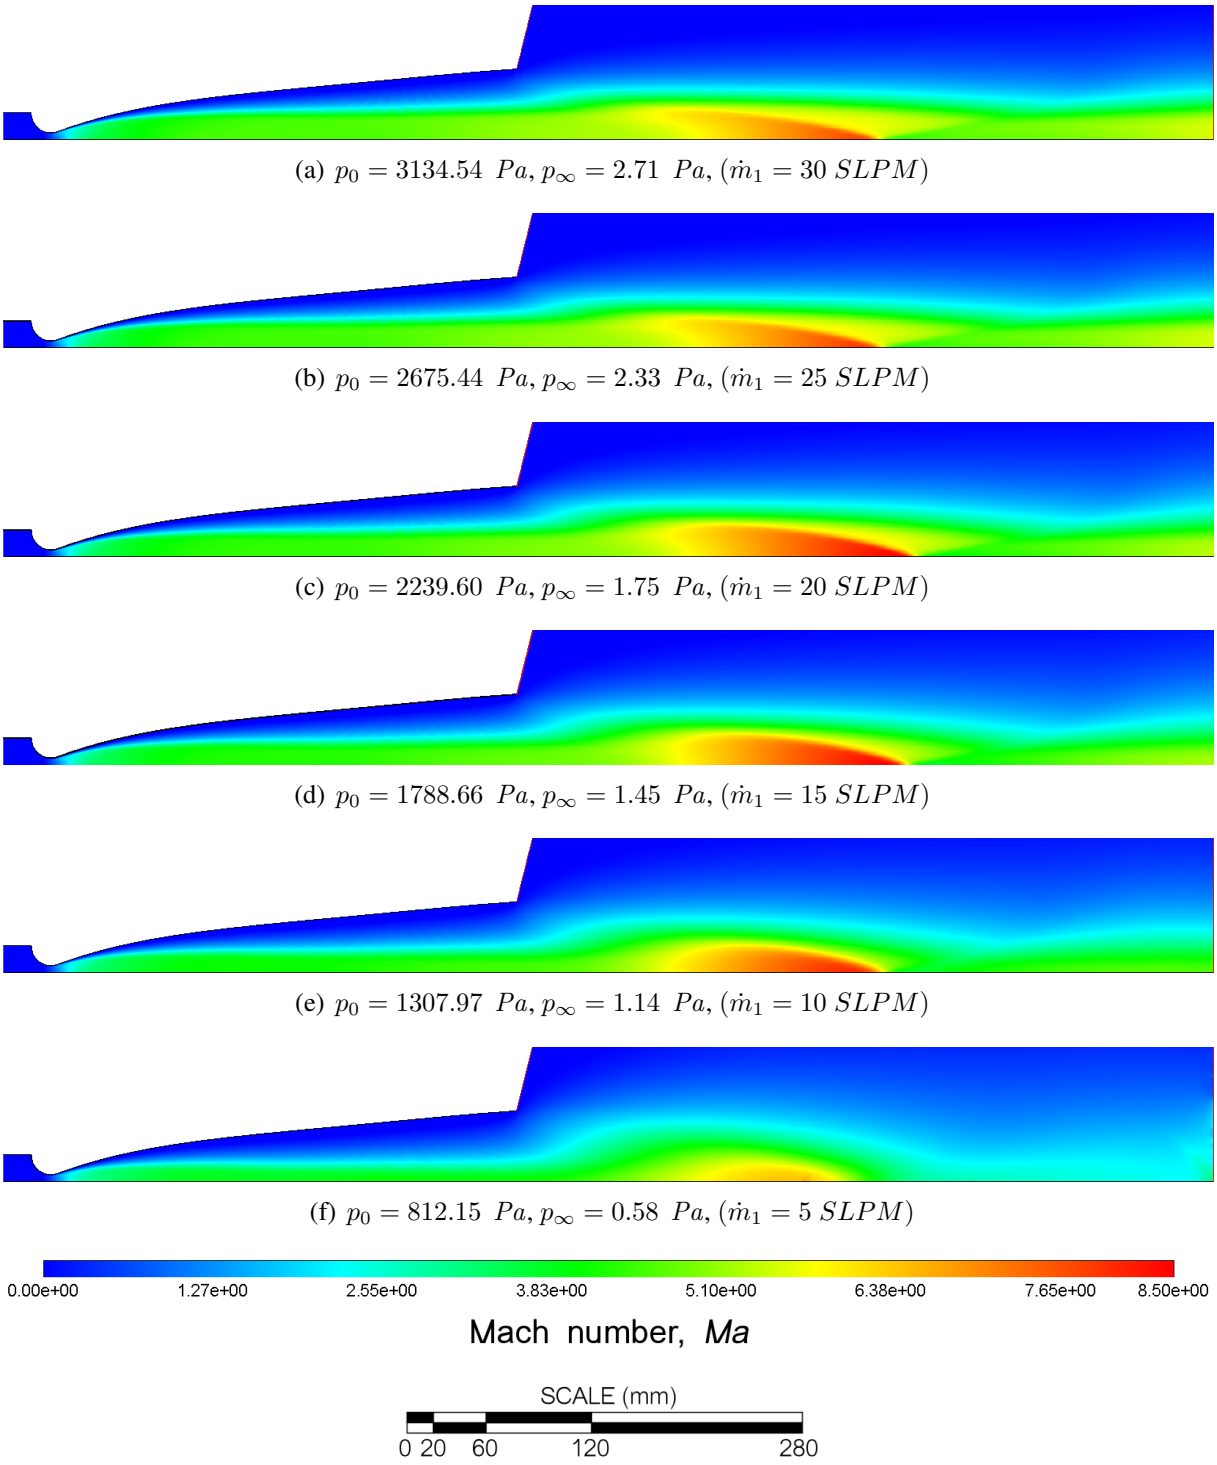
(c) $p_0 = 2476.51\ Pa, p_\infty = 1.79\ Pa, (\dot{m}_1 = 20\ SLPM)$

(d) $p_0 = 1961.45\ Pa, p_\infty = 1.52\ Pa, (\dot{m}_1 = 15\ SLPM)$

(e) $p_0 = 1430.32\ Pa, p_\infty = 1.17\ Pa, (\dot{m}_1 = 10\ SLPM)$

(f) $p_0 = 872.32\ Pa, p_\infty = 0.70\ Pa, (\dot{m}_1 = 5\ SLPM)$

Figure 4.20: Axisymmetric 2D CFD simulations of the plume from the contoured nozzle using a no-slip boundary condition, and a stagnation temperature, T_0, of 385 K. Plotted contours represent Mach number.



(a) $p_0 = 3134.54\ Pa, p_\infty = 2.71\ Pa, (\dot{m}_1 = 30\ SLPM)$

(b) $p_0 = 2675.44\ Pa, p_\infty = 2.33\ Pa, (\dot{m}_1 = 25\ SLPM)$

(c) $p_0 = 2239.60\ Pa, p_\infty = 1.75\ Pa, (\dot{m}_1 = 20\ SLPM)$

(d) $p_0 = 1788.66\ Pa, p_\infty = 1.45\ Pa, (\dot{m}_1 = 15\ SLPM)$

(e) $p_0 = 1307.97\ Pa, p_\infty = 1.14\ Pa, (\dot{m}_1 = 10\ SLPM)$

(f) $p_0 = 812.15\ Pa, p_\infty = 0.58\ Pa, (\dot{m}_1 = 5\ SLPM)$

Figure 4.21: Axisymmetric 2D CFD simulations of the plume from the contoured nozzle using a no-slip boundary condition, and a stagnation temperature, T_0, of 293 K. Plotted contours represent Mach number.

The results of the no-slip simulations were subsequently compared to the Pitot survey data in Figures 4.22 and 4.23, in the radial and axial directions, respectively.

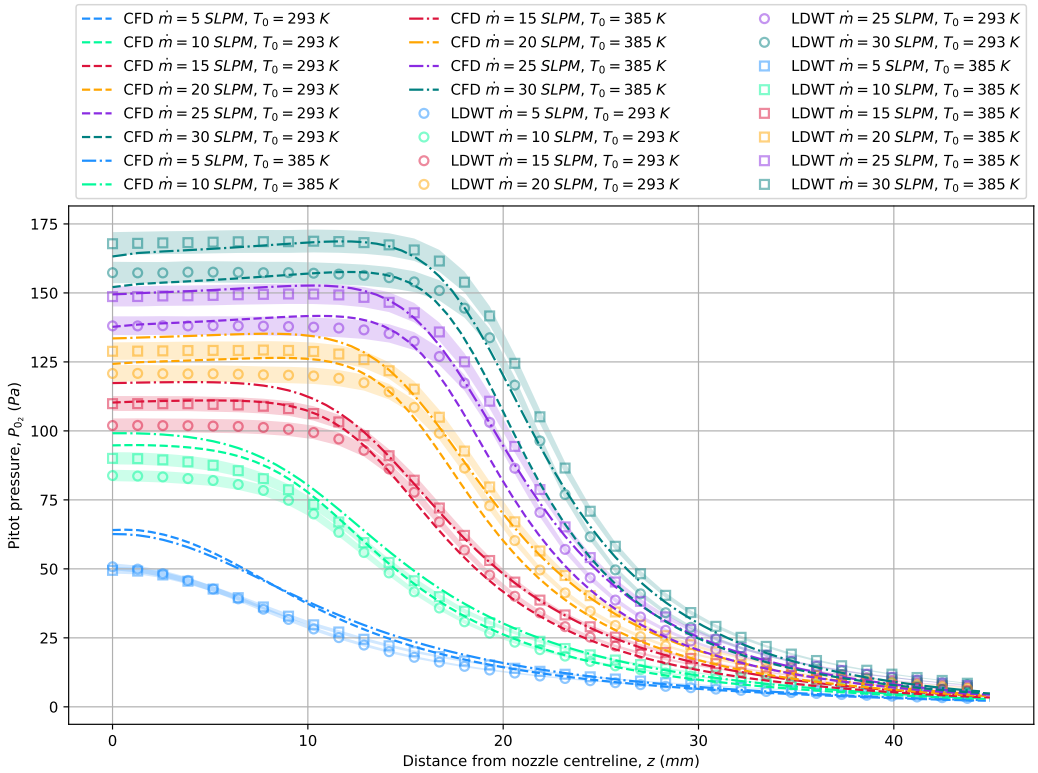


Figure 4.22: Radial Pitot survey of the contoured nozzle plume ($x = 0$). Measurement error is indicated by shaded area around each series of data.

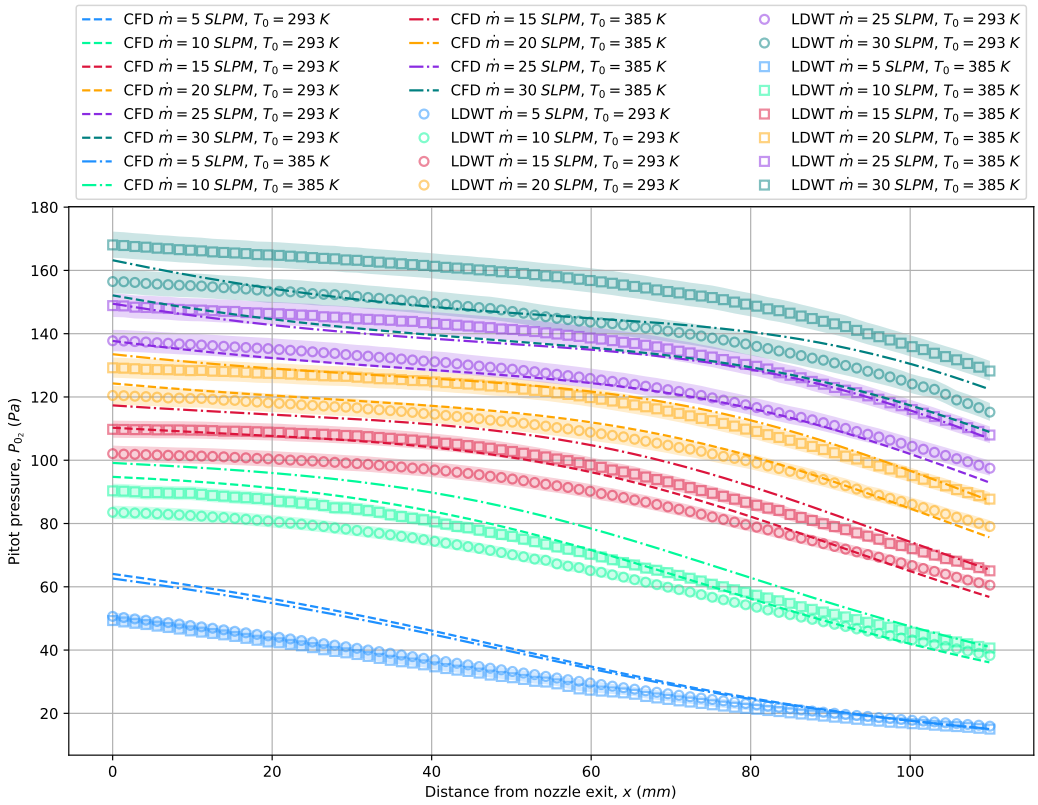


Figure 4.23: Axial Pitot survey of the contoured nozzle plume ($y = 0$, $z = 0$). Measurement error is indicated by shaded area around each series of data.

4.3.5 Mach and Knudsen number calculation

The Mach and Knudsen numbers in the experimental plume are of particular interest when examining the results of the CFD simulations, as these quantities are not directly measured by the present characterisation activities. Flowfield Mach numbers are output automatically by FLUENT, and Knudsen numbers were evaluated using Eq. (2.3) (where the characteristic length, L, was taken to be $0.01\ m$). These data are presented in Figures 4.24 and 4.25.

The characteristic length of $0.01\ m$ used in Eq. (2.3) was selected due to it being a desirable minimum diameter for thermal flux probes in the LDWT. Although probes smaller than this have previously been used in the LDWT, this limit was specified in order to improve the accuracy of post-processing when using thermochromatic coatings on probes in the LDWT (which rely on visual measurement).

Referring to Figure 4.24, it may be observed that a region of near uniform Mach number in the core flow region is present near the exit of the nozzle. This region increases in length with increasing mass flow rate, as do the Mach number magnitudes. Conversely, the overall extent of the variation in Mach number decreases, implying that a completely uniform flow field could potentially be obtained at suitably high mass flow rates. Increasing the total temperature of the flow can be seen to generate increases in core Mach number toward the exit of the nozzle which diminish further downstream (although this trend is reversed for mass flow rates lower than $15\ SLPM$ due to boundary layer growth).

Another approach for generating this uniform region would involve connecting the mass flow controller to pressure transducers at the nozzle mouth and some distant extremity of the test chamber (where the background pressure could be measured without influence fro the nozzle plume). This would allow the ratio between these pressures to automatically be adjusted in order to achieve a value of 1, which would result in a perfect expansion, i.e. a nozzle plume without the Mach diamond structure or expansion waves predicted by CFD results.

At inlet mass flow rates larger than $15\ SLPM$, the minimum Mach number is always higher than 5.0, allowing these regions to be considered beyond the hypersonic limit. Although lower mass flow rates show larger maximum Mach numbers towards the downstream extent of the Pitot surveys, the high rate of change ($\frac{\partial}{\partial x}(Ma)$) in this region would likely lead to larger variations in

freestream conditions proximate to test pieces, hence requiring additional care when designing and post-processing experiments such that uncertainties in results may be mitigated.

The predicted Knudsen numbers of the flowfield shown in Figure 4.25 can be seen to grow geometrically with increasing distance from the nozzle exit. This is caused by the vast drops in pressure compared to a relatively small decreases in temperature generated by the expansion of the plume (see Eq. (2.3)). As one would expect, raising the flow total temperature can be seen to significantly increase the flow Knudsen number, an effect that is more evident at lower mass flow rates. The relationship between mass flow rate and Knudsen number is the inverse of that between the former and Mach number, i.e. lower values of $\dot{m}_1$ lead to higher ranges of Kn.

A feature of note in Figure 4.25 is the progression of Knudsen number related to a mass flow rate of $10\ SLPM$. These data indicate that the flow Knudsen number reaches 0.1 at just over $100\ mm$ downstream of the nozzle exit, a difference of almost a full order of magnitude when compared to the Knudsen numbers at the nozzle exit for higher mass flow rates. This region represents the upper limit of the slip regime prior to the onset of the transition regime, and is as such a condition of interest in the study of re-entry aerothermodynamics.

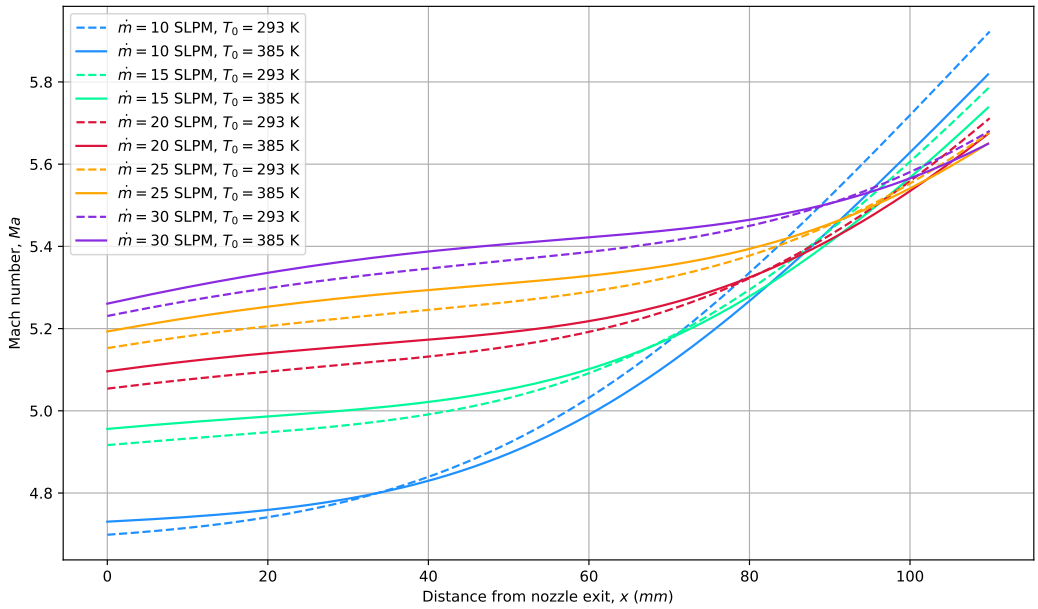


Figure 4.24: Flow Mach number downstream of the nozzle as predicted by CFD simulations using the no-slip nozzle wall.

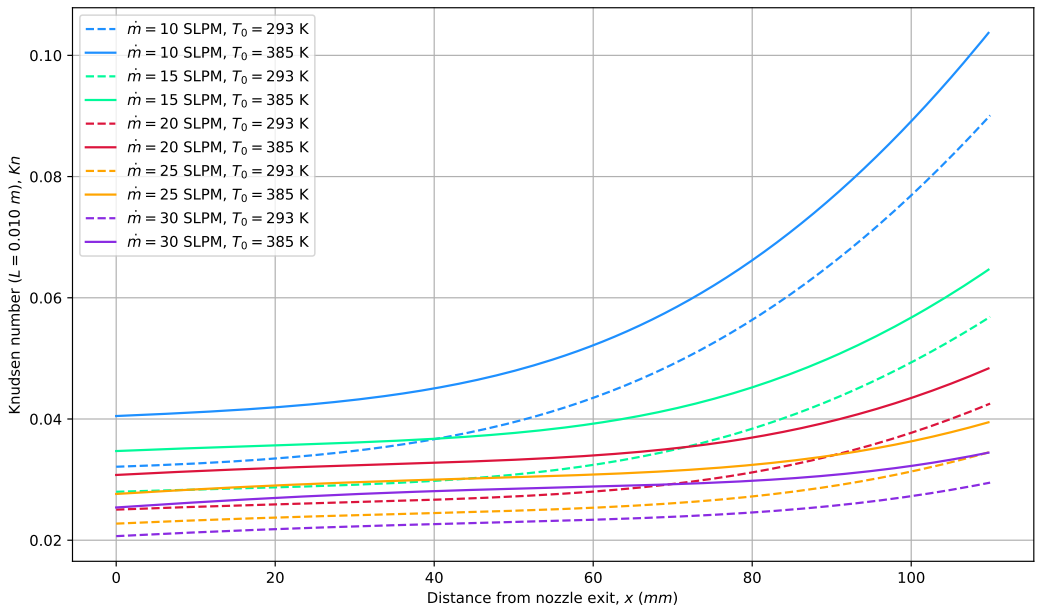


Figure 4.25: Flow Knudsen number downstream of the nozzle as predicted by CFD simulations using the no-slip nozzle wall. The characteristic length used for all calculations was 0.01m.

4.3.6 Surveys of regions of interest in the plume

As this section of the present work is a characterisation of an experimental test facility, one of its primary purposes is to assist in the identification of suitable test conditions in the LDWT plume for upcoming studies of atmospheric re-entry heating. Two such regions and conditions have been identified from the CFD simulations, with their comparison to Pitot data serving as a validation of the models used. Contour plots of the simulated flowfield's Mach number at each condition are presented in Figure 4.26, with labels identifying the relevant locations of interest.

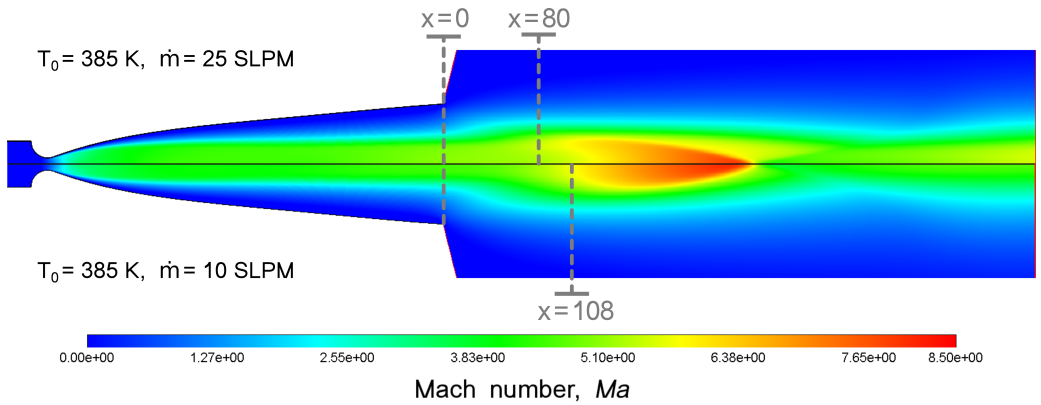


Figure 4.26: Contour plots of Mach number calculated using axisymmetric no-slip CFD simulations at LDWT conditions (see Table 4.2). The spatial coordinates of flow conditions which are of interest to the present study are labelled in terms of their distance from the nozzle exit in millimetres.

The first condition that has been chosen for subsequent studies represents a Knudsen number of 0.022, and is generated by using an inlet mass flow rate of $25\ SLPM$, and a gas total temperature of $385\ K$. The condition exists $80\ mm$ downstream of the nozzle exit in the plume, and has been selected in order to capture heat transfer distribution over satellite debris analogues in the lower section of the slip regime.

A radial Pitot survey of this location in the nozzle plume was performed and compared to the no-slip CFD simulations. These data are presented in Figure 4.27, with data from the case of interest shown in isolation in Figure 4.28. It is clear that the modelling of the core flow Pitot pressure magnitudes agree well with the experimental values, with an error of 1.3% at $z = 0\ mm$, rising to 1.6% at $z = 5\ mm$, and 3.2% at $z = 10\ mm$. The predicted variation of Mach number in this region ($0\ mm \leqslant z \leqslant 10\ mm$) is of the order of $\sim 1.4\%$, with a corresponding Knudsen number variation of $\sim 6.6\%$. Note that these variations were calculated using isentropic flow theory (Eqs. (4.1) and (B.3) in the case of Mach number), and kinetic gas theory (Eq. (2.3)in the case of Knudsen number).

As was observed at the nozzle exit (see Figure 4.22), the boundary layer is poorly modelled at higher mass flow rates (likely as a result of using a fully no-slip wall boundary in the nozzle), resulting in increased inaccuracy further downstream. However, this does not seem to significantly impact the core flow predictions at higher mass flow rates. From the apparent overprediction of boundary layer thickness that the no-slip wall formulation gives, it could be inferred that the area of relative uniformity in the core flow region will in reality be larger than predicted by CFD simulations. The Pitot surveys of the flowfield, however, showed that this is not the case, and that the uniformity of the core flow region is more significantly affected instead by the underexpansion of the plume as it exits the nozzle.

A significant deviation in the predicted trend is visible in the $30\ SLPM$ case at $z \approx 22\ mm$. This apparent "bulge" in the Pitot data has previously been identified by Owen [87] during his Pitot surveys of mass flow rates in the LDWT of the order of $50\ SLPM$, and is caused by an expansion fan emanating from the edges of the nozzle exit, a byproduct of plume underexpansion. Despite the no-slip model's previously mentioned inability to fully predict the nozzle boundary layer, this deviation in the plume's radial trends is still evident. Additionally, it may be observed

that heating of the flow mitigates this effect, as high mass flow rate cases with a stagnation temperature of $385\ K$ show a significantly smoother distribution than those run at $T_0 = 293\ K$.

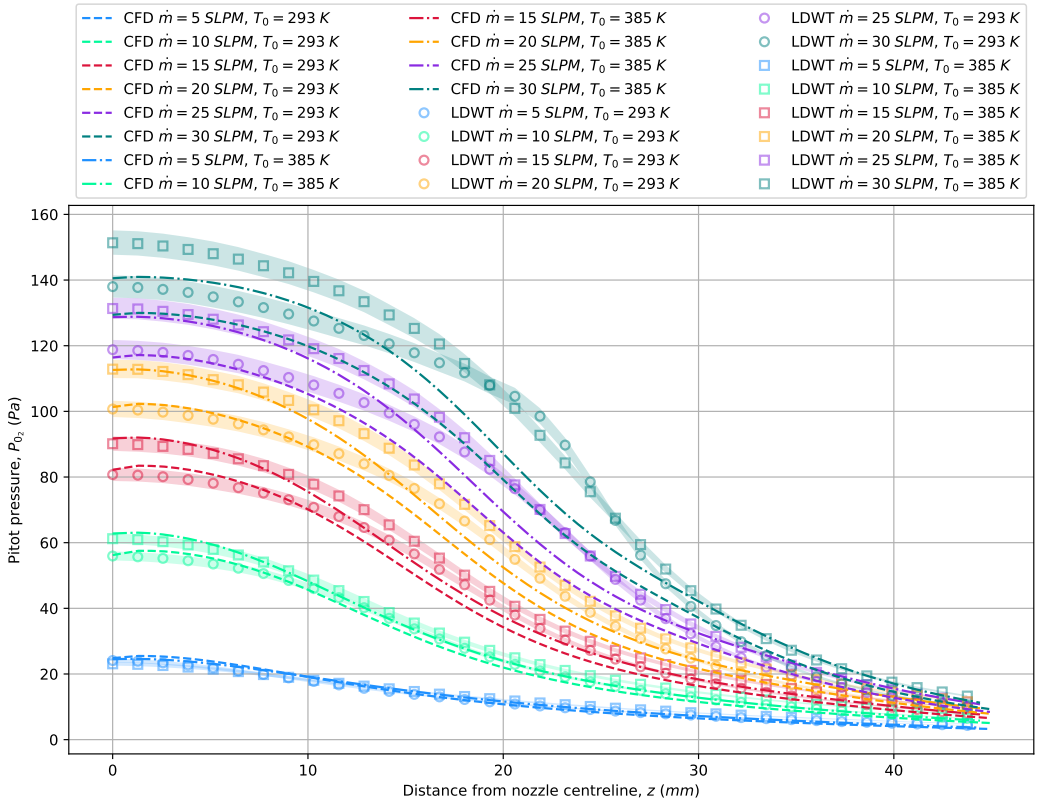


Figure 4.27: Virtual Pitot data from CFD simulations 80mm downstream of the nozzle exit using a no-slip condition at the nozzle wall. Experimental values are overlaid.

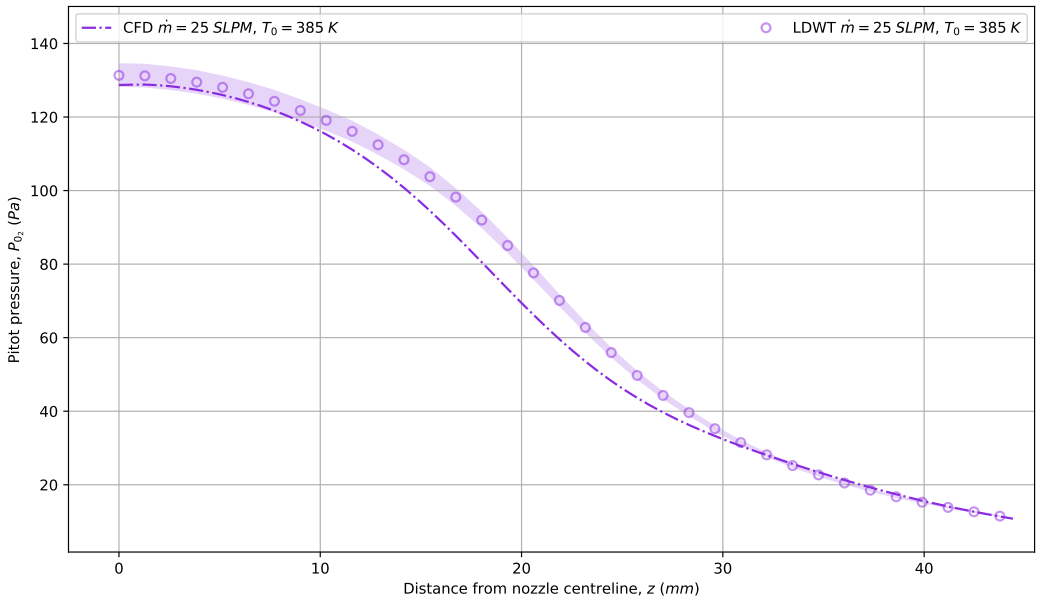


Figure 4.28: Virtual Pitot data from CFD simulations and experimental data 100mm downstream of the nozzle exit using a no-slip condition at the nozzle wall. Only data from the relevant case of interest are shown ($T_0 = 385\,K$, $\dot{m}_1 = 25\,SLPM$, $x = 80\,mm$).

The second region of interest in the LDWT plume is located at $x = 108\ mm$, and is produced using a total temperature of $385\ K$ and an inlet mass flow rate of $10\ SLPM$. Referring to Figure 4.25, it can be observed that this condition represents a predicted Knudsen number of 0.1, a value that, as is mentioned earlier in this section, represents the upper limit of the aerothermodynamic slip regime.

At the time of writing, the closest radial Pitot survey to this location was performed at $x = 100\ mm$. As such, these measurements were used for comparison with data extracted from the corresponding location in the CFD results. A plot of this comparison is presented in Figure 4.29, with data from the relevant case of interest shown in isolation in Figure 4.30.

Predictions of the Pitot pressure at $z = 0\ mm$ agree well with experimental values, with relative errors for moderate to high mass flow rates not exceeding 7%. The relative error at this location for the case of interest ($10\ SLPM$, $T_0 = 385\ K$) is low, at 2.9%. The predicted variation of Mach and Knudsen number in the core flow region ($0\ mm \leqslant z \leqslant 10\ mm$) was calculated using the same equations as in the previous location of interest. These variations are

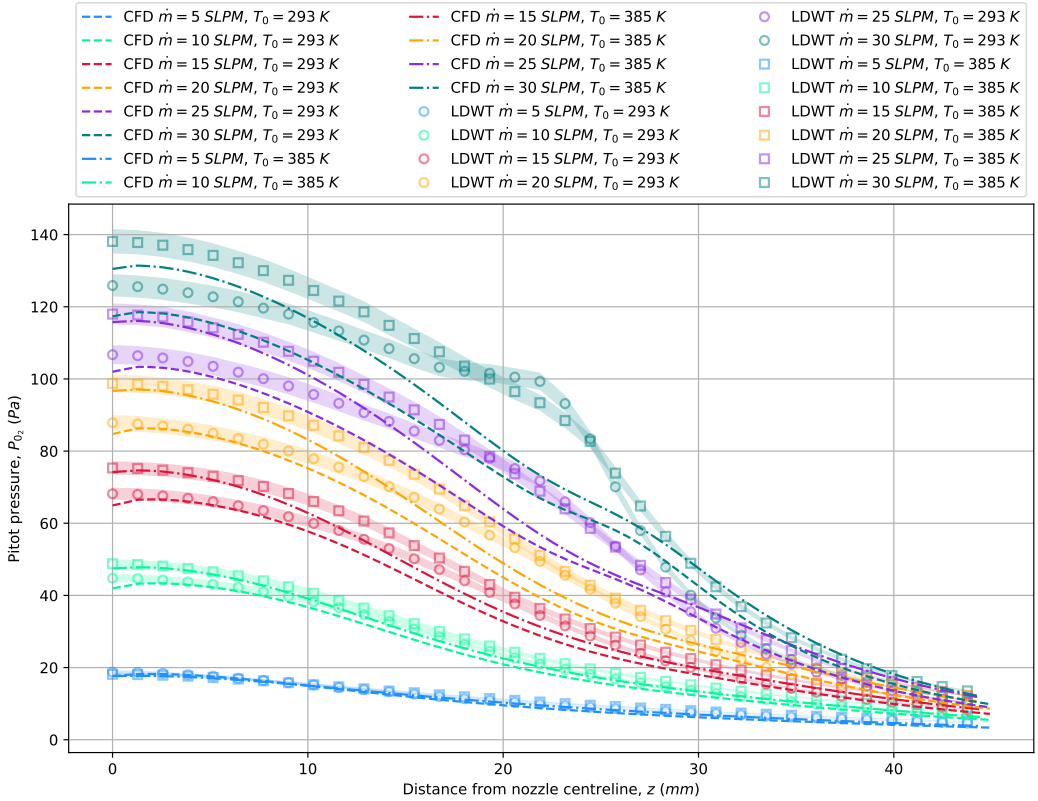


Figure 4.29: Virtual Pitot data from CFD simulations 100mm downstream of the nozzle exit using a no-slip condition at the nozzle wall. Experimental values are overlaid.

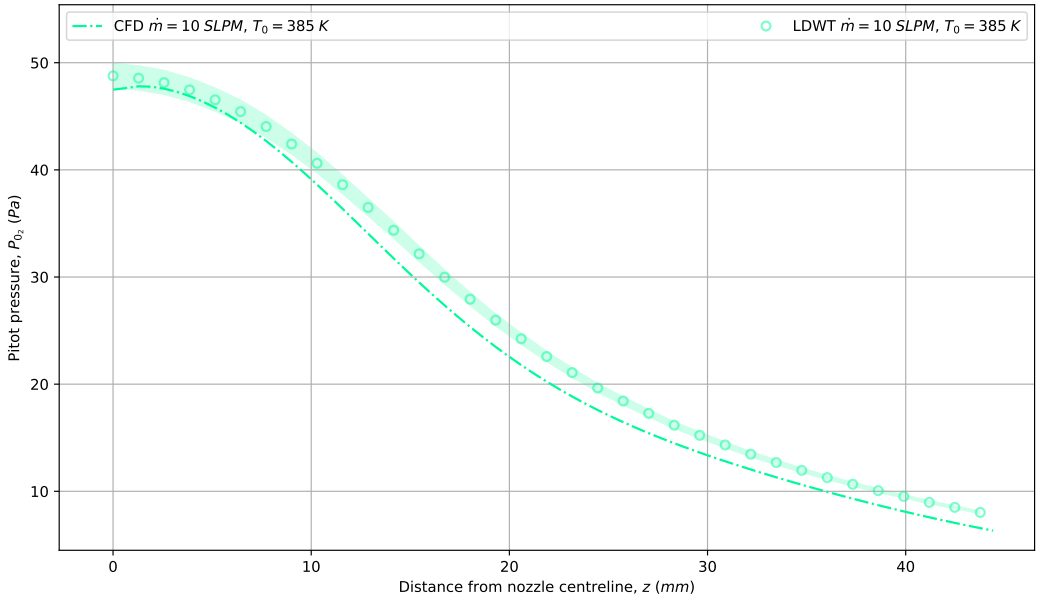


Figure 4.30: Virtual Pitot data from CFD simulations and experimental data 100mm downstream of the nozzle exit using a no-slip condition at the nozzle wall. Only data from the relevant case of interest are shown ($T_0 = 385\,K$, $\dot{m}_1 = 10\,SLPM$, $x = 100\,mm$).

of the order of $\sim 1.3\%$ and $\sim 6.5\%$ for the Mach and Knudsen numbers, respectively.

With this low error manifesting proximate to the location of interest in the flow (and considering the similarly good agreement at $x = 80\ mm$), it is assumed that the CFD simulations will yield accurate predictions of the flow conditions at $x = 108\ mm$. This conclusion is also reinforced by the excellent agreement in boundary layer prediction at low mass flow rates.

4.3.7 Temporal variability of test conditions

In order to characterise periodic variations in test conditions, the environment in both the test and stagnation chambers were monitored over a period of one hour with the LDWT in its high vacuum mode, and without any input mass flow i.e. $\dot{m}_1 = 25\ g/s$. The results of this monitoring are presented in Figure 4.31.

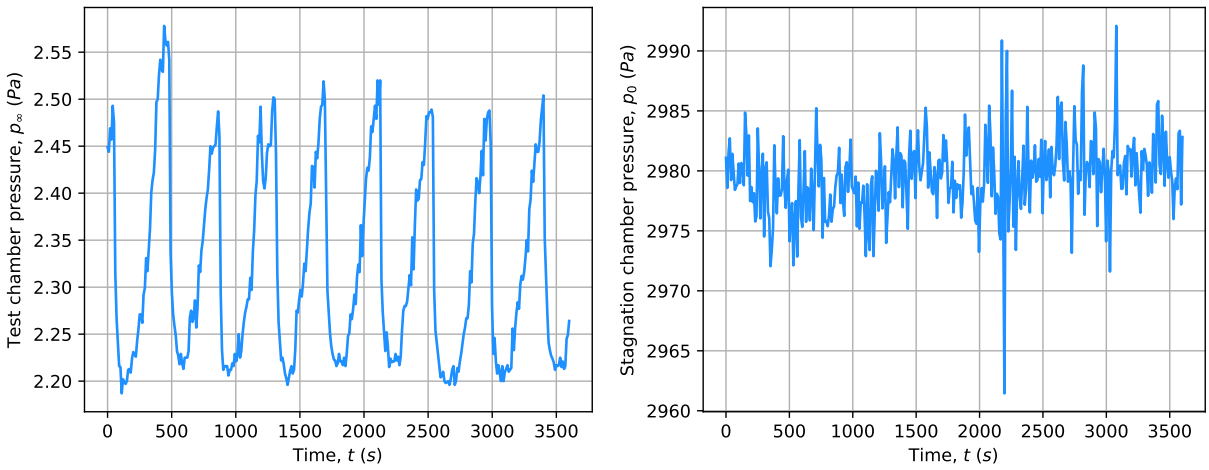
It is immediately apparent that the variations in stagnation and test chamber pressures (p_0 and p_∞, respectively) occur on two separate timescales; the former having a period of ~ 7 minutes, and the latter having a period of ~ 40 seconds. This may be explained by considering the mechanisms which control the flow in and out of the LDWT. In the case of chamber pressure, p_∞, the periodicity of the readings coincides with the automatic operation of the vapour booster heaters, which may only be switched on or off by a PID controller (emitting either $90\ kW$ or

$0\ kW$ during continued operation).

The pressure in the stagnation chamber, p_0, varies on a much smaller timescale, and is set by the mass flow controller that is installed upstream of the stagnation chamber's inline heater. As the powered valve installed in this mass flow controller is adjusted at a frequency of $\sim 10\ Hz$, the input mass flow rate into the stagnation chamber, $\dot{m}_1$, may be maintained with a high degree of accuracy.

The magnitude of these variations are of the order of 10% and 1% for p_∞ and p_0, respectively. Although this variance in p_∞ may seem large, it must be kept in mind that it is the ratio between p_0 and p_∞ that affects the plume Mach number. The effect of this variance was calculated using isentropic flow theory (Eq. (B.3)) based upon the quantities measured in the LDWT. The resulting change in Mach number was of the order of 2%, allowing the plume core to consistently remain above the hypersonic limit.

The Pitot probe used in surveys of the plume was also used to record the variations in test chamber conditions. The probe was situated at the centre of the nozzle exit and data recorded both with and without flow passing through the nozzle (inlet conditions for the former of these were $\dot{m}_1 = 25\ SLPM$, $T_0 = 385\ K$). These data are presented in Figure 4.32.



(a) Test section background pressure

(b) Stagnation chamber pressure

Figure 4.31: Variation of LDWT pressures during a typical test. The inlet mass flow rate and stagnation temperature setpoints used for these measurements were $\dot{m} = 25\ SLPM$, $T_0 = 385\ K$, respectively.

The variation in the pressure transducer zero-offset voltage (shown in Figure 4.32(a)) was of the order 0.025% of the full scale output ($5V$). A low level of variability in this value is

important when taking pressure readings, as it is used as a reference during post-processing, and cannot be measured directly whilst data is being recorded. A downwards trend is evident in these data, and was found to be linked to the variations in p_∞ caused by the aforementioned operation of the vapour booster heaters. This conclusion was further supported by the similarity between the timescales shown in Figures 4.31(a) and 4.32(a).

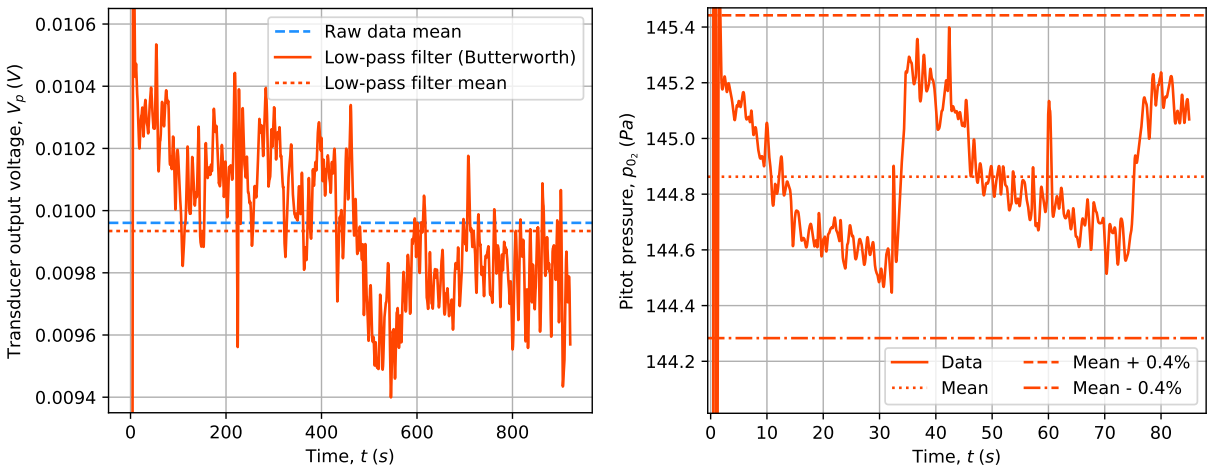
Variations in Pitot pressure in the plume are presented in Figure 4.32(b) for an inlet mass flow rate, $\dot{m}_1$, of 25 SLPM. The data show periodic sharp increases in pressure followed by slow relaxation, behaviour indicative of a control loop like that incorporated in the heaters of the gas conditioning system. The pressure fluctuations occur on the same timescale as the operation of the inline heater PID controller, and represent a variation of $\sim 0.5\%$.

When assessing the flow's Knudsen number, it is the static pressure in the plume itself which is most important for calculations (see Eq. (2.3)), as opposed to the static pressure of the test chamber. As such, the data presented in Figure 4.32(b) is more relevant in this regard than that shown in Figure 4.31(a). Although these data represent the post-shock total pressures at the Pitot, the relevant variance in static pressure, and hence Knudsen number, may be readily calculated.

$$\frac{p}{p_0} = \left(1 + \frac{\gamma - 1}{2} Ma^2\right)^{\frac{-\gamma}{\gamma - 1}} \tag{4.2}$$

$$\frac{T}{T_0} = \left(1 + \frac{\gamma - 1}{2}\, Ma^2\right)^{-1} \tag{4.3}$$

First, the extremes of p_0 and p_{0_2} from Figures 4.31(b) and 4.32(b), respectively, may be used in conjunction with Eq. (4.1) to find the Mach number variance, which for the stated flow conditions was $5.29(88) \leqslant Ma \leqslant 5.33(66)$, a variance of $\sim 0.7\%$. Then, Eqs. (4.2) and (4.3) are used to calculate the freestream static pressure, p, and static temperature, T, at these extremes of Mach number. Finally, Eq. (2.3) is used to find the final variance in Knudsen number. Following this calculation, the variance in Kn was found to be of the order of $\sim 3\%$, demonstrating the stability of the nozzle plume for the duration of a typical experiment.



(a) Transducer output voltage variation with no input flow ($\dot{m}_1 = 0\ g/s$)

(b) Transducer Pitot pressure, p_{02}, variation with an inlet mass flow of 25 SLPM ($\dot{m}_1 \approx 0.5\ g/s$)

Figure 4.32: Variation of LDWT transducer readings at the beginning of a typical test. Readings are presented following the application of a 4th order Butterworth filter.

Figure 4.33 shows the ratio of Pitot pressure, p_{0_2}, and stagnation chamber pressure, p_0. It must be noted that, due to the difference in sampling rates between the two data sets ($0.1\ Hz$ versus $1000\ Hz$), and the inherent noise encountered in the measurement of stagnation chamber pressure (see Figure 4.31(b)), changes in the ratio p_{0_2}/p_0 could not be accurately assessed over timescales shorter than 10 seconds.

The variability of these data is $< 0.8\%$, thereby supporting the previous conclusions that the Mach number of the core flow is temporally stable over the duration of a typical liquid crystal experiment ($\sim$ 60 seconds). Additionally, during the course of the above calculations, it was found that the predicted static pressure in the test plume varied between $3.81(22)\ Pa$ and $3.97(69)\ Pa$. Taking into account a variance in test chamber background pressure between $2.18(7)\ Pa$ and $2.57(8)\ Pa$ (see Figure 4.31(a)), it is clear that the plume is constantly underexpanded, as the ratio between the nozzle exit static and the chamber background pressure is always greater than unity.

Overall, the conditions in the test section can be seen to behave consistently throughout a typical test, with the notable exception of the test section pressure. These variances have been shown to have a negligible effect on the flow Mach and Knudsen numbers, and the structure of the nozzle plume. Hence, it may be concluded that the plume in the LDWT test section can be considered stable for the required test times, and that tests performed during separate runs will

be comparable.

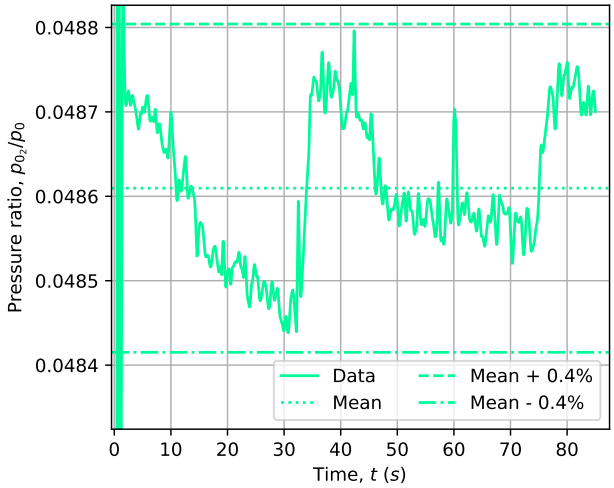


Figure 4.33: Variation of pressure ratio, p_{0_2}/p_0, with an inlet mass flow rate of 25 SLPM ($\dot{m}_1 \approx 0.5\ g/s$). The ratio was calculated using the same data that shown in Figures 4.32(b) and 4.31(b)

Although the above findings support these conclusions, further precautions were taken during experimental runs in order to minimise inaccuracies in the recorded data. Readings from the sensors shown in Figure 4.4 (and Table 4.1) were recorded throughout all subsequent experiments, and utilised during post-processing in conjunction with the semi-analytic model described in Appendix C to better assess the freestream conditions (and the variance thereof) at which a particular test would be performed (see Section 5.1).

4.4 CHAPTER SUMMARY

In this chapter, the layout of the LDWT has been presented, along with a characterisation of the facility's abilities. This was accomplished via radial and axial Pitot surveys of total pressure and total temperature measurements at the centre of the nozzle exit. The former of these were then compared to CFD simulations in order to draw conclusions regarding the accuracy of the modelling approaches employed. A semi-empirical model was then generated from experimental data in conjunction with the CFD results in order to allow fast calculation of the conditions in the LDWT experimental plume both for the present work, and for future users of the facility (see Appendices B and C). This model was expected to yield slightly more accurate predictions of flow variables due to the inclusion of experimental values, and showed a small ($\sim 2.8\%$) increase in Mach number in the core flow region over CFD on its own.

It was found that CFD solutions utilising no-slip boundary conditions at the nozzle wall were well suited to modelling the hypersonic flow in the LDWT test chamber. This was confirmed by calculating the expected Pitot pressures within the plume, and then comparing these data with values gathered during experimental Pitot surveys of the same locations. Errors between simulations and experimental data were found to be low in most cases, with a maximum of 7% in the core flow region, hence validating the CFD results. It was clear, however, that the boundary layer thickness was overpredicted in cases with high mass flow rates, and the effects of this became more evident with increasing distance from the nozzle exit. This did not have a large effect on predictions near the centre of the nozzle plume though, which as previously mentioned, showed consistent agreement with experimental data. Mach and Knudsen numbers in the plume were subsequently calculated from the CFD results, and regions of interest for future heat flux experiments identified.

The temporal variability of the flow in the LDWT was characterised using data from various pressure transducers (including the Pitot probe) recorded over sustained operation of the LDWT. From these data, it was determined that the test chamber and stagnation chamber pressures varied periodically with time, with only the variation in the test chamber being significant (of the order of 15%). The consequent variations that the periodicity of these variables caused in the flow's Mach and Knudsen numbers over time were then calculated, and found to be minimal ($\sim 0.8\%$ and $\sim 3\%$, respectively).

Overall, the capabilities of the LDWT have been simulated effectively, and have been validated using experimental measurements. Errors in the predictions of the spatial variance of the flow were found to be acceptably small, as were the effects of pressure variances over time in the LDWT on the properties of the experimental plume. From these findings, it may be concluded that the LDWT facility, as well as the test conditions presented herein, are suitable for atmospheric entry heating experiments in the hypersonic slip regime.

EXPERIMENTAL METHODS 5

This chapter deals with experimental setup for the present work. These activities comprised the selection of appropriate operating conditions in the LDWT, wherein a summary of the freestream variables are presented and uncertainties based upon LDWT instrumentation are calculated. This is followed by a description of the liquid crystal probes manufactured for hypersonic experiments, and the relevant process required to calibrate them for post-processing.

The data analysis procedure involved in pre and post-processing the experimental results is then described. This section begins with the initial processing of the experimental videos, and is followed by the rationale for the chosen analysis techniques. Specifically, this involves a description of the one dimensional (1D) analysis of the thermal fluxes, why this type of analysis may be applied to only some of the experiments performed during the present work, and the iterative FEA procedure that is incorporated instead for those experiments which cannot be analysed using the 1D method.

Finally, the algorithms used for the pre and post-processing of experimental data are described. These are split into "Experimental video analysis" and "FEA simulation and iterative HTC correction" sections, respectively.

5.1 OPERATING CONDITIONS

The operating conditions that were selected for the present work's experiments are based upon the CFD results described in Sections 4.3.5 and 4.3.6. The properties of these two conditions are listed in Table 5.1, and were computed using linear interpolation of the data extracted from the aforementioned simulations in order to select specific Knudsen numbers. This interpolation was performed by first calculating the mean free path, λ, of the CFD flowfield and, hence, the Knudsen number (by dividing through by various candidate probe diameters). The desired values of Knudsen number were then selected, and the spatial locations in the flow where they occurred recorded for use in subsequent experiments.

Uncertainties in the chosen freestream conditions were calculated using the manufacturer's measurement uncertainty for each sensor used to characterise conditions in the LDWT (see Table

4.1. Other uncertainties, i.e. those which pertain to values that were calculated as opposed to measured, were obtained using isentropic flow equations and gas kinetic theory (Eqs. (2.1), (2.3), (2.5), (2.11), and (B.3), amongst others).

As has been mentioned previously, a difference of one order of magnitude between experimental Knudsen numbers was desired, with the ideal values for the selected operating conditions being $Kn = \{0.01,\ 0.1\}$.

However, owing to the physical limitations of the nozzle plume produced by the LDWT (i.e. a core flow diameter of less than 20 mm), a corresponding limit on model size had to be imposed, thus restricting the effective range of achievable Knudsen numbers (since $Kn = \lambda/L$, and λ is similarly restricted by the flowrate limit of the LDWT's vapour booster pump — see Appendix B). As a result of this, the actual Knudsen number values of the selected conditions are $Kn = \big\{0.02(16),\ 0.10(12)\big\}$.

The Knudsen numbers at these conditions are calculated using two different values of characteristic length, L, which in the case of axially symmetric geometries is taken to be the diameter, $\varnothing$. For condition 1 ($Kn = 0.02$), a value of $\varnothing = 0.015$ m was used, such that the freestream Knudsen number could be brought as close to the desired value of 0.01 without the probe's outermost sections being outside of the LDWT's core flow during experiments. For condition 2, a value of $\varnothing = 0.01$ m was used, which allowed the desired Knudsen number of 0.1 to be selected with little error whilst once again allowing the probe to remain in the near-homogeneous region at the core of the nozzle plume in the LDWT.

Figure 5.1 shows the simulated flowfield generated by the LDWT nozzle at a stagnation temperature, T_0, of $385\,K$, and inlet mass flow rates, $\dot{m}_1$, of $25\,SLPM$ (upper contour plot) and $10\,SLPM$ (lower contour plot). The outlines of two hemisphere-capped cylinders scaled to the diameters derived above are overlaid in black at their respective positions of interest. As may be observed from the figure, both locations of interest identified above lie within the expansion region immediately downstream of the nozzle exit. The selected probe diameters are also far from the boundaries of the shock diamond structure which develops at the radial extremities of the experimental plume, thereby providing an acceptable degree of core flow uniformity (see Section 4.3.6).

Parameter	Symbol	Unit	Nominal value		Nominal + pos. uncertainty		Nominal + neg. uncertainty	
			Condition 1	Condition 2	Condition 1	Condition 2	Condition 1	Condition 2
Distance from nozzle exit	x	m	0.08	0.108	–	–	–	–
Knudsen number	Kn	–	0.0216	0.1012	0.0244	0.1458	0.0193	0.0759
Mach number	Ma	–	5.394	5.787	6.2	8.6	4.7	4.2
Mean free path	λ	m	3.242×10^{-4}	1.012×10^{-3}	3.654×10^{-4}	1.458×10^{-3}	1.458×10^{-3}	1.012×10^{-3}
Probe diameter	$\varnothing$	m	0.015	0.01	–	–	–	–
Freestream velocity	U_∞	$m \cdot s^{-1}$	813.53	822.75	920.63	1196.87	722.23	610.74
Speed of sound	a_∞	$m \cdot s^{-1}$	150.82	142.17	153.8	145.29	147.94	139.08
Mass flow rate (inlet)	$\dot{m}_1$	$SLPM$	25	10	25.1	10.1	24.9	9.9
Total temperature	T_0	K	385	385	387.2	387.2	382.8	382.8
Total pressure	p_0	Pa	2826.16	1231.21	2892.66	1297.71	2759.66	1164.71
Static temperature	T_∞	K	56.67	50.34	58.87	52.54	54.47	48.14
Static pressure	p_∞	Pa	3.395	0.966	3.661	1.232	3.129	0.7
Mass density	ρ_∞	$kg \cdot m^{-3}$	2.087×10^{-4}	6.687×10^{-5}	2.342×10^{-4}	8.917×10^{-5}	1.852×10^{-4}	4.642×10^{-5}
Number density	n	m^{-3}	4.339×10^{21}	1.390×10^{21}	4.868×10^{21}	1.854×10^{21}	3.850×10^{21}	9.650×10^{20}

Table 5.1: Freestream conditions and associated uncertainties used for experimental studies and accompanying DSMC simulations (derived from CFD simulations of the plume in the LDWT).

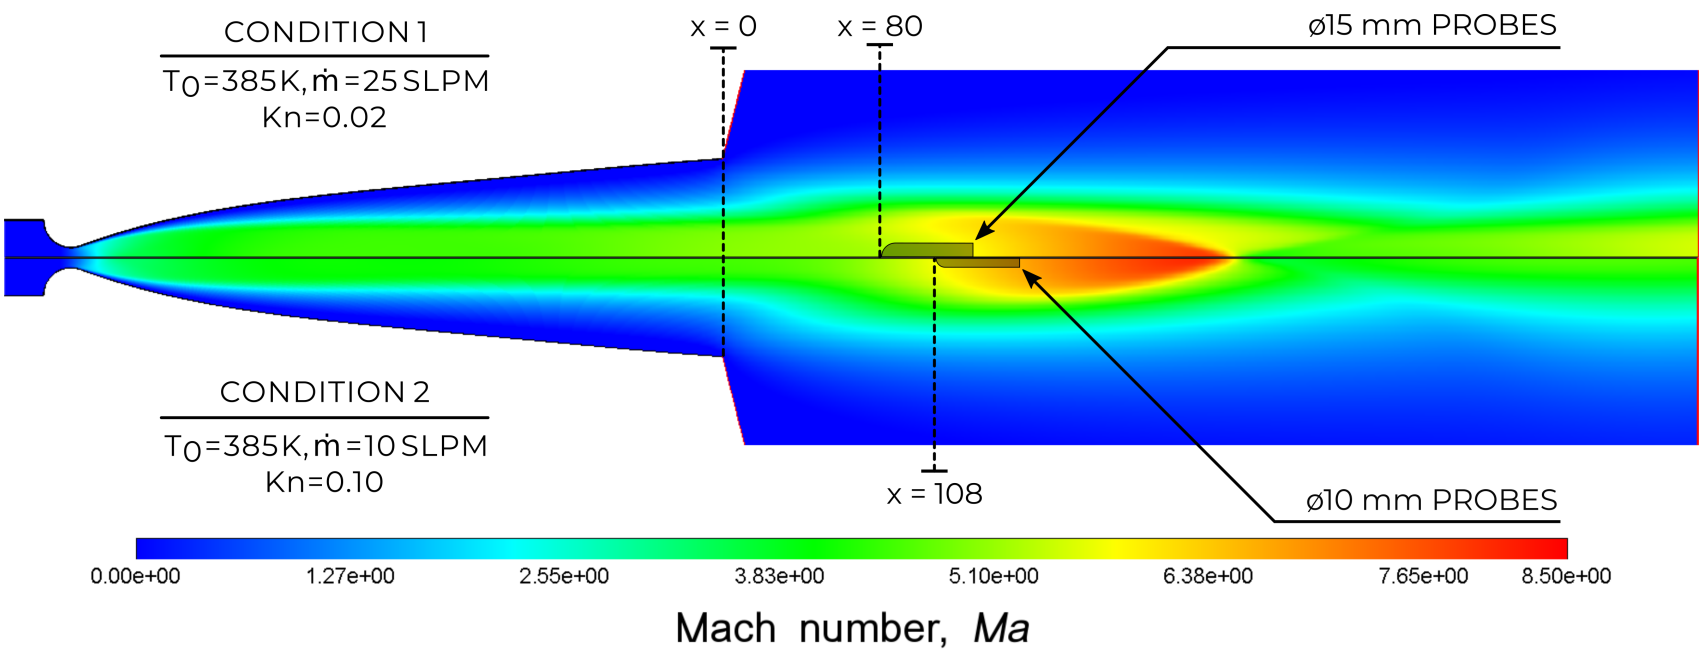


Figure 5.1: Contour plots of Mach number calculated using axisymmetric no-slip CFD simulations at LDWT conditions (see Table 4.2). The spatial coordinates of flow conditions which are of interest to the present study are labelled in terms of their distance from the nozzle exit in millimetres. Silhouettes of both selected hemispherical probes are included at the appropriate regions of interest in the flow (cylindrical probe silhouettes are not shown, but would have the same outer diameters and positions as the hemispherical probes).

5.2 LIQUID CRYSTAL PROBES

Several liquid crystal probes were developed for the present work, with four final versions eventually being utilised for experimental studies. Diagrams of these probes are presented in Figure 5.2.

Of these, two feature flat circular tips, while the other two are tipped with hemispheres. Two variants of each probe shape were manufactured: two with an outer diameter of 15 mm, and another two with an outer diameter of 10 mm. The rear sections of the probes all have an outer diameter of 12 mm such that they could be mounted to the traversing mechanisms in the LDWT.

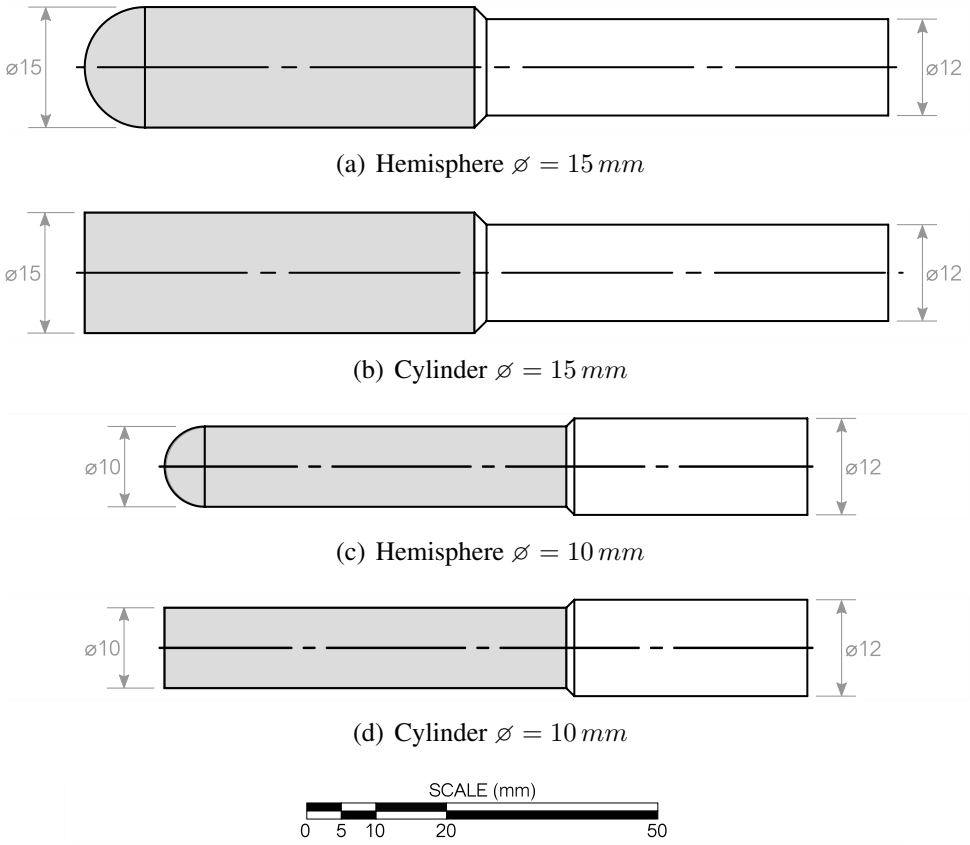


Figure 5.2: CAD diagrams of liquid crystal probe bodies. Shaded regions indicate surfaces which are treated with thermochromic liquid crystal slurries during manufacture.

The probe bodies are manufactured by turning rods of Perspex (polymethyl methacrylate, abbreviated PMMA) on a lathe, thereby ensuring that the probes are axially symmetric. This material was selected for its low cost, wide availability, and well understood thermal properties (which, as Ryley [4] notes, are known to remain stable at the temperatures expected during the experiments described herein).

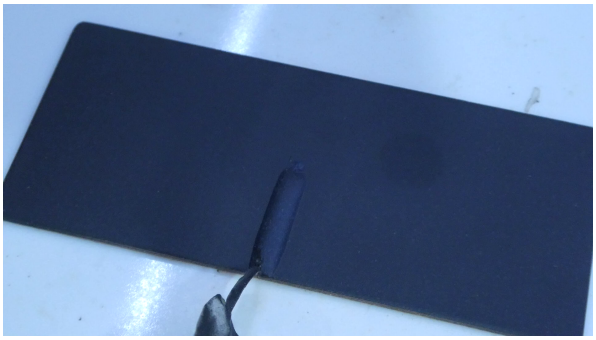
Once the probe bodies have been manufactured, they are sprayed with a thin undercoat of black paint in order to improve image quality during experiments, which is followed by the liquid crystal (LC) slurry mixture. The LC mixture used for the present work comprises three separate slurries manufactured by LCR Hallcrest Ltd., with transition temperatures of 25^oC, 30^oC, and 35^oC (see Section 2.5.3). Owing to the high cost of the liquid crystal slurries, they are only applied to those sections of the probes which are of interest (as shown in Figure 5.2).

5.2.1 Calibration procedure

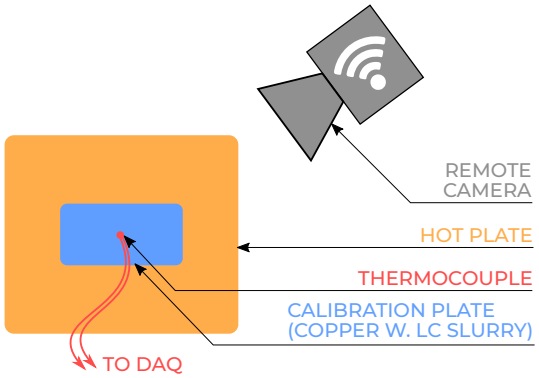
The liquid crystal calibration procedure involves characterising the visible response of an LC mixture with respect to the surface temperature of its substrate. These data are essential during the post-processing of experimental results as they allow temperatures to be derived from the visible colour play produced by an experiment (or indeed, vice versa, whereby variations of temperature over time may be used to generate a simulated colour play). Note that a direct correlation between temperature and pixel intensity is only required for the latter of these cases – when analysing colour play data (whether from experiments or simulations), it is only the location of the transition peaks in time that is of concern, and not which specific value of temperature corresponds to a given pixel intensity. This is advantageous since the pixel intensities do not increase monotonically with temperature.

Although the transition temperature of a given LC slurry is stated by the manufacturer, mixtures of these solutions often have slightly different proportions of their constituent slurries. As such, a calibration procedure is required in order to define the expected visual response of an LC mixture to a variation in temperature.

To obtain these data, an LC solution is mixed thoroughly and applied to a thermally conductive substrate (a copper plate was utilised during the course of the present work) using a pneumatic paint sprayer in order to ensure an even coat. A thermocouple is then applied to this calibration plate, and steady, even heating is performed while both the colour play and temperature of the calibration plate's surface are simultaneously recorded. In the case of the present work, a PID controlled hot plate was used as the heat source for calibration. This setup is pictured in Figure 5.3.



(a) Photograph of copper plate treated with LC slurry, and with a thermocouple affixed to the surface



(b) Diagram of calibration setup

Figure 5.3: Liquid crystal calibration setup

Due to the inherent variability of the slurry proportions in LC mixtures, it is vital to the calibration process that a mixture is applied to all calibration pieces and experimental probes simultaneously. This ensures that the same mixture (and, hence, the same proportions of the constituent slurries) is used both for experiments and calibration procedures.

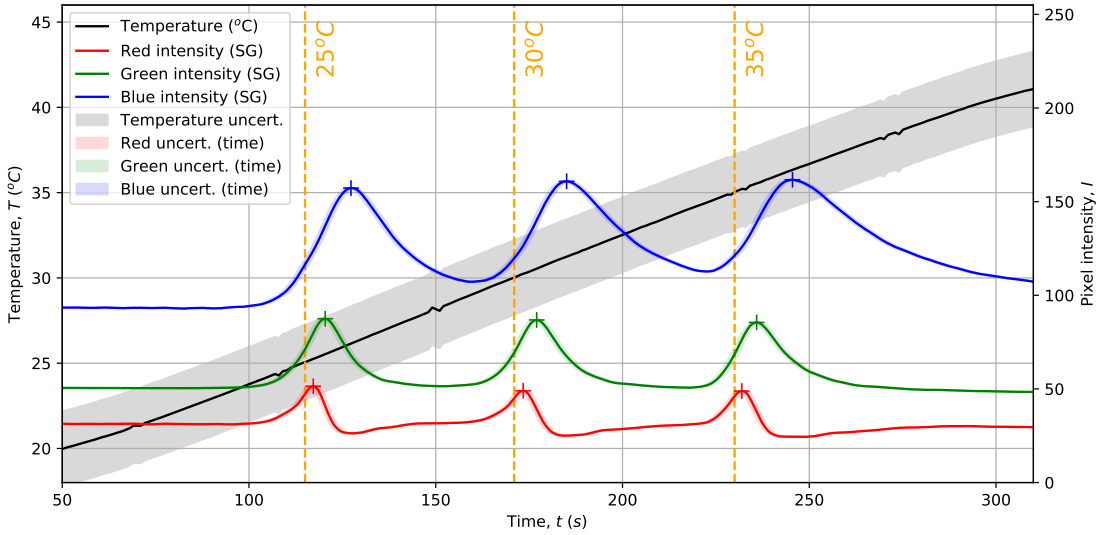


Figure 5.4: Liquid crystal calibration data measured during even heating of a copper plate. Video colour intensities and plate temperatures are plotted, and transition temperature crossing times (at 25^oC, 30^oC, and 35^oC) are shown as dotted lines. The term"SG" refers to those data which have undergone Savitzky-Golay filtering.

The results of such a calibration are shown in Figure 5.4. The presented data were sampled from a 40×40 pixel window in the calibration video proximate to the thermocouple (which, as is mentioned above, was affixed to the top surface of the calibration plate). The intensity of each colour channel of this video and the temperature recorded by the thermocouple are plotted

against time.

A lookup table is then generated from each colour channel with respect to the thermocouple temperature, allowing interpolation between a given temperature and the intensities of each colour channel. For the present work, the green channel is utilised exclusively. It was observed during early tests that the red channel was too faint to effectively analyse, and that rapid heating near probe edges caused transitions in the blue channel to blend into one another, further complicating post-processing of experimental videos.

5.3 DATA ANALYSIS

The results of transient thermal experiments generally require post-processing in order to obtain useful data, which often takes the form of heat transfer coefficients, h_{conv} (also abbreviated as HTC). In the case of liquid crystal experiments, this procedure involves calculating h_{conv} for each point on the surface of a test piece. For the present work, this calculation is performed using one of two different methods (depending on the probe geometry used): using either a comparative, one dimensional heat transfer analysis, or an iterative, three dimensional finite element solution. These procedures are elaborated upon in Section 2.5.3, while the technical aspects of their implementation and use are described below.

Additionally, the experimental video itself must also undergo image stabilisation and filtering in order for subsequent analyses to yield accurate results. This is due to minor vibrations induced by the LDWT pumping system and flow shield motor, which propagate to the camera used to film the experiments. As such, pixel locations do not remain constant over the course of the resulting videos. The procedures involved in this process are also described below:

5.3.1 Video processing

The procedure for readying an experimental video for analysis is a five stage process comprising stabilisation, trimming, frame rotation, colour channel separation, and cropping. Filtering of the video data to reduce noise is explicitly not performed at this stage in order to reduce computational expense (and allow greater flexibility later on in the analyses). Perspective distortion correction may also be employed at this stage of processing, but this was found to be

largely unnecessary during the course of the present work and, as such, was not performed. Note that while selecting the orientation and viewing angle of the camera for each experiment, areas of the probes that were expected to experience significant heating were prioritised over those that were not.

The first step — stabilisation — eliminates camera and probe motion relative to a video's first frame, hence ensuring spatial consistency between pixels across the duration of the video. The tool used to perform this step is called Blender [134], and is an open source program designed for creative professionals and hobbyists. As such, it contains a great many features for editing videos, including an interactive stabilisation dialogue.

This tool operates by having the user select a series of tracking points in a video around each of which a patch of pixels is selected and stored. Subsequent frames in the video are then searched for these patterns (within specified bounds in order to reduce processing time), and the position at which the best match is found stored as the tracked feature's new position. Once motion within the video has been characterised using this method, subsequent frames are translated such that the trackers are aligned. In the case of multiple trackers being used, this translation is calculated using the mean of all tracker positions for any given frame.

Trimming is then performed whereby frames prior to the start of heating are removed from the video. As such, the first frame of a video may be taken to be the start of an experiment, thereby simplifying subsequent parts of the post-processing procedure. Frames toward the end of the videos, i.e. those that occur following the meaningful portions of the colour play on the surface of a probe, are also removed.

The next part of the video processing procedure is frame rotation, which is performed in order to align a probe's longitudinal axis with the horizontal axis of the video frame. This greatly simplifies visual inspection of the video such that the positions of the centre and edges of the cylinder probes' circular faces may be identified. It also enables the determination of Euler rotation angles for the probe, which are essential for mapping of the video data onto a computational mesh during subsequent stages.

The final two steps — colour channel separation and cropping — are performed to reduce computational expense during processing. Cropping of the videos in particular is extremely

useful in this regard, as it vastly reduces the amount of RAM required during subsequent analyses.

5.3.2 One dimensional (1D) heat transfer analysis

Analysing transient heat transfer in one dimension is an extremely efficient method for calculating h_{conv}, as it relies solely upon Eqs. (2.44) and (2.43). However, as is mentioned in Section 2.5.4, this method is only applicable in cases where the thermal penetration depth, x_{tp}, remains negligible for a significant portion of an experiment.

Referring to Eq. (2.45), it may be observed that the time at which a thermal pulse begins to interact with itself near the centre of some medium with thermal diffusivity, α, is t_{tp}. If the dimensions of an experimental test piece are approximately comparable in all directions, Eq. (2.45) may be re-arranged for t_{tp}, with x_{tp} taking the smallest of these dimensions. This then allows the calculation of the time at which the semi-infinite substrate assumption will be invalidated.

In the case of the hemispherical probes used for the present work, x_{tp} was set to the radius of each test piece, i.e. $\varnothing/2$. The results of these calculations were values of $25.31\,s$ and $11.25\,s$ for the $\varnothing = 15\,mm$ and $\varnothing = 10\,mm$ probes, respectively. As such, these times represent the point in an experiment after which the semi-infinite substrate assumption becomes invalid, and hence, the point up to which one-dimension thermal analyses may be performed.

In the case of the cylindrical probes, the value of t_{tp} approaches zero at the shoulder of the cylinder. This is due to the small thermal mass of these sections of the probes, coupled with the fact that heating occurs on both the side and front surfaces of the cylinders. It may then be concluded that these probes experience three dimensional heat transfer immediately upon the start of an experiment, and as such may not be analysed using the semi-infinite substrate assumption.

The tool used for one dimensional analysis of the hemisphere cases is called THTAC (Transient Heat Transfer Analysis Code) [135], and is an in-house code developed at OTI using Matlab. THTAC operates by calculating a series of temperature rises against time using a range of user-specified possible values of h_{conv} in conjunction with Eqs. (2.44) and (2.43). The liquid crystal calibration lookup table (described in Section 5.2.1) is then used to generate a simulated

intensity profile for each specified value of h_{conv}, and these are compared to the actual values from the experimental video. Once comparisons of experimental and simulated intensities have been made, the value of h_{conv} which produced the lowest error between the two series is selected as that pixel's value.

5.3.3 Iterative FEA solution

A new experimental post-processor was developed for the present work, and is based upon the method previously developed by Ryley [4]. This code (which will simply be referred to as the "3D HTC solver" from here onwards, or "3D-HTCS" for short) is written in MATLAB, and utilises the LiveLink interface for COMSOL Multiphysics in order to run FEA simulations.

As is explained in Section 2.5.4, the thermal environment within the cylindrical probes was expected to be complex; so much so that the assumptions on which 1D analyses are based would be invalid. The 3D HTC solver was developed in order to address this limitation, allowing a video of such an experiment to be analysed such that the surface heat transfer coefficients could be correctly determined. The method utilised by the code is described below:

Experimental video analysis

The following preprocessing operations are performed at the beginning of each run of the 3D HTC solver. As neither the input video data nor the array indexing order used by COMSOL change between iterations, these steps are only performed once.

1. **Video import** A single colour channel of an experimental video is read into memory as individual frames.

2. **Normalisation** The colour intensity of each pixel is normalised to the interval $[0, 1]$, allowing direct comparison between different video data and simulated pixel intensities. Since it is *events* in the LC colour play data that are analysed, the absolute values of pixel intensity are not required (and can in fact hinder certain algorithms such as those used for peak finding).

3. **Mesh reorientation** A precomputed computational surface mesh (generated using

COMSOL) is loaded into memory and transformed using Euler rotation matrices to match the orientation of the probe as observed in the experimental video.

4. **Mesh shielding check** Following the reorientation of the FEA surface mesh, it is necessary to reject nodes which are not coincident with the probe's surface as viewed through the camera in the LDWT. The visual panel identification method described in Section 3.3.4 is used to detect those surface mesh nodes which would be obscured, allowing them to be ignored during the course of the HTC correction process. Any nodes that are detected as being "shielded" in this fashion, have their HTC values set to $0\,W{\cdot}m^{-1}{\cdot}K^{-1}$ in subsequent simulations. Note that the setup of the cylindrical probes and traverse, i.e. in which only the pitch, α, of the probe is altered between experiments, were designed to be symmetrical about the vertical mid plane of the LDWT nozzle, with the camera used to capture the colour play seeing one full half of the cylinder. In this fashion, the colour play could be mapped onto the FEA mesh, and a symmetric boundary condition used to assume the same thermal behaviour on the opposite side.

5. **Video interpolation and resampling** In order to utilise the LC colour play data in simulations, nodes in the FEA mesh are assigned the intensity history of the closest (spatially) video pixel. Linear interpolation is then used to calculate the video colour intensity against time at each unshielded node in the FEA mesh. These interpolated video intensity histories are then resampled using the same array of times utilised in the FEA solutions to generate output data. Hence, direct numerical comparisons between simulated and experimental data may be made. It should be noted that these data are operated on instead of the original video in all subsequent processing steps, thereby greatly reducing the time required to run the analyses (this is because there are ~ 4000 unshielded nodes in the FEA surface meshes, compared to $\sim 1 \times 10^5$ pixels per cropped video). An example of this interpolation is presented in Figure 5.5 for the $\varnothing = 15\,mm$ probe when oriented at $\alpha = 45^o$ in the LDWT's hypersonic plume.

6. **High activity detection** Regions of high activity in the experimental video, i.e. nodes where all colour play transitions occur in a short time, are determined and stored, hence

allowing additional analyses to be performed automatically in the event that peak finding is found to be insufficient. This categorisation is performed by calculating the cumulative sum of all activity for a given node. The 90th percentile of this cumulative activity is then calculated, and those nodes whose percentiles occur before a user-defined time threshold are classed as featuring high activity. An example of high activity detection is presented in Figure 5.6, and was performed on the $\varnothing = 15\,mm$ probe when oriented at $\alpha = 45^o$ in the LDWT.

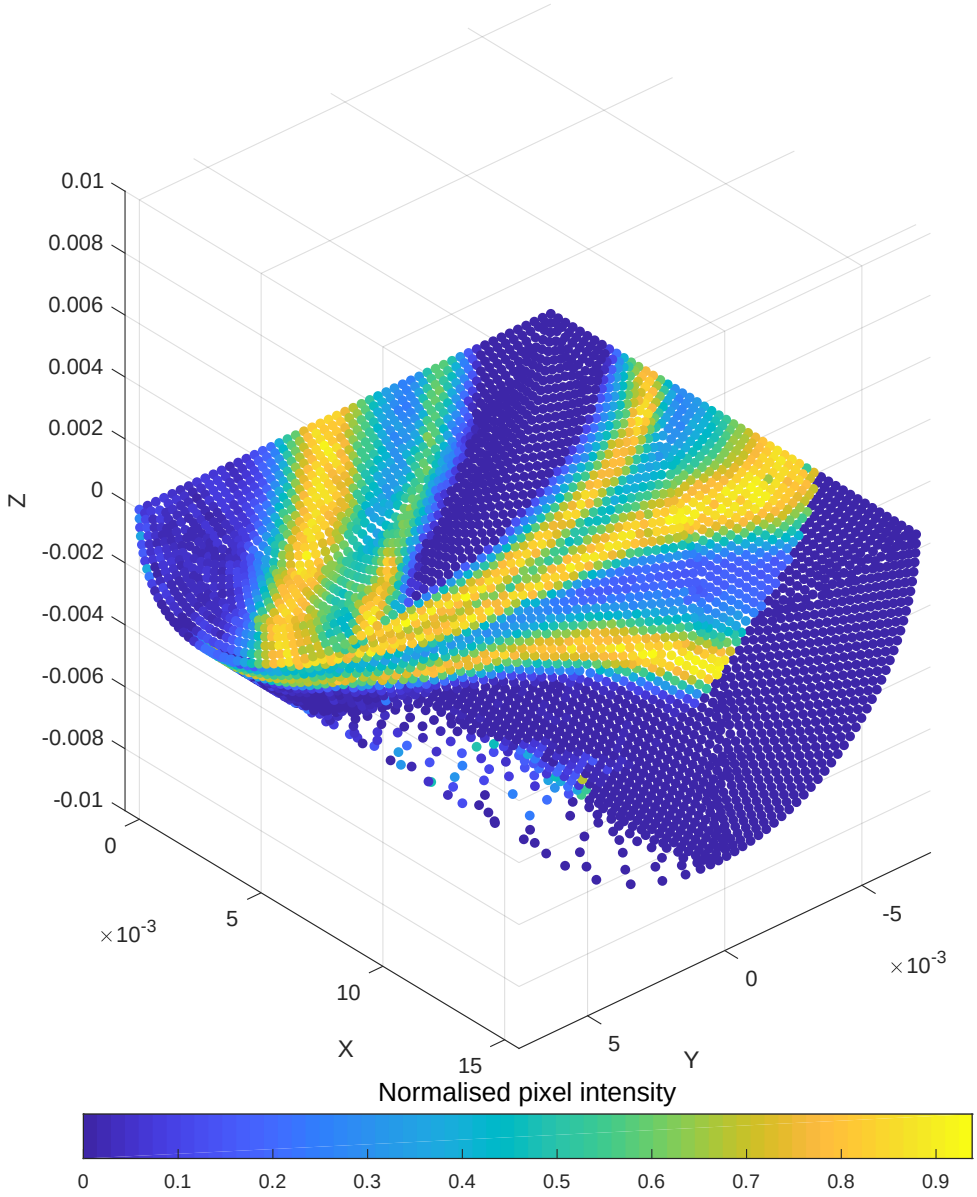


Figure 5.5: Normalised node activities following video signal interpolation and mapping onto the FEA mesh. Data from the $\varnothing = 15\,mm$, $\alpha = 45^o$ case are shown at $t = 23.3\,s$.

7. **Low activity detection** Regions of low activity represent areas where there is either

excessive noise, or no colour play transitions occur. These regions are identified by calculating the standard deviation of the experimental video signal and comparing it to a user-defined limit. Any nodes whose standard deviations fall below this limit are classed as having low activity, and as such have their HTC values set to $0\,W{\cdot}m^{-1}{\cdot}K^{-1}$. An example of low activity detection is presented in Figure 5.6, and was performed on the $\varnothing = 15\,mm$ probe when oriented at $\alpha = 45^{o}$ in the LDWT.

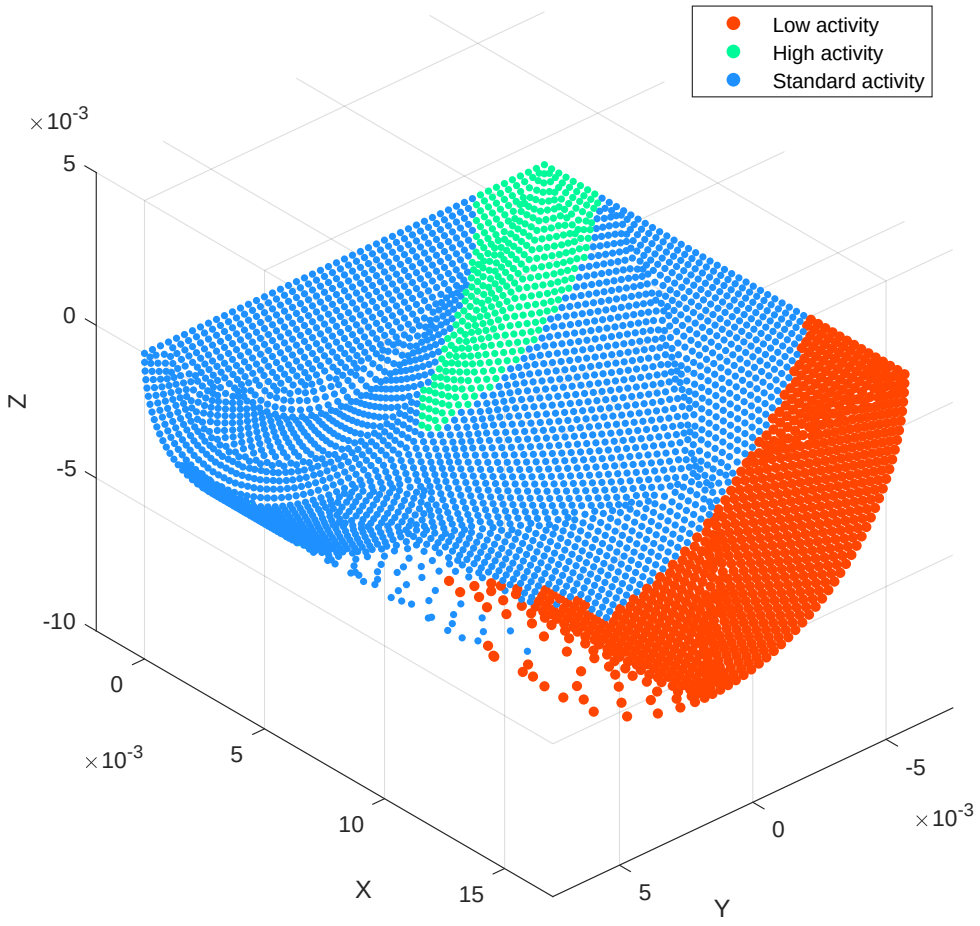


Figure 5.6: Node activity thresholds as calculated by the 3D HTC solver for the $\varnothing = 15\,mm$ case at $\alpha = 45^{o}$

8. **Savitzky-Golay (SG) filtering** This method of signal filtering was developed by Savitzky and Golay in 1964 [136], and operates by fitting a series of polynomials to successive subsets of points in a dataset. The 3D HTC solver applies an SG filter of low order and window length multiple times to the interpolated video signal in order to remove noise whilst avoiding shifting of the signal in time. Less aggressive filtering is applied to those nodes which have previously been identified as featuring high activity. This is due to

their features often being significantly less well defined than other nodes whose colour play transitions do not occur over such short timescales. Figure 5.7 shows the effects of Savitzky-Golay filtering on a typical video activity signal.

9. **Peak detection** As is demonstrated by Figure 5.4, peaks in the colour intensity signal of a pixel correspond to temperature transitions of the liquid crystal coating on the experimental probes. As such, identifying the times at which these peaks occur is a highly accurate method for comparing the colour play between experimental videos and those generated using FEA simulation data. Peak finding in the 3D HTC solver is performed using the MATLAB `findpeaks` command, where the settings for the peak finding algorithm are manually tuned for each separate experimental video. The results of peak finding on a typical video activity signal are presented in Figure 5.7.

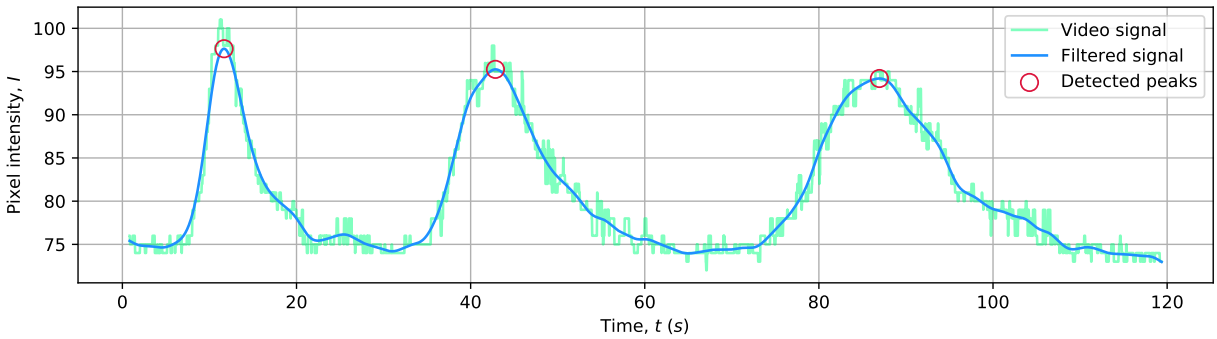


Figure 5.7: Example of peak finding applied to the filtered signal of a single pixel of an experimental video

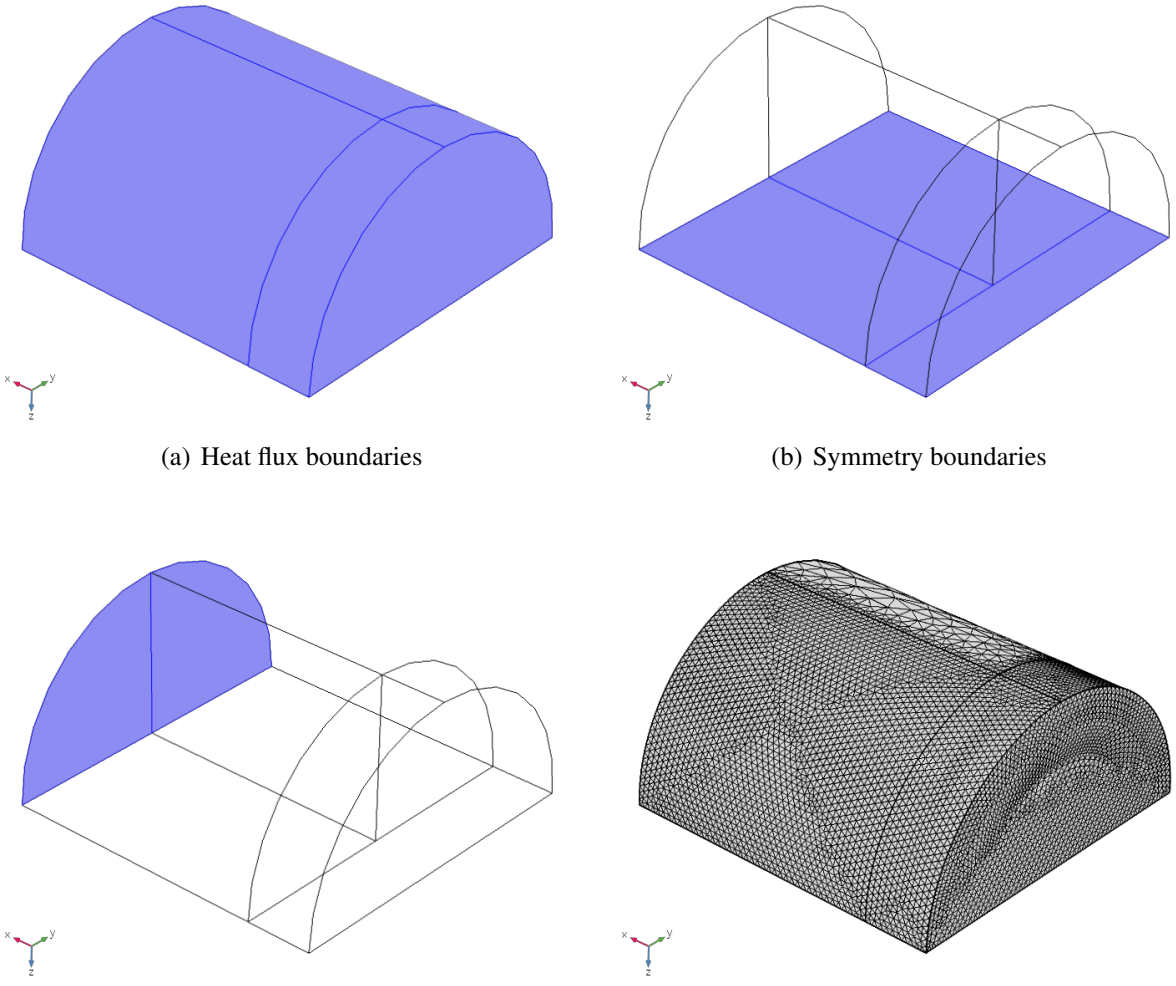
FEA simulations and iterative HTC correction

The following operations are performed during every iteration of the 3D HTC solver. The underlying assumptions of this algorithm are twofold. The first is that the heating experienced by a probe is constant during an experiment, i.e. there are no times at which the temperature of the probe surface is not increasing. Secondly, it is implicitly assumed that any sections of the probe not imaged by the camera do not contribute significantly to the thermal behaviour of the probe.

1. **Run FEA simulation** The MATLAB-COMSOL LiveLink interface is called in order to start an FEA simulation using the current estimate of HTC for each node. The initial simulation in an analysis is run using an HTC distribution calculated using the DSMC

simulations described in Section 3.4.4. The adiabatic wall temperature for each node is also determined using DSMC results; specifically, a series of adiabatic simulations run using the same boundary conditions as the aforementioned cases.

Each COMSOL simulation utilises the same geometry, mesh, boundary types, and material properties. The former three of these are shown in Figure 5.8, while the latter are the same as those used by Ryley for Perspex ($\rho = 1180\,kg{\cdot}m^{-3}$, $C_P = 1450\,J{\cdot}kg^{-1}{\cdot}K^{-1}$, $k = 0.2\,W{\cdot}m^{-1}{\cdot}K^{-1}$).



(a) Heat flux boundaries

(b) Symmetry boundaries

(c) Thermal insulation boundaries

(d) Surface mesh

Figure 5.8: Boundary conditions and surface mesh used for FEA simulations during iterative HTC solutions.

The mesh shown in Figure 5.8(d) was generated for the $\varnothing = 15\,mm$ case and contains $90,755$ nodes, ~ 4500 of which lie on the visible portions of the outer boundary. The mesh

generated for the $\varnothing = 10\,mm$ case is similar, containing $91,225$ nodes, with approximately the same number lying on the visible portions of the outer boundary.

The symmetry boundary condition is incorporated in order to both reduce the computational expense of the FEA simulations, and satisfy the assumption of symmetric thermal behaviour on both sides of the probe (hence allowing the use of only a single camera - see the "Experimental video analysis" preprocessing steps earlier in this chapter). The thermal insulation boundary was selected in order to simulate the far end of the probe which was situated in the wake region of the flow during all experiments. As such, it was assumed that negligible heat transfer would occur, which could be ignored in the FEA solutions without impacting the physical accuracy of the simulations.

Each simulation is run up to the end time of the current experimental video, and reports its results in 0.025 second intervals for the initial portion of an anaylsis. This is in order to provide sufficient resolution for the comparison of high activity regions in the video to the FEA results. Following the end of the high activity period in the experimental video (which is typically 5 to 15 seconds, depending on the case), results are reported in 0.2 second intervals.

Finally, once an FEA simulation has been completed, the temperature history for each surface node is written to disk, and the main MATLAB script is allowed to continue with subsequent steps.

2. **Generate simulated experimental response** Temperature histories from the most recent FEA simulation are read into memory and are used to generate simulated pixel intensities for each node. These calculations are performed by using linear interpolation to query the liquid crystal calibration lookup table described in Section 5.2.

3. **Normalisation** The simulated pixel intensities are normalised to the interval $[0, 1]$ such that direct comparisons may be made by the code between the simulated intensities and the experimental video data.

4. **Peak finding / upper percentile calculation** The same peak finding algorithm used to analyse the experimental video is utilised in identifying peaks in the simulated intensities

for each node. This procedure is greatly simplified by the fact that, since all simulated intensities are generated using the same lookup table, the simulated intensity profiles are all of similar shape.

For those nodes that reside in high activity regions, the time associated with the occurrence of the 90th percentile of their intensity activity is calculated. This is performed using the same strategy as the initial video analysis, i.e. calculating the cumulative sum of all activity over time (per node), and finding the time which corresponds to 90% fo this activity having occurred.

5. **Correction factor calculation** The correction factor, ϵ_I, between simulated and experimental signals is calculated using Eq. (5.1), where t_{FE} and $t_{exp.}$ are the times which correspond to comparable events in the FEA and experimental results, respectively. This event may either be the time at which a detected peak occurs in the data, or (in the case of high activity regions) may be taken as the time at which the 90th percentile of colour play activity occurs. When using peak finding, the last detected peak is used for this calculation, as this allows finer corrections to HTC values as opposed to earlier peaks (this has previously been observed by Ryley).

$$\epsilon_I = \sqrt{\frac{t_{FE}}{t_{exp.}}} \tag{5.1}$$

6. **Sign flipping** Sign flipping is performed on those elements whose HTC values tend towards $0\,W{\cdot}m^{-1}{\cdot}K^{-1}$ without improvement in the agreement between the simulated and experimental intensities. Although the 3D HTC solver allows this flip to occur in either direction, in the vast majority of cases where it is employed, it is used to introduce negative HTC values to nodes which are in post-shock or wake regions, and hence are far more likely to radiate heat than nodes on which hypersonic flow is incident. Essentially, flipping of the HTC sign acts to reverse the direction of heat transfer for a given node.

   For a node to undergo a sign flip, it must have an HTC magnitude less than $0.1\,W{\cdot}m^{-2}{\cdot}K^{-1}$, its previous HTC values must predominantly tend towards $0\,W{\cdot}m^{-2}{\cdot}K^{-1}$, and the scaling factor calculated using Eq. (5.1) must be larger than

1%, i.e. $0.99 > \epsilon_I > 1.01$, indicating a non-negligible error.

7. **Update of HTC profile** The HTC profile for the next FEA simulation is calculated using Eq. (5.2) (which is identical to that used by Ryley - see Eq. (2.47)). In this expression, the term $h_{FE_{old}}$ refers to the HTC values from the last FEA simulation performed, while and $h_{FE_{new}}$ refers to the HTC profile for the next simulation. The symbol ϕ is a damping factor that may be employed to combat divergence in the 3D HTC solver's iterative solution, and ϵ_I is employed as the correction factor (as described above).

$$h_{FE_{new}} = \phi\ \epsilon_I\ h_{FE_{old}} + h_{FE_{old}}\,(1-\phi) \tag{5.2}$$

The damping factor, ϕ, was set to 1 throughout all calculations performed using the 3D HTC solver, as Eq. (5.2) proved to be self-damping (this was also observed by Ryley). As such, Eq. (5.2) may be reduced to the form shown in Eq. (5.3).

$$h_{FE_{new}} = \epsilon_I\ h_{FE_{old}} = \sqrt{\frac{t_{FE}}{t_{exp.}}}\ h_{FE_{old}} \tag{5.3}$$

Finally, any nodes whose HTC values fall outside the interval $(-20, 150)\,W{\cdot}m^{-2}{\cdot}K^{-1}$ during this part of an analysis are set to $0\,W{\cdot}m^{-2}{\cdot}K^{-1}$ and tagged such that they remain at zero throughout all subsequent iterations. This measure prevents divergence of the solution by removing problem nodes from the FEA simulations, and is primarily employed when peak finding or percentile calculation have failed to correctly estimate a node's ϵ_I value. Note that the range of acceptable values for HTC, i.e. the range $(-20, 150)\,W{\cdot}m^{-2}{\cdot}K^{-1}$ was selected based upon early DSMC results. As such, over or underestimations in HTC values could be accommodated during the course of the iterative process prior to a solution being converged upon, but divergence (which would severely disrupt the FEA solution) would be mitigated.

Examples of this iterative adjustment of HTC values are presented in Figures 5.9 (which shows the error between simulated and experimental signals at the beginning of an analysis), and 5.10 (which shows convergence, i.e. agreement between simulated and experimental

signals).

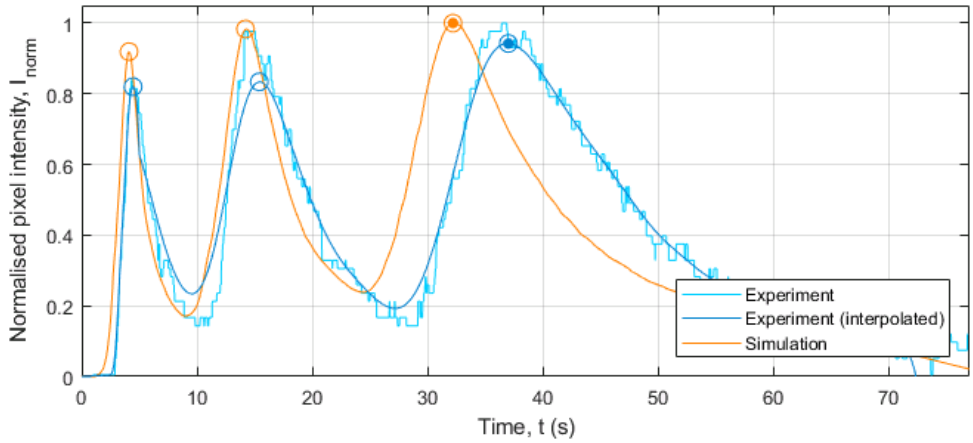


Figure 5.9: An example of the error between initial simulated LC activity and experimental values. Detected peaks and Savitzky-Golay filtered signals are visible in the plot.

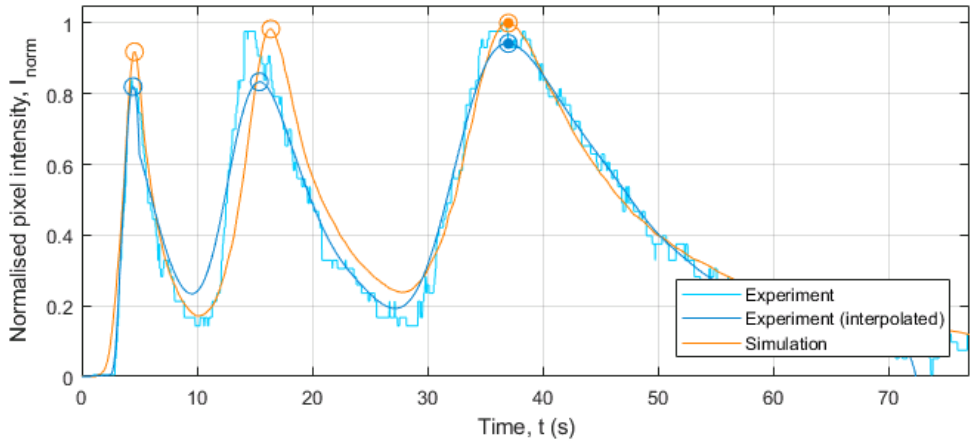


Figure 5.10: An example of convergence in the 3D HTC solver showing agreement between simulated LC activity and experimental values. Detected peaks and Savitzky-Golay filtered signals are visible in the plot.

8. **Generate plots / save iteration data** Inspection plots are generated and all variables from the current iteration are saved to disk. As such, the history of a calculation performed using the 3D HTC solver may be examined for problems or to monitor convergence. A GUI was written for this purpose using the MATLAB `guide` tool (this interface was heavily inspired by that of THTAC [135]). A screenshot of this tool showing its use in the post processing of an experimental case is presented in Figure 5.11.

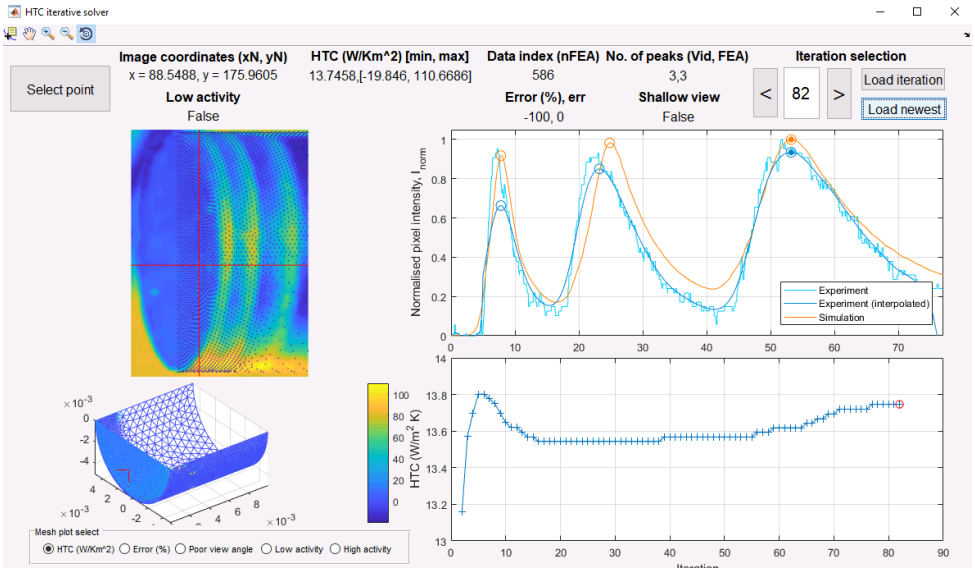


Figure 5.11: Screenshot of the GUI for the 3D HTC solver showing the successful fitting of a pixel signal for an experimental case.

9. **Start next FEA iteration** The newly updated HTC profile is saved to disk, and the MATLAB-COMSOL LiveLink interface is called in order to start another FEA simulation. The preceding steps are then looped until either a specified number of iterations have been completed, or the halt flag (a small text file which is read at the beginning of each iteration) is modified by the user.

5.4 CHAPTER SUMMARY

In this chapter, the experimental techniques used to measure heat transfer in the LDWT and the methods used to post process these data are discussed.

The locations of interest in the experimental plume are identified from CFD results presented in Chapter 4 and are subsequently discussed. The freestream conditions at these locations are calculated, with their spatial positions in the flowfield being expressed in terms of their axial distance from the nozzle exit. Uncertainties in the calculated freestream properties are also presented; these are derived from the accuracy of the flow measurement equipment installed in the LDWT (such as thermocouples and pressure transducers), as stated by their respective manufacturers. As the aforementioned calculations are based upon measurements recorded in the LDWT using these sensors, it is expected that the presented uncertainties will be a conservative

estimate of the actual range of conditions generated in the experimental plume of the LDWT.

The shapes and major dimensions of the PMMA probe bodies designed for the present work are then listed, and the calibration procedure of the liquid crystal coatings is described. The calibration data for the batch of liquid crystal slurries mixed for the present work is presented, showing the three distinct colour transitions in the red, green, and blue channels of the calibration video. The relationship between these components of the LC colour response and the surface temperature is also clearly shown, with the three transition temperatures of 25^oC, 30^oC, and 35^oC clearly corresponding to the fringes of the LC colour bands observed during the course of the calibration video.

Finally, procedures used to analyse the experimental videos are described, and the rationale for their selection discussed. A series of preprocessing steps which must only be performed once prior to a video being analysed are listed. These are then followed by descriptions of the algorithms used to derive surface distributions of heat transfer coefficient (HTC, or h_{conv}) from the video data.

EXPERIMENTAL RESULTS 6

This chapter deals with the presentation, analysis, and discussion of post-processed results from transient thermal experiments performed in the LDWT. These studies comprise 16 individual experiments, utilising four separate LC probes. The first two of these were performed using the appropriately sized hemispherical probes at an attitude of $\alpha = 0^o$ in both conditions of interest in the LDWT plume (see Section 4.3.6).

The latter 14 experiments were performed using the two cylindrical geometries at seven different attitudes $\alpha = \{0^o, 15^o, 30^o, 45^o, 60^o, 75^o, 90^o\}$. Like the hemisphere experiments, each of the cylinder attitudes were run at both locations of interest in the LDWT plume, utilising the appropriate size of probe (since both the probe diameters and the locations of interest in the LDWT plume were selected based upon desired freestream Knudsen numbers).

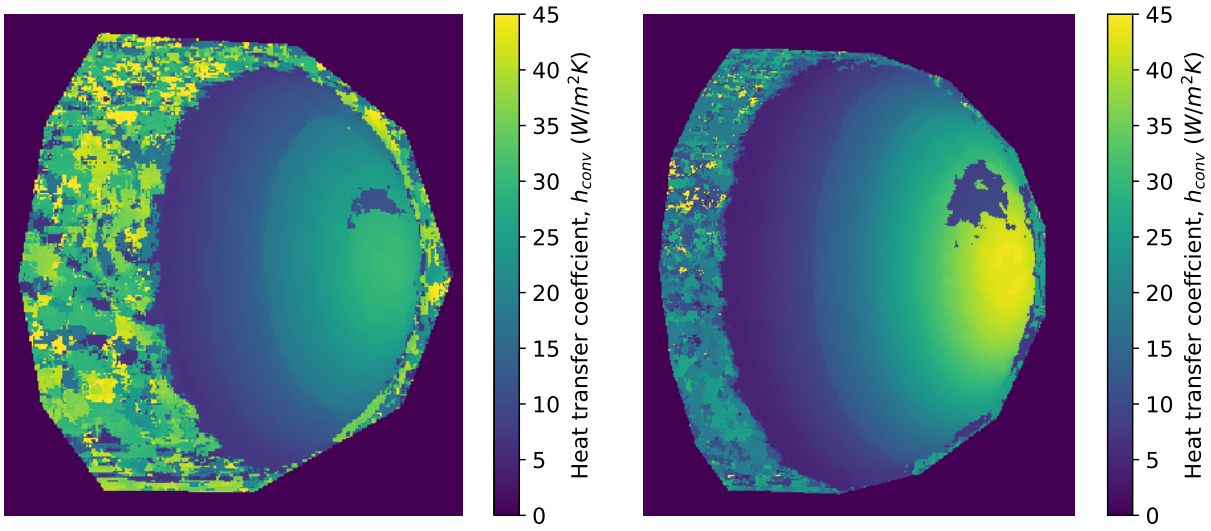
6.1 HEMISPHERE

The hemisphere cases were run in the LDWT at a single angle of incidence ($\alpha = 0^o$) at both examined flow conditions (as defined in Table 5.1). The resulting experimental videos were analysed using THTAC, and were processed up to the point that the semi-infinite principle was invalidated ($25.31\,s$ for the $\varnothing = 15\,mm$ probe, and $11.25\,s$ for the $\varnothing = 10\,mm$ probe). Contour plots of the resulting HTC distributions are presented in Figure 6.1.

As can be observed from the THTAC results, the derived distributions of HTC are extremely smooth across both cases, with few asymmetric or erroneous features. As expected, the HTC values are noticeably higher for the $Kn = 0.0216$ case, owing to the high number densities at the examined location in the hypersonic plume.

The HTC distributions on the cylindrical sections of the probes show significant noise, and are the result of there being no colour play data at these locations (since the experimental video was cut short in order to preserve the semi-infinite substrate assumption), resulting in what is essentially random fitting of HTC values by THTAC. Similarly, the small patches of low HTC on the upper portions of the hemispheres are the result of glare on the probe surfaces caused by the internal lighting of the LDWT test chamber. This causes the pixel intensities in these locations

to become saturated, thereby disrupting THTAC's ability to analyse the colour play on the probe surfaces. As such, both of these features may be considered erroneous, and may be disregarded when considering the solutions.



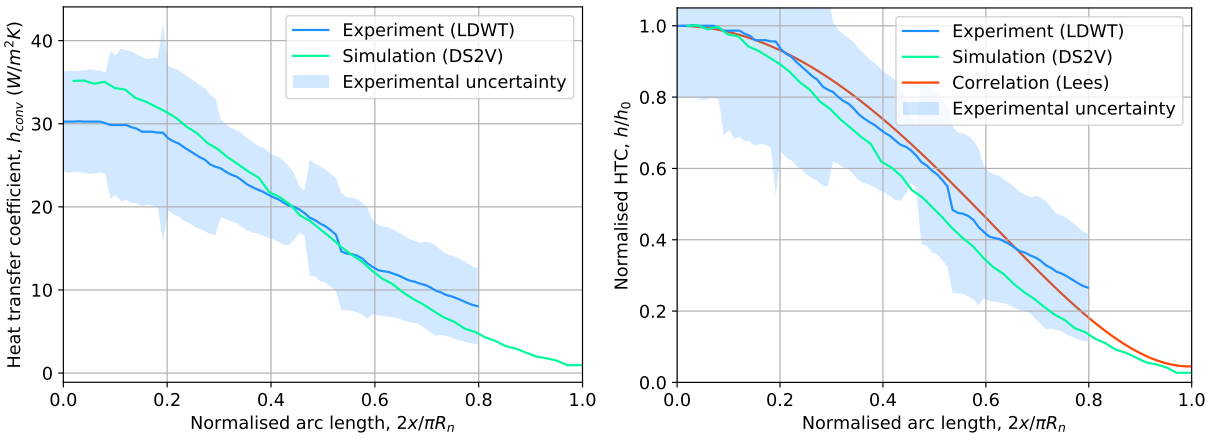
(a) $\varnothing = 10\,mm$, $Kn = 0.1012$ (b) $\varnothing = 15\,mm$, $Kn = 0.0216$

Figure 6.1: Contour plots of experimental convective heat transfer coefficient, h_{conv}, over hemispherical probes. These data were obtained using THTAC (one dimensional heat transfer analysis).

A circular arc with a central angle of $\pi/2$ and comprising 100 sampling points was defined in order to extract data from the THTAC results for further examination. These points were translated such that the end of the arc was coincident with each hemisphere's stagnation point, and were then rotated using Euler matrices to match the orientation of the probes in space. Note that the arc length of this sampling series is defined in the same fashion as those in Section 3.4.1 (see Figure 3.22).

The data extracted using this process, as well as the results of 2D axisymmetric DSMC simulations run using *DS2V* (see Section 3.4.3), are shown in Figures 6.2 and 6.3. The simulated HTC values presented here are generated by first obtaining the heat transfer rate distribution, $\dot{q}$, using an isothermal *DS2V* simulation. These data are then divided by the difference between the flow stagnation temperature, T_0 (as calculated by an adiabatic simulation using the same geometry and freestream conditions), and the wall temperature (which was set to $293\,K$).

$$\epsilon_{RSS} = \sqrt{\sum_{i=1}^{N} \epsilon^2} = \sqrt{\sum_{i=1}^{N} \big(y_i - f(x_i)\big)^2} \tag{6.1}$$



(a) Heat transfer coefficient, h_{conv} (b) Normalised heat transfer coefficient, h_{conv}/h_0

Figure 6.2: Plots of experimental and simulated convective heat transfer coefficient, h_{conv}, over a hemisphere with diameter 10 mm at LDWT conditions ($Kn = 0.1012$).

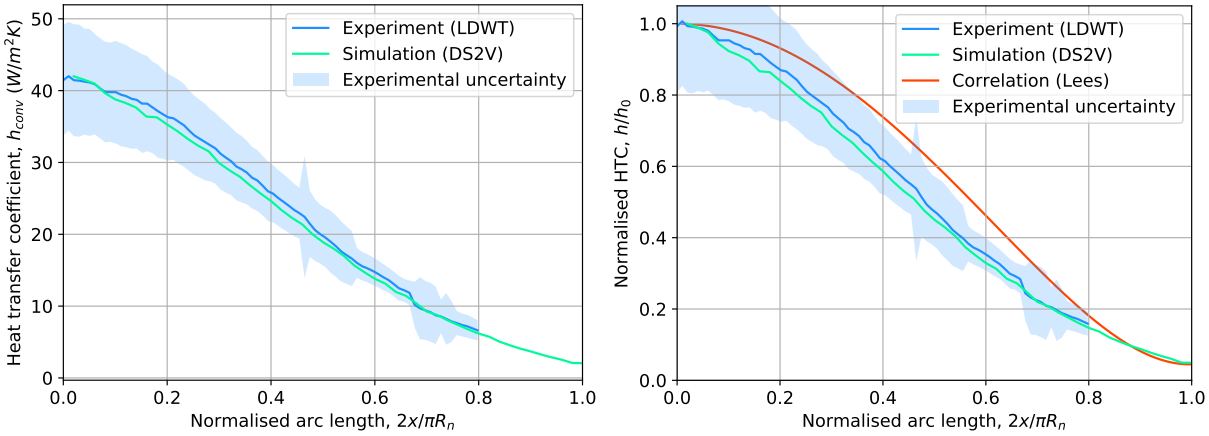
An error envelope is also plotted in these figures, showing the overall RSS (root of sum of squares) uncertainty in the HTC values output by THTAC. These uncertainties were characterised internally by THTAC for each individually processed pixel, and take into account possible variations in initial probe temperature, gas temperature rise time, material thermal product, liquid crystal calibration temperatures, and video time offset. The overall RSS error, ϵ_{RSS}, which is defined in Eq. (6.1) as the sum of the squares of all other errors, is also calculated by THTAC. A summary of the calculated uncertainties at the hemisphere stagnation point is presented in Table 6.1.

Parameter	**Variation**	**Contribution to HTC uncertainty**	
		$\varnothing = 10\,mm$	$\varnothing = 15\,mm$
Initial temperature	$\pm 2\,K$	18.3%	14.1%
Video start time offset	$\pm 1\,s$	6.9%	8.3%
Thermal product (Perspex)	$\pm 29\,J{\cdot}s^{-0.5}{\cdot}K^{-1}{\cdot}m^{-2}$	5.5%	4.8%
LC calibration temperatures	$\pm 0.25\,K$	2.7%	2.0%
Gas temperature rise time	$\pm 2\%$	2.1%	2.0%
HTC discretisation error	–	0.7%	0.7%
Overall RSS uncertainty	–	20.6%	17.3%

Table 6.1: Contributions to experimental uncertainty at the stagnation point of the hemispherical probes.

The largest source of uncertainty at the stagnation point of the hemisphere is the probe's initial

temperature, T_i, which was perturbed by $\pm 2\,K$ during uncertainty quantification. This resulted in uncertainties of 18.3% and 14.1% for the $\varnothing = 10\,mm$ and $\varnothing = 15\,mm$ probes, respectively. The next greatest source is the temporal offset of the experimental video, which is assumed to coincide with the start of heating in an experiment. This variable was perturbed by $\pm 1\,s$ and resulted in uncertainties of 6.9% and 8.3% for the $\varnothing = 10\,mm$ and $\varnothing = 15\,mm$ probes, respectively. This is followed by variation in the thermal product of the probe substrates, $\sqrt{\rho\, C_P\, k}$, which was perturbed by $\pm 29\,J{\cdot}s^{-0.5}{\cdot}K^{-1}{\cdot}m^{-2}$, an uncertainty value determined by Ireland [137]. The experimental uncertainty caused by this variation is 5.5% and 4.8% for the $\varnothing = 10\,mm$ and $\varnothing = 15\,mm$ probes, respectively.



(a) Heat transfer coefficient, h_{conv} (b) Normalised heat transfer coefficient, h_{conv}/h_0

Figure 6.3: Plots of experimental and simulated convective heat transfer coefficient, h_{conv}, over a hemisphere with diameter $15\ mm$ at LDWT conditions ($Kn = 0.0216$).

Variations in the liquid crystal calibration temperatures, gas temperature rise time, and HTC discretisation introduce far smaller uncertainties than the other contributors, at 2.7%, 2.1%, and 0.7% (respectively) for the $\varnothing = 10\,mm$ probe, and 2.0%, 2.0%, and 0.7% for the $\varnothing = 15\,mm$ probe.

It is clear from these data that the largest sources of uncertainty are the probes' initial temperatures, the starting times of experimental videos, and the material of which the probes are constructed. Although the latter of these is difficult to control, the former two are far easier to mitigate against. As such, the initial probe temperatures were measured prior to the start of experiments using a thermocouple affixed to each probe's afterbody (the unshaded sections with $\varnothing = 12\,mm$ shown in Figure 5.2) using polyimide (Kapton) tape. Although these single

measurements were made in situ, and each probe was allowed ~ 2 hours to equilibrate in the LDWT test chamber, they were still taken at a point relatively far from those sections of the probes under examination, hence the large value of $\pm 2\,K$.

Additionally, the flow shield described in Section 4.2.4 (which is responsible for blocking the hypersonic plume, hence preventing it from interacting with the probes prior to the start of an experimental run) is visible at the beginning of each of the experimental videos, and so could be observed on a frame-by-frame basis during its transit out of the flow. This, coupled with the experimental recording frame rate of 60 FPS, allowed the start time of each experiment to be measured relatively accurately. The perturbation value of $\pm 1\,s$, then, may be considered somewhat conservative.

Considering Figure 6.2(a), it may be observed that the $\varnothing = 10\,mm,\ Kn = 0.1012$ case shows good agreement in terms of absolute HTC values, with a relative error of 13.7% near the stagnation point, and errors which subsequently tend towards zero as $2x/\pi R_n \rightarrow 0.4$. A region of excellent agreement follows this at $0.4 \leqslant 2x/\pi R_n \leqslant 0.6$, finally leading to a maximum error of $\sim 45\%$ towards the probe's cylindrical section (where $2x/\pi R_n \rightarrow 0.8$). It should be noted, however, that the absolute error between experimental and simulated values here is in fact smaller than at the stagnation point.

When normalised to their respective local stagnation values (as in Figure 6.2(b)), the trends of HTC show similarly good agreement between simulation and experiment, with steadily increasing errors as $2x/\pi R_n \rightarrow 0.8$, up to a maximum of 50.2% at $2x/\pi R_n = 0.8$. The correlation of Lees [126] is also plotted alongside these normalised data, but disagrees slightly with the *DS2V* simulation results where $0.1 \leqslant 2x/\pi R_n \leqslant 0.8$. Lees' curve appears to predict the experimental HTC trend with greater accuracy than the *DS2V* simulation, although it is important to note that all disagreements between these data fall within the experimental uncertainty interval.

As is apparent from Figures 6.3(a) and 6.3(b), both the absolute and normalised results from the $\varnothing = 15\,mm,\ Kn = 0.0216$ case show excellent agreement between simulated and experimental HTC measurements. Considering the normalised data, it is evident that Lees' correlation does not predict the normalised trend nearly as well as *DS2V*, as the maximum absolute error (in terms of h/h_0) between correlation and experiment is 0.124. This is

significantly higher than the error between the *DS2V* simulation results and experimental values, which, for this normalised case, has a maximum magnitude of 0.05. The maximum percentage error, δ, between the absolute values of simulated and experimental HTC is $\sim 7\%$ at $2x/\pi R_n \approx 0.67$, while the maximum absolute error, ϵ, is $\sim 1.96\,W{\cdot}m^{-2}{\cdot}K^{-1}$ at $2x/\pi R_n \approx 0.16$.

It may be concluded from these cases that DSMC simulations run using *DS2V* produce excellent overall agreement with experimental results for hemispheres at Knudsen numbers in the slip and lower transition regimes. It is also evident that Lees' correlation for the normalised heat flux trend on blunt bodies is unsuitable for fine predictions in this regime. This conclusion is supported by the fact that Lees formulated this correlation based upon bounding cases in hypersonic flow, where the stagnation temperature is significantly higher than the wall temperature i.e. $T_0 \gg T_w$. This difference is significantly lower for the cases examined herein than in those studied by Lees [126] ($65.3\,K$ and $75.6\,K$ for the $\varnothing = 10\,mm$ and $\varnothing = 15\,mm$ cases, respectively, compared to a stagnation temperature of $385\,K$), hence providing an explanation for the discrepancy. (Using the assumption that $T_0 \gg T_w$ in re-entry flows allowed significant simplifications to be made by Lees, and adapting his expressions for comparable values of T_0 and T_w is outwith the scope of the present work.)

6.2 CYLINDER

The cylinder probes were examined at both flow conditions defined in Table 5.1, and were also inclined relative to the freestream direction in increments of 15^o in the range ($0^o \leqslant \alpha \leqslant 90^o$). The resulting experimental videos were analysed using the 3D HTC solver (3D-HTCS) described in Section 5.3.3 in order to determine the experimental heat transfer coefficients. The processed data were then sampled for plotting using the locations defined in Figure 3.22.

Simulated heat transfer coefficients were calculated for comparison with these cases using the DSMC simulations described in Section 3.4.4 in conjunction with Eq. (6.2). This expression requires two separate DSMC simulations as inputs: an isothermal case (with a fixed wall temperature, $T_w = 293\,K$) which is used to determine heat transfer rate, $\dot{q}$, and an adiabatic case (with zero heat transfer, i.e. $\dot{q} = 0\,W{\cdot}m^{-2}$) which is used to determine recovery temperature, T_r,

(also called adiabatic wall temperature).

$$h_{conv} = \frac{\dot{q}}{T_r - T_w} \tag{6.2}$$

In the following sections, these data are presented as absolute values in Figures 6.8 through 6.20 for both flow conditions. In addition, normalised HTC values are presented for those cases run using the $\varnothing = 15\,mm$ probe (due to the non-optimal simulation parameters utilised in the accompanying DSMC runs - see Section 3.4.2). The normalisation factor for this flow condition is the stagnation point HTC value, h_0, at $\alpha = 0^o$ (sampled at the centre of the cylinder's circular face). A summary of these values is presented in Table 6.2 (for completeness, both flow conditions are listed). In each subsequent section, the results from the sampling locations defined in Figure 3.22 are presented in the following order: front stagnation line, side stagnation line, front circular edge, side circular edge.

	Flow condition	
Stagnation point HTC, h_0	**1**	**2**
Experimental	20.55	13.79
Simulated	23.44	12.14

Table 6.2: Stagnation point HTC values (h_0) sampled at the centre of the cylinder probe's circular face at $\alpha = 0^o$. Values are presented in W/m^2K

The purpose of this normalisation strategy is to allow the attitude dependant trends of heat transfer coefficient (HTC) to be correlated in terms of a single stagnation heat flux at $\alpha = 0^o$. As such, the trends for a given geometry may subsequently be reconstructed for any set of freestream conditions that feature the same (or a similar, i.e. same order of magnitude) Knudsen number, as long as the appropriate stagnation heat flux, adiabatic wall temperature, and hence HTC at $\alpha = 0^o$ can be accurately predicted (see Section 2.3.3).

Uncertainties in the HTC results were calculated based upon the $\alpha = 45^o$, case run at $Kn = 0.1$ (condition 2 in Table 5.1) and applied to all cylinder cases. In total, six additional 3D HTC solver runs were performed, introducing uncertainties in video start time, t_i, initial probe temperature, T_i, and material thermal product, $\sqrt{\rho C k}$. The uncertainties applied during

this process are listed in Table 6.3, and the individual errors were combined for plotting using Eq. (6.1).

Variable	Symbol	Units	Nominal value	Minimum value	Maximum value
Video start time	t_i	s	0	-0.15	0.15
Probe initial temperature	T_i	^oC	21.1	18.9	23.3
Thermal product	$\sqrt{\rho C k}$	$J{\cdot}s^{-0.5}{\cdot}K^{-1}{\cdot}m^{-2}$	569	540	598

Table 6.3: Uncertainties applied to the $\alpha = 45^o$, $Kn = 0.1$ 3D HTC solver simulations.

6.2.1 Performance of the 3D HTC solver code

The 3D HTC solver developed for the present work proved to be invaluable during the course of the present work, as it allowed the post processing of colour play videos generated by LC experiments which could not be treated using the semi-infinite assumption. However, as a result of several factors — which will be discussed shortly — multiple artefacts and errors manifested during the course of extracting these data. These issues are explored here such that subsequent sections may focus more on discussing the results themselves.

Figure 6.4 shows the HTC distributions generated by both DSMC simulations and the 3D HTC solver (derived from an experimental video). This particular case was run at flow condition 2 (see Table 5.1) with the $\varnothing = 10\,mm$ cylinder probe oriented at $\alpha = 45^o$ relative to the freestream direction. The video was processed up to a point $60\,s$ after the start of the experiment using both peak finding and percentile matching (see Section 5.3.3).

It is clear that the simulated HTC distributions (presented on the bottom half of the cylinder in Figure 6.4) are significantly smoother than those derived from experimental results. This phenomenon comes about as a result of the 3D HTC solver's formulation, whereby individual nodes are only capable of affecting one another during the course of an FEA calculation. Whilst analysing data from experimental videos, however, the colour intensity histories for individual nodes are processed in isolation (see Section 5.3.3). As such, issues like poor video quality, rapid colour transitions brought on by high heat transfer rates, or inadequate signal filtering can hamper the 3D HTC solver's ability to properly calculate errors between simulated liquid crystal

intensities and actual experimental video values. This, in turn, can lead to large disagreements between adjacent nodes, and hence, inaccuracies in the final HTC values. In order to mitigate the effects of these errors, the colour scales of Figures 6.7 and 6.16 in subsequent sections have been adjusted to conform to the DSMC results, thereby framing the available experimental data in a far more useful manner, i.e. the limits of the scale have been set so that severe HTC overestimations from experimental data are outside of the plotted colour range.

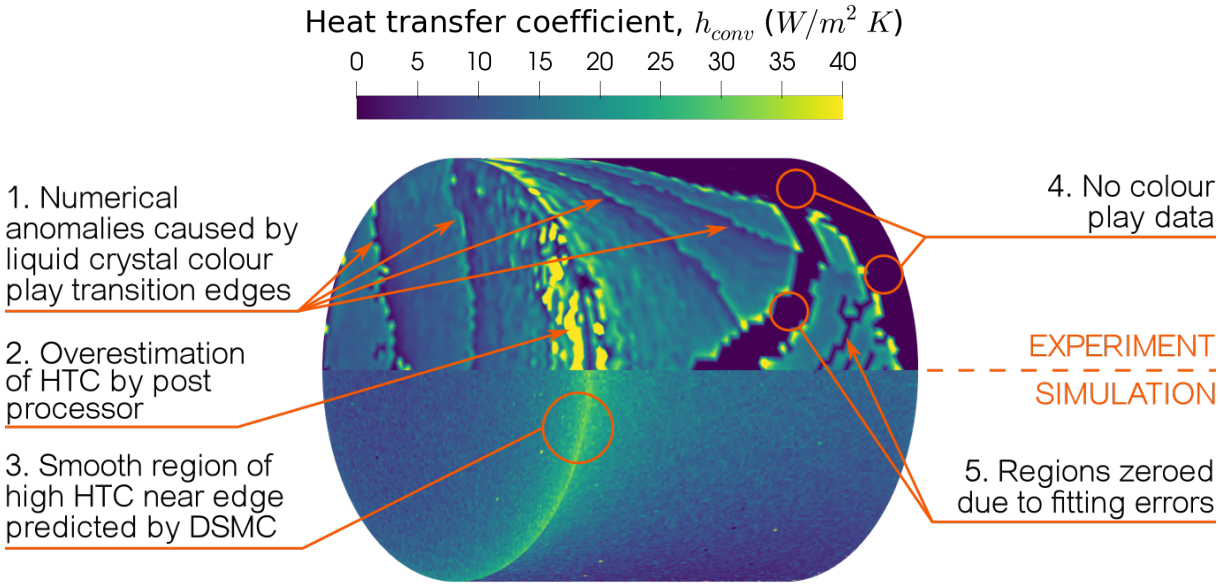
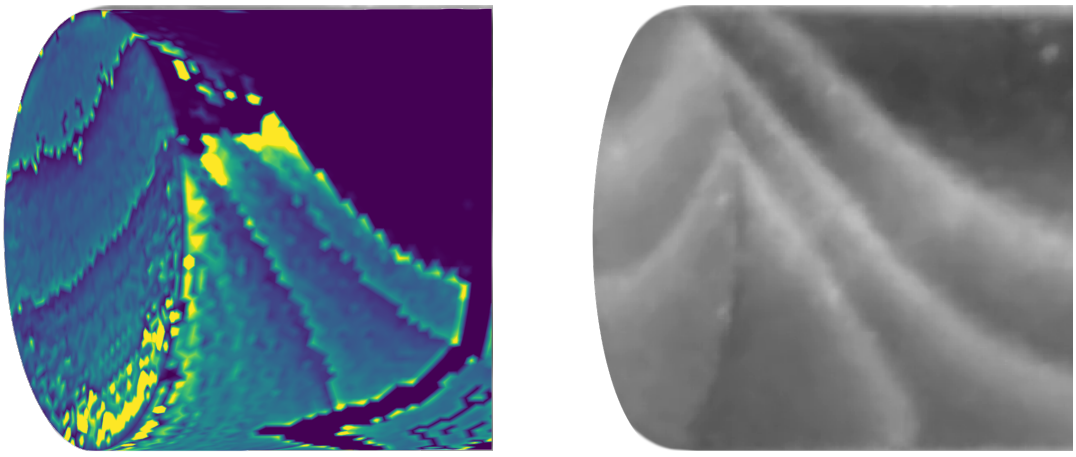


Figure 6.4: Diagram showing common errors encountered during use of the 3D HTC solver code. Heat transfer coefficient data from cases run at condition 1 ($Kn = 0.1012$), and oriented at $\alpha = 45^o$ are shown. DSMC results are presented on the bottom half of the cylinder, with results from the 3D HTC solver on top.

Several distinct artefacts are visible in the experimental data presented in Figure 6.4 which are not reproduced in the DSMC results. These are numbered in the aforementioned figure, and their causes discussed in depth below:

1. **Influence of colour play transition edges** As the 3D HTC solver adjusts the HTC values of individual nodes in the FEA mesh, there are inevitable discontinuities between those nodes that see different numbers of LC transitions. These LC activity thresholds can manifest in the FEA solution as anomalous ridges in the resulting HTC distribution.



(a) HTC distribution produced by 3D HTC solver (b) Final frame ($t = 60\,s$) of experimental video

Figure 6.5: Comparison of FEA results from the 3D HTC solver and the final processed frame of the associated experimental video.

As can be observed in Figure 6.5, the locations and shapes of these ridges correspond exactly to the final positions of the LC transitions in the experimental video. This effect is caused by proximate nodes having different numbers of LC transitions to compare with simulation results.

In the case shown, nodes close to the lower left circular edge of the cylinder see three transitions, while those at the top right of the figure see none. Nodes between these regions see one or two transitions depending on their positions. By default, the 3D HTC solver uses the final detected peak at a given node to determine the error between video and simulation (and, hence, corrections to the local HTC values). As a result of this behaviour, there are inevitable points in the solution where the solver must switch from using one peak consistently to another, e.g. utilising the third peak near the edge where heating is greatest, but only utilising the second peak slightly further downstream (since a third peak does not manifest at these locations). Figure 6.6 shows the structure of the FEA mesh used by the 3D HTC solver proximate to these ridges. It is clear from this image that the effects of switching from one set of peaks to another is consistently isolated to single nodes, further supporting the conclusion that these effects are non-physical.

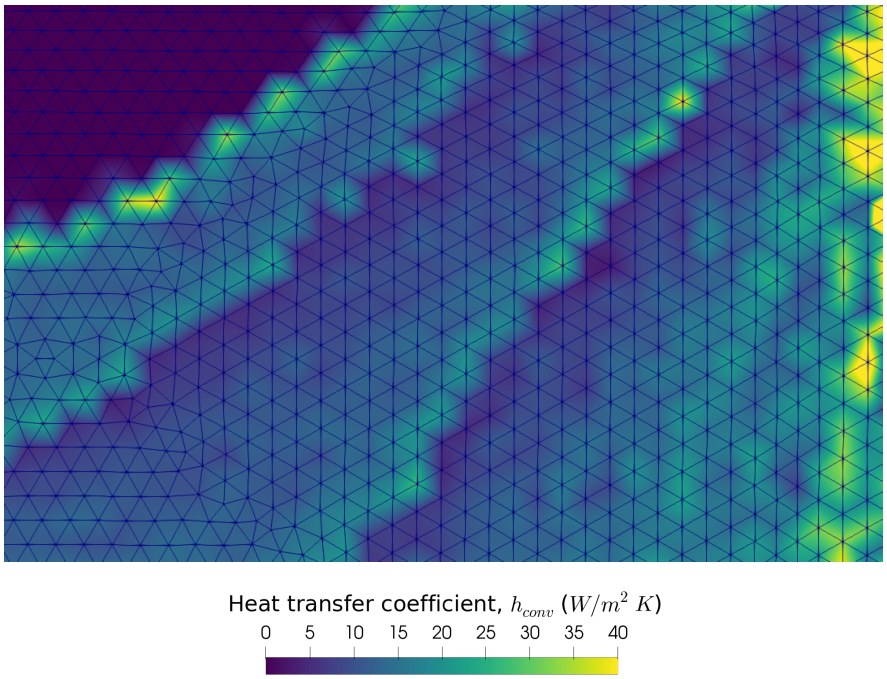


Figure 6.6: Enlarged view of the HTC distribution proximate to anomalous ridges. The FEA mesh used for in the 3D HTC solver is overlaid.

These ridges can also manifest when the 3D HTC solver switches between peak finding and percentile matching in high activity areas. This is a side effect of the percentile matching routine itself, which, although more robust, is less accurate than peak finding. As a result, regions in which this switchover occur can demonstrate a small step change in HTC.

These observations show that the aforementioned ridge-like features are non-physical numerical anomalies, and hence may be ignored.

2. **Overestimation of HTC values** During the correction factor calculation step of the 3D HTC solver (see Section 5.3.3), errors between the simulated (FEA) thermochromic LC response and that which is extracted from the experimental video are calculated. These data are then used in a subsequent step to correct the prediction of HTC at each individual node. If the 3D HTC solver fails to properly analyse a region in an experimental video (potential causes of which include video noise, low light intensity, incorrect filtering, etc.), the result of the error calculation between video and simulation is likely to be incorrect. As such, some nodes are set to consistently anomalous HTC values, producing corresponding anomalies in the calculated HTC distribution.

3. **Smoothness of HTC distribution** There are clear differences in the smoothness of the HTC distributions generated by experimental post processing versus simulations. However, these are simply side effects of the different methods used to calculate these distributions.

   The DSMC method (which produced the simulation data in Figure 6.4) utilises stochastic modelling of many millions of particulate collisions, the effects of which are averaged over time. This results in an extremely smooth HTC distribution, with small amounts of stochastic noise between cells which diminish in intensity as the simulation continues to perform averaging. The results of the 3D HTC solver, however, rely upon corrections to HTC values for individual, isolated nodes (as is mentioned above). Taking the disconnected nature of these nodes during this step of the calculations into account, it is clear that visible discrepancies in the smoothness of these results are non-physical, numerical anomalies. An FEA mesh which is too sparse could also contribute to this effect manifesting.

4. **Lack of colour play data** Groups of nodes exhibiting low video activity (as identified by the standard deviation check described in Section 5.3.3) have their HTC values set to zero during preprocessing. Low video activity is indicative of regions in the experimental video where there are no thermochromic transition events of a probe's LC coating. Therefore, such regions must logically undergo negligible heating, and are explicitly zeroed in order to simplify subsequent iterations of the 3D HTC solver.

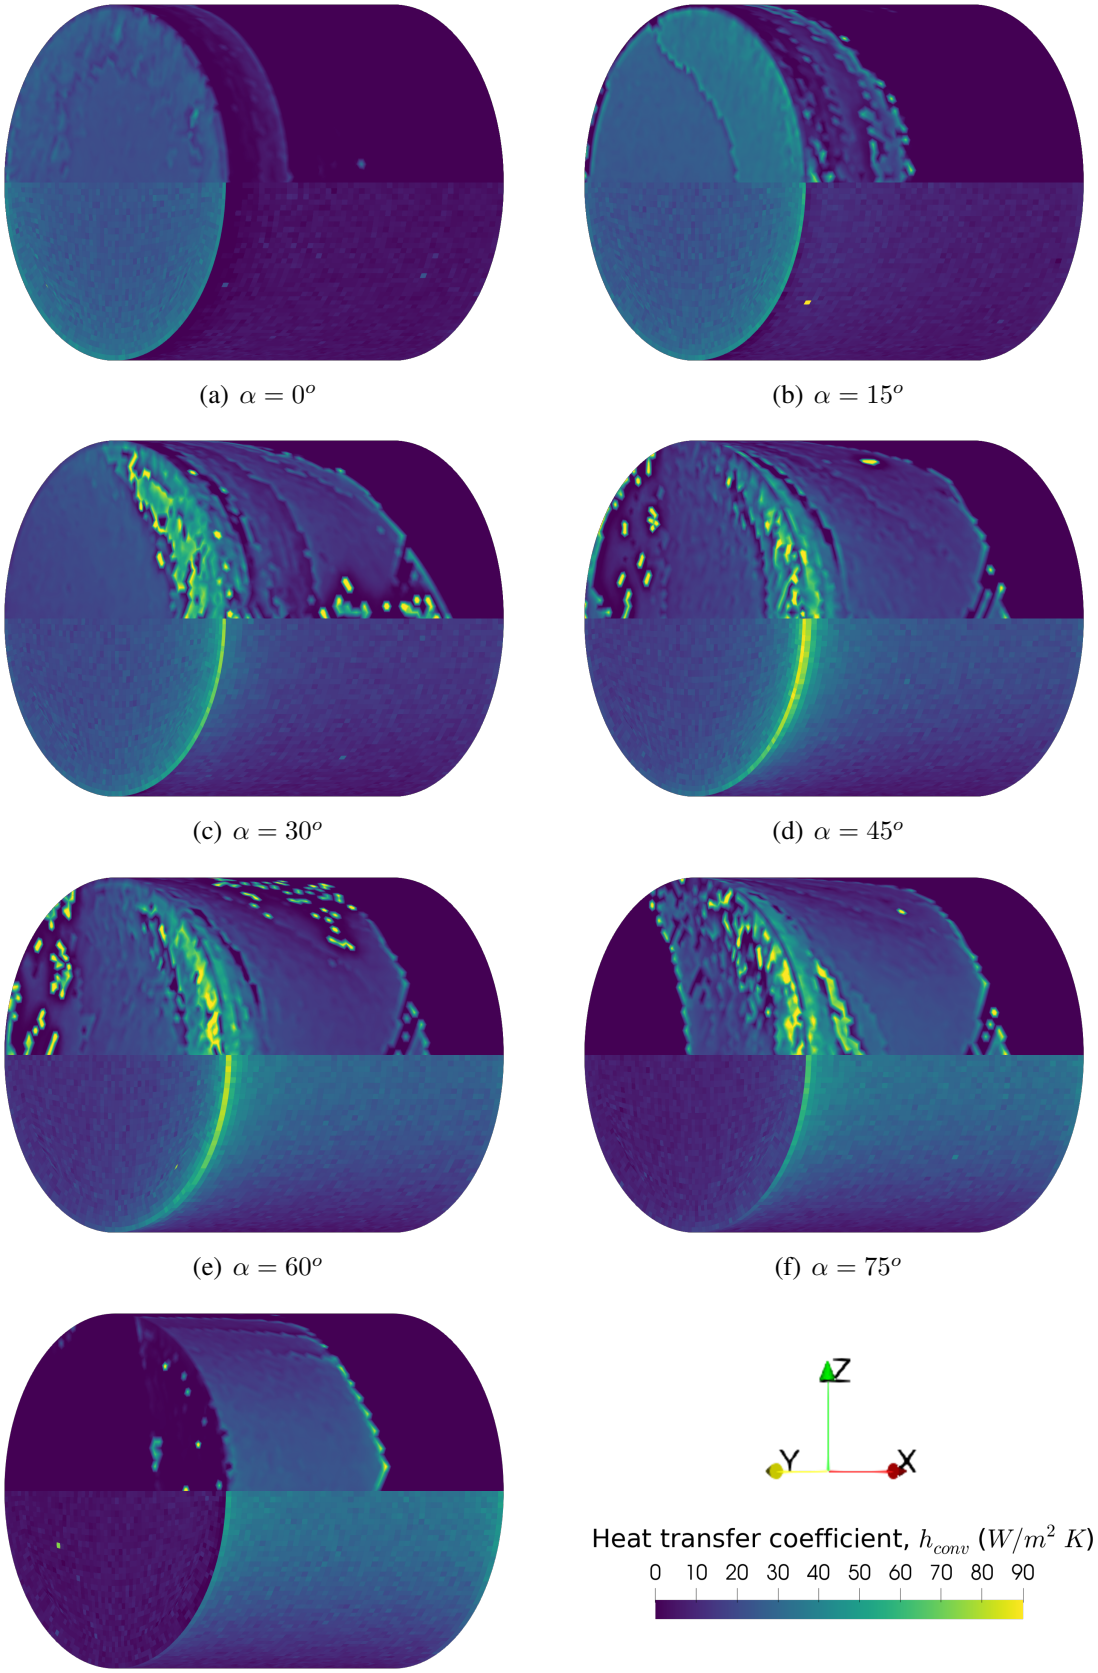
5. **Zeroing of HTC values** During the aforementioned correction factor step of the 3D HTC solver, any adjustment to a node's HTC value which would cause it to fall outside of the interval $(-20,\ 150)\, W \cdot m^{-2} \cdot K^{-1}$ is ignored. The node's HTC is then automatically set to $0\, W \cdot m^{-2} \cdot K^{-1}$, and the node eliminated from subsequent iterations of the 3D HTC solver. This is to prevent numerical divergence in the FEA simulations, whereby small groups of nodes with erroneously high HTC values dominate the thermal behaviour of the FEA simulations. Zeroing is usually triggered by regions of an experimental video in which colour transitions occur very quickly (due to a high local heat transfer rate), or where uneven lighting was employed during an experiment.

6.2.2 Flow condition 1 ($\varnothing = 15\,mm$, $Kn = 0.02[16]$)

The first set of tests utilised the $\varnothing = 15\,mm$ cylinder probe at a freestream Knudsen number of $Kn = 0.02(16)$. The results of these experiments, as well as those generated by DSMC simulations, are presented in Figures 6.7 through 6.15. Figure 6.7 shows the surface HTC distributions generated by both experiments and DSMC simulations. In each subfigure, experimental results are presented on the upper half of each cylinder, while simulation results are presented on the lower half. Figures 6.8, 6.10, 6.12, and 6.14 show the absolute HTC values, while Figures 6.9, 6.11, 6.13, and 6.15 show these HTC values normalised to the local stagnation HTC, h_0, as sampled from the associated $\alpha = 0^o$ case (see Table 6.2).

As is elaborated upon in Section 3.4.4, the DSMC simulations for this case were performed using *dsmcFoam+*, and utilised sub-optimal DSMC parameters in order to reduce computational expense. As such, exact matching of absolute HTC values was not expected. However, normalising these data with the procedure described in Section 3.4.2, i.e. dividing by the stagnation value of heat flux, $\dot{q}$, or HTC, h_{conv} (as appropriate) at $\alpha = 0^o$, was expected to produce good agreement (as has previously been demonstrated using *DS2V* simulations in the aforementioned section).

Regarding Figure 6.7, it is clear that overall agreement between simulations and experiments is generally good. Trends of HTC are reproduced well, with the highest values manifesting towards the edge of the cylinders, with these maxima becoming more pronounced as $\alpha \rightarrow 45^o$. Similarly, on the lateral faces, trend agreement is also favourable.



(a) $\alpha = 0^o$ (b) $\alpha = 15^o$

(c) $\alpha = 30^o$ (d) $\alpha = 45^o$

(e) $\alpha = 60^o$ (f) $\alpha = 75^o$

(g) $\alpha = 90^o$

Figure 6.7: Surface plots of HTC for flow condition 1 ($Kn = 0.0216$) derived from simulated and experimental data. Experimental data is shown on the top of each subfigure, while DSMC results are shown on the bottom.

Many of the artefacts described in Section 6.2.1 manifest in these cases, with the notable exceptions of the $\alpha = 0^o$ and $\alpha = 90^o$ cases, which only suffer from a lack of colour play data on the side faces of the cylinder. The most prominent errors of any consequence occur close to the edge of the cylinder in cases with pitch angles in the range $30^o \leqslant \alpha \leqslant 90^o$. A lack of colour play data is evident in all cases, resulting in zeroing of the affected regions. Colour play edge effects can be seen in the all cases expect those run at $\alpha = 0^o$ and $\alpha = 90^o$. The transition of the 3D HTC solver from its percentile matching routine to peak finding also produced a series of ridges in the HTC distribution. These are most visible on the front circular faces of the cases with pitch angles in the range $15^o \leqslant \alpha \leqslant 75^o$. This effect can also be seen to manifest on the lateral curved faces of those cases run at pitch angles in the range $30^o \leqslant \alpha \leqslant 75^o$.

Agreement is generally excellent between those regions of the experimental and simulated HTC distributions which do not exhibit errors from the 3D HTC solver. Random nodes proximate to the stagnation line on the front and lateral faces (see Figure 3.22) of each case were sampled, as were the corresponding cells in the DSMC results. Errors between the two datasets were then calculated for each pair using Eq. (2.50). This strategy of selecting random node/cell pairs and rejecting problem nodes in the 3D HTC results was utilised in order to mitigate the effects of noise in the HTC distributions whilst limiting the influence of nodes which demonstrate erroneous behaviour. It also serves to show the accuracy of the 3D HTC solver when optimal conditions are achieved during processing.

Nodes coincident with the circular edge were not selected due to high concentrations of processing artefacts (mostly overestimations of HTC). These occurred due to the node-based formulation of COMSOL, which meant that thermal calculations were performed directly on the edge of the cylinder. As a result of this, heat transfer results at these nodes were influenced by both the front circular and lateral curved faces simultaneously. By sampling the nodes one row removed from the cylinder edge, the individual thermal environments experienced by each surface could be examined and characterised. Additionally, this also facilitated comparison with the *dsmcFoam+* results, which were calculated using a cell-based formulation. Like the COMSOL results, this would have resulted in heat transfer data from both sides of the cylinder edge interacting with another via interpolation during post-processing. Since these areas were

also expected to experience the highest heating rates, they are considered as regions of particular interest to the present study, as thermal concentrations near sharp features can lead to early demise of satellite components during re-entry. As such, data from these nodes are plotted in Figures 6.8 through 6.15 and examined separately.

In the case of the $\alpha = 0^o$ case, the error on the front face is 4.85%, and 4.92% on the lateral face. The $\alpha = \{15^o, 30^o, 45^o, 60^o\}$ cases show even better agreement, with front face errors of 1.2%, 3.7%, 2.7%, and 1.1%, respectively. On the lateral faces of these cases, the errors are 3.1%, 3.5%, 1.3%, 2.4%, respectively. Errors on front and lateral faces of the $\alpha = 75^o$ case are 28.9% and 6.2%, respectively. Finally, the $\alpha = 90^o$ case exhibits a 17.3% error in HTC on the lateral face, while the vast majority of its front circular face exhibits zeroing due to a lack of experimental colour play data. This was caused by the camera angle utilised during this case's experiment, and as such, errors in this region cannot be estimated.

The $\alpha = 60\%$ and $\alpha = 75\%$ cases show the most significant errors compared to others in that their highest HTC values occur consistently on the front faces, while in the associated simulation results these maxima occur on the lateral face. This is likely a side effect of poor viewing angle or lighting during the experimental run itself, as the majority of the hypersonic flow in these cases is incident on the lateral face. Hence, the highest HTC values should occur proximate to the circular edge on the lateral face.

Considering Figure 6.8, excellent agreement can be observed at $0 \leqslant x/R_n \leqslant 0.6$ for the $\alpha = 30^o$ case, and good agreement is visible in the $\alpha = 45^o$ and $\alpha = 60^o$ cases. As $x/R_n \rightarrow 1.0$, however, larger errors begin to manifest, likely as the result of poor intensity fitting during the course of the 3D HTC solver's analyses (see Section 5.3.3). The $\alpha = 0^o$ and $\alpha = 75^o$ cases exhibit poorer agreement, while the $\alpha = 90^o$ case is the most inaccurately predicted. Each of these cases exhibit similar spikes in HTC as $x/R_n \rightarrow 1$, which is likely a symptom of poor lighting in the LDWT during these experiments. This, in turn, hampers the 3D HTC solver's ability to properly calculate errors between simulated liquid crystal intensity values and actual experimental video values.

It should also be noted that poor intensity fitting can lead to severe numerical disruption in an FEA solution. If, for example, a node were to have its HTC value set extremely high, the activity

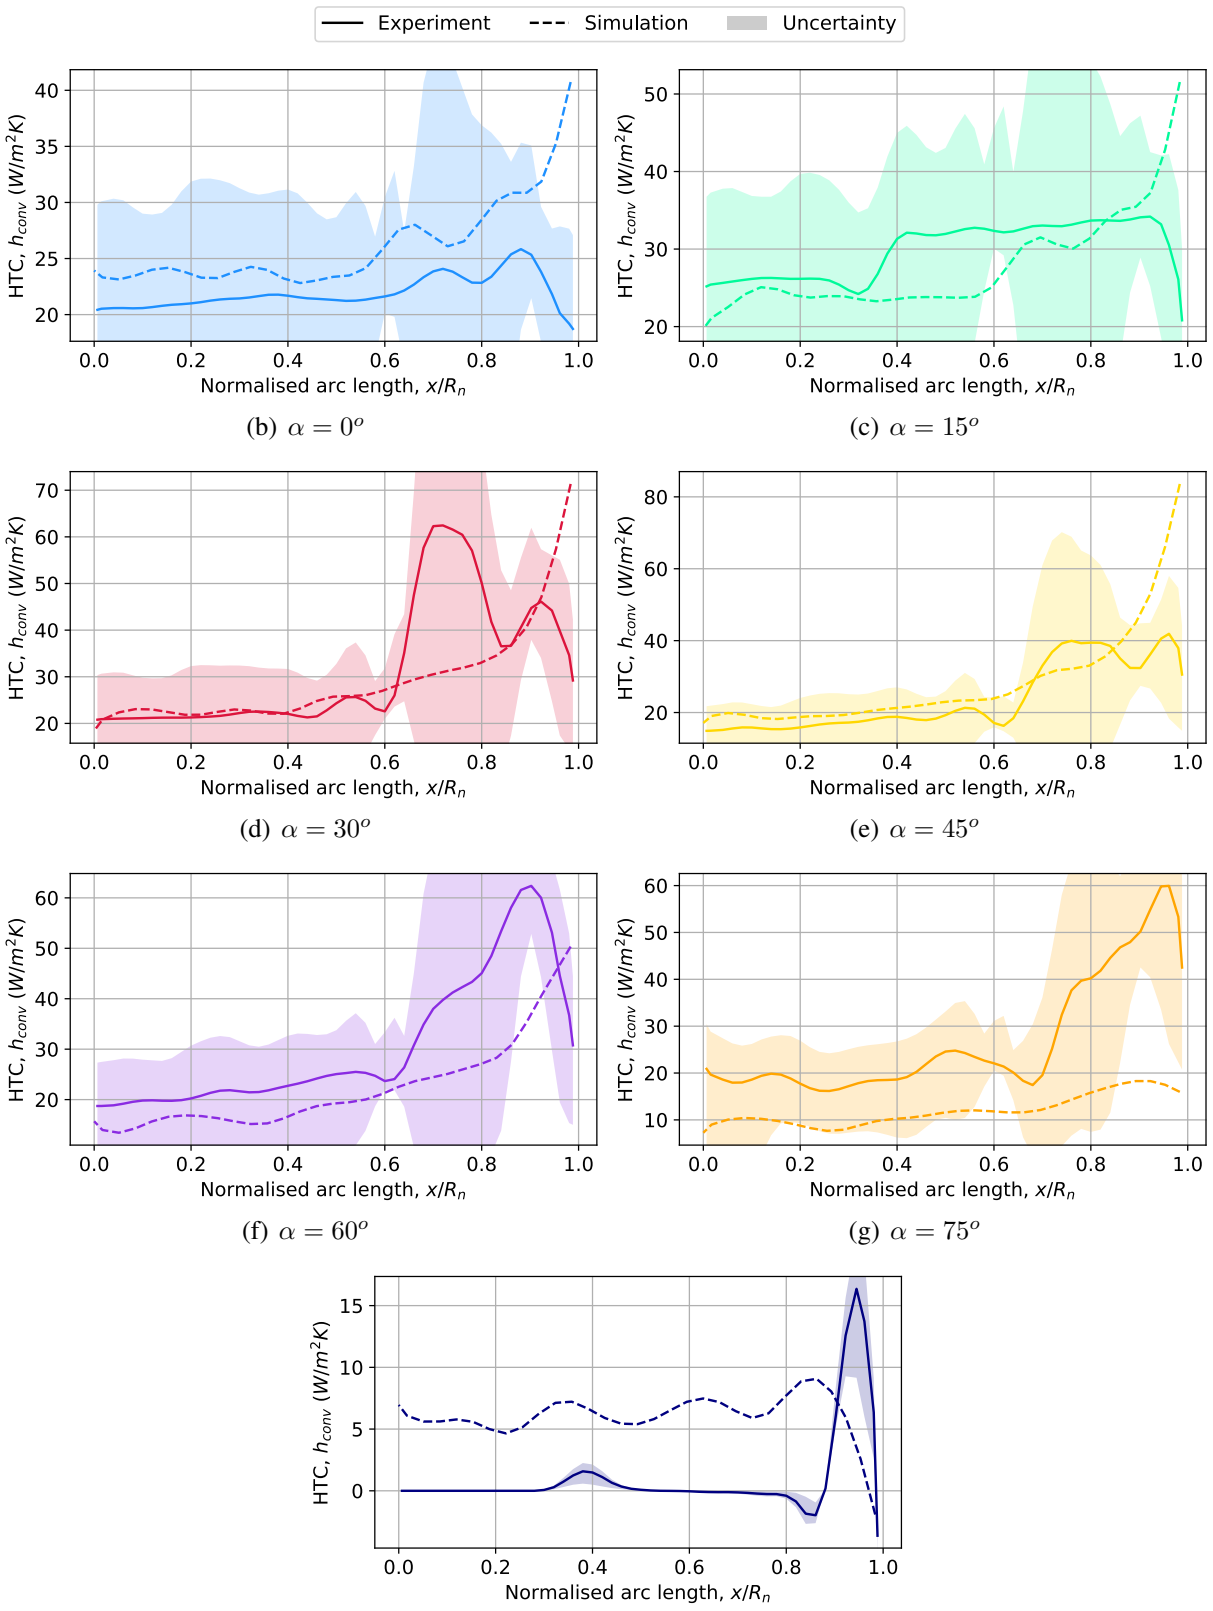
of this node would dominate the solution, massively skewing the results from other nodes in the simulation. As such, the 3D HTC solver is programmed to permanently set any nodes whose HTC values fall outside of the interval $(-20, 150)\, W{\cdot}m^{-2}{\cdot}K^{-1}$ to zero (see Section 5.3.3). The experimental HTC distribution for the $\alpha = 90^o$ case, then, may be observed to be the result of this HTC limiting.

Moving on to Figure 6.9, it is clear that agreement between normalised front stagnation line values is superior to that of the absolute measures in the previous Figure. This agreement is excellent for the $\alpha = 0^o$, $\alpha = 30^o$, and $\alpha = 45^o$ cases, and reasonably good for the $\alpha = 60^o$ and $\alpha = 75^o$ cases. As expected, the $\alpha = 90^o$ case shows the poorest agreement once again, owing to poor intensity characterisation and hence, HTC limits in the FEA solution being enforced.

The $\alpha = 15^o$ case is an outlier in this context, as it appears to feature a step in HTC at $x/R_n \approx 0.4$. This is due to the transition between high activity percentile fitting and standard peak fitting in the 3D HTC solver. As such, these data for this case may be taken to be mostly erroneous.

The side stagnation line results (presented in Figure 6.10) show reasonably good agreement for the $\alpha = 30^o$ case, although a great deal of noise is clearly present in the experimental data. Once again, this is likely due to poor fitting of experimental LC intensities brought on by poor lighting conditions. The $\alpha = 45^o$, $\alpha = 60^o$, and $\alpha = 75^o$ cases show good agreement in the range $0 \leqslant x/R_n \leqslant 1.2$, with reductions to $0\, W{\cdot}m^{-2}{\cdot}K^{-1}$ enforced by the 3D HTC solver manifesting shortly thereafter. The $\alpha = 90^o$ case also exhibits this behaviour, although agreement between simulated and experimental HTC measures is inferior to other values of α. Finally, the $\alpha = 0^o$ and $\alpha = 15^o$ cases show reasonable agreement in the ranges $0 \leqslant x/R_n \leqslant 0.4$ and $0 \leqslant x/R_n \leqslant 0.8$, respectively, after which they reduce to zero in order to prevent numerical divergence in the FEA solver.

The normalised HTC values from the side stagnation line are presented in Figure 6.11, and show slight improvements in agreement between simulated and expeirmental values. As one would expect from the previous figure, the $\alpha = 30^o$, $\alpha = 45^o$, $\alpha = 60^o$, and $\alpha = 75^o$ cases show the best improvement. Agreement for the $\alpha = 90^o$ case is also slightly improved, however the $\alpha = 0^o$ and $\alpha = 15^o$ cases continue to exhibit the largest errors.



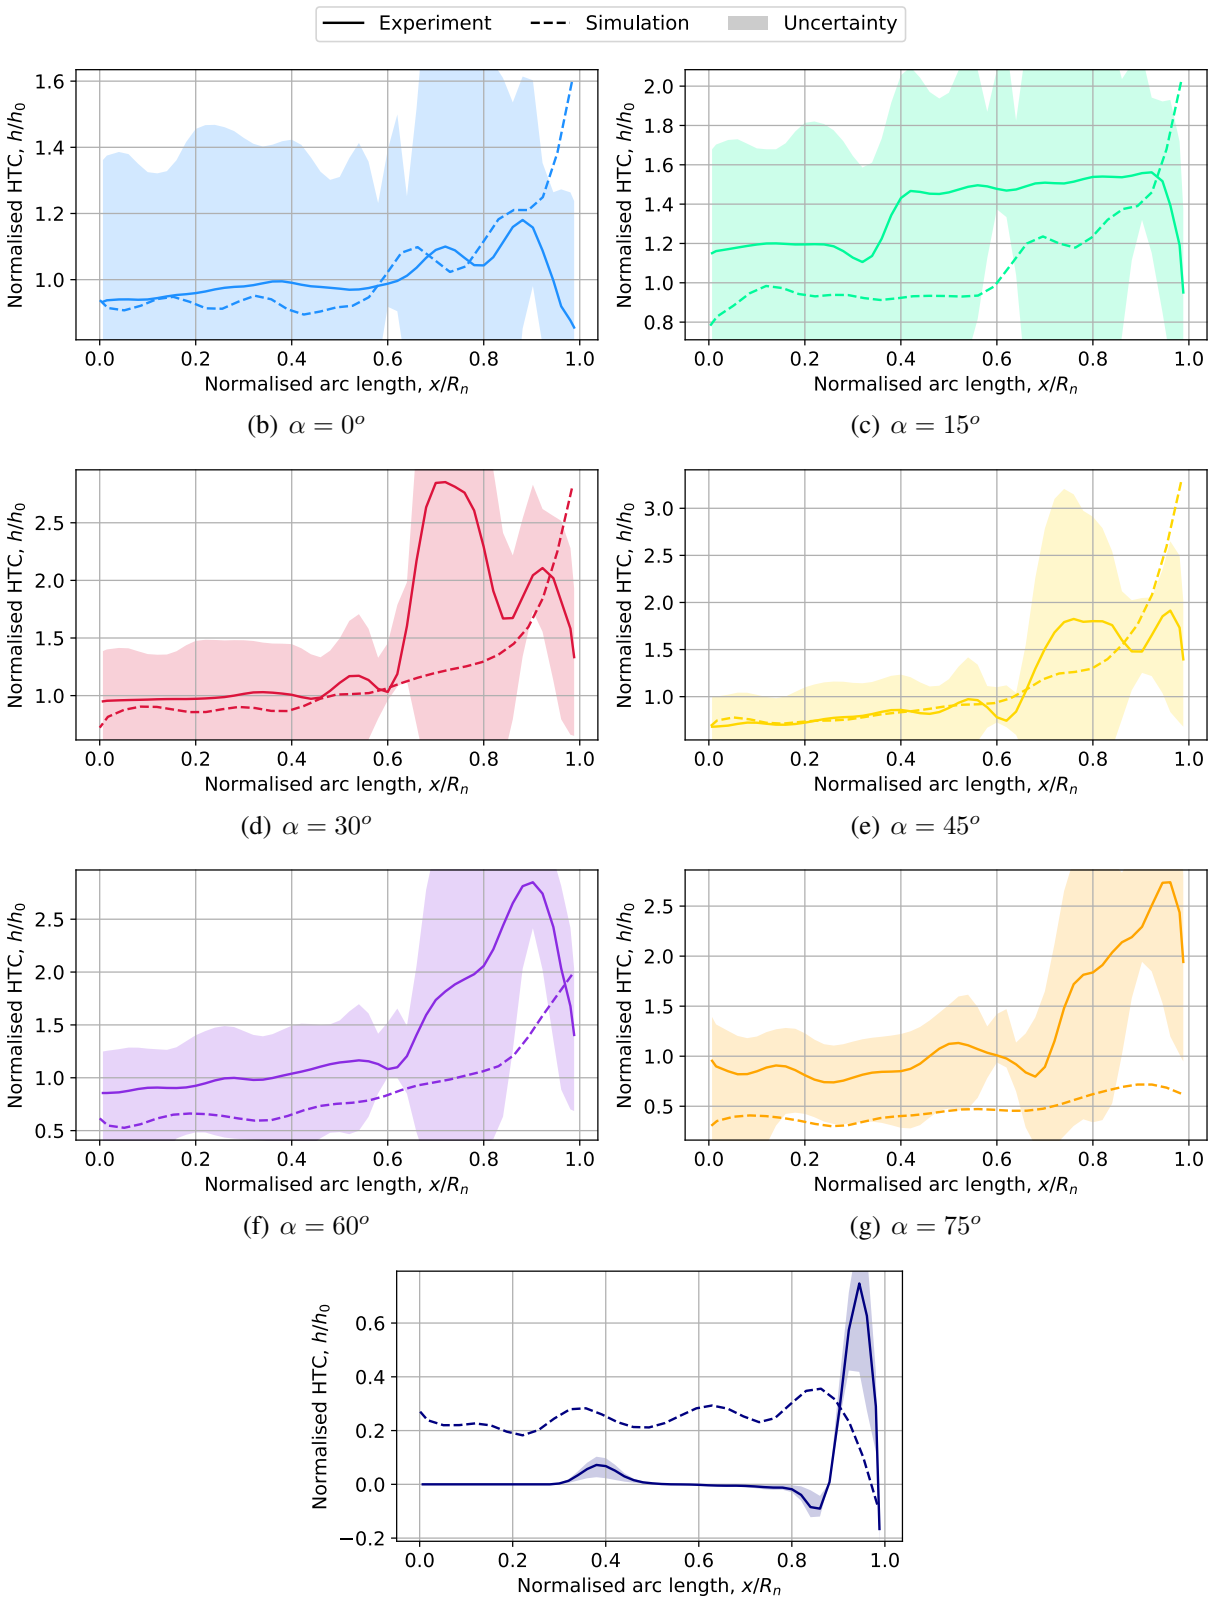
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.8: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the front stagnation line sampling location are plotted.



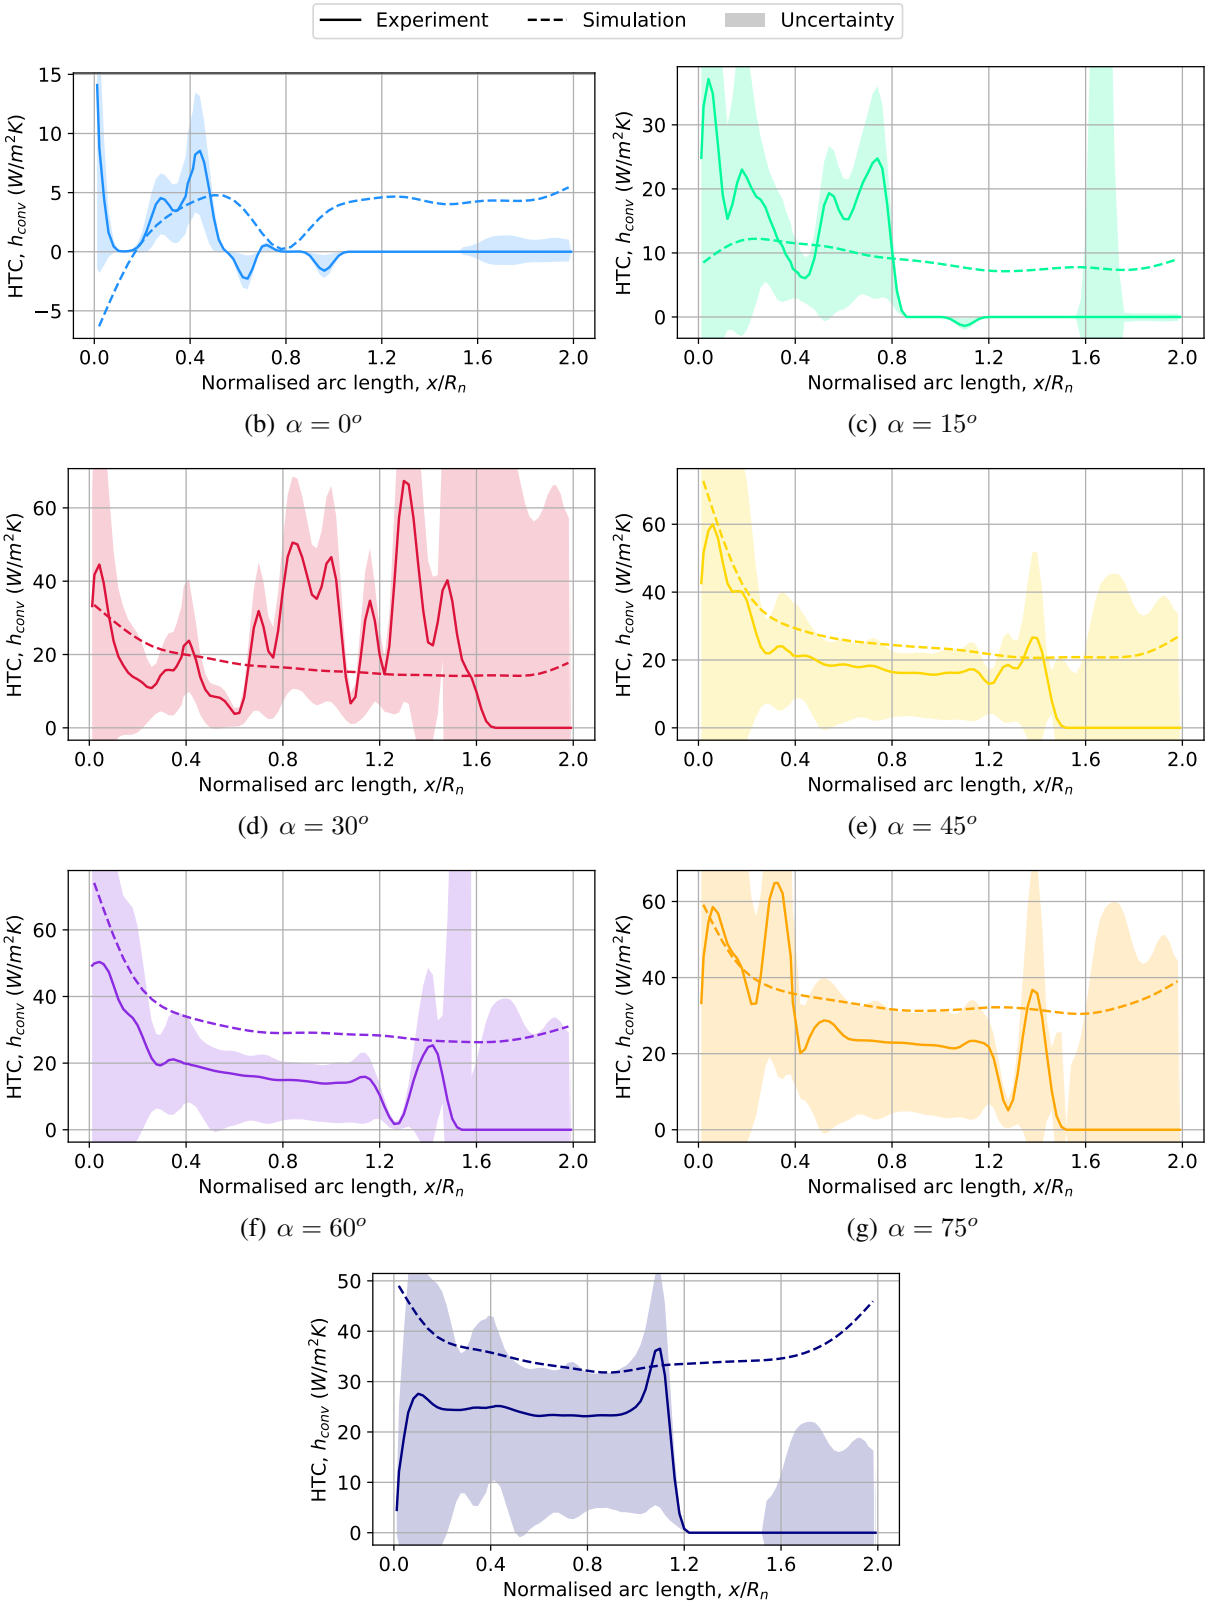
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.9: Normalised results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the front stagnation line sampling location are plotted.



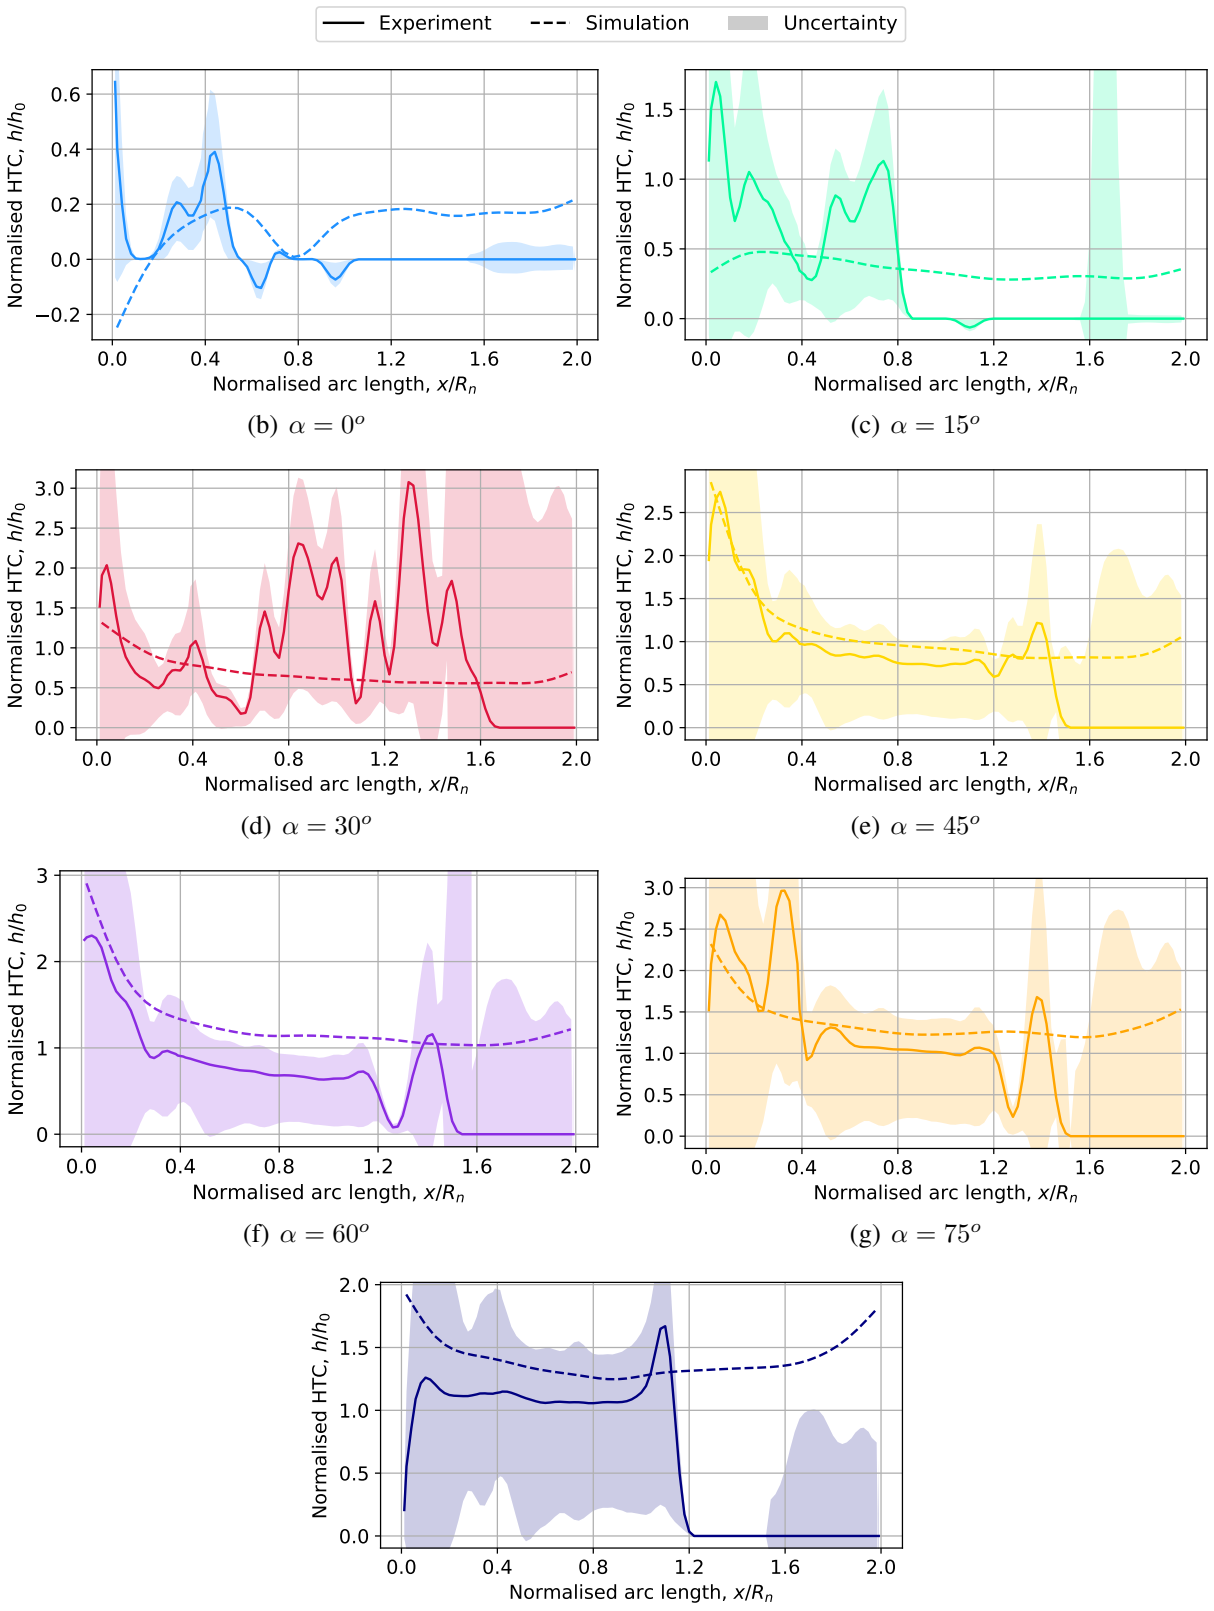
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.10: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the side stagnation line sampling location are plotted.



(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

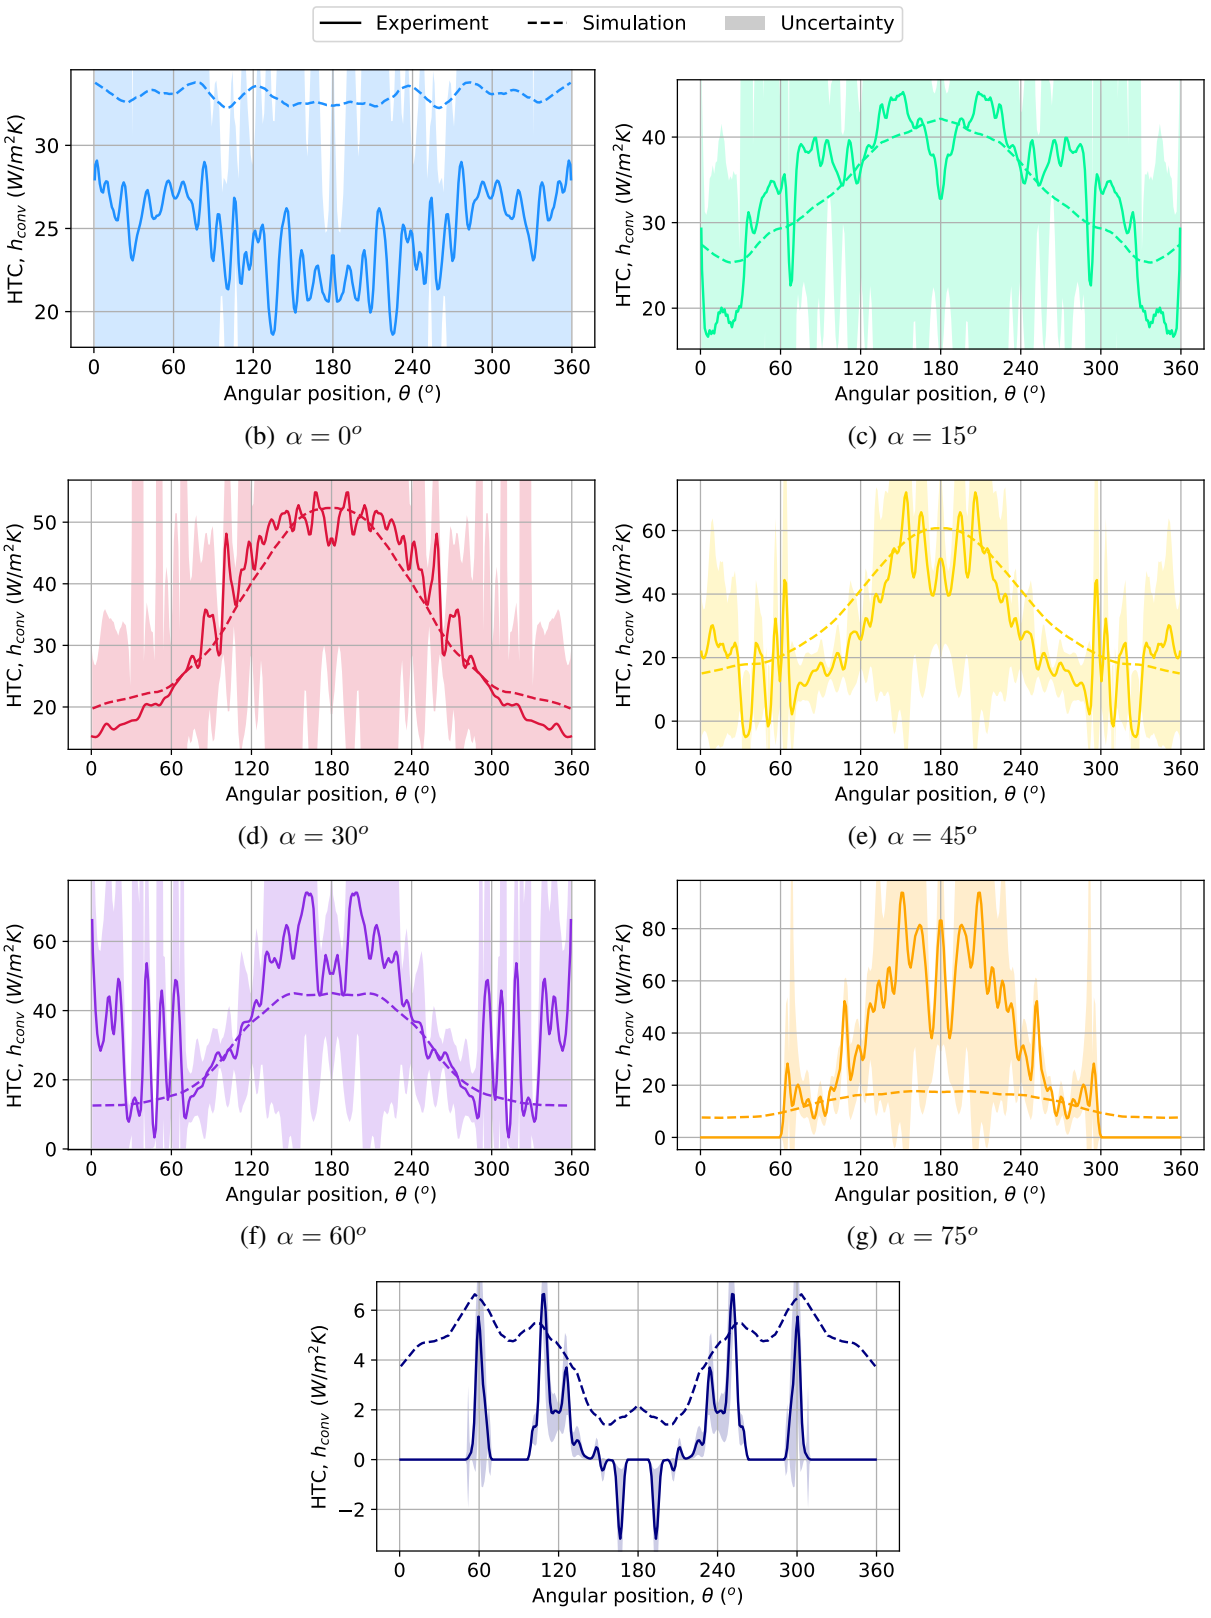
Figure 6.11: Normalised results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the side stagnation line sampling location are plotted.

Data extracted from the front circular edge of the $\varnothing = 15\,mm$ probe are presented in Figure 6.12. Immediately, it is apparent that there is excellent agreement between simulated and experimental values in the $\alpha = 15^o$, $\alpha = 30^o$, $\alpha = 45^o$, $\alpha = 60^o$, and $\alpha = 90^o$ cases, despite significant noise in the latter two of these (which was caused by numerical overestimation errors in the 3D HTC solver - see Section 6.2.1). The $\alpha = 0^o$ and $\alpha = 75^o$ cases both show poorer agreement, particularly in the latter case, whose central peak is the result of poor signal matching by the 3D HTC solver.

Normalised HTC values for this sampling location are presented in Figure 6.13. These data do in fact show slightly poorer agreement than the absolute values discussed previously, particularly in the $\alpha = 15^o$, $\alpha = 30^o$, $\alpha = 45^o$, and $\alpha = 60^o$ cases. Agreement in the $\alpha = 90^o$ case remains consistent, and improves significantly for the $\alpha = 0^o$ case. The experimental results from the $\alpha = 75^o$ continue to exhibit an anomalous rise around $\theta = 180^o$.

Results from the side circular edge are presented in Figure 6.14, and show some of the best agreement thus far for the $\varnothing = 15\,mm$ probe studies. Simulated and experimental values in the $\alpha = 45^o$ and $\alpha = 60^o$ cases show very small errors throughout, while the $\alpha = 15^o$ and $\alpha = 30^o$ cases show similarly excellent agreement in the range $90^o \leqslant \theta \leqslant 270^o$ (signal noise evidently dominates outwith these bounds, except in the $\alpha = 15^o$ case, where a further erroneous peak is visible near $\theta = 180^o$). The $\alpha = 75^o$ and $\alpha = 90^o$ cases exhibit reasonable agreement, with the largest absolute error being present in the range $90^o \leqslant \theta \leqslant 270^o$. Here, the $\alpha = 0^o$ case shows the poor agreement throughout.

The normalised side circular edge results are presented in Figure 6.15, and demonstrate the same excellent agreement for the $\alpha = 45^o$ and $\alpha = 60^o$ cases. The $\alpha = 15^o$ and $\alpha = 30^o$ cases also maintain their low errors in the range $90^o \leqslant \theta \leqslant 270^o$ with both the $\alpha = 0^o$ and $\alpha = 90^o$ cases also showing improvements in their agreement (particularly in the $\alpha = 0^o$ case, whose mean relative error falls to $\sim 10\%$ when normalised). The $\alpha = 75^o$ case shows slightly poorer agreement when normalised compared to its absolute values.



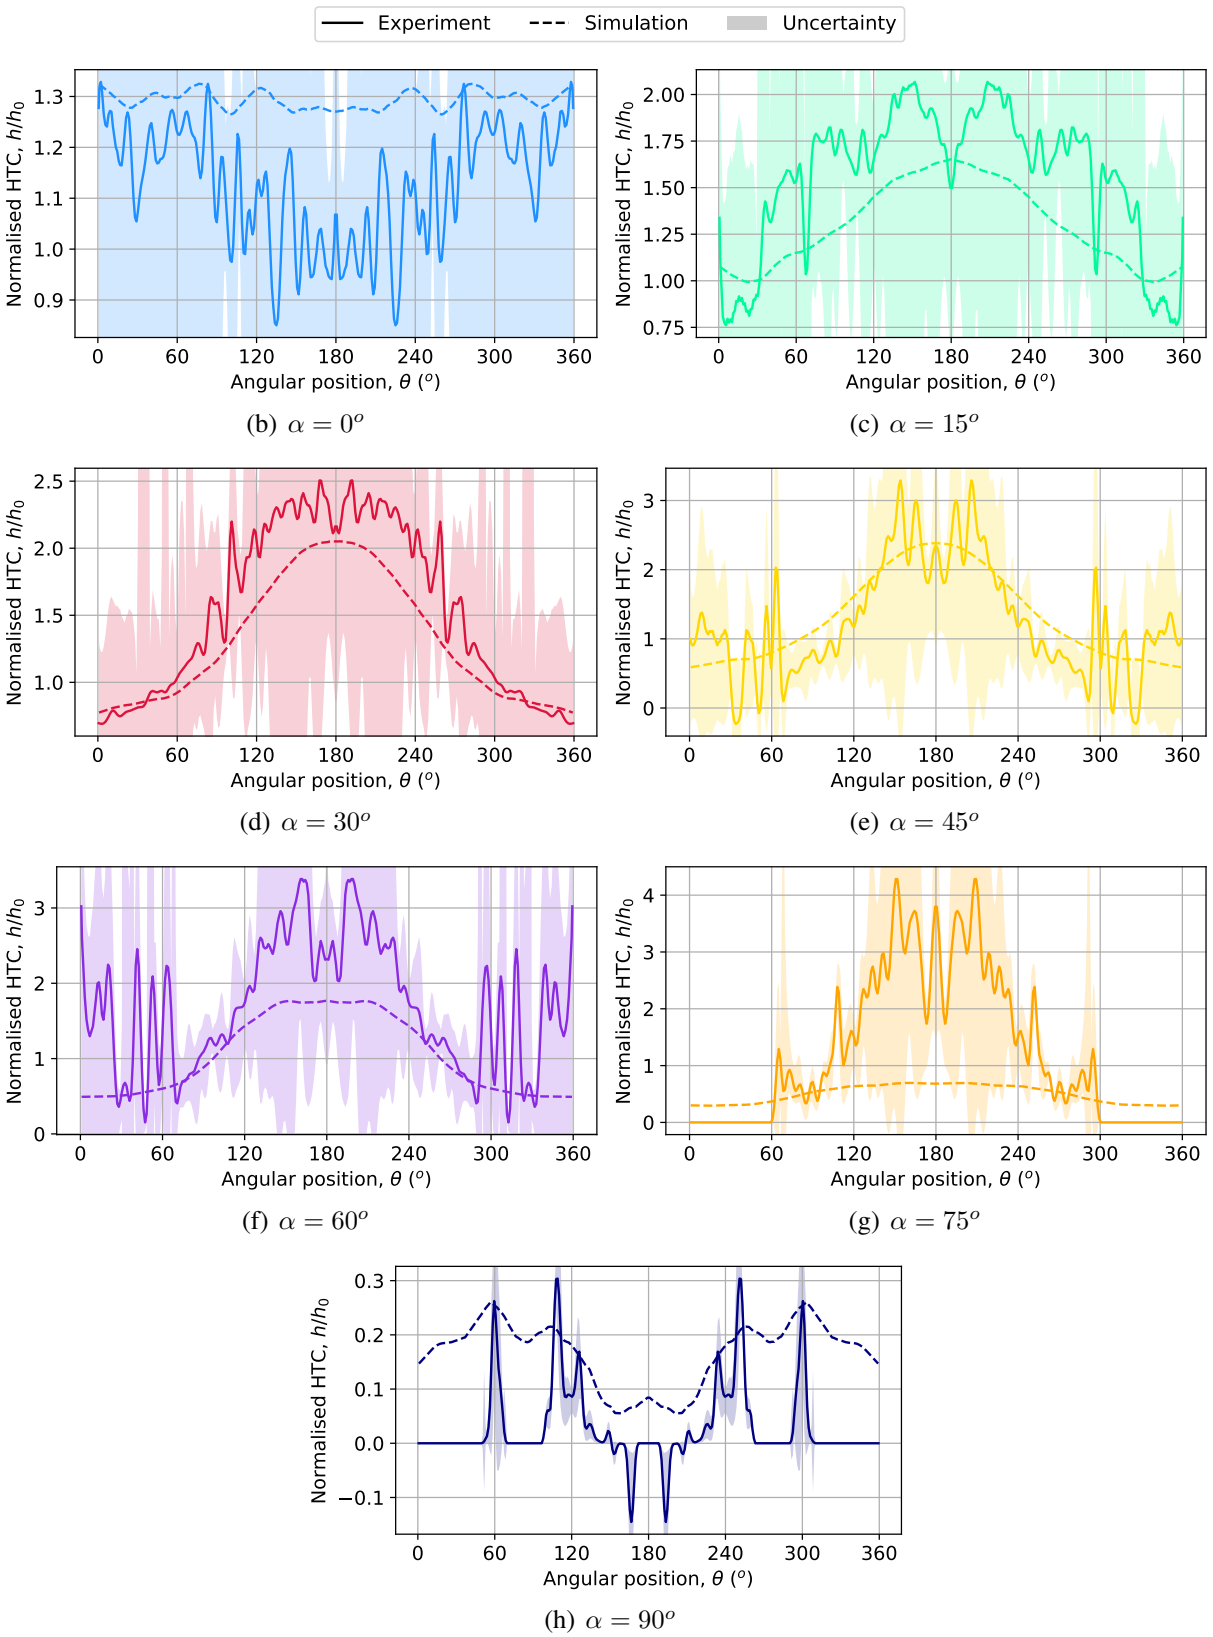
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.12: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the front circular edge sampling location are plotted.



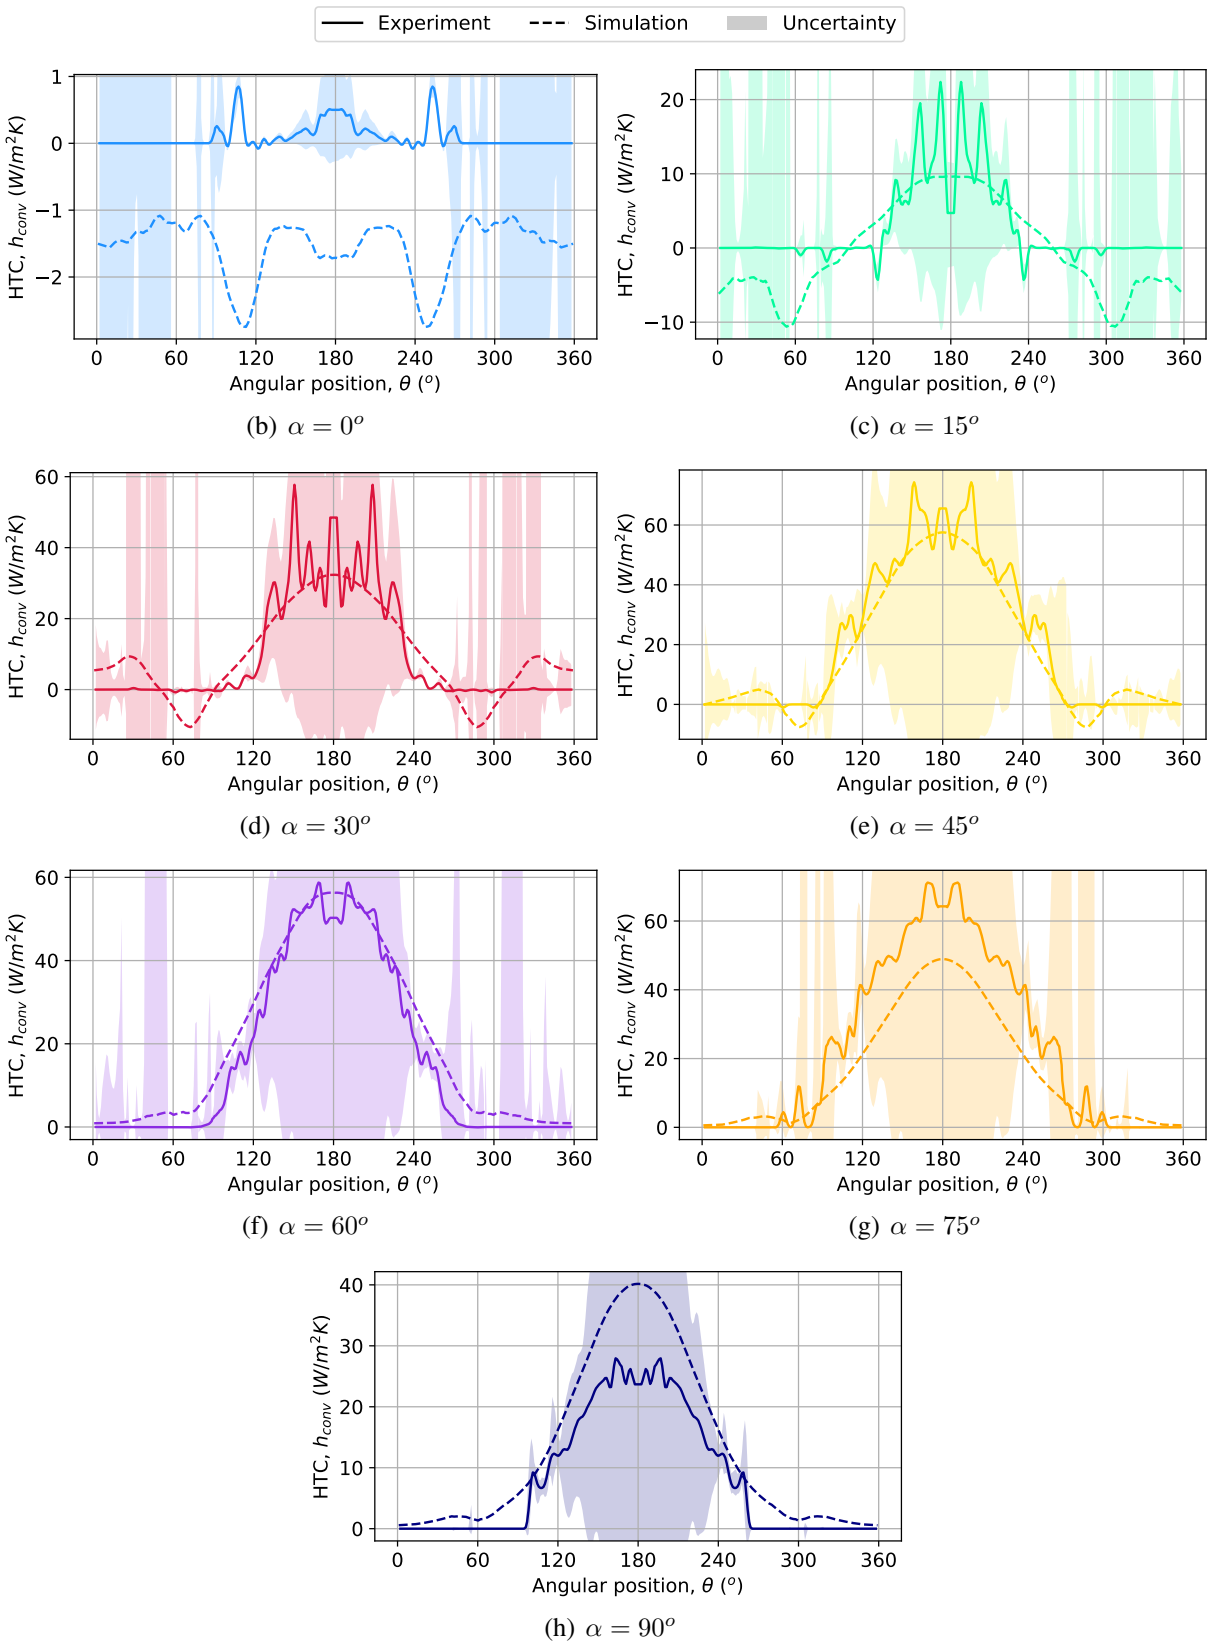
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.13: Normalised results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the front circular edge sampling location are plotted.



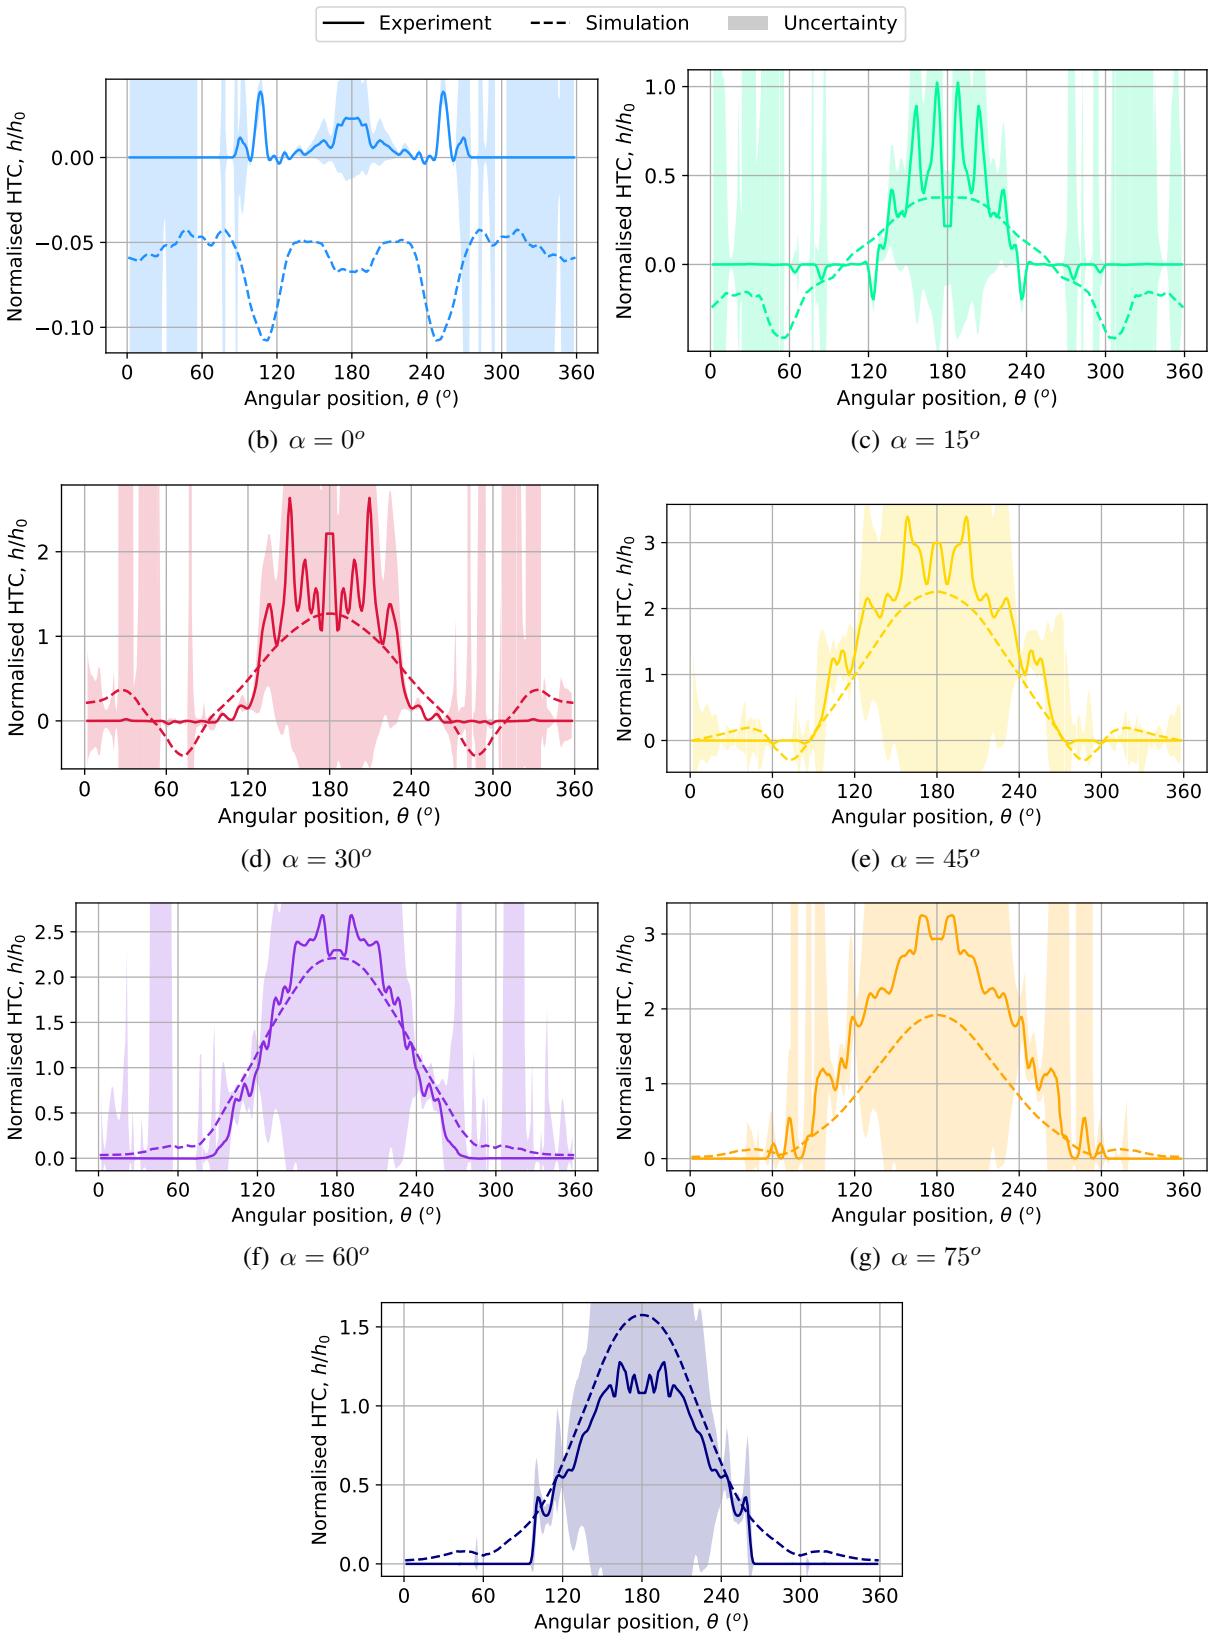
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.14: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the side circular edge sampling location are plotted.



(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

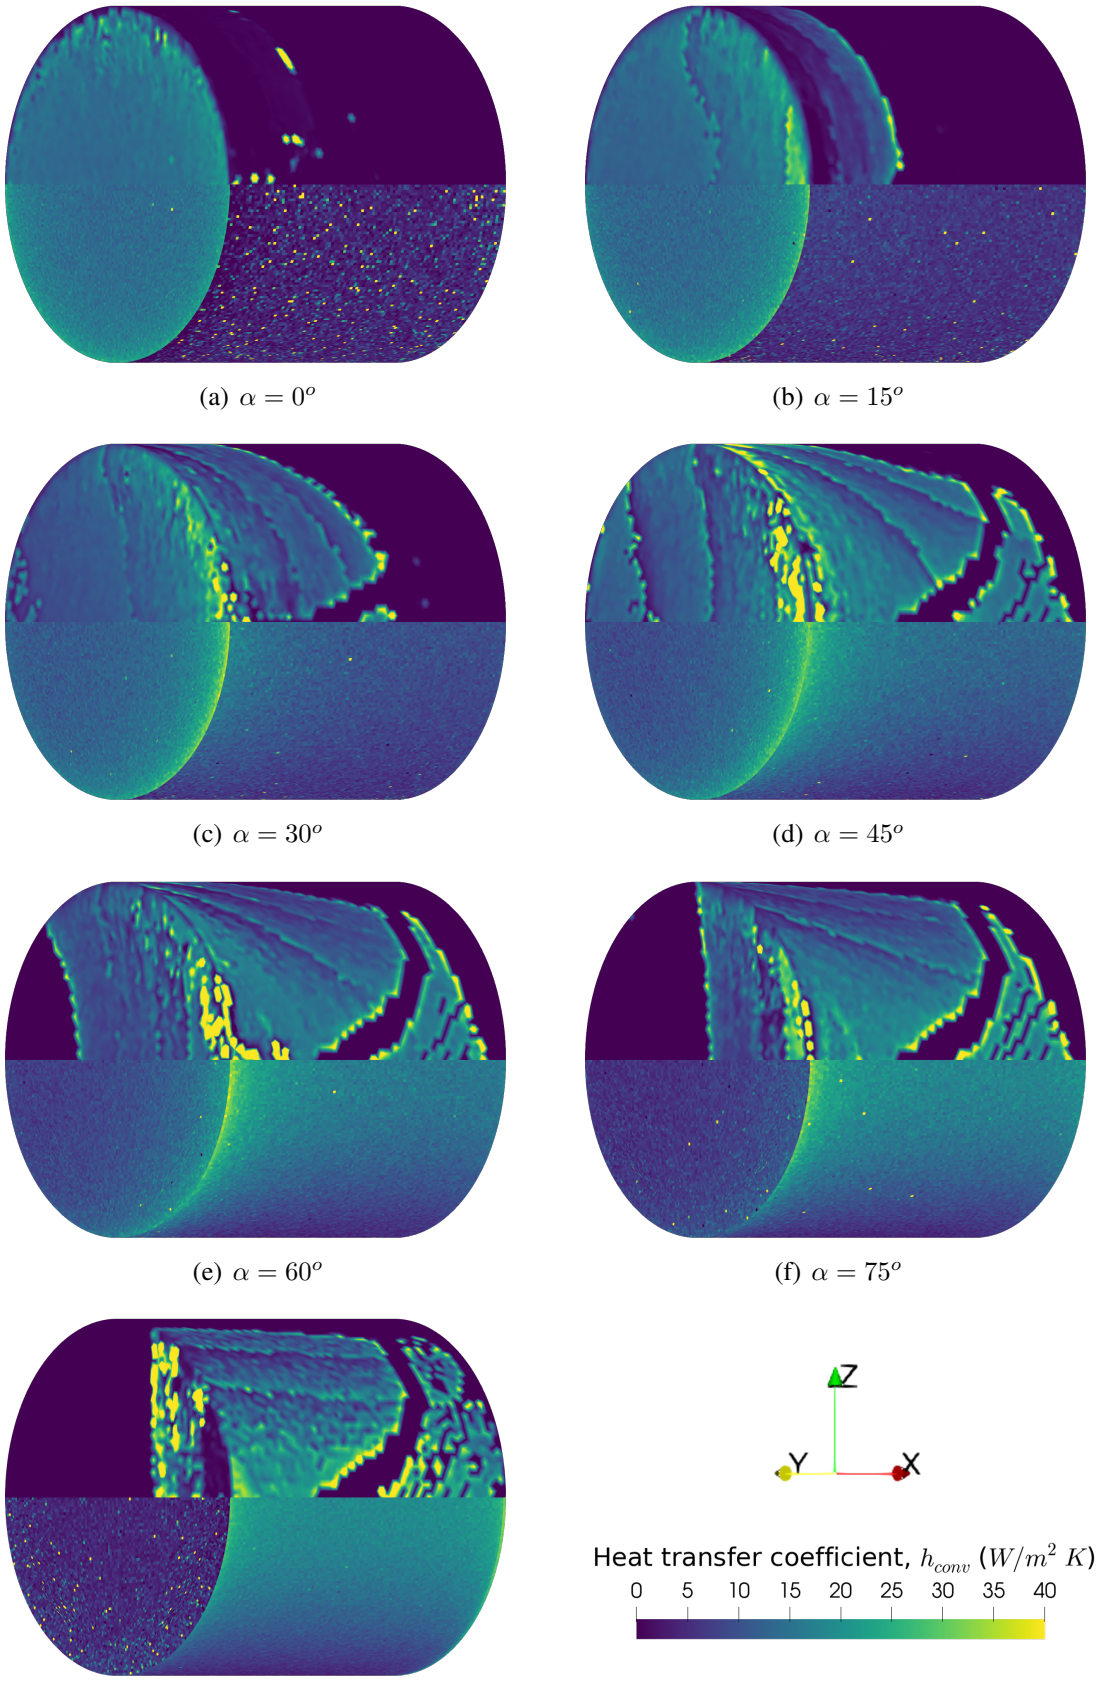
Figure 6.15: Normalised results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 15\,mm$ probe at various angles of incidence. Data from the side circular edge sampling location are plotted.

6.2.3 Flow condition 2 ($\varnothing = 10\,mm,\ Kn = 0.10[12]$)

The second set of tests utilised the $\varnothing = 10\,mm$ cylinder probe at a freestream Knudsen number of $Kn = 0.1012$. The results of these experiments, as well as those generated by DSMC simulations, are presented in Figures 6.16 through 6.20. Unlike the results shown in Section 6.2.2, the cases presented herein are not normalised to local stagnation values at $\alpha = 0^o$. The normalisation technique (see Section 3.4.2) was not utilised here as the *dsmcFoam+* simulations were performed using appropriate DSMC variables for a "good" simulation (as described in Section 2.3.4).

Regarding Figure 6.16, it is clear that, much like the cases run at condition 1 (see Section 6.2.2), overall agreement between simulations and experiments is generally good. Trends of HTC are reproduced well, with the highest values manifesting towards the edge of the cylinders, with these maxima becoming more pronounced as $\alpha \rightarrow 45^o$. Similarly, on the lateral faces, trend agreement is also favourable, with an overall smoother distribution of HTC. This is to be expected from a flow condition with a higher freestream Knudsen number (condition 2 is roughly one order of magnitude closer to the free molecular limit than condition 1). Despite these favourable features, however, results produced by the 3D HTC solver at condition 2 demonstrate more frequent anomalous artefacts (see Section 6.2.1).

Many of the aforementioned artefacts described in Section 6.2.1 manifest in these cases, with the notable exceptions of the $\alpha = 0^o$ case, which only suffers from a lack of colour play data on the side faces of the cylinder. The most prominent errors of any consequence occur close to the edge of the cylinder in cases with pitch angles in the range $30^o \leqslant \alpha \leqslant 90^o$. A lack of colour play data is evident in all cases, resulting in zeroing of the affected regions. Highly pronounced colour play edge effects can be seen in the all cases expect $\alpha = 0^o$. The transition of the 3D HTC solver from its percentile matching routine to peak finding also produced a series of ridges in the HTC distribution. These are most visible on the front circular faces of the $\alpha = 15^o$ and $\alpha = 30^o$ cases.



(a) $\alpha = 0^o$ (b) $\alpha = 15^o$

(c) $\alpha = 30^o$ (d) $\alpha = 45^o$

(e) $\alpha = 60^o$ (f) $\alpha = 75^o$

(g) $\alpha = 90^o$

Figure 6.16: Surface plots of HTC for flow condition 2 ($Kn = 0.1012$) derived from simulated and experimental data. Experimental data is shown on the top of each subfigure, while DSMC results are shown on the bottom.

Like the results from condition 1, agreement is generally excellent between those regions of the experimental and simulated HTC distributions which do not exhibit errors from the 3D HTC solver. The same sampling and error calculation strategy as used in Section 6.2.2 is utilised here to characterise this agreement.

In the case of the $\alpha = 0^o$ case, the error on the front face is 1.4%, but was not calculated on the lateral face due to zeroing by the 3D HTC solver. The $\alpha = \{15^o, 30^o, 45^o, 60^o\}$ cases also show reasonable agreement, with front face errors of 2.8%, 4.3%, 7.7%, and 6.7%, respectively. On the lateral faces of these cases, the errors are 5.3%, 4.7%, 9.0%, 22.5%, respectively. Errors on front and lateral faces of the $\alpha = 75^o$ case are 13.4% and 22.5%, respectively. Finally, the $\alpha = 90^o$ case exhibits an 11.63% error in HTC on the lateral face, and significant overestimation errors (of the order of 90%) on the front circular face.

A prominent processing artefact is visible in cases with pitch angles in the range $30^o \leqslant \alpha \leqslant 90^o$ towards the rear of the 3D HTC results (see item 5 in Figure 6.4). These stripe-like regions of low HTC are the results of zeroing performed by the solver in order to prevent large HTC estimates from dominating the FEA thermal solutions. The error itself is almost certainly non-physical in that it manifests in approximately the same position across several cases, despite significant differences in α. It is thought that these artefacts occur as a result of consistently poor lighting or camera viewing angle during the experiment itself. This conclusion is further supported by recalling that the camera was fixed to the traversing mechanism during experiments, meaning that its viewing angle relative to the cylinder model did not change with α.

The highest concentrations of overestimation errors occur in the $\alpha = 45^o$ and $\alpha = 90^o$ cases, with sparse clusters of nodes near the circular edge showing extremely high HTC values. The $\alpha = 60\%$ and $\alpha = 75\%$ cases once again show the most significant overall errors compared to others in that their highest HTC values occur consistently on the front faces, while in the associated simulation results these maxima occur on the lateral face. This is likely a side effect of poor viewing angle or lighting during the experimental run itself, as the majority of the hypersonic flow in these cases is incident on the lateral face. Hence, the highest HTC values should occur proximate to the circular edge on the lateral face.

The front stagnation line results are presented in Figure 6.17 and show excellent agreement

in all cases except $\alpha = 90^o$, which features a pair of large erroneous peaks at $x/R_n = 0.4$ and $x/R_n = 0.7$. The other cases begin to manifest larger errors as $x/R_n \to 1$, which are the result of poor characterisation of the experimental LC response by the 3D HTC solver. The fact that all of these features appear near the same region suggests that a slight misalignment occurred between the FEA mesh and the experimental video data during processing.

In the $\alpha = 0^o$, $\alpha = 15^o$, $\alpha = 30^o$, and $\alpha = 45^o$ cases, however, the anomalous peaks that manifest around $x/R_n = 0.9$ are mirrored by the DSMC results, suggesting that these may in fact be caused by statistical scatter in the adiabatic DSMC results (which were used as input conditions for the 3D HTC solver in lieu of an estimated bulk value of T_r). It is also likely that the low thermal mass of the FEA model at these locations causes heat to be lost to the surrounding material in downstream portions of the probe as opposed to remaining in the corner of the cylinder.

Results from the side stagnation line are presented in Figure 6.18 and show significant noise compared to the front stagnation line data. Despite this, however, the $\alpha = 45^o$, $\alpha = 60^o$, $\alpha = 75^o$, and $\alpha = 90^o$ cases show excellent agreement in the ranges $0 \leqslant x/R_n \leqslant 0.8$ and $1.2 \leqslant x/R_n \leqslant 2$. The region between these two parts of the probe have been reduced to $0\,W{\cdot}m^{-2}{\cdot}K^{-1}$ by the 3D HTC solver in order to prevent divergence of the FEA solutions. The $\alpha = 0^o$, $\alpha = 15^o$, and $\alpha = 30^o$ cases show similarly good agreement in the range $0 \leqslant x/R_n \leqslant 0.4$, after which the zeroing behaviour of the 3D HTC solver once again becomes apparent (except in the $\alpha = 30^o$ case, which features a large anomalous peak at $x/R_n \approx 1.1$).

The front circular edge results are presented in Figure 6.19, and once again exhibit excellent agreement between simulated and experimental HTC values. The $\alpha = 15^o$, $\alpha = 30^o$, $\alpha = 45^o$, and $\alpha = 60^o$ cases feature the best agreement, with both trends and magnitudes of HTC comparing favourably with one another, despite significant noise in places. The $\alpha = 75^o$ and $\alpha = 90^o$ cases both show small regions of agreement with DSMC results (at $90^o \leqslant \theta \leqslant 130^o$ and $230^o \leqslant \theta \leqslant 170^o$ for the former, and $150^o \leqslant \theta \leqslant 210^o$ for the latter).

The $\alpha = 0^o$ case features an error unlike any encountered at this sampling location in the $\varnothing = 15\,mm$ probe results. Specifically, this is a large peak that manifests in the range $120^o \leqslant \theta \leqslant 240^o$, which suggests a minor angular misalignment of the probe relative to the

nozzle centreline (as "hot spots" centred around $\theta = 180^o$ tend only to manifest when $\alpha > 0^o$). As such, it is likely that the other cases (which already show excellent agreement) would experience a further improvement in accuracy were this angular misalignment not present.

Results from the side curved edge of the $\varnothing = 10\,mm$ probe are presented in Figure 6.20 and, like the front edge data, show excellent agreement between simulated and experimental HTC values. The $\alpha = 0^o$, $\alpha = 30^o$, and $\alpha = 60^o$ cases exhibit very low errors in both the HTC trends and magnitudes, once again despite significant noise in places. The $\alpha = 45^o$ and $\alpha = 90^o$ cases also show reasonably good agreement, but are disrupted by comparatively large levels of noise in the experimental HTC distributions (which, like other errors of this kind, are caused by poor characterisation of the experimental LC activity signal).

Finally, the $\alpha = 15^o$ and $\alpha = 75^o$ cases exhibit the largest errors at this sampling location, and both feature significant dips in experimental HTC in the range $90^o \leqslant \theta \leqslant 270^o$. The reasons for these errors are not identical though; the $\alpha = 15^o$ case has clearly been zeroed by the 3D HTC solver in order to prevent large values of HTC from skewing the FEA results, while the $\alpha = 75^o$ case suffers from poor characterisation of the experimental LC activity signals.

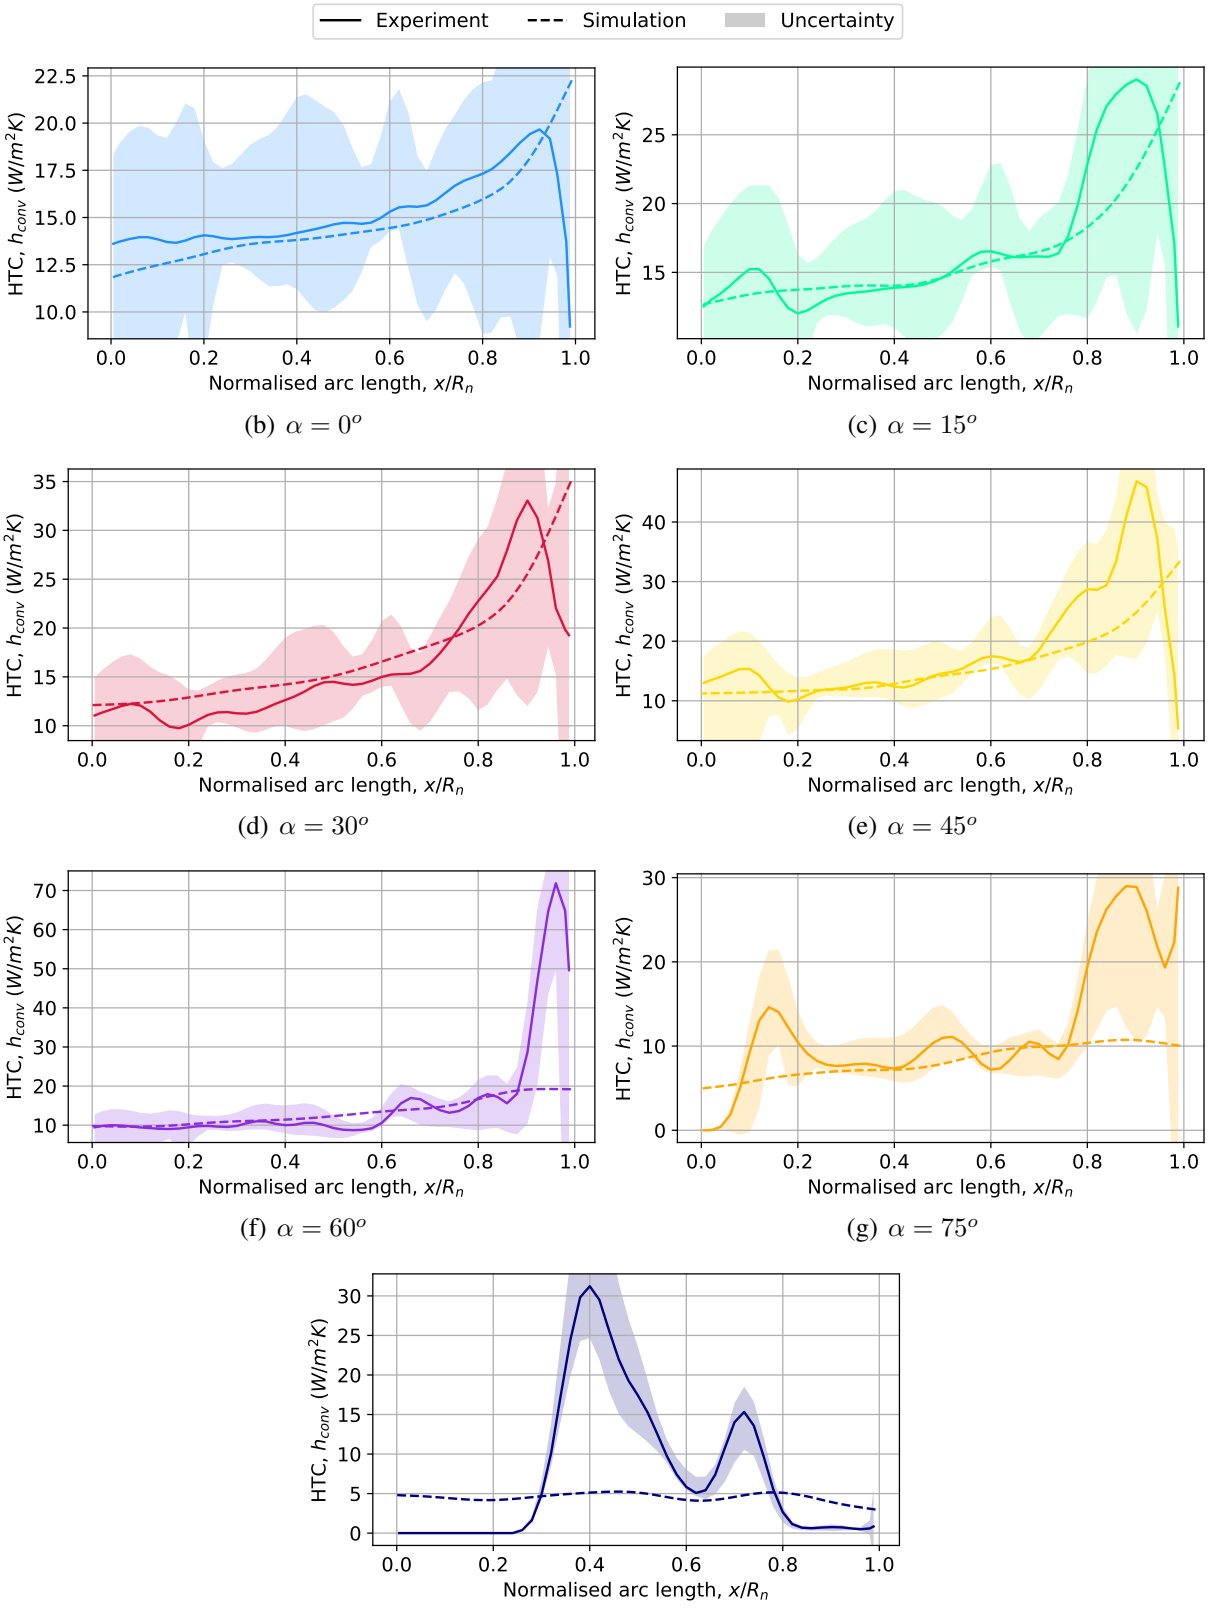
6.2.4 Comparison of both flow conditions

Experimental results from both sets of flow conditions behaved as expected relative to one another.

Heat flux magnitudes were lower for flow condition 2 ($Kn = 0.1$) than condition 1 ($Kn = 0.02$); a result of the similar freestream velocities between conditions despite there being a difference in Knudsen number of approximately one order of magnitude (see Table 4.2). As heat flux magnitude is a function of both freestream velocity and density, and density is itself related to the mean free path length (and hence Knudsen number), lower fluxes were expected – and subsequently observed – for condition 2.

In terms of heat flux *distributions*, both conditions also behaved as expected. Referring to Figures 6.7 and 6.16, it is clear that the the heat flux trend towards the corner of the cylinder geometry is flatter for condition 2 than condition 1. This was expected owing to the higher level of rarefaction in the flow, whose relationship with heat flux distribution is described in Section

2.2.4.



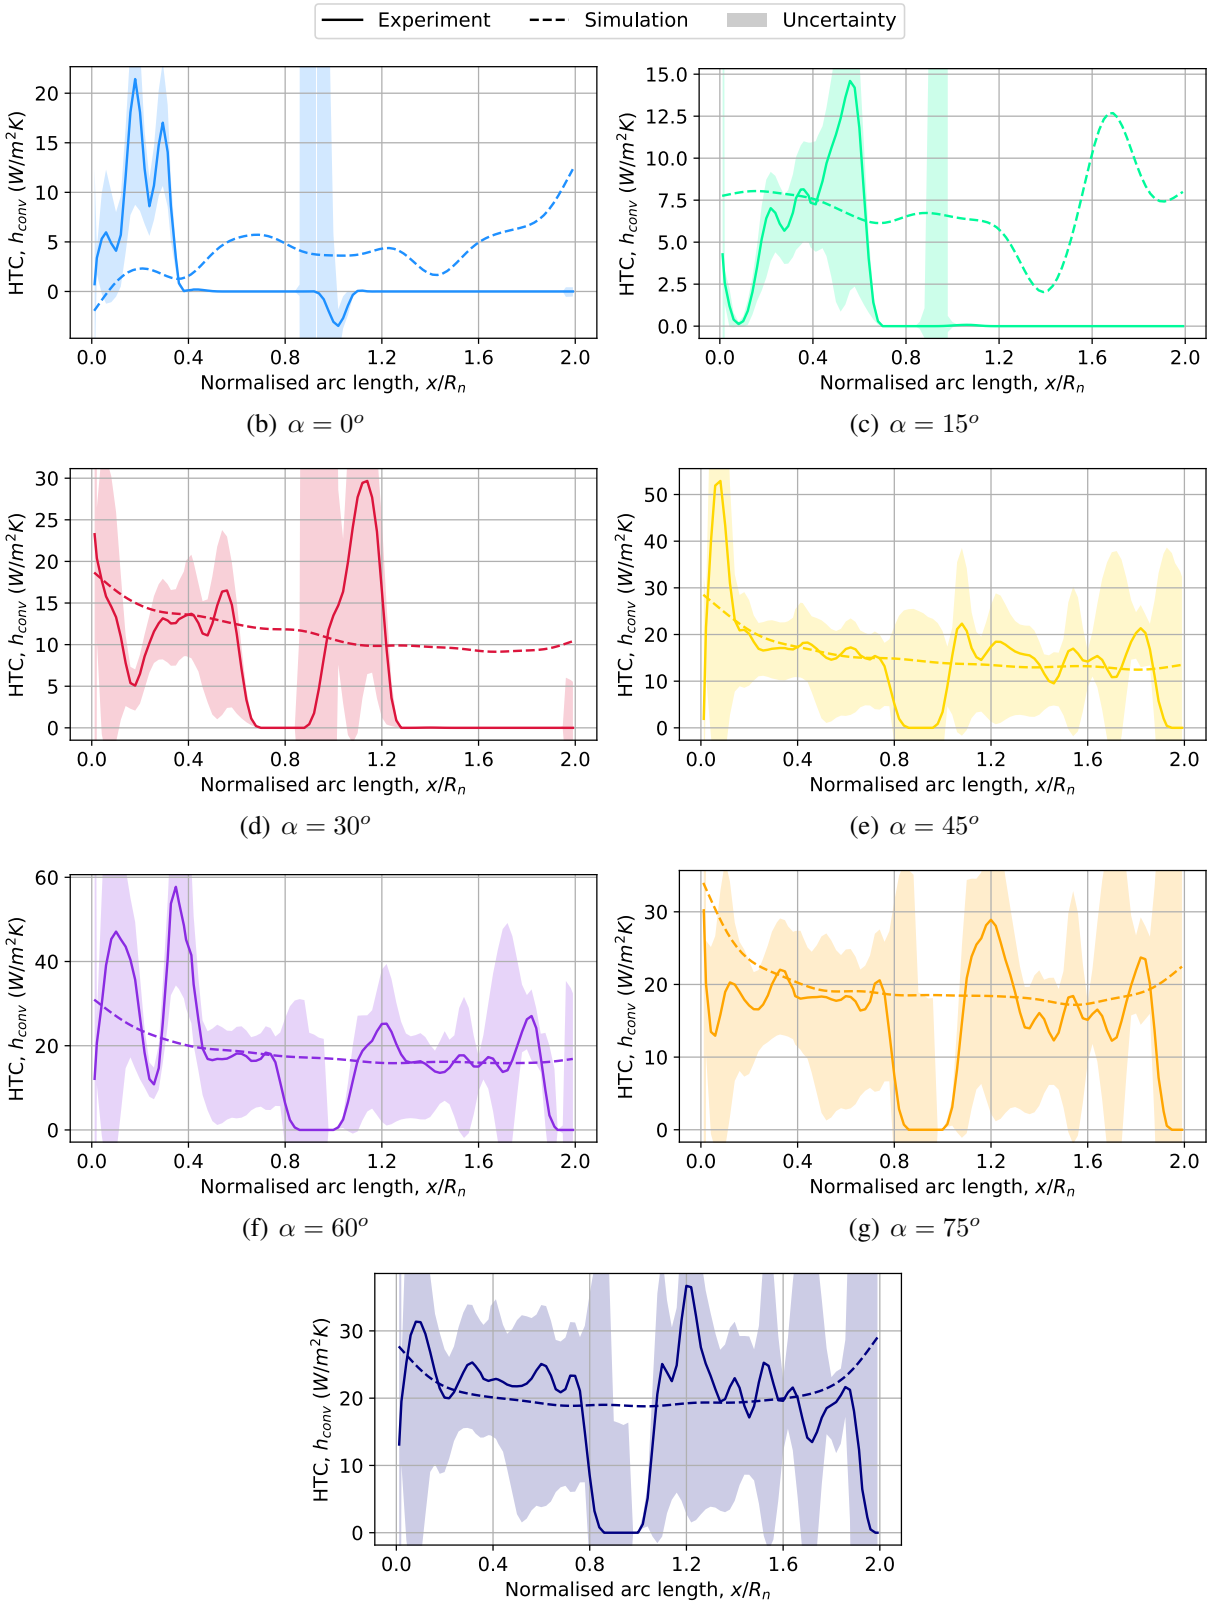
(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.17: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 10\,mm$ probe at various angles of incidence. Data from the front stagnation line sampling location are plotted.



(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.18: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 10\,mm$ probe at various angles of incidence. Data from the side stagnation line sampling location are plotted.

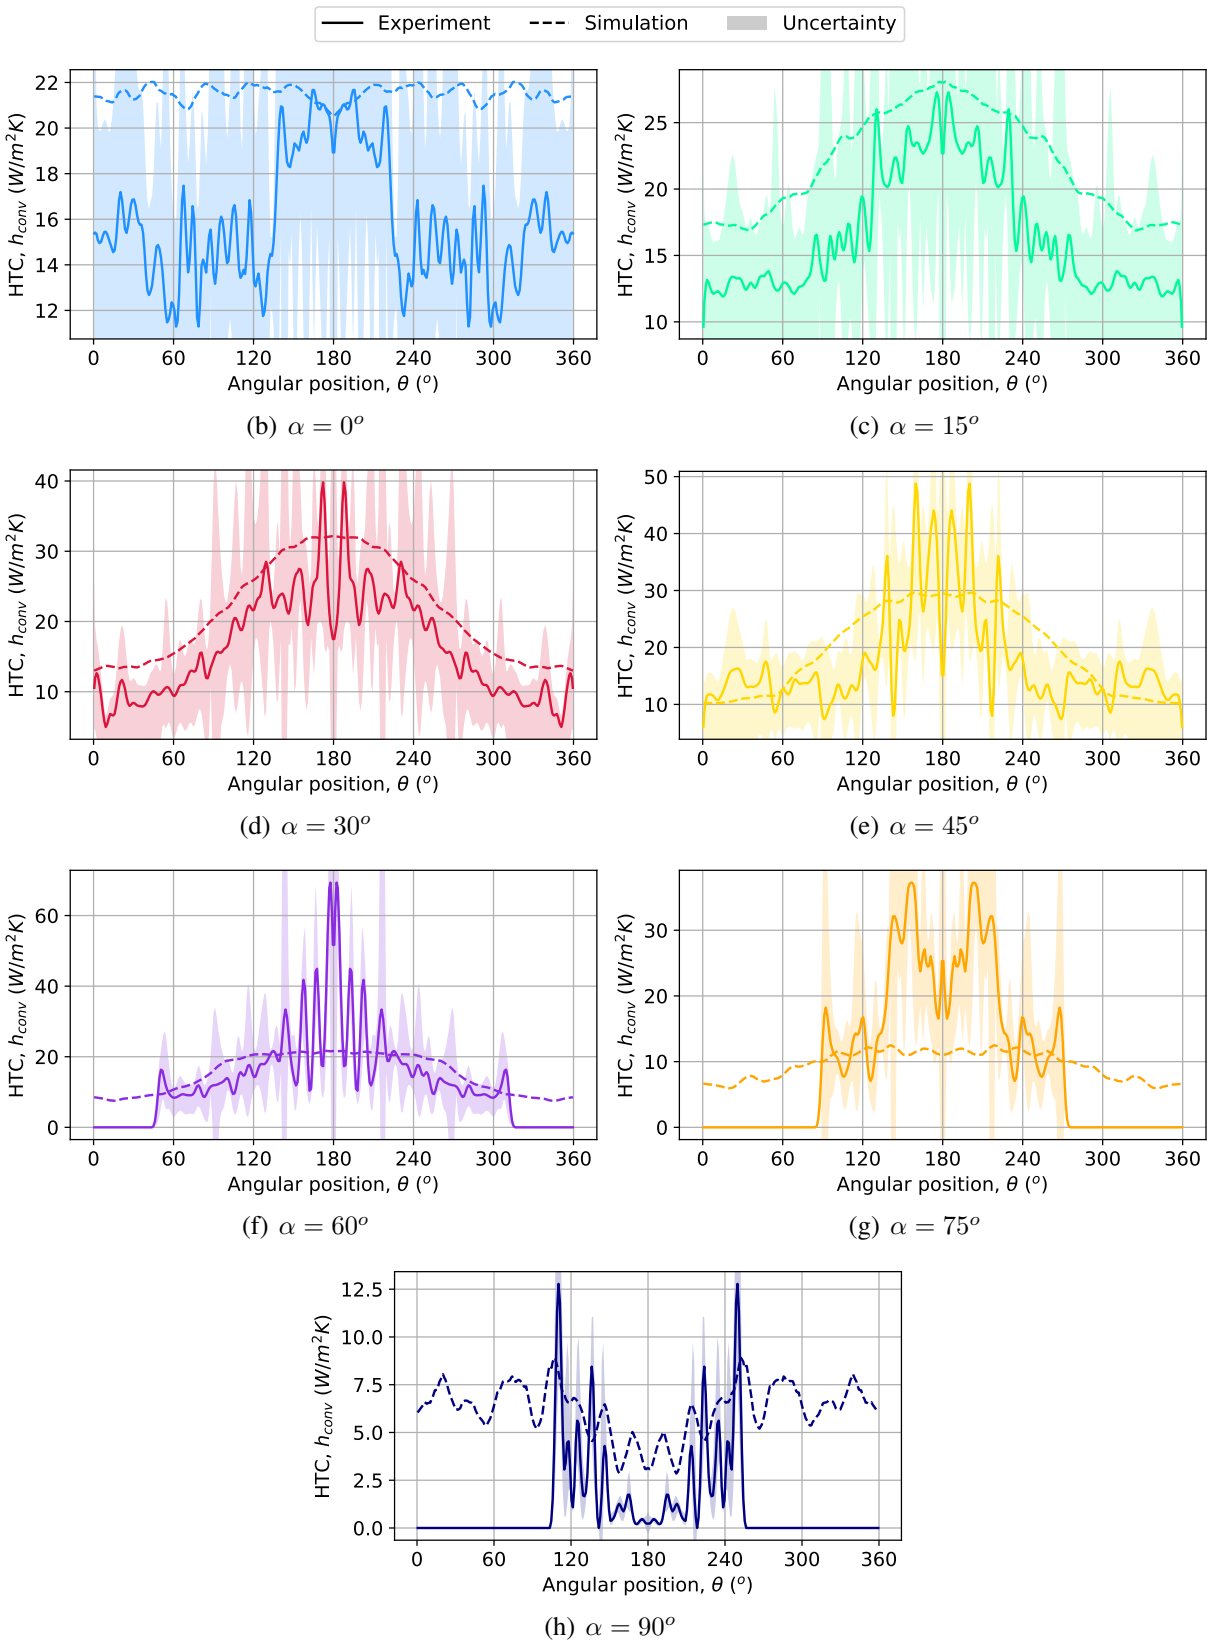
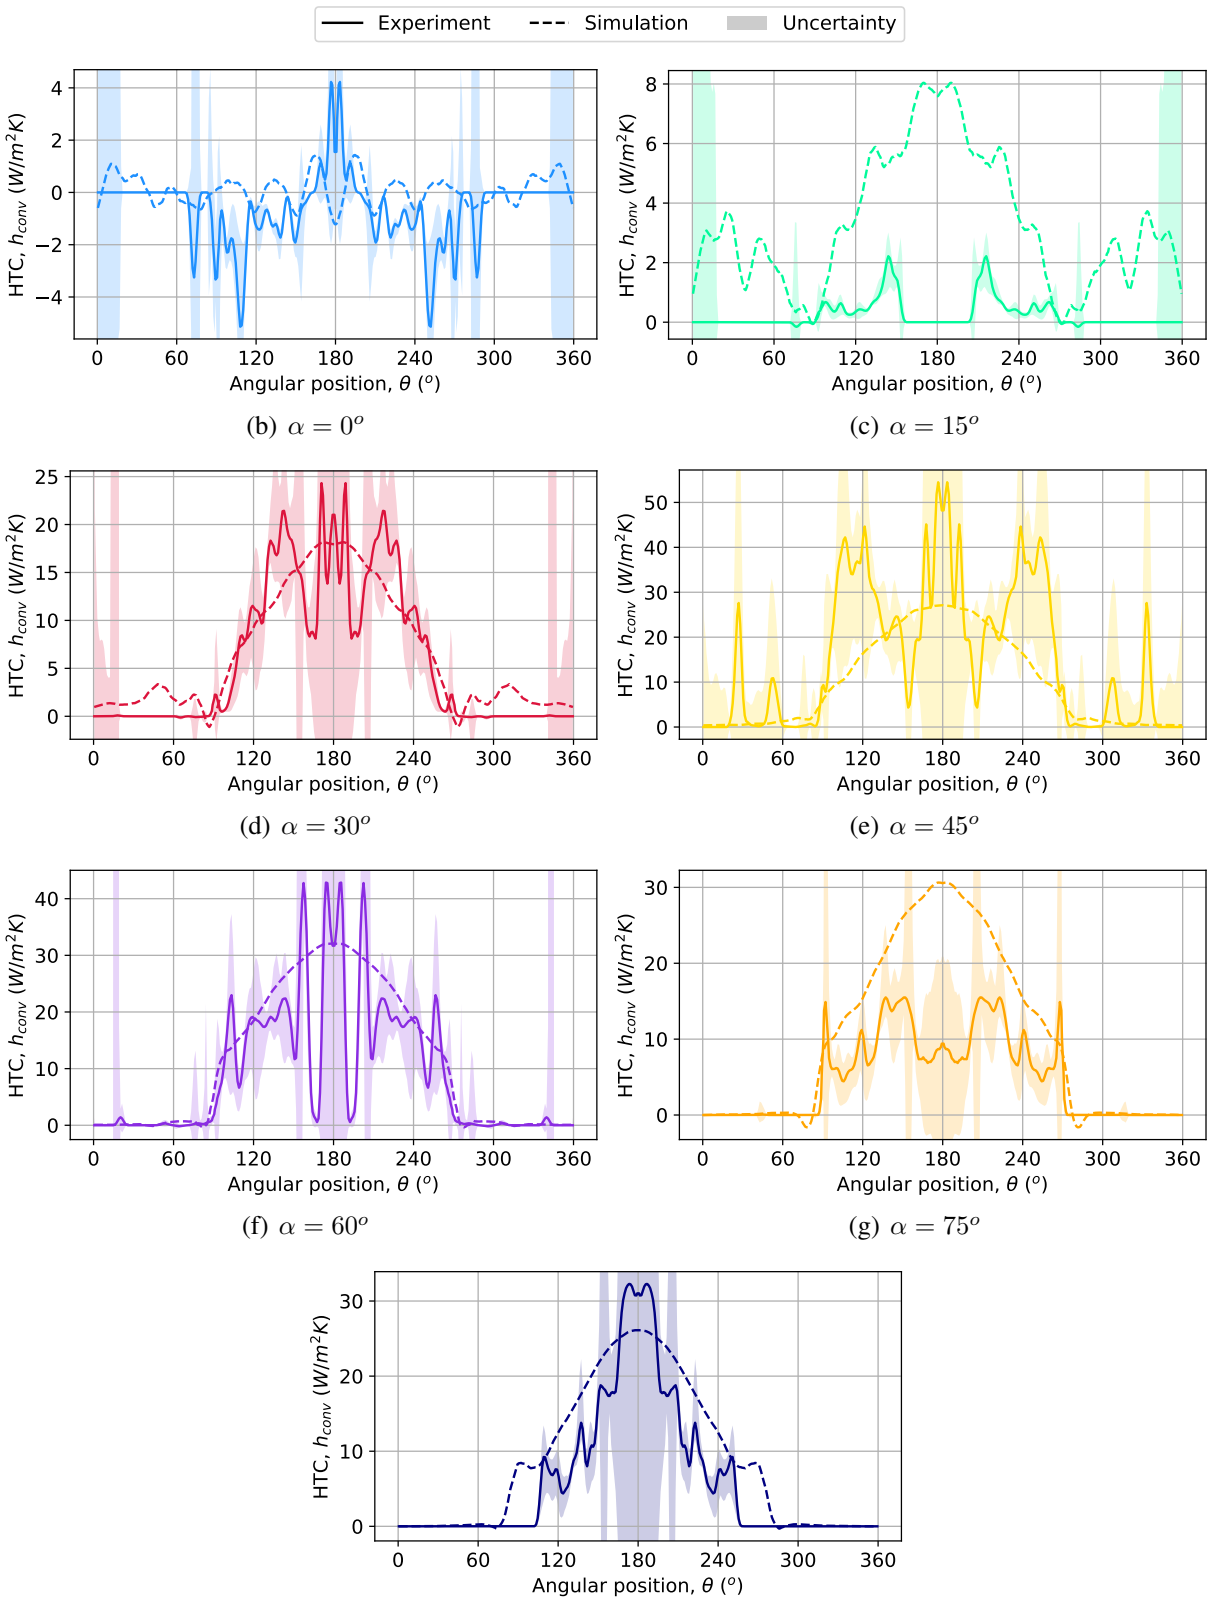


Figure 6.19: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 10\,mm$ probe at various angles of incidence. Data from the front circular edge sampling location are plotted.



(b) $\alpha = 0^o$ (c) $\alpha = 15^o$

(d) $\alpha = 30^o$ (e) $\alpha = 45^o$

(f) $\alpha = 60^o$ (g) $\alpha = 75^o$

(h) $\alpha = 90^o$

Figure 6.20: Results from experimental studies processed using the 3D HTC solver and *dsmcFoam+* simulations for the $\varnothing = 10\,mm$ probe at various angles of incidence. Data from the side circular edge sampling location are plotted.

6.3 CHAPTER SUMMARY

In this chapter, hypersonic experiments performed in the LDWT are described, and associated results discussed. Experiments were performed on a pair of shapes used as satellite debris analogues: a hemisphere, and a cylinder.

As these probes are heated by the hypersonic plume in the LDWT, the colour play which manifests on their exterior surfaces is recorded by a camera, and the resulting video analysed in order to derive HTC distributions. This is performed using two separate methods, depending on the probe used in a given experiment.

The first of these is THTAC, a MATLAB code developed in-house at the Osney lab. This software uses the semi-infinite substrate assumption to calculate a series of possible liquid crystal colour plays which may manifest on a test piece's surface. These are then compared to the actual colour play recorded during an experiment, and the HTC values which produce the lowest error between experimental and simulated activities is selected as the correct value for a given video pixel. THTAC also allows the automatic perturbation of setup variables such that uncertainties may be calculated for an experimental run.

THTAC is used to analyse the two hemisphere cases which were run in the LDWT (both at $\alpha = 0^o$; one at each of the two different flow conditions selected for experiments). The resulting HTC distributions are extremely smooth, and have little in the way of errors. These data are then compared to 2D axisymmetric DSMC simulations run using the same geometry and freestream conditions, and demonstrate excellent agreement.

The second method of processing a liquid crystal experiment is reserved for use in cases where the semi-infinite substrate principle may not be applied. This method is required to process those experiments which utilise the cylindrical probes, as their corners are heated from both sides during experiments and have a low thermal mass compared to the rest of the probe, hence presenting a comparatively complex three dimensional heat transfer problem.

A new experimental processing code, simply called the "3D HTC solver" (or "3D-HTCS" for short) has been developed for this purpose, and is based upon previous work by Ryley [4]. This code is written in MATLAB, and makes use of the MATLAB-COMSOL LiveLink interface

to allow scripted FEA multiphysics simulations to be run automatically. The method that this program utilises to process experimental videos is explained in depth in Chapter 5, and so will only be summarised here.

The 3D HTC solver operates by first spatially aligning an FEA mesh to the video of an LC experiment. Initial HTC and adiabatic wall temperature distributions are input as a "first guess" for the code to operate on, and a thermal FEA simulation is run. At the conclusion of this calculation, a simulated experimental colour play is generated by the 3D HTC solver, and this is then compared to the experimental colour play. A combination of peak finding and signal activity calculation (upper percentile characterisation) is used to determine the error between the simulated and experimental colour plays, and adjustments to the HTC profile made accordingly. A new FEA simulation is then run with the updated HTC profile is then run, and the process repeated until the simulated and experimental colour plays are matched in time. Hence, the true HTC distribution on the surface of a probe during an experiment may be determined.

This code is applied to a series of experiments using the cylindrical probes (which total 14 runs: 7 angles of incidence at 2 flow conditions). Agreement between these data and the associated 3D DSMC simulations is generally very good, but is hampered in many of the sampled HTC distributions by significant noise. These discrepancies come about as a result of errors in the characterisation of difference between simulated and experimental colour plays by the 3D HTC solver. In cases where this occurs, the code is programmed to set a node's HTC to $0\,W{\cdot}m^{-2}{\cdot}K^{-1}$ such that it may not introduce divergent behaviour in the FEA solution, thereby preventing the results from being skewed by extreme outliers.

In the case of experiments run using the $\varnothing = 10\,mm$ probe, HTC values produced by DSMC simulations and the 3D HTC solver are compared directly. For experiments run using the $\varnothing = 15\,mm$ probe, however, both absolute and normalised values of HTC are compared. This is due to the setup of the DSMC, to which these data are compared. The $\varnothing = 15\,mm$ cases were run using non-optimal DSMC variables in order to reduce computational expense, such that normalisation of the resulting HTC distributions would have to be used to derive useful data from the results.

Ultimately, the agreement between experimental and simulated HTC values for these cases

was good, and was only slightly improved by the normalisation procedure (wherein HTC distributions are divided by the local stagnation HTC, h_0, at $\alpha = 0^o$). Those data sampled at the side stagnation line showed the largest errors when compared to simulated values, with the other sampling locations showing excellent agreement overall (even without normalisation). Although this is advantageous to the present study, it is important to consider that the normalisation procedure also produced excellent agreement in many cases, hence proving that this strategy (as described in Section 3.4.2) is suitable for the study of hypersonic heat transfer in the slip and lower transition regimes.

Those experiments which utilised the $\varnothing = 10\,mm$ probe were also shown to agree extremely well with DSMC simulations performed at the same freestream conditions, in terms of both HTC trends and magnitudes (note that these calculations utilise optimal DSMC simulation variables, unlike the $\varnothing = 15\,mm$ cases). Like the $\varnothing = 15\,mm$ cases, it may be concluded from the data presented herein that the DSMC simulations performed for the present work are suitable for the study of hypersonic heat transfer in the slip and lower transition regimes (notably without the use of a large scale cluster computer).

The distributions of HTC over both the $\varnothing = 15\,mm$ and $\varnothing = 10\,mm$ probes agreed well with DSMC results, but are hampered in several locations (including key areas of interest near the edges) by processing errors produced by the 3D HTC solver. These errors are identified, and their causes characterised, hence allowing them to be ignored when considering agreement between experimental and simulation results. Nodes-cell pairs (in the 3D HTC solver and DSMC results, respectively) were selected randomly along the stagnation line and compared in order to characterise errors in regions where the 3D HTC solver had operated correctly. These errors were found to be similar for both sets of experiments, with mean errors of 8.54% and 5.52% on the front and lateral faces of the $\varnothing = 15\,mm$ probe, receptively. For condition 2, which was run using the $\varnothing = 10\,mm$ probe, mean errors on the front and lateral faces were found to be 6.07% and 10.37%, respectively.

An important finding which supports the suitability of both the DSMC codes used and the new 3D HTC solver is the fact that no single experimental case shows consistently poor results. For example, the results from the $\varnothing = 15\,mm$ case run at $\alpha = 0^o$ demonstrate poor agreement

between normalised simulated and experimental values at the side stagnation line and side circular edge. The front stagnation line, however, shows excellent agreement, and the side curved edge demonstrates reasonably good agreement, with a maximum error of $\sim 20\%$. This suggests that the vast majority of the discrepancies between HTC values derived from experiments and simulations are the result of experimental or numerical error. As such, it may ultimately be concluded that both of the DSMC codes used for the present study, the normalisation procedure developed for non-optimal DSMC simulations, the experimental facility and method, and the 3D HTC solver code are all suitable for the study of hypersonic heat transfer in the slip and lower transition regimes.

In terms of the demisability of the examined debris geometries, it is apparent that edge heating effects manifest in a similar fashion in the transition regime as predicted in the continuum regime by the work of Klett [51]. This heating is expected to contribute significantly to the demisability of cylindrical debris items, as their sharp geometric features (i.e. corners) are more likely to melt first than any other part of the examined geometries. As such, during the course of an atmospheric re-entry trajectory, these features demising first would allow heated hypersonic flow to enter the object itself, exposing its internal structure to significant thermal fluxes, potentially hastening the demise of the object. Similarly, if this heating were to cause the exterior of a cylindrical enclosure to demise earlier than expected, its contents (e.g. electrical components, structural elements, etc.) could potentially be scattered at a much higher altitude than would be predicted by methods which do not take edge heating into account, increasing the casualty footprint of the object.

Finally, since both the data presented herein and the methods used to generate them have now been validated, the space debris demise community can utilise them to improve the treatment of object-oriented space debris re-entry simulations. Hence, uncertainties in the calculations pertaining to the destructive re-entry of dangerous space debris may be reduced.

SUMMARY AND CONCLUSIONS 7

In this thesis, a study of the aerothermodynamics of geometric space debris analogues under hypersonic flow conditions has been performed. This study comprises a review of the current state of the art in rarefied gas dynamics experimentation and simulation, a summary of current space debris demise estimation tools and methods, the refurbishment and characterisation of an existing experimental facility, the development of new low-order models for re-entry aerothermodynamics, a detailed numerical study of heat transfer to selected axisymmetric geometric primitives during atmospheric entry, and an experimental campaign wherein both new and existing tools are validated.

A summary of the work described previously herein is presented in this chapter, along with the major conclusions drawn from the preceding chapters.

7.1 EXPERIMENTAL FACILITY

An existing experimental facility — the Oxford low density wind tunnel (LDWT) — has been refurbished and its performance characterised for the present work. This continuous flow, vacuum facility is capable of maintaining pressures less than $10\ Pa$ for extended periods (of the order of hours). It is also equipped with a Mach 6 contoured nozzle, which allows the study of hypersonic gas dynamics in the slip and lower transition regimes.

New data acquisition and experimental support equipment has been designed, manufactured, and installed in the LDWT, bringing it to a point where it may once again be considered a state of the art experimental facility. This equipment includes new traversing mechanisms (for performing pitot studies of the hypersonic plume and orienting test pieces for aerothermodynamic experimental studies), new data acquisition hardware, and a large volume of LabView code which allows a user to automate a large portion of the facility's operation.

Pitot rakes and total temperature surveys of the hypersonic plume have also been performed, culminating in the characterisation of the LDWT's experimental capabilities, and the selection of flow conditions for subsequent experimental work. Numerical simulations utilising a no-slip CFD model were also performed using newly developed low temperature gas property correlations

for air. These simulations were then validated against the pitot data, and were subsequently used to obtain accurate estimations of the plume structure and properties without the need for a complex partial slip model at the nozzle wall. Overall errors in these calculations were low near the nozzle exit, but increased further downstream due to inaccurate modelling of boundary layer growth (due to the aforementioned no-slip model). Despite these inaccuracies, however, the CFD models were found to agree well with core flow properties in the hypersonic plume (especially for those conditions that were selected for subsequent experimental studies).

The data from these simulations and pitot surveys was then used to generate a reduced order model of the facility's performance, hence allowing future users of the LDWT to quickly estimate the required input flow conditions, i.e. total temperature and mass flow rate, for a desired flow condition in the hypersonic plume (in terms of Mach and Knudsen numbers).

The temporal variability of the conditions in the LDWT test chamber were also characterised, and were found to vary on two separate timescales. One of these is related to the action of the facility's mass flow controller and gas heating system, which rely on PID controllers to maintain the stagnation conditions during an experiment. The second source was found to be the vapour booster pump, the facility's primary high vacuum pumping system. Both of these variations were studied, and the reduced order model used to characterise their effects on the stability of conditions in the experimental plume. Errors caused by these variations in predicted flow variables were estimated to be of the order of 2%, hence confirming the LDWT's suitability for the present work's experimental campaign.

7.2 NUMERICAL STUDIES

Numerical studies of re-entry aerothermodynamics comprised the use of 3DOF trajectory simulations and hypersonic panel methods in order to characterise the aerodynamic coefficients of geometric primitives during atmospheric entry. Following these, high fidelity DSMC simulations were performed, utilising freestream variables derived from the LDWT, and also from the results of the aforementioned trajectory simulations.

7.2.1 Re-entry aerothermodynamics

A new code has been developed using open source software and libraries for the purpose of quickly estimating the aerodynamic properties of re-entering space debris items: *RAC* (Re-entry Aerothermal Calculator). This code has been used in conjunction with Gaussian process regression (GPR) models and a series of novel correlative expressions to construct new predictive models for quickly calculating aerodynamic coefficients at hypersonic conditions typical of entry into a planetary atmosphere. These models are designed to provide significant advantages in terms of both accuracy over standard tumble-averaged correlations, and computational efficiency over conventional engineering panel methods, such as those currently used by many analysts in this field.

These models were found to be approximately one order of magnitude faster to query than the panel method code used to provide their training data (RAC). The models were subsequently found to be extremely accurate when compared to test data, with uniformly high coefficients of determination, R^2, indicating an excellent fit to training data. When comparing the predictions of these models to those made by RAC, it was found that the overall errors between the "true" values and the GPR predictions were consistently small, with a maximum mean-normalised RMS error, $\delta_{NRMS(\mu)}$, of 5.886%, and a maximum range-normalised RMS error, $\delta_{NRMS(R)}$ of 0.26%.

As such, it has been shown that models like these may be utilised in future debris demise codes (or, indeed, any code that requires a large number of aerodynamic calculations to be performed). By correlating an object's orientation, geometric aspect ratio, and freestream Mach and Knudsen numbers, the GPR models allow the prediction of aerodynamic coefficients for multiple common space debris shapes with both excellent accuracy, and a negligible penalty in terms of computational expense.

Of particular note in this regard is the aforementioned geometric aspect ratio, which is used as a correlation parameter. The inclusion of this information in the GPR models significantly improves their versatility, as it allows them to be used not just for unit cylinders and cubes (which may be used as analogues of tanks, electronics enclosures, etc.), but also for long, thin rods (such as structural truss elements) or wide, squat objects (like reaction wheels or exterior panels).

Additionally, the simplicity of having a maximum of six input arguments, i.e. Ma, Kn, α, β, ψ, and $L/\varnothing$, also greatly improves the usability of the GPR models, allowing them to be completely generalised for use with a wide variety of freestream conditions.

7.2.2 DSMC modelling of re-entry heating

DSMC simulations were performed using two different codes: Bird's *DS2V* code, and the University of Strathclyde's *dsmcFoam+*. Hemisphere geometries were examined using the former (2D, axisymmetric), and cylindrical geometries with the latter (full 3D). Simulations were performed over a wide range of Knudsen numbers, culminating in an examination of heating that spanned the majority of the transition regime ($0.01 \leqslant Kn \leqslant 10$).

A novel normalisation strategy has been developed, whereby the results of a DSMC analysis computed using non-optimal simulation parameters may be divided by the local stagnation value of heat flux, $\dot{q}$, or heat transfer coefficient, h_{conv}, at an angle of incidence, α, of 0^o. The resulting data may then be used to reconstruct the heat flux trends on arbitrary cylinders at comparable freestream Knudsen numbers, thus vastly reducing computational expense whilst providing improved accuracy over heating correlations that currently see widespread use.

These normalised data (which were generated using the *dsmcFoam+* simulations) were correlated using GPR modelling in order to generate predictive models of the heating on the sharp edges of cylinders. These data are important to the study of such objects during atmospheric entry, as they act to concentrate heat during their transit through the atmosphere. As such, the demise of a cylindrical piece of space debris during re-entry may be inaccurately predicted if this heating is not well understood. The GPR models were also used in lieu of traditional interpolation techniques such that the minimum number of computationally expensive DSMC simulations would have to be performed.

7.3 EXPERIMENTAL STUDIES

Experimental studies carried out using the newly refurbished LDWT comprised 16 cases: 2 utilising a hemisphere geometry, and a further 14 utilising a cylindrical geometry. Two flow conditions were examined, one in the slip regime ($Kn = 0.0216$) and another in the lower

transition regime ($Kn = 0.1012$). The hemispheres were examined at one angle of incidence ($\alpha = 0^o$), while the cylinders were examined at 7 angles of incidence ($0^o \leqslant \alpha \leqslant 90^o$, varying in increments of 15^o).

These probes were coated in a layer of thermochromic liquid crystal, allowing temperature changes on the probe surfaces to be visualised in situ. The videos of the hemisphere experimental studies were analysed using the Osney Lab's THTAC code, which operates using the semi-infinite substrate principle. These data were found to agree extremely well with the results of *DS2V* simulations, and were also found to be poorly predicted by existing hypersonic blunt body heating correlations.

Experiments utilising the cylinder geometries had to be processed in a different manner, owing to the invalidity of the semi-infinite substrate assumption near the corners of the probes. As such, a new post-processing code was developed using work previously undertaken by Ryley. This new program, simply called the 3D HTC solver, was written in MATLAB and utilises the MATLAB-COMSOL LiveLink interface to run FEA multiphysics simulations automatically using scripted instructions. The code operates by mapping the activity from an experimental video onto an FEA mesh, which it then uses (in conjunction with an initial estimate of HTC distribution) to generate a simulated liquid crystal colour play. These data are then compared to the actual colour play as recorded suring an experiment, and adjustments made to the HTC distribution. The FEA solver is then called once again, and the process looped until agreement between simulated and experimental colour plays is observed.

The 3D HTC solver code was used to analyse the remaining cylinder experimental videos, and the resulting data were compared to DSMC results generated using *dsmcFoam+*. Overall agreement between the simulated and experimental HTC distributions was excellent in many cases, with others suffering from anomalous peaks in HTC (likely caused by poor characterisation of colour play errors during processing). Additional errors came about as a result of code in the 3D HTC solver designed to prevent divergence in the FEA solver by limiting HTC values. Any nodes whose HTC values exceed a set of pre-defined limits are assumed to be divergent, and are zeroed for all subsequent iterations of the code.

Overall though, it was clear that the DSMC simulations and 3D HTC solver results agreed

extremely well for many cases at various angles of incidence. Mean percentage errors, $\overline{\delta}$, on the cylinder probes were characterised by sampling data from locations which did not suffer from common errors exhibited by the 3D HTC solver (as characterised in Section 6.2.1). These data were then compared to the corresponding locations in the *dsmcFoam+* results. Mean errors in the HTC distribution of the $\varnothing = 15\,mm$ probe (which was used in experiments at condition 1 — see Table 5.1) were found to be 8.54% and 5.52% on the front and lateral faces of the cylinder, respectively. Agreement between the experimental and simulated HTC distributions of the $\varnothing = 10\,mm$ probe (which was used in experiments run at condition 2) was similar, with mean errors of 6.07% and 10.37% on the front and lateral faces, respectively.

As such, the suitability of these tools, i.e. the 3D HTC solver and *dsmcFoam+*, for the study of hypersonic heat transfer in the slip and lower transition regimes was demonstrated. The normalisation procedure described previously was also validated experimentally. This was performed by using the procedure to normalise DSMC results for the $\varnothing = 15\,mm$ probe, which utilised non-optimal simulation parameters. These data were subsequently compared to experimental results processed using the 3D HTC solver, and were found to agree well at many of the examined sampling locations.

The experimental technique developed herein is expected to greatly improve the understanding of external thermodynamics in the slip and transition regimes for non-aerodynamically styled blunt bodies. Consequently, the experimental data generated in support of the present work have granted insight into the complex heating phenomena around sharp edges on cylindrical bodies in hypersonic conditions. In the context of space debris demise, these data are extremely valuable, as they will facilitate an improvement in the accuracy of thermal load predictions for cylindrical components during re-entry, thereby reducing uncertainties in destructive re-entry calculations.

7.4 FUTURE WORK

The present work has provided insight into hypersonic heating phenomena in the slip and transition regimes for geometric primitives. However, many areas where improvement is necessary were also identified.

In the case of hypersonic rarefied gas studies in the LDWT, it is recommended that future users of the tunnel improve the facility's maximum total temperature, thus increasing the range of available flow conditions (possibly even extending the facility's capabilities far into the transition regime). This could be performed by replacing the seals in the gas conditioning system with high temperature silicon or FFKM. An alternative approach would be the recommissioning of the water cooling system, which has not been utilised in experimental studies since the 1980s. Additionally, it is also recommended that the vapour booster pump's high voltage relays are replaced with a variable current alternative. This modification would allow the facility to operate continuously without having to switch its oil heaters on and off to maintain the boiler temperature (this was a source of uncertainty in the characterisation of the LDWT). The gas conditioning system should also be outfitted with a dedicated PID controller for the working gas heaters, as the current system uses a software implemented controller, which has been shown to introduce small periodic uncertainties in flow conditions.

The experimental work carried out herein demonstrated the suitability of the 3D HTC solver code, but also revealed its sensitivity to the quality of experimental videos. For this reason, future users of this code utilising liquid crystal probes for their experiments should ensure that adequate lighting is incorporated in their experimental setup. Also, it is recommended that regions of interest on experimental probes are oriented as close to perpendicular to the viewing plane of the camera as possible in order to eliminate shallow viewing angles and shadows/glare.

Although the experiments performed for the present work produced favourable results, it must be acknowledged that only a small set of freestream conditions and geomtries were examined. As such, further studies geared towards improving the understanding of hypersonic heating in the slip and transition regimes are recommended. The use of GPR modelling and experiment design via Latin Hypercube Sampling (LHS) is also recommended, as these methods have been shown to be extremely useful in reducing the number of demanding tasks (such as experiments or large scale simulations) required for studies of this kind.

Future users of the LDWT who perform liquid crystal studies should incorporate an additional probe with embedded thermocouples in their experimental design. This probe would be machined to the same specifications as their experimental test pieces, and would allow the equilibration of

internal temperatures to be studied during the initial evacuation of the LDWT test chamber (as, during this process the remaining gas in the chamber expands, causing the internal temperature to drop drastically for a short time). As such, measures to better ensure an isothermal initial condition in future test pieces may be implemented, such as providing a heater or cooling system (as appropriate) to counteract the initial temperature drop during pump down of the test chamber. Future users of the LDWT should also ensure that they are able to accurately measure the initial temperature of their test pieces, as this was found to be a significant source of uncertainty during the present work. Possible solutions include sub-surface thermocouples installed in test pieces, or an infrared camera which would allow the surface temperature to be studied and characterised during evacuation of the test chamber.

Finally, it is recommended that development of the *RAC* code used in the present work is continued. The experiments and DSMC simulations performed for the present work should also be utilised as either validation data for any new methods incorporated into *RAC*, or as training data for new surrogate models of geometry-specific heating.

The following improvements to the code as a whole should be treated as priorities:

- Develop a streamline calculator for evaluating the local continuum Reynolds number (and hence facilitate calculation of viscous forces in the continuum regime).
- Develop and implement an improved shielding algorithm which is less computationally intensive than the visual shielding process.
- Modify the method used to save temporary values to a hard drive such that the code may be executed more efficiently (hard drive access is a performance bottleneck at the time of writing).
- Incorporate or couple the *traj_calc* trajectory code into *RAC* such that full trajectories may be evaluated in a single simulation.
- Include more up-to-date surface heat flux calculation methods such as those implemented in the *FOSTRAD* [16] and *PAMPERO* [15] codes.
- Implement an object-oriented spacecraft breakup model that may be used to study debris demise scenarios.
- Implement or couple a thermal finite element solver such that internal heat transfer during atmospheric entry may be simulated.
- Implement or couple a ablation/partial melting solver such that external material loss during atmospheric entry may be simulated.

- Incorporate GPR models derived from the DSMC simulations and experiments described herein into an object-oriented breakup model (mentioned above).

7.5 CONCLUSION

Overall, this thesis has studied the aerothermal environment of the atmospheric slip-transition regime in the context of space debris re-entry. The so-called "object-oriented" modelling strategy for geometric space debris analogues has been studied, with cylindrical and cuboidal analogues being characterised numerically in terms of aerodynamic coefficients, and with spherical and cylindrical analogues being characterised both numerically and experimentally in terms of convective surface heat transfer. The experimental facility and supporting equipment utilised for these heat transfer studies have also been described, accompanied by details of refurbishment activities performed for the present work, and the results of flow surveys and simulations intended to characterise the facility's experimental capabilities.

The numerical models developed during the present work are intended to improve the accuracy of future debris demise codes by introducing computationally inexpensive surrogate models of the aerodynamic behaviour of debris analogues, thereby reducing uncertainty in trajectory calculations without increasing computational expense. The study of heat transfer to the aforementioned debris analogues has allowed the characterisation of the thermal fluxes on both blunt, spherical bodies, and sharp-edged cylindrical bodies. These experiments, and their accompanying simulations, serve to both validate the numerical and experimental procedures developed and utilised herein, and provide validation data for future heating correlations and studies. Finally, the numerical codes, surrogate models, and experimental infrastructure developed during the course of the present work are expected to greatly benefit the space debris demise community by facilitating further research into the aerodynamics and heating of common space debris geometries in the slip-transition regime.

BIBLIOGRAPHY

[1] ESA, *Space debris by the numbers*, 2019. [Online]. Available: https://m.esa.int/Our%7B%5C_%7DActivities/Space%7B%5C_%7DSafety/Space%7B%5C_%7DDebris/Space%7B%5C_%7Ddebris%7B%5C_%7Dby%7B%5C_%7Dthe%7B%5C_%7Dnumbers.

[2] ——, *ESA Clean Space*. [Online]. Available: https://www.esa.int/Our%7B%5C_%7DActivities/Space%7B%5C_%7DSafety/Clean%7B%5C_%7DSpace.

[3] J. Gelhaus *et al.*, "DRAMA Final Report," ESOC, Darmstadt, Germany, Tech. Rep., 2014. [Online]. Available: files/289/Gelhaus%20-%202014%20-%20DRAMA%20Final%20Report.pdf.

[4] J. Ryley, "Turbine Blade Mid-Chord Internal Cooling," Thesis, University of Oxford, 2014.

[5] IADC, "IADC Space Debris Mitigation Guidelines," Inter-agency Space Debris Coordination Committee, IADC action item IADC-02-01, Sep. 2007. [Online]. Available: http://www.iadc-online.org/Documents/IADC-2002-01,%20IADC%20Space%20Debris%20Guidelines,%20Revision%201.pdf.

[6] J. Gelhaus, "DRAMA Final Report," ESA, ESOC, Darmstadt, Germany, Tech. Rep. ILR/DRAMA/FR, Jun. 2014.

[7] M. McWinnie, "Health Monitoring of Re-entry Vehicles," Ph.D. dissertation, Cranfield University, Cranfield, England, UK, Aug. 2011. [Online]. Available: https://pure.ltu.se/portal/files/36189234/LTU-EX-2012-36180467.pdf.

[8] T. Lips *et al.*, "Comparison of ORSAT and SCARAB Reentry Survival Results," Darmstadt, Germany: ESA, Apr. 2005. [Online]. Available: http://adsabs.harvard.edu/abs/2005ESASP.587..533L.

[9] H. Krag, *ESA Space Debris Tool Workshop*, ESOC, Jun. 2014.

[10] G. Koppenwallner, B. Fritsche, T. Lips, and H. Klinkrad, "SCARAB - A Multi-Disciplinary Code for Destruction Analysis of Space-craft During Re-entry," Cologne, Germany, Feb. 2005. [Online]. Available: http://articles.adsabs.harvard.edu/cgi-bin/nph-iarticle_query?2005ESASP.563..281K&data_type=PDF_HIGH&whole_paper=YES&type=PRINTER&filetype=.pdf.

[11] G. Koppenwallner, B. Fritsche, and T. Lips, "Multidisciplinary analysis tools for orbit and reentry," in *Hyperschall Technologie Goettingen, Kallenburg-Lindau, Germany, presented at the 3rd International Workshop on Astrodynamics Tools and Techniques, ESTEC, Noordwijk, Netherlands*, 2006. [Online]. Available: http://trajectory.estec.esa.int/Astro/3rd-astro-workshop-presentations/Multidisciplinary%20Analysis%20Tools%20for%20Orbit%20and%20Re-entry.pdf (visited on 07/01/2014).

[12] R. L. Kelley, N. M. Hill, W. C. Rochelle, N. L. Johnson, and T. Lips, "COMPARISON of ORSAT and SCARAB Reentry Analysis Tools for a Generic Satellite Test Case," English, NASA, Bremen, Germany, COSPAR Scientific Assembly presentation PEDAS1-0021-10, Jul. 2010. [Online]. Available: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20100005304.pdf.

[13] J. Merrifield, "Simplified Aerothermal Models for S / C Components (Sam):" Tech. Rep. July, 2015, pp. 1–24.

[14] J. Merrifield, J. Beck, G. Markelov, and R. Molina, "Simplified aerothermal models for destructive entry analysis," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal: ESA, 2015.

[15] J. Annaloro, P. Omaly, V. Rivola, and M. Spel, "ELABORATION OF A NEW SPACECRAFT-ORIENTED TOOL : PAMPERO," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal: ESA, 2015.

[16] A. Falchi, V. Renato, E. Minisci, and M. Vasile, "FOSTRAD : An Advanced Open Source Tool for Re-entry Analysis," in *15th Reinventing Space Conference*, Glasgow, Scotland, UK, 2017.

[17] A. Falchi, E. Minisci, M. Vasile, D. Rastelli, and N. Bellini, "DSMC-based Correction Factor for Low-fidelity Hypersonic Aerodynamics of Re-entering Objects and Space Debris," in *7th European Conference for Aeronautics and Aerospace Sciences (EUCASS)*, 2017.

[18] A. Falchi, E. Minisci, M. Vasile, and M. Kubicek, "Aero-thermal re-entry sensitivity analysis using DSMC and a high dimensional model representation-based approach.," in *7th European Conference on Space Debris*, 2017.

[19] F. Sourgen and E. Laroche, "MUSIC / FAST , A PRE-DESIGN AND PRE-MISSION ANALYSIS TOOL FOR THE EARTH ATMOSPHERIC RE-ENTRY OF SPACECRAFT , CAPSULES AND DE-ORBITED SATELLITES," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal: ESA, 2015.

[20] P. Omaly and M. Spel, "DEBRISK, a Tool for Re-Entry Risk Analysis," *5th IAASS Conference A Safer Space for Safer World*, vol. 2011, no. October 2011, 2012.

[21] J. Gelhaus, "DRAMA Software User Manual," ESOC, Darmstadt, Germany, Tech. Rep., 2014.

[22] G. Koppenwallner, B. Fritsche, T. Lips, and H. Klinkrad, "SCARAB - A Multi-Disciplinary Code for Destruction Analysis of Space-craft During Re-entry," in *Fifth European Symposium on Aerothermodynamics for Space Vehicles*, 2005. [Online]. Available: `http://articles.adsabs.harvard.edu/cgi-bin/nph-iarticle%7B%5C_%7Dquery?2005ESASP.563..281K%7B%5C&%7Damp%20data%7B%5C_%7Dtype=PDF%7B%5C_%7DHIGH%7B%5C&%7Damp%20whole%7B%5C_%7Dpaper=YES%7B%5C&%7Damp%20type=PRINTER%7B%5C&%7Damp%20filetype=.pdf%20files/62/Koppenwallner%20et%20al.%20-%202005%20-%20SCARAB%20-%20A%20Multi-Disciplinary%20Code%20for%20Destruction.pdf`.

[23] R. L. Kelley, N. M. Hill, W. C. Rochelle, N. L. Johnson, and T. Lips, "COMPARISON of ORSAT and SCARAB Reentry Analysis Tools for a Generic Satellite Test Case," English, Bremen, Germany, Tech. Rep., 2010. [Online]. Available: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20100005304.pdf%20files/13/Kelley%20et%20al.%20-%202010%20-%20COMPARISON%20of%20ORSAT%20and%20SCARAB%20Reentry%20Analysis%20To.pdf.

[24] T. Lips *et al.*, "Comparison of ORSAT and SCARAB Reentry Survival Results," in *Fourth European Conference of Space Debris*, ESA, 2005. [Online]. Available: http://adsabs.harvard.edu/abs/2005ESASP.587..533L%20files/50/Lips%20et%20al.%20-%202005%20-%20Comparison%20of%20ORSAT%20and%20SCARAB%20Reentry%20Survival%20Re.pdf.

[25] Astromaterials Research and Exploration Science Directorate *et al.*, *Debris Assessment Software User's Guide (Version 2.0)*, Nov. 2007. [Online]. Available: http://orbitaldebris.jsc.nasa.gov/library/DAS2%7B%5C_%7D0/DAS2.0%7B%5C_%7DUsersGuide.pdf%20files/43/DAS2.0%7B%5C_%7DUsersGuide.pdf.

[26] W. C. Rochelle, R. E. Kinsey, E. A. Reid, R. C. Reynolds, and N. L. Johnson, "Spacecraft Orbital Debris Aerothermal Analysis," Lockheed Martin/NASA, Houston, TX, USA, Tech. Rep., 1997. [Online]. Available: https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19970040121.pdf.

[27] J. Beck and J. Merrifield, "A heat balance integral method for destructive re-entry codes," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal: ESA, 2015.

[28] Y. Prévereaud, J.-l. Vérant, J. Annaloro, and S. Galera, "Development of New Analytical Models of Pressure and Heat Transfer Distributions on Space Debris Uncontrolled Atmospheric Entry : Planar Bodies," in *Proceedings of the 7th European Conference for Aeronautics and Space Sciences (EUCASS)*, Milan, Italy, 2017, p. 5959. DOI: 10.13009/EUCASS2017-603.

[29] J. Annaloro *et al.*, "Space Debris Atmospheric Entry Prediction With Spacecraft-Oriented Tools," in *Proceedings of the 7th European Conference on Space Debris*, Darmstadt, Germany: ESA Space Debris Office, 2017, pp. 18–21.

[30] P. M. Mehta, E. Minisci, M. Vasile, A. C. Walker, and M. Brown, "An Open-source Hypersonic Aerodynamic And Aerothermodynamic Modeling Tool," *8th European Symposium on Aerothermodynamics for Space Vehicles*, no. 1, p. 9, 2015.

[31] N. Donaldson, *traj_calc - 3DoF atmospheric entry ODE solvers*, 2016. DOI: 10.5281/zenodo.49930. [Online]. Available: http://dx.doi.org/10.5281/zenodo.49930.

[32] N. L. Donaldson and P. Ireland, "A Panel Method Aerodynamic Preprocessor for Planetary Entry Trajectory Simulations," in *21st AIAA International Space Planes and Hypersonics Technologies Conference*, 2017, pp. 6–9, ISBN: 978-1-62410-463-3. DOI: 10.2514/6.2017-2379. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2017-2379.

[33] G. A. Bird, *Molecular gas dynamics and the direct simulation of gas flows*, ser. Oxford engineering science series. Oxford: Oxford : Clarendon Press, 1994.

[34] J. D. Anderson, *Hypersonic and high-temperature gas dynamics*, 2nd ed.., I. ebrary, Ed., ser. AIAA education series. Reston, Va.: Reston, Va. : American Institute of Aeronautics and Astronautics, 2006, ISBN: 978-1-56347-780-5.

[35] *Molecular sizes*. Encyclopedia Britannica, 2020. [Online]. Available: https://www.britannica.com/science/gas-state-of-matter/Molecular-sizes.

[36] W. Sutherland, "LII. The viscosity of gases and molecular force," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 36, no. 223, pp. 507–531, 1893, ISSN: 1941-5982. DOI: 10.1080/14786449308620508.

[37] G. T. Chapman, "Theoretical laminar convective heat transfer and boundary-layer characteristics on cones at speeds to 24 km/sec," no. August, 1964. [Online]. Available: http://hdl.handle.net/2060/19640017893.

[38] C. R. Wilkie, "A viscosity equation for gas mixtures," *The Journal of Chemical Physics, Volume 18, Issue 4, p.517-519*, vol. 18, no. 4, pp. 517–519, 1950. DOI: 10.1063/1.1747673.

[39] D. G. Goodwin, H. K. Moffat, and R. L. Speth, *Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes*, http://www.cantera.org, Version 2.2.1, 2016.

[40] ——, *Cantera: Mixtransport class reference*, https://cantera.github.io/docs/doxygen/html/classCantera_1_1MixTransport.html, Version 2.2.1, 2016.

[41] G. A. Bird, "Monte Carlo simulation of gas flows," *Annual Review of Fluid Mechanics*, vol. 10, no. 1, pp. 11–31, 1978. [Online]. Available: http://www.annualreviews.org/doi/pdf/10.1146/annurev.fl.10.010178.000303 (visited on 03/04/2014).

[42] B. J. Alder and T. E. Wainwright, "Studies in Molecular Dynamics. I. General Method," en, *The Journal of Chemical Physics*, vol. 31, no. 2, p. 459, 1959, ISSN: 00219606. DOI: 10.1063/1.1730376. [Online]. Available: http://scitation.aip.org/content/aip/journal/jcp/31/2/10.1063/1.1730376 (visited on 08/19/2014).

[43] G. A. Bird, *The DSMC Method*, 1.2. 2013, ISBN: 9781492112907.

[44] J. D. Anderson, *Hypersonic and high-temperature gas dynamics*, 2nd ed.., I. Ebrary, Ed. Reston, Va.: Reston, Va. : American Institute of Aeronautics and Astronautics, 2006, ISBN: 978-1-56347-780-5.

[45] R. Grabow, C. Pannabecker, A. Plate, and R. Rice, "The Determination of Hypersonic Drag Coefficients for Cones and Spheres," Research and Advanced Develpoment Division Avco Corporation, Redstone Arsenal, AL, USA, Tech. Rep., 1965. [Online]. Available: http://www.dtic.mil/dtic/tr/fulltext/u2/487365.pdf.

[46] D. J. Masson, D. N. Morris, and D. E. Bloxsom, *Measurements of Sphere Drag From Hypersonic Continuum to Free-Molecular Flow*, 1960. [Online]. Available: http://www.rand.org/content/dam/rand/pubs/research%7B%5C_%7Dmemoranda/2008/RM2678.pdf.

[47] S. A. Schaaf and L. Talbot, "Handbook of Supersonic Aerodynamics - Volume 5, Section 16 - Mechanics of Rarefied Gases," University of California, Berkely, CA, Tech. Rep. NAVORD REPORT 1488 (Vol. 5), 1959.

[48] J. B. Parham and L. a. Barba, "Finding the Force - Consistent Particle Seeding for Satellite Aerodynamics," pp. 1–14, 2014. arXiv: arXiv:1312.3691v1.

[49] D. Kliche, C. Mundt, and E. H. Hirschel, "The hypersonic Mach number independence principle in the case of viscous flow," *Shock Waves*, vol. 21, no. 4, pp. 307–314, 2011, ISSN: 09381287. DOI: 10.1007/s00193-011-0318-y.

[50] E. Buckingham, "On physically similar systems; illustrations of the use of dimensional equations," *Physical Review*, vol. 4, pp. 345–376, 1914.

[51] R. D. Klett, "Drag Coefficients and Heating Ratios for Right Circular Cylinders in Free-Molecular and Continuum Flow from Mach 1 to 30," English, Albuquerque, NM, USA, Tech. Rep., 1964.

[52] P. Agrawal, S. A. Sepka, J. F. Aliaga, E. Venkatapathy, and J. A. Samareh, "Thermal Soak Analysis of Earth Entry Vehicles," *43rd AIAA Thermophysics Conference*, no. June, pp. 1–16, 2012. DOI: 10.2514/6.2012-3010.

[53] A. Viviani and G. Pezzella, "Nonequilibrium Aerothermodynamics for a Capsule Reentry Vehicle," *Engineering Applications of Computational Fluid Mechanics*, vol. 3, no. 4, pp. 543–561, 2009, ISSN: 1994-2060. DOI: 10.1080/19942060.2009.11015290. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/19942060.2009.11015290.

[54] AGARD, "Capsule Aerothermodynamics," von Karman Institute for Fluid Dynamics (VKI), Rhode-Saint-Genese, Belgium, Tech. Rep., 1995. [Online]. Available: https://www.cso.nato.int/Pubs/rdp.asp?RDP=AGARD-R-808.

[55] F. K. Lu, H. Kim, and L. N. McNay, "A simplified trajectory analysis model for small satellite payload recovery from low Earth orbit," *Aerospace Science and Technology*, vol. 7, no. 3, pp. 231–237, 2003, ISSN: 12709638. DOI: 10.1016/S1270-9638(02)00016-0.

[56] J.-p. Baumann, A. S. Pagan, and G. Herdrich, "AEROTHERMODYNAMIC RE-ENTRY ANALYSIS OF THE CUBESAT-SIZED ENTRY VEHICLE MIRKA2," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal: ESA, 2015.

[57] H. J. Allen and A. J. Eggers, "A study of the motion and aerodynamic heating of ballistic missiles entering the earth's atmosphere at high supersonic speeds," NACA, Tech. Rep., 1958, p. 18. [Online]. Available: http://hdl.handle.net/2060/19930091020.

[58] Z. R. Putnam and R. D. Braun, "Extension and Enhancement of the Allen–Eggers Analytical Ballistic Entry Trajectory Solution," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 3, pp. 414–430, 2015, ISSN: 0731-5090. DOI: 10.2514/1.G000846. [Online]. Available: http://dx.doi.org/10.2514/1.G000846.

[59] E. Gentry, Arvel, D. N. Smyth, and W. R. Oliver, *The Mark IV Supersonic-Hypersonic Arbitrary Body Program; Volume II - Program Formulation*, English, 1973.

[60] C. Rosema, J. Doyle, L. Auman, M. Underwood, and W. B. Blake, "AFRL-RB-WP-TR-2011-3071 Missile DATCOM User's Manual - 2011 Revision.," vol. 3071, 2011.

[61] M. Pontani and P. Teofilatto, "Simple method for performance evaluation of multistage rockets," *Acta Astronautica*, vol. 94, no. 1, pp. 434–445, 2014, ISSN: 00945765. DOI: 10.1016/j.actaastro.2013.01.013. [Online]. Available: http://dx.doi.org/10.1016/j.actaastro.2013.01.013.

[62] D. Kinney, "Aerothermal Anchoring of CBAERO Using High Fidelity CFD," pp. 1–28, 2007. [Online]. Available: http://www.aric.or.kr/treatise/journal/content.asp?idx=96691.

[63] R. Wuilbercq and R. E. Brown, "RAPID AERO-THERMODYNAMIC ANALYSIS FOR HYPERSONIC AIR VEHICLES," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal, 2015.

[64] R. Wuilbercq, A. O. Ahmad, T. Scanlon, and R. Brown, "Towards Robust Aero-Thermodynamic Predictions for Re-Usable Single-Stage to Orbit Vehicles," *AIAA*, vol. 22, pp. 1753–1759, 2008, ISSN: 03601323. DOI: 10.1016/j.buildenv.2006.10.027. arXiv: 0812.0143v2.

[65] P. M. Mehta, E. Minisci, M. Vasile, A. C. Walker, and M. Brown, "AN OPEN-SOURCE HYPERSONIC AERODYNAMIC AND AEROTHERMODYNAMIC," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal, 2015.

[66] L. Morsa *et al.*, "Analysis of Bridging Formulae in Transitional Regime," in *AIP Conference Proceedings-American Institute of Physics*, vol. 1333, 2011, p. 1319. [Online]. Available: http://www2.aero.psu.edu/RGD_2010/PaperSubmission_Full_Uploads/Morsa,%20Luigi_30_PDF_Final-Morsa-Zuppardi-Schettino-Votta-30.pdf (visited on 08/22/2014).

[67] A. Owen, "Experimental Studies of the Hypersonic, Low Density, Aerodynamics of Re-entry Vehicles," Ph.D. dissertation, University of Oxford, Oxford, UK, 1997.

[68] R. G. Wilmoth, R. C. Blanchard, and J. N. Moss, "Rarefied transitional bridging of blunt body aerodynamics," *21st International Symposium on Rarefied Gas Dynamics*, pp. 26–31, 1998.

[69] J. A. Fay and F. R. Riddell, "Theory of Stagnation Point Heat Transfer in Dissociated Air," *Journal of the Aeronautical Sciences*, vol. 25, no. 2, pp. 73–85, Feb. 1958. [Online]. Available: http://pdf.aiaa.org/downloads/TOCPDFs/36_373-386.pdf.

[70] R. W. Detra and H. Hidalgo, "Generalized heat transfer formulas and graphs for nose cone re-entry into the atmosphere," *ARS Journal*, vol. 31, pp. 318–321, 1961.

[71] A. M. Brandis and C. O. Johnston, "Characterization of Stagnation-Point Heat Flux for Earth Entry," in *AIAA Aviation*, Atlanta, Georgia, USA: AIAA, Jun. 2014.

[72] *The Problem*. [Online]. Available: http://www.math.uwaterloo.ca/tsp/problem/index.html (visited on 08/19/2014).

[73] C. White *et al.*, "dsmcFoam+: An OpenFOAM based direct simulation Monte Carlo solver," *Computer Physics Communications*, vol. 224, pp. 22–43, 2017, ISSN: 00104655. DOI: 10.1016/j.cpc.2017.09.030. [Online]. Available: https://doi.org/10.1016/j.cpc.2017.09.030.

[74] C. White, Private communication.

[75] T. Scanlon, E. Roohi, C. White, M. Darbandi, and J. Reese, "An open source, parallel DSMC code for rarefied gas flows in arbitrary geometries," *Computers & Fluids*, vol. 39, no. 10, pp. 2078–2089, Dec. 2010, ISSN: 00457930. DOI: 10.1016/j.compfluid.2010.07.014. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0045793010001891 (visited on 02/24/2014).

[76] G. A. Bird, "Sophisticated dsmc," in *Notes prepared for a short course at the DSMC07 meeting Santa Fe*, 2007. [Online]. Available: http://www.ae.su.oz.au/dsmc_gab/Resources/DSMC07notes.pdf (visited on 04/10/2014).

[77] C. White, "Benchmarking, Development and Applications of an Open Source DSMC Solver," Ph.D. dissertation, Glasgow, Scotland, UK, 2013. [Online]. Available: files/239/White%20-%202013%20-%20Benchmarking,%20Development%20and%20Applications%20of%20an%20O.pdf.

[78] G. A. Bird, "The Q-K model for gas-phase chemical reaction rates," en, *Physics of Fluids*, vol. 23, no. 10, p. 106 101, 2011, ISSN: 10706631. DOI: 10.1063/1.3650424. [Online]. Available: http://scitation.aip.org/content/aip/journal/pof2/23/10/10.1063/1.3650424 (visited on 08/22/2014).

[79] Z. Shang, "3d DSMC Simulation of Rarefied Gas Flows around a Space Crew Capsule Using OpenFOAM," *Open Journal of Applied Sciences*, vol. 03, no. 01, pp. 35–38, 2013, ISSN: 2165-3917, 2165-3925. DOI: 10.4236/ojapps.2013.31005. [Online].

Available: http://www.scirp.org/journal/PaperDownload.aspx?DOI=10.4236/ojapps.2013.31005 (visited on 02/24/2014).

[80] J. N. Moss, R. C. Blanchard, R. G. Wilmoth, and R. D. Braun, "Mars Pathfinder Rarefied Aerodynamics: Computations and Measurements," AIAA report AIAA 98-0298, 1998. [Online]. Available: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19980030782.pdf.

[81] J. N. Moss, R. G. Wilmoth, and J. M. Price, "DSMC Simulations of Blunt Body Flows for Mars Entries: Mars Pathfinder and Mars Microprobe Capsules," English, NASA Langley Research Center, Hampton, Virginia, USA, Tech. Rep., 1997. [Online]. Available: http://www.cs.odu.edu/~mln/ltrs-pdfs/NASA-aiaa-97-2508.pdf.

[82] F. W. Vogenitz, G. A. Bird, J. E. Broadwell, and H. Rungaldier, "Theoretical and Experimental Study of Rarefied Supersonic Flows about Several Simple Shapes," *AIAA Journal*, vol. 6, no. 12, pp. 2388–2394, Feb. 1968.

[83] A. Walker, "DSMC Modeling of Drag Coefficients," Los Alamos National Laboratory, Tech. Rep. LA-UR-13-20529.

[84] C.-H. Cheng and F.-L. Liao, "DSMC analysis of rarefied gas flow over a rectangular cylinder at all Knudsen numbers," *Journal of fluids engineering*, vol. 122, no. 4, pp. 720–729, 2000. [Online]. Available: http://fluidsengineering.asmedigitalcollection.asme.org/article.aspx?articleid=1429274 (visited on 08/28/2014).

[85] M. Mcgilvray, L. Doherty, R. Morgan, and D. Gildfind, "T6: the oxford university stalker tunnel," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal, 2015.

[86] M. Mcgilvray *et al.*, "The oxford high density tunnel," in *Proceedings of the 8th European Symposium on Aerothermodynamics for Space Vehicles*, Lisbon, Portugal, 2015.

[87] A. Owen, "Experimental Studies of the Hypersonic, Low Density, Aerodynamics of Re-entry Vehicles," Ph.D. dissertation, Oxford, UK, 1997.

[88] V. Lago, R. Joussot, and J. D. Parisse, "Influence of the ionization rate of a plasma discharge applied to the modification of a supersonic low Reynolds number flow field around a cylinder," *Journal of Physics D: Applied Physics*, vol. 47, no. 12, 2014, ISSN: 00223727. DOI: 10.1088/0022-3727/47/12/125202.

[89] T. Ozawa, T. Suzuki, and K. Fujita, "Aerodynamic Measurements and Computational Analyses in Hypersonic Rarefied Flows," *AIAA Journal*, vol. 53, no. 11, pp. 3327–3337, 2015, ISSN: 0001-1452. DOI: 10.2514/1.J053889. [Online]. Available: http://arc.aiaa.org/doi/10.2514/1.J053889.

[90] W. H. Beck, U. Henne, and C. Hirschen, "Application of the Pressure Sensitive Paint Technique to a Scramjet Nozzle Flow in the DLR Hypersonic Wind Tunnel," no. October, pp. 7–10, 2008.

[91] L. d. Luca, G. Cardone, G. M. Carlomagno, D. A. d. l. Chevalerie, and T. A. d. Roquefort, "FLOW VISUALIZATION AND HEAT TRANSFER MEASUREMENT IN A HYPERSONIC WIND TUNNEL," en, *Experimental Heat Transfer*, vol. 5, no. 1, pp. 65–78, Jan. 1992, ISSN: 0891-6152, 1521-0480. DOI: 10.1080/08916159208946433. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/08916159208946433 (visited on 08/08/2014).

[92] Y. Le Sant, M. Marchand, M. Millan, and J. Fontaine, "An overview of infrared thermography techniques used in large wind tunnels," *Aerospace Science and Technology*, vol. 6, pp. 355–366, May 2002. [Online]. Available: http://ac.els-cdn.com/S1270963802011720/1-s2.0-S1270963802011720-main.pdf?_tid=4c77bdd6-1f2f-11e4-a9e9-00000aab0f01&acdnat=1407525053_b8c423166e826ebceb774719e07a2235.

[93] A. Henckels and F. Maurer, "Applications of infrared thermography in a hypersonic blowdown wind tunnel," in *Instrumentation in Aerospace Simulation Facilities, 1989. ICIASF'89 Record., International Congress on*, IEEE, 1989, pp. 516–524. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=77708 (visited on 08/08/2014).

[94] C. O. Asma and V. Van der Haegen, "Footprint Analysis of the EXPERT Vehicle Winglet at Hypersonic Conditions," Chicago, Illinois, USA: AIAA, Jul. 2010. [Online]. Available: http://adsabs.harvard.edu/abs/2009ESASP.659E...1R.

[95] D. L. Schultz and T. V. Jones, "Heat-Transfer Measurements in Short-Duration Hypersonic Facilities," English, Oxford, UK, Tech. Rep., 1973. [Online]. Available: http://ftp.rta.nato.int/public//PubFullText/AGARD/AG/AGARD-AG-165//AGARD-AG-165.pdf%20files/212/Schultz%20and%20Jones%20-%201973%20-%20Heat-Transfer%20Measurements%20in%20Short-Duration%20Hyper.pdf.

[96] M. Collins, K. Chana, and T. Povey, "New technique for the fabrication of miniature thin film heat flux gauges," *Measurement Science and Technology*, vol. 26, no. 2, p. 025 303, 2015, ISSN: 0957-0233. DOI: 10.1088/0957-0233/26/2/025303. [Online]. Available: http://stacks.iop.org/0957-0233/26/i=2/a=025303?key=crossref.6427c35b77a5dcb749286d815a9cc561.

[97] M. L. G. Oldfield, "Impulse Response Processing of Transient Heat Transfer Gauge Signals," *Journal of Turbomachinery*, vol. 130, no. April, p. 021 023, 2008, ISSN: 0889504X. DOI: 10.1115/1.2752188.

[98] P. T. Ireland and T. V. Jones, "Liquid crystal measurements of heat transfer and surface shear stress," *Measurement Science and Technology*, vol. 11, no. 7, Jan. 2000. [Online]. Available: http://iopscience.iop.org/0957-0233/11/7/313%20files/32/0957-0233%7B%5C_%7D11%7B%5C_%7D7%7B%5C_%7D313.pdf.

[99] T. Liu, B. T. Campbell, J. P. Sullivan, J. Lafferty, and W. Yanta, "Heat transfer measurement on a waverider at Mach 10 using fluorescent paint," *Journal of Thermophysics and Heat Transfer*, vol. 9, no. 4, pp. 605–611, Oct. 1995, ISSN: 0887-8722, 1533-6808. DOI: 10.2514/3.714. [Online]. Available: http://arc.aiaa.org/doi/abs/10.2514/3.714 (visited on 01/20/2014).

[100] T. Liu, Z. Cai, J. Lai, J. Rubal, and J. P. Sullivan, "Analytical Method for Determining Heat Flux from Temperature-Sensitive-Paint Measurements in Hypersonic Tunnels," *Journal of Thermophysics and Heat Transfer*, vol. 24, no. 1, pp. 85–94, Jan. 2010, ISSN: 0887-8722, 1533-6808. DOI: 10.2514/1.43372. [Online]. Available: http://arc.aiaa.org/doi/abs/10.2514/1.43372 (visited on 01/20/2014).

[101] J. P. C. W. Ling, "Development of Heat Transfer Measurement Techniques and Cooling Strategies for Gas Turbines," Ph.D. dissertation, University of Oxford, 2005.

[102] P. M. Mehta, M. Kubicek, E. Minisci, and M. Vasile, "Debris re-entry modeling using high dimensional derivative based uncertainty quantification," in *Advances in the Astronautical Sciences*, vol. 156, American Astronautical Society, 2016, pp. 3993–4011, ISBN: 9780877036296. [Online]. Available: https://pureportal.strath.ac.uk/en/publications/debris-re-entry-modeling-using-high-dimensional-derivative-based-.

[103] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006, ISBN: 026218253X. DOI: 10.1142/S0129065704001899. arXiv: 026218253X.

[104] J. Sacks, W. J. Welch, J. S. B. Mitchell, and P. W. Henry, "Design and Analysis of Computer Experiments," *Statistical Science*, vol. 4, no. 4, pp. 409–423, 1989, ISSN: 08834237. DOI: 10.2307/2245858. [Online]. Available: http://dx.doi.org/10.2307/2245858.

[105] A. Forrester, A. Sobester, and A. Keane, *Engineering design via surrogate modelling: a practical guide*. 2008, ISBN: 0470770791. [Online]. Available: http://eprints.soton.ac.uk/64699/.

[106] D. K. Duvenaud, "Automatic Model Construction with Gaussian Processes," no. June, p. 144, 2014. [Online]. Available: https://www.cs.toronto.edu/%7B~%7Dduvenaud/thesis.pdf.

[107] Sheffield Machine Learning Group, *GPy: A gaussian process framework in python*, http://github.com/SheffieldML/GPy, Department of Computer Science, University of Sheffield, 2012.

[108] T. Lips and B. Fritsche, "A comparison of commonly used re-entry analysis tools," en, *Acta Astronautica*, vol. 57, no. 2-8, pp. 312–323, Aug. 2005, ISSN: 00945765. DOI: 10.1016/j.actaastro.2005.03.010. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0094576505000767%20files/246/Lips%20and%20Fritsche%20-%202005%20-%20A%20comparison%20of%20commonly%20used%20re-entry%20analysis%20to.pdf.

[109] J. Picone, A. Hedin, and D. Drob, "NRLMSISE-00 Empirical model of the atmosphere statistical comparison and scientific issues," *Journal of Geophysical Research*, p. 70, 2001. [Online]. Available: https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20020038771.pdf.

[110] DeepHorizons, *Python-nrlmsise-00*, https://github.com/DeepHorizons/Python-NRLMSISE-00. [Online]. Available: https://github.com/DeepHorizons/Python-NRLMSISE-00.

[111] D. G. Tetzman, "Simulation and Optimization of Spacecraft Re-entry Trajectories," Ph.D. dissertation, University of Minnesota, May 2010. [Online]. Available: http://hdl.handle.net/11299/93005.

[112] A. J. Hodges, "The Drag Coefficient of Very High Velocity Spheres," *Journal of the Aeronautical Sciences*, no. 3, pp. 755–759, 1958.

[113] S. A. Schaaf and P. L. Chambré, *Flow of rarefied gases*, ser. Princeton aeronautical paperbacks 8. Princeton University Press, 1958.

[114] P. Moore and A. Sowter, "Application of a satellite aerodynamics model based on normal and tangential momentum accommodation coefficients," *Planetary and Space Science*, vol. 39, no. 10, pp. 1405–1419, 1991, ISSN: 00320633. DOI: 10.1016/0032-0633(91)90020-B.

[115] S. F. Rafano Carná and R. Bevilacqua, "High fidelity model for the atmospheric re-entry of CubeSats equipped with the Drag De-Orbit Device," *Acta Astronautica*, vol. 156, no. September 2017, pp. 134–156, 2019, ISSN: 00945765. DOI: 10.1016/j.actaastro.2018.05.049. [Online]. Available: https://doi.org/10.1016/j.actaastro.2018.05.049.

[116] C. B. Barber, D. P. Dobkin, and H. T. Huhdanpaa, ""the quickhull algorithm for convex hulls"," *ACM Trans. on Mathematical Software*, vol. 22, no. 4, pp. 469–483, Dec. 1996.

[117] T. K. F. Da, "2D alpha shapes," in *CGAL User and Reference Manual*, 4.9, CGAL Editorial Board, 2016. [Online]. Available: http://doc.cgal.org/4.9/Manual/packages.html#PkgAlphaShape2Summary.

[118] *Polygon Area*, 2020. [Online]. Available: http://mathworld.wolfram.com/PolygonArea.html (visited on 02/05/2020).

[119] P. M. Mehta, G. B. Arnao, D. Bonetti, E. Minisci, and M. Vasile, "COMPUTER GRAPHICS FOR SPACE DEBRIS," in *ICATT*, vol. 1, 2015, ISBN: 9788578110796. DOI: 10.1017/CBO9781107415324.004. arXiv: arXiv:1011.1669v3. [Online]. Available: https://indico.esa.int/indico/event/111/session/14/contribution/58/material/paper/0.pdf.

[120] P. Ramachandran and G. Varoquaux, "Mayavi: 3D Visualization of Scientific Data," *Computing in Science & Engineering*, vol. 13, no. 2, pp. 40–51, 2011, ISSN: 1521-9615.

[121] E. Gentry, Arvel, D. N. Smyth, and W. R. Oliver, *The Mark IV Supersonic-Hypersonic Arbitrary Body Program; Volume II - Program Formulation*, English, 1973.

[122] D. Kinney, "Aero-thermodynamics for conceptual design," *AIAA Paper*, no. January, pp. 1–11, 2004. [Online]. Available: http://www.aric.or.kr/treatise/journal/content.asp?idx=53932.

[123] C. I. Cruz and G. J. Sova, "Improved Tangent-Cone Method for the Aerodynamic Preliminary Analysis System (APAS) Version of the Hypersonic Arbitrary-Body Program," NASA, Tech. Rep., 1990. [Online]. Available: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19900006574.pdf.

[124] E. Bonner, W. Clever, and K. Dunn, "Aerodynamic Preliminary Analysis System II, Part I - Theory," Rockwell International Corporation, Los Angeles, CA, USA, Tech. Rep., 1991, p. 32. [Online]. Available: https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19910013767.pdf.

[125] Ames research staff, "Equation, Tables and Charts for Compressible Flow," English, NACA, Tech. Rep. 1135, 1953, p. 73. DOI: 10.1097/00152192-198911000-00004. [Online]. Available: https://www.grc.nasa.gov/www/BGH/Images/naca1135.pdf.

[126] L. Lees, "Laminar Heat Transfer Over Blunt-Nosed Bodies at Hypersonic Flight Speeds," *Journal of Jet Propulsion*, vol. 26, no. 4, pp. 259–269, 1956, ISSN: 1936-9980. DOI: 10.2514/8.6977. [Online]. Available: http://arc.aiaa.org/doi/10.2514/8.6977.

[127] E. Eckert and O. Drewitz, "The Heat Transfer to a Plate in Flow at High Speed," Washington, USA, Tech. Rep., 1943.

[128] J. Merrifield, Private communication.

[129] J. Ahrens, B. Geveci, and C. Law, "ParaView: An end-user tool for large-data visualization," in *Visualization Handbook*, 2005, ISBN: 9780123875822. DOI: 10.1016/B978-012387582-2/50038-1. arXiv: arXiv:1301.7609v1.

[130] G. Dahlen, "Cone Drag in the Transition from Contiuum to Free Molecular Flow," England, UK, Tech. Rep., 1984.

[131] ——, "Cone Drag in the Transition from Contiuum to Free Molecular Flow," University of Oxford, England, UK, OUEL report 1538/84, Mar. 1984.

[132] A. Widodo and D. Buttsworth, "Stagnation temperature in a cold hypersonic flow produced by a light free piston compression facility," *Experiments in Fluids*, vol. 54, no. 4, 2013, ISSN: 07234864. DOI: 10.1007/s00348-013-1486-6.

[133] H. S. Carslaw and J. C. Jaeger, *Conduction of heat in solids*, 2nd ed. London: Oxford Science Publications, 1959.

[134] *Blender Reference Manual*, Blender Foundation, 2019.

[135] B. Tang, M. McGilvray, and D. Gillespie, *Transient Heat Transfer Analysis Code (THTAC) Instructions for End Users*, Osney Thermofluids Institute, University of Oxford, 2018.

[136] A. Savitzky and M. J. E. Golay, "Smoothing and Differentiation of Data by Simplified Least Squares Procedures.," *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, Jul. 1964, ISSN: 0003-2700. DOI: 10.1021/ac60214a047. [Online]. Available: http://pubs.acs.org/doi/abs/10.1021/ac60214a047.

[137] P. Ireland, "Internal Cooling of Turbine Blades," Ph.D. dissertation, University of Oxford, 1987.

[138] P. J. Linstrom and W. G. Mallard, *NIST Chemistry WebBook, NIST Standard Reference Database Number 69*. Gaithersburg, MD: National Institute of Standards and Technology, 2018. DOI: doi.org/10.18434/T4D303. [Online]. Available: https://webbook.nist.gov.

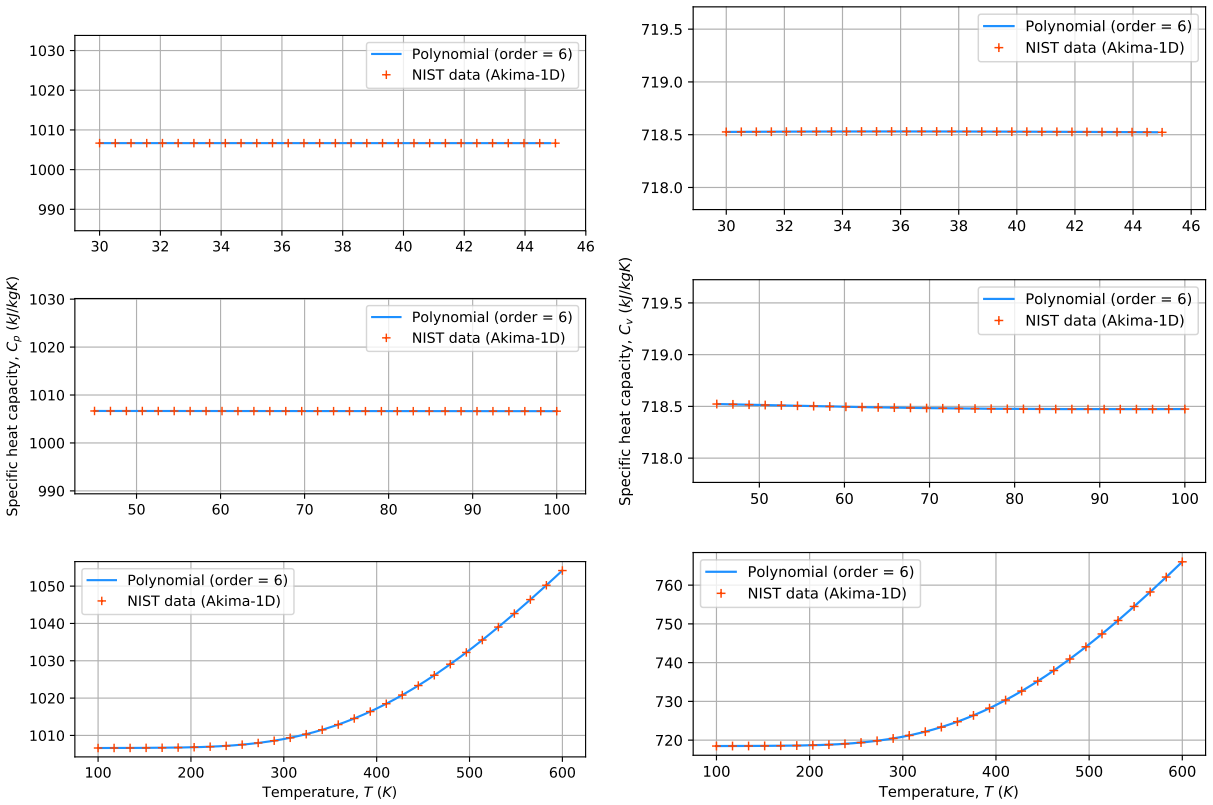
[139] H. Akima, "A New Method of Interpolation and Smooth Curve Fitting Based on Local Procedures," *Journal of the ACM*, vol. 17, no. 4, pp. 589–602, 1970, ISSN: 00045411. DOI: 10.1145/321607.321609. [Online]. Available: http://portal.acm.org/citation.cfm?doid=321607.321609.

APPENDIX A



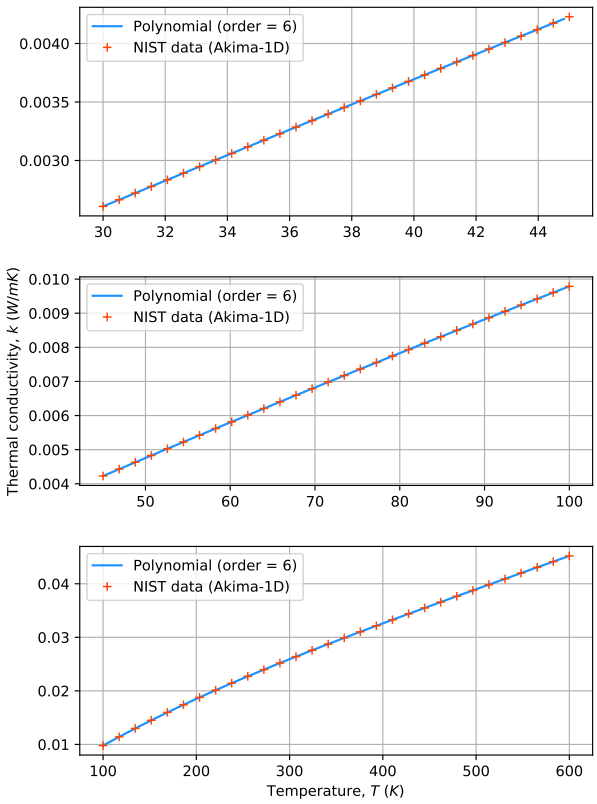
THERMOCHEMICAL PROPERTIES OF AIR AT LOW TEMPERATURES

Due to the low temperatures expected following the expansion of the flow into the test chamber (of the order of $40K$ according to isentropic calculations), it was expected that the standard thermochemical models available in FLUENT would be inadequate, and would likely produce inaccurate results. As such, specialised thermochemical relations were developed by extrapolating from NIST thermochemical data [138] for the full range of expected temperatures throughout the plume ($30 \leqslant T \leqslant 600$). This extrapolation was performed using an Akima 1D spline scheme [139], and a mass fraction weighted mixture of the NIST properties for Nitrogen, Oxygen, and Argon ($78\% N_2$, $21\% O_2$, $1\% Ar$). Polynomials were then fitted to these data such that they could be automatically queried during CFD simulations. The data produced by this extrapolation and the polynomial fits are presented in Figure A.1, below.

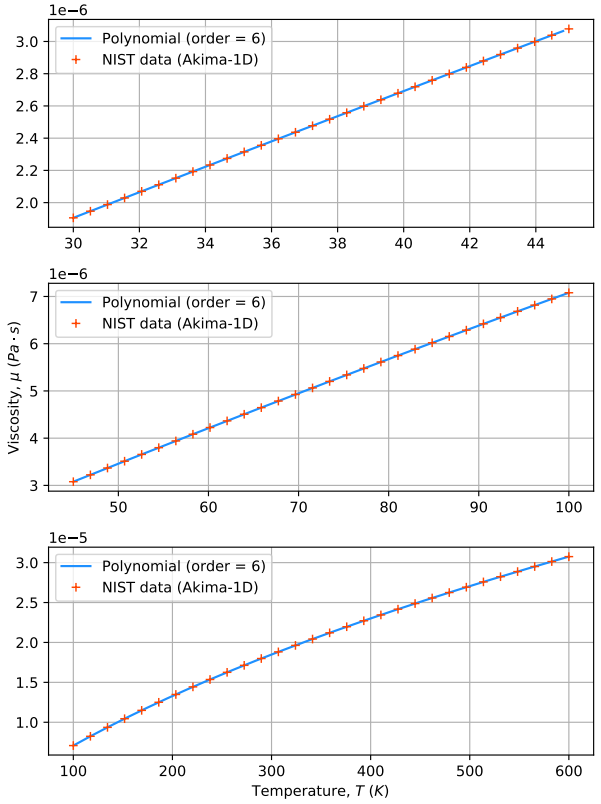


(a) Specific heat capcity, C_p

(b) Specific heat capcity, C_v



(c) Thermal conductivity, k



(d) Viscosity, μ

Figure A.1: Thermochemical polynomial fits produced using NIST reference data for air in the temperature ranges of $(30 \leqslant T < 45\ K)$, $(45 \leqslant T < 60\ K)$, and $(60 \leqslant T \leqslant 1000\ K)$

APPENDIX B

B

ANALYTIC LDWT PERFORMANCE MODEL

A basic analytical model of the LDWT's performance was devised based upon the parameters of the flow that an operator can freely change during the setup of an experiment. These variables are the stagnation temperature, T_0, and the inlet mass flow rate, $\dot{m}_1$. All other variables in the system are related to the pumping speed of the vacuum equipment, Q_{LDWT}, which cannot be adjusted nearly as easily due to the nature of the vapour booster technology. A simplified diagram of the LDWT and its main sections is presented in Figure B.1, followed by the derivation of the analytic model. The diagram also shows the primary quantities of interest to an operator, and their approximate measurement locations.

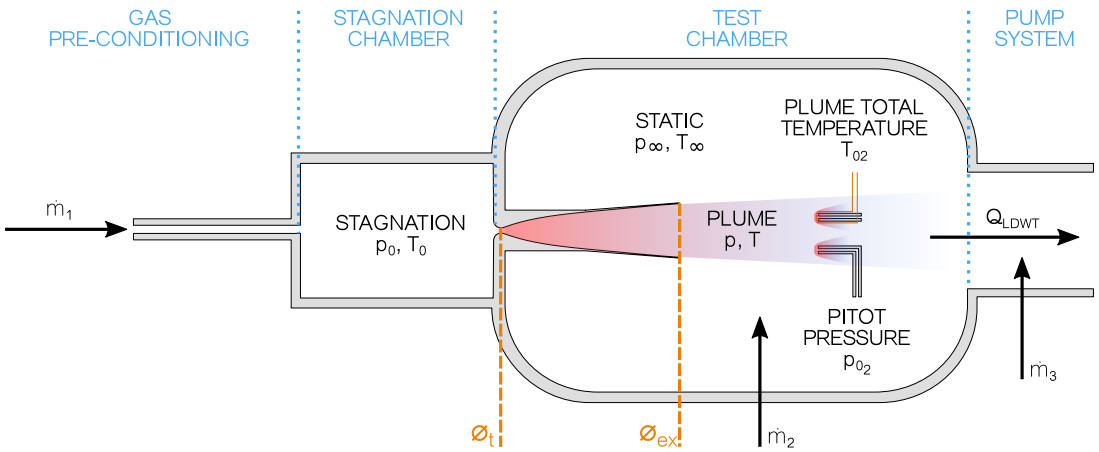


Figure B.1: Diagram of the LDWT showing the main variables of interest.

To begin, the mass flow rate of gas entering the pumping system is taken to be equal to the inlet mass flow rate, $\dot{m}_1$, plus a secondary mass flow rate introduced in order to represent leaks into the test chamber, $\dot{m}_2$. A third mass flow rate term, $\dot{m}_3$, is also introduced in order to allow for the consideration of leaks in the pumping system. This expression is then related to the system's pumping speed, Q_{LDWT}, i.e. the volume flow rate of the pumping system when the vapour booster is fully engaged. Note that the pumping speed of the facility does not vary when internal pressures are below $\sim 1\ Pa$, so a steady state value is assumed here for simplicity. No leaks were measured or modelled in the stagnation chamber, as the state of the gas is directly measured here prior to being expanded through the contoured nozzle. As such, the state of the

gas at this point in the system is known, and thus both leaks into the chamber, and the input gas flow are combined into $\dot{m}_1$.

$$\dot{m}_{total} = \dot{m}_1 + \dot{m}_2 + \dot{m}_3 = \rho\, Q_{LDWT} = \frac{p_\infty Q_{LDWT}}{RT_\infty} \tag{B.1}$$

Mach number, Ma, may now be introduced by substituting in expressions for the local speed of sound, $a = \sqrt{\gamma RT}$, and density, $\rho = \frac{p}{RT}$ (where R is the specific gas constant for the test gas, and γ is the ratio of its specific heats, $\frac{C_P}{C_V}$). This yields:

$$\dot{m}_1 = \rho A V = \left(\frac{p}{RT}\right) A\, Ma\, a = \frac{pA\, Ma\, \sqrt{\gamma RT}}{RT} = \frac{pA\, Ma\, \sqrt{\gamma}}{\sqrt{RT}} \tag{B.2}$$

In order to simplify this expression, the conditions at the nozzle throat may be considered. The ratio between static and stagnation pressure for an isentropic nozzle in terms of Mach number is:

$$\frac{p}{p_0} = \left(1 + \frac{\gamma - 1}{2} Ma^2\right)^{\frac{-\gamma}{\gamma - 1}} \tag{B.3}$$

Rearranging for p, and setting $Ma = 1$ and $\gamma = 1.4$ (conditions at the throat for a choked flow of air) yields the static pressure at the nozzle throat:

$$p_{th} = 0.5283\, p_0 \tag{B.4}$$

Combining equations (B.4) and (B.2), expanding out the expression for the nozzle's throat area, and setting them equal to equation (B.1) then yields:

$$\dot{m}_{total} = \frac{0.5283\, p_0 \pi {\varnothing_t}^2 \sqrt{\gamma}}{4\sqrt{RT_0}} + \dot{m}_2 + \dot{m}_3 = \frac{p_\infty Q_{LDWT}}{RT_\infty} \tag{B.5}$$

The mass flow rates, $\dot{m}_1$, $\dot{m}_2$, and $\dot{m}_3$, are known quantities, the former being set by the tunnel operator, and the latter two having been determined by measurements described in Section 4.2.8. Hence, expression (B.1) may be rearranged into expressions for the expected test chamber static pressure (otherwise known as the test chamber background pressure), p_∞, and the stagnation chamber pressure, p_0:

$$p_\infty = \frac{RT_\infty(\dot{m}_1 + \dot{m}_2 + \dot{m}_3)}{Q_{LDWT}} \tag{B.6}$$

$$p_0 = \frac{4\dot{m}_1\sqrt{RT_0}}{0.5283\pi {\varnothing_t}^2\sqrt{\gamma}} \tag{B.7}$$

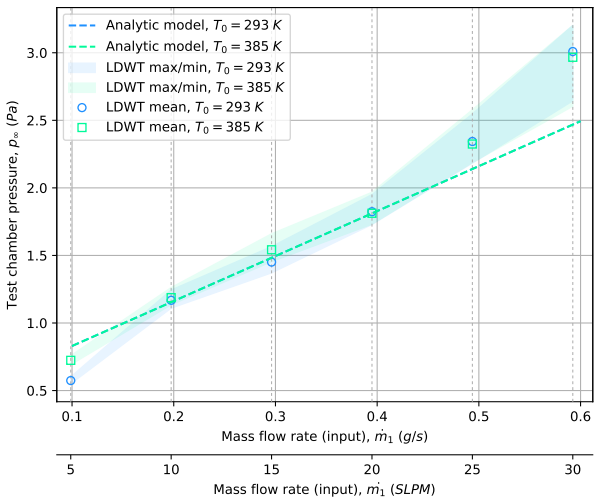
Combining equations (B.6), (B.7), and (B.3), and rearranging for Ma yields the expression for the nozzle exit Mach number, Ma_{ex}. This expression is intended to allow future users of the LDWT to quickly assess the expected Mach number for a given total temperature, T_0, and input mass flow rate, $\dot{m}_1$ (or indeed calculate what combination of these variables would produce a desired nozzle exit Mach number).

$$Ma_{ex} = \sqrt{\left[\left(\frac{0.1321\pi T_\infty {\varnothing_t}^2(\dot{m}_1+\dot{m}_2+\dot{m}_3)\sqrt{R\gamma}}{\dot{m}_1 Q_{LDWT}}\sqrt{T_0}\right)^{\frac{1-\gamma}{\gamma}} - 1\right]\left[\frac{2}{\gamma-1}\right]} \tag{B.8}$$

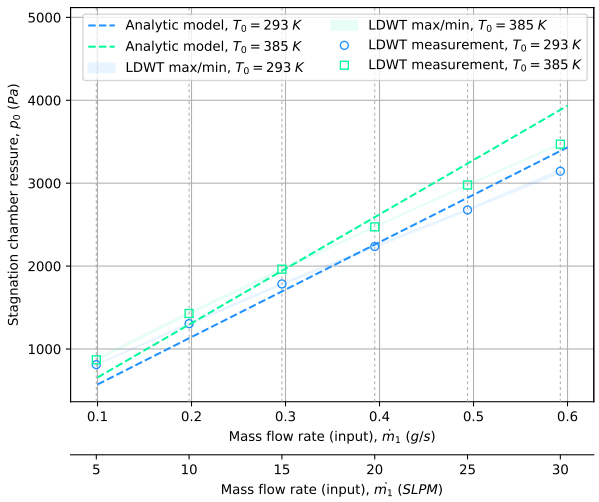
The results of these calculations are plotted against input mass flow rate in Figure B.2, where they are accompanied by measurements from pressure transducers P0, P2, and P5 (stagnation chamber, test chamber, and pitot pressures, respectively - see Figure 4.4). Predictions of Mach number, Ma, and Knudsen number, Kn, using the analytic model are presented in Figure B.3. Note that, for convenience, mass flow rates are stated in both g/s (for use in calculations), and SLPM (the unit used by the LDWT mass flow controller). Since the unit SLPM represents a volumetric flow rate at "standard" conditions, it is in fact a measure of mass flow rate. The conversion factor between these two measurements is: $1\,SLPM = 0.01975\,g/s$

As can be observed from Figures B.2(a) and B.2(b), the analytic model predicts the test section background, and stagnation chamber pressures (p_∞ and p_0, respectively) well. Inaccuracies in these predictions are evident towards the extremities of the range of examined mass flow rates. In the case of lower values of $\dot{m}_1$, it is expected that these discrepancies are caused by instabilities in the plume which vary over time, resulting in an inconsistent structure. This effect is examined in Section 4.3.7.

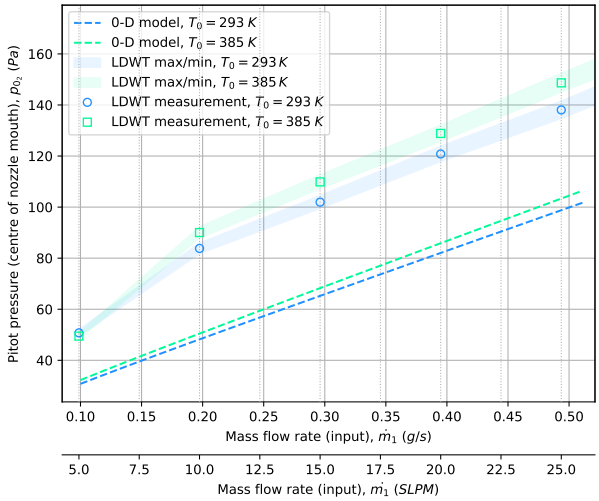
It was discovered empirically that the vapour booster's pumping action was fully impeded beyond an input mass flow rate, $\dot{m}_1$, of $\sim 32\,SLPM$. As such, the discrepancies in p_∞ and p_0 prediction with increasing $\dot{m}_1$ likely represent the onset of this phenomenon, which results in a gradual reduction in the effectiveness of the LDWT vacuum pumping system.



(a) Test chamber background pressure against mass flow rate

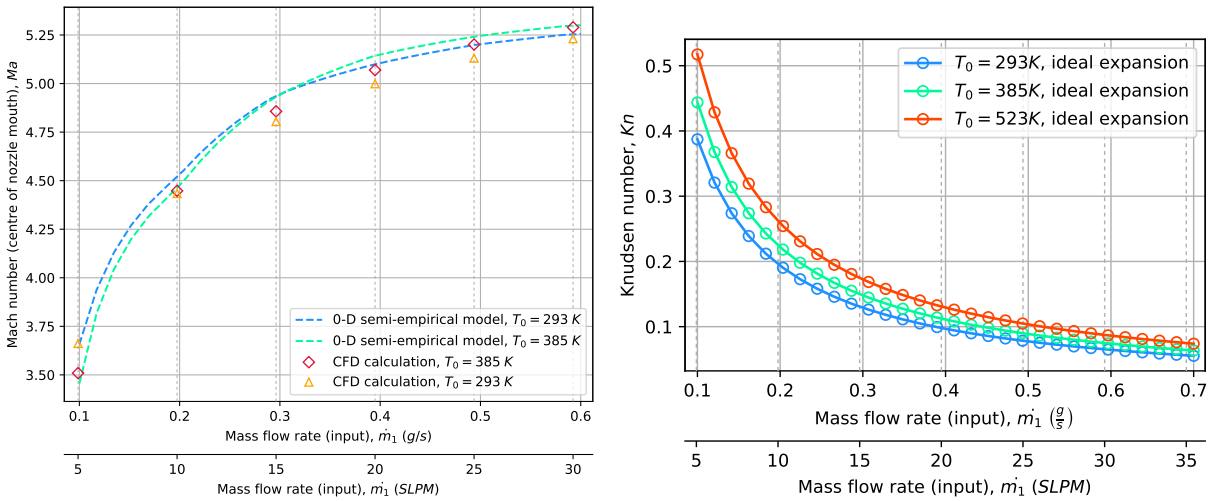


(b) Stagnation chamber pressure against mass flow rate



(c) Plume pitot pressure at the centre of the nozzle exit against mass flow rate

Figure B.2: Results of LDWT analytic model compared to pressure transducer readings.



(a) Mach number against stagnation temperature (b) Knudsen number against mass flow rate ($L = 0.01$m)

Figure B.3: LDWT modelling capabilities predicted using analytic model.

From Figure B.2(c), it is clear that the analytic model is capable of predicting the trend of pitot pressure, p_{0_2}, with $\dot{m}_1$ at the centre of the nozzle exit correctly. However, the model can also be seen to consistently underpredict the magnitude of these pressures. This may be explained by considering that the analytic model is formulated using isentropic expansion equations. As such, the model's predictions represent a perfectly expanded flow with no boundary layer. Owing to the rarefied nature of the flow, a large boundary layer is expected to form inside the nozzle. It is also expected that ideal expansion of the plume (whereby the exit pressure of the nozzle would exactly equal the background pressure of the LDWT test chamber) will not occur. As such, the predictions of Ma and Kn presented in Figure B.3 are expected to have similar characteristics, i.e. correct trend prediction with incorrect magnitudes. From these trends, it may be inferred that varying the inlet mass flow rate, $\dot{m}_1$, and total temperature, T_0, would have a far more significant effect on the Knudsen number, Kn, of the flow than on the Mach number, Ma, which would consistently remain above the hypersonic limit.

In conclusion, the analytical expressions presented above accurately describe the measured dependence of stagnation pressure, p_0, and test chamber pressure, p_∞, on inlet mass flow rate; however, they underpredict the test chamber pitot pressure by approximately 50%. To further explain this, it is necessary to examine spatial variation within the flow in order to compensate for boundary layer growth and include this in the analytic model.

APPENDIX C

C

SEMI-EMPIRICAL LDWT PERFORMANCE MODEL

The analytic model used to initially characterise the performance of the LDWT (see Appendix B) was based upon isentropic flow theory. As such, it was able to model the centre region of the flow at the nozzle exit, but could not directly predict boundary layer thickness or plume structure (nor, consequently, the effects of these phenomena on the flow). These calculations (derived in Appendix B) were compared with the results of the no-slip CFD simulations and the experimental Pitot surveys. These data are presented in Figures C.1 and C.2.

The original analytic model was found to predict the test section pressure, p_∞, and stagnation chamber pressure, p_0, well, as is shown in Figures B.2(a) and B.2(b), respectively (see Section B). Discrepancies in the prediction of p_∞ with increasing $\dot{m}_1$ are thought to represent the reduction in effectiveness of the vapour booster's pumping action, which was fully impeded by inlet mass flow rates above 32 SLPM (this was marked by a sharp increase in p_∞).

Due to the boundary conditions selected for the CFD simulations (specification of static pressure at the exit and stagnation pressure at the inlet), it is necessary to calculate the actual mass flow rate being simulated. This is accomplished by extracting profiles of density, ρ, and flow velocity, U, at the nozzle exit and multiplying them together to obtain the mass flux profile. These data are then integrated with respect to the annular area of the nozzle and summed to find the CFD mass flow rate (see Eq. (C.1)). A comparison between this calculated mass flow rate and that which was defined during experiments using the LDWT mass flow controller is presented in Figure C.1(d). It can be seen that both the simulations and experiments agree well when $15 \lessapprox \dot{m}_1 \lessapprox 20\ SLPM$, but diverge beyond that value. This serves to confirm the conclusion made in Section 4.3.4.4; namely that the CFD simulations over-predict the thickness of the boundary layer at higher mass flow rates, thus reducing the effective nozzle cross-sectional area being simulated (and hence the mass flow rate, since pressure boundaries were utilised in the CFD simulations). Despite this apparent over-prediction of the boundary layer effects, however, the CFD predictions of p_∞ and p_0, (and hence p_{0_2} and Ma) have been shown to agree

fairly well with experimental data.

$$\dot{m} = \int_0^r U_r \rho_r 2\pi r.dr = \sum_{n=1}^{N} U_n \rho_n \pi \left(r_n^2 - r_{n-1}^2\right) \tag{C.1}$$

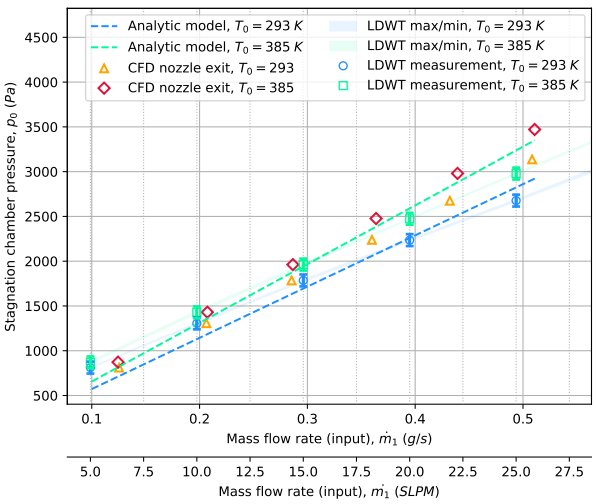
While p_∞ and p_0 are predicted with reasonable accuracy by the analytic model, it is evident that it is far less effective at predicting Pitot pressure, p_{0_2}, at the nozzle exit, as is shown in Figures C.1(b) and C.2(a). This is likely due to the inability of the analytic model to take the nozzle's large boundary layer into account.

In order to remedy this inaccuracy, a semi-empirical model based upon both the experimental Pitot data and CFD predictions was formulated using 2D linear interpolation of $\dot{m}_1$ and T_0. As such, the post-shock total pressure, p_{0_2}, at the Pitot was defined using experimental data (see Eq. (C.2)). The interpolations of p_{0_2} based upon data from the LDWT are shown in Figures C.1(b) and C.2(a).

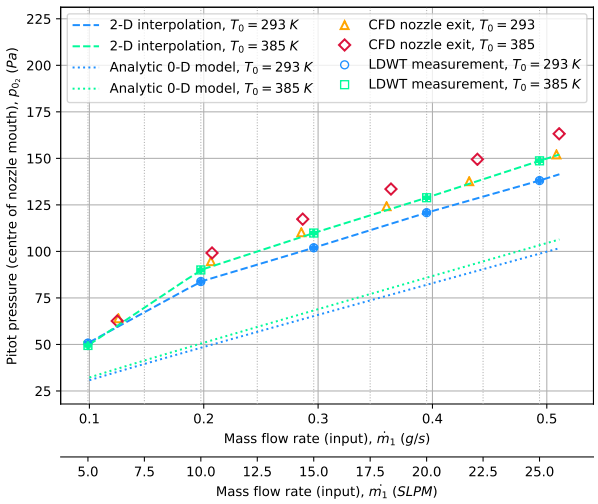
$$p_{0_2} = f(\dot{m}_1, T_0)_{LDWT} \tag{C.2}$$

$$p_0 = f(\dot{m}_1, T_0)_{CFD} \tag{C.3}$$

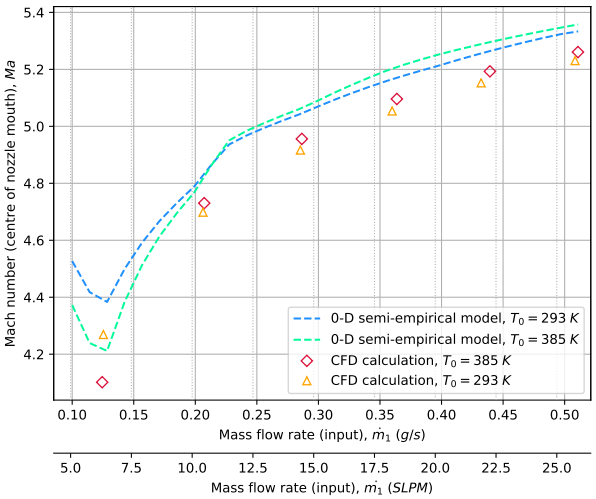
The prediction of Mach number at the nozzle exit, Ma_{ex}, was also a desired output of the semi-empirical model. Using the semi-empirical p_{0_2} correlation described above in conjunction with Equation (4.1) was expected to be the most expedient method to obtain these data. However, it was not possible to directly measure the total pressure upstream of the normal shock, p_{0_1}, so a second linear interpolation model was developed based upon the results of the no-slip CFD simulations (see Eq. (C.3)). This was then utilised in the aforementioned manner to predict the Mach number at the nozzle exit. The results of this prediction are shown in Figures C.1(c) and C.2(b).



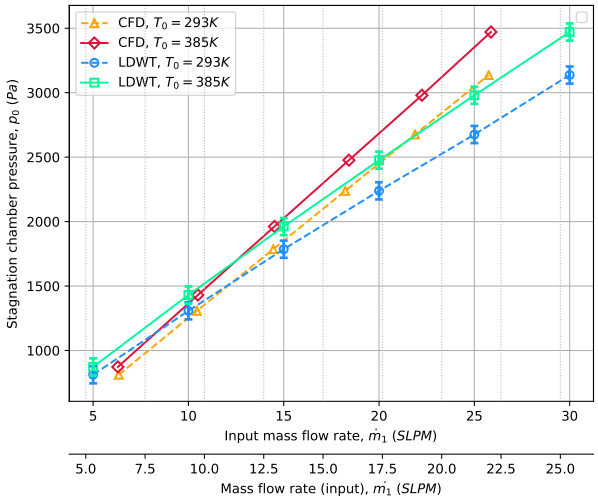
(a) Stagnation chamber pressure, p_0



(b) Nozzle exit Pitot pressure, p_{0_2}

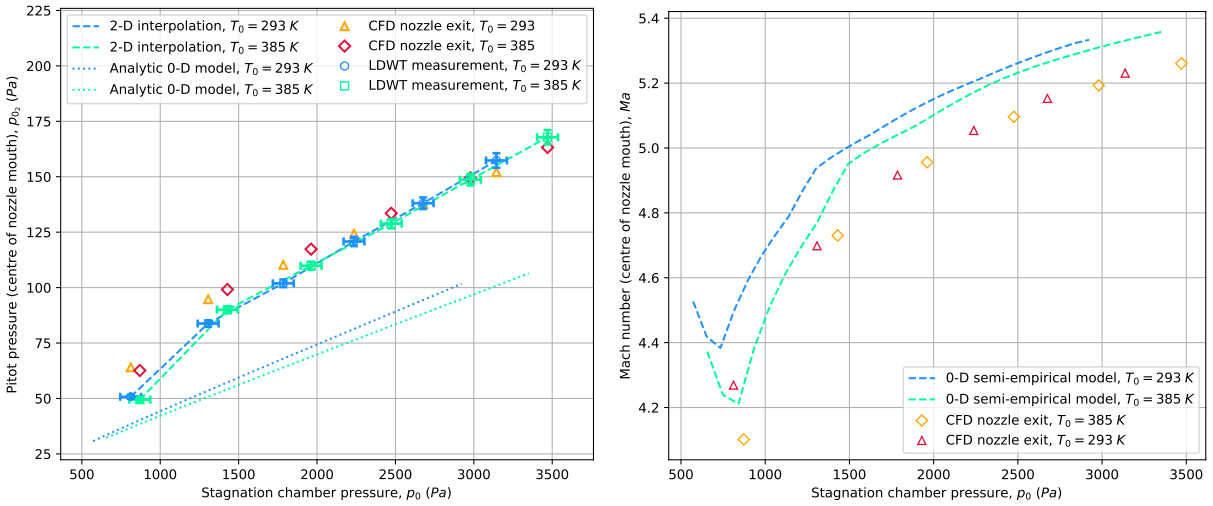


(c) Nozzle exit Mach number, Ma



(d) Stagnation pressure plotted against both inlet mass flow rate, $\dot{m}_1$, (for LDWT data) and integrated mass flow rates, $\dot{m}$, at the nozzle exit (for CFD data).

Figure C.1: Comparison between analytic model predictions, CFD calculations, and measured quantities from the LDWT with varying inlet mass flow rate, $\dot{m}_1$, and stagnation chamber temperature, T_0.



(a) Pitot pressure, p_{0_2}

(b) Mach number, Ma

Figure C.2: Comparison between analytic model predictions, CFD calculations, and measured quantities from the LDWT at the centre of the nozzle exit, with varying stagnation chamber temperature, T_0.

The trend of Ma_{ex} against $\dot{m}_1$ can be seen to agree very well with CFD predictions, with the proximity between the two examined total temperatures, T_0, being reproduced with the greatest accuracy. The crossing of the two trends (at $\sim 0.22\,g/s$ in Figure C.1(c)) is also reproduced well, although it occurs at a higher mass flow rate than predicted by the CFD simulations alone. An increase in Mach number is predicted at low mass flow rates, although this is expected to be a non-physical effect, and a byproduct of extrapolation in the interpolation algorithm used.

Overall, this method is expected to produce slightly more accurate predictions over the CFD simulations alone, with the maximum difference between the resultant data being $\sim 2.8\%$. Although this method is much more convenient for quickly assessing the capabilities of the LDWT at a given mass flow rate and total temperature, the minor improvement in accuracy was not deemed sufficient to apply this method to the other CFD simulations described above.

APPENDIX D

NOZZLE THROAT DIAMETER MEASUREMENT

Of particular concern during the process of defining the CFD domain was the nozzle throat diameter, $\varnothing_{th}$, which had to be measured using progressively larger steels rods. From isentropic nozzle theory, it is known that the ratio between the throat and exit areas of a nozzle has a significant impact on the exit Mach number of the nozzle. As such, it was expected that this dimension would likely have a large impact on the accuracy of simulations. In order to characterise this uncertainty, an initial set of two simulations was performed using domains with throat diameters of $9.60\ mm$ and $9.77\ mm$ (the boundary conditions used were those of the $25\ SLPM$, $385\ K$ case in Table 4.2). These were the sizes found to border the actual throat diameter during measurement (the smaller of the two was able to pass cleanly through the throat, while the larger was not). The results of these simulations were then compared with the relevant Pitot data (see Figure D.1).

Marked divergence between the simulations and experimental data can be observed towards the nozzle exit, and from $\sim 90\ mm$ downstream. In the case of the nozzle exit, this is likely due to incorrect prediction of the nozzle boundary layer thickness. Further downstream, it was expected that a Mach disk would form due to underexpansion of the plume, and that the CFD would likely under or overpredict its location owing to the high gradients and minor uncertainties in the measurement of the nozzle contour. It was also considered that the wall boundary condition model (no-slip) could have been a factor in this inaccuracy, and this is explored in Section 4.3.4.2).

Aside from these slight inaccuracies, it can be seen that there is a region of relatively stable flow in the centre of the surveyed region. The CFD simulations border the experimental data as expected, indicating that a throat diameter approximately halfway between the simulation dimensions would be most appropriate. The mean difference between those sections of the experimental and simulation results which showed the best agreement was subsequently evaluated, and a throat diameter, $\varnothing_{th}$, of $9.65\ mm$ selected for all subsequent CFD simulations.

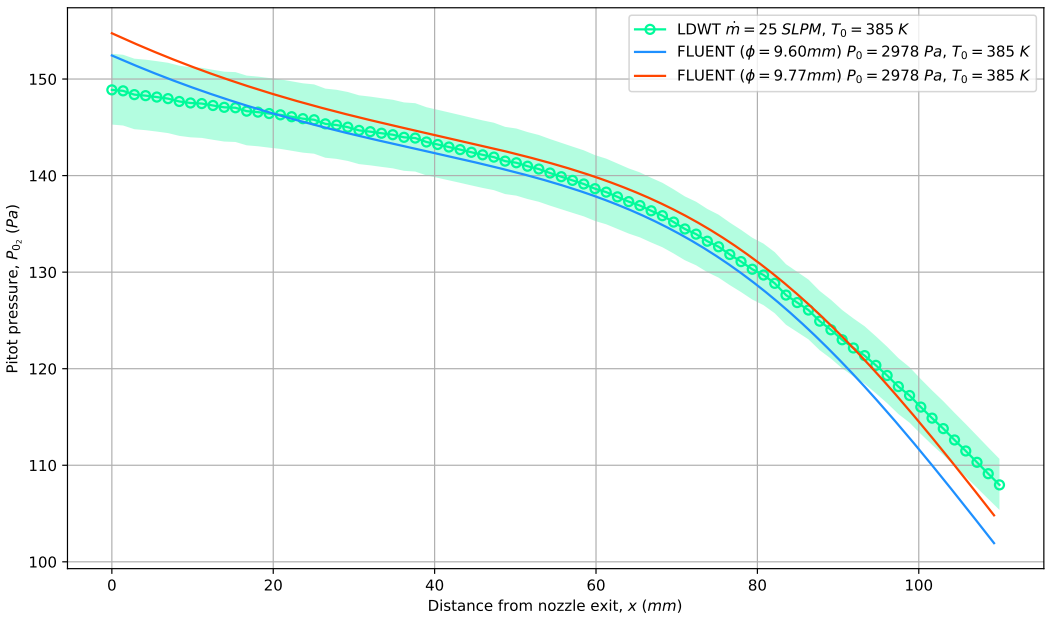


Figure D.1: Comparison of axial Pitot survey readings downstream of the nozzle exit and CFD results for varying nozzle throat diameter, $\varnothing_{th}$. The variable plotted is the post-shock total pressure, p_{0_2}, and the shaded area is the uncertainty in the experimental data.

APPENDIX E

ORIGINAL LDWT SPECIFICATIONS

This appendix comprises scanned copies of original documentation for the LDWT and the individual components of its pumping system.

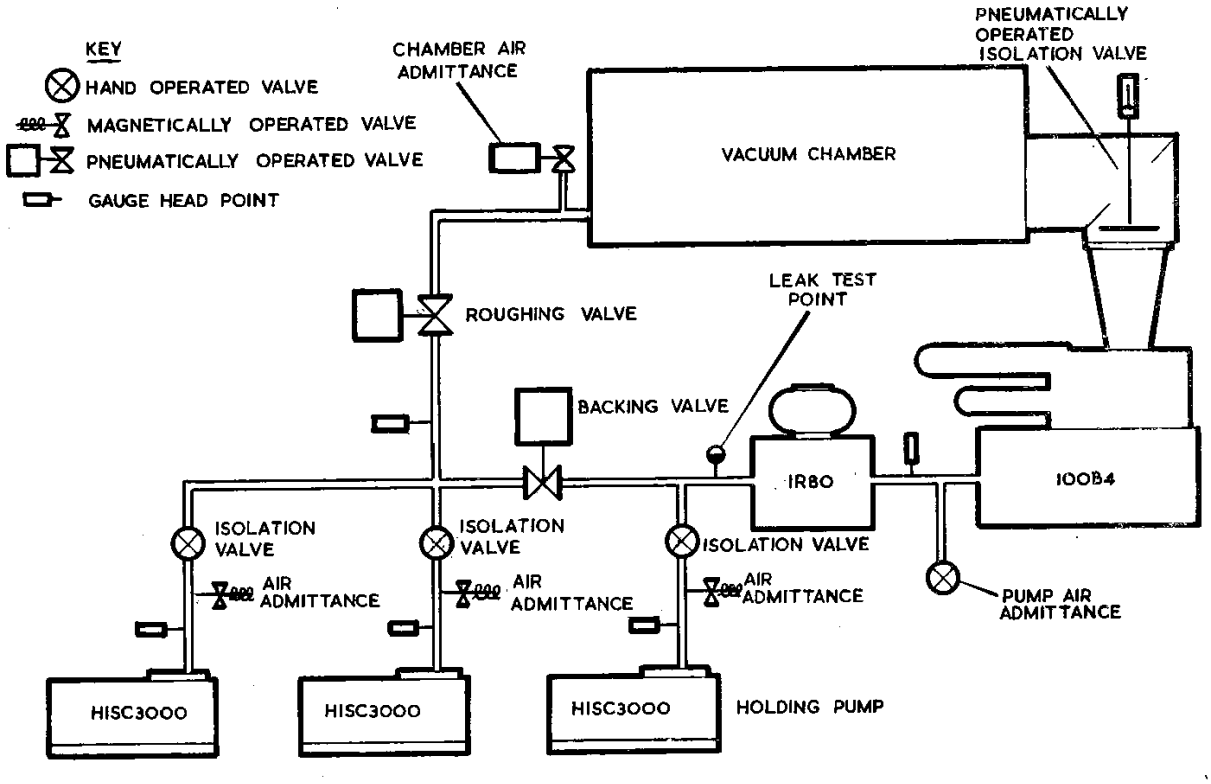


Figure E.1: System level block diagram of the LDWT pumping plant produced by Edwards

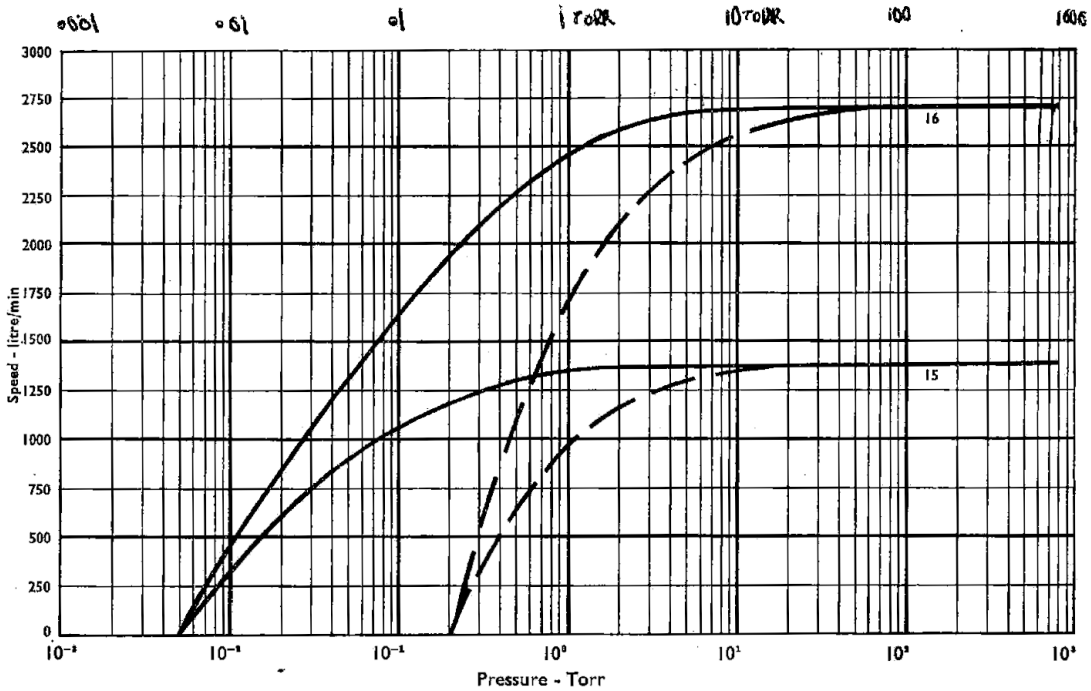
SPECIFICATION		1SC1500	1SC3000	H1SC1500	H1SC3000	ES7500
* Ultimate vacuum without gas ballast*, better than	(torr)	0.01	0.01	0.01	0.01	0.01
* Ultimate vacuum with auxiliary gas ballast, better than	(torr)	0.02	0.02	0.02	0.02	0.02
* Ultimate vacuum with full gas ballast, better than	(torr)	0.5	0.5	0.5	0.5	0.5
Normal speed of rotation	(rpm)	350	350	490	490	350
Displacement (swept volume) at normal speed of rotation	(ft^3/min)	56.5	113	71	145	275
	(litres/min)	1600	3200	2000	4100	7800
	(m^3/hr)	96	192	120	246	470
Max. gas ballast flow	(ft^3/min)	5	10	5	10	18-20
	(litres/min)	141	283	141	283	510-560
	(m^3/hr)	8.5	17	8.5	17	30-33
Approx. max. safe inlet pressure of water vapour at full gas ballast	(torr)	70	70	70	70	56
Safe water vapour pumping rate	(lb/hr)	12	24	15	28	46
Normal running temp. without gas ballast at ambient 20°C	(°C)	50-60	50-60	55-65	55-65	50-60
*Normal running temp. with gas ballast	(°C)	80-90	80-90	85-95	85-95	80-90
Normal oil charge	(Imp. galls)	1.5	3.5	1.5	3.5	7
	(litres)	6.8	16	6.8	16	30
Recommended grade of Edwards oil		17 or 15	17 or 15	17 or 15	17 or 15	17 or 15
Horse power of motor		3	5	5	7.5	15
Min. cooling water flow at 20°C	(litres/min)	1	2	2.5	2	2.3

*Air is a suitable ballast for most operation

*ULTIMATE vacuum measured with McLeod gauge.

*Running temp. should not exceed 110°C on all pumps.
See page 2 figure 1. Point A.

Table E.1: Specifications of various Edwards pumps available in 1963. Note that the LDWT utilises a trio of H1SC3000 model roughing pumps.



Typical speed curves for rotary piston vacuum pumps: 16 = 1SC3000; 15 = 1SC1500.
with ballast — — —

Figure E.2: Plot of pressure versus pumping speed for Edwards 1SC3000 and 1SC1500 series roughing pumps (both with and without gas ballasts). Note that the LDWT roughing pumps are variants of the 1SC3000 model and do not feature gas ballasts.

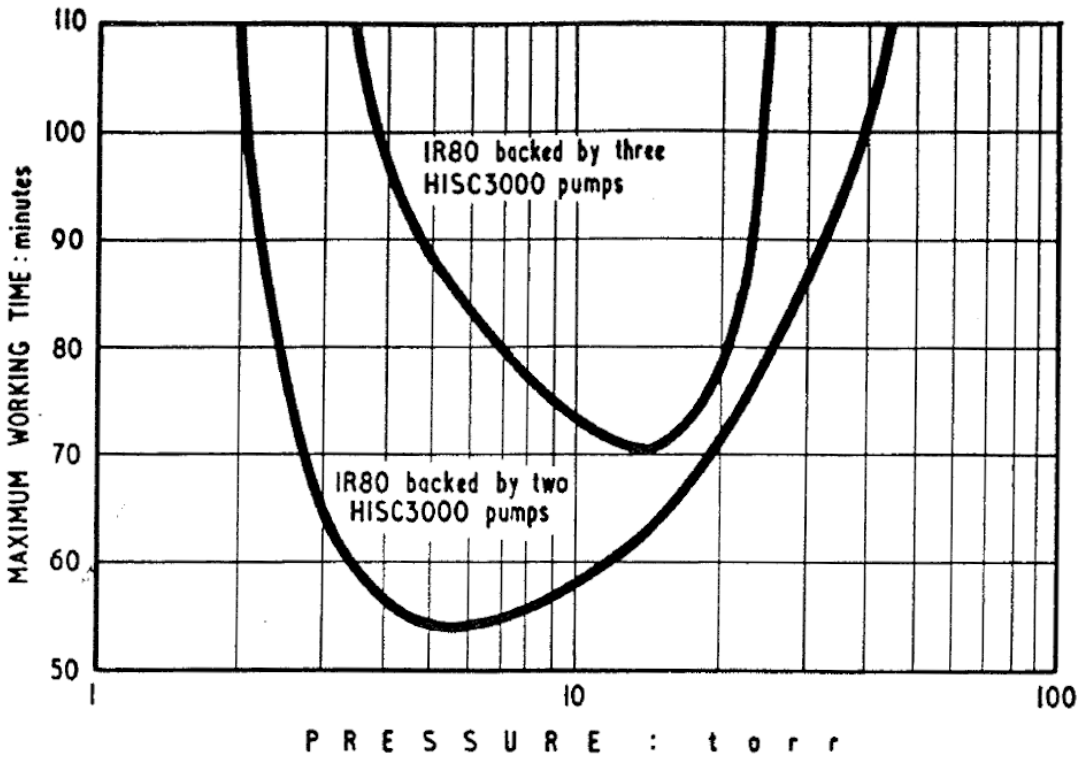


Figure E.3: Plot of pressure versus working time for Edwards 1R80 Roots type mechanical booster pump. Curves for backing of the 1R80 by both two and three H1SC3000 roughing pumps are shown.

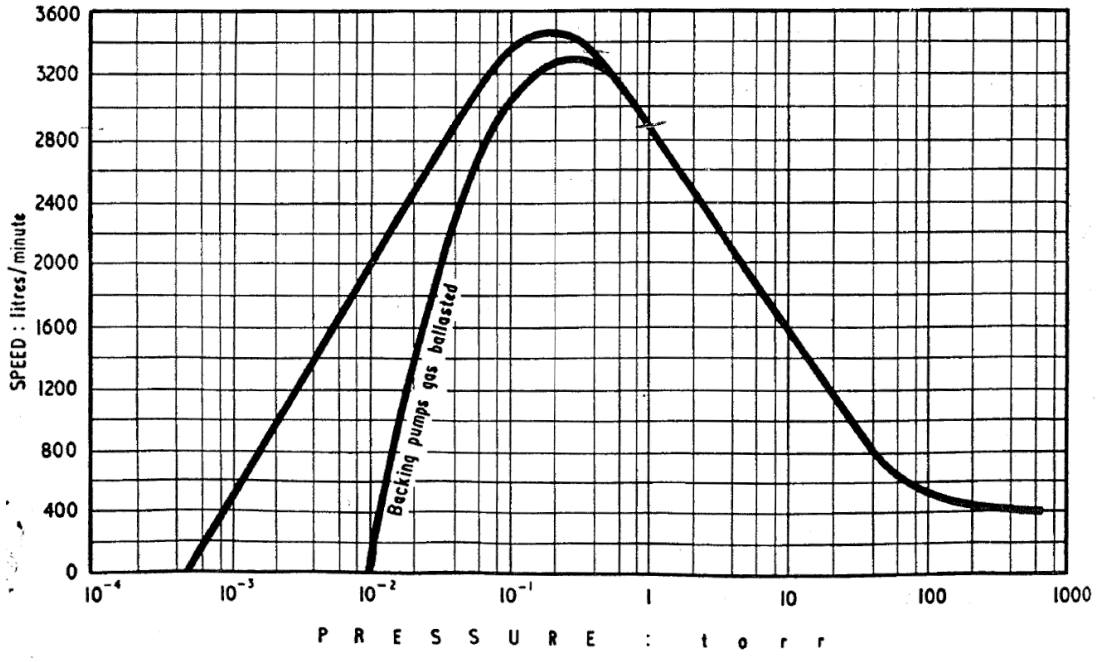
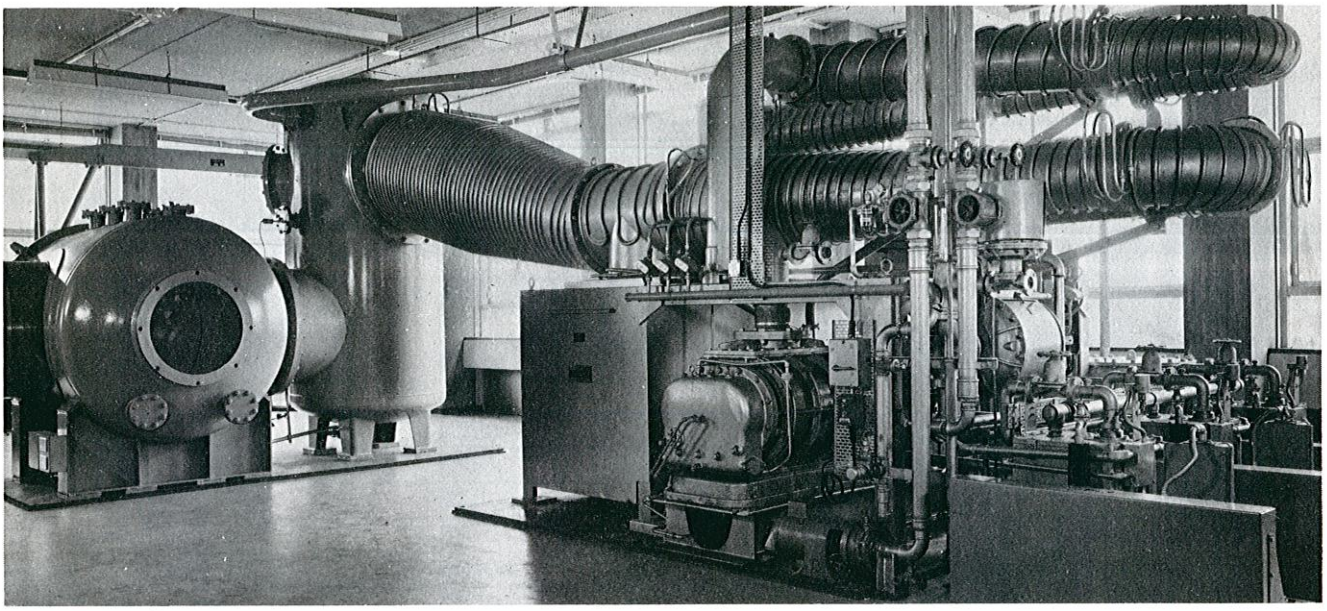


Figure E.4: Plot of pressure versus pumping speed for Edwards 1R80 Roots type mechanical booster pump. Curves for backing of the 1R80 roughing pumps both with and without gas ballasts are shown.

Booster pumps



Edwards Model 100B4 vapour booster installed at Oxford University with the top body cone swung through 90° to clear low headroom.

Data

MODEL		9B1	9B3	9B4	18B3A	18B4A	30B4A	30B5A	100B4
Total ultimate vacuum torr, using Edwards booster pump fluid		Better than 10^{-4}							
Peak speed for air, unbaffled	litres/second	40	850	2800	1300	2800	4500	10000	22000
Peak speed for hydrogen, unbaffled	litres/second	120	—	4200	2400	6000	9000	15000	—
Limiting backing pressure‡	torr	2–5	2–3	2–2.5	2–3	2–3	5	4–6	1.5–2
Recommended backing pump displacement*	litres/min	450	450	450	2400	2400	5000	8000	65000
Recommended backing pump type		1SC450	1SC450	1SC450	1SC3000	1SC3000	ES7500	ES7500	1R80+ 3×H1SC3000
Backing connexion, Edwards vacuum union or flange		1½ in (38.1 mm)	1½ in (38.1 mm)	1½ in (38.1 mm)	2 in (50.8 mm)	2 in (50.8 mm)	6 in (152 mm)	6 in (152 mm)	special flange
Recommended operating fluid		Edwards booster pump fluid							
Operating fluid charge	gallons litres	⅓ 1.5	⅓ 1.5	⅓ 1.5	2.2 10	2.2 10	12.0 55	12.0 55	70 320
Heater loading (maximum)	kW	1.85	2.25	2.25	6.3	6.3	22.5	22.5	90
Warming-up time for full performance at maximum heater input		25 min	25 min	25 min	60 min**	60 min**	60 min	60 min	60 min
Cooling water connexions, standard water unions		⅜ in BSP	⅜ in BSP	⅜ in BSP	⅜ in BSP	⅜ in BSP	1 in BSP	1 in BSP	1½ in BSP
Minimum water flow†	litres/min	2–3	2–3	2–3	7	7	25	25	90
Overall height of pump	in mm	32 813	45¼ 1149	57⅜ 1457	62⅞ 1597	74 1880	105 2667	128 3251	164 4166

**The displacement specified is the minimum necessary for the efficient operation of the pump. If rapid cycles on large volumes are required, a larger backing pump may be desirable.*

***Pump begins to operate after 40 minutes.*

†The outlet water temperature should be kept below 35°C.

‡The limiting backing pressures of the pumps vary as the ultimate pressure is approached. The higher figure in the table applies at high and the lower at low mass flows.

Edwards research leads to continuous improvement; details of products supplied may therefore vary from those described.

EDWARDS HIGH VACUUM LIMITED
A member of Edwards International Group
MANOR ROYAL CRAWLEY SUSSEX ENGLAND
Telephone: Crawley 28844 Telex: 87123 Edhivac Crawley

Printed in England GD WLP 867

© 1967 Edwards High Vacuum Ltd.

Vapour booster pumps—22000 l/s

Model 100B4

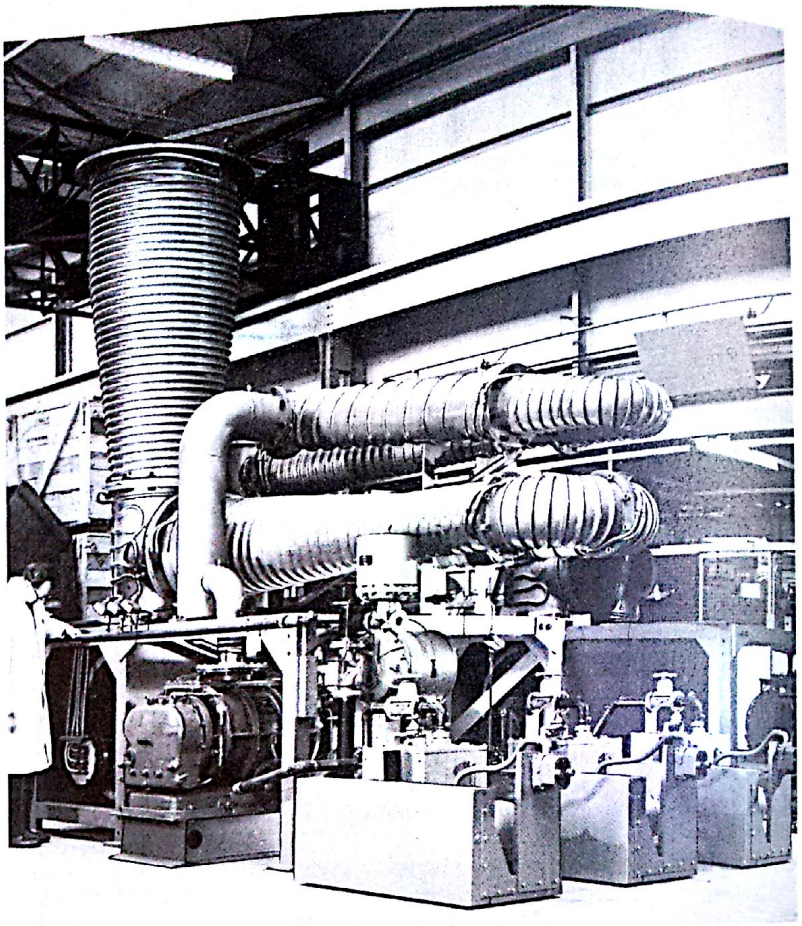
Pumping speed (air)	22000 l/s
(hydrogen)	—
Ultimate vacuum	10^{-4} torr
Limiting backing pressure	1.2 to 1.5 torr
Recommended backing pump displacement	65000 l/min
Recommended backing pump	ER2000+ER600+ ES4000
Fluid capacity	320 litre
Heater loading (maximum)	90kW
Warming-up time for full performance at max heater input	60 min
Minimum water flow	90 l/min
Water connexion	1½in BSP
Weight	6800 lb (3084kg)

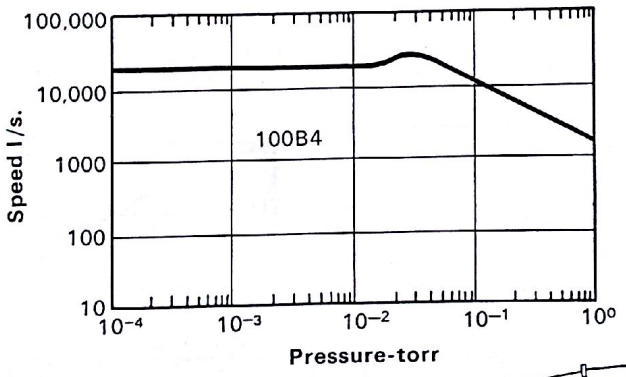
Product description	**Ordering number**
Pump 340–380V 3 ph 50/60Hz 4 wire	**05-B067-01-380**
400–440V 3 ph 50/60Hz 4 wire	**05-B067-01-440**

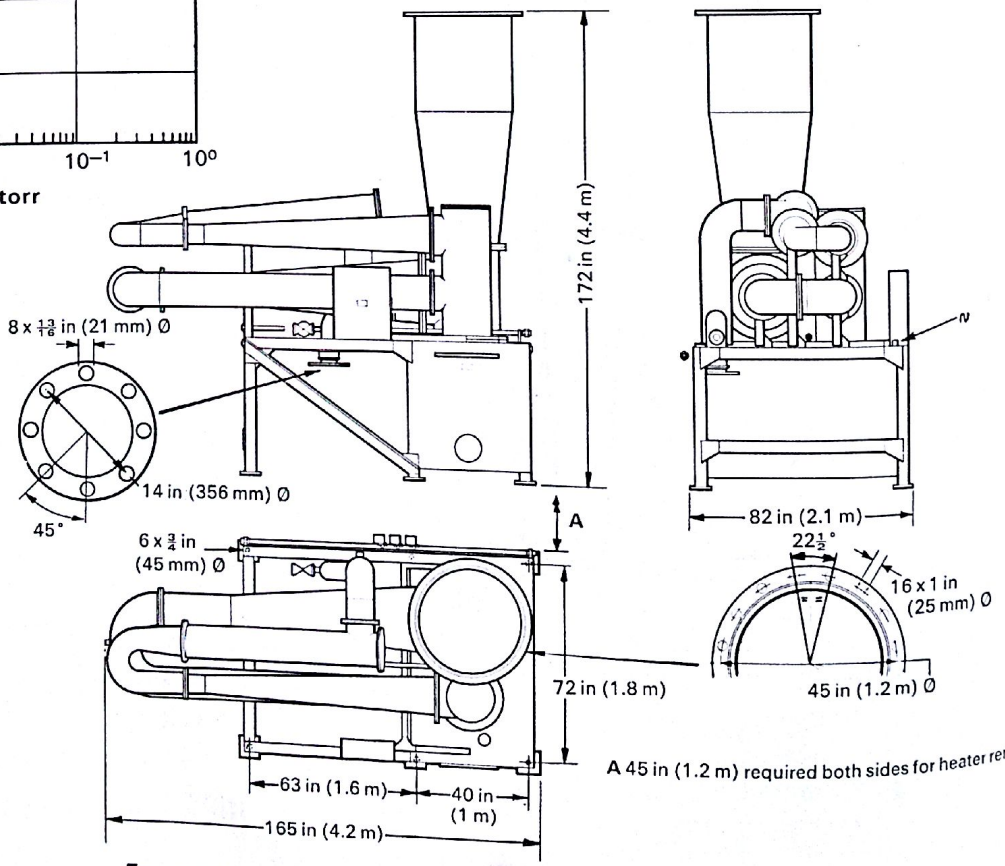
Spares

Kit comprising seals, fuses, etc.	**05-B067-01-801**
Heater (Tubalox bank type)—state voltage	**Dev 1512-10A**
Fluid AP201 22.7 litre	**09-H026-00-052**

For full range of vapour booster pump fluids, see fluids and sealants section.







For new products and prices current at time of ordering, see latest price list

Generated by CamScanner from intsig.com