THE UNIVERSITY OF QUEENSLAND

AUSTRALIA

Mixing and Combustion Enhancement in a Mach 12 Shape-Transitioning Scramjet Engine

James Eric Barth
BE (Mechanical Engineering) with Great Distinction
M.Sc (Mechanical Engineering)
M.Sc (Space Studies) *cum laude*

A thesis submitted for the degree of Doctor of Philosophy at
The University of Queensland in 2014
School of Mechanical and Mining Engineering,
Centre for Hypersonics

Abstract

Despite being an active topic of research for over 50 years, scramjet technology has only recently matured to a point where flight tests are being successfully carried out at the lower end of the hypersonic regime. While this progress is encouraging, a renewed interest in low-cost, reliable, and environmentally responsible access to space has identified scramjets capable of accelerating to speeds as high as Mach 12 as desirable. One class of scramjets thought to be capable of hypervelocity performance are those that employ three-dimensional streamtraced compression inlets to efficiently compress captured air. One promising example of this type of scramjet is the Mach 12 Rectangular-to-Elliptical Shape-Transitioning (REST) engine. The aims of this study are to investigate and characterize the flow physics behind the Mach 12 REST engine's current performance, and then attempt to improve its combustion performance by tailoring the engine's fuel injection to its internal flow field without otherwise modifying the engine geometry. To meet these aims, the engine was studied both numerically and experimentally.

The first-ever combusting simulations of a REST scramjet operating at Mach 12 conditions were performed for the Mach 12 REST engine using the CFD research code US3D. The simulations covered a range of conditions, including: unfuelled engine flow, inlet-fuelled flow, and various combined inlet/combustor fuelling configurations. The simulations were found to match well with the experiments they were designed to reproduce and be compared against. A comparison of simulations with experimental inflow conditions and their equivalent flight conditions on an otherwise identical engine showed that experiments in the T4 Stalker tube reproduce engine pressure and heat flux distributions well. The tunnel condition tends to capture less incoming flow than the engine at flight conditions, which leads to the ground-tested engine over-predicting the engine equivalence ratio.

The Mach 12 REST inlet was found to produce a thick "bubble-shaped" boundary layer along its bodyside compression surface, due to the compression effects of the inlet sidewalls acting on a thick turbulent boundary layer ingested from the vehicle forebody. This thick boundary layer forces the majority of inlet-captured air into a high-density, high Mach number flow region along the engine cowlside wall . The inlet also produces a symmetric pair of high-temperature swept separations that enter the engine isolator along the sidewalls of the engine. When fuel is injected from the bodyside surface of the inlet, it remains trapped inside the thick, turbulent boundary layer, where it becomes well-mixed and begins to burn just upstream of the inlet throat.

As much as 50% of this fuel is burned by the time it enters the engine isolator, while its injection and burning increases the inlet's drag by less than 5%. This burning bodyside flow region thermally compresses the remaining air flow within the engine isolator, and provides a source of heat and combustion radicals for the ignition of fuel injected further downstream. Overall combustion efficiency of inlet-injected fuel at the engine exhaust plane was found to be nearly 80% at high equivalence ratios.

Flow within the Mach 12 REST combustor is strongly shock-dominated. This is caused by both the cowl closure shock train transmitted from the inlet, and a strong recompression shock generated at the start of the combustor. This recompression shock is generated by the flow passing over a backward step at the entrance to the combustor, and is reinforced on the cowlside of the engine by compression caused by the engine flow path turning to realign with the nominal direction of flight. Fuel injected from the face of the combustor step was combined with inlet injection in an attempt to reduce skin friction drag through boundary layer combustion. This was found to be ineffective: the individual fuel jets never coalesced into a single continuous flow structure, and the fuel layer was quickly disrupted by the pressure gradients induced by the engine turn. The step-injected fuel was quickly ignited by the bodyside combustion region spreading circumferentially around the combustor wall, and its final combustion efficiency was found to be 82.9%.

Boundary layer injection from the step was replaced with tailored fuel injection, in which fuel was injected directly into the cowlside core flow, and the swept separation regions along the engine sidewall. This was found to improve combustion efficiency to 84.9%, verifying the validity of the tailoring approach. The tailored-injection Mach 12 REST engine was found to produce the same level of thrust as the boundary layer injection case despite having a lower fuel mass flow rate. The uninstalled specific impulse of the tailored injection engine was 5% higher than that of the step-injection configuration. Improvements to the tailored fuel injection geometry, combined with small modifications to the engine's geometry and operation may be sufficient to boost thrust and reduce internal viscous drag enough to allow a flight-model scramjet to achieve net thrust.

Declaration by Author

This thesis is composed of my original work, and contains no material previously published or written by another person except where due reference has been made in the text. I have clearly stated the contribution by others to jointly-authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my research higher degree candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the General Award Rules of The University of Queensland, immediately made available for research and study in accordance with the *Copyright Act 1968*.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis.

James Barth

Publications During Candidature

Journal Publications

Barth, J.E., Wheatley, V., & Smart, M.K., "Hypersonic Turbulent Boundary-Layer Fuel Injection and Combustion: Skin-Friction Reduction Mechanisms," *AIAA Journal*, Vol. 51 No. 9, 2013, pp. 2147-2157. DOI: 10.2514/1.J052041.

Conference Papers

Barth, J.E., Wheatley, V., & Smart, M.K., "Inlet Fuel Injection in a Mach 12 Shape-Transitioning Scramjet," presented to the *AIAA SciTech 52nd Aerospace Sciences Meeting*, National Harbor MD, 13-17 January 2014. AIAA 2014-1159.

Barth, J.E., Wheatley, V., Smart, M.K., Petty D.J., & Basore, K.D., "Flow Physics Inside a Shape-Transitioning Scramjet Engine," presented to the *18th AIAA/3AF International Space Planes and Hypersonic Systems and Technologies Conference*, Tours, France, 24-28 September 2012. AIAA 2012-5888.

Barth, J.E., Wheatley, V., & Smart, M.K., "Effects of Ethylene Combustion in a Hypersonic Turbulent Boundary Layer," presented to the *11th Australian Space Science Conference*, Canberra Australia, 26-29 September 2011.

Conference Posters

Barth, J.E., Wheatley, V., & Smart, M.K., "Streamwise Porthole Fuel Injection for Boundary-Layer Combustion Inside a Scramjet Engine," presented to the *28th International Congress of the Aeronautical Sciences*, Brisbane QLD Australia, 23-28 September 2012. Paper ICAS2012-P4.16.

Publications Included in this Thesis

Barth, J.E., Wheatley, V., Smart, M.K., Petty D.J., & Basore, K.D., "Flow Physics Inside a Shape-Transitioning Scramjet Engine," presented to the *18th AIAA/3AF International Space Planes and Hypersonic Systems and Technologies Conference*, Tours, France, 24-28 September 2012. AIAA 2012-5888.

Parts of this paper are included in Chapter 4. Barth was responsible for 50% of the simulation work, 90% of the analysis and interpretation of data, and 90% of the drafting and writing. Wheatley and Smart were responsible for 10% of the analysis and interpretation of data, and 10% of the drafting and writing. Petty and Basore each responsible for 25% of the simulation work, as Barth was not allowed to access the flow solver utilized in this study.

Barth, J.E., Wheatley, V., & Smart, M.K., "Hypersonic Turbulent Boundary-Layer Fuel Injection and Combustion: Skin-Friction Reduction Mechanisms," *AIAA Journal*, Vol. 51 No. 9, 2013, pp. 2147-2157. DOI: 10.2514/1.J052041.

In addition to being referenced in Chapter 2, parts of this paper appear in Chapter 6, and the complete paper is included as Appendix G. Barth was responsible for 100% of the analysis, 90% of the interpretation of data, and 85% of the writing. Wheatley provided 5% of the interpretation of data, and 10% of the writing. Smart provided 5% of the interpretation of data, and 5% of the writing.

Barth, J.E., Wheatley, V., & Smart, M.K., "Inlet Fuel Injection in a Mach 12 Shape-Transitioning Scramjet," presented to the *AIAA SciTech 52nd Aerospace Sciences Meeting*, National Harbor MD, 13-17 January 2014. AIAA 2014-1159.

Parts of this paper are included in Chapter 4. Barth was responsible for 100% of the simulation work, 90% of the interpretation and analysis of data, and 85% of the writing. Wheatley was responsible for 10% of the writing, and 5% of the interpretation and analysis. Smart contributed 5% of the writing and 5% of the interpretation and analysis.

Contributions by Others to the Thesis

No contributions by others.

Statement of Parts of the Thesis Submitted to Qualify for the Award of Another Degree

None.

Acknowledgements

I've often ruminated on the idea that a PhD is usually hailed as a hallmark of personal achievement, and yet I would never have been able to finish this work alone. I wish to express my gratitude and thanks to those who provided support, assistance, and encouragement.

First and foremost, my thanks goes to my PhD supervisor, Dr. Vincent Wheatley, and my associate supervisor Prof. Michael Smart. Your advice, insight, guidance, and willingness to share your experience in the field of scramjet propulsion have been invaluable throughout this three and a half year period.

I owe a further debt of gratitude to The University of Queensland for the scholarship support, as well as to Dr. Russell Boyce and the SCRAMSPACE program for attracting the computational and financial resources that allowed me to complete this research and present some of it at the 18th Hypersonics and Spaceplanes Conference in 2012. I also owe a great deal to the NCI National Facility at ANU, and the UQ Barrine cluster team for their respective provisions of the computational resources used in the CFD work for this study, as well as their technical support when things went sideways.

To Dylan Wise, thank you for agreeing to let me tag my experiments on to the end of his campaign. Without his help, I'd never have been able to test my own tailored injection configuration in the T4 Stalker Tube. I owe further thanks to Keith Hitchcock, our ever-helpful T4 technician, and my T4 operators: Kevin, Tamara, and Tristan. Thank you all for making my brief experimental campaign a success.

To Wilson, Rowan, Luke, Sarah, Chris, Fabs, Rolf, Philippe, Sandy, Melrose, Matt, Dawid, Dave Petty, Dave Peterson, and Rowan. Thank you all for the many fruitful discussions, ideas, advice, help, and laughs over the course of my time at UQ.

To my parents and family back in Canada: thank you for making me the person I am today. Thank you for supporting me when I announced I was quitting my job to "go to Space University," and for believing in me when my dream of getting into the space sector took me away from you and around the world. Twice.

To my amazing wife, Sharine, for her constant patience, love, and support; and to my new family-in-law for their support and understanding. I can never thank you enough for agreeing to let me join my life to yours. I love you more than I have words to use.

Finally, to my Lord and Saviour Jesus Christ, from whom all blessings flow. I simply pray that this work is a pleasing offering in your eyes.

Keywords

scramjet, hypersonics, hypervelocity, airbreathing propulsion, boundary layer combustion, tailored fuel injection

Australian and New Zealand Standard Research Classifications (ANZSRC)

ANZSRC code: 090107, Hypersonic Propulsion and Hypersonic Aerodynamics, 100%

Fields of Research (FoR) Classification

FoR code 0901, Aerospace Engineering, 100%

Contents

List of Figures

List of Tables

Nomenclature

Roman Symbols

a	=	Millikan-White coefficient, $a = 0.00116\varphi^{1/2}\theta^{4/3}$; or,
	=	Local speed of sound, m/s
A	=	Area, m^2; or,
	=	Pre-exponential factor in the chemical kinetic rate equation; or,
	=	Coefficient in the jet penetration equation of Portz and Segal (2006)
B	=	Temperature exponent in the chemical kinetic rate equation; or,
	=	Coefficient in the jet penetration equation of Portz and Segal (2006)
b	=	Millikan-White coefficient, $b = 0.015\varphi^{0.25}$
c_{b1}, c_{b2}	=	Constants in the Spalart-Allmaras turbulence model
c_p	=	Specific heat at constant pressure, $J/kg \cdot K$
c_{w1}	=	Constant in the Spalart-Allmaras turbulence model
C	=	Coefficient in the jet penetration equation of Portz and Segal (2006)
d	=	Distance to the nearest wall in the Spalart-Allmaras turbulence model
D	=	Diameter, mm
E	=	Activation Energy, $J/kmol$; or,
	=	Total energy in the Navier-Stokes equations; or,
	=	Coefficient in the jet penetration equation of Portz and Segal (2006)
f_w	=	Coefficient in Spalart-Allmaras turbulence model
f_{v1}	=	Coefficient in Spalart-Allmaras turbulence model
F	=	Force, N; or,
	=	Coefficient in the jet penetration equation of Portz and Segal (2006)
G	=	Coefficient in the jet penetration equation of Portz and Segal (2006)
h	=	Static Enthalpy, MJ/kg
H_0	=	Stagnation Enthalpy, MJ/kg
I	=	Identity matrix
k	=	Reaction rate, mol/s
L	=	Length, m or mm
$\dot{m}$	=	Mass flow rate, kg/s
M	=	Mach Number
p	=	Order of Convergence
P	=	Pressure, kPa
Pr	=	Prandtl number
q	=	Dynamic Pressure, kPa

$\vec{q}$ = heat flux vector in the Navier-Stokes equations,
Q = Measured heat flux, W
r = Coefficient in the Spalart-Allmaras turbulence model; or,
= Coefficient used in the reference enthalpy equation
R = Universal Gas Constant, 8313 $J/kmolK$
Re = Reynolds Number, $1/m$
s = Entropy, J/K
$\hat{S}$ = Coefficient in the Spalart-Allmaras turbulence model
$\bar{S}$ = Coefficient in the Spalart-Allmaras turbulence model
Sc = Schmidt number
t = Thickness, mm; or
= Time, s
T = Temperature, K
T_v = Vibrational Temperature, K
u^+ = Friction velocity, m/s: $u^+ = \sqrt{\tau_w/\rho}$
U = Velocity magnitude, m/s
$\vec{U}$ = Velocity vector, m/s
$\vec{v}_s$ = species-specific diffusion velocity, m/s
W = Source term in the Navier-Stokes equations
x = Position in the streamwise direction, mm
y = Position in the lateral (side wall to sidewall) direction, mm
y^+ = Dimensionless wall distance, $y^+ = u^+ y/\nu$
Y = Mass fraction, ρ_s/ρ
z = Position in the vertical (body to cowl) direction, mm

Greek Symbols

α_j = Fuel system calibration coefficient
γ = Ratio of specific heats
δ = Boundary Layer Thickness, mm
Δ = Variance
η = Efficiency
θ = Characteristic Vibrational Temperature, K; or,
= angle, $^\circ$
κ = von Kármán constant, 0.41
μ = Dynamic viscosity, $kg/m{\cdot}s$
ν = Kinematic viscosity, $\nu = \mu/\rho$, m^2/s

$\hat{v}$	=	Transported turbulent viscosity in the Spalart-Allmaras turbulence model
ρ	=	Density, kg/m^3
ϕ	=	Equivalence Ratio, $\frac{\dot{m}_f}{\dot{m}_{f,stoich}}$
φ	=	Equivalent molecular weight between particles in vibrational relaxation, g/mol
Y	=	Uncertainty %
$\tilde{\tau}$	=	Viscous stress tensor in the Navier-Stokes equations
σ	=	Constant in the Spalart-Allmaras turbulence model
$\tau_i d$	=	Ignition delay time, s
τ_{react}	=	Chemical reaction time, s
τ_o	=	Time offset for normalizing experimental pressures, s
τ_v	=	Vibrational relaxational time, s
τ_w	=	Wall Shear Stress, Pa
χ	=	Molecular Mass, $kg/kmol$; or,
	=	a term in the SA turbulence model
ψ	=	Measurement uncertainty %

Subscripts

ad	=	adiabatic
aw	=	adiabatic wall
$calc$	=	calculated
$comb$	=	combustion
cr	=	critical
$drag$	=	acting along the vehicle drag axis
e	=	control volume exit
eff	=	effective
est	=	estimated
f	=	fuel
$flight$	=	condition is for engine in flight
i	=	control volume inlet
j	=	jet
KE	=	Kinetic Energy
LT	=	Ludweig Tube
mix	=	mixing
nom	=	nominal
NS	=	nozzle supply condition

p	=	penetration
$plenum$	=	condition inside fuel plenum
r	=	rotational; or,
	=	reference
s	=	species
sim	=	condition is from simulation
$stoich$	=	stoichiometric
ST	=	Shock Tube
t	=	throat condition; or,
	=	translational
$thrust$	=	acting along the vehicle thrust axis
tr	=	transition
u	=	unit
v	=	vibrational
w	=	wall
x	=	acting in the cartesian x-direction
z	=	acting in the cartesian z-direction
0	=	stagnation condition
∞	=	freestream condition

Acronyms

CFL	=	Courant-Friedrichs-Lewy number
EP	=	Enrichment Percentage
FAP	=	Freestream Addition Percentage
JF	=	Jet mass Fraction parameter
M12	=	Mach 12
REST	=	Rectangular-to-Elliptical Shape Transition
scramjet	=	supersonic combustion ramjet

1

Introduction

"It is difficult to say what is impossible, for the dream of yesterday is the hope of today and the reality of tomorrow."

– Robert H. Goddard, *"On Taking Things for Granted,"* (June 1904)

All levels of society are becoming ever more dependent on space-based technology. Individuals, corporations, and governments make use of orbiting satellites for diverse reasons, be they navigation, scientific inquiry, communication, or the proliferation of media. Marching in lock-step with the ongoing increase in space asset use is a growing interest in developing launch systems that might soon offer access to space that is more economical, environmentally responsible, and reliable than current systems. Among the most promising of these so-called "new space" technologies is the air-breathing supersonic combusting ramjet (scramjet) engine. By burning its fuel with atmospheric oxygen instead of on-board (carried) oxidizer, scramjets may potentially offer a much higher specific impulse, and therefore higher efficiency, engine alternative to rocket technology (Mehta, 1996).

1.1 Scramjets

What, then, makes a scramjet distinct? The answer may be obtained by examining the basic scramjet configuration shown in Figure 1.1. The reader could be forgiven for thinking that the contoured tube they are seeing is unusually simple for what is ostensibly a piece of advanced technology. As is the case with many things, the apparent simplicity of the scramjet is deceiving. Flying through the air at hypersonic speeds, (that is, speeds exceeding five times the local speed of sound), the scramjet relies on its speed and the shape of its inlet to compress the air it is flying through sufficiently to sustain combustion, despite the air the engine processes never dropping to a subsonic

speed. In an access-to-space applications, a scramjet-powered vehicle would follow a trajectory that maximizes the dynamic pressure, in order to achieve to higher internal pressures and therefore engine thrust (Billig, 1992).

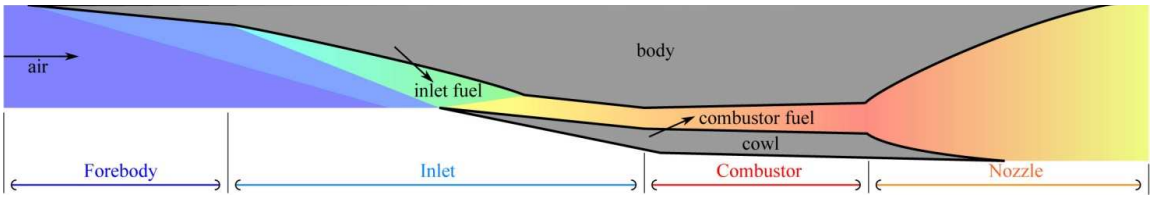

Figure 1.1: Schematic of a typical scramjet engine (Wheatley, 2014)

After incoming air is compressed by the inlet surfaces, it typically passes through a constant-area duct, known as an isolator, that serves to prevent the high pressures generated by combustion further downstream from propagating upstream into the inlet and causing the engine to unstart. At some point along the flow path, fuel is injected and allowed to mix and react with the airstream. The subsequent heat release inside the engine combustor causes a local increase in pressure, which is translated into a thrust force as the flow is expanded through the exhaust nozzle. Clearly, while the scramjet may be *structurally* quite plain to look at, the nature of the combusting air flow a scramjet creates is anything but; the engine flow is characterized by shock waves, regions in which temperatures are in the thousands of kelvin, and pressures that can be 100 times greater than that of the air the engine is pushing a vehicle through.

Of course, scramjet technology is no panacea for efficient launch vehicles: Weber and MacKay were the first to show that a well-designed scramjet only becomes efficient and capable of net thrust at speeds just below Mach 6 (Weber and MacKay, 1958). Since scramjet operation is limited to high speeds within a sufficiently dense atmosphere, scramjets alone cannot provide access to space. However, using scramjets inline with other propulsion technologies, a hybrid launcher with greater fuel efficiency than rockets alone can be achieved. This efficiency translates directly into great mass availability for increased carrying capacity or added system redundancy, leading to higher reliability and/or more economical delivery to payload to orbit. In order to achieve this, however, a scramjet capable of delivering net thrust at speeds of up to Mach 12 or higher will likely be required (Smart and Tetlow, 2009).

There are multiple technical hurdles to be cleared in order to achieve a practical scramjet capable of operating at the high speeds required for space access. Some problems, such as the high structural heating loads common at these velocities (Anderson, 2006), may be only solvable through the development and application of new materi-

als that can withstand such a harsh environment. Others can be influenced by altering the nature of the engine's internal flow.

While nearly any scramjet engine design could theoretically power a launch system through the middle part of its trajectory to orbit, one of the most promising scramjet designs is the Rectangular-to-Elliptical Shape-Transition (REST) engine. This type of engine features a low-loss, mixed-compression inlet with a rectangular capture area that gradually morphs along the length of the inlet to yield an elliptical cross section at the combustor. Such an engine is therefore easy to integrate into a practical vehicle configuration while maintaining a low wetted combustor surface area that requires only a thin wall to withstand the structural loads generated by combustion pressure rise. Since this type of engine has fixed geometry, its inlet geometry must be chosen such that high performance is achievable across the entire access-to-space Mach number range. A REST engine developed at The University of Queensland (UQ) with an on-design flight speed of Mach 12 (Suraweera and Smart, 2009) could potentially operate at speeds as low as Mach 6-7, and as high as Mach 12: the entire range of Mach numbers necessary for access-to-space.

In order to be viable as a space-access technology, significant improvements to the Mach 12 (M12) REST engine are required. This has been true of scramjets in general for much of the half century that they have been a topic of active research. It was 37 years after the pioneering study of Weber and MacKay in 1958 that the first successful successful demonstration of a scramjet producing net thrust was achieved at UQ in 1995 (Paull et al., 1995). In 2010, that the X-51 scramjet vehicle sustained a hypersonic flight, accelerating from Mach 4.5 to Mach 5 over its 200 second test time (Warwick, 26 May 2010). While encouraging, the body of literature to date suggests a more sobering picture: the establishment of robust supersonic combusting flow in a scramjet at high flight Mach number is no simple task (Curran, 2001). After all this time, hypersonic air-breathing vehicles remain experimental only and are most commonly tested at the lower end of the hypersonic speed range. Outside of a single 10s flight by the X-43 at Mach 9.68 (Marshall et al., 2005), sustained long-duration scramjet operation at hypervelocity speeds (greater than 3000 m/s or about Mach 10) has yet to be demonstrated.

A REST-style engine could, however, be one of the first to achieve such a feat: Turner and Smart's Mach 8 REST engine was able to exceed its expected combustion efficiency, burning nearly 80% of the injected fuel at near-stoichiometric conditions (Turner, 2010). Smart has recently estimated that this combustion efficiency is likely the minimum required for a scramjet to generate enough thrust to become a viable engine for an access-to-space launch system (Smart, 2012). In contrast, the original

design for a M12 REST engine currently had an estimated combustion efficiency of approximately of 60%.

Thus, there is room for significant improvement on the M12 REST engine's performance characteristics. Achieving a high combustion efficiency in this engine will eventually raise the maximum flight speed at which net thrust can be expected to speeds in excess of Mach 12, which is a requirement for air-breathing access-to-space propulsion to be a viable option (Smart and Tetlow, 2009). There remain, however, several unanswered questions as to the nature of the flow inside the engine, as well as how best to improve combustion efficiency in the engine.

1.2 Thesis Objectives

The purpose of the work described in this thesis is to improve the performance of the M12 REST engine, and prove its viability as a good candidate for use in an access-to-space system. This can only be achieved by first developing a detailed understanding of the current engine design's flow behaviour, and then tailoring fuel injection to exploit features of the engine flow path. The general aim of this thesis can therefore be summed up as an attempt to answer the following two-part question:

What flow physics cause the M12 REST scramjet engine to perform as it does when operating at on-design conditions? Can the internal flow observed be exploited through tailored fuel injection to improve its mixing and combustion efficiency?

The lessons learned from this investigation should be applicable to any practical stream-traced hypervelocity scramjet, of which the M12 REST engine serves as an example configuration that used as the engine geometry throughout this study. The investigation was primarily computational in nature, with some experimental work carried out to provide validation data for the simulations.

1.2.1 Research Questions

The overall objective of improved combustion performance in the M12 REST has a clearly defined end, but did little to elucidate the means by which this end could actually be achieved. To provide direction on this point, a series of smaller questions were posed. The answers to these questions, posed below, build into the overall goal given above.

1. What flow structures are generated by the structural geometry of the M12 REST, and by what mechanisms?

The answer to this question is critical to understanding engine performance: shock and vortical structures generated by incoming flow interacting with the M12 REST geometry undoubtedly have a role in the mixing and ignition of fuel inside the engine.

2. What changes to the flow structure occur when fuel is injected in the inlet, and what portions of the flow does the injected fuel interact with? Where does the inlet-injected fuel ignite, and by what mechanism?

As the inlet and jet flow structures interact, changes to their expected behaviour may occur. The addition of fuel jet structures into the boundary layer may have a disruptive effect, and was believed to cause some increase in the drag forces exerted on the inlet.

3. What happens to fuel injected from the face of the backward-facing step that marks the junction between the engine isolator and combustor sections? By what mechanism is the combustor-injected fuel ignited?

Fuel was injected from the face of a backward step at the entrance of the combustor, as it was hoped this would reduce the skin friction generated the boundary layer of the M12 REST combustor. Answering this question will show how effective the injected fuel performs in this role, or if it stays in the boundary layer at all!

4. Can fuel injection tailored to the engine's internal flow structures provide the significant increase in mixing and combustion efficiencies required to bring the M12 REST engine closer to being a viable access-to-space technology?

Answering the first three questions should provide knowledge of if, and where, tailored fuel injection would be required in the Mach 12 REST engine. Actually attempting to implement a bespoke fuelling configuration for the M12 REST can provide answers to what approach is the most effective at promoting fuel/air mixing and robust combustion.

5. In addition to improving engine performance through tailored injection, is it possible to identify other avenues for improving the engine performance?

Lastly, during the course of this investigation, there will be a great deal of attention paid to what is causing the engine to perform as it does. There may be specific aspects of the engine design or operation that, if changed, could yield further improvements to the engine's performance. Estimating the rough magnitude of the effects of these

changes can be accomplished through a simplified analysis, giving guidance to future design and investigation of three-dimensional hypervelocity scramjets.

1.3 Thesis Overview

The remainder of this thesis is organized into seven further chapters, not including the appendices to the main document that contain additional technical information. These seven chapters are described briefly below.

Chapter 2 - Literature Review provides an overview of the relevant literature, to give an understanding of the current research environment in the field of hypersonic propulsion. Topics examined include the three main challenges facing scramjet designers, hypersonic inlets, a detailed accounting of the state of research into REST scramjet engines, and the nature and uses of fuel injection inside a scramjet engine.

Chapter 3 - Methodology This chapter presents the exact geometry of the engine studied, as well as the methods by which on-design flow through the engine was examined. Both the computational tools used to analyze the engine flow field and experimental apparatus used to obtain validation data are described. Both methods are analyzed to give an appreciation of the uncertainty in either approach, and the flow conditions used are presented.

Chapter 4 - Shape-transitioning Inlet Flow Physics and Fuel Injection Effects The first of the results chapters, Chapter 4 focuses on the flow inside of the M12 REST inlet. The flow structures generated, and the resulting inlet performance, are presented through simulations of both an unfuelled inlet, and an inlet in which inlet fuel injection is employed. The simulated inlet is contrasted with available experimental data. The impact on inlet performance of injected fuel beginning to combust upstream of the inlet throat is investigated, and the net forces acting on the inlet are contrasted with those of the unfuelled inlet.

Chapter 5 - Combustor and Nozzle flow physics and performance Much like the combustor and nozzle sections are downstream of the engine inlet, so Chapter 5 follows on from the analysis begun in Chapter 4. Chapter 5 presents a simulation of the unfuelled flow through the engine combustor and nozzle, then follows this with an examination of how the engine flow changes with the inclusion of inlet fuel injection. Finally, the "baseline" engine is investigated, in which fuel injection into the combustor boundary layer is combined with inlet injection in an attempt to reduce drag. In all cases, the forces acting on the entire engine are presented in order to determine the net

performance benefits that are gleaned from fuel injection and the resulting combustion efficiency of the engine.

Chapter 6 - Tailored Fuel Injection Armed with the knowledge obtained from earlier results, this chapter presents an examination of attempts made to tailor fuel injection to take advantage of the flow fields inside the M12 REST engine with the inclusion of inlet injection. Two tailoring injection cases are presented, with the second being compared to the baseline engine presented in the previous chapter. After comparing the second case that was simulated with validation data obtained through experiments, the flow field of the engine operating at a true flight condition is compared with that of the tunnel conditions simulated previously, in order to gain an appreciation of how well the experiments match to the expected flight conditions.

Chapter 7 - Improving Tailored Injection and Engine Performance Based on the results presented in the three previous chapters, an analysis of ways in which the current best-case tailored injection engine could be improved is presented. The impacts of fuel heating and oxygen enrichment on the current configuration are examined before ways in which the injection geometry, and the geometry of the entire engine, might be modified to provide further improvements to engine performance are presented.

Chapter 8 - Conclusions and Recommendations The final chapter of the thesis summarizes the findings and analysis of the previous chapters, and ties them back to the research questions posed in Section 1.2. The chapter closes with a set of recommendations on how study of the M12 REST engine may be extended.

1.4 Summary

This chapter was intended to provide the reader with a brief introduction to the problem of hypervelocity scramjet operation, and the motivation for the work that follows was defined. The chapter also sought to elucidate the general structure of this investigation, as well as the research questions that would be answered. With a complete roadmap of the thesis work defined, attention can now shift to the body of scientific research which was drawn upon in order to properly define the scope of the project, and inform design decisions made throughout the work.

2

Literature Review

"Later-on, I realized that this mission had to end in a let-down, because the real barrier wasn't in the sky, but in our knowledge and experience of supersonic flight..."

– Chuck Yeager, *Yeager: An Autobiography* (1986)

Scramjet propulsion has been a field of active study for more than 50 years, and a substantial body of literature exists to provide the background and basis for the present study. The intent of this review is to establish the context in which this thesis fits into the domain of the scramjet research as a whole, and how it will provide new insights to the design and operation of an efficient scramjet engine. As such, the chapter begins with a brief overview of the primary challenges facing scramjet designers today. This overview provides a launching point for investigating hypersonic inlets, particularly the REST scramjet inlet. The body of literature available regarding REST scramjet engines is examined, including the present gaps in knowledge that need to be addressed to improve their efficiency and potential for generating net thrust at hypervelocity conditions. Following this study of the REST engine, the issue of fuel injection and mixing is covered in detail, focusing on the various ways fuel can be injected and mixed inside a scramjet engine. Particular attention is paid to not only the phenomenon of injection itself, but introducing the concept of tailored fuel injection, and the flow structures such injection might target.

2.1 Key Challenges in Scramjet Propulsion

While the principles of scramjet propulsion are well understood as a matter of theory, there are several practical issues that have prevented the development and deployment of scramjet-powered vehicles. Chief among these are the problems of how to achieve stable, efficient combustion in a supersonic flow; how to predict and manage

the heating a vehicle would encounter while travelling at hypersonic speeds; and how to manage the problem of skin friction drag, the importance of which increases with flight speed. Each of these issues is examined in turn to identify the state of current knowledge on each issue, and suggest paths for dealing with them in the context of a practical scramjet engine.

2.1.1 Supersonic Combustion

The process of combustion in a supersonic flow condition is a complicated affair: fuel must mix, and then burn, with captured air inside the confines of the engine, a process made difficult by the fact that flow moving at thousands of meters per second will have a very short residence time within an engine that is (at best) a few meters in length. This problem is magnified if the engine is operating at the high flight Mach numbers an access-to-space scramjet might operate at. At this range of speeds, the fuel residence time can often approach the timescale of the complete combustion chemical reaction (McGilvray, 2008).

Supersonic combustion can be broadly classified into two main types: mixing-limited and reaction-limited combustion. In the former, flame propagation is controlled by the speed of fluid transport phenomena such as diffusion, whose time-scales are longer than chemical kinetic timescales (Ferri, 1973). By contrast, in reaction-limited combustion the flame propagates as quickly chemistry permits the fuel-air mixture ignite and burn, a common feature in detonation-combustion systems (e.g. Dunlap et al. (1958), Sislian (2000)).

In a typical hydrogen-fuelled scramjet, the combustion process is mixing-limited and strongly coupled to the flow structures present. In addition to the presence of shock waves and recirculation zones promoting ignition (Ben-Yakar and Hanson, 2001), large flow field gradients can make combustion a mixing-limited process in some parts of the combustor, and reaction-limited in others (Segal et al., 1995). Flow compressibility in isolation can lead to significant decreases in mixing. However, some baroclinic and dilatational flow effects (e.g. shock-flow density gradient interactions) can generate shear, and consequently turbulent eddies that enhance the mixing and burning of fuel (Ingenito and Bruno, 2010).

Combustion can, in turn, change the flow structures observed. Pressure rises associated with combustion have been shown to reduce the three-dimensionality of flows (Segal et al., 1995), and can introduce new flow features (Malo-Molina et al., 2010). The

heat release from combustion can also reduce turbulent shear stresses and total turbulent kinetic energy in the flow, as well as decrease the magnitude of vorticity in affected regions (McMurtry et al., 1989). How such flow structures affect mixing is discussed further in Section 2.4.

2.1.2 Aerothermodynamic Heating

As flight Mach number increases, the difficulty in achieving combustion heat release is not only due to the high kinetic energy of captured air rushing through the engine faster than it can burn; there are also effects due to increased static temperatures in the air interacting with the engine. This is primarily effects seen due to the dissociation of constituent molecules in the air. At high static temperatures, such as those encountered at flow stagnation points along scramjet inlets, molecular oxygen and nitrogen in the air can begin dissociating into their atomic form. This endothermic process in turn absorbs heat from the flow, reducing net heat release in the engine. While this lost energy would ideally be recovered by the recombination of the atomic species into a molecule, optimized nozzle designs often include a rapid expansion at the start of the nozzle, which "freezes" the chemistry. The dissociation energy is therefore lost from the engine, reducing net thrust (Anderson et al., 2000).

Aerodynamic heating is also an important consideration in the design of hypersonic vehicles, engine capture shapes, and leading edges (Anderson, 2006). Because scramjets can only produce net thrust in relatively dense atmospheric conditions (Smart, 2007), the aerodynamic shear imposed by the no-slip condition along vehicle surfaces results not only in increased viscous drag (which is discussed in detail in the next subsection), but in increased wall heat transfer levels (Kutschenreuter, 2000).

At hypersonic speeds, the heat transfer at a stagnation point (such as an inlet leading edge) is inversely proportional to the square root of the edge radius (Anderson, 2006). As scramjet engine leading edges tend to be relatively 'sharp,' with typical leading edge radii on the order of 1 mm or less, stagnation heat loads can quickly exceed the natural limitations of most modern materials used in aerospace vehicles. In hypersonic inlets, at least, increasing leading edge radii above 1 mm can have an adverse effect on inlet performance (Drayna et al., 2006). It is therefore likely that a scramjet operating at high Mach numbers will require active cooling, and/or advanced materials, to allow these sharp leading edges to survive long-duration flights at these speeds.

2.1.3 Hypersonic Skin Friction Drag

The third problem that makes scramjet propulsion difficult is related in some way to both of the previous two problems: skin friction drag is a problem that is tied to the phenomenon of aerodynamic heating, and is often a particular challenge to flow control inside of a scramjet combustion chamber. The drag experienced by any intra-atmospheric vehicle is composed of two components: an inviscid (pressure) drag force, and a viscous (skin friction) drag. Hypersonic vehicle design tends toward aircraft that are long and slender. Coupled with the fact that skin friction drag increases with speed and that hypersonic vehicle boundary layers are predominantly fully turbulent in their flight envelope, this may lead to as much as 50% of the net drag force acting on a typical hypersonic vehicle being caused by skin friction (Anderson, 2006).

The dominance of skin friction drag is also true on the engine subsystem level: much of the total drag in a scramjet engine can be skin friction drag (Paull et al., 1995). As shown in Figure 2.1, the coupling of high pressures and temperatures with high flow velocity inside a supersonic combustor can lead to as much as half of the total viscous drag acting on a scramjet being generated by the combustor alone (Tanimizu et al., 2009). Producing a scramjet capable of producing net thrust at hypervelocity speeds is therefore not only a matter of improving combustion efficiency: reducing skin friction drag is arguably of equal importance to improving engine performance (Czysz and Bruno, 2009).

Several options have been investigated for reducing skin friction. These include laminar flow control, passive turbulent drag reduction, and active turbulent drag reduction. Laminar flow control seeks to keep the boundary layer laminar, as a laminar boundary layer will have a lower coefficient of skin friction c_f relative to a turbulent boundary layer at a given Reynolds number. Passive turbulent drag reduction includes the addition of riblets or convex curvature to a surface, or static devices that modify the outer layer of the boundary layer. Finally, active turbulent drag control can include methods such as boundary layer capturing, MHD manipulation, and local heat and/or mass addition. All of these techniques are described in detail in the literature review of Suraweera (2006), and in the work of Bushnell and Hefner (1990).

Most of these techniques may be applied to the problem of external flows, but are not necessarily appropriate to the internal flow of a scramjet. As will be shown in Section 2.3.1, REST inlets require turbulent boundary layers, which in turn implies that internal viscous drag will be higher than an engine without this requirement. Passive turbulence control devices may suffer from high heat transfer loads, and many of the

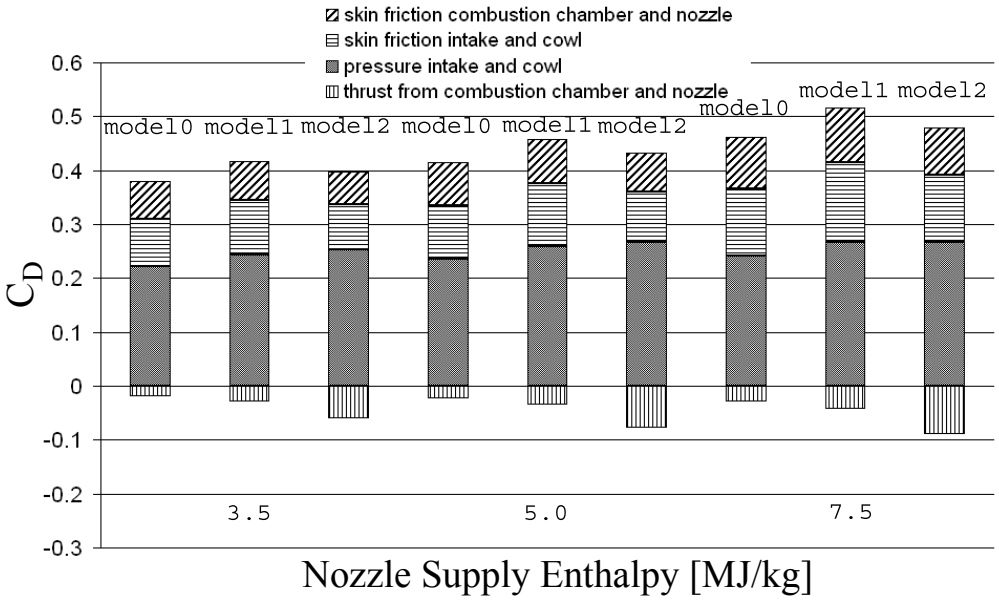

Figure 2.1: Typical sources of drag in a scramjet, and their relative magnitudes. Reproduced from Tanimizu (2008).

active control methods are complex enough to warrant a complete independent study. One method, however, is quite suitable for application to scramjets (and particularly scramjet combustion chambers): fuel mass addition to the boundary layer of scramjet combustors has been well studied, and will be presented in Section 2.4.3.

Many of the problems presented above are influenced in no small part by the size and shape of a given scramjet's flow path. Much of a scramjet's flow path is determined by the type of inlet it employs. As the inlet provides most of the compression the flow encounters during operation, inlets are among the most studied (and most difficult to design) components of a scramjet flow path, and are often used to categorize scramjets. To provide context for where the present study fits into the general landscape of scramjet research, each of these categories are briefly presented along with their limitations.

2.2 Hypersonic Inlets

As their name implies, hypersonic inlets are designed to compress air entering a scramjet engine at hypersonic conditions, defined by convention as flow conditions at speeds higher than a Mach number of 5-7 (Heiser and Pratt, 1994). The simplest, and most thoroughly-researched scramjets those that employ a simple two-dimensional or axisymmetric geometry. The relative simplicity of designing these engines make them

attractive for study and, especially in the case of planar scramjets, easy to integrate into standard vehicle designs. While well-studied, these types of engines are not without their limits. Often, they are limited in the amount of compression their inlets can do before becoming unstarted (Van Wie, 2000).

2.2.1 Three-Dimensional (3D) Inlets

A great deal of work has been done investigating the utilization of an inlet's sidewalls as an additional compression surface. Gruhn and Gühlan, for example, recently investigated the effects of adding of sidewall compression to a planar inlet, as shown in Figure 2.2. This addition was shown to all the inlet to remain started at higher contraction ratios, along with imparting improved ability to handle adverse yaw angle and combustion chamber back-pressure effects (Gruhn and Gülhan, 2011). Increases in inlet performance relative to their planar equivalents have been demonstrated at speeds as high as Mach 10 (Holland, 1995).

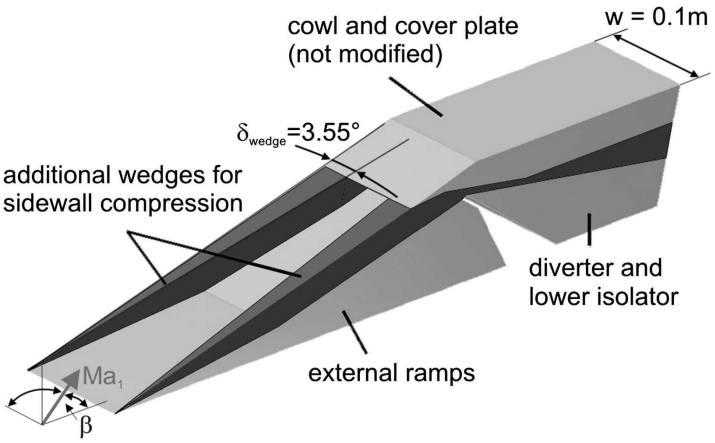

Figure 2.2: Schematic of a planar scramjet inlet with sidewall compression. Reproduced from Gruhn and Gühlan (2011).

Another type of inlet that inherently uses sidewall compression are stream-traced inlets. The simplest of this type of engine is an axisymmetric Busemann-type inlet (Mölder and Szprio, 1966). However, there is no reason an inward-turning inlet must remain axisymmetric: any desired capture shape can be superimposed over a truncated Busemann flow-field, and then stream-traced to produce the total inlet geometry. The Busemann flow-field is typically truncated, as even a small truncation angle can lead to a significant reduction in inlet length, which helps to minimize inlet drag (Drayna et al., 2006) Through stream-tracing, a scramjet inlet can thus be designed to fit perfectly into a proposed vehicle design, and still move into a combustion chamber of any desired shape. Such engines can thus have the desirable characteristics of

rectangular inlets with sidewall compression, but with several improvements: lower viscous losses, higher pressure recovery, and fewer requirements for the design of the engine isolator (Billig and Jacobsen, 2003).

Inlets such as the "sugar scoop" (Gaitonde et al., 2007) shown in Figure 2.3, and "jaws" inlet (Croker, 2007)) are examples of geometries defined by stream-tracing through a low-loss axisymmetric compression flow field, and show the desirable characteristics listed previously. Their improved performance comes with a significant drawbacks: their unique geometry can make vehicle integration a challenge, and the flow entering the combustor is no longer uniform. Another example of a streamtraced inlet geom-

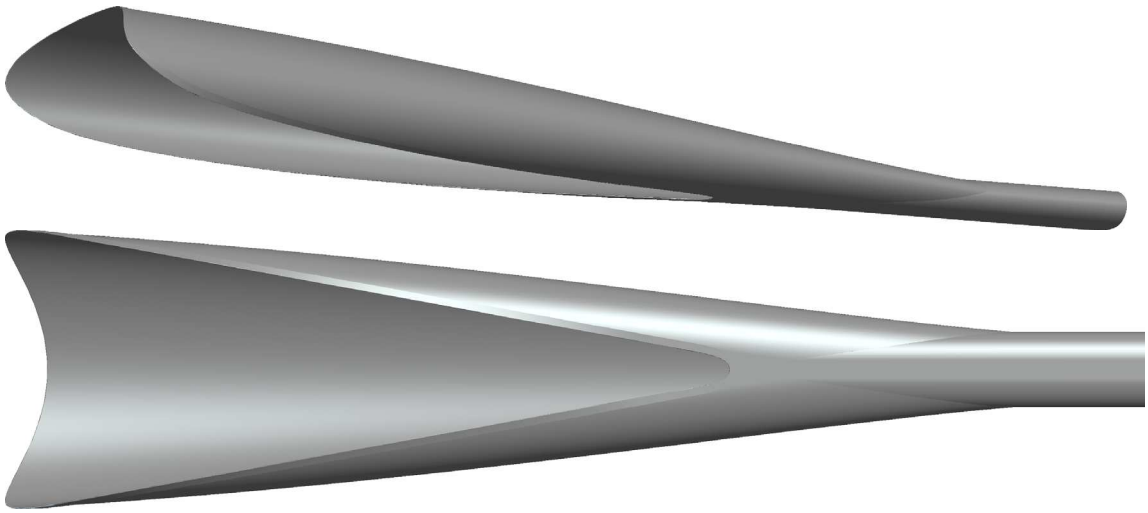

Figure 2.3: The stream-traced "sugar scoop scramjet inlet. Reproduced from Drayna (2006).

etry, one which is far less challenging to integrate into a hypersonic vehicle, are the Rectangular-to-Elliptical Shape Transitioning (REST) inlets, one of which was examined over the course of this study.

2.3 The REST Scramjet Engine

The REST scramjet inlet design was first described in 1999 by Michael Smart, who used a streamline tracing technique to blend a rectangular capture area to an elliptical combustion chamber (Smart, 1999). This combined the ease of integrating a rectangular low-loss inlet into a vehicle with the structural strength and low wetted area of a rounded combustion chamber cross-section. While the REST inlet itself is designed with a fixed geometry based on a desired flight Mach number, previous studies have shown that a REST inlet can operate efficiently over a wide range of Mach numbers

near its design point (Smart and Trexler, 2004). This is in large part due to the cowl closure shape, and highly swept leading edges shown in Figure 2.4, that allow for overboard flow spillage under starting and off-design conditions (Turner, 2010).

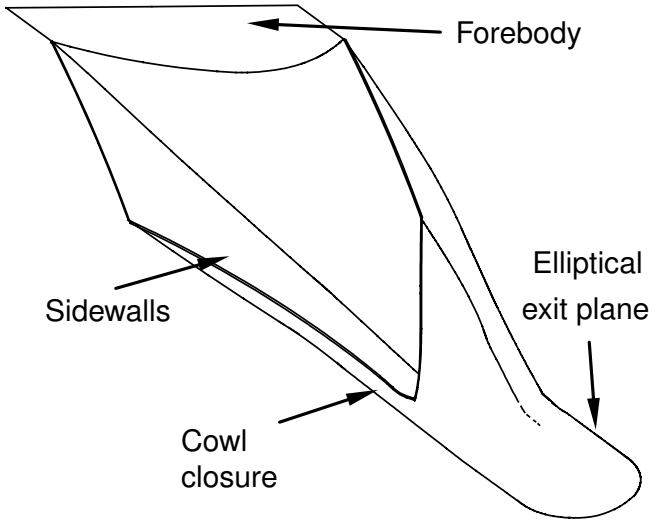

Figure 2.4: A schematic of the basic REST inlet geometry. Reproduced from Turner (2013).

Past studies have investigated REST engines with design flight Mach numbers of 7, 8, and 12. In all cases, the incoming freestream flow was assumed have a dynamic pressure of 50 kPa, and is processed by a 6-degree leading edge shock from the vehicle forebody before encountering the engine inlet (Smart, 1999). These studies are summarized in the sections that follow.

2.3.1 Experimental Investigations

The first experimental tests of the REST engine were for an engine with a design flight condition of Mach 7.1, which corresponds to Mach 6 flow entering the engine behind the vehicle forebody shock. The experiments of Smart (2001) tested the engine with mechanically-induced back pressures to simulate a pressure rise due to combustion downstream. The engine was found to be self-starting, even under mechanically induced engine unstart conditions, and remained started with back-pressure ratios as high as 5.7. Further testing with the injection and combustion of hydrogen fuel under Mach 5.3 flight conditions showed that the engine maintained robust combustion for fuel equivalence ratios between 0.5 and 1.5 (Smart and Ruf, 2006). The engine remained self-started under a variety of adverse conditions, though at the lower range of Mach numbers (i.e. Mach 4) this engine was found to need additional flow spillage to remain started (Smart and Trexler, 2004).

A REST engine with a design flight Mach number of 8, that utilized the inlet injection technique developed at the University of Queensland (Stalker et al., 2005), was investigated by Turner and Smart at a Mach 8.1 flight condition. The intent of using inlet injection was to provide extra fuel residence time within the engine in hopes of improving its overall mixing and combustion efficiency, as well as to mitigate skin friction drag by shortening its combustor length (Turner, 2010). The inlet at a slightly oversped flight condition in order to allow the leading edge shock to be swallowed into the engine, in an attempt to promote combustion via the so-called 'radical farming' technique of Gardner et al. (2002). It was found that with this configuration, combustion efficiencies as high as 60% could be achieved (Turner and Smart, 2010) when the engine was fuelled on the inlet to an equivalence ratio of up to 0.8. Interestingly, Turner also found that combining inlet injection with symmetric porthole fuel injection in the engine combustor could yield combustion efficiencies as high as 80% (Turner, 2010). This was especially promising in light of the fact that the combustion chamber was designed to only be sufficient for ignition and reaction of inlet-injected fuel to occur, with no allowance made for mixing of the additional fuel injected in the combustor.

More recently, an engine with a design point of Mach 12 has undergone extensive testing at off-design Mach 9 conditions. This engine also made use of inlet injection, but had the ability to inject fuel in the streamwise direction along the combustor wall, from an array of portholes in the face of a backward step located at the junction of the engine isolator and combustion chamber. This injection technique was intended to promote boundary layer combustion of injected fuel, for reasons which will become clear in Section 2.4.3 (Suraweera and Smart, 2009). As was the case with the Mach 8 engine, the M12 REST engine's best performance came when both injection modes were combined, with an increase in thrust potential of nearly 50% and a combustion efficiency of 65% at an equivalence ratio of 1.2 (Moule and Smart, 2013).

Tests of a 0.32-scale version of the M12 REST engine were conducted at a Mach 10 flight condition (Doherty, 2013). These free-jet experiments were the first to test if an airframe-integrated engine could produce sufficient thrust to overcome not only the internal engine drag, but drag on the forebody and engine external surfaces. It was found that the engine is relatively insensitive to differences in dynamic pressure, as demonstrated by the lack of ρL-scaling (Hornung, 1988) for this model relative to the flight condition. The engine did not achieve net thrust, but with a basic conical thrust nozzle and 55% combustion efficiency, it was surmised that improvements in both these areas could make the engine produce net thrust (Doherty, 2013).

A half-scale version of the M12 REST engine was tested at Mach 12 flight conditions

in semi-free jet mode near the end of 2013 (Wise, 2014). As with the previous tests, its best performance occurred when the engine was fuelled at an equivalence ratio well above unity, achieving an apparent global combustion efficiency of 65%. Like the work of Doherty before it, it was also determined that the combustor boundary-layer-injected fuel did not ignite in the absence of inlet injection, suggesting the importance of inlet injection to the overall engine operation.

Examining the results of the Mach 8 and Mach 12 engine tests together, it appears that most of the combustion efficiency gains from combined fuel injection came from the injection of fuel into regions where inlet fuel did not penetrate, using heat release from the combustion of inlet-injected fuel to light the boundary-layer-injected fuel. This assumption is supported by the numerical studies done to date, which are described in the following subsection.

2.3.2 Numerical and Analytical Work

The primary body of numerical work existing for REST engines is largely unpublished; while many of the experimental papers cited in the Section 2.3.1 make mention of basic flow simulations, little detail is supplied. There is, however, the simulation work of Smart and White (2002), which investigated the flow fields in a REST engine designed for Mach 7 flight conditions at on-design conditions using the VULCAN flow solver. The flow was reproduced relatively well with and without artificially induced back-pressures modelled, though the model did not appear to capture the internal shock train observed in experiments. A strong influence of the external angle of the cowl leading edge on inlet performance was also noted (Smart and White, 2002).

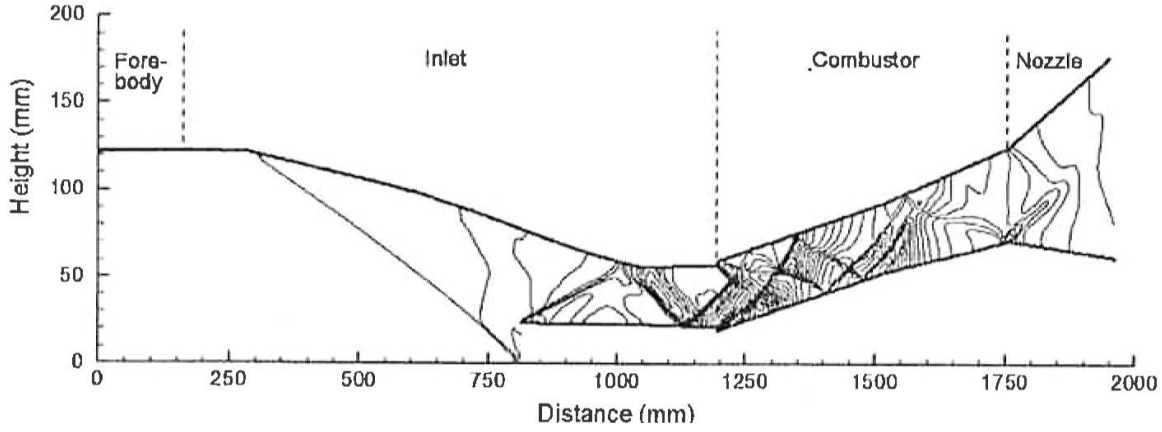

Figure 2.5: A typical unfuelled centerline pressure distribution in the M12 REST engine. Reproduced from Suraweera (2009).

Much of the other published numerical results are similar to the plot given in Figure 2.5: pressure contours taken at the vertical symmetry plane that runs down the center of the engine. These contours are instructive as to the condition of the flow inside a REST engine. A large shock train is seen to propagate through the engine, originating at the cowl notch.

Plots such as Figure 2.5 are useful in determining general characteristics of the flow field, but do little to elucidate the presence of 3D flow structures that certainly exist in this flow. There is a great deal more information that can be gained from these relatively simple models of fluid flow through the engine without fuel injection. The work of Gollan and Ferlemann investigated the effects of parametric changes to REST class inlets (Gollan and Ferlemann, 2009),(Gollan and Ferlemann, 2011). These studies yielded a number of interesting results, one of the most notable being the presence of interesting 3D flow structures across the span of the engine isolator.

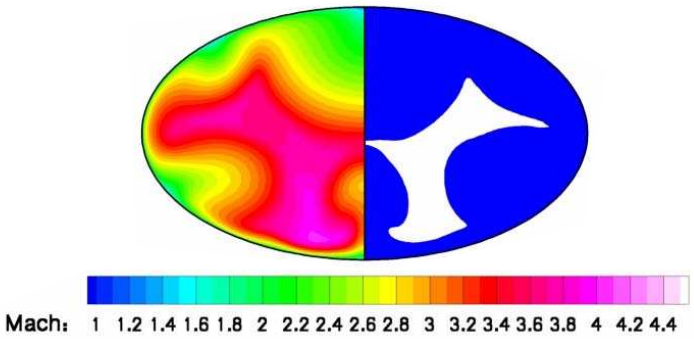

Figure 2.6: Mach number (left) and boundary layer contours (right) at the throat of a REST inlet. Reproduced from Gollan and Ferlemann (2009).

As demonstrated in Figure 2.6, the boundary layer appears to grow into significant low-speed "bubbles" at various locations around the perimeter of the isolator. Turner (2010) demonstrated the existence of a swept separation that forms inside of a Mach 8 REST inlet, but there is little characterization of the flow contained in this flow structure. The origin of such structures as these have not yet been fully explained, and they could be responsible for the mixing of cowl-side-injected fuel that lead to the excellent combustion efficiency observed in Turner (2010).

The act of injecting fuel will undoubtedly change the flow structures that can be observed in simulations like the one presented above. This brings to light a current deficiency in the literature regarding REST scramjets: to date, only one combusting flow simulation has been successfully carried out for this engine: the work of Moule et al. (2014), which was conducted at an off-design condition of Mach 9. While the results

agreed well with the estimated combustion efficiency of 65% estimated from experiments at this Mach number, the simulation lacked sufficient resolution to resolve and describe the detailed flow physics within the engine. High-resolution, fully reacting simulations will be critical to understanding the flow physics inside the REST engine, and to any attempt at improving performance through additional injection and/or mixing of fuel.

One final note may be made regarding current analysis of REST engine performance. Smart has estimated that a scramjet engine with 80% combustion efficiency may produce thrust sufficient for use in an access-to-space applications (Smart, 2012), as illustrated in Figure 2.7. The challenge is therefore to develop an engine that can meet this efficiency requirement across a wide range of Mach numbers.

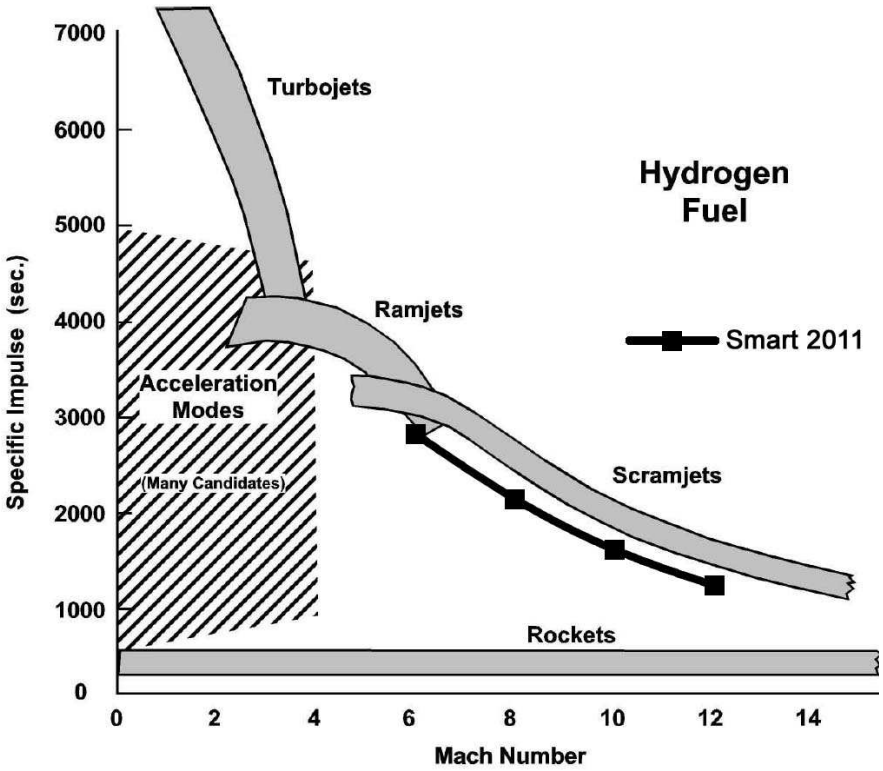

Figure 2.7: Expected uninstalled performance of a scramjet with 80% combustion efficiency. Reproduced from Smart (2012).

Since flow residence times within the engine decrease with increased flight speed, it is believed that developing a M12 REST engine that meets this performance criterion is the ideal choice. Given REST-type engines' strong performance at off-design conditions, a M12 REST engine achieving high combustion efficiencies and net thrust at its design flight condition should be sufficient to produce net thrust at speeds across the access-to-space regime. Seeing as much of this improvement will need to come from

fuel injection and mixing within the engine, our attention thus shifts to the literature available regarding these processes.

2.4 Fuel Injection and Mixing

The current design for the M12 REST engine employs two forms of fuel injection: transverse sonic porthole injection along the inlet surface, as well as boundary layer injection around the circumference of the combustor. The former was chosen to improve fuel residence time inside the engine to improve combustion efficiency. The latter was introduced in an effort to reduce skin friction via film cooling and boundary layer combustion. All mixing of the fuel was passive; the interactions of the fuel jets with the surrounding flow field determined the rate of mixing. If improvements are to be made to combustion efficiency and engine performance, they will most easily come from modifying the injection and mixing arrangements. In order to most appropriately choose injection types and placement, a survey of existing techniques for injecting and mixing fuel is required.

Methods for injecting and mixing fuel in a scramjet may be either active, or passive in nature. Active methods, with the exception of pulsed fuel injection, do not deal with fuel injection so much as they attempt to enhance mixing through forced excitation of the flow (Seiner et al., 2001). Since employing an active method would require significant modifications to the engine and/or moving parts, such methods were excluded from the scope of this project, and will not be mentioned further. Passive methods for injecting and mixing fuel include coaxial/shear layer fuel injection, transverse porthole injection, non-circular porthole geometries, porthole arrays, backward steps and cavities, ramps, strut injection, and film injection.

The geometric mixing methods (e.g. ramps and cavities) can often be combined with, or used as, fuel injectors as well. A good example of such a combination is the injection from the face of a 30° ramp into the recirculating shear flow it generates, enhancing the mixing of the fuel (Bonanos et al., 2009). All of the methods for injection of fuel are wall-based, with the exception of strut injection, and are illustrated in Figure 2.8.

Many early studies into the injection, mixing, and combustion of fuel in the context of supersonic combustion were based on a mixing shear layer. In this arrangement, a stream of fuel was injected parallel to the airstream. This was typically achieved by injection from the base of a strut inside the flow (e.g. Berglund et al. (2010)), or in a coaxial configuration (Baurle and Edwards, 2010). The problem has also been studied

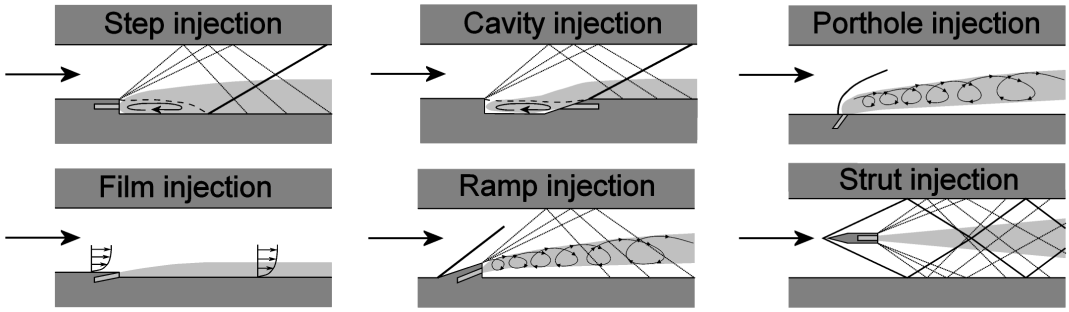

Figure 2.8: Potential methods for injecting and mixing fuel in a scramjet combustor. Reproduced from Turner (2010).

extensively both numerically (the work of Im et al. (1996) for example) and analytically (Pulsonetti et al., 1988)as a problem in supersonic flow. It has been generally found that the mean growth rate of the turbulent shear layer is linear, and the mixing layer grows at a slower rate relative to the shear layer (Curran et al., 1996).

Generally speaking, this shear layer mixing was found to be unable to achieve near-field mixing of air with hydrogen to stoichiometric conditions. The lateral/span-wise vorticity generated by this method is unlikely to give good fuel-air mixing for a reasonable energy cost (Dimotakis, 1991). Streamwise vorticity is therefore necessary to achieve passive mixing of the fuel into the airstream. Such vortices are produced by porthole injection, and ramp injection.

Strut-based porthole injectors (e.g.(Gerlinger et al., 2008),(Genin and Menon, 2010)) and ramp injectors (e.g.(Donohue et al., 1994),(Eklund et al., 1997), (Schumacher, 2000)) have been shown to be effective at mixing fuel, though at a cost of increased flow entropy (Abdel-Salam et al., 2003). Injection into a cavity (Ben-Yakar and Hanson, 2001) has similarly been shown to enhance mixing with available air, and increase fuel residence time within the engine. However, these mixing methods based on physical structures interacting with the flow may be inappropriate for a scramjet operating at a Mach 12 flight condition: at Mach numbers greater than 10, high heating loads and flow losses may render techniques involving cavities or physical protrusions into the mainstream flow infeasible (Billig, 1992). These methods are therefore not examined further in the context of this study, and wall-flush injection methods are emphasized.

2.4.1 Transverse Sonic Porthole Injection

Sonic injection into a supersonic crossflow through a wall-flush circular porthole has been a topic of active research for decades, and is still an area of active research. The

averaged flow-field generated by this injection arrangement is complex, as illustrated in Figure 2.9. Such jets are typically under-expanded, and a barrel shock forms around the core of the jet as it exits the wall and expands into the flow field. This barrel shock is fully contained within the fuel jet, and terminates in a Mach disk. This structure recompresses the fuel after it has been over-expanded by the expansion fans originating at the edge of the porthole, and ultimately matches the static pressure of the jet to the surrounding flow field. This barrel shock structure is, other than an inclination due to the cross-flow velocity, virtually unchanged relative to those observed for an under-expanded jet injected into a quiescent medium (Viti et al., 2009). Surrounding

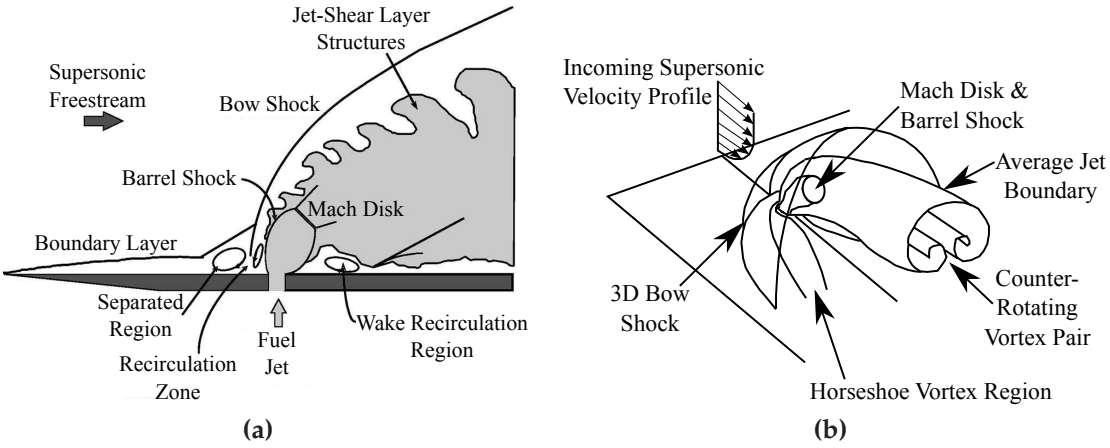

Figure 2.9: Near-field flow structures surrounding sonic porthole injection into crossflow. Adapted from Ben-Yakar et al. (2006)

this barrel shock are several flow features well described by Ben-Yakar et al. (2006) each of which play a role in fuel mixing. These include an upstream bow shock, jet-shear layer vortices above and downstream of the barrel shock; a counter-rotating vortex pair inside the jet plume; a reattachment shock at the terminus of a near-field wake recirculation region behind the barrel shock; a horseshoe vortex that originates between the jet and the bow shock; and additional downstream wake vortices as described by Viti et al. (2009). A separation zone upstream of the bow shock usually causes a second shock to form. This recirculating separation is typically fuel-rich, often containing products from the chemical reaction of fuel with the incoming air stream.

The mixing of a jet into crossflow is efficient enough that, despite the appearances in Figure 2.9, the jet structure is often only coherent 8-10 porthole diameters downstream of the injection point (Gruber et al., 1997). Past this point, the jet structure breaks down into discrete pockets of fuel and a well-mixed fuel-air region. The same phenomenon also occurs approximately 10 porthole diameters downstream of inclined

fuel jets (Ferrante et al., 2011). Despite being a well-characterized flow phenomenon, the complexity of the turbulent flow interactions between the jet plume and the mainstream continue to be a subject of ongoing study (Watanabe et al., 2011), (Peterson and Candler, 2011).

Jet Penetration

When designing fuel injection via wall porthole, one of the most important parameters of interest is how far the jet will penetrate into the cross-stream air flow. The extent of jet penetration directly influences the ability of fuel to mix into the air prior to combustion. The study of jet penetration as it relates to supersonic combustion goes as far back as the 1960s, beginning with the first published study of wall-normal jet injection into a supersonic crossflow (Schetz and Billig, 1966). Several studies followed, each attempting to determine the critical parameters upon which jet penetration depends via differing methods. Billig et al. (1971) developed an analytical theory of jet penetration, while studies such as McClinton (1972), Gruber et al. (1995), and Portz and Segal (2006) have employed various experimental means to independently develop their own correlations for jet penetration.

While the studies above are only a sampling of the total number of studies done (Portz and Segal (2006), for example, compared their own study to several others, and provide direction to several more), a clear agreement in the literature emerges: jet penetration is predominantly influenced by the ratio of the jet and mainstream dynamic pressures; the ratio jet diameter to the local boundary layer thickness; freestream Mach number; and downstream distance from injection. Other variables, such as the molecular weight of the fuel relative to the air, or the jet-to-air Mach number and Reynolds number ratios, are thought to have weak effects on penetration distance. Most correlations thus take a form similar to that of Portz and Segal (2006), reproduced in the Equation 2.1 and Table 2.1 below.

$$\frac{z_p}{D_j} = A\left(\frac{q_j}{q_\infty}\right)^B \left(\frac{x}{D_j} + C\right)^E \left(\frac{\delta}{D_j}\right)^F \left(\frac{\chi_j}{\chi_\infty}\right)^G \tag{2.1}$$

Most correlations, like the one above, are for jets entering the airstream at a 90 degree angle to the direction of flow. There are, however, some studies that deal with predicting the effects of inclining the jet to give it a velocity component in the streamwise direction. The effects of doing so are somewhat counter-intuitive, and are presented below.

Table 2.1: Coefficients for Eqn. 2.1 (Portz and Segal, 2006)

Coefficient	Definition
A	$1.049M_\infty$-0.192
B	-$0.08M_\infty$+0.615
C	-$2.34/M_\infty$
E	$0.406M_\infty^{-0.823}$
F	-$0.067M_\infty$+0.325
G	-0.025

Influence of Injection Angle

Sonic porthole injection need not be confined to injection normal to the wall surface. Several studies into inclined porthole injection into a cross-stream have been conducted, with injector angles as shallow as 15° to the wall surface, and with the jet angled in both streamwise and cross-stream injector planes. In all cases, near-field jet penetration into the freestream is decreased relative to an otherwise identical normal jet. McClinton (1972) showed that further downstream (x/D > 20) from the injector, shallow injection angles experienced enhanced mixing and penetrated further into the freestream relative to an otherwise identical normal jet injector. This was corroborated by the experiments of Kovachevich (2010), and it is thought to be due to the jet retaining a greater portion of its wall-normal momentum since less freestream momentum input is required to turn the fuel plume downstream.

Injection angle may also play a role in the total thrust and specific impulse developed by the engine. At speeds greater than Mach 10, stream thrust analyses have showed that the streamwise component of the injected gas' velocity can give a significant contribution to net engine thrust (Drummond et al., 2002).

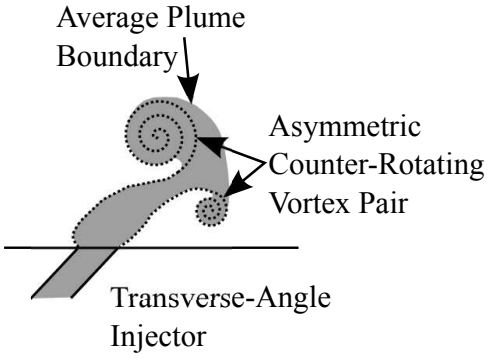

Figure 2.10: A typical fuel plume for sonic fuel inlined in the cross-stream direction

Fuel injection can also be angled in the cross-stream direction, as shown in Fig-

ure 2.10. In the case of such a laterally-inclined jet, the counter-rotating vortex pair becomes significantly displaced from the location of injection. The plume becomes asymmetric, with a smaller vortex close to the wall, and a larger vortex in the mainstream of the flow. This asymmetry increase with shallower injection angles, to the point that the smaller vortex is virtually undetectable (Beresh et al., 2007).

Influence of Porthole Shapes and Arrangements

Up until this point, all discussion of porthole fuel injection has focused on a single circular porthole injecting fuel from a flat wall into crossflow. There are, however, many possible ways of modifying this arrangement to achieve improvements in fuel mixing.

One such modification is to change the shape of the porthole. Examples of this include elliptic portholes (Gruber et al., 2000), diamond-shaped portholes (Srinivasan et al., 2008), streamwise or spanwise rectangular portholes (Wheatley and Jacobs, 2010), stinger portholes (Kouchi et al., 2013), or portholes with even more exotic shapes (Hariharan and Babu, 2014). In all cases, lateral spreading and mixing of the jet appears to be improved, often as a result of the different shock structures caused by the change in porthole shape. However, this improved mixing often comes with decreased jet penetration.

Another method for improving jet penetration and/or mixing is to move away from a single porthole to an array of portholes of equal size. The simplest of these arrays, two portholes aligned in the streamwise direction, has been shown to vastly improve penetration of the downstream jet. This is due to the 'shielding' of the second injection plume by the upstream injector (Lee, 2006). Arrays of multiple jets, such as those examined by Pudsey et al. (2010), showed that dividing a single jet into up to four smaller in-line injectors could give a substantial increase in mixing efficiency for the same fuel mass flow rate. Moving to more than four jets was found to be detrimental to overall fuel mixing, since the scaling of jet penetration with injector diameter ultimately overcame the shielding effects caused by upstream portholes.

Another type of array, the "aerodynamic ramp" porthole array of Fuller et al. (Fuller et al., 1998), was designed to mimic the mixing effects of injection behind a physical ramp, and showed good mixing performance at a lower overall pressure loss relative to its physical counterpart. Transverse injection in conjuction with physical barriers such as a backward-facing step have also been examined (Berman et al., 1983), (Ali and Sadrul Islam, 2006). Most of these studies have focused on normal injection from

the wall downstream of a step, typically located at the start of a model scramjet combustor. In these cases, recirculation zones often form between the jets and the step, and as a result there is often an increase in mixing efficiency relative to injection from a flat plate (Kim et al., 2003). However, fuel injection need not be limited to the scramjet combustor, as will be presented in the next subsection.

2.4.2 Inlet Injection

As its name implies, inlet injection involves the implementation of fuel injection from the compression surfaces of a supersonic inlet. It has often been used in the study and development of detonation-wave and shock-induced combustion ramjets (Sislian, 2000), both of which seek to minimize engine length by initiating a detonation reaction of a well-mixed fuel-air stream using a strong shock wave virtually eliminating the need for a combustor section in the engine.

Inlet injection has also been proposed as a means of improving the performance of traditional (i.e. non-detonation wave) scramjets: by moving injection upstream, the air and fuel are at least partially premixed prior to entering the combustion chamber. This reduces the required combustor length necessary to contain the complete fuel combustion reaction, relative to an in-combustor-only fuelling arrangement (Guoskov et al., 2001). Since skin friction drag in the combustion chamber and thrust nozzle can account for more than 25% of total scramjet engine drag (Tanimizu et al., 2009), reducing this length can provide a positive contribution to engine net thrust. However, care must be taken to avoid robust combustion on the inlet, as this may cause engine unstart. Inlet injection can also be advantageous in that the fuel plume mixing may be enhanced by strong shock interactions with reflected inlet leading edge shocks at the combustor entrance (Peterson et al., 2013).

Gardner et al. (2002) were among the first to investigate scramjet inlet porthole injection experimentally. Hydrogen fuel injection was added to a 2D single-compression-ramp scramjet operating in a Mach 6.5, 3.0 MJ/kg total enthalpy test flow, with a 'cold' engine wall temperature of 300 K, as shown in Figure 2.11.

It was concluded from the experimental pressure data that fuel injected on the inlet did not ignite on the inlet surfaces, but did ignite inside the combustion chamber. Gardner also found, by testing injection angles of 25°, 45°, and 90°, that the engine's thrust levels varied little with change in inlet injection angle. This suggests oblique injection angles are highly desirable for inlet injection; as was mentioned in Section 2.4.1 the

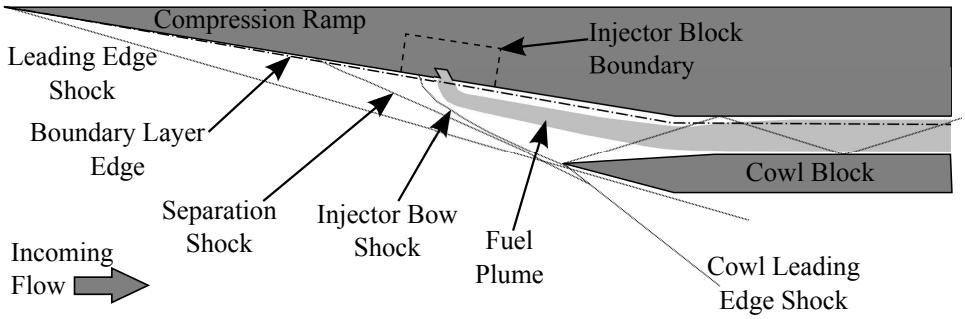

Figure 2.11: The experimental geometry of Gardner et al. (2002)

fuel jet itself could provide additional thrust without adversely impacting combustion efficiency.

Further testing of this engine was done with a heated wall, to ensure that the absence of combustion on the inlet implied by pressure traces was indeed the case, and also to check if combustion was being suppressed by the wall's cold temperature. It was shown conclusively that no OH combustion radicals were formed on the inlet with wall temperatures ranging from 300-700K (Kovachevich et al., 2006). The fuel jet penetrated far enough into the mainstream to be effectively isolated from the high-temperature flow in the boundary layer, explaining the lack of combustion on the inlet. Many subsequent studies have made use of inlet injection across a range of flight Mach numbers (McGuire, 2007), (Lorrain et al., 2012), (Brieschenk et al., 2013), including the inlet of the SCRAMSPACE flight experiment conducted at the University of Queensland (Peterson et al., 2013). In most of these studies inlet injection was the only method of fuelling, and fuelling was done at global equivalence ratios as close as possible to unity without causing the inlet to unstart.

While this fuelling arrangement may ultimately lead to shorter combustion chamber lengths, any improvements in scramjet performance may be nullified by another consequence of inlet injection: significant entropy production may occur when stoichiometric fuelling is done on a scramjet inlet, thanks to the strong bow shock generated (McGilvray and Morgan, 2009). In some cases, the performance loss associated with this entropy gain may match or exceed the benefits gained from a shorter combustion chamber. Moving to a combined fuelling method, in which a smaller ($\phi < 0.5$) fraction of the fuel is injected on the inlet could provide significant benefits with a smaller performance loss: the inlet-injected fuel could ignite early enough to pilot ignition of fuel injected further downstream, generating robust combustion without the large penalty.

An inlet-only fuelled test was conducted on a Mach 8 REST inlet for mildly over-

sped flight conditions (Turner and Smart, 2010). By ensuring adequate porthole spacing in order to avoid fuel interaction with the high-temperature zones created by injector bow shock interactions between the injected fuel jets (Schwartzentruber et al., 2005), the inlet remained started for hydrogen fuelling up to an equivalence ratio of 0.92, though inlet drag was increased. Beyond this equivalence ratio limit, the pressure rise due to combustion was high enough to feed back into the inlet and unstart the engine. The combined fuelling tests of the M12 REST engine at Mach 9 flight conditions (Suraweera and Smart, 2009) showed similar behaviour: the inlet remained started for inlet equivalence ratios up to approximately 0.41, with a pressure rise near the inlet throat possibly indicating fuel ignition upstream of the combustor.

2.4.3 Boundary Layer Fuel Injection

As was alluded to in Section 2.1, mass addition to, and heat release inside of, a boundary layer are methods of active drag control which are readily applicable to the flow inside a scramjet combustion chamber. The action of mass addition with minimal penetration into the cross-flow is often referred to as film cooling, and was originally developed as a means to control heat transfer into the surfaces of turbo-machinery. This mass addition was found to have an additional effect: the lower-momentum injected flow reduced the velocity gradients and momentum near the wall, causing a reduction in skin friction (Parthasarathy and Zakkay, 1970), (Cary and Hefner, 1972), (Schetz and van Overeem, 1975). Heat addition to a hypersonic turbulent boundary layer has been investigated by both Levin and Larin (2003) and Denman (2007). Both studies showed that boundary layer thickness increased with heat release. Both also showed that heat release in a turbulent boundary layer weakened the boundary layer Reynolds (turbulent) shear stresses. The Reynolds stress are responsible for turbulent transport of momentum to the wall, and as much as 90% of the turbulent contribution to skin friction (Gomez et al., 2009). It was also found that, in contrast to film cooling, the release of heat in the boundary layer could as much as double local heat transfer to the walls (Levin and Larin, 2003).

If the mass addition to the boundary layer is done with a fuel, it is possible to combine the effects of film cooling and boundary layer heat release. By doing so in a combustor, the additional fuel injection and combustion can also be harnessed to provide heat release and compression. Boundary layer combustion in a scramjet combustor was first demonstrated experimentally by Goyne et al. (2000), who injected hydrogen tangentially along the wall surface into the turbulent boundary layer of a 1m-long rectangular combustion chamber via a slot in the face of a backward-facing step, as shown

in Figure 2.12. It was found that when the hydrogen burned inside the boundary layer, local skin friction reductions between 70% and 80% were achieved.

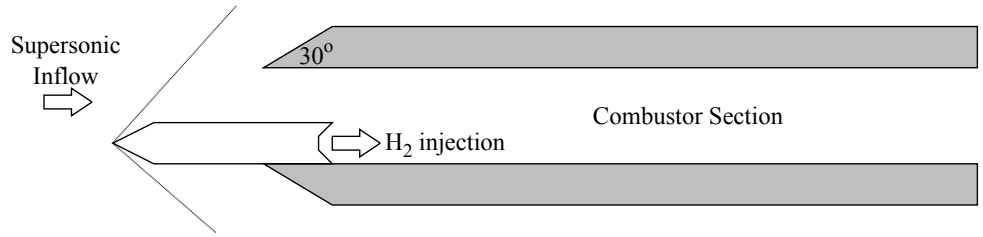

Figure 2.12: The experimental geometry of Goyne et al. (2000)

In contrast to cases in which skin friction reduction is achieved by heat release alone, this experiment showed no noticeable increase in wall heat transfer. In his paper on an analytical model for boundary layer combustion, Stalker (2005) postulated that this is due to the combined effects of film cooling and boundary layer heat release. The cooling effect of fuel being injected along the wall surface counteracts the additional heat flux to the wall due to combustion. Stalker's theory was developed by using Shvab-Zeldovich flow variable coupling (Zeldovich, 1951) with the popular Van Driest II model for compressible turbulent boundary layers (Van Driest, 1956) to model both the injection and combustion of hydrogen in a turbulent boundary layer. The resulting model showed excellent agreement with the data of Goyne (2000) and showed that both film-cooling and heat release effects can combine to create a higher skin friction reduction than either could achieve alone. This model was later extended to show that similar effects could be achieved with hydrocarbon fuels, albeit with smaller film cooling effects (Barth et al., 2013).

Besides the work of Goyne, multiple experimental campaigns have been undertaken at the University of Queensland on slot injection of hydrogen into combustion chamber boundary layers. A parametric study of ignition showed that for low Mach numbers, ignition distance increases with Mach number, but begins to decrease again at higher Mach numbers thanks to viscous heating (Im et al., 1993).The work of Rowan and Paull (2006) showed that tangential fuel injection can be disrupted by the flow structures generated upstream due to tangential porthole injection. They also showed that fuel injected into a boundary layer will not necessarily burn without an external heat source. It was postulated that autoignition of boundary-layer-injected fuel requires some minimum boundary layer thickness to have sufficient thermal energy to cause the fuel to ignite within the length of the combustion chamber (Rowan, 2003). This was confirmed by the work of Kirchhartz et al. (2010), which showed that thick entropy layers and turbulent boundary layers both promote ignition and combustion

of fuel injected into the boundary layer. A later study on the sensitivity of boundary layer injection to upstream disturbances confirmed that film-cooling effects are sensitive to upstream disturbances, but proved that boundary layer combustion effects are relatively insensitive to disruptions caused by upstream flow features (Chan, 2012). It should be noted that the study of Chan generated these disturbances in part by using a REST-type of inlet, and that the inlet and combustor were parallel. This stands in contrast to the REST engine that is examined in this study, in which the combustor is turned relative to the inlet axis. This difference will be shown to have some importance in the present study.

In the work of Chan (2012), and indeed in most of the work done on boundary layer combustion, fuel injection into the boundary layer was done via a slot. However, the M12 REST engine as it is currently designed utilizes injection from several discrete portholes equally spaced along the face of a backward step, as was mentioned in Section 2.3 and illustrated in Figure 2.13. This stands in contrast to the injection behind a backward step described in Section 2.4.1. While no in depth studies have been completed on this geometry to date, preliminary work has indicated that boundary layer combustion performance is sensitive to porthole spacing: portholes spaced 2D or less apart will form a coherent sheet of fuel similar to that of a slot injector, while increased spacing may improve mixing and combustion efficiency of the fuel in the boundary layer (Barth et al., 2012).

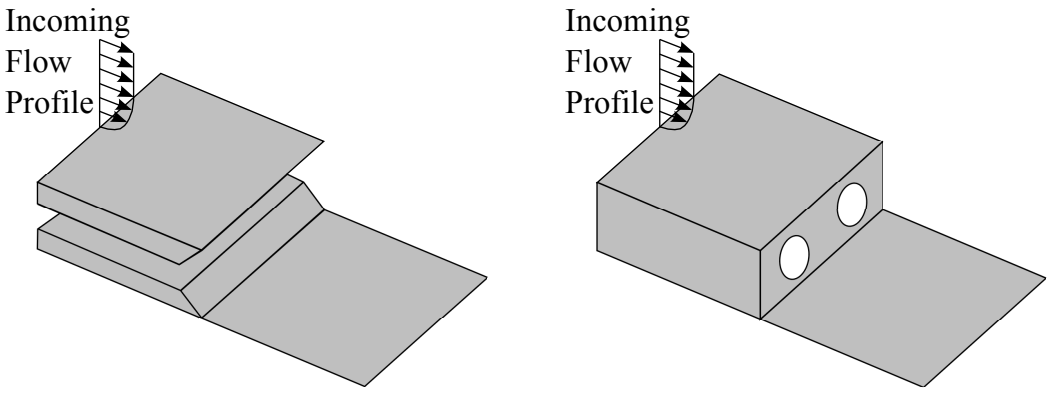

(a) Slot Injection **(b)** Porthole Injection

Figure 2.13: A comparison of boundary layer injection via slot versus porthole

2.4.4 Tailored Fuel Injection

Much like other 3D stream-traced inlets, it is believed that the compression flow fields generated by the M12 REST inlet should lead to asymmetric combustion. One study of

a sugar-scoop inlet concluded that injection strategies should be developed to exploit the flow structures developed by the inlet (Malo-Molina et al., 2010). This "tailoring" of the fuel injection to the detailed engine flow field is an as-yet unresearched area in scramjet combustion, likely because much of the literature is still primarily focused on the basic physics of specific flow structures or interactions, rather than on system-level studies of entire engines.

An examination of some of the flow physics expected to occur in the REST inlet was undertaken to inform decisions made when attempting to tailor injection to the inlet flow. The effects of combustion on a turbulent flow field were presented in Section 2.1.1. This effect is certainly important, likely particularly in the case of the M12 REST engine with inlet injection: the combustion of this fuel likely has a significant effect on the combustor flow field. Other effects that should be taken into account are those caused by 3D flow structures inside the engine, as well as shock wave-related effects.

3D Flow Structures

Many of the flow structures that are expected to be typical of flow in a REST-class engine are a product of the interactions of bodyside and sidewall compression fields, and shock waves generated by the leading edges. The flow is likely dominated by streamwise vorticity and high helicity(Ingenito and Bruno, 2010). These effects can be modelled well using the analogous flow between two symmetric fins mounted on a plate, as shown in Figure 2.14a.

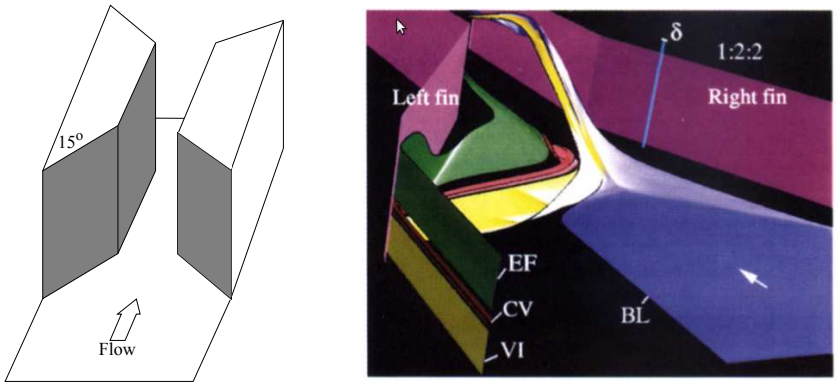

(a) Geometry **(b)** Flow Structure

Figure 2.14: Double fin geometry and streamlines. Reproduced from Gaitonde et al. (1995)

In this flow, the incoming boundary layer interacts with the oblique shock waves generated by the fin leading edges. The boundary layer separates across the entire width of the channel, with the greatest degree of separation at the channel centreline.

Rather than forming a recirculation zone, flow from the side walls is entrained into the zone behind the separation, ultimately forming a counter-rotating pair of streamwise vortices along the centreline shown in Figure 2.14b (Gaitonde et al., 1995). This geometry does not account for compression on the wall surface, but experiments on a 2D scramjet inlet with the addition of sidewalls shows that much the same flow structures are observed(Häberle and Gülhan, 2008).

Separations of this sort are also present in stream-traced inward turning inlets such as the "sugar scoop" circular capture area inlet (Gaitonde et al., 2007). This can lead to a curious flow pattern, in which flow through the engine core may be moving slower than the flow along the sides. If such a zone exists in the M12 REST engine, this could be beneficial to the penetration of fuel into the core of the engine flow.

Shock Wave Effects

The interaction of shock waves with other flow structures could be a major feature of the M12 REST engine flow path. Much of the combustor flow is expected to consist of mixing regions in which fuel and air are being combined and eventually reacting. It has been previously shown that shock waves impinging onto mixing regions tends to improve the overall mixedness of the flow; flame lengths have been observed to decrease significantly(Huh and Driscoll, 1996). The same study noted that the best results occur when the primary shocks are positioned to interact with the flame base, such that downstream recompression shocks interact with the central portion of the flame, providing additional mixing between reacting fuel and available unreacted air.

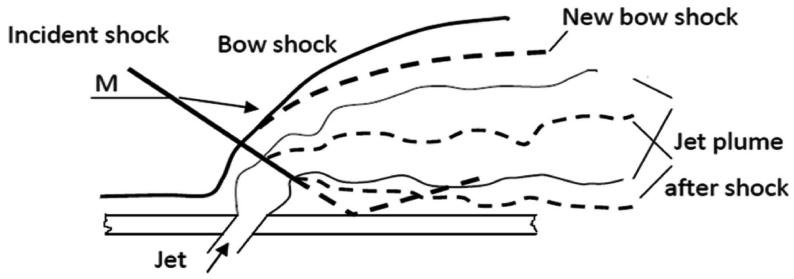

Figure 2.15: Schematic of an oblique shock/porthole injector interaction. Reproduced from Schetz et al. (2010).

Similar effects have been observed for shocks interacting with fuel injection stations, shown in Figure 2.15. The addition of shock waves immediately upstream or downstream of a transverse porthole injection significantly reduces jet penetration into the mainstream, but also increases mixing (Schetz et al., 2010). This could also be impor-

tant for ignition, as the shock will increase the local temperature and pressure in the vicinity of the injector.

Shock-boundary layer interactions (SBLI) can also have a strong effect on flow structures and combustion. In the case where an oblique shock impinges onto a fuel-rich boundary layer, ignition has been shown to occur just behind the shock reflection point (Burtschell and Zeitoun, 2004). In addition to the interaction described previously, the interaction of a shock and boundary layer can have multiple effects, including causing boundary layer separation. An excellent summary of research into both two and three dimensional SBLI can be found in the work of Zheltovodov (1996).

2.5 Summary

It is clear that undertaking the study of an entire engine flow path relies upon a wide swath of fundamental research on hypersonic flow, supersonic combustion, and those phenomena specific to fuel injection and mixing. In order to improve the combustion efficiency and performance of the M12 REST engine, the fundamental flow field of the REST inlet must be understood (both with and without fuel injection), as well how the flow evolves through the combustor of the engine. Specific flow structures within the M12 REST need to be identified and characterized, in order to allow for fuel injection tailored to exploiting these structures to be attempted.

The study that follows therefore undertook not only to understand the combusting flow field of the engine as it is currently designed, but also to attempt to improve the engine's performance by investigating the feasibility of tailored fuel injection. In order to undertake this study, both advanced computational simulations and experimental validation were required. The methodology by which this study was undertaken are presented in the next chapter.

3

Methodology

"A good rule for rocket experimenters to follow is this: always assume that it will explode."

– *Astronautics Magazine,* Issue 38 (October 1937)

This chapter presents the methods by which data were collected for the analysis that follows in later chapters. The study was primarily numerical in nature, utilizing computational fluid dynamics (CFD) tools to simulate the reacting flow inside of the M12 REST engine. The study also involved an experimental investigation of the M12 REST in order to provide validation data for comparison with the numerical solutions obtained. These two methodologies are both presented in the sections that follow, preceded by a description of the M12 REST flow path, and followed by a description of how the flow conditions for both methods were obtained.

3.1 Flow Path Geometry

As was mentioned in Chapter 2, the flow path geometry of a REST engine is developed for a desired flight Mach number and dynamic pressure, which for this study were Mach 12 and q = 50kPa. The true design Mach number of the inlet was then determined by the assumption of a vehicle forebody shape and angle. In the case of the M12 REST, it was assumed that the engine was to be mounted on a vehicle with a planar forebody at a 6° angle relative to the incoming air flow when flying with zero vehicle angle of attack. Designed according to the procedure of Smart (1999), the resulting stream-traced flow-path is shown in Figure 3.1, and is labelled using the terms defined in Table 3.1. This terminology contained in this table will be used from this point forward.

As was mentioned in the previous chapter, this engine flow path was identical to the geometry of Wise (2014) (save for the geometric details of fuel injection in the

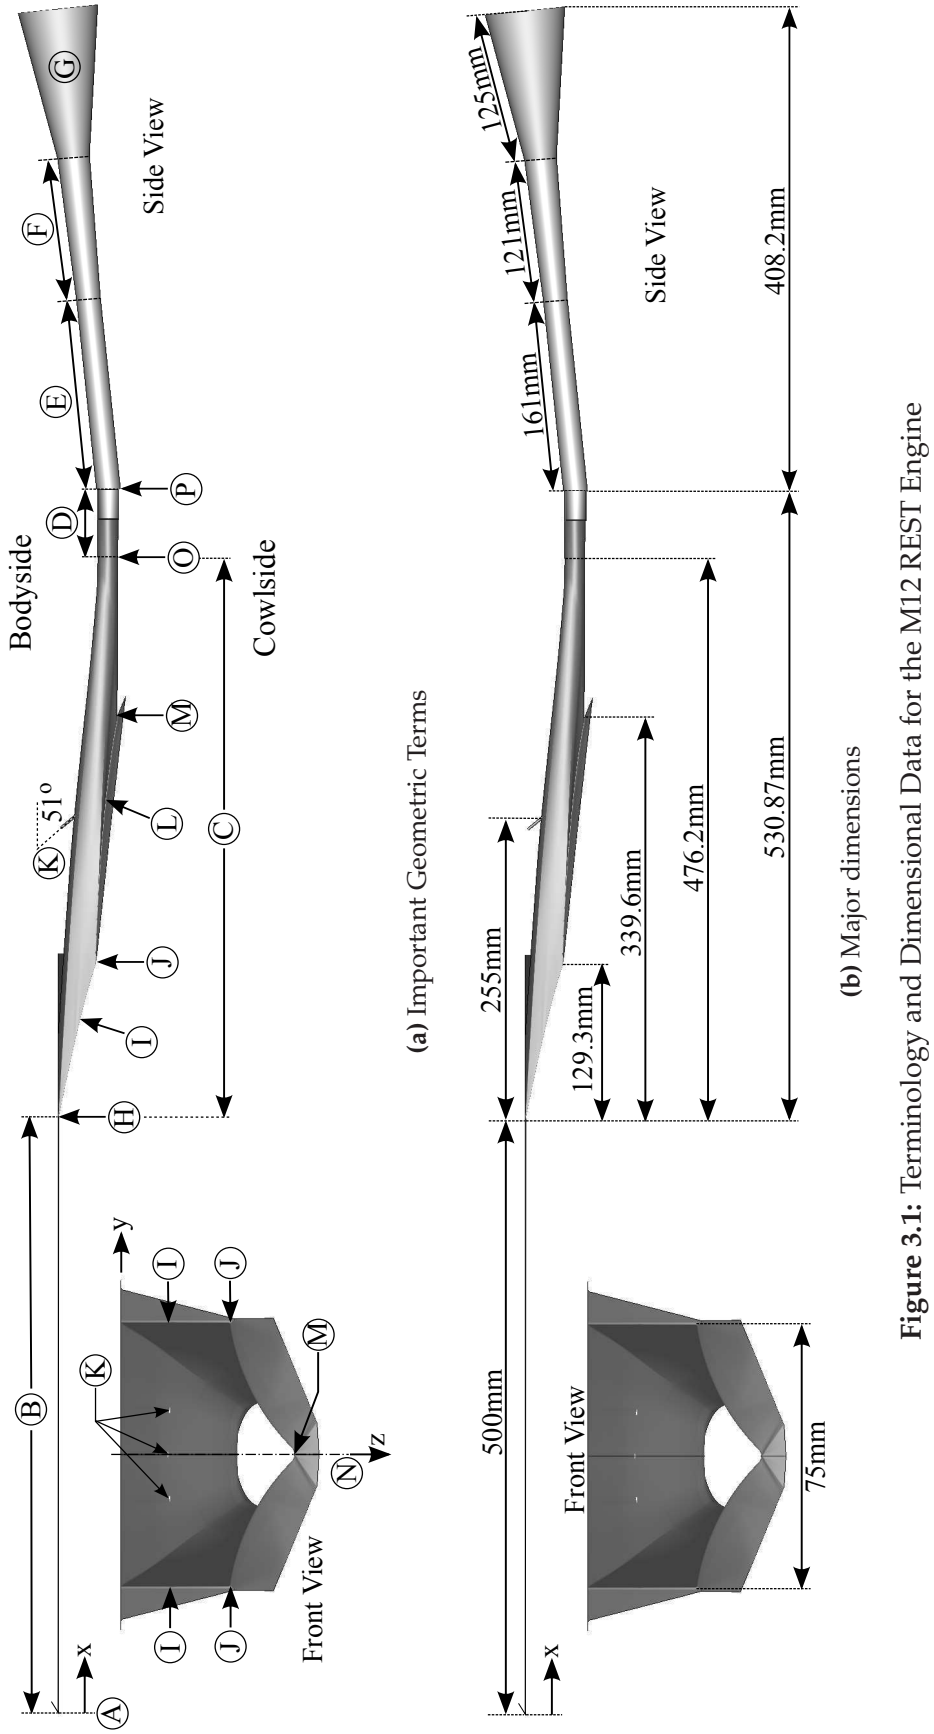

(a) Important Geometric Terms

(b) Major dimensions

Figure 3.1: Terminology and Dimensional Data for the M12 REST Engine

Table 3.1: Geometric terminology associated with labels in Figure 3.1a

Label	Term Used	Label	Term Used
A	Leading Edge	I	Sidewall Leading Edge
B	Forebody	J	Side-Cowl Corner
C	Inlet	K	Inlet Injectors
D	Isolator	L	Cowl Leading Edge
E	Constant-Area Combustor	M	Cowl Notch
F	Diverging Combustor	N	Centerplane
G	Nozzle	O	Throat
H	Inlet Leading Edge	P	Combustor Step

combustor), which itself was a half-scale version of the geometry of Suraweera and Smart (2009). Including the planar forebody, the entire geometry was 1435 mm in length. The scale of the engine was chosen in order to meet the limitations of the experimental facility, as it was determined that the length of the original engine was too long to achieve fully-established flow within available test time (Wise, 2014).

On the highest level, the M12 REST engine can be broken down into four sections: a planar forebody, the inlet section, the combustor, and the nozzle. The forebody was 500 mm in length, and significantly wider than the inlet capture area to mitigate the possibility of edge effects affecting the inlet-captured air during experiments. The inlet section was 476.2 mm in length, beginning at the leading edge of the body and sidewall compression surfaces (the “inlet leading edge”) and terminating at the point of smallest cross sectional area ("the throat"). The shape transition is not entirely complete at this point, and continues into the isolator. The isolator reaches its final elliptical cross-section shape 505.8 mm from the inlet leading edge, and remains constant until the start of the combustor 25 mm downstream.

The lip of the inlet sidewall turns inward at the side-cowl corner to form the start of the cowl wall 129.3mm downstream of the inlet leading edge. The inlet closed completely at the cowl notch, 339.6 mm downstream of the inlet leading edge. The inlet capture width is 75 mm and contracts to an ellipse 31.05 mm wide with an aspect ratio of 1.76, giving the inlet a total geometric compression ratio of 6.61, with an internal compression ratio of 2.25. As was mentioned in the previous chapter, this highly swept geometry allows the engine to spill excess air mass overboard, permitting the inlet to be self starting across a range of Mach numbers.

The leading edges of the inlet were rounded with a double radius. The rounding of the leading edge into the internal flow path had a near-sharp radius of 0.25 mm.

The radius of curvature applied to the leading edge on the external surface side had a radius of 0.5 mm. This external radius was a part of the original design of the REST flow path, whereas the internal radius was given a typical tool-nose radius in order to facilitate computational mesh generation. The leading edge radii were both small enough that they were expected to have a negligible effect on inlet mass capture and efficiency (Drayna et al., 2006).

The combustor section began at the end of the elliptical isolator. The isolator had an aspect ratio of 1.76, and a total cross-sectional area of 430 mm^2. This isolator joined to the combustor 530.86 mm downstream of the inlet leading edge at the lip of the 1.25 mm high circumferential backward-facing combustor step. This coincided with a 6° turn in the flow path in order to redirect the flow toward the nominal direction of flight (recall, the vehicle forebody this engine was assumed to be 6° to the nominal flight direction at zero angle of attack). This step was originally included to facilitate the injection of fuel into the combustor boundary layer, which was employed in the experiments of Wise (2014). The combustor consisted of two sections: a 161 mm long constant-cross section elliptical duct, followed by a diverging section 121 mm in length that expanded the flow uniformly at 1.6° for a total expansion ratio of 2 relative to the throat. The diverging combustor section then fed into the 125 mm long nozzle section, which was a uniform conic expansion with an upstream-to-downstream area ratio of 5. This gave the engine thrust surfaces a total expansion ratio of 10 relative to the throat by the time the flow exited the engine. This area ratio was constrained by space available in the experimental facility, which limits the potential for thrust production by the nozzle.

3.1.1 Fuel Injection Geometries

The M12 REST engine utilized multiple forms of fuel injection, depending on the case examined. However, there were some commonalities to all engine configurations examined. The most notable of these commonalities was the consistent use of inlet injection, shown in Figure 3.2.

In the M12 REST inlet, three 2mm diameter portholes were located on the body-side inlet compression surface, approximately 255mm downstream of the inlet leading edge. The injectors were angled at 45° to the local surface normal (51° to the geometry x-axis), in the streamwise direction. One of the three portholes was located on the engine centerplane, while the other two were located on either side of the central injector, with a center-to-center spacing of 12.5mm to ensure minimal jet interaction and/or blockage of the incoming boundary layer.

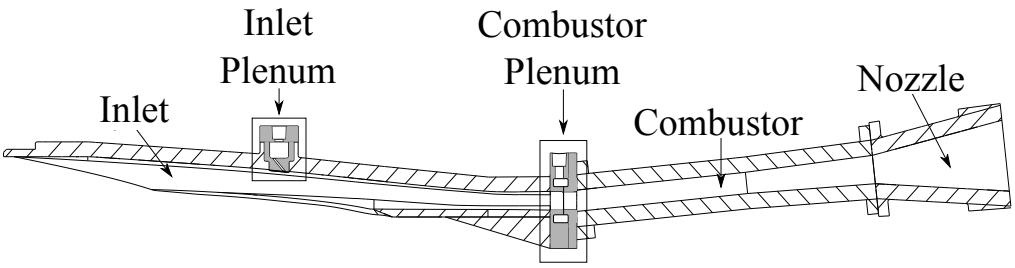

(a) Plenum Locations

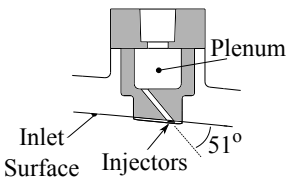

(b) Inlet Plenum

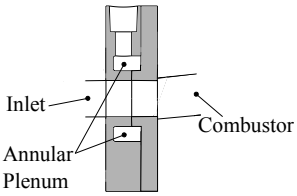

(c) Combustor Plenum

Figure 3.2: Injection Plenum Locations and Details, adapted from Wise (2014)

In experiments, the three portholes were fed by a common fuel plenum, as shown in Figure 3.2b. There was also a second fuel plenum: the elliptical annular plenum shown in Figure 3.2c which wrapped around the isolator section just upstream of the engine combustor. This plenum was originally designed to feed 36 x 0.9mm portholes placed circumferentially around the backward-facing step of the combustor entrance, which provided boundary layer injection of fuel in the experiments of Wise (2014). This plenum was also used as a design constraint during the tailored injection study: the tailored mainstream injection attempted in this study had to be fed from this plenum.

3.2 Computational Approach

The majority of the investigations undertaken in examining the M12 REST for this study were computational, and conducted with a state of the art compressible flow solver, US3D. This section describes the flow solver, the details of the chosen discretization parameters, meshing of the flow path geometry, and the grid convergence behaviour of the meshed M12 REST engine.

3.2.1 The US3D Flow Solver and Simulation Details

The numerical simulations run to provide the results presented in the chapters that follow were performed using the US3D flow solver, developed at the University of Minnesota (Nompelis et al., 2004). US3D is a highly-scalable parallel CFD solver that

was developed to simulate hypersonic flows, and is capable of solving the Navier-Stokes equations on structured, unstructured, and hybrid meshes. The first of these equations is the species-specific mass conservation equation

$$\frac{\partial \rho_s}{\partial t} + \nabla \cdot (\rho_s \vec{U}) = -\nabla \cdot (\rho_s \vec{v}_s) + W_s, \tag{3.1}$$

where ρ_s is the density of an individual chemical species s, $\vec{v}_s$ is the species' diffusion velocity, and W_s is the rate of species production due to chemical reactions in the flow. If all species mass conservation equations are summed, the species densities will add up to the mixture density (ρ) and the source terms on the right hand side of equation 3.1 sum to zero. The traditional continuity equation is thus recovered in the solver.

The conservation of linear momentum is expressed as

$$\frac{\partial (\rho \vec{U})}{\partial t} + \nabla \cdot (\rho \vec{U} \otimes \vec{U} + PI) = \nabla \cdot \tilde{\tau}, \tag{3.2}$$

where P is the thermodynamic pressure, I is an identity matrix vector, and $\tilde{\tau}$ is the viscous stress tensor. As it is written in vector form, this equation represents the momentum equation in each spatial dimension in the flow solver. The solver assumes no external body forces are exerted on the fluid control volumes, and that the diffusion velocities sum to a null contribution as stated previously.

The conservation of total energy of the fluid mixture is formulated as

$$\frac{\partial E}{\partial t} + \nabla \cdot ((E + P)\vec{U}) = \nabla \cdot (\tilde{\tau}\vec{U}) - \nabla \cdot (\vec{q_t} + \vec{q_r} + \vec{q_v}) - \nabla \cdot \sum_{s=1}^{ns} (\rho_s h_s \vec{v}_s), \tag{3.3}$$

where $\vec{q_t}$, $\vec{q_r}$, $\vec{q_v}$ are the translational, rotational, and vibrational heat flux vectors, and h_s is the total enthalpy per unit mass of species s. US3D assumes that vibration-vibration energy coupling is strong, such that all vibrational modes are in equilibrium and can be described with a single vibrational temperature. The equation for vibrational energy per unit volume is

$$\frac{\partial E_v}{\partial t} + \nabla \cdot (E_v \vec{U}) = -\nabla \cdot \vec{q_v} - \nabla \cdot \sum_{s=1}^{ns} (\rho_s e_{v,s} \vec{v}_s) + W_v, \tag{3.4}$$

where $e_{v,s}$ is the vibrational energy per unit mass of species s, and W_v is a source term. More about this equation, and the various terms presented in all the equations presented above, as they are solved in US3D can be found in by consulting the work of Nompelis (2004).

Inviscid fluxes were computed using the modified Steger-Warming flux vector splitting method, while viscous fluxes were computed using the central difference MUSCL scheme on conserved variables, turbulent viscosity, and vibrational temperature. Both inviscid and viscous fluxes were second-order accurate schemes. Cell-centered gradient reconstruction was also performed by the code, using a weighted-least-squares reconstruction method.

The solver was run to convergence using an implicit time stepping scheme that was first order in time, and made use of the Full Matrix Point Relaxation method. Time steps were calculated for each iteration based on the CFL number used, and the smallest combination of physical cell size, flow velocity magnitude and sound speed, using the relation

$$\Delta t = CFL \times \frac{\Delta x}{U + a}. \tag{3.5}$$

The simulation CFL number was changed dynamically to allow short time steps during startup and/or restart of the simulations, and larger time steps once the large gradients common in startup (or during changes to simulation boundary/solver conditions) had been smoothed out. In practice, as the simulation converged, effort was made to ensure the last entire flow-through of a model had time steps on the order of 1-5x10^{-8} s, as ignition delay time has been found to become inaccurate when larger time intervals are used (Gerlinger et al., 2010).

US3D is able to run simulations in fully laminar, fully turbulent, or tripped turbulent modes. While the solver is capable of solving turbulent flows using unsteady Large-Eddy Simulation turbulence modelling, this method was too computationally expensive for a study of this scope. As will be discussed later, the full flow path RANS simulations conducted in this experiment consumed all available computational resources. US3D was therefore run using the Reynolds-Averaged Navier-Stokes (RANS) method, to provide time-averaged flow solutions for this study.

To close the discretized compressible Reynolds-Averaged Navier-Stokes equations, US3D employs the Spalart-Allmaras one-equation turbulence model (Spalart and Allmaras, 1992) modified by the compressibility corrections of Catris and Aupoix (2000), which in vector notation is

$$\frac{\partial \rho\hat{\nu}}{\partial t} + u_j \frac{\partial \rho\hat{\nu}}{\partial x_j} = c_{b1}\hat{S}\rho\hat{\nu} - c_{w1} f_w \rho\left(\frac{\hat{\nu}}{d}\right) + \frac{1}{\sigma}\frac{\partial}{\partial x_j}\left(\mu\frac{\partial \hat{\nu}}{\partial x_j}\right) + \frac{1}{\sigma}\frac{\partial}{\partial x_j}\left(\sqrt{\rho}\hat{\nu}\frac{\partial \sqrt{\rho}\hat{\nu}}{\partial x_j}\right) + \frac{c_{b2}}{\sigma}\frac{\partial \sqrt{\rho}\hat{\nu}}{\partial x_i}\frac{\partial \sqrt{\rho}\hat{\nu}}{\partial x_i} \tag{3.6}$$

where $\hat{v}$ is the transported coefficient of turbulent eddy viscosity, d is the distance from the field point to the nearest wall, and c_{b1}, c_{b2}, σ, and f_w are constants (or functions of constants, d, and a sub-quantity r). $\hat{S}$ and r were defined according to the Edwards correction, which holds that

$$\hat{S} = \bar{S}^{0.5}\left(\frac{1}{\chi} + f_{v1}\right) \tag{3.7}$$

and

$$r = \frac{\tanh[\hat{v}/(\hat{S}\kappa^2 d^2)]}{\tanh(1.0)}, \tag{3.8}$$

where $\chi = \hat{v}/\nu$, $\kappa = 0.41$ is the von Kármán constant, and f_{v1} is a function of χ and a constant (Edwards and Chandra, 1996). It should be noted that this method does not make use of the vorticity magnitude Ω, but rather,

$$\bar{S} = \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)\frac{\partial u_i}{\partial x_j} - \frac{2}{3}\left(\frac{\partial u_k}{\partial x_k}\right)^2. \tag{3.9}$$

The Edwards correction is primarily used to improve the near-wall numerical behaviour of the model.

In practice, simulations were run in fully turbulent mode, with solutions initialized to have a low-level of freestream turbulence: turbulent viscosity μ_t was set to be 1% of laminar dynamic viscosity μ. This was the default value for turbulent viscosity in the version of US3D used in this study (RC 14), and also matched reasonably well with the calculated turbulence intensity at the exit of the Mach 10 nozzle (Doherty et al., 2012a). The effects of this modelling choice are discussed in Chapter 4. The turbulent Schmidt and Prandtl numbers were set at 0.7 and 0.9, respectively.

US3D is also capable of including vibrational excitation and chemical kinetics in its simulations. Vibrational non- equilibrium is modelled through Millikan-White two-temperature (T-Tv) collisional relaxation model:

$$p\tau_v = exp[a(T^{-\frac{1}{3}} - b) - 18.42] \tag{3.10}$$

where $a = 0.00116\varphi^{\frac{1}{2}}\theta^{\frac{4}{3}}$ and $b = 0.015\varphi^{\frac{1}{4}}$. In these equations, τ_v is the characteristic relaxation time, φ is the equivalent molecular weight between the colliding particles, and θ is the characteristic vibrational temperature of the relaxing particle (Millikan and White, 1963).

Where appropriate, values from experimental/theoretical studies were substituted for the coefficients a and b in the characteristic vibrational relaxation time equation.

The substitutions made are listed in Table 3.2 below. While Park (1994) includes many potential substitutions for the coefficients, those actually used in the model were found by Gehre et al. (2012) to accurately model the collisional vibrational relaxation of a Nitrogen-Oxygen mixture.

Table 3.2: Particle-Partner specific Millikan-White constants

Molecules Colliding	*a*	*b*	Reference
H_2 - H_2	80.6	0.01975	Dove and Teitelbaum (1974)
N_2- O	72.4	0.01500	Park *et al.*(1994)
O_2 - N	72.4	0.01500	Park et al.(1994)
O_2 - O	47.7	0.05900	Park et al.(1994)
NO - NO	49.5	0.04200	Park et al.(1994)

The Dove and Teitelbaum (1974) substitution for H_2-H_2 collisional vibrational relaxation was made based on the fact that the default characteristic temperature for molecular hydrogen differed from both calculated collisional exchange rates for other particles and experimental values for hydrogen relaxation time by two orders of magnitude. This difference is likely due to the fact that the vibrational relaxation model was originally developed for molecules far heavier than hydrogen, such as Nitrogen and Oxygen, which have far lower characteristic vibrational temperatures. When the new values for Hydrogen self-relaxation were implemented, it was observed that the solutions for cold hydrogen injection into a vibrationally excited Nitrogen/Oxygen crossflow became more resistant to numerical errors.

Flow chemistry was modelled using the standard finite-rate chemistry formulation, where the reaction rates are typically calculated from the expanded Arrhenius form

$$k = AT^B e^{-\frac{E}{RT}} \tag{3.11}$$

where A is a coefficient, B is a temperature exponent, E is the activation energy for the reaction, R is the universal gas constant, and T is the reacting mixture's temperature. There were no turbulence-chemistry interaction models implemented in the version of US3D utilized, and as such chemistry was calculated on the assumption that the flow inside any cell in the domain was perfectly mixed. The finite rate scheme used to calculate chemical reactions was the 13-species, 33 hydrogen-air reaction model proposed by Jachimowksi(1992). The forward reaction rate coefficients for this model are listed in Appendix B. Backward rates were calculated dynamically by US3D using the model's forward rates and equilibrium rates calculated from the NASA Lewis thermochemi-

cal database, via NASA's Chemical Equilibrium with Applications software (McBride et al., 2002).

3.2.2 The M12 REST Engine Mesh

The M12 REST engine geometry was modelled with a fully structured mesh, using the Program Development Company's grid generating software GridPro (2014). Taking advantage of the lone symmetry plane, the model consisted of half of the full engine flow path. GridPro's unique software allows for structured grids to be rapidly built and smoothed around complex geometries, and was ideally suited to the challenge of meshing the M12 REST's highly swept inlet geometry. Despite the power of this software, creating a mesh for the M12 REST inlet that was suitable for use with a low-dissipation CFD solver like US3D was still a challenging task. Details on the procedure followed to generate the inlet mesh can be found in Appendix A. The resulting M12 REST mesh is presented in Figure 3.3.

The structured mesh sweeps through the inlet, and gets finer in the vicinity of leading edges and injectors through the creation and destruction of blocks outlined in Appendix A. The inlet mesh was held constant once its quality was established, and contained 20,643,712 cells. Due to the variety of injection cases investigated, the combustor mesh shown in Figure 3.3c above is presented without any injection included. The meshes used for each specific injection case are presented in subsequent chapters, when each unique configuration is introduced. Typically, a combustor mesh including injection contained between 20 and 25 Million cells, leading to a full flow-through mesh on the order of 45 Million cells.

Cell Spacing

In general, the M12 REST meshes were generated such that the average hexahedral cell zones located away from walls and injectors was 0.25 mm on each side, though there was significant variance in cell size due to the smoothing algorithms employed by GridPro. Inside the fuel injectors tubes, cell sizing was reduced such that each cell was approximately 0.05 mm on a side. In order to adequately resolve the flow features around the jet injectors, cell edge length was held to 0.1 mm inside a region that extended upstream 2 jet diameters (D) to capture the recirculation zone, and 30D downstream to capture the evolution of the jet plume. The zone extended away from the wall to a distance based on the nominal calculated jet penetration height, with a "safety factor" of 10% added to ensure the entire plume was captured.

(a) Symmetry Plane View

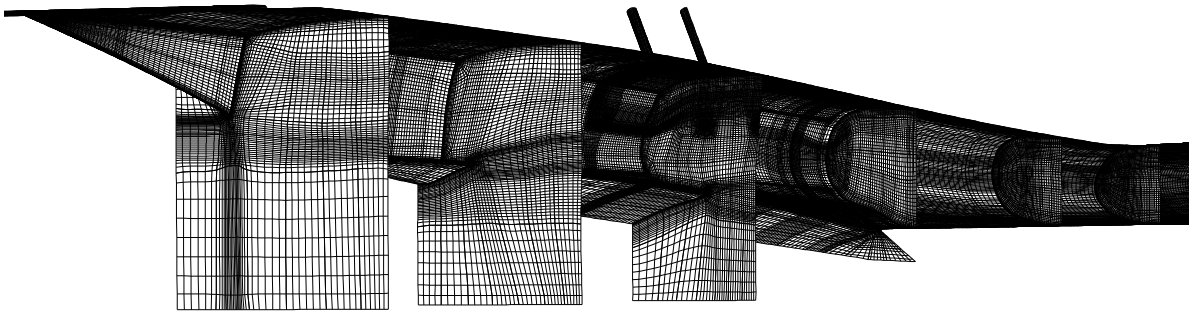

(b) Inlet Mesh Detail

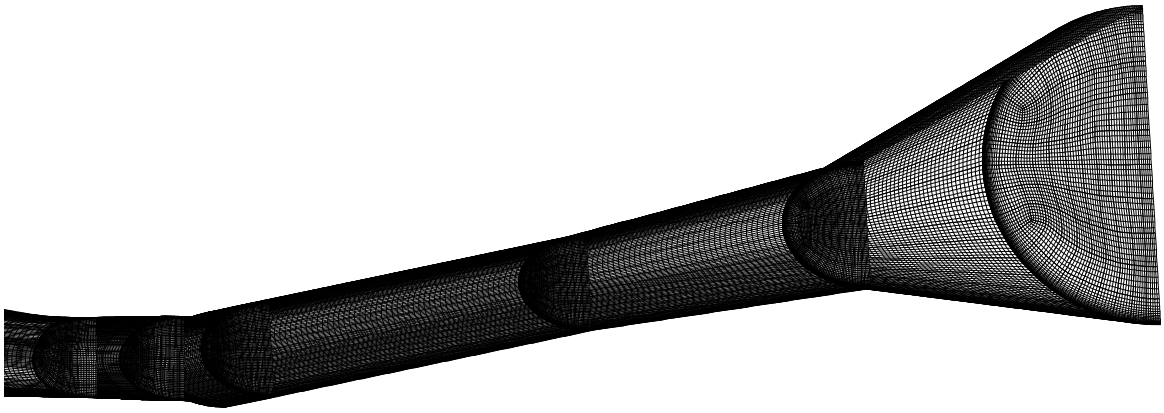

(c) Combustor Mesh Detail

Figure 3.3: Details of the M12 REST Engine Computational Mesh

The mesh was originally meshed as an Eulerian grid, ignoring viscous grid requirements adjacent to walls. Once this inviscid grid had been smoothed, the GridPro utility "clu" was used to add viscous near-wall cell-spacing to the grid. In the inlet, wall-adjacent cells had an off-wall heigh of 0.9μm. Downstream in the combustor, the first-cell spacing started at the same value, but blended down to 0.7μm midway through the constant-area combustor section in order to account for the thinning of the boundary layer along the cowl-side wall of the combustor section.

This cell-size at the wall was sufficient to achieve a y^+ of unity or less in most of the M12 REST engine. There were some exceptions: typically, in the vicinity of a shock reflection or a region of strong compression like the leading edges, y^+ could occasionally reach a value as high as 2.5 (though typically it remained less than 2 except near the reflection of the strongest shocks). The biggest region in which it was impossible to keep y^+ at, or below, two was at the cowl closure notch. This region was both a stagnation zone, and downstream of multiple shock reflections. As such, static temperatures in this region were as high as 5500K, with densities two orders of magnitude larger than the incoming flow. As such, in order to get a y^+ less than 2, cell heights next to the wall would need to be on the order of 1 nm. Properly resolving the mesh to the wall in this region was deemed intractable. Fortunately, the affected region around the notch is quite small; 2.41 x 10^{-6} m^2, or 0.008% of the inlet area. Flow residence times within this region are therefore short, and only a small portion of the flow would pass through it. It was therefore decided that the first-cell spacing from the rest of the inlet would be maintained here. Any effects of this mesh design decision will be determined when examining grid convergence.

Boundary Conditions

The boundary conditions applied to the M12 REST included a symmetry condition at the engine centerplane, and multiple supersonic outflow conditions on bounding surfaces around the inlet and at the exit of the engine thrust nozzle. The walls were given a no-slip, isothermal boundary condition with a wall temperature of 300 K. This constant temperature was chosen to reflect the fact that almost all the simulations were conducted for tunnel conditions, in which the test-time is short enough that viscous heating of the walls are thought to be negligible.

The inflow boundary condition that provided the mainstream air flowing into and over the engine was given a characteristic supersonic inflow condition based on flow velocity, density, temperature, vibrational temperature, and chemical composition. The values for this inflow condition varied for each simulation, and the method by which

they were determined will be discussed in Section 3.4. Fuel injector inflow boundaries were specified using a sonic mass flow boundary condition that required flow velocity, and the stagnation pressure, temperature and vibrational temperature were specified for an assumed stagnation region upstream of the inflow boundary. This condition as it is a good analogue to the fuel plenums that feed the injectors in the experiment, and because the mass flow rate through the inlet plane scales linearly with assumed stagnation pressure, allowing the mass flow rate to be easily set to match an experiment. Inflow velocity for the injector inflow boundary condition was specified as the speed of sound for hydrogen at choked conditions, and was calculated using the plenum temperature and the isentropic flow relation:

$$\frac{T_{sonic}}{T_0} = \left(1 + \frac{\gamma - 1}{2}\right)^{-1}. \tag{3.12}$$

When using this equation, it was assumed that the specific heat ratio, γ, was equal to 1.4. It was further assumed that the fuel inside the fuel plenum was in vibrational equilibrium, and as such, the vibrational temperature was set equal to the plenum static temperature of 300 K.

3.2.3 Grid Convergence and Experimental Validation

In order to have confidence in the results provided by computational tools such as US3D, it was necessary to prove that the grid resolution is sufficient to minimize the effects of discretization on the simulated flow properties. For RANS simulations such as those undertaken in this study, this is often accomplished by conducting a grid refinement study to prove that the solution variables of interest are asymptoting toward a theoretical value that would be generated by an infinitely fine mesh. Wherever possible, solutions are also compared to experiments as another form of validation. Both methods were employed in this study, and merit further discussion.

Grid Convergence Behaviour

A grid refinement study was performed on the mesh of the forebody and inlet of the M12 REST. The baseline mesh for this portion of the engine consisted of a structured mesh containing 20.6 million cells. In order to determine grid convergence behaviour, coarser meshes of 5 million and 12 million cells were created. All three meshes were run to convergence with the same inflow and inlet fuelling conditions, with chemical reactions and vibrational non-equilibrium effects turned on. The convergence parameters for the fuelled inlet performance were taken to be the fuel jet penetration height

(taken 10D downstream of the injection point); combustion efficiency at the inlet throat; and the integrated inlet viscous (skin friction) drag. The three solutions appeared to show monotonic behaviour, and the grid convergence was estimated normally. The results of this analysis are presented in Table 3.3.

Table 3.3: Grid Convergence Indices

Variable (f)	p_{calc}	p_{est}	GCI_{12}	f_{exact}
y_p	25.65	2	0.0218	7.27mm
$F_{drag,viscous}$	8.32	2	0.0002	12.9N
η_{comb}	5.95	2	0.0049	53.01

The calculated orders of convergence, p_{calc}, were shown to be universally found to be higher than the formal order of the numerical scheme employed in US3D. This was a result of non-uniform, mixed-order behaviour caused by flow discontinuities, as previously mentioned. Following the recommendation of Stern et al. (2001), the order of convergence was estimated to be identical to the order of the solver. Based on this method, the Grid Convergence Indices (GCIs) calculated suggest that the 20 Million cell grid should match an infinitely fine grid to within 2% or less.

Because the calculated orders of convergence for each variable suggested that the behaviour may not be entirely non-monotonic, a finer inlet grid was created. This mesh contained 40 Million cells, and was run with the same solver and flow conditions as the previous three grids. The results from the four experiments are compared in Figure 3.4.

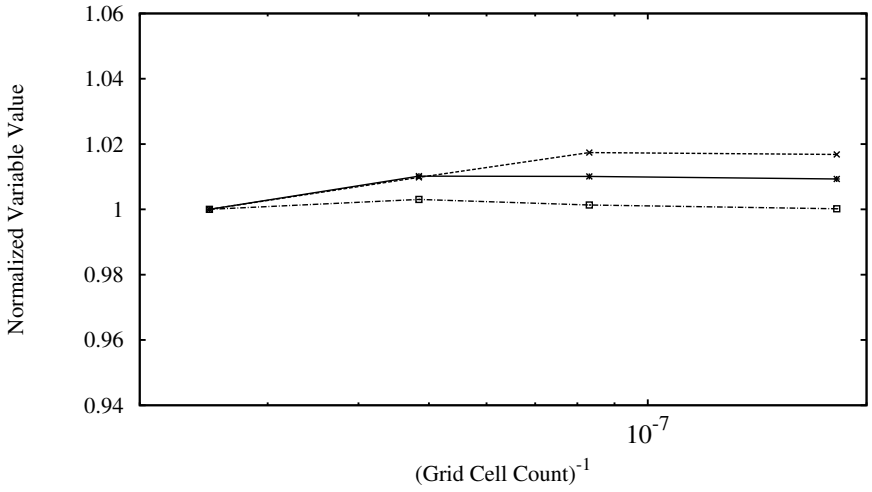

Figure 3.4: Grid Convergence Behaviour

From this figure it can be seen that as the grid is refined further, none of the variables appear to show good convergence according to the accepted definition of GCI (Roache, 1994). This was likely due to the complexity of the flow physics. It has been noted in literature that in the presence of discontinuities, non-monotonic behaviour such as that shown in the Figure 3.4 is to be expected (Oberkampf and Roy, 2010) and no matter what the solver's formal solution order, the solution order becomes linear in the vicinity of shock waves (Banks et al., 2008).

Attempts at using a blended-order convergence method, such as the one proposed by Roy (2003) showed the variables were converging in an oscillatory manner. In order to continue using this method, finer-still meshes would have been required, which was beyond the ability of the computational resources available. However, it should be noted that none of the variables vary by greater than 2.5%, which is in good agreement with the GCI behaviour calculated previously. As such, the variation in each grid convergence index may be deemed acceptably small, and any errors in the numerical solutions of the inlet flow field can be considered acceptably small.

Combustor Grid Considerations

The M12 REST combustor was meshed in the same manner as the inlet, in an attempt to keep grid convergence behaviour on the same order as that of the inlet. There were difficulties in running simulations in which the combustor grid was decoupled from the inlet, and limited computational resources prevented a regular grid convergence study from being carried out on the entire engine. Fortunately, there is an alternate method by which simulations may be validated: comparison with experimental results. This study included experimental investigation of one engine configuration which, along with the experiments of Wise (2014), was compared with the simulations presented to determine their accuracy.

3.3 Experimental Approach

This section details the experimental apparatus and methodology used to obtain validation data. All experiments were undertaken in partnership with Dylan Wise (2014), and as such, this section only provides a brief overview of the facility and the process by which data were collected.

3.3.1 The T4 Stalker Tube

Named in honour of its inventor, the late Professor Ray Stalker, the T4 Stalker Tube (T4) is an impulse flow facility located at the University of Queensland. A free-piston driven reflected shock tunnel, T4 is capable of generating stagnation pressures of up to 90 MPa and stagnation enthalpies as high as 12 MJ/kg. Such energetic conditions come at a cost: steady test times in T4 are typically limited to a few milliseconds, though this is usually enough time for steady conditions to be established over the model being tested.

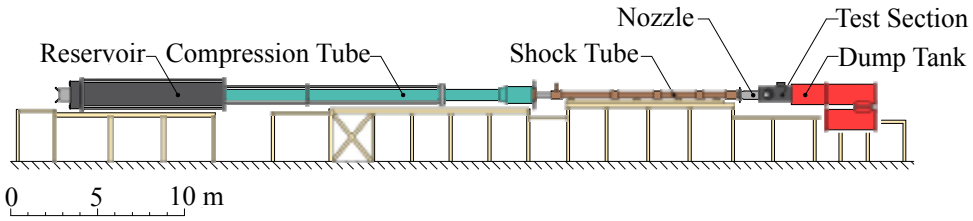

Figure 3.5: Schematic of the T4 Stalker Tube. Reproduced from Doherty (2013)

As shown in Figure 3.5, T4 consists of 6 main sections: a high-pressure annular reservoir, a compression tube, a shock tube, an expansion nozzle, the test section, and the dump tank. The compression tube is 26 m in length and 0.23 m in diameter, and is separated from the reservoir by a 90.4 kg free piston prior to the start of an experiment. The shock tube is 0.076 m in diameter and 10 m long, and is isolated from the compression tube upstream by an unscored rolled steel primary diaphragm, and from the downstream nozzle and test section by a 0.1 mm thick mylar secondary diaphragm.

Prior to the start of an experiment, the entire facility is evacuated down to a pressure below 133 Pa (1 torr). The shock tube is then filled with the test gas, either air or nitrogen, and the compression tube is filled with a driver gas, usually Argon, Helium, or some mixture of the two. Lastly, the reservoir is filled with air to a high pressure. The fill pressures of the test gas, driver gas, and reservoir air are determined by the flow conditions desired. Everything downstream of the secondary diaphragm remains under vacuum until the experiment begins.

When the experiment begins, air from the high pressure reservoir is allowed to leak into an evacuated space behind the free piston. The resulting pressure difference across the piston accelerates it down the compression tube. The driver gas in front of the piston is compressed in a near-adiabatic process due to the speed of the piston stroke, reaching the rupture pressure of the primary diaphragm as the piston reaches the end of the compression tube. As this diaphragm opens, the high pressure driver

gas expands into the shock tube, generating a normal shock that travels down the shock tube. This process is tuned such that after the primary diaphragm opens, the driver gas enters the shock tube at a rate equal to that of the displacement of gas by the now decelerating piston. The major processes that occur downstream are illustrated in Figure 3.6.

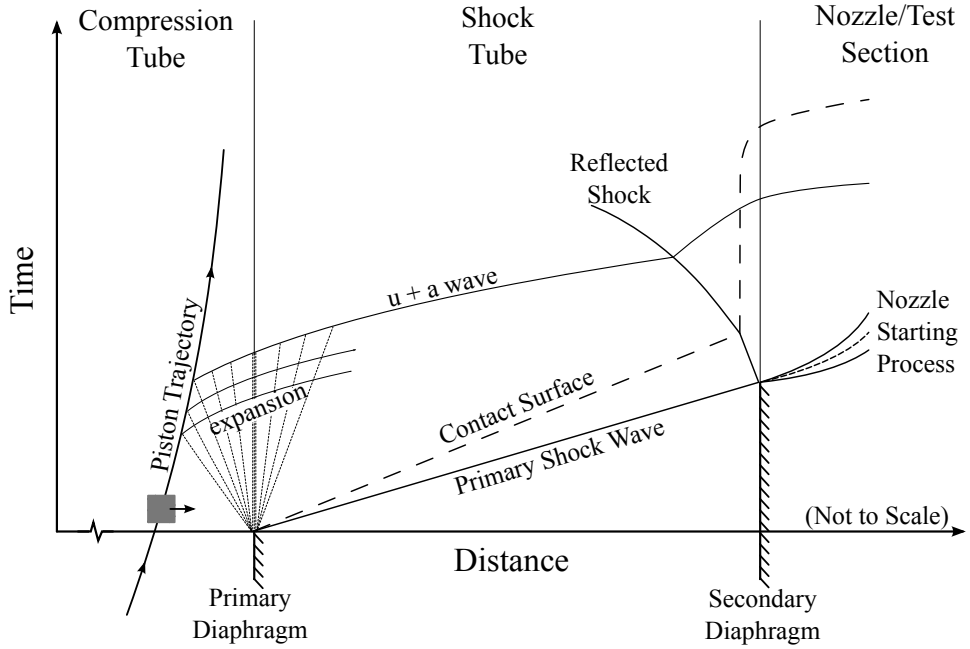

Figure 3.6: An x-t diagram depicting the major flow processes in the T4 Stalker Tube. Reproduced from Doherty (2013)

As the normal shock travels down the shock tube, it processes the quiescent test gas, imparting it velocity in the streamwise direction, as well as increasing its temperature and pressure. The shock passes a series of pressure sensors as it travels down the tube. From these signals, it is possible to extract a shock speed, which assists with the calculation of nozzle exit flow conditions, which will be discussed in Section 3.4. When the shock reaches the end of the shock tube it ruptures the secondary diaphragm and reflects. As the reflected shock moves upstream, it stagnates the test gas it passes through. The stagnated test gas reaches extremely high temperatures and pressures in the nozzle supply region at the end of the shock tube, which then expands outward through the facility nozzle and over the model located in the test section. The pressure in this "nozzle supply" region of stagnated gas at the end of the shock tube is measured with a pair of pressure transducers. After a brief unsteady expansion, the nozzle flow stabilizes into a steady flow condition, which marks the start of useful test time. This test time lasts until either the pressure upstream of the nozzle drops sufficiently to reduce the mass flow through the nozzle below useful levels, or when the test gas

becomes sufficiently contaminated by arriving driver gases. Further discussion, and a more in-depth analysis of the operation of a free-piston driven reflected shock tunnel, can be found in the work of Stalker et al. (2005) or Doherty (2013).

3.3.2 Experimental Model and Apparatus

The experimental model and apparatus consisted of the engine model, the fuel injection system, and multiple pressure transducers and heat transfer gauges to capture information about the experimental flows. The experimental model was built to the specifications and dimensions described in Section 3.1. The engine model was manufactured out of 6061-T6 Aluminium, with the inlet being machined in two halves that were held together by bolts and dowel pins. The combustor and nozzle sections were manufactured out of solid pieces of aluminium using a 3D CNC wire cutter to create the internal surfaces. Figure 3.7 shows the experimental model at the beginning of the experimental campaign.

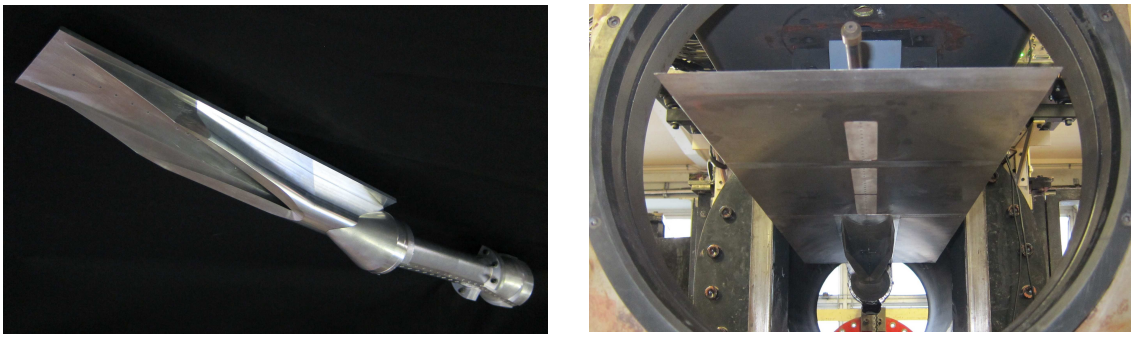

Figure 3.7: The M12 REST engine experimental model prior to experiment (left) and mounted in T4's test section (right). Reproduced from Wise (2014)

In addition to the engine surfaces, the model included a fuelling system that will be discussed later on in this section, and metal shielding to protect both electronic instrumentation and the fuelling system from the incoming flow. The model also was manufactured with multiple portholes throughout the engine, in which the pressure and heat transfer instrumentation were mounted.

The model was instrumented with thin-film heat transfer gauges (hereafter referred to as "HT gauges") and pressure transducers mounted flush with the internal surfaces of the engine. The HT gauges were manufactured in-house at the University of Queensland, and provided the fast response times (on the order of a microsecond) required to take accurate measurements in an impulse flow facility such as T4. More information

on the manufacture and calibration of both the engine model and the HT gauges is contained in Wise (2014).

The pressure transducers chosen were Kulite® XTEL-190M series pressure transducers (hereafter referred to as "Kulite transducers"). These transducers use a piezoresistive silicon sensor to measure absolution pressure and, much like the HT gauges, have a fast response time. The method by which the Kulite transducers were mounted into the model, and calibrated can be found in Kirchhartz (2009).

The inlet contained 11 Kulite transducers and 14 HT gauges, all mounted into the bodyside surface of the inlet, along the engine centerplane. A total of 28 Kulite transducers and 28 HT gauges were mounted into the combustor, with 14 of each mounted on the bodyside and cowlside surfaces of the combustor. Because all this gauges would not fit along the centerline of either combustor wall, each set of instrumentation was mounted on each surface along a streamwise line offset 4mm to either side of the engine centerplane. The nozzle section of the engine model was instrumented with 8 Kulite transducers in a row along the centerline of both its bodyside and cowlside surfaces. The relative streamwise locations of these gauges are shown in Figure 3.8. For their exact streamwise location, refer to the tables in Appendix D.

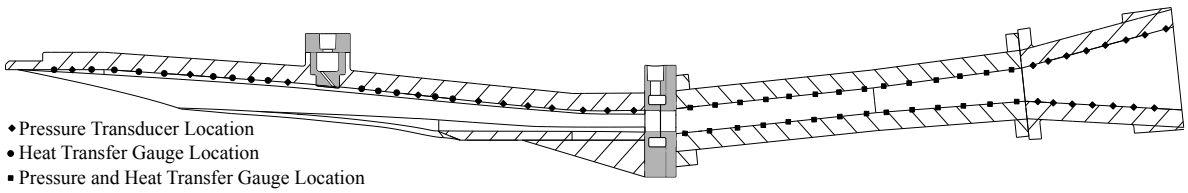

Figure 3.8: Locations of Pressure and Heat Transfer Gauges in the M12 REST Engine

Measurement Uncertainty

As with all experimental measurements, the data compared with each simulation all had uncertainties that needed be quantified. These uncertainties included the uncertainty in measurements made by the Kulite transducers, the HT gauges, and measurements of the T4 facility conditions prior to, and during experiments. The list of uncertainties also included the uncertainty in the nozzle flow exit conditions that were calculated based on the measured facility conditions during an experiment. The exit conditions were used as inflow conditions for many of the simulations. More on how these flow conditions were obtained is described in Section 3.4.

For the sake of brevity, more details of how the various experimental uncertainties were determined are presented in Appendix D. The uncertainties for each measurement type are summarized in Table 3.4. As indicated in the table, the fuel mass flow

Table 3.4: Experimental Measurement Uncertainties

Measurement	Device	Uncertainty
Absolute Pressure	Kulite®	±2.2%
Combustor Heat Flux	HT gauge	±11.9%
Fuel mass flow rate	Kulite®/Calculation	±5.1%

rate was calculated based on a pressure measurement. How and why the mass flow rate was determined in this way is discussed in the next subsection.

3.3.3 Fuel Injection System

The experimental apparatus required that the fuel system be capable of injection at two different streamwise locations in the engine. As demonstrated in Figure 3.9, this meant that two separate fuel plena were incorporated into the engine design. The first was mounted on top of the inlet in order to feed the inlet injectors, while a second annular plenum was installed surrounding the engine isolator to feed the various combustion injection configurations investigated.

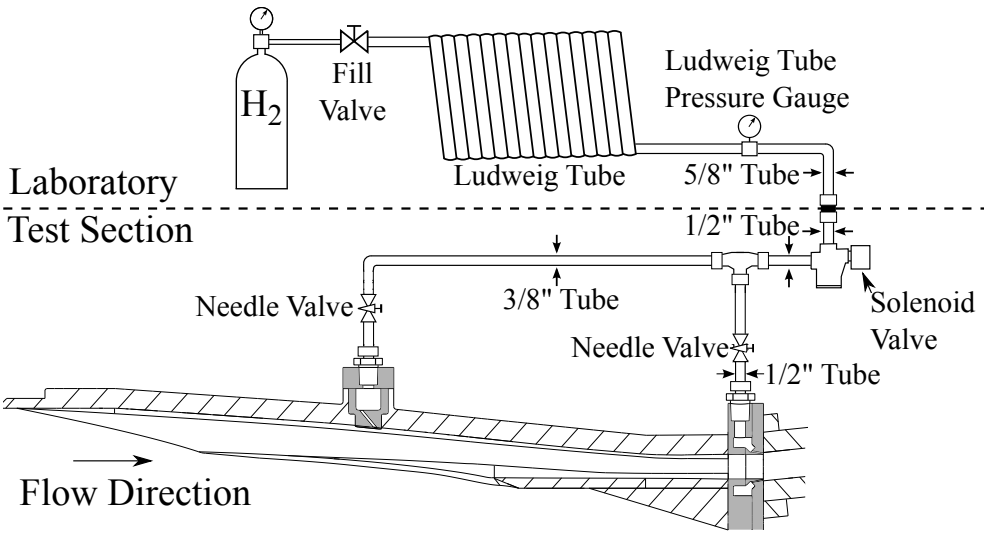

Figure 3.9: Schematic of the experimental fuel delivery system. Adapted from Wise (2014)

Both plena were supplied with high pressure hydrogen gas from a single 14.4 m long Ludweig tube, which was isolated from the engine by a fast-acting Joucomatic™ ASCO solenoid valve. The valve was triggered by the recoil of the tunnel during the experiment. The timing of the valve opening relative to the tunnel recoil was tuned to avoid flow blockage in the model due to excessively early fuelling. More information

on fuel timing is available in Wise (2014). The Ludweig tube length was chosen such that the reflected expansion wave caused by the solenoid valve opening would not reach either fuel plenum during the test time. Fuel supplied from the Ludweig tube could thus stagnate in each plenum at high pressure, and provide steady, choked flow through the injector portholes for the duration of experimental test time.

Between the solenoid valve and each plenum, a manual needle valve was installed. These valves were capable of moderating the flow feeding into the plenum, allowing the mass split between each plenum to be controlled. During the boundary layer injection experiments of Wise, the pressures required in each plenum were similar, and the valves were left open. This left the fuel split in the engine at a nominal value of 31:69 (Wise, 2014). However, during the tailored injection experiments done for experimental validation in this study, the pressure in the combustor plenum needed to be an order of magnitude larger than that of the inlet plenum to achieve the same inlet:combustor fuel ratio. The needle valve controlling flow into the inlet plenum was therefore closed partially, allowing the inlet plenum pressure to drop to the level necessary to achieve the fuel split required to compare tailored injection performance to boundary layer injection.

Throttling the mass flow in the inlet chamber had an unintended consequence: the pressure rise in the inlet occurred at a slower rate relative to the combustor plenum. By the time the combustor plenum had reached a steady operating pressure, the inlet pressure was still rising. However, due to the short test time duration relative to the speed of the pressure rise, the plenum pressure could be assumed as constant across the test time, thus effectively providing an equivalent constant mass flow rate. The needle valve was set such that the average inlet fuel flow rate across the test time had the desired ratio with the combustor mass flow rate. How the fuel mass flow rate was determined from the plenum pressure, and the resulting uncertainty in mass flow rate, is detailed in Appendix D.4.

3.3.4 Test Time Determination

For high enthalpy conditions, the test time of an experiment in T4 is usually limited by driver gas contamination of the test gas (Boyce et al., 2005). At these conditions, it is possible for driver gas to move forward through the shock tube boundary layer and mix with the test gas earlier than expected. The time after secondary diaphragm rupture at which driver gas contamination reaches 10% can be estimated using the correlation

$$t = 62.129 H_0^{-1.7183} \pm 38\%, \tag{3.13}$$

where H_0 is the flow stagnation enthalpy (Boyce et al., 2005). Using this correlation for the average tunnel condition gives an estimated driver gas contamination time of 2.2 ms. This short time made selecting a suitable test time difficult.

Normally, it is assumed that the flow through a model will stabilize within time it takes steady flow from the nozzle exit to move a length equal to three times the model length (Jacobs et al., 1992). A suitable start to the test time (which for T4 is typically 1 ms in length) is then chosen between the time the engine flow becomes steady, and the driver contamination time. For the M12 REST campaign, assuming a nozzle flow establishment time of 1 ms, steady test time would have nominally begun 2.2 ms after the secondary diaphragm ruptured. However, inlet fuelling draws out the inlet starting process, because the test gas must sweep the fuel injected prior to the flow's arrival through the engine before steady flow can establish, pushing flow establishment time out by more than 50%.

This led to the establishment of viable test time at the same time driver gas contamination was expected to have risen above 10%. It should be noted that this nominal time is the point at which the nozzle exit flow is expected to reach 10% contamination. There is therefore a finite time before the flow reaches the engine model.

This lag is illustrated visually in Figure 3.10. All pressure traces inside the engine are normalized against the nominal forebody pressure P_1, which is defined in Appendix F. Test time inside the engine taken at a point later than the nominal nozzle supply test time, where the pressure traces throughout the engine were steady on average. In order to avoid as much contamination as possible, the test time duration was shortened to 0.5 ms. In doing so, the test time was therefore taken inside the engine over a time at which flow was steady and the driver gas contamination would be minimized.

The presence of this inert driver gas was expected to inhibit combustion, as the captured mass fraction of molecular oxygen will drop. Experiments did show increases in combustor pressure when the engine was fuelled, so the contamination inside the engine during test time was not sufficient to suppress combustion. The data collected therefore can be considered only a mildly conservative estimate for the combustion performance of the engine. More details on the determination of the engine test time can be found in Wise (2014).

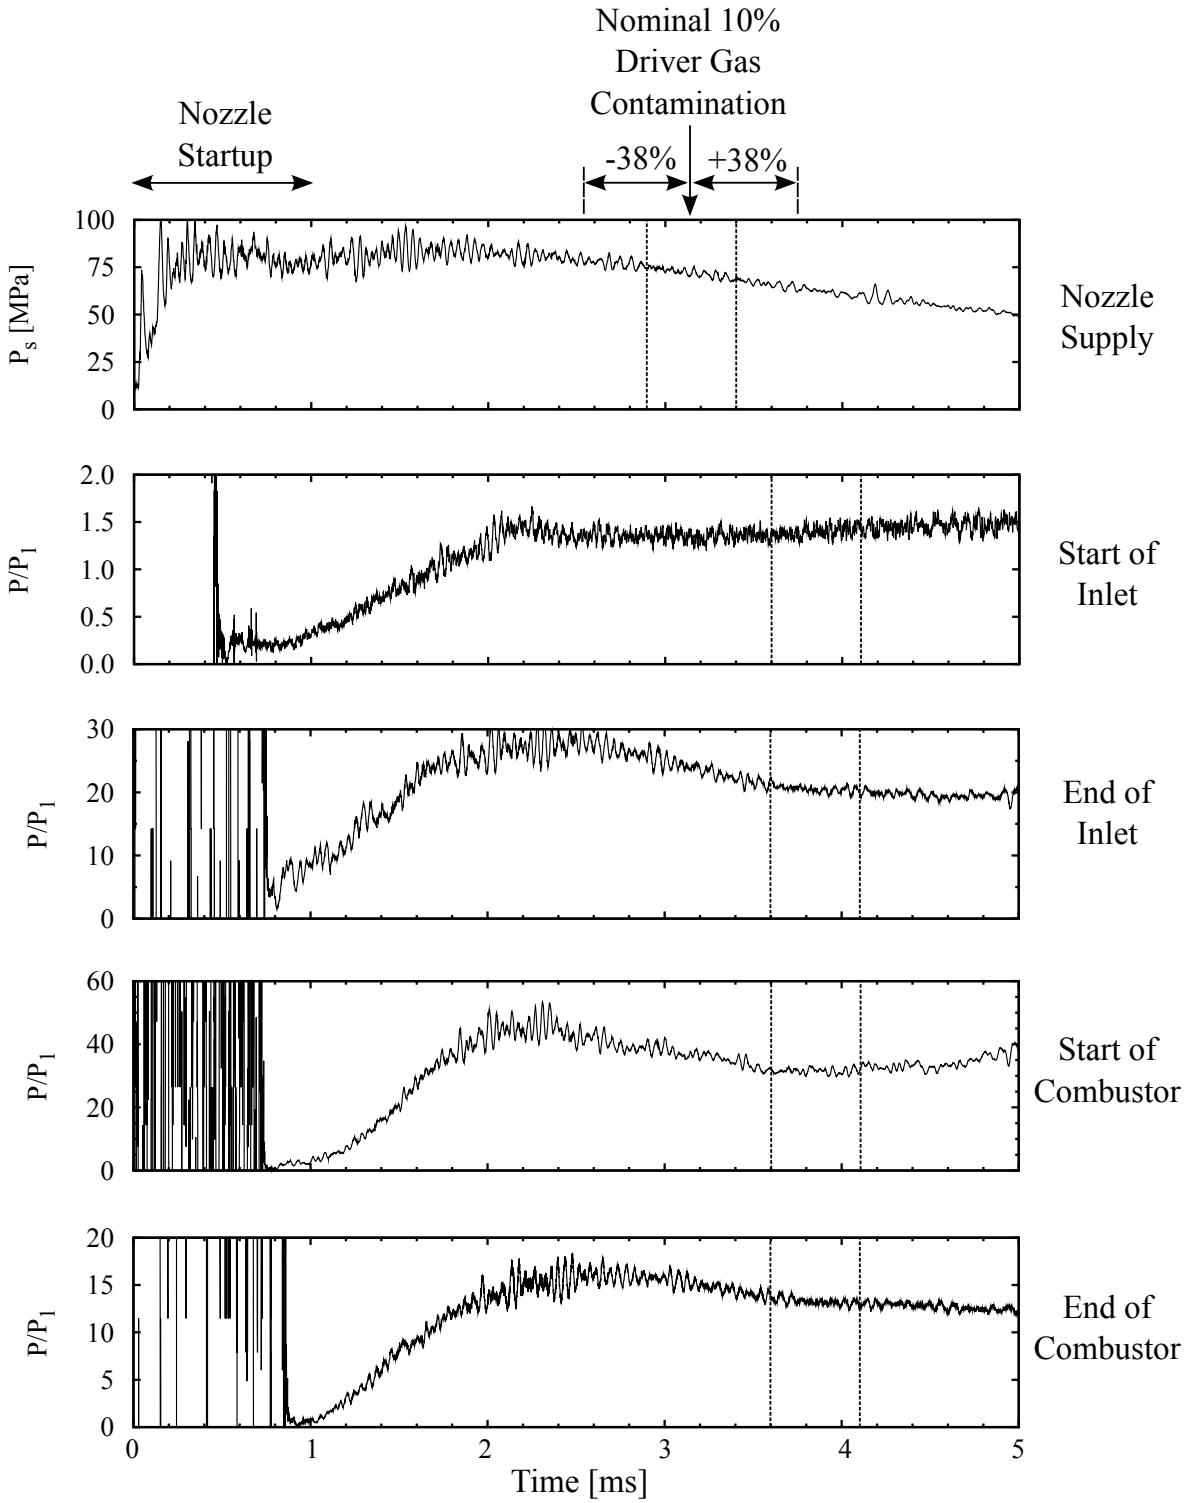

Figure 3.10: Typical transient Nozzle Supply Pressure showing nozzle startup and bounds on 10% driver gas contamination. Test time shown in vertical dotted lines. Also shown are normalized pressure distributions for various pressure taps in the M12 REST engine.

3.3.5 Experimental Repeatability

When testing in an impulse facility, it is important to prove that the variability between experiments is within the tolerance of experimental measurements. Due to the nature of the T4 facility, there is always some shot-to-shot variance in the flow conditions supplied to the test section, despite the best efforts of the operator. In order to check that the validation experiments were repeatable, one condition was repeated at conditions as close to identical as possible. The results of this exercise are presented in Figure 3.11.

Owing to the inexact nature of fuelling the inlet plenum while pressure was rising, the equivalence ratios of these two otherwise identical shots were ϕ=0.75 and ϕ=0.8. Despite this mismatch in fuelling, and some natural variance in the nozzle exit condition between experiments, remarkable agreement was seen in the pressure values across much of the experimental engine length, on both the bodyside and cowlside. This match in behaviour was strongest along the forebody and inlet surfaces, as well as in the last part of the diverging combustor and the nozzle. The absolute pressure measurements showed some variance in the combustor, though the differences are mostly quite small considering the differences in fuel equivalence ratio and natural shot-to-shot variation in tunnel conditions.

A better test of the repeatability of the conditions in the combustor was conducted by comparing the wall heat flux measurements taken in both experiments. While the heat flux gauge located 1184.6 mm from the leading edge on the cowlside combustor wall stopped functioning between the two tests, the combustor heat flux measurements were found to show excellent repeatability. This good shot-to-shot agreement is demonstrated clearly in Figure 3.12: all measured points show agreement to within the measurement error of the gauges.

The repeatability found by comparing these shots agrees well with the findings of Wise (2014), who tested the repeatability of the experiments with an unfuelled engine. This repeatability lends confidence to the subsequent experiments used for comparison to simulation: it is likely that the numerical results are being compared with the true behaviour of the M12 REST inlet operating at on-design conditions. The last exercise remaining in describing the methodology of this study is to describe the flow conditions used, and how they were determined.

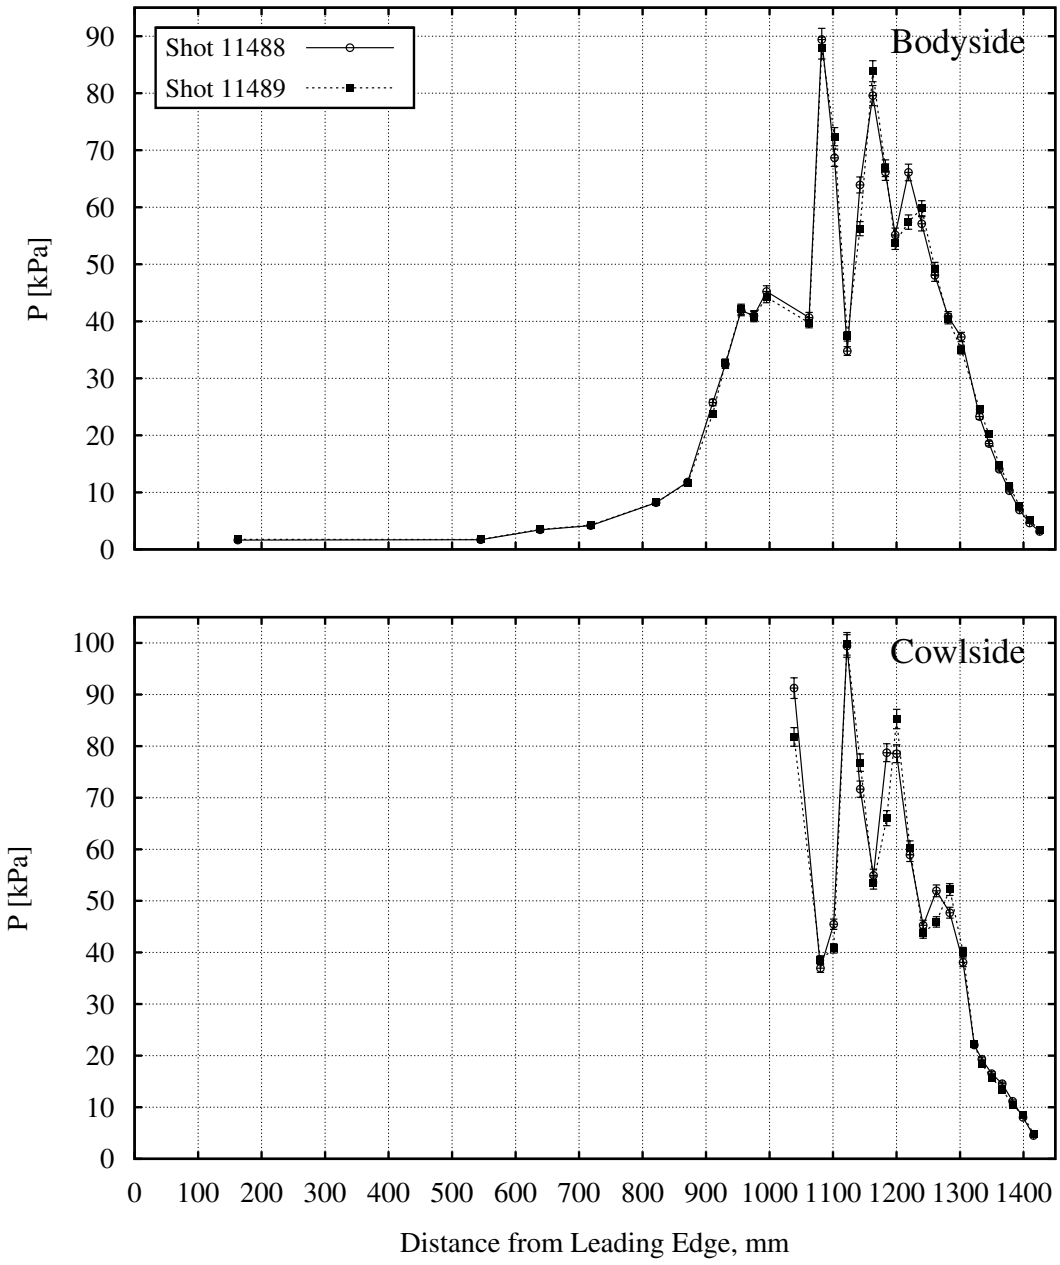

Figure 3.11: Comparison of pressure measurements for two shots at near-identical fuelling conditions

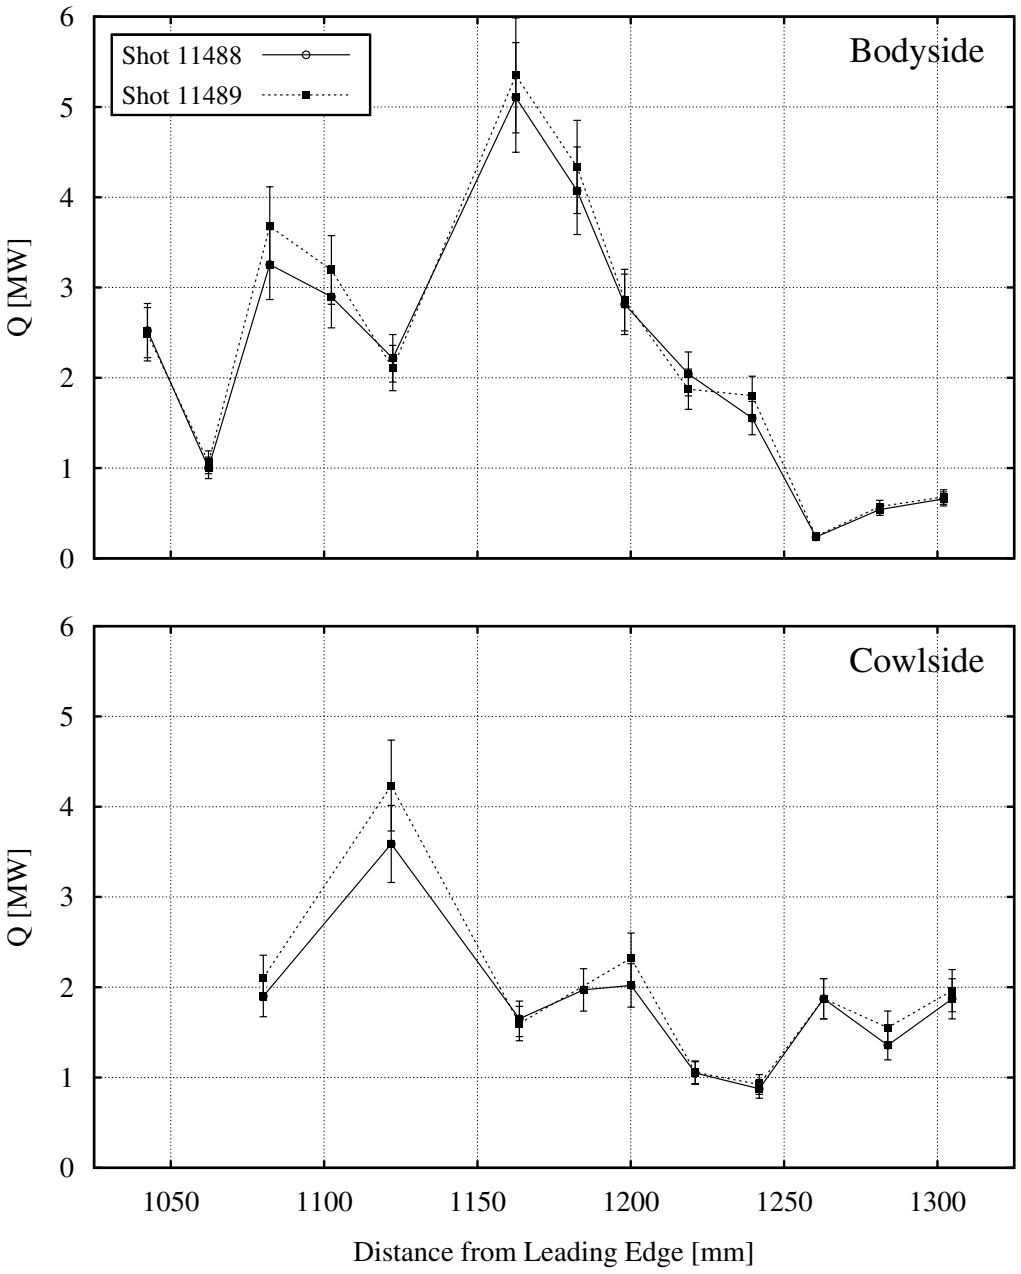

Figure 3.12: Comparison of wall heat flux measurements for two shots at near-identical fuelling conditions

3.4 Experimental and Numerical Flow Conditions

As much as possible, the study undertaken was targeted toward quantifying the flow physics and performance of the M12 REST at on-design conditions. This target flow condition implies that ideally the study would examine the engine's performance when mounted to a vehicle flying at Mach 12 along a constant dynamic pressure ascent trajectory of q $\cong$ 50 kPa. According to the US standard atmosphere model (NOAA, 1976), this would correspond to a geometric altitude of $\sim$36 km, with a flight stagnation enthalpy $H_0 \cong 7.16 MJ/kg$.

In practice the flow conditions have been limited by the capabilities of the T4 Stalker Tube. Because the nozzle supply pressures required to match the true flight dynamic pressure were above facility operating limits, the experiments were conducted as "semi-free-jet" tests, in which the flow was expanded through a Mach 10 nozzle, providing a matching post-forebody-shock flow condition. The flight stagnation enthalpy could therefore be matched, but the target dynamic pressure still could not be matched.

Two separate average test conditions were considered. Because simulations were undertaken prior to the start of the joint experimental campaign, the first test condition matches the calculated average nozzle outflow of the flat plate experiments of Wise (2014). The second condition is average tunnel condition used during the joint campaign. These two conditions are summarized in Table 3.5. Since the shock tube was typically filled 15 minutes or more before the experiment occurred, it has been assumed that the test gas had sufficient time to equilibrate to the ambient temperature, taken as 300 K $\pm$ 2%.

Table 3.5: T4 fill conditions for each campaign

Campaign	Fill Pressures			Gas Composition		Diaphragm
	Reservoir	Comp. Tube	Shock Tube	Driver	Test	Thickness
	MPa	kPa	kPa	Ar/He %	Gas	mm
Flat Plate	12.5	156	150	35/65	Air	6
Joint	12.0	156	160	35/65	Air	6

Using these fill conditions together with the measured mean shock speed and nozzle supply pressure, the two average nozzle exit conditions were calculated using UQ's in-house code NENZFr (Doherty et al., 2012a). This code expanded the measured nozzle supply conditions through the T4 Stalker Tube's Mach 10 nozzle geometry in thermochemical non-equilibrium utilizing a 5-species air finite-rate chemistry model. The

resulting calculated freestream conditions at the nozzle exit are shown in Table 3.6. Doherty and Wise (2013) showed that these simulations generally agree with pitot survey data of T4's Mach 10 nozzle to within experimental uncertainty. The method used to calculate the uncertainty in the values calculated by NENZFr is described in Appendix D. For the purposes of the M12 REST simulations run, nozzle outflows were assumed to be thermochemically frozen, and the flow vibrational temperature was set to be equal to the flow static temperature.

Table 3.6: NENZFr-calculated freestream conditions and uncertainties

Flow Variable	Flat Plate Condition	Joint Condition	Variable Uncertainty
P, Pa	855.3	1152.1	$\pm$7.62%
ρ, kg/m^3	0.0089	0.0101	$\pm$8.98%
U, m/s	3516	3655.5	$\pm$3.46%
T, K	333.2	394.67	$\pm$10.01%
M, K	9.61	9.15	$\pm$1.73%
Re_u, $1/m$	1.570 x 10^6	1.613 x 10^6	$\pm$11.87%
q, kPa	55.01	67.33	$\pm$5.74%
H_0, MJ/kg	6.54	7.12	$\pm$7.52%
Y_{N_2}	0.728491	0.72837	$\pm$0.03%
Y_{O_2}	0.183630	0.18258	$\pm$1.24%
Y_{NO}	0.082676	0.082921	$\pm$0.49%
Y_O	0.005203	0.006129	$\pm$36.44%

As is readily visible in Table 3.6, the over-driven piston caused the flat plate nozzle flow to have a lower than normal stagnation enthalpy and dynamic pressure. By contrast, the joint campaign flow had approximately the desired stagnation enthalpy and an improved dynamic pressure. Despite this, neither condition was able to match the desired flight dynamic pressure, though both fall within the recommended scramjet dynamic pressure ascent corridor (25-100 kPa) of Heiser and Pratt (1994).

Table 3.6 also shows that NENZFr predicts a significant conversion of the test gas' molecular Nitrogen and Oxygen into NO and O radicals. This makes physical sense: stagnation pressures of 80-85 MPa, and stagnation temperatures in excess of 5000 K are generated in the nozzle supply region of T4, and molecular nitrogen and oxygen are likely to dissociate and react at these conditions. As such, the air captured by the engine will contain less molecular oxygen than true flight conditions, compounding the under-prediction of combustion due to flow contamination that was discussed in Section 3.3.2.

All simulations conducted before the start of the joint campaign used the flat plate nozzle exit condition as their inflow. This includes all of the simulations presented in Chapter 4. Any simulations run concurrently to, and after, the joint campaign were run with either the average conditions presented above, or with inflow conditions calculated from NENZFr to match individual shots as closely as possible to allow for a good comparison between the simulation and its matching experiment. A record of all validation experiments conducted can be found in Appendix C, while the shot-matched inflow conditions are tabulated in Appendix E and will be referred to each time a new engine configuration is presented.

3.4.1 Equivalent Flight Conditions

Because both test conditions had Mach numbers higher than would be expected from a Mach 12 flow processed by a 6° shock, both experiments and simulations were conducted with the engine inclined at a 1.6° angle of attack relative to the incoming flow. To better understand what flight conditions the flows were equivalent to, the nozzle exit flow conditions were input into the classical $\theta - \beta - M$ oblique shock relations (Anderson, 2003) assuming a 1.6° turning angle to determine the nominal conditions behind the experimental forebody shock. This condition was then used to calculate backward through the same relations assuming a 6° forebody angle, arriving at the nominal freestream conditions they would correspond to.

These freestream conditions were then matched to the 1976 US Standard Atmosphere (NOAA, 1976) by matching static pressure of the flow to atmospheric pressure. This gave the nominal altitude at which flight would be taking place, and the properties of the quiescent air. The velocity was then calculated to ensure that the stagnation enthalpies matched. This was done by taking the formula for stagnation enthalpy

$$H_0 = \frac{1}{2}U^2 + c_p T. \tag{3.14}$$

Using atmospheric values for T and c_p, velocity could be back-calculated to match the flight and tunnel stagnation enthalpies. Because temperatures between the assumed altitude and the tunnel freestream conditions are close enough to assume constant c_p, and the static enthalpy of the fluid is relatively small compared to the kinetic component, the difference between the velocity in flight computed from H_0 and that which would be obtained from the oblique shock relations using the experimental velocity and static temperature is small, typically less than 1% at this flight condition. Substituting the velocity and atmospheric temperature into the ideal gas formulation for

Mach number,

$$M = \frac{U}{\sqrt{\gamma RT}} \tag{3.15}$$

the equivalent flight Mach number was thus calculated. The results of these calculations are shown in Table 3.7.

Table 3.7: Equivalent flight conditions for each test condition

Variable	Flat Plate	Joint
P, Pa	274	392
ρ, kg/m^3	0.0038	0.0056
U, m/s	3546	3708
T, K	251.3	244.05
M, K	11.16	11.84
Re_u, $1/m$	0.84 x 10^6	1.32 x 10^6
q, kPa	23.9	38.5
Altitude, km	40.1	37.5

It is evident that each equivalent flight condition falls short of the desired dynamic pressure of 50 kPa. They do, however, also fall within the recommended constant-dynamic-pressure flight envelope, as demonstrated in Figure 3.13. Both conditions match to within 10% the on-design Mach number, and are useful examples of the engine operating at near on-design conditions.

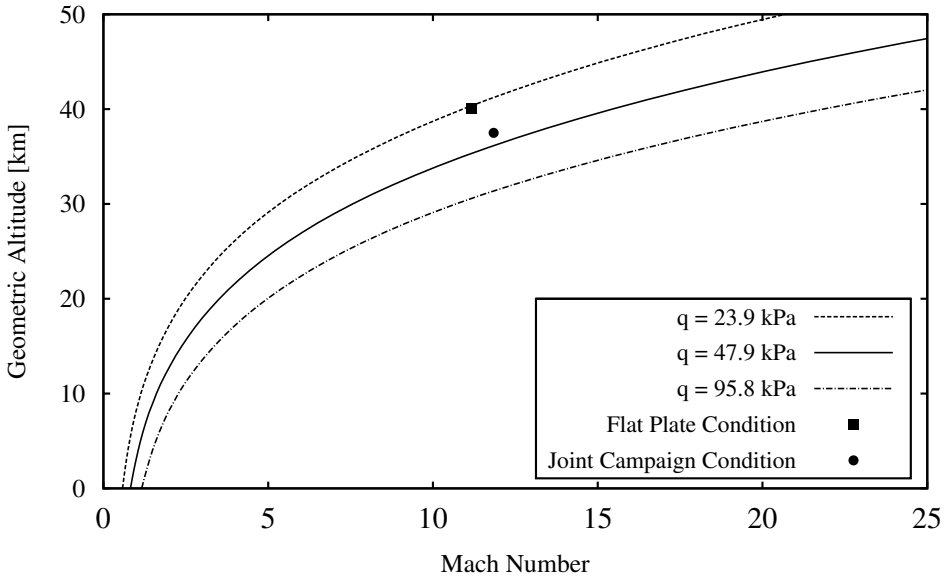

Figure 3.13: Tunnel-equivalent flight conditions relative to the recommended constant dynamic pressure trajectory corridor of Heiser and Pratt (1994)

3.5 Summary

Both the numerical and experimental methodology have been presented, and the areas in which they complement each other has been identified. Using the methodologies presented above, a detailed study of the M12 REST engine's flow physics and performance at on-design conditions was carried out. The analysis begins with a detailed study of the M12 REST inlet, which is the subject of the next chapter.

4

Shape-transitioning Inlet Flow Physics and Fuel Injection Effects

"I am, and ever will be, a white-socks, pocket-protector, nerdy engineer, born under the second law of thermodynamics, steeped in steam tables, in love with free-body diagrams, transformed by Laplace and propelled by compressible flow."

– Neil Armstrong, *Address to the National Press Club* (2000)

The investigation into hypervelocity scramjet flow physics begins with an examination of the of the M12 REST inlet flow field. In this chapter, an examination the basic (unfuelled) flow field will demonstrate the principal flow structures and performance of the inlet as it captures and compresses incoming air. How the inlet flow field changes when fuel is injected along the inlet bodyside surface, and what effects that has on inlet performance are then examined.

The CFD simulations presented are of the M12 REST inlet only, terminating at the nominal throat 976.2 mm downstream of the leading edge. This location is upstream of the completed shape transition, but is the streamwise location where inlet cross-sectional area is minimized. In all cases, the simulations were conducted with inflow conditions equal to the flat plate condition presented in Section 3.4 of the previous chapter. In the simulations where fuel was injected from the inlet bodyside compression surface, the mass flow rate of hydrogen injected was set to give an equivalence ratio (ϕ) of 0.33, based on the mass capture of the unfuelled inlet and the mass fractions of available atomic and molecular oxygen in the inflow condition. In all cases, the simulations' RMS residuals converged by at least 7 orders of magnitude, and the mass flow balance across their inflow and outflow boundary conditions was less than 1×10^{-6} kg/s. In addition, all momentum balances were less than 10^{-3} N.

4.1 The Unfuelled Inlet

There are several flow structures generated by the M12 REST inlet that are of some significance. The first and most obvious feature of the inlet flow field is the thick bodyside boundary layer, which originates from the vehicle forebody, whose Mach number contours are shown in Figure 4.1. The boundary layer originates at the leading edge of the vehicle forebody (left hand side of the figure), 500mm upstream of the inlet leading edge. By the time the flow reaches the start of the inlet it has developed into a fully turbulent hypersonic boundary layer, approximately 7 mm thick. The boundary layer contains a high-temperature zone in its log layer caused by the strong viscous heating characteristic of compressible turbulent boundary layers (Van Driest, 1951).

Figure 4.1: Mach number contours showing boundary layer growth along the forebody upstream of the inlet

Downstream of this point, the inlet begins to capture the oncoming flow and the bodyside boundary layer's development changes, as shown in the cross- stream contour slices shown in Figure 4.2. As a note on convention, this figure and all inlet figures that follow assume a flow direction from left to right within the figure. The slices shown are located 87, 135, 177, 216, 253, 287, 318, 346, 372, 395, 417, 437, 457, and 476 mm downstream of the inlet leading edge.

Most clearly shown in Figure 4.2a, the boundary layer takes on a rounded "bubble" shape as it develops within the confines of the inlet, much like the Mach 8 REST flow fields examined by Gollan and Ferlemann (2009). This is a result of the compressive force exerted on the flow by the inlet sidewalls as they contract inward toward the inlet symmetry plane. The lower speed and momentum of the boundary layer makes it particularly sensitive to the change in inlet geometry, forcing air near the top of the sidewalls upward and inward along the bodyside wall. This ultimately forces the bodyside boundary layer to bulge outward as extra mass enters this region of the flow.

The density contours shown in Figure 4.2b demonstrate another key characteristic of the inlet flow field: by the time it exits the inlet, the majority of the captured air

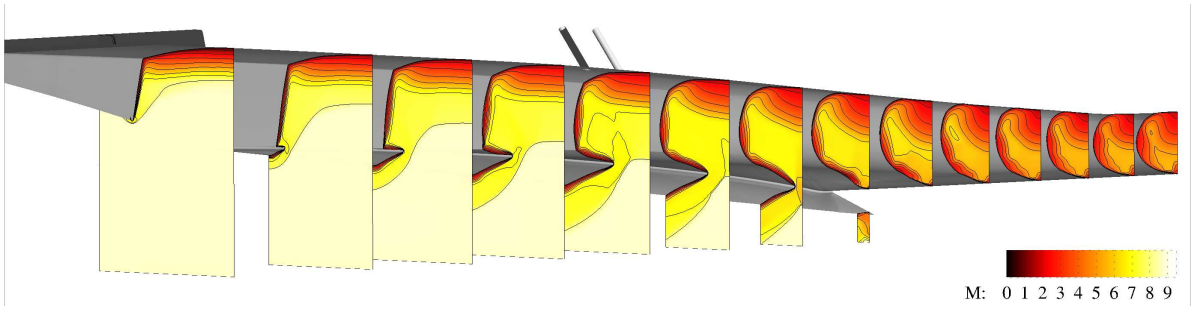

(a) Mach Number

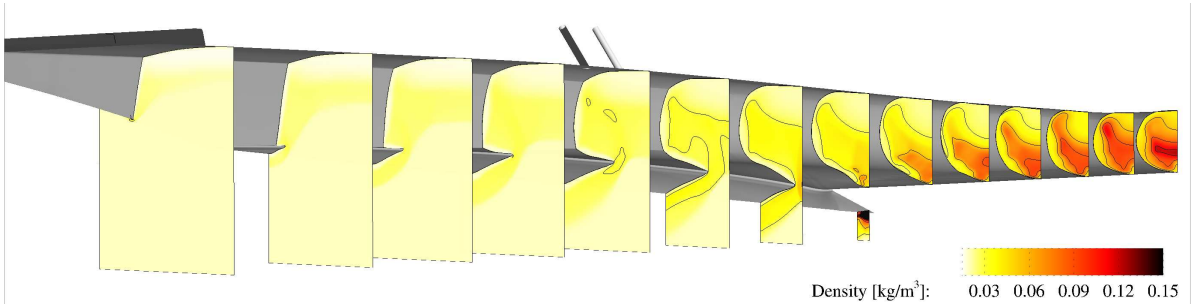

(b) Density

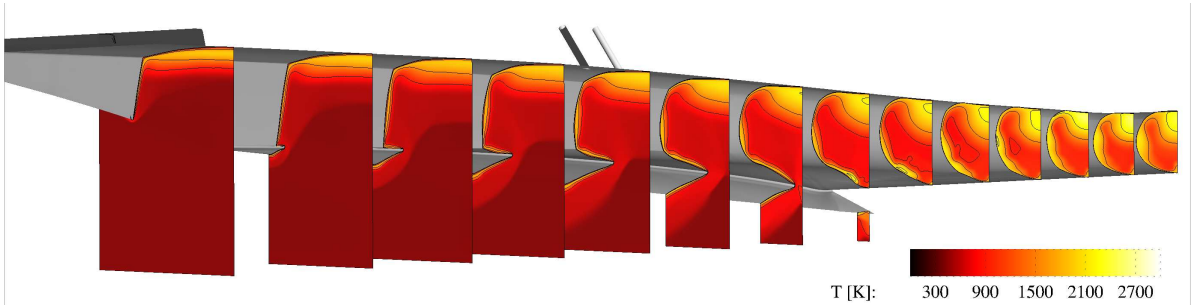

(c) Temperature

Figure 4.2: Flow characteristics of the unfuelled inlet flow field

mass flow ends up in a high Mach number, higher density, and relatively low temperature zone centered on the cowlside surface of the inlet. This high-density zone extends upward and away from the inlet symmetry plane along the edge of the rounded bodyside boundary layer. As was previously mentioned in Section 2.2.1, body-cowl flow asymmetry is common in mixed-compression inlets such as this one. This feature is of particular interest because it implies that fuel injected along the inlet's bodyside compression surface might not interact with the majority of available oxygen in the captured air; it could instead stay trapped inside the bodyside low-density bubble.

Temperature contours of the inlet flow field show the aforementioned high-temperature region developing in the bodyside boundary layer, which could serve to ignite any injected fuel that remains inside the boundary layer. Another important flow structure

can be inferred from Figure 4.2c: a symmetric pair of large, coherent, high-temperature regions appear downstream of the cowl closure notch, and propagate outward and upward along the cowlside and sidewall surfaces of the inlet. With temperatures that are comparable to the bodyside boundary layer, and a lower local Mach number and density within these regions, the most logical conclusion is that a pair of swept separations had formed, much like those that formed in the Mach 8 REST inlet of Turner (2010). This will be explored further in Section 4.1.3.

4.1.1 Pressure Contours and Experimental Comparison

The density and temperature contours in Figure 4.2 hint at regions where there are likely to be high or low pressures. These pressures are examined by plotting the pressure field along the engine centerplane, as well as the pressure distribution along the bodyside centerplane surface pressures, as shown in Figure 4.3.

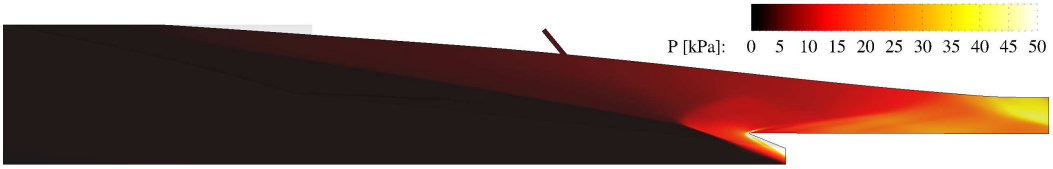

(a) Symmetry Plane Structure

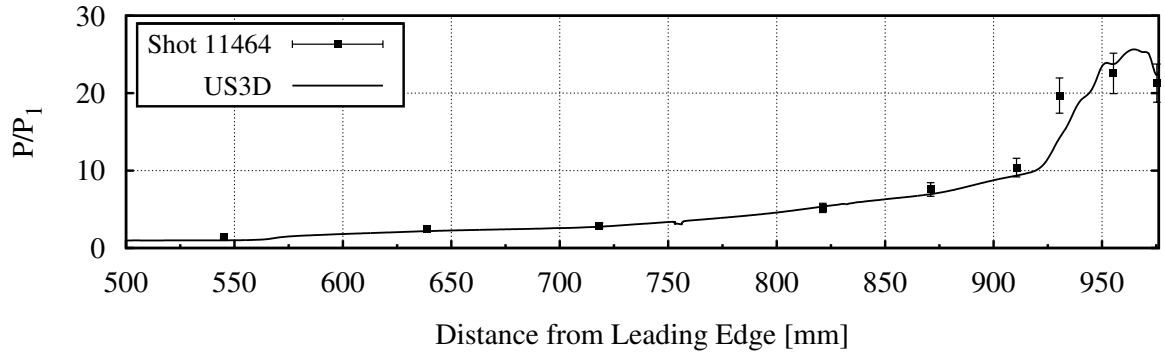

(b) Body-Side Centreline Pressure Distribution

Figure 4.3: Unfuelled inlet centerplane pressure contours and bodyside pressure distribution

The pressure field shown in Figure 4.3a gives a good summary of how the inlet compresses the flow. The cowl closure shock is plainly visible from the pressure rise in the flow downstream of the cowl lip, as is the strong flow compression that occurs behind the reflection of the cowl closure shock off of the bodyside surface. The bodyside surface pressure distribution is compared with experimental normalized pressure levels, as shown in Figure 4.3b, and also clearly shows the reflection of the cowl closure

shock at the location of the sudden bodyside pressure rise. The experimental data is taken from Shot 11464 of Wise (2014). There is good agreement between the simulated normalized pressure contours and the experimental data. The relative strength and location of the bodyside pressure rises are well-captured by the simulation, with only one pressure measurement not lying directly on the simulated line.

The pressure contours also show an interesting flow feature that has escaped mention until now: there is a sudden compression in the center of the inlet flow field 30 mm upstream of the cowl closure. This compression region begins well away from any physical surfaces, and as such, is likely due to an as-yet unexamined shock structure.

4.1.2 Inlet Shock Structure

The basic shock structure of the M12 REST inlet therefore needs examining, and is presented as Figure 4.4. These contours are of $\partial_1\rho$, which is defined as the dot product of the local density gradient and local velocity unit vector, which has been shown to provide an excellent representation of shock and expansion waves (Ma et al., 1996). By only plotting values of $\partial_1\rho$ greater than 0, by definition only compression gradients (and therefore, shocks) are rendered visible.

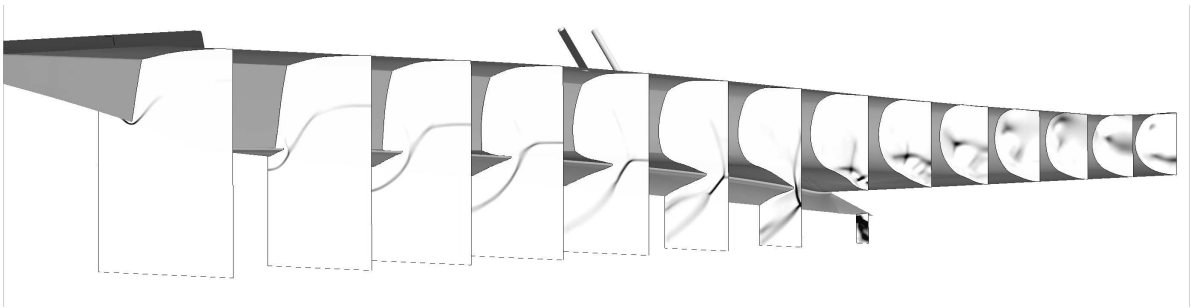

(a) Cross-plane Structure

(b) Symmetry Plane Structure

Figure 4.4: $\partial_1\rho$ contours showing the shock structures of the unfuelled inlet.

Immediately visible in Figure 4.4 is the quasi-conical shock that was generated by the bodyside and sidewall leading edges. This shock first begins propagating downward away from the bodyside compression surface. At a point 129.3 mm downstream

of the inlet leading edge (midway between the first and second slices in Figure 4.4a, the cowl begins to close as the cowlside leading edge moves inward toward the inlet centerplane. The sudden growth of the cowlside surface inward generates a new shock, which travels inward toward the centerplane along the quasi-conical shock. This forms the triple point weakly visible in the cross-plane slices of Figure 4.4a, beginning at the second slice from the left side of the figure. The lateral shock continues inward until a point 30 mm upstream of the cowl closure, where it reflects off the symmetric shock generated by the other half the inlet. This marks the beginning of the high compression region upstream of the inlet closure visible in Figure 4.3.

The reflected lateral shock then travels outward, encountering the cowlside leading edge of the inlet 15 mm downstream of the reflection location, in the last contour slice upstream of the inlet closure. At this point, the transmitted portion of the shock continues outward along the inlet's cowlside walls, both internally and externally. Meanwhile, a small reflected shock (which originates at the portion of the cowl leading edge that is normal to the shock) travels back inward toward the engine centerplane and the cowl closure. This small reflected shock meets its symmetric partner just upstream of the cowl closure, amplifying the density and temperature of the stagnation zone that forms at the cowl notch. Not far downstream of the cowl closure, the laterally-travelling shock forms a lambda foot as it moves outward toward the sidewall.

The closure of the cowl generates another shock, which stands off slightly from the surface of the notch. The cowl closure shock rapidly propagates outward, away from both the inlet centerplane and the cowlside wall. As is visible in both parts of Figure 4.4, several weak compression waves also form just downstream of the cowl closure as the flow that was locally travelling downward is redirected by the cowlside wall.

The cowl closure shock and the transmitted lateral shock interact in a complex pattern. The outward-travelling lateral shock begins to turn up the sidewall at approximately the same time as the cowl closure shock encounters the thick bodyside boundary layer bubble. Because of a local speed of sound gradient corresponding to the change in temperature across the bodyside bubble, the transmitted cowl shock travels at different rates across the top half of the inlet, leading to the diffuse pressure gradients observed in Figure 4.4a and the non-linear shape of the centerplane cowl shock in the top half of the inlet (Figure 4.4b). The cowl closure shock then reflects off the bodyside wall and returns to the cowlside of the inlet, generating the observed high-density flow region adjacent to the cowlside wall.

4.1.3 Inlet Vortex Structures

The shock structure alone does not give any further indication of the presence of a swept separation along the sidewall of the inlet as the flow approaches the throat, nor does it betray the presence of any other vortical structures in the flow. The presence of streamwise vortices, including the swept separation, can be confirmed by examining the features presented in Figure 4.5.

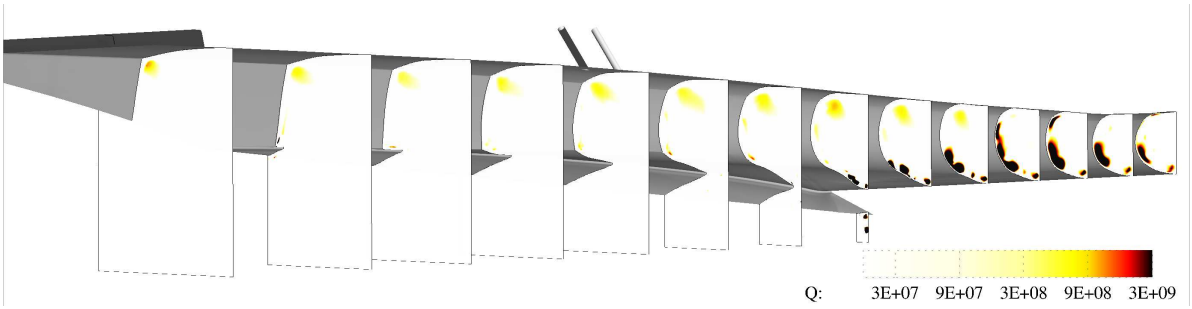

(a) Cross-plane Q criterion magnitude

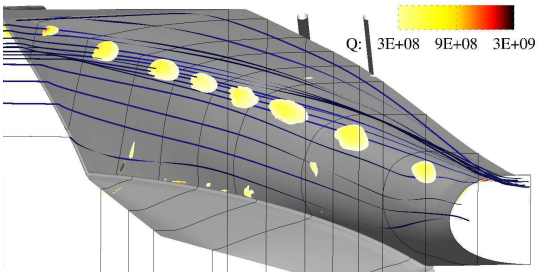

(b) Streamlines through top corner vortex

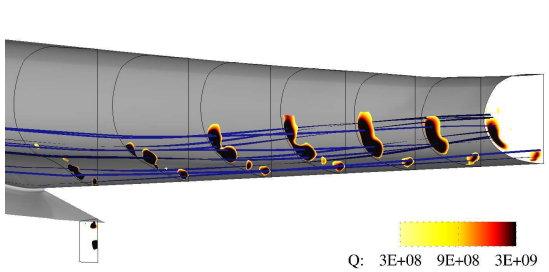

(c) Streamlines through swept separation vortices

Figure 4.5: Unfuelled inlet vortical structures

By employing the Q-criterion (Hunt et al., 1988), it is possible to see what major vortical structures are propagating downstream through the inlet, as illustrated in Figure 4.5a. These regions have a high rate of fluid rotation relative to fluid shear, and indicate the presence of a pair of vortices created by the sidewall leading edge (one internal, the other external) that travel along the top "corners" of the inlet flow, at the junction between the bodyside and sidewalls of the inlet. Unsurprisingly, the Q contours also identify a vortical region at the throat that coincides with the high-temperature region in Figure 4.2c previously identified as a swept separation. More interestingly, the Q-criterion also implies that the swept separation is formed by the coalescence of two individual zones of high vorticity. The strength of this rotational zone diminishes as it approaches the throat. This is due to the compression and weakening of the coherent vortex structure by the passage of the bodyside reflected shock as it approaches the

cowlside wall.

The findings from the examination of Q-criterion are confirmed by streamtracing through the high-Q regions of the flow field. The streamlines passing near the top corner streamwise vortex are shown in Figure 4.5b. The rotation of these streamlines prove that the top vortex is, along with the contraction of the inlet surfaces, responsible for the convection of flow inward along the bodyside wall. As the flow encounters the sidewall leading edge, the leading edge shape attempts to force the flow downward. Given the considerable momentum of the freestream, the flow instead splits around the leading edge. Since the inlet is at an angle of attack to the flow, the near-wall flow travels "upward" toward the bodyside surface and is then turned inward, inducing a clockwise rotation in the flow near the corner.

The streamlines shown in Figure 4.5c confirm that the swept separation moving up the sidewall at the throat forms from two individual vortices that coalesce as the flow approaches the throat. These individual vortices are formed by the passage of the laterally-travelling shock wave through the cowlside boundary layer, which induces a strong baroclinic torque. The inner-most of the two vortical structures forms behind the lambda foot of the laterally transmitted shock as it travels outward along the wall. The second region of vorticity appears after the same shock travels further around the curvature of the inlet surface from the cowlside wall to the sidewall. These two regions of strong rotation eventually merge into a single swept separation, with the streamlines passing through each region ending up clustered together along the sidewall at the throat.

Finally, there is a small counter-rotating vortex pair formed along the centerline of the cowlside wall. This vortex forms when the streams of fluid being forced inward by the sidewall contraction converge downstream of the cowl notch. Being bounded by the cowlside wall, the meeting streams are forced to divert upward, and the generated vorticity rolls the fluid into the observed vortex pair.

4.1.4 Unfuelled Inlet Performance

Having examined the vortical structures of the inlet flow field, the flow physics of the inlet can be said to be well-characterized. This is not the entire story, however. There are parameters frequently used to characterize the performance of a scramjet inlet that should not be ignored. Such parameters include the adiabatic kinetic energy efficiency and mass capture efficiency of the inlet, and the inlet static compression and contraction ratios. As reported in Chapter 3, the total contraction ratio of the inlet is

6.61. Using the mass-flux-weighted average static pressure at the inlet throat and the static pressure of the inflow, the compression ratio of the inlet is 37.2.

Knowing the total contraction ratio of the inlet is based on the projected capture area of the inlet (that is, the projection of the leading edges along the axis of the forebody), it is possible to define the simulated mass capture efficiency of the inlet as

$$\eta_{capture} = \frac{\dot{m}_{simulated}}{\rho U A_{projected}} \tag{4.1}$$

where the density ρ and velocity U are flow properties calculated behind the oblique forebody shock. For the flow condition simulated, the engine was found to have a mass capture of 47.2 g/s, which gives a mass capture efficiency of 84%. This lower mass capture is partly due to the growth of the bodyside boundary layer, which lowers the cross-sectional area through which high speed flow can pass. It is also due to the semi-free-jet tunnel inflow condition, which gives a different post-shock flow profile than an engine at the on-design flight condition. How much of a difference in mass capture exists between tunnel and flight inflow conditions will be examined in Chapter 6.

The adiabatic kinetic energy efficiency is frequently reported when real gas effects are important (Van Wie, 2000), as it provides a measure of the kinetic energy of the flow at the throat relative to the expected kinetic energy at the throat of an isentropic inlet. This quantity is defined mathematically as

$$\eta_{KE,ad} = \frac{H_{0,t} - h(P_\infty, s_t)}{H_{0,\infty} - h_\infty} \tag{4.2}$$

where the subscript t denotes the average conditions at the throat, and the subscript ∞ denotes average freestream conditions. The variable $h(P_\infty, s_t)$ is the static enthalpy obtained by expanding air isentropically from the mean throat conditions to the freestream pressure. Taking the properties of the simulation inflow, and the mass-flux-weighted average quantities at the throat, the M12 REST inlet is calculated to have an adiabatic kinetic energy efficiency of $\eta_{KE,ad} = 96.84\%$. This high efficiency is typical of stream-traced inlets, where the flow compression done by the inlet is as close as possible to the isentropic compression flow field its shape was derived from. This value also falls within the range of values estimated by Smart (2012), and agrees well with Smart's correlation between the compression ratio and efficiency when the compression ratio in this simulation is adjusted to account for the fact that the simulation inflow condition is a semi-freejet condition, and thus corresponds to the pressure behind a vehicle leading-edge shock wave.

It should be noted that because these quantities are calculated from mass-flux-averaged quantities, it has effectively been assumed that the flow is mixed into a uniform, quasi-one-dimensional state. While the preceding sections have shown this assumption to be untrue, the values above can be thought of as a lower bound for the true performance of the inlet, since the mixing assumption inherent to the averaged quantities produces an entropy rise that is higher than would otherwise be expected.

It is worthwhile noting at this point that many of the features discussed so far are also found in the Mach 10 sugar scoop inlets. Despite lacking a long forebody upstream of their inlet compression surfaces, these stream-traced inlet designs also form thick, hot bodyside boundary layers, as well as a similar high-temperature structure sweeping outward from their centerlines (Drayna et al., 2006). Based on these similar large-scale flow structures, it is likely that whatever flow changes are observed when inlet injection is employed in the M12 REST inlet can be expected to occur in sugar scoop inlets as well, with performance changes that would likely be of the same order of magnitude.

4.2 The Fuelled Inlet

With a full understanding of the unfuelled inlet, it is now possible to examine the impact of injecting fuel at a low equivalence ratio along the bodyside compression surface of the inlet. As stated previously, simulations for fuel injection at an equivalence ratio of 0.33 have been carried out with both air and Nitrogen inflow conditions. Due to the similarity in their flow structures, this section primarily focuses on the more interesting case of fuel injection into air, where the hydrogen was able to react with free oxygen in the airstream. The Nitrogen inflow case is used primarily in comparison with the reacting case in order to gain a better understanding of specific phenomena observed.

4.2.1 Inlet Shock Structures

The most noticeable changes to the inlet flow caused by fuel injection are in the shock structures of the inlet, shown in Figure 4.6. The under-expanded fuel jets generate a pair of bow shocks which propagate outward and away from the bodyside wall, merging into the "double-humped" shape visible in the cross-stream slice just downstream of the injection point. As the combined bow shocks travel downstream, the curvature smooths into a curved shock wave that begins reflecting off the cowlside wall just downstream of the point where the reflected lateral shock passes over the cowl leading edge.

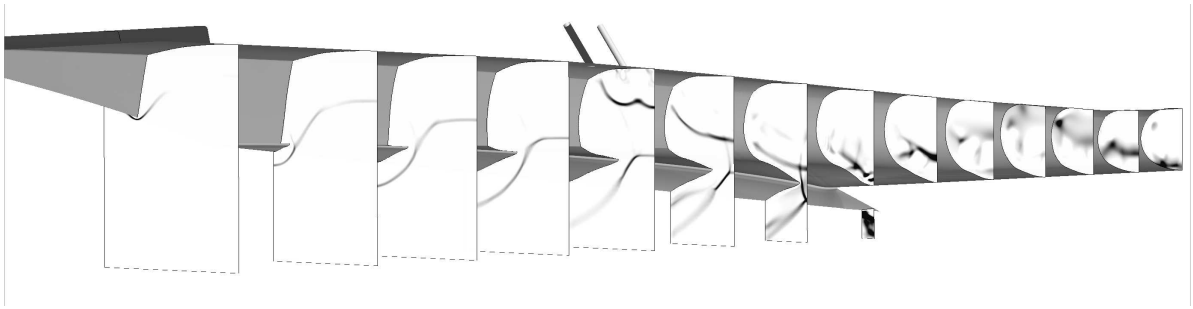

(a) Cross-plane Structure

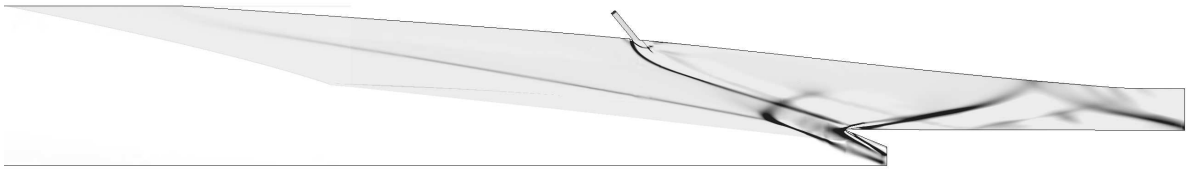

(b) Symmetry Plane Structure

Figure 4.6: Shock structures in the fuelled inlet

The center of the combined injector bow shock reflects off the cowlside wall just downstream of the formation of the cowl closure shock, and merges with it to form a single strong shock wave. The slightly higher angle of this shock, and the change in local speed of sound in the hydrogen-rich boundary layer, causes this combined shock to reflect off the bodyside wall upstream of the point that the closure shock reflects in the unfuelled inlet. This in turn leads to the observed impingement of the bodyside reflected shock off the cowlside wall at the end of the inlet, whereas in the unfuelled simulation it is still transiting between the bodyside and cowlside surfaces at the throat.

A second, weaker shock is also observed in the presence of inlet fuelling, and is most clearly observed in the symmetry plane shock structures shown in Figure 4.6b. This weak shock is formed by the recompression of the expanded flow in the wake region just downstream of the fuel injector plume. Much like the inlet bow shock, this recompression shock propagates away from the bodyside wall, and reflects off the cowlside wall downstream of the cowl notch. At the point where this weak shock intersects with the combined cowl-closure/bow shock reflection off the bodyside wall, a number of weak compression waves are observed. These waves are caused by the inlet turning, and are also present in the unfuelled case.

4.2.2 Inlet Pressure Contours and Experimental Comparison

The pressure contours shown in Figure 4.7 are closely coupled to the inlet shock structures. The contours also bear a strong resemblance to those of the unfuelled inlet, though with a noticeably higher pressure near the inlet throat. With a pressure gain of nearly 20% relative to the unfuelled case, this indicates that the fuel may be burning near the throat.

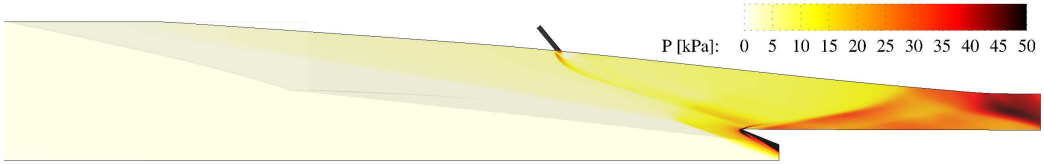

(a) Symmetry Plane Structure

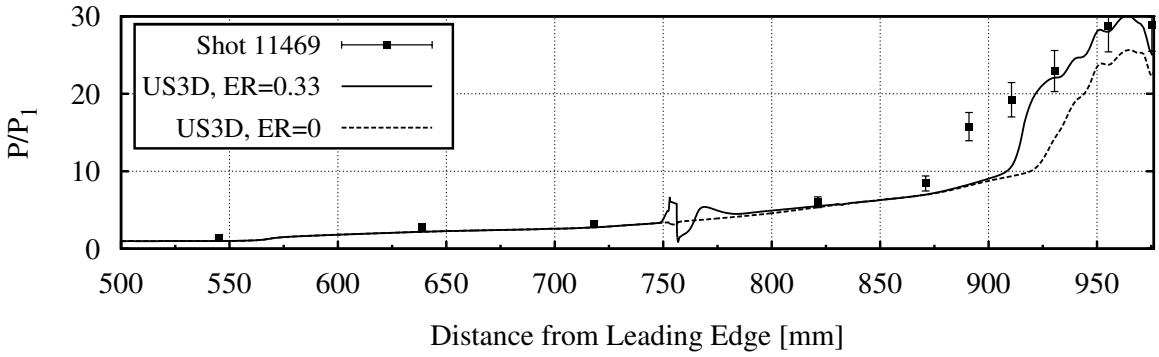

(b) Body-Side Centreline Pressure Distribution

Figure 4.7: Fuelled inlet centerplane pressure contours and bodyside pressure distribution

The small jump in pressure just downstream of the injection point in Figure 4.7b gives a clue as to the extent of the injector wake region. The pressure distribution downstream of this wake region then returns to the same levels observed in the unfuelled case, up until the shock reflection point. Both plots of pressure confirm that the reflection of the cowl closure shock off the bodyside wall occurs further upstream than it does in the unfuelled case. This earlier shock compression of the flow may also contribute to the higher pressures in the isolator.

This pressure rise near the inlet throat shown in Figure 4.7b also agrees well with Shot 11469 of Wise (2014), which has the same inlet equivalence ratio as this simulation. The experimental data shows a more gradual pressure rise beginning 870 mm from the leading edge, whereas the simulation shows a steeper pressure rise occurring at x = 910 mm. This mismatch is observed in the nitrogen-fuelled experimental data as well, and it may be related to a long-time-scale transient effect during testing in the T4

facility. By x = 929 mm, the reacting simulation and the experiments once again show good agreement.

This return to close agreement in the last 50 mm of the inlet length is likely due to the mixing behaviour inherent to RANS modelling. The RANS method being used in this study tends to under-predict the near-field mixing of fuel compared to the more accurate mixing behaviour a Large-Eddy Simulation (LES) model would predict in the first 50D downstream. However, a previous study comparing the same RANS method with wall-modelled LES of fuel injection and mixing in a turbulent, axisymmetric Mach 8 scramjet inlet found that the RANS model predictions begin to over-predict mixing further downstream, but still provide a reasonably accurate estimate of mixing (Peterson et al., 2013). Once the RANS model begins to match or over-predict the true mixing behaviour, the uniformly distributed fuel allows for rapid combustion to occur, giving the RANS model an accurate magnitude of the pressure rise at the inlet throat.

4.2.3 Inlet Vortex Structures

The addition of fuel along the bodyside surface of the inlet also has an effect on the vortical structures generated by the inlet, as illustrated in Figure 4.8. One of the most noticeable changes is the appearance of the counter-rotating fuel plume vortices immediately downstream of the injector location. The vortices generated by the off-center injectors do not remain parallel to the bodyside wall in the vertical (body-cowl) direction: as seen in Figure 4.8a, the inner vortex moves closer to the bodyside wall, while the outer vortex moves outward away from the wall. This indicates that the off-center fuel plume is inclined toward the centerplane as a result of the transverse pressure gradient imposed by the sidewall contraction.

The contours of Q-criterion in Figure 4.8a also show that the top corner vortex is temporarily disrupted when the fuel is injected. This is a result of the passage of the injector bow shocks outward along the bodyside wall. The predominantly rotational region is not entirely disrupted, however, and the Q contours suggest that a coherent vortex re-establishes further downstream. The motions of the streamlines in Figure 4.8b appear to confirm this: flow that passes through the high-Q top corner region downstream of the injectors continues to be swept inward toward the centerplane.

One final difference between the fuelled and unfuelled cases is in the behaviour of the swept separation. Unlike the unfuelled case, there appears to be no strong regions of rotational flow at the inlet throat. This is due to the earlier arrival of the reflected shock at the cowl wall relative to the unfuelled case: the shock passes through the

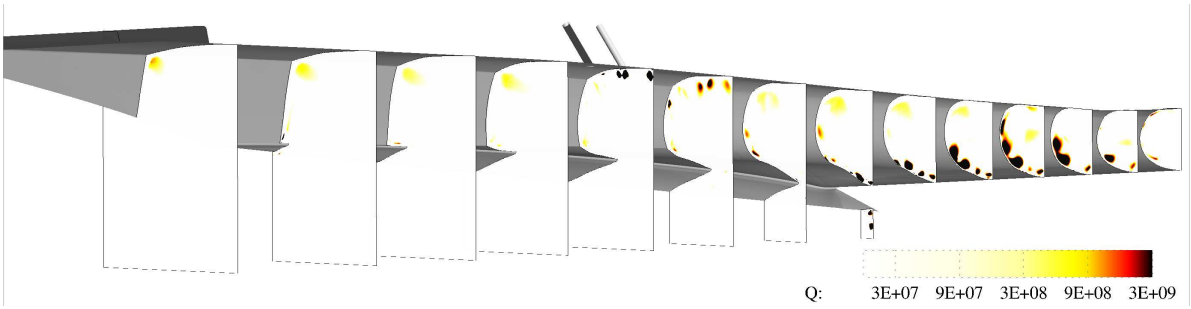

(a) Cross-plane Q criterion magnitude

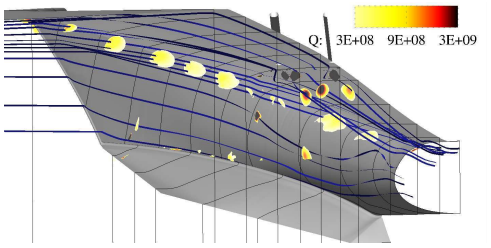

(b) Streamlines through top-corner vortex

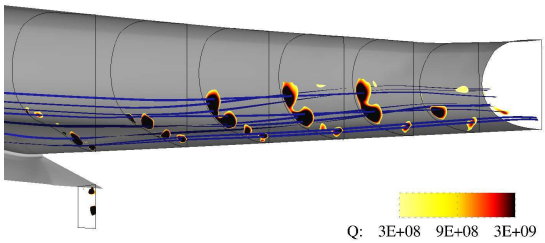

(c) Streamlines through swept separation vortices

Figure 4.8: Fuelled Inlet Vortical Structures

individual vortical regions before they have a chance to merge. These smaller vorticity regions are compressed and weakened by the arriving shock wave, and remain distinct at the throat: the streamlines in Figure 4.8b do not coalesce to the same extent as the unfuelled case.

4.2.4 Flow Variable Contours

The changes in the shock and vortical structures also have varying effects on the fundamental flow field variables. In comparison to the unfuelled case, the effects of injecting hydrogen on the boundary layer structure of the inlet flow field are remarkably small. This is likely due to the fact that the fuel is unable to escape the thick boundary layer bubble, as demonstrated by the Mach number contours shown in Figure 4.9a. Despite some distortion of the boundary layer profile, which is caused by the individual fuel injector bow shocks, the "bubble" shape is quickly re-formed as the jets mix and combine into a single structure inside the turbulent boundary layer. This lack of significant affect on the mainstream flow has an additional effect: inlet mass capture is unaffected by the addition of fuel injection.

An examination of the density contours is less informative: since the overall mass of the injected fuel is low, the fuel had little effect on the bodyside boundary layer

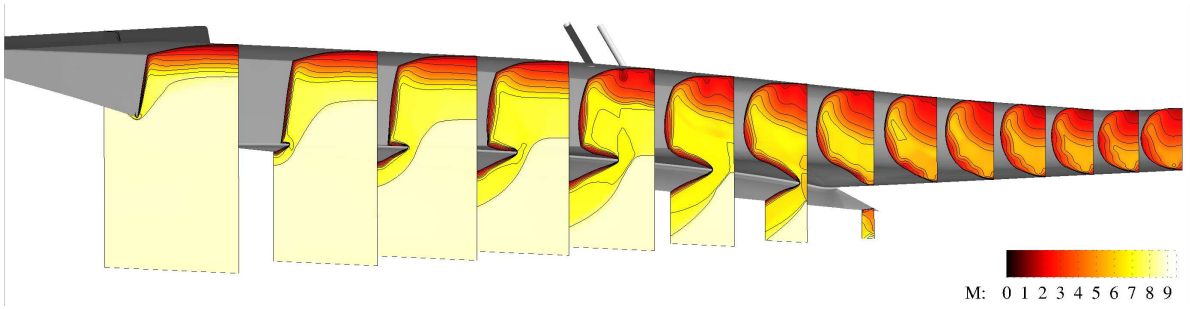

(a) Mach Number

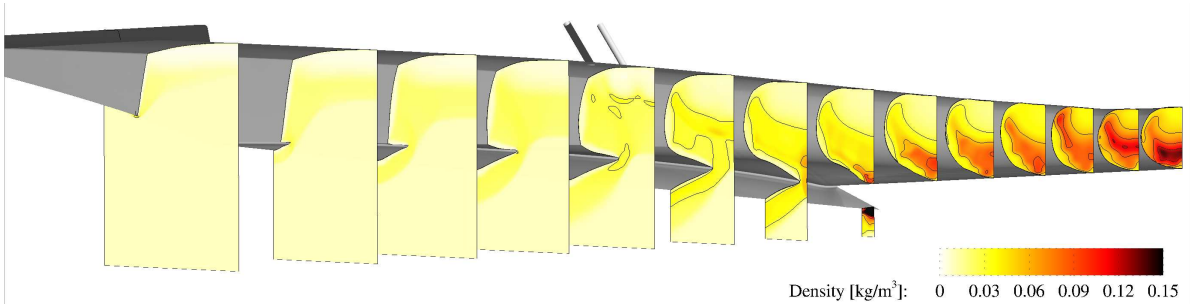

(b) Density

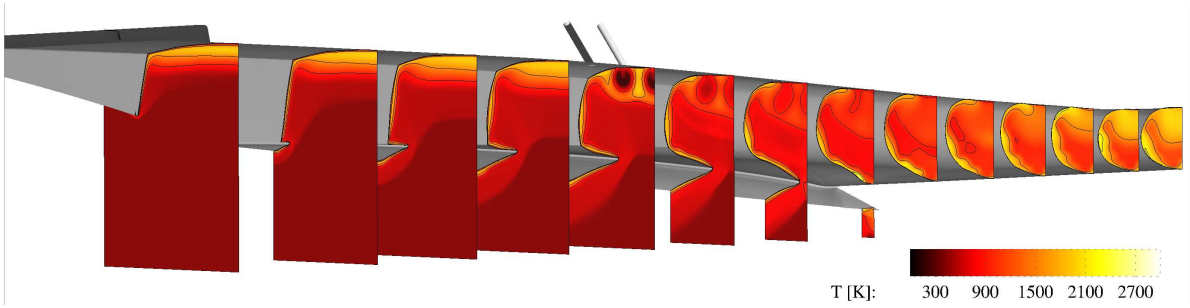

(c) Temperature

Figure 4.9: Flow Characteristics of the Fuelled Inlet Flow field

density profile. Figure 4.9b does show that the high density region at the throat is somewhat closer to the cowl wall, due to the bodyside reflection of the cowl closure shock moving further upstream. There may also be some increases in density due to thermal compression if the injected fuel is, in fact, burning.

The temperature field shown in Figure 4.9c, on the other hand, yields a wealth of information about the effects of fuel injection on the inlet. The highly under-expanded fuel jets rapidly cool to a temperature below 100 K as they expand out of the injector orifices, which leads to a significant decrease in the temperature of the bodyside boundary layer. This low-temperature region persists for a significant distance downstream of the injection point. The boundary layer temperature begins to recover first at a point along the edge of the bodyside boundary layer, and then rises rapidly across

the boundary layer as the flow reaches the throat. It is interesting to note that the temperature across much of the boundary layer bubble remains at 1000 K for a significant time before the boundary layer temperature rapidly rises to over 2000 K. This mode of temperature rise downstream of injection, combined with the rapid rise in pressure shown in Figure 4.7, suggests the fuel cannot ignite until the inlet compression process is nearly complete. It is only after the cowl closure shock interacts with and compresses the hydrogen-rich boundary layer that the fuel begins to burn. It is therefore unlikely that an increase in the wall temperature to a "flight condition" temperature much higher than the current 300 K wall temperature will have a significant impact on the ignition of the injected fuel.

4.2.5 Ignition and Combustion of Injected Hydrogen

In order to confirm whether or not the fuel is burning, and to determine how the fuel ignites, it is necessary to investigate the production of reaction product species along the length of the inlet. While the absence of H_2 implies the initiation of combustion, and the presence of H_2O confirms that some of fuel has reacted to completion, the relative concentration of OH radicals can be used to determine the location where combustion reactions are taking place.

The mass fraction of OH radicals is contrasted with the mass fraction of molecular hydrogen in Figure 4.10. The hydrogen jets mix rapidly downstream of the injectors, and quickly blend into a uniform, outward-diffusing structure. This is due to the turbulence of the boundary layer enhancing fuel mixing with the surrounding air.

More interestingly, the distribution of the hydrogen shown in Figure 4.10a does not mirror the shape of the boundary layer bubble exactly: the contours move inward toward the symmetry plane near the body-side wall. This inclusion in the hydrogen-containing region matches well with the production of OH shown in Figure 4.10b. While production of OH is expected where hydrogen interacts with the injector shock structures, there is also a persistent OH “hot-spot” near the corner where the bodyside and sidewalls meet.

This persistent structure aligns well with the boundary between the thick, hydrogen-cooled bodyside boundary layer and the thin high-temperature boundary layer being swept inward by the top corner vortices. Given that the inward-moving boundary layer has a temperature in excess of 1500 K, it is not surprising that the hydrogen reacts with the available oxygen in this region. The OH-rich zone grows with downstream distance until some critical point at which the high-OH mass fraction region expands

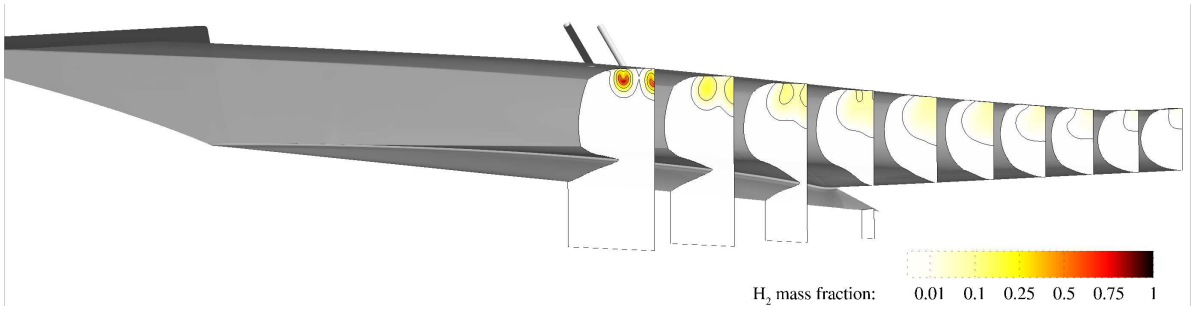

(a) H_2 Mass Fraction

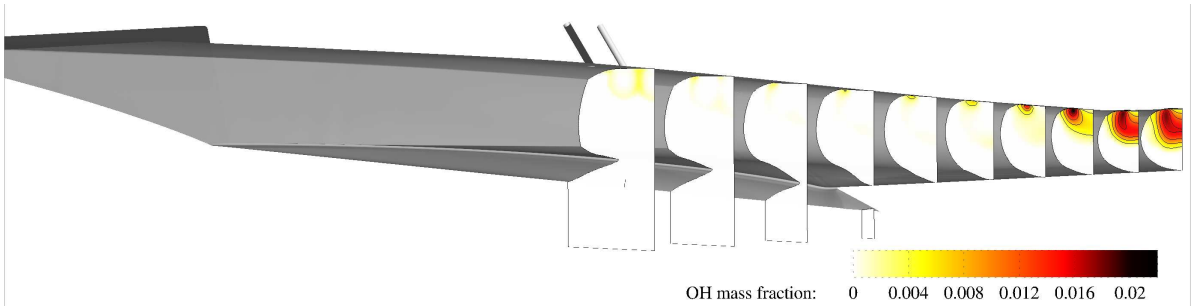

(b) OH Mass Fraction

Figure 4.10: Distribution of H_2 molecules and OH radicals in Fuelled Inlet.

rapidly across the fuel-rich boundary layer, indicating the onset of robust combustion near the inlet throat.

This onset in combustion appears to correlate with the Shock-Boundary Layer Interaction (SBLI) that occurs as the cowl shock crosses the bodyside boundary layer to reflect off the bodyside wall. While it may be tempting to assume the SBLI is also responsible for igniting the fuel, the previously mentioned 'hot-spot' along the bodyside wall is what truly drives the ignition process, as demonstrated in Figure 4.11. At the point shown in Figure 4.11a the shock has just begun to process the boundary layer flow, and the greatest concentration of OH radicals is present in the hot-spot.

In this persistent hot-spot, the local temperature is sufficient to permit hydrogen reactions to take place. This process is sustained as reaction products diffuse toward the engine centreline, and unreacted fuel and oxygen move into the hot-spot. At a point 10 mm downstream of the point at which the center of the shock has reflected off the bodyside wall, the hot-spot has begun to grow considerably in size, as shown in Figure 4.11b. This local growth is an effect of the increase in boundary layer pressure due to the SBLI, though there may also be some local thermal compression effects from the on-going combustion reactions. Regardless, the entire boundary layer sees a dramatic increase in temperature and pressure due to its interaction with the impinging shock.

Downstream of the plane shown in Figure 4.11b, the flame spreads rapidly from the

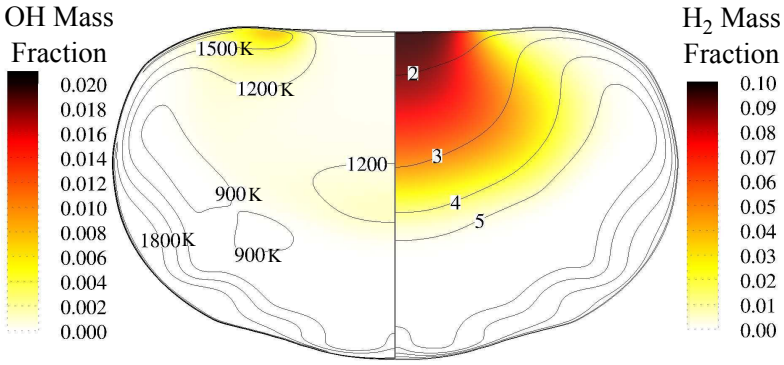

(a) x = 0.83

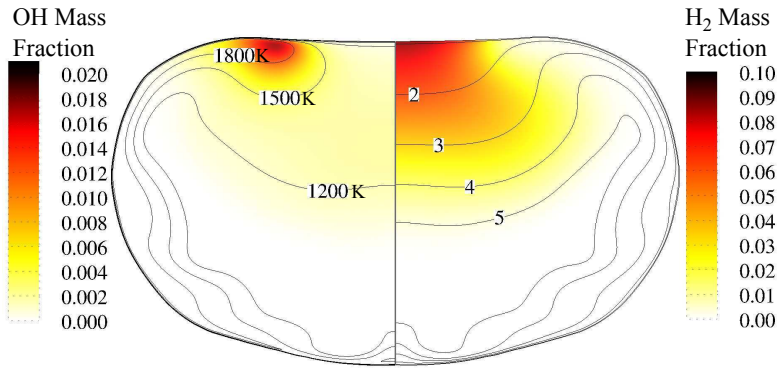

(b) x = 0.9

Figure 4.11: Cut plane views of inlet fuel ignition. Slices show OH mass fraction and Temperature (lines) on the left, and H_2 mass fraction contours and Mach number (lines) on right.

hot-spot across the boundary layer as temperatures in excess of 1200 K allow for rapid combustion of the hydrogen with available oxygen. Robust chemical reactions are no longer confined to the small hot-spot, but are occurring along the entire bodyside surface of the inlet. While the shock is not responsible for igniting the fuel, it provides the conditions that allow for widespread burning at the inlet throat.

4.3 Fuelled Inlet Performance

Having examined the flow physics of the fuelled inlet, the focus now turns to what effect the observed changes have on the inlet's performance. However, with fuel injection included, the efficiencies presented in Section 4.1.4 are no longer an appropriate measure of performance. Since the adiabatic kinetic energy efficiency is based on flow enthalpies, accounting for the sudden, irreversible change in enthalpy due to fuel injection makes this performance metric problematic to interpret and solve. While the inlet compression ratio is still easily calculated, alone it will not determine how much

additional compression is from the act of injecting fuel alone, and how much is due to the fuel burning at the throat.

However, several other measures of inlet performance can be extracted for an inlet with fuel injection: the combustion and mixing efficiency of inlet fuel injection are of interest, as this gives some measure of how much additional heat and combustion radicals are available to pilot combustion of fuel injected further downstream in the engine. This heat release may also lead to thermal compression of the captured air, further enhancing downstream combustion. The amount of thermal compression can be extracted by comparing the average pressure at the throat of reacting flow simulation with that of a nitrogen-inflow simulation in which combustion is suppressed. Finally, what effect low equivalence ratio fuel injection has on the inlet drag is of some importance, as additional drag will have an opposing effect to the combustion benefits inlet injection can provide.

4.3.1 Mixing and Combustion Efficiencies

It wis logical to begin by examining the efficiency with which fuel is mixed and burned in the inlet-captured air, as this has implications for the other performance measures in the scope of this study. Combustion efficiency is defined as the fraction of hydrogen fully converted into water vapour,

$$\frac{0.1119\dot{m}_{H_2O}}{\dot{m}_{H_{2,total}}}, \tag{4.3}$$

while the mixing efficiency was defined as the ratio of the mixed fuel mass flow rate, and the total fuel mass flow rate, as adapted from Axdahl et al.(2012). This ratio is defined mathematically as

$$\frac{\dot{m}_{H_2,mix}}{\dot{m}_{H_2,total}} = \frac{\int Y_R \rho u dA}{\int Y \rho u dA}, \tag{4.4}$$

where

$$Y_R = \begin{cases} Y_{H_2} & \text{if } Y \leq Y_{stoich}, \\ Y_{stoich}\frac{1-Y_{H_2}}{1-Y_{stoich}} & \text{if } Y > Y_{stoich}. \end{cases} \tag{4.5}$$

Unfortunately, scalar variance is not tracked in these RANS simulations, so what can be computed with the relation above is the mixedness of the time-averaged solution, rather than the true mixing efficiency. This computed quantity can reasonably be thought of as a macroscopic stirring efficiency. This is still useful since, as was mentioned previously, a study of fuel mixing in an axisymmetric turbulent Mach 8

inlet demonstrated that mixing and stirring efficiencies approach the same value well downstream of the point of injection (Peterson et al., 2013).

The mixing and combustion efficiencies are plotted in Figure 4.12, along with a line that indicates the closed area ratio of the inlet upstream of the throat. The conversion of hydrogen into water progresses slowly for some distance downstream of the injection point. However, at x = 900mm the change in combustion efficiency suddenly becomes non-linear, and rises rapidly across the last 10% of the inlet.

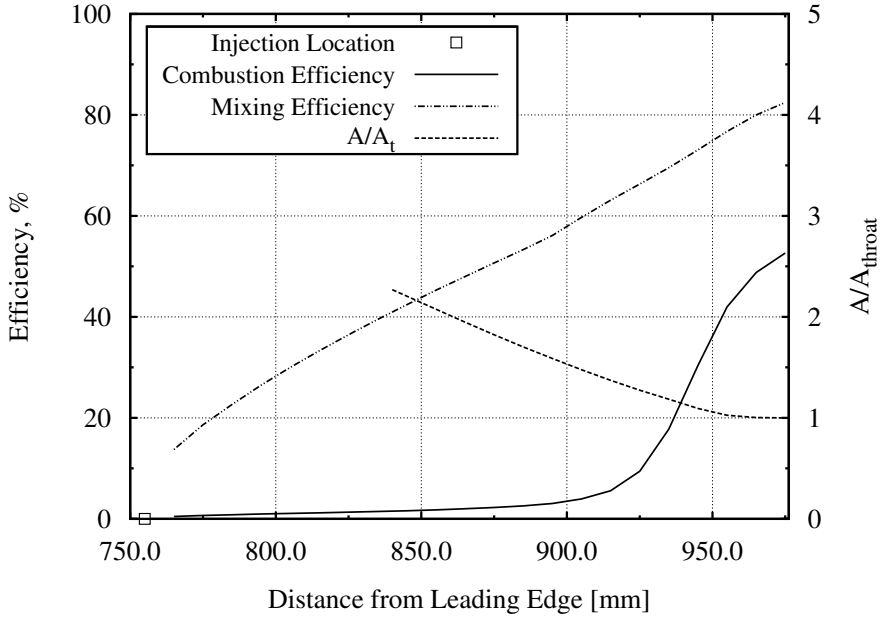

Figure 4.12: Fuel-based combustion and mixing efficiency in the M12 REST inlet

This rapid growth suggests that the well-mixed fuel-air boundary layer undergoes a reaction-limited combustion process, which helps to explain the rapid growth of OH and H_2O mass fractions at the inlet throat. Where the rise in combustion efficiency tapers in the last 5% of the inlet length, the readily available oxygen has been consumed and the reactions revert to a mixing-limited process. By the time the flow reaches the throat, 52.64% of the hydrogen ($\phi = 0.33$) is completely burned.

The rapid jump in combustion efficiency also coincides with the inflection visible in the mixing efficiency curve shown in Figure 4.12. This inflection in the mixing efficiency can be explained by referring back to the shock structures shown in Figure 4.6: this is the point at which the combined cowl/reflected bow shock begins to interact with the bodyside boundary layer. It is worth noting that the improvement in mixing efficiency due to the shock processing the boundary layer is a relatively minor effect.

The highly turbulent boundary layer has already achieved a mixing efficiency in excess of 50% before the shock encounters the boundary layer. The SBLI only enhances mixing further, yielding a mixing efficiency of 82.4% for this $\phi = 0.33$ fuelled flow at the inlet throat.

This 'shock-enhanced' combustion of the fuel, despite its intensity, occurs without causing the inlet to unstart. This could be beneficial to integrated engine performance: flow entering the combustor does so containing a high-temperature region full of partially-burned fuel. Such a feature can effectively act as a pilot flame to ignite fuel injected in the isolator or combustor downstream.

4.3.2 Thermal Compression

The heat release due to combustion also provides thermal compression of the inlet flow, as demonstrated by the temperature and pressure fields shown in Figure 4.13. By comparing the temperature and pressure fields for inlet fuel injection into nitrogen (where combustion is suppressed,) versus injection into air side-by-side, it is readily visible that combustion has a strong effect on not only the fuel-rich regions of the inlet flow, but also on the rest of the captured flow mass.

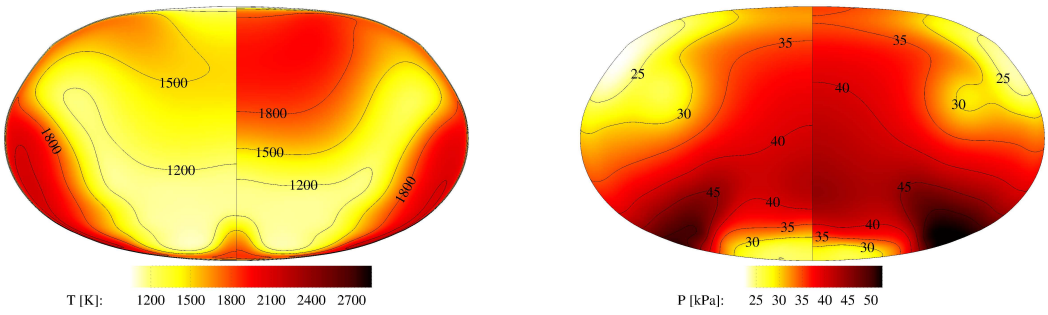

(a) Temperature Fields **(b)** Pressure Fields

Figure 4.13: Comparison of T and P fields at inlet throat for fuel into Nitrogen (left side of each subfigure) and fuel into air (right side)

While the temperature effects shown in Figure 4.13a are most dramatic in the burning bodyside boundary layer, there is also a small increase in the average temperature across the dense, unfuelled air flowing along the cowlside and sidewall surfaces of the throat. The pressure fields in Figure 4.13b show similar behaviour: there are higher pressures not only across the combusting region, but across the entire inlet cross-section. The horizontal pressure contours at the bottom center of the throat plane suggest that the reflected shockwave arrives at the cowl side of the inlet slightly faster when the fuel is allowed to react with the captured flow.

Table 4.1: Average Temperature and Pressure of unfuelled air at the inlet throat

Inlet Flow Case	T_{avg} [K]	P_{avg} [kPa]
Unfuelled	1380	31.8
Fuel into N_2	1396	38.2
Fuel into Air	1474	42.4

The extent of the thermal compression effect is estimated by taking the mass-flow-weighted averages of the temperature and pressure fields outside of the regions that the fuel reached in the flow. These results are summarized and compared with the unfuelled inlet in Table 4.1. The average temperature of the unfuelled air at the throat is approximately 5.6% higher when the fuel is injected and allowed to burn. The pressure change is also greater when the fuel burns, with a pressure increase of 33% over the unfuelled case, and 11% over the suppressed combustion case. Inlet fuel injection and burning therefore not only acts as a pilot flame for combustion downstream, it also provides additional compression of the incoming air, making conditions downstream even more ideal for rapid combustion.

4.3.3 Inlet Drag Forces

One common objection to fuel injection along the inlet surfaces of a scramjet is that the additional compression from fuel bow shocks, the additional mass of fuel in the inlet, and increased pressure due to combustion would all increase the drag generated by the internal surfaces of the inlet, where internal surfaces are defined as all those surfaces defined by the streamtracing technique used to design this inlet that compress the captured air. A large drag increase could easily overwhelm any gains in generated thrust caused by the longer reaction length available to inlet-injected fuel for combustion prior to exiting the engine.

In the case of the M12 REST inlet, such fears are unwarranted, at least at an equivalence ratio of 0.33. As demonstrated in Table 4.2, the total net increase in inlet drag due to fuel injection (excluding forebody drag) is between 4-5%, regardless of whether the fuel burns. The pressure drag of the combusting and fuel-into-nitrogen cases are virtually identical, varying by less than 0.5%. Viscous (skin friction) drag increases in the presence of fuel injection, and even more so in the presence of combustion.

There are a number of reasons that the increase in drag due to inlet injection is much lower than is commonly feared. First among these reasons is the fact that the fuel jet

Table 4.2: Increase in inlet internal drag components, normalized to unfuelled total drag

Flow Case	Inviscid Drag	Viscous Drag	Total Drag
Unfuelled	0.6358	0.3642	1.0
Fuel into N_2	0.6764	0.3666	1.0431
Fuel into Air	0.6734	0.3750	1.0484

is injected into the thick boundary layer bubble along the bodyside wall. Since the jet is shielded from the much higher-Mach-number mainstream flow, the corresponding bow shock that develops is substantially weaker than it would have otherwise been, causing a lower pressure drag increase than might otherwise be expected. The pressure rise due to mass addition is also low; the amount of additional mass flowing into the inlet from fuel injection is less than 1%, thanks to the low equivalence ratio and small molecular mass of hydrogen.

The thermal quenching effect caused by the expansion of the hydrogen jets inside the boundary layer could also have played a significant role in limiting the drag increase: while the increase in viscous drag suggests that any film-cooling effects due to the hydrogen are small, the cooled boundary layer allows the ignition of the fuel to be delayed until the flow is just upstream of the throat. In the portion of the inlet where combustion actually takes place, the upstream component of the wall surface normal vector is small, and trending towards zero. Most of the pressure rise due to combustion along the bodyside compression surface is therefore directed into a lifting force rather than drag, leaving the increase in skin friction due to the higher flow temperature being the primary effect of combustion in the inlet throat.

The slight combustion-induced increase in viscous drag seems to be the dominant cause of the increase in total drag forces between the suppressed combustion and fully combusting cases. This is likely because the heat-release induced increase in viscous drag occurs almost exclusively near the throat, where the viscous drag vector is parallel to the direction of overall inlet drag. The areas in which skin friction is generated are plotted for both unfuelled and combusting inlet flow in Figure 4.14.

The shear stress acting along the bodyside of the inlet can readily be seen to increase when fuel is injected and burned in the inlet. The high-stress region along the cowlside is seen to grow in extent and intensity thanks to the thermal compression effects of combustion, as shown in Figure 4.14a. Some additional viscous drag is caused by the presence of the inlet bow shocks, but as Figure 4.14b shows, the primary increase in bodyside shear stress occurs downstream of the reflection of the cowl closure

shock. This region is slightly expanded thanks to the shock reflection moving further upstream in the combusting flow case, and somewhat stronger due in part to the combusting flow in this region.

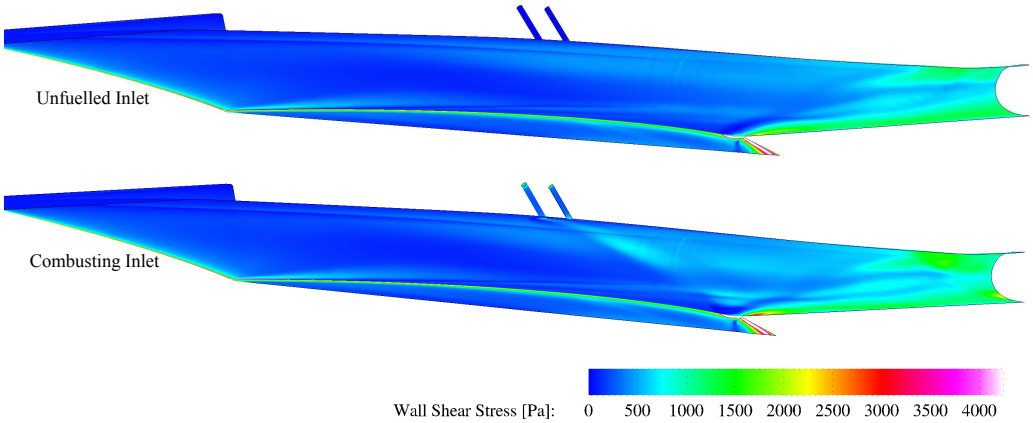

(a) Primary View

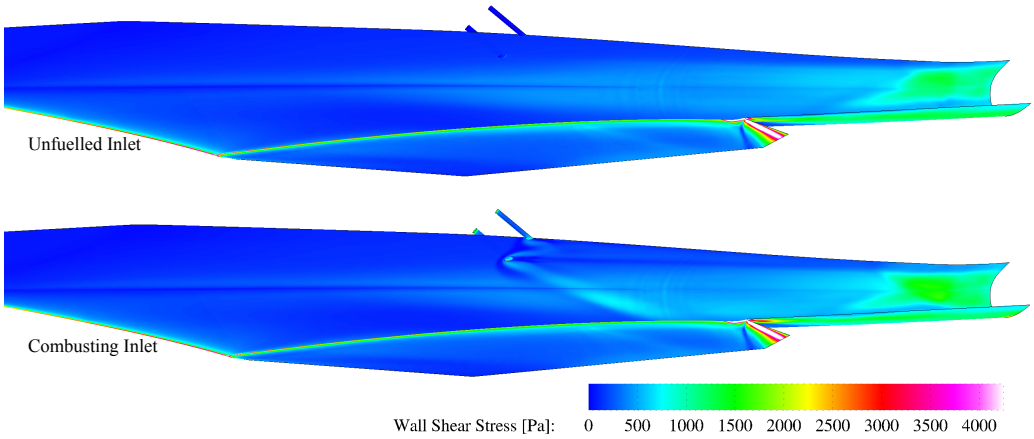

(b) Bodyside Detail

Figure 4.14: Two views of the inlet wall (viscous) shear stress distribution

The drag forces in Table 4.2 may be low relative to the drag induced in a real engine, as is indicated by the earlier pressure rise in the experimental data the simulations are compared with in Section 4.2.2. Therefore, the calculated drag provides a lower bound for the drag increase due to fuel injection. However, even with an earlier pressure rise, the region in which the experimental pressure-rise took place is still only a small portion of the inlet compression surfaces, and it is likely that the drag penalty would still be small relative to the total drag of the unfuelled inlet. This also suggests that increasing inlet fuelling ϕ above 0.33 may not be beneficial: an increase in fuelling would give rise to a greater induced drag on the inlet.

4.4 Summary

Having examined the flow physics present in the M12 REST inlet at near-design conditions, it is clear that the inlet efficiently captures and compresses the captured air, both with and without hydrogen fuel injection along the bodyside compression surface. The flow leaving the inlet and continuing onward into the isolator and combustor is found to be asymmetric along the body-cowl (z) axis. This non-uniformity is advantageous when fuel is injected in the inlet. The fuel jet is confined to the bodyside boundary layer, ultimately allowing it to ignite just upstream of the throat. This creates a source of heat, pressure and combustion radicals that are transmitted to the downstream components of the M12 REST, while causing only a small penalty to the inlet drag. How the examined flows leaving the inlet developed as they travel into and through the engine combustor, and how they affect any fuel injected into the combustor boundary layer, are the subjects of the next chapter.

5

Combustor and Nozzle Flow Physics and Performance

"Combustion is the hidden principle behind every artefact we create... Like our bodies and like our desires, the machines we have devised are possessed of a heart which is slowly reduced to embers."

– W.G. Sebald, *The Rings of Saturn* (1995)

Having examined the flow field of the M12 REST inlet in detail, both with and without fuel injection, the next step in this study is to examine the combustor and nozzle that process the air captured and compressed by the inlet. Much like Chapter 4, this chapter opens with an examination of the combustor and nozzle flow field in the absence of any fuel injection, and then examines how the combustor flow field changes when fuel is injected in the inlet upstream. The full baseline fuelling case, in which boundary layer injection is employed inside the combustor, is then examined. Where appropriate, the simulations are compared with experimental data, and conclusions are drawn regarding the global combustion efficiency of the M12 REST engine.

5.1 The Fundamental Unfuelled Flow Field

Knowing from the previous chapter that the unfuelled inlet flow field generates multiple flow features that are transmitted downstream into the isolator, a simulation of the entire engine without fuel injection is in order. The simulation inflow conditions have been set to match the nozzle exit conditions of experimental Shot 11486. These conditions are listed in Appendix E. While the inflow conditions have changed from those used in Chapter 4, the computational mesh used is identical to the mesh geometry presented in Chapter 3, with a total cell count of 34,789,776 cells.

The unfuelled simulation was allowed to run for 62000 iterations, at a typical time step of 2 x 10^{-8} s. This was sufficient to converge the RMS and energy residuals by 6 orders of magnitude, with a net inflow/outflow mass balance of 0.555 x 10^{-6} kg/s. The momentum balance in all three cartesian coordinate directions was less than 0.5 x 10^{-3} N. The mass capture of the inlet was found to be 57.85 g/s.

The combustor grid has sufficient resolution to reconstruct the velocity profile of the boundary layers in the combustor and nozzle sections. The non-dimensional height, y^+, of the first cell next to the wall is less than unity throughout nearly the entire engine downstream of the combustor step, and is never higher than 1.2. It can therefore be stated with confidence that the viscous sublayer of the wall-bounded flow is adequately resolved through the entire combustor and nozzle, and there are negligible grid effects on the development of the combustor boundary layer and flow field.

5.1.1 The Structure of the Combustor Flow Field

As with the inlet flow field before it, the flow field of the unfuelled combustor and nozzle can be described by examining the development of the Mach number, temperature and density fields downstream of the inlet. The contours for these variables are shown in the cross-stream slices shown in Figure 5.1. As a convention, this figure and all subsequent figures have their cross-stream slices are taken in the same place. The first slice is taken at x =1019.5 mm, just upstream of the combustor step. The streamwise locations of the centroid of the constant-area combustor slices downstream of the combustor step are 1037.12, 1056.35, 1075.60, 1094.85, 1114.10, 1133.35, 1165.26, 1171.85, and 1191.10 mm from the leading edge. The slice centroids in the diverging combustor are taken 1215.0, 1239.0, 1263.0, 1287.0, and 1311.26 mm from the leading edge. The four nozzle slices have their centroids placed 1335.0, 1365.0, 1400.0, and 1436.5 mm from the leading edge. Because the engine surfaces change direction at the combustor step to realign with the nominal direction of vehicle flight, the slices inside the combustor and nozzle are inclined 6° relative to the solution domain x-axis, such that they remain normal to the direction of bulk flow in these sections of the engine.

The Mach number contours presented in Figure 5.1a behave much as they would be expected to, based on the inlet flow field. The large, low-Mach-number regions along the bodyside and sides of the isolator continue into the combustor, with a high-speed core of flow moving along the cowlside of the combustor. This high-speed core appears to fluctuate between Mach 4 and Mach 5, which is likely due to the shock wave reflections and their associated changes in temperature and pressure. The con-

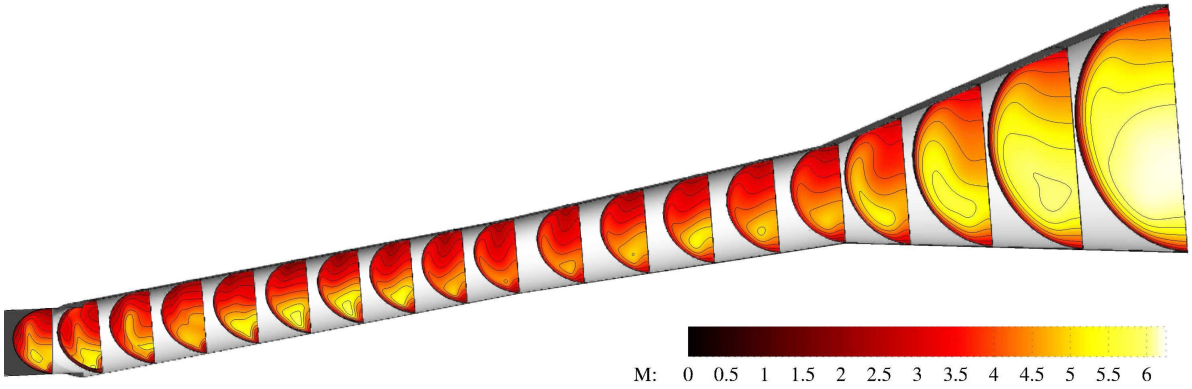

(a) Mach number

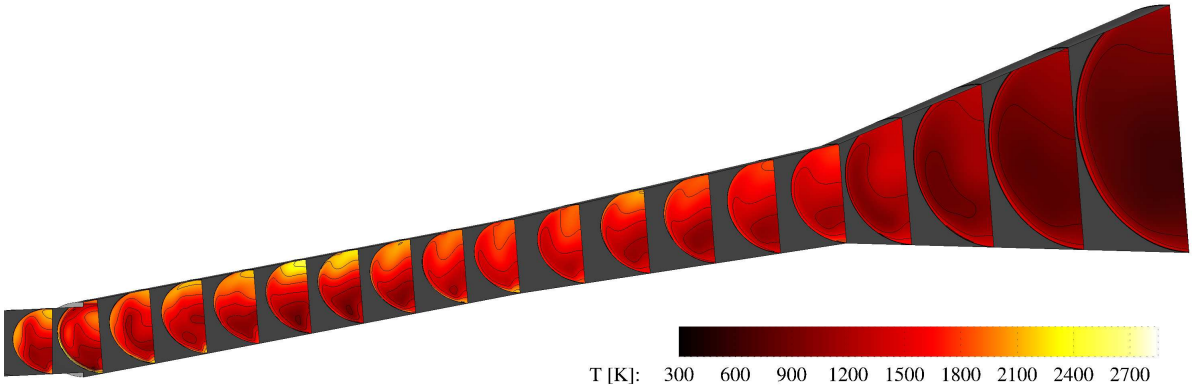

(b) Temperature

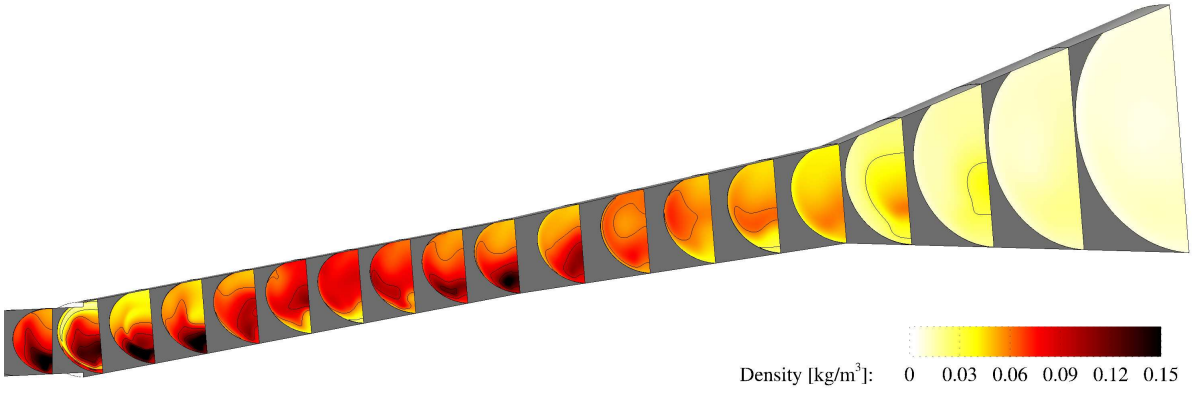

(c) Density

Figure 5.1: Cross-stream planes showing the development of the bulk variables of the combustor flow field.

tours also indicate a comparatively thin boundary layer on most of the cowlside wall surfaces of the combustor sections.

Visible along the cowlside wall at the centerplane is a low Mach number region corresponding to the counter-rotating vortex pair described in Section 4.1.3 of the previous chapter. While the coherent vortex has been compressed and weakened by the cowlside reflection of the inlet closure shock, the low-Mach number region remains intact. This structure appears to elongate at first, driving upward along the centerline of the engine and into the flow. The extent of this region then appears to fluctuate, moving closer the wall, and spreading outward at the exit of the diverging combustor. The flow remains entirely supersonic throughout the combustor, with the core of high-speed flow accelerating to speeds in excess of Mach 6 at the exit of the Nozzle.

The temperature fields in Figure 5.1b show that the flow temperature remained higher in the thick, turbulent boundary layer flowing along the bodyside wall, with a temperature in excess of 1800 K for the entire length of the combustor. Meanwhile, the flow on the cowlside appears to remain at temperatures between 600 K and 900 K. The remains of the high-temperature swept separation, whose discrete vortex structure is weakened upstream in the inlet, joins to and then merges with the bodyside high temperature region in the first half of the constant area combustor.

The average temperature in both the bodyside and cowlside regions fluctuates along the length of the constant-area and diverging combustor sections. Moreover, the peaks in temperature in the bodyside and cowlside flow regions appear to alternate. For example, the combustor side sees an increase in temperature just downstream of the combustor step, coinciding with the turn in the flow path and its corresponding decrease in Mach number. As the cowlside flow begins to cool downstream, a spike in bodyside core flow temperature up to 2400 K is observed. This alternating temperature behaviour is a strong indicator that the flow is shock-dominated through the combustor.

The density flow field shows considerably more variation than it does in the inlet. As demonstrated in Figure 5.1c, the high density, high speed core of flow continues from the isolator into the combustor. As this flow is forced to turn by the cowlside wall, the high-density region appears to lose coherence. By the midpoint of the constant-area combustor, the most dense region of flow is located along the center of the engine cross-section. This density distribution does not last, with another high density region forming along the cowlside wall just upstream of the junction between the constant-area and diverging combustor sections.

Downstream of this point, the density flow field gradually diminishes in magnitude as the engine cross-section expands. This process accelerates rapidly as the flow enters the nozzle. While there appears to be a consistent density bias toward the cowlside of the engine, much of the density contour shapes suggest that the density flow field is strongly coupled to the combustor shock structure downstream of the point where the flow is turned and compressed by the cowlside wall.

5.1.2 The Unfuelled Shock Structure

One of the most notable flow structures exiting the unfuelled inlet in the previous chapter was a strong, curved oblique shock that reflects off of the bodyside wall of the inlet. To examine how the shock structure develops downstream of the throat and interacts with the shock-dominated flow suggested by the flow variable contours, the symmetry plane shock structure of the entire engine (excluding the forebody) and cross-stream shock contour slices are examined. Where appropriate, dashed lines have been added to Figure 5.2 to indicate the location of strong expansion waves in the engine.

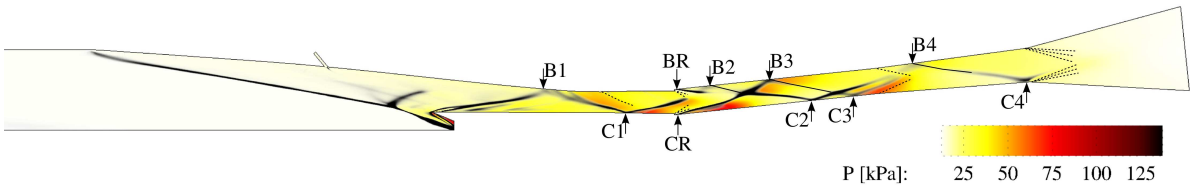

(a) Symmetry Plane

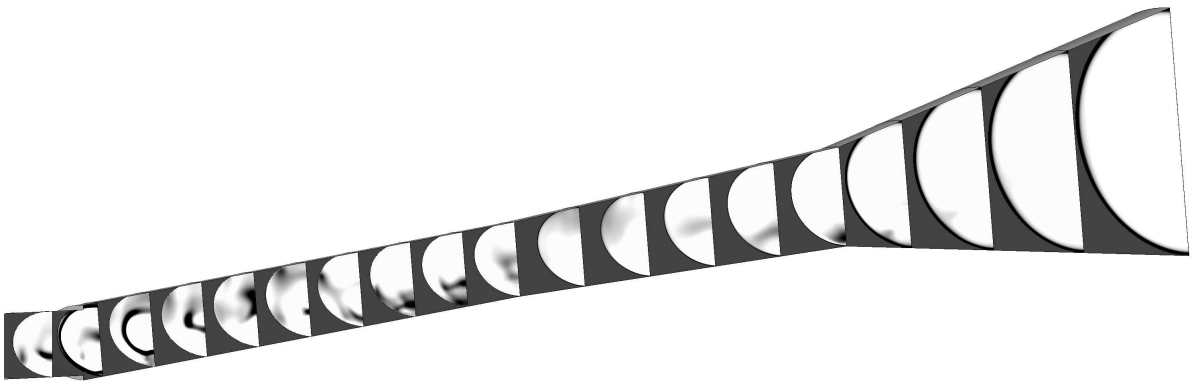

(b) Cross-stream planes

Figure 5.2: Contours of the combustor shock structure and symmetry plane pressure field.

The symmetry plane shock structure, shown in Figure 5.2a, demonstrates that after its bodyside reflection (B1), the cowl closure shock reflects off the cowlside wall of the isolator just downstream of the engine throat at location C1. There is a high pressure region behind the shock extending from B1 to C1, which is relieved by the expansion

wave generated by the final turn of the captured flow into the isolator. Downstream of the cowlside shock reflection, the curved shock crosses the engine for a third time, and is midway across the engine's body-cowl span as the flow exits the isolator and passes over the combustor step.

The sudden area change behind the combustor step causes a strong rarefaction fan to form. This expansion in turn weakens the shock transmitted from the isolator into the combustor, as is visible in the first slice downstream of the step in Figure 5.2b. Just downstream of the step, a quasi-symmetric recompression shock forms, which is labelled as points BR and CR in in the symmetry plane shock structure image. While it appears to be symmetric at first, this elliptical recompression shock weakens quickly on the bodyside of the engine, while it remains strong on the cowlside. This is due to the direction change of the flow inside combustor, which occurs at the combustor step; the cowlside wall acts as a compression ramp, while the bodyside flow experiences an expansion process as the wall turns away from the dominant direction of the flow in the isolator in order to realign with the direction of nominal vehicle travel.

Midway between the generation of the elliptical recompression shock, and its convergence near the body-cowl centerplane of the combustor, the inlet-transmitted shock reflects weakly off the bodyside wall, shown as point B2 in Figure 5.2a. All the bodyside shock reflections that occur from this point downstream are sufficiently weak that lines are superimposed on the image to clarify their paths on the symmetry plane. This weakly-reflected shock crosses the engine at the same point where the elliptical recompression shock converges, leading to the strong pair of shocks that emerge from this region: one travelling toward the cowlside surface, and the other toward the bodyside. Travelling through the hotter (and therefore higher speed of sound) flow along the bodyside of the combustor, the upward-travelling shock impinges off the bodyside wall quickly (B3). Following this reflection, it travels across the combustor, ultimately reflecting off the cowlside wall (C3) shortly after the the downward-travelling portion of the post-convergence recompression shock reaches the cowlside wall centerline (C2).

Downstream of CR, the centerplane pressure field can be seen to weaken significantly before its recompression by the arrival of a shock at C2. This corresponds closely to the observed increase in local Mach number and decrease in local density in this part of the constant-area combustor, which were remarked on in Section 5.1.1. This is caused by a re-expansion of the flow downstream of the shock structures presented so far.

The two shocks reflecting at C2 and C3 remain distinct until the flow reaches the en-

trance to the diverging combustor. At this point, the reflected shocks both pass through the expansion wave created by the divergence of the combustor walls. Downstream of this point, they appear to merge into a single shock that reflects off the bodyside wall at B4, and then again off the cowlside wall just upstream of the nozzle entrance (C4). Downstream of this point the shock passes through the strong nozzle expansion, and has weakened sufficiently to be no longer visible in the shock contour plot.

A shock-dominated flow such as this has no stable, coherent vortex features, as the reflecting shock waves rapidly destroy any streamwise rotating structures that start to form. Mixing can therefore be expected to be dominated by shear layer mixing and the local vorticity generated by the passage of a shock wave, which temporarily induces lateral rotation and mixing in the regions immediately upstream and downstream of shock and expansion wave structures. Having a strong understanding of the nature of the unfuelled flow through the M12 REST combustor and nozzle, the net drag performance of the unfuelled engine can be examined.

5.1.3 Wall Shear Stress and Aerodynamic Forces

The total aerodynamic forces acting on the M12 REST internal flow path is of great interest, as this provides a lower bound for the amount of thrust the engine would produce simply to maintain a Mach 12 cruise condition. Net uninstalled thrust, and therefore acceleration, requires thrust levels greater than this force. By examining the relative contribution of each engine component, it is therefore also possible to identify the areas of the engine where the greatest improvements can be made. Internal forces are the primary parameter of interest. While some external surfaces were modelled in the simulation, these are not be taken as indicative of true external engine surfaces, which would undoubtedly be unique to the vehicle such an engine is to be integrated onto.

The aerodynamic force data are obtained through post-processing subroutines of US3D, which allow for the viscous (skin friction) and inviscid (pressure) drag forces to be calculated for each surface defined in the computational domain. For this simulation, the wall boundaries are broken into three parts: the inlet/forebody surfaces, an isolator section terminating in the combustor step (to facilitate easy changes to injection geometry during the tailored injection study), and the combustor and nozzle wall sections. The total integrated forces are also extracted using the integration routines of Tecplot, which allows the individual force contributions of each section of the combustor and nozzle to be examined separately. This also makes it possible to separate

the contributions of the internal inlet surfaces from those of the forebody and other external surfaces.

Because the forces are integrated for each axial component, the solution domain x- and z-components of the aerodynamic forces acting on the internal surfaces must combined into a single drag or thrust force component in the direction of vehicle travel (the streamwise axis of the combustor and nozzle section). This calculation takes the form

$$F_{drag} = -\sum F_x \cos(6°) + \sum F_z \sin(6°), \tag{5.1}$$

where a net drag force has a negative magnitude, and a net thrust force has a positive magnitude. (Recall from Chapter 3 that the positive x-direction is defined in the primary direction of flow through the simulation domain, and the positive z-axis is defined as normal to the forebody surface, pointing away from the vehicle.) The resulting breakdown in the aerodynamic forces is summarized in Table 5.1.

Table 5.1: Aerodynamic thrust (+) and drag (-) forces acting on the unfuelled M12 REST engine

Engine Component	Inviscid Forces [N]	Inviscid Forces %	Viscous Forces [N]	Viscous Forces %	Combined Forces [N]	Combined Forces %
Inlet Internal Surfaces	-11.942	25.19	-10.919	23.04	-22.861	-66.39
Forebody	-5.133	10.83	-3.117	6.58	-8.250	-23.96
Isolator and Combustor Step	0.764	5.89	-1.395	2.94	-0.631	1.83
Constant Area Combustor	0.138	1.06	-7.233	15.26	-7.095	-20.61
Diverging Combustor	4.109	31.69	-5.016	10.58	-0.907	-2.63
Nozzle	7.956	61.36	-2.644	5.58	5.312	15.43
Total Drag Forces	-17.075	36.02	-30.324	63.98	-47.399 N	
Total Thrust Forces	12.967	100.0	-	-	12.967 N	
Total Internal Force	1.025 N		-27.207 N		-26.182 N	
Total Force	-4.108 N		-30.324 N		-34.432 N	

The total internal force in Table 5.1 is calculated excluding the forebody surface, while the "total force" includes this contribution. Fractional (percentage) values of individual forces are normalized based on the total drag force or thrust force as appropriate. The combined forces are normalized to the total force, with a sign to indicate a net drag or thrust contribution by that engine component.

While the combustor and nozzle surfaces experience some pressure forces acting in the direction of thrust, the unfuelled engine generates a large amount of drag globally, as expected. Viscous effects are responsible for nearly two-thirds of the drag forces generated, with the constant-area combustor accounting for 20% of the total drag force

acting on the surfaces considered. This underscores the importance skin-friction reduction techniques will likely have in producing a viable hypersonic vehicle at the hypervelocity conditions this engine is designed for.

While some external surfaces of the engine were modelled, they did not cover the entire length of the engine. While this means a meaningful installed engine drag cannot be calculated, it may be remarked that the external surfaces contained in the simulation produced nearly as much pressure drag as the inlet internal surfaces did. There may, therefore, be benefits to be gained in an installed engine from a careful design of the external inlet surfaces to minimize drag, particularly the external pressure drag generated downstream of the cowl notch.

Returning focus to the unintegrated (internal) flow path, it can be seen that while the inlet is the dominant source of drag, the isolator, combustor, and nozzle sections together produce 34% of the internal drag forces through skin friction alone - nearly half again as much viscous drag as is produced by the inlet's internal surfaces. Half of the combustor/nozzle drag is generated by the constant-area combustor, with a further third coming from the skin friction generated by the diverging combustor. The breakdown in combustor skin friction drag can be explained by examining the wall (viscous) shear stress magnitude distribution, shown in Figure 5.3.

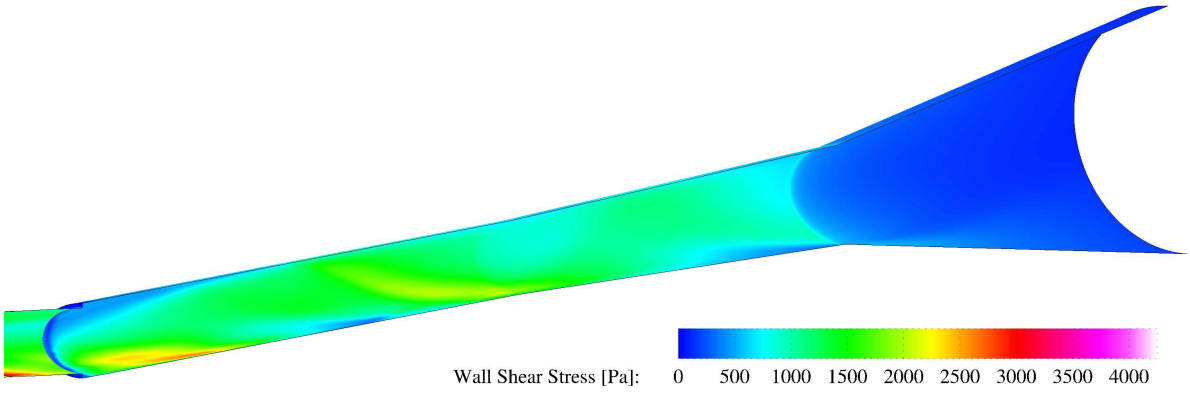

Figure 5.3: The distribution of local wall shear stress magnitude in the unfuelled combustor and nozzle.

The highest magnitude shear forces occur just downstream of the combustor step, along the centerline of the constant-area combustor cowlside wall. This coincides with the point at which the high-density, high-speed core of flow entering the combustor is forced to turn upward as it encounters the cowlside combustor wall. The regions of high wall shear are mostly confined to the cowlside and side walls in the constant-area combustor, with peaks in the shear stress occurring near shock impingement locations.

The shear stress contours also mirror the locations of expansion regions, including the zone that lies between the cowlside recompression shock origin, and the first cowlside reflection downstream. The nozzle has low levels of shear stress, due to the lower density and temperature of the expanded flow.

5.1.4 Comparison between the simulation and experiment

As was mentioned at the start of this section, the CFD simulation used inflow conditions calculated from the tunnel conditions of experimental Shot 11486 of the joint campaign. This in theory allows the simulation to be compared directly to the experimental data without the need for the normalization which was used in Chapter 4. However, for reasons that are explained in detail in Appendix F, this is not possible for experiments conducted using the under-tailored Mach 12 tunnel condition. The correct method for analyzing and comparing the experimental data with a steady state, time-averaged RANS simulation is indeed using the normalized pressure data. The simulated and experimental pressure distributions are therefore normalized to the nominal forebody pressure, P_1, and compared as shown in Figure 5.4.

There is very good agreement between the experimental data and the simulation along the bodyside of the inlet, both in the location of the pressure rise and in its magnitude. The simulated and experimental pressure peaks associated with shock reflections within the combustor are, for the most part, in roughly the same location, and of the same magnitude. However, the pressure rise of the last bodyside and cowlside shock reflections in the combustor appear to be significantly higher in the experiment.

Given the length of the engine, it is possible that this last shock wave reflection is a transient effect caused by the difficulties in choosing an appropriate test time in the experiment. As explained in Appendix F, this shock would likely have diminished in strength if the test time were chosen later, however this would have lead to a much higher level of driver gas contamination. The slight difference in the location of the bodyside shock reflection may also have some effect: the shock in the simulation passes through the diverging combustor corner expansion fan, which would weaken the shock relative to the experiment, where it reflects upstream of this corner. There may also be an effect due to mesh resolution in this region: if the first cell has a y^+ greater than unity, or the cell spacing is too wide in the region of shock reflection, the pressure would be under-predicted by the simulation.

Along the bodyside of the nozzle, the experimental data suggests a pause in the expansion of the flow. This is due to a reflection of the stronger shock that reflects off the

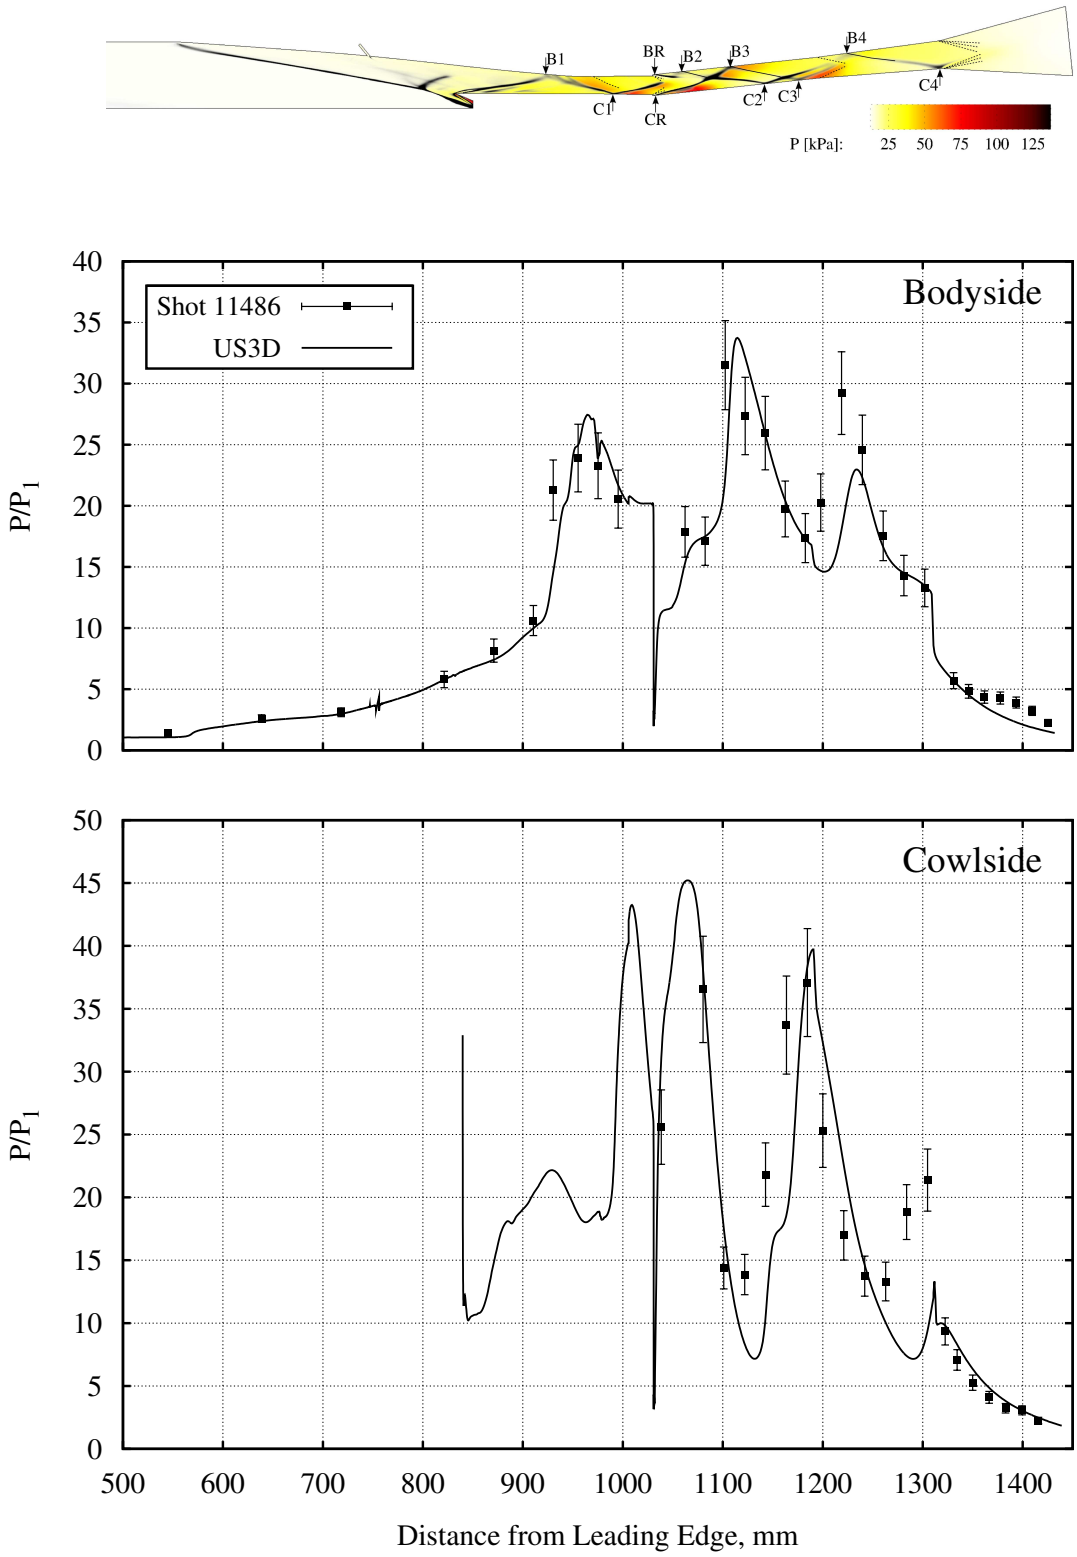

Figure 5.4: Comparison of the normalized simulation and experimental unfuelled engine pressure distributions

cowlside wall just upstream of the nozzle entrance in the experimental data. The nozzle pressure distribution is otherwise well-matched by the simulation. This suggests that, notwithstanding the bodyside shock reflection, the nozzle pressure distributions are relatively insensitive to the strength of the cowlside shock reflection upstream.

Overall, the pressure data shows that the simulation provides a reasonable reconstruction of the pressure distribution observed in the experiment. Before declaring the simulation fully validated, the wall heat flux distributions from the simulation should be compared to the values measured in the experiment. While the simulation mesh was not specifically designed to accurately reproduce wall heat flux, the heat flux into the walls is of interest for estimating the heating loads for design of a flight-ready version of the engine. The results of this comparison are shown in Figure 5.5.

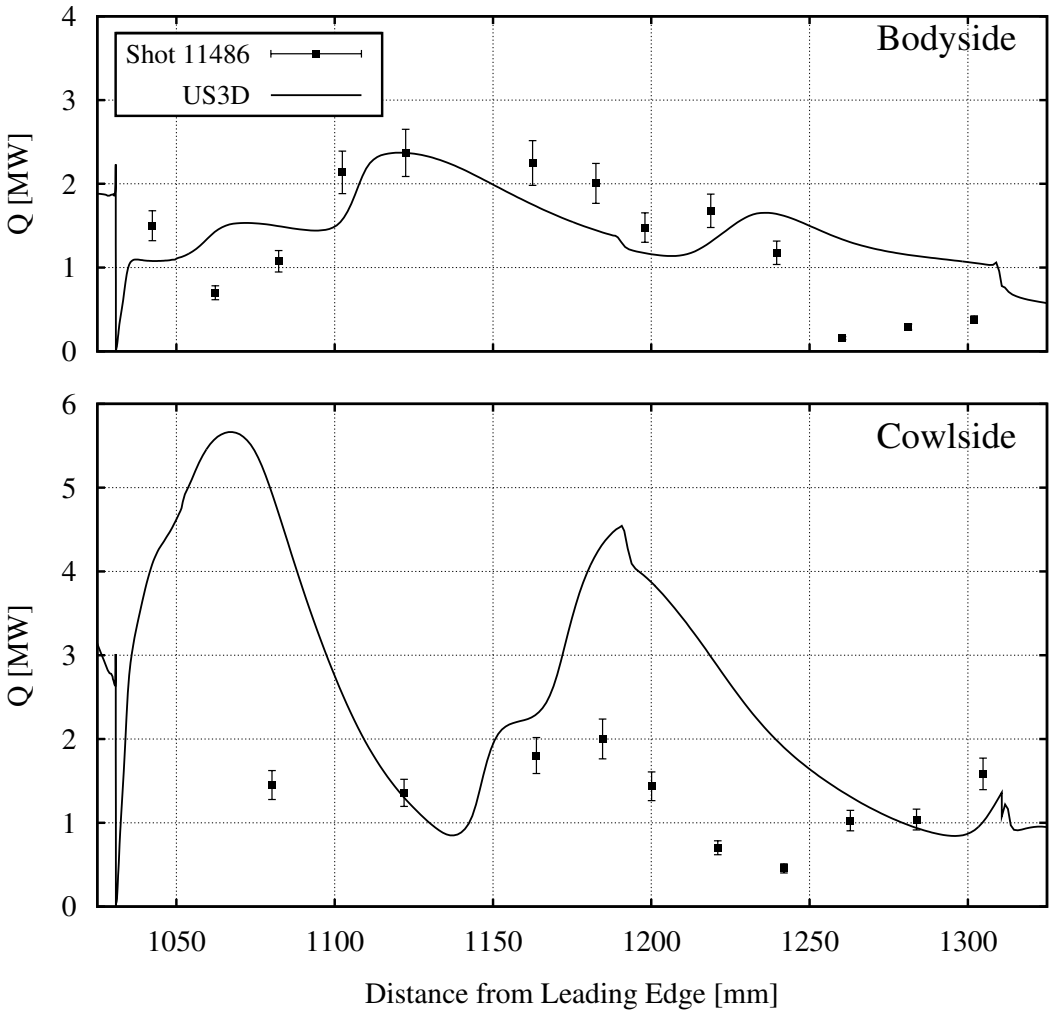

Figure 5.5: Comparison of simulated and experimental unfuelled engine heat flux distributions

The bodyside wall heat flux distribution shows reasonably good agreement to the

experiment across the constant-area combustor, and much of the diverging combustor. Both the magnitude and trend of the wall heat flux appears to be captured accurately along the bodyside of the combustor. The cowlside wall heat flux distribution, on the other hand, appears to show a rather strong lack of agreement. The simulation greatly over-predicts the magnitude of the heat flux into the wall, though it captures the general shape of the distribution throughout the diverging combustor. At the end of the diverging combustor, the heat transfer on the cowlside seems to under-predict the wall heat flux slightly, though given there being a disparity in the pressure measurements due to the strength of the local shock reflection, this is not surprising. Some of the differences observed may be due to the simulation wall temperature being held to 300 K, whereas the actual surface temperatures of the heat transfer gauges was measured in experiment to increase by as much as 60 K by the end of the test time. This temperature difference is accounted for in the experimental measurement uncertainty represented by the error bars, and so still does not account for why the bodyside heat flux seems to agree well, whereas the cowlside simulated heat flux appears to be greatly over-predicted.

There is, however, one variable in the simulation whose effects could have been substantial: the non-dimensional height (y^+) of the wall-adjacent cells in the simulation. The distribution of y^+ is plotted in Figure 5.6.

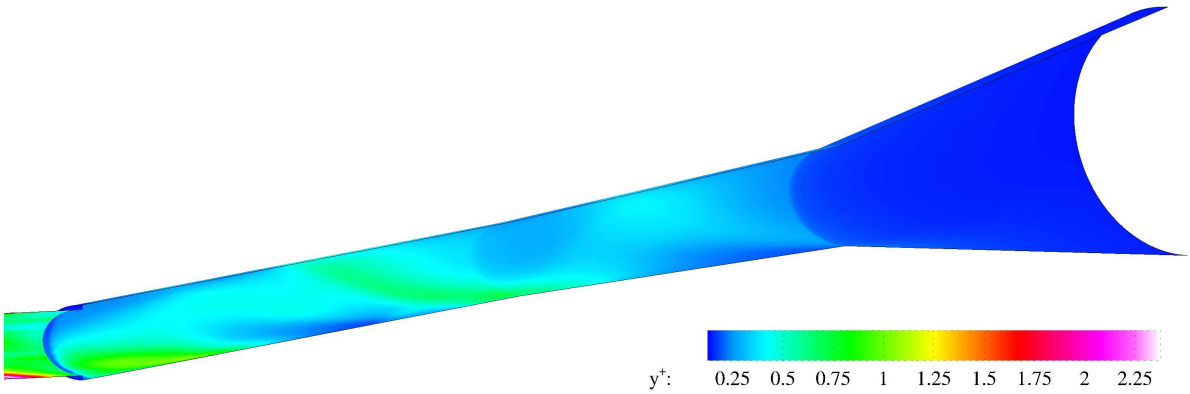

Figure 5.6: Distribution of y^+ through the M12 REST combustor and nozzle

The value of y^+ in the combustor is generally less than one, and never greater than 1.2. This wall- adjacent cell height is considered to be sufficient for a wall-resolved flow simulation (that is, a CFD solution that does not rely on wall functions to capture gradients across the viscous sublayer). As such, while resolving the value of y^+ along the wall to below 0.3 (as it typically is along the bodyside of the combustor) may improve resolution of the near-wall temperature gradient, it's unlikely to completely

resolve the lack of agreement along the cowlside of the combustor. Another factor must therefore be recognized to explain the incongruence between the the simulation and experiments.

It has been noted in literature that RANS models tend to have difficulty in accurately reproducing wall heat flux observed in scramjet experiments (see, for example, Chapuis et al. (2013)). The ability of a RANS simulation to reproduce wall heat flux has been shown to depend heavily on the choice of both turbulent Schmidt and Prandtl numbers (Pečnik et al., 2012), which in turn suggests a strong dependence on the turbulence model used. While the smaller values of y^+ may in some cases mitigate the inaccuracies in wall heat flux caused by the turbulence model, changing the values of Sc_t and Pr_t, or using another turbulence model may also be necessary to improve agreement between the simulation and experiment. Alternatively, the implementation of dynamically calculated turbulent Schmidt and Prandtl numbers, such as the model developed by Xiao et al. (2007), may be more appropriate.

Overall, taking the pressure and heat flux contours together, it can be concluded that the majority of the flow physics shown in the simulated contour plots are a reasonably accurate representation of what occurs during an unfuelled engine experiment. While the heat flux magnitudes are inaccurate along the cowlside combustor wall, the bodyside heat flux and the pressure distributions on both sides of the engine show good agreement with the experimental data. With this encouraging result for the validity of the unfuelled simulations, the effects of inlet fuelling on the combustor flow field and the engine's performance are examined next.

5.2 Combustor Flow Field with Inlet Injection

Having examined the unfuelled flow field, the effects of inlet fuel injection on the combustor flow field are examined. A robust inlet injection and combustion case has been chosen, Shot 11464 of Wise (2014). The simulation conducted for the full engine matches the inflow and inlet fuel mass flow conditions of the experiment, which can be found in Appendix E.

Much like the unfuelled simulation in the previous section, this study was carried out using the engine mesh presented in Chapter 3. The simulation was run for 70,000 iterations, achieving 6 orders of magnitude convergence in the RMS residual and 5 orders of convergence in the energy residual. The mass flow balance across all inflow and outflow boundaries was 1.67×10^{-7} kg/s, and the momentum balance in the x-, y-, and z-directions were all below 1×10^{-3} N.

The final equivalence ratio for this simulation is based on the mass capture of the engine, which is 54.201 g/s according to the simulation. This inlet mass capture is 93.7% of the mass capture calculated for the unfuelled simulation. This difference in inlet mass capture is believed to be caused by the difference in inflow conditions between each simulation. It is assumed that the oxygen available for combustion in the captured flow is equal to the sum of the free molecular and atomic oxygen mass fractions of the inflow condition. With an inlet fuel mass flow rate of 0.70663 g/s and the mass capture given above, the equivalence ratio is calculated to be 0.55. Despite the lower inlet mass capture, and nearly 1.66 times more fuel than was injected in the inlet simulations examined in Chapter 4, the general form of the inlet flow field remains unchanged, and as such, will not be re-examined in detail in this section.

The examination of the effects of inlet injection on the combustor flow field begins with an investigation into how the combustor shock structure is affected by the presence of reacting flow along the bodyside wall of the engine. The resulting shock contours are presented in Figure 5.7.

The presence of inlet injection has a strong effect on the shock structures inside the isolator, and within the combustor and nozzle sections. The first major difference between the unfuelled shock structure in Figure 5.7a, and the inlet-fuelled shock structure in Figure 5.7b, is the movement of the second bodyside reflection of the cowl closure shock (B2) upstream into the isolator. This is likely due to the increased speed of sound along the bodyside wall of the inlet and isolator due to the injection of hydrogen fuel, and its ignition at the inlet throat, as demonstrated in Chapter 4. The shock angle was also likely changed by the merging of the injector bow shocks with the cowl closure shock, which produces a single strong shock that propagates at a higher angle. Both the presence combustion and the upstream shock reflection significantly increase the pressure of the flow entering the combustor.

The elliptical recompression shock, visible originating at points BR and CR in Figure 5.7b as well as in the first slices of the constant area combustor in Figure 5.7c, propagates far more quickly across the bodyside of the combustor. This in turn means that the elliptical recompression shock converges far earlier in the constant area combustor than before, and does so closer to the cowlside wall. This convergence also occurs just downstream of the cowlside reflection of the inlet-transmitted shock (C2). The reflection of the inlet-transmitted shock appears to merge with the bodyside-travelling component of the post-convergence elliptical shock, and reflects off the bodyside wall of the combustor at B3.

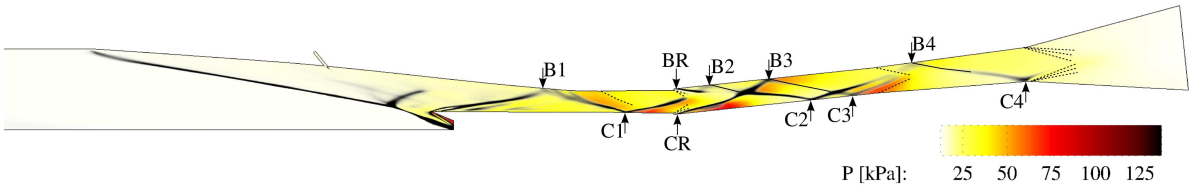

(a) Unfuelled Symmetry Plane Structure

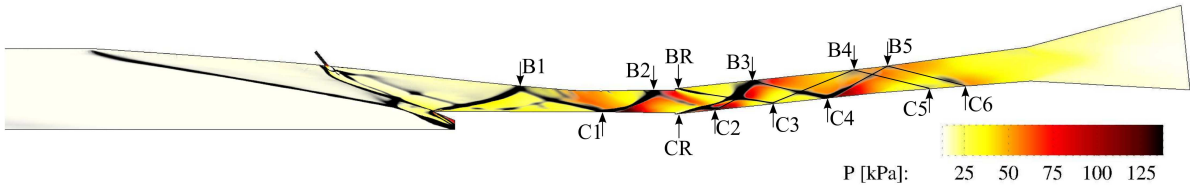

(b) Inlet-fuelled Symmetry Plane Structure

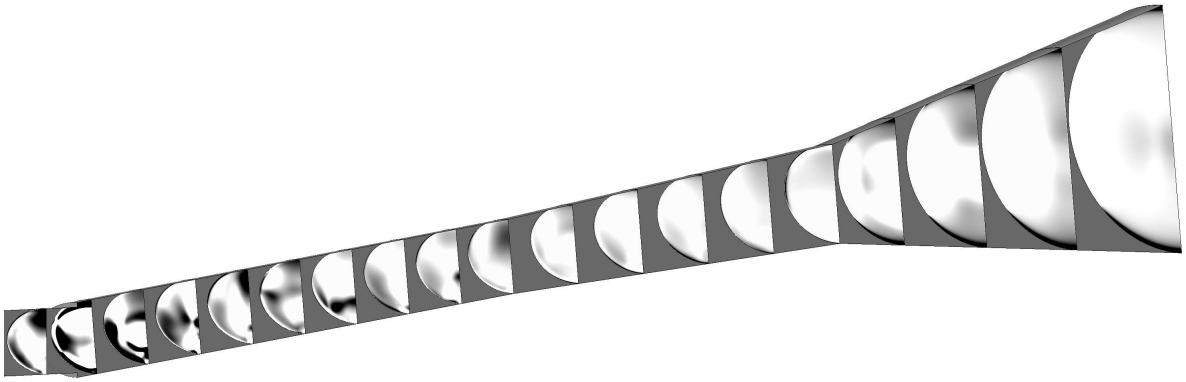

(c) Cross-plane Structure

Figure 5.7: Comparison of Inlet-fuelled and Unfuelled Shock Structures

Unlike C2 and C3 in the unfuelled case, the pair of cowlside reflections C3 and C4 in Figure 5.7b have a considerable distance between them. This is due to the convergence of the elliptical recompression shock occurring far closer to the cowlside wall: the cowlside-travelling component of the post-convergence elliptical shock is able to reach the cowlside wall far more quickly. This has an additional effect: the shock that reflects at C3 is able to reach the bodyside wall (B4) before the end of the constant area combustor, while the reflected shock from C4 does not reach the bodyside wall (B5) until after it enters the diverging combustor. Unlike the unfuelled case, the two shock trains do not merge upstream of the nozzle; both reflect off the cowlside wall of the diverging combustor at C5 and C6. Both shock trains appear to weaken as the flow travels downstream through the diverging combustor, and aren't readily visible in the contours of the nozzle. This weakening of the shocks is caused by the expansion of the flow, and possibly magnified by the presence of combustion driving up temperatures inside the combustor.

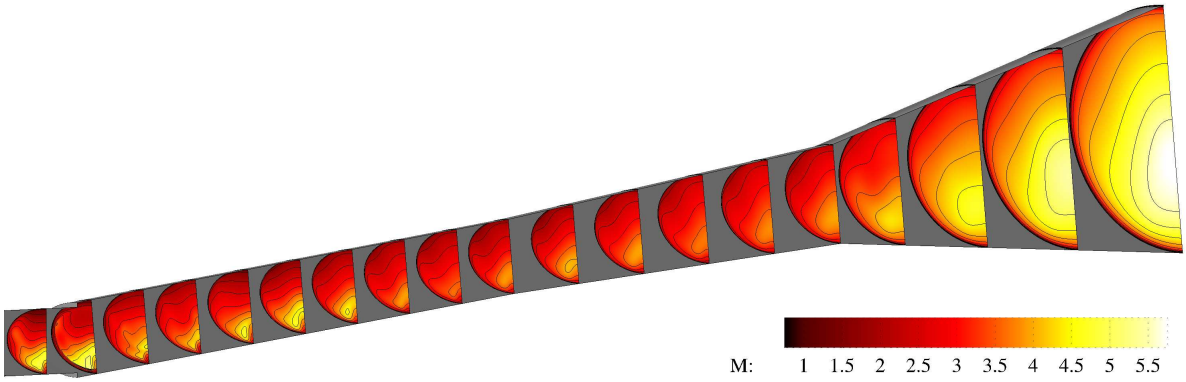

(a) Mach number

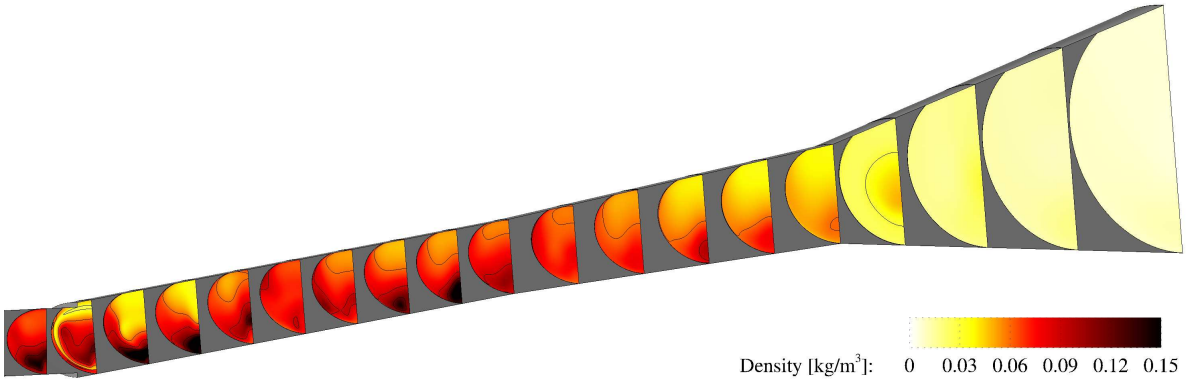

(b) Density

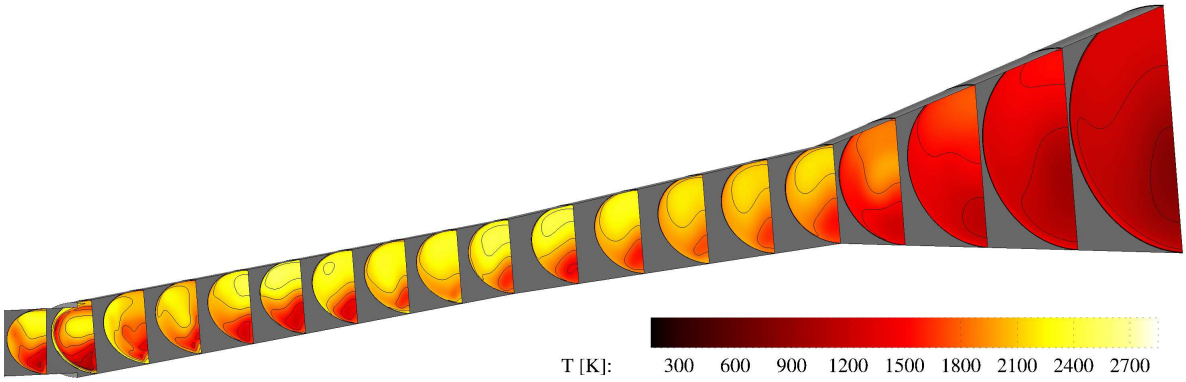

(c) Temperature

Figure 5.8: Mach number, Density, and Temperature Contours in the combustor with inlet fuelling

The effects of the shock structure presented, as well as an explanation of the major effects that inlet fuel injection has on the combustor flow field can be found by examining Figure 5.8. The existing Mach number mismatch between the bodyside and cowlside of the engine continues into the combustor. With the addition of the now-burning inlet-injected fuel in the thick bodyside boundary layer, the Mach number contours shown in Figure 5.8a show that this mismatch is further intensified: the core of the cowlside flow remains at Mach numbers between 3.5 and 5 throughout the length of the combustor, while the bodyside flow is typically moving at Mach 2 through the length of the combustor. This represents a significant reduction in the mean combustor flow Mach number from the unfuelled case, and is due to both the increased flow pressure entering the combustor, as well as temperature and pressure increases caused by combustion of the fuel in the combustor itself.

There is also a significant slowing of the cowlside flow in the diverging combustor section relative to the unfuelled combustor. This coincides with the region downstream of the second strong cowlside shock reflection (C4) near the end of the constant-area combustor, which is visible in the density contours of Figure 5.8b. The density field remains remarkably uniform throughout much of the combustor, with all observed fluctuations remaining strongly coupled to the shock- and expansion-wave structures in the constant-area and diverging combustor sections. By the time the flow reaches the nozzle, much of the density anisotropy has been smoothed out by the expansion of the gas and the subsequent weakening of the reflected shocks. In contrast with the Mach number of the flow, the bias of the density flow field toward the cowlside of the engine as the flow is exhausted out of the nozzle is significantly smaller.

The temperature flow fields shown in Figure 5.8c suggest that in the first half of the constant-area combustor, the bodyside flow actually cools as the flow is deflected and expanded by the combustor step rarefaction fan. Downstream, the temperature along the cowlside is raised by the transmitted inlet shock and recompression shocks both impinging on the cowlside surface (C3 and C4). The temperatures of both the bodyside and cowlside flows recover downstream of these reflections. In the case of the bodyside flow, the temperature is universally higher than it is in the unfuelled engine, due to the heat release from combustion of the inlet-injected hydrogen. The bodyside temperature also increases a great deal in the second half of the constant-area combustor, indicating not only the reflection of the combined inlet/recompression shock at B3, but possibly also an increase in combustion heat release downstream of the shock due to SBLI-enhanced mixing of the bodyside flow.

5.2.1 Flow Chemistry and Combustion

The flow remains divided into a high-temperature bodyside flow region and a much cooler cowlside flow region, even through the nozzle expansion at the back of the engine. This temperature distribution, and indeed many of the observed combustor flow features, are driven by the combustion of the inlet-injected fuel. This can be confirmed by examining the species mass fraction contours presented in Figure 5.9.

The distribution of OH throughout the combustor matches closely not only with the temperature contours, but with the density contours as well. A region of low OH production occurs in the first third of the constant-area combustor, as the low-density/low-pressure region downstream of the combustor-step expansion locally slows chemical reactions. Much like the temperature contours, OH production peaks in the vicinity of both bodyside shock reflections (B3 and B4) in the constant area combustor, reinforcing the fact that the temperature rise observed in these regions is driven by both the shock reflection and the heat released by burning hydrogen.

Combined with the slow diffusion of OH across the combustor cross-section, these peaks in OH production confirm that combustion of the inlet-injected fuel is mixing-limited in the combustor. The shocks act to locally improve mixing of the fuel with remaining oxygen, and the local increases in pressure and temperature caused by the shock, allow reactions to proceed at a rapid pace. Meanwhile, the OH that migrates toward the cowlside wall appears to react slowly, if at all. The unreacted oxygen (Figure 5.9b) and water vapour (Figure 5.9c) contours appear to reinforce this idea: neither shows any behaviour other than that consistent than the slow mixing of each species across the engine, and the reaction of larger intermediate species along the bodyside of the combustor.

Heat release due to combustion slowly raises the temperature across the entire combustor with downstream distance. The cowlside flow is processed by the second strong cowlside shock reflection just as it enters the diverging combustor section. The expanding cowlside flow is unable to cool as much as before, as heat release from combustion is convected and/or conducted across the combustor. Though the cowlside temperature remains low relative to the bodyside, this overall warming due to combustion explains the drop in local Mach number observed in the flow along the cowlside of the engine.

The cooler temperatures along the cowlside wall inhibit the reaction of incomplete combustion radicals diffusing from the bodyside of the combustor with available oxygen along the cowlside of the engine. The eventual disappearance of any significant

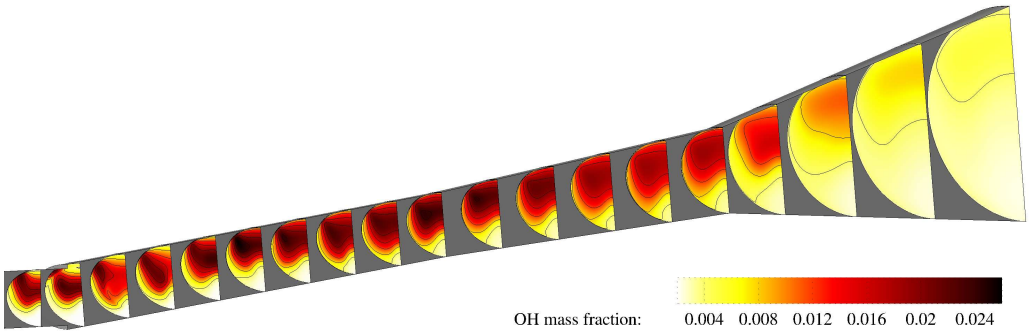

(a) OH radical distribution

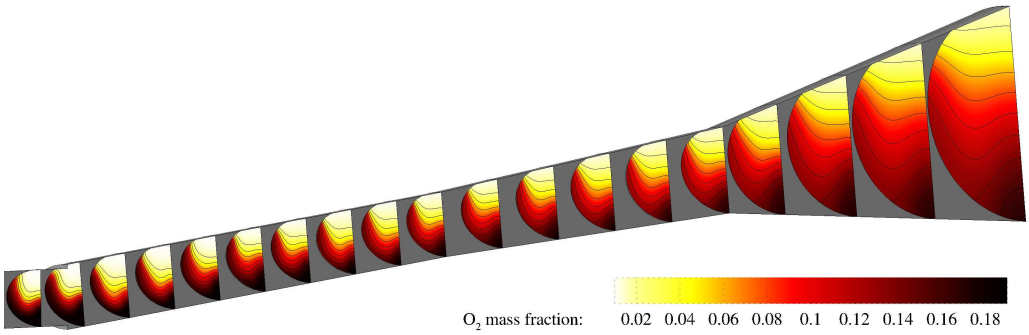

(b) Free Oxygen distribution

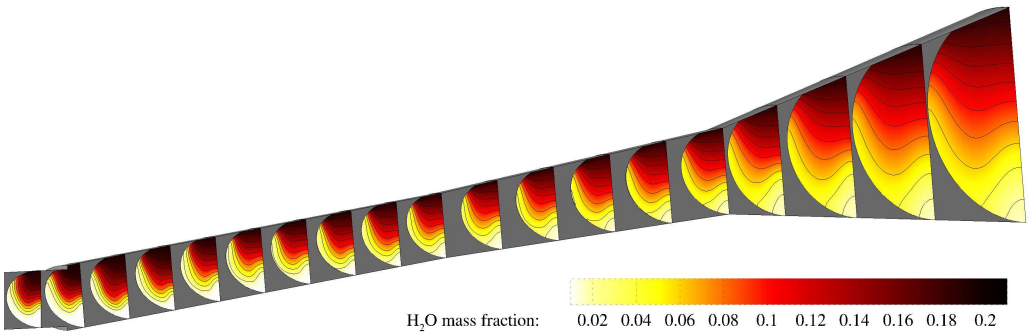

(c) Water vapour distribution

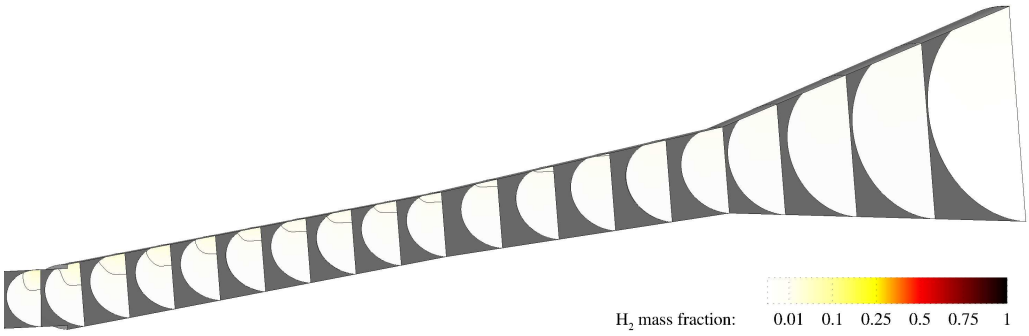

(d) Unconsumed Hydrogen distribution

Figure 5.9: Species mass fraction contours in the M12 REST combustor with inlet fuelling

mass fraction of molecular hydrogen in Figure 5.9d, which remains distributed near the bodyside wall of the combustor, suggests that while there is still hydrogen and intermediate radicals diffusing toward the cowlside wall, low temperatures in this region prevent their reaction with OH to form water vapour, driving down combustion efficiency. The OH radicals instead combine to form H_2O_2 as the flow passes through the nozzle.

5.2.2 Engine Performance with Inlet Injection

The mass fraction contours presented previously indicate both a strong level of water production in the engine and a disappearance of molecular hydrogen. However, they also show a significant amount of free OH radicals exiting the engine nozzle. Based on this behaviour, before examining combustion-induced changes in the aerodynamic forces acting on the engine surfaces, the progression of combustion over the length of the engine will first be investigated.

Just how well the engine burns this inlet-injected hydrogen is examined by plotting how the mixing and combustion efficiencies evolve with streamwise distance in the engine. As this simulation is for hydrogen fuel injection at an equivalence ratio less than unity, the mixing and combustion efficiencies defined in the previous chapter are employed. The resulting mixing and combustion efficiencies of the inlet-injected fuel across the entire region downstream of injection are plotted in Figure 5.10.

The behaviour of the efficiency curves in the inlet are very similar to those of the the lower-equivalence ratio case presented in Chapter 4: the injected fuel mixed into the available captured air at a nearly-linear rate in the first 150 mm (75D) downstream from injection. In this same region, the fuel reacts only slowly, until a sudden rise in combustion efficiency occurs downstream of the bodyside reflection of the cowl closure shock (B1), at a point 925 mm from leading edge.

The mixing and combustion efficiencies both grow at a slower rate as the flow enters the isolator. Inflection points can be seen in the combustion efficiency curve at the start of the combustor, and when the flow enters the diverging combustor section and the nozzle. These inflections occur due to the reaction rates slowing temporarily as the flow is turned around expansion corners in the flow. The reaction rates quickly recover after each inflection, giving this $\phi = 0.55$ fuelled engine a final mixing efficiency of 96.7% at the exit of the nozzle, and a global combustion efficiency of 77.92%.

This agrees well with the findings of Wise (2014), who determined combustion efficiency in the experiment by performing a quasi-one-dimensional cycle analysis of

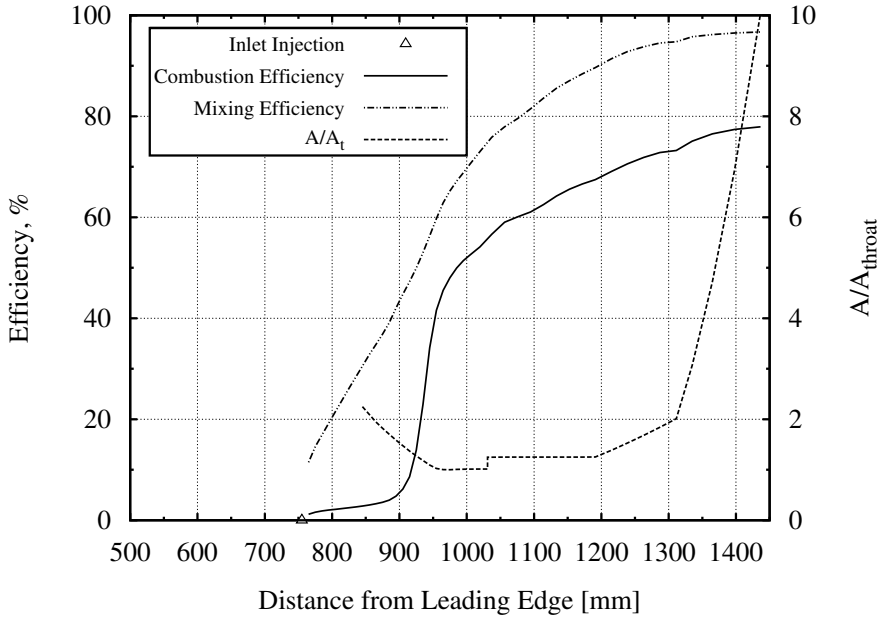

Figure 5.10: Fuel-based mixing and combustion efficiencies for the inlet-fuelled M12 REST engine at an equivalence ratio of 0.55

the engine, based on the average of the experimental bodyside and cowlside pressure distributions. By fitting a one-dimensional pressure distribution to the nozzle and diverging combustor pressures, the experimental combustion efficiency was estimated to be between 70 and 75%. The CFD combustion efficiency prediction also agrees well with the off-design operation performance of Moule et al. (2014), who found that the combustion efficiency of inlet-injected fuel in the M12 REST engine operating at Mach 9 flight conditions was 78%.

It is also worth comparing the combustion and mixing efficiencies at the inlet throat to the efficiencies found at the lower equivalence ratio studied in Chapter 4. The $\phi = 0.33$ inlet injection case shows mixing and combustion efficiencies at the throat of 82.4% and 52.64%, respectively. These are both higher than the mixing and combustion efficiencies at the throat of the $\phi = 0.55$ inlet injection case considered here, which are 65.2% and 48.1%, respectively. It is likely that the final combustion efficiency of the fuel injected on the inlet examined in the previous chapter would also be higher if that simulation were carried on through the entire engine flow path.

Based on this decrease in combustion efficiency at the throat, the argument may be made that $\phi = 0.33$ should be considered an upper bound for inlet injection. While a greater amount of fuel is burned in the higher inlet equivalence ratio case (by mass), when inlet injection is combined with mainstream injection elsewhere in the engine,

the goal should be to completely burn as much fuel as possible. If a smaller fraction of the inlet-injected fuel burns, there would be a greater benefit to shifting more fuel into mainstream injection downstream. The lower inlet mass fraction will still provide heat release and hydrogen radicals to pilot ignition in the combustor, while its higher combustion efficiency means less of the inlet-injected fuel would leave the engine unburned.

Aerodynamic Forces acting on the internal flow path

As with the unfuelled case, the thrust and drag forces of this inlet injection case are calculated using US3D's post-processing abilities, and broken down into subsystem contributions using Tecplot. The individual inviscid and viscous component forces are normalized against the net thrust and drag forces, while the combined component forces are expressed as a positive (thrust) or negative (drag) percentage of the total force acting on the engine. The results of this force accounting are presented in Table 5.2.

Table 5.2: Aerodynamic thrust (+) and drag (-) forces acting on the inlet-fuelled M12 REST

Engine Component	Inviscid Forces [N]	%	Viscous Forces [N]	%	Combined Forces [N]	%
Inlet Internal Surfaces	-11.863	24.06	-12.053	24.44	-23.916	-81.48
Forebody	-4.946	10.03	-3.089	6.26	-8.035	-27.37
Isolator and Combustor Step	0.261	1.31	-1.489	3.02	-1.228	-4.18
Constant Area Combustor	0.210	1.05	-7.716	15.65	-7.506	-25.57
Diverging Combustor	6.330	31.71	-5.185	10.51	1.145	3.90
Nozzle	13.163	65.93	-2.975	6.03	10.188	34.71
Total Drag Forces	-16.809	34.08	-32.507	65.92	-49.316 N	
Total Thrust Forces	19.964	100.0	-	-	19.964 N	
Total Internal Force	8.101 N		-29.418 N		-21.317 N	
Total Force	3.155 N		-32.507 N		-29.352 N	

As is the case for the unfuelled engine, more than half the total drag forces acting on the flow path is produced by viscous effects, with the total skin friction drag on the engine increasing 1.94% to account for 65.92% of the total drag. Most of this additional skin friction drag increase occurs on the internal inlet and isolator surfaces, where temperature and pressure increases due to combustion have an effect on the bodyside boundary layer. The internal net (drag) force acting on the engine decreases, primarily due to 54% increase in thrust forces generated by the nozzle and combustor from the combustion of hydrogen. While the engine is nowhere near producing net thrust, it

was not expected to do so under lean fuelling conditions in tunnel operating conditions (where a great deal of heat is lost to the engine walls) and without an ideal thrust nozzle design.

It is difficult to compare the increase in inlet drag to the unfuelled case, as each simulation uses a slightly different inflow condition. In order to estimate the internal drag increase on the inlet, the combined internal inlet forces are normalized by the combined forebody forces in each case. Comparing the unfuelled and inlet fuelled normalized forces leads to an estimated inlet drag increase of 7.4% when fuel is injected on the inlet at this equivalence ratio. This drag increase is greater than the 4.8% increase calculated for the $\phi = 0.33$ inlet injection case studied in Chapter 4. This gives further indication that there is an optimal inlet fuel injection equivalence ratio, and it is likely less than 0.33.

The isolator, combustor and nozzle sections' skin friction contributions are slightly higher than the unfuelled case, rising to 35.21% of the total engine drag. The constant area combustor remains the largest contributor to drag downstream of the inlet, accounting for 15.65% of the net internal drag. This increase in the combustor's share of the drag generation is actually relatively modest, considering the extra energy in the combustor flow due to combustion. Why the rise in skin friction is smaller than expected could be due to flow energy being lost by thermal conduction into the cold engine walls in every experiment and simulation. It can also be explained in part by examining the differences between the inlet fuelled and unfuelled wall shear stress distributions, as plotted in Figure 5.11.

It is apparent that the presence of inlet fuelling has a strong effect of the distribution of shear (viscous) stress in the engine. While the unfuelled shear is relatively evenly distributed around the cowlside wall and well up the sides of the combustor (Figure 5.11a), it is far more concentrated along the cowlside wall when inlet injection is employed (Figure 5.11b). While the magnitude of the wall shear stress acting along the cowlside surface is stronger than in the unfuelled case, the shear is far weaker along the bodyside half of the combustor. It may well be that the slower, combusting boundary layer moving along the bodyside surface provides some local reductions to shear stress, which are counterbalanced by the strong increases seen along the cowlside walls. If viscous drag control were implemented along the cowlside walls in the combustor sections, significant benefits to engine performance would likely be realized.

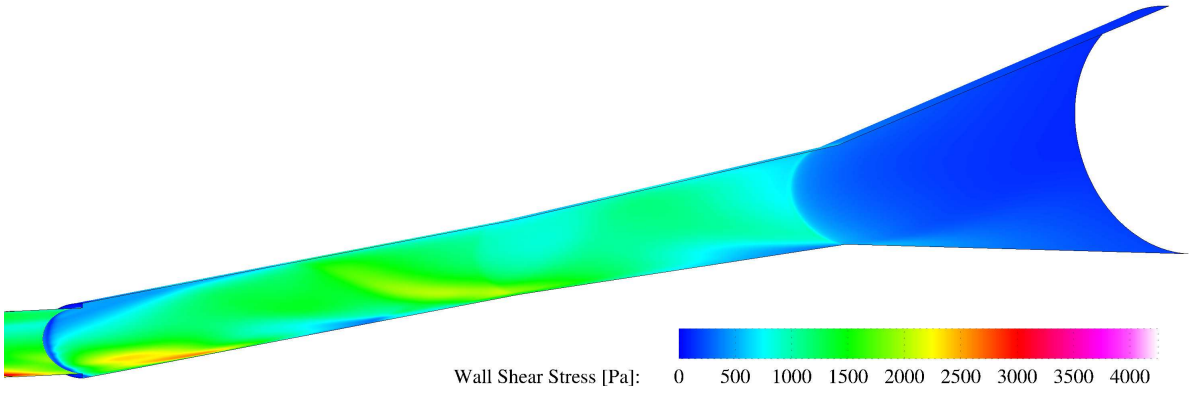

(a) Unfuelled

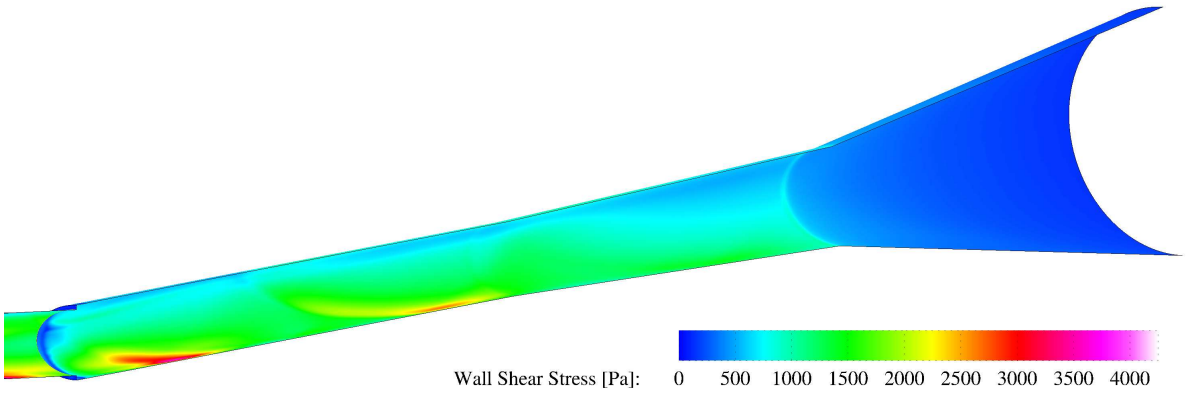

(b) Inlet fuelled, $\phi = 0.55$

Figure 5.11: Comparison of the unfuelled and inlet-fuelled M12 REST combustor wall shear stress distributions

5.2.3 Comparison between Experimental Data and Simulation

The bodyside and cowlside pressure distributions from the simulation are compared to the pressure distribution measured in Shot 11464. This makes it possible to assess the quality of the simulation, and check how well it represents the physics of the experiment it is intended to replicate. This comparison is shown in Figure 5.12.

There is very little variation between the CFD simulation and the experiment along the inlet bodyside surface, and in the diverging section of the combustor. In the constant area combustor, there appears to be a shift in the shock reflection locations in the simulation relative to the experiment, though the post-shock reflection pressure rises are generally captured well. One exception to this appears to be the shock reflection corresponding to point C3, which induces a far stronger pressure rise in the experiment than it does in the simulation. The shift in shock locations can be explained by a number of factors: the heat release in the simulation may happen earlier than it does in the

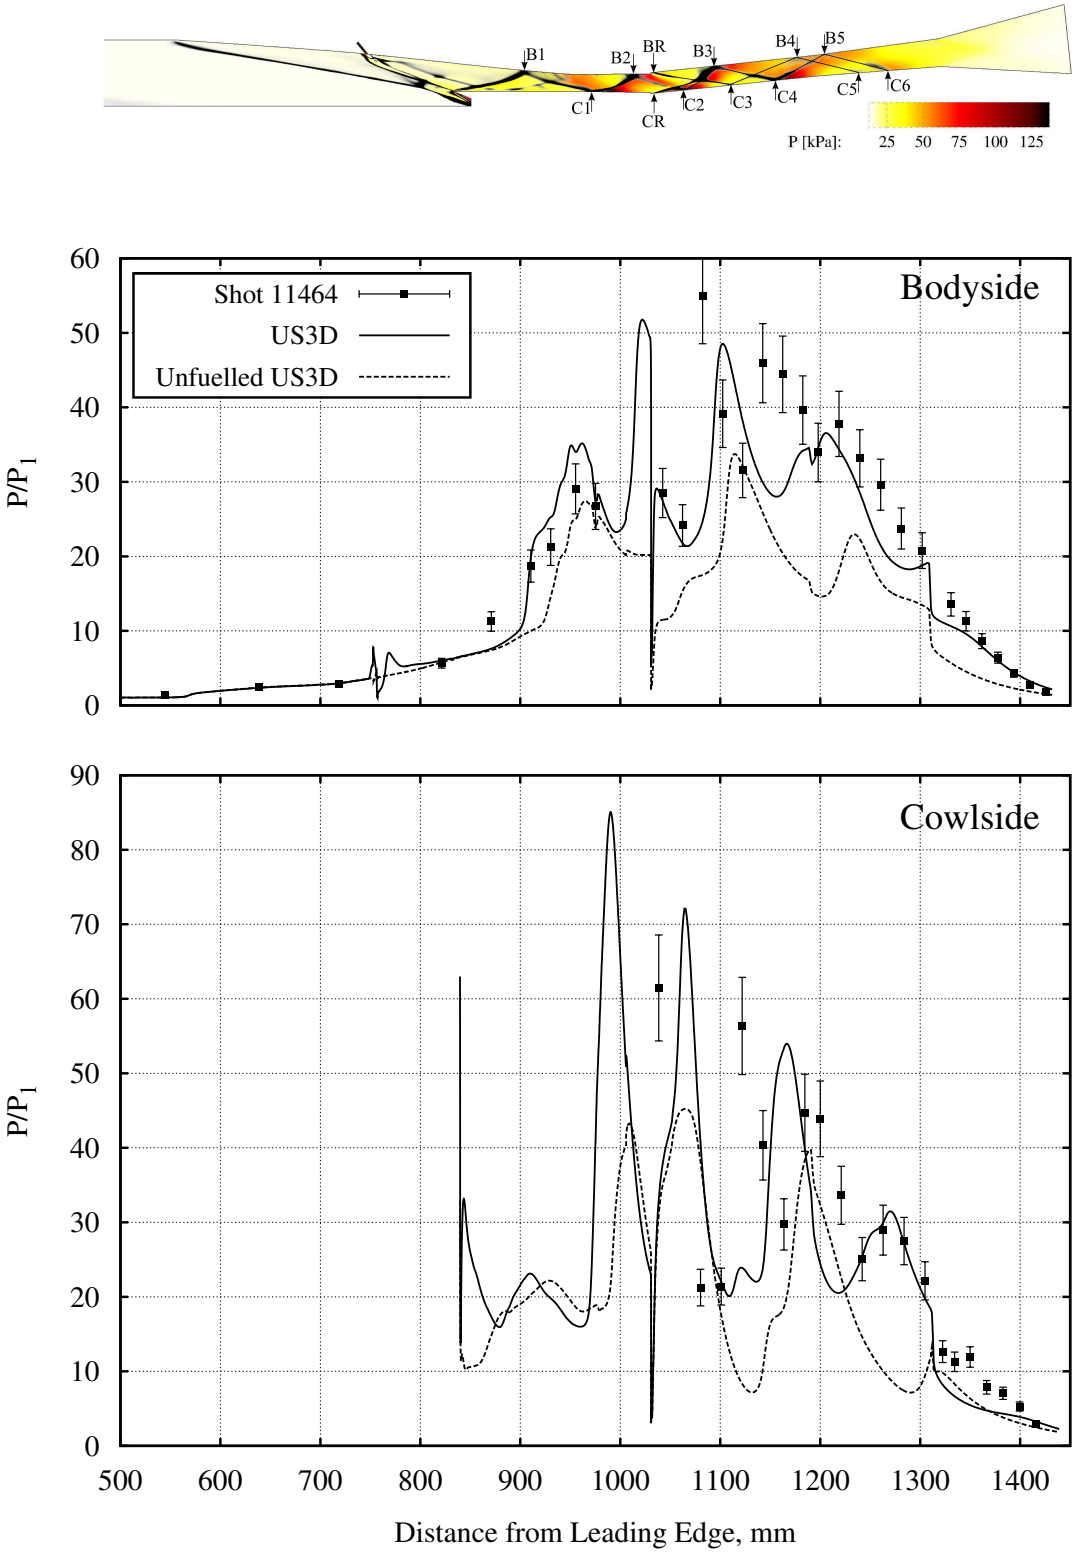

Figure 5.12: Comparison of the normalized experimental and CFD pressure distributions for the inlet-fuelled M12 REST

experiment, the bulk temperature entering the combustor may be somewhat different, and there may be transient effects in the pressure data that could not be resolved by the normalization procedure used.

The bodyside nozzle pressure distribution appears to have been captured accurately by the simulation, but there is some difference between simulation and experiment on the cowlside of the nozzle. The behaviour of the experimental cowlside nozzle pressure is likely caused by an asymmetric expansion effect, brought about by a bodyside shock reflection just upstream of the nozzle entrance. As this reflected shock passes back through the rarefaction fans it weakens slightly, but remains strong enough to compress flow along the first 40 mm of the cowlside nozzle surface. This shock is far weaker in the CFD (though some suggestion of its bodyside reflection can be seen in the sudden pressure rise just prior to the bodyside nozzle expansion), and is therefore not seen in the simulated pressure field.

Overall, the simulation appears to match the experimental pressure distributions. While no combustor wall heat fluxes are available for comparison, the pressure distributions suggest that the RANS simulation captures the time-averaged combustion process. Similarly, the pressures on the inlet and in the nozzle both suggest that the simulated thrust levels are a reasonably accurate reproduction of the thrust generated in the experiment.

5.3 Combustor Flow Field with Baseline Combined Fuel Injection

With a better understanding of inlet injection and its effects on the flow field of the M12 REST combustor and nozzle sections, the "baseline" configuration of the M12 REST internal flow path is the next to be examined. In this configuration, inlet fuel injection is combined with fuel injection from the face of the combustor step into the boundary layer of the constant area combustor. The nominal ratio of fuel injected in the inlet and the combustor is 3:7.

Injection into the boundary layer is employed in the baseline configuration in an attempt to actively reduce the skin friction drag generated by the combustor sections. As presented in Section 2.4.3 of the literature review, the injection and combustion of fuel inside a turbulent compressible boundary layer has been demonstrated to provide local skin friction reductions of greater than 70%. While in the baseline configuration

the fuel is injected through a set of discrete portholes, rather than the single continuous slot geometry used to demonstrate these drag reductions, it has also been shown that portholes spaced sufficiently close together should have their fuel jets form into a continuous sheet of fuel across the wall surfaces downstream (Barth et al., 2012). If even a fraction of the potential skin friction reduction mentioned above is achieved using boundary layer injection in the M12 REST combustor, it could provide a 5-10% improvement in the uninstalled engine thrust.

While 70% of the fuel is injected from the step, not all of it is necessarily expected to remain in the boundary layer. Due to the design constraints placed on the annular fuel plenum by the change in engine direction at the combustor step, the injectors are set at an angle of 10° away from the local isolator wall. This in turn means that the injectors are angled anywhere between 4° (on the cowlside) and 16° (on the bodyside) away from the downstream combustor wall. Having some component of velocity normal to the wall means that fuel can more easily escape the boundary layer and mix into the mainstream flow. The question is therefore whether enough of the fuel injected from the combustor step will remain in the boundary layer to allow sufficient boundary layer combustion to achieve a drag reduction.

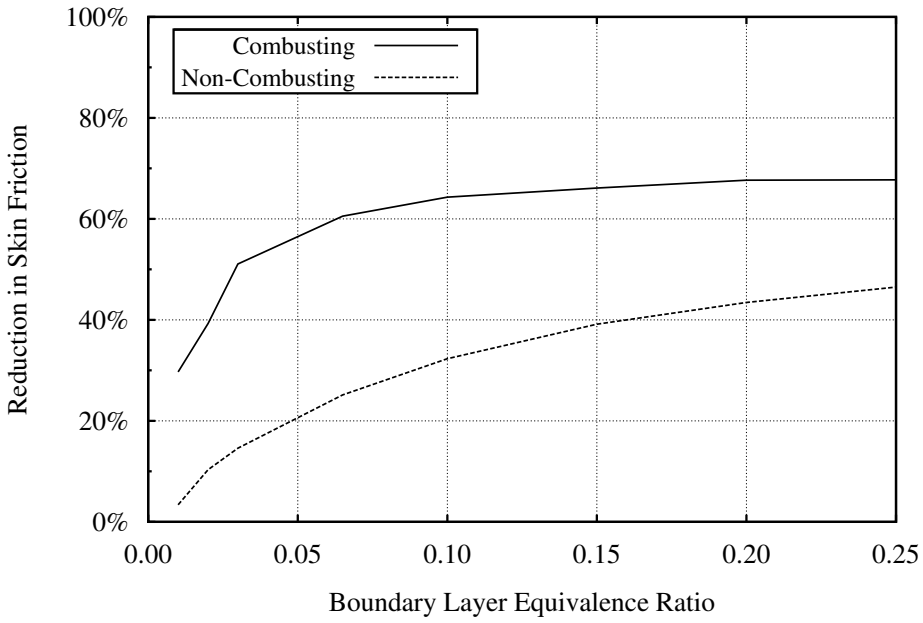

Figure 5.13: The effects of hydrogen injection and combustion in the boundary layer of a scramjet combustor. Reproduced from Barth et al. (2013)

To answer the question of how much fuel will need to stay inside the combustor boundary layer for drag reduction in a scramjet combustor to occur, an extension to

Stalker's (2005) analytical model is employed. As illustrated in Figure 5.13, for a scramjet combustor in which the mainstream flow enters the combustor with an average Mach number of 4, it was found that above an equivalence ratio of 0.15, there is little additional benefit to injecting more fuel into the boundary layer (Barth et al., 2013). (More on this analytical model and its treatment of boundary layer combustion and skin friction reduction can be found in Appendix G.) It is therefore assumed that so long as 20% of the fuel injection from the combustor step remains in the constant-area combustor's boundary layer, the engine should experience a substantial net reduction in combustor skin friction.

5.3.1 Simulation Details

In order to simulate the baseline configuration, the combustor mesh is be modified to accommodate half of the 36 portholes that are spaced evenly around the circumference of the backward step in the full engine. This entails 17 whole injector tubes, plus the two half tubes of the injectors that are placed on the engine centerplane. The resulting engine mesh contains 41,928,944 cells, with the grid structure in the vicinity of the combustor step having changed to the form shown in Figure 5.14.

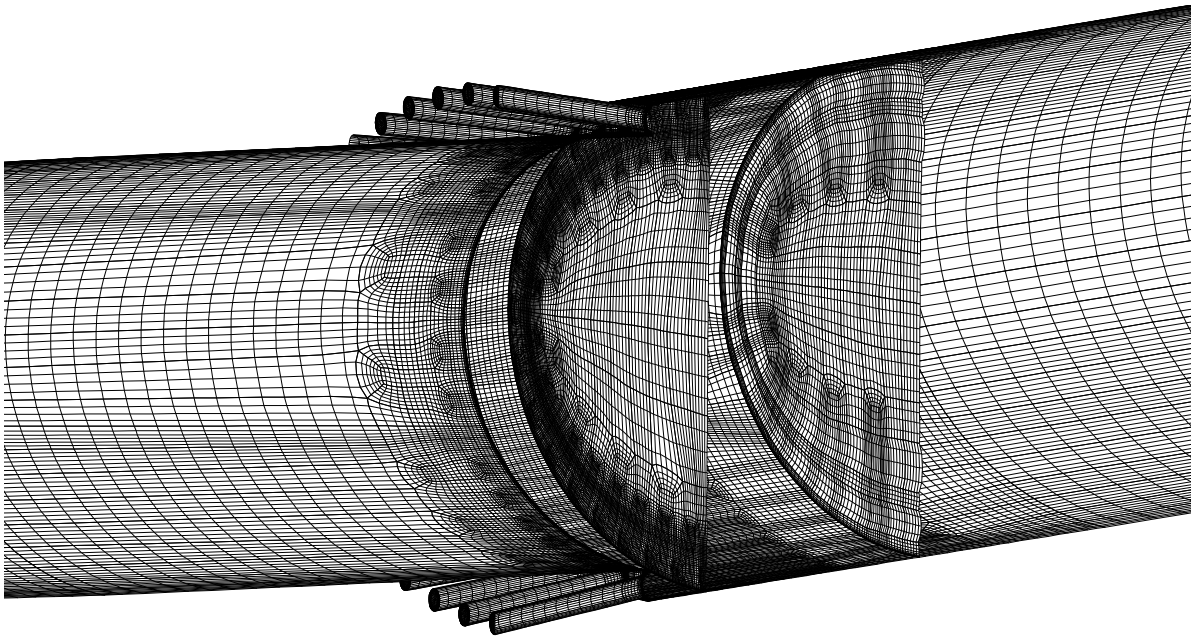

Figure 5.14: The computational mesh used to simulate injection and near-field mixing behind the combustor step

The resulting mesh just before and just behind the step is quite dense in order to accurately capture the expansion and interaction of the individual fuel jets with the flow

coming over the step, and with each other. Ten jet diameters (10D) downstream, this high-density mesh is then successively coarsened over the next 30D downstream, as it has been shown previously that the primary interactions between jet shock structures are typically finished by this point. Unlike the combustor mesh without any injection, the first cell height along the combustor surface is constrained to be 0.7 μm everywhere in the combustor. This has been done to account for any effects that injection from the step might have that could lead to thinner than normal boundary layers, earlier in the combustor.

The inflow to the simulation has been chosen to match the nozzle exit flow calculated from the tunnel conditions of experimental Shot 11471 from Wise (2014). The details of this nozzle exit/inflow condition, and its equivalent flight condition, can be found in Appendix E. Fuel injection is set such to match the fuel mass flow rates of the experiment. This leads to a total simulated hydrogen mass flow rate of 1.85764 g/s, with a near-nominal 31:69 mass flow split between the inlet and combustor step injectors.

The solution converged after being run for a total of 73,000 iterations, for a total in-simulation flow time of 1.499 ms. The RMS residual dropped by 6 orders of magnitude, as did the energy residual. The momentum balance was less than 10^{-4} N in each of the simulation coordinate axes, and the solution mass flow balance was 3.6 x 10^{-8} kg/s. The mass capture of the engine is 59.35 g/s, which is 2.6% higher than the unfuelled engine mass capture presented earlier in Section 5.1. Given that the mass capture increases, the difference in mass capture is almost certainly due to the difference in inflow conditions. As such, much like the ϕ = 0.33 case presented in Chapter 4, the inlet injection has no effect on inlet mass capture. Based on the mass capture of the inlet, the mass fractions of atomic and molecular oxygen of the inflow condition, and the hydrogen mass flow rate, the equivalence ratio of the M12 REST in this simulation is calculated to be 1.33.

5.3.2 The Baseline Combined-Fuelling Flow Field

An examination of the combustor flow field for the baseline combined fuelling case is undertaken to elucidate the flow physics of porthole boundary layer injection from the combustor step face, and to compare its performance to the unfuelled and inlet-fuelled cases. As was done for the inlet-fuelled case presented in Section 5.2, the changes to the shock structure caused by the baseline combined fuel injection geometry are the first part of the flow field considered. The combined fuelling shock structures are compared

with the unfuelled shock structures in Figure 5.15. Where the shock contours are weak and difficult to see in the plots, solid lines are added for clarity. Major expansions are denoted by a dashed line.

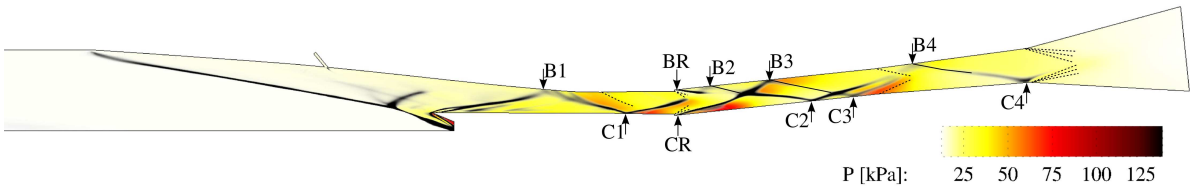

(a) Unfuelled Symmetry Plane Structure

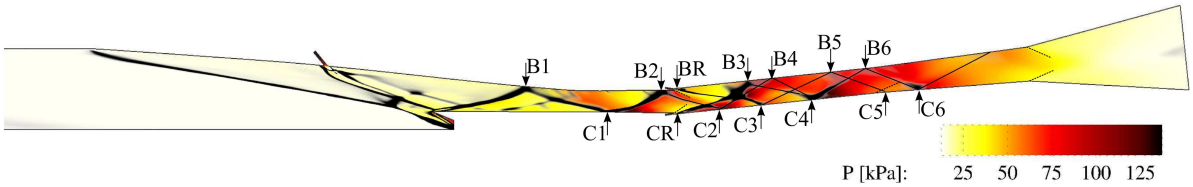

(b) Combined-Fuelling Symmetry Plane Structure

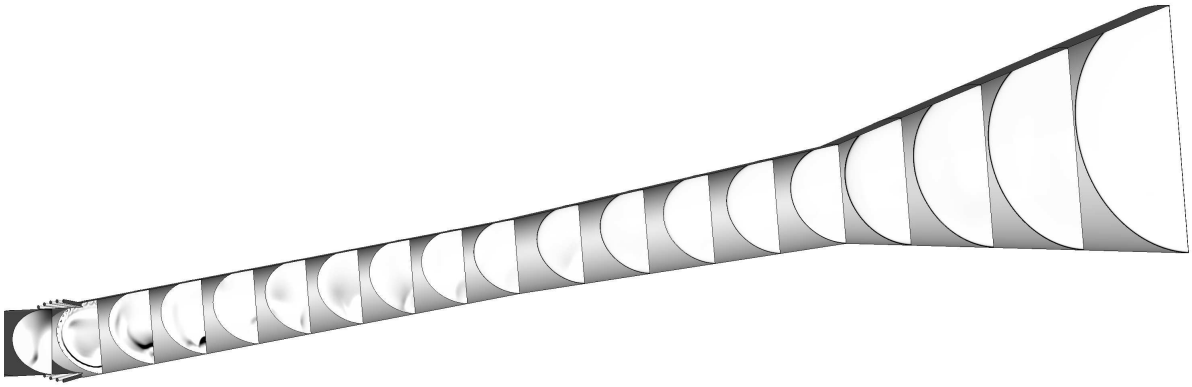

(c) Cross-plane Structure

Figure 5.15: Comparison of M12 REST Combined-Fuelling and Unfuelled Shock Structures

Like the inlet-fuelled case before it, the pressure rise behind the first bodyside reflection (B1) of the cowl closure shock is much higher than that of the unfuelled case, indicative of the combustion of inlet-injected fuel. The pressure rise is also slightly higher than what is observed in the inlet-only fuelling case in Section 5.2. This indicates that the higher equivalence ratio case presented in the previous section (0.55 vs. the current 0.41) did not have as strong an effect on the flow entering the isolator. This is likely caused by higher combustion efficiency in the isolator of the baseline case, which in turn means more thermal compression of the flow inside the isolator.

The cowl closure shock of the combined fuelling case reflects off the bodyside wall of the engine a second time (B2) upstream of the combustor step. This reflection occurs

closer to the combustor step in the combined injection case than it does in the inlet-injection-only case. This is explained by both the difference in inlet fuel equivalence ratio, and the slightly different inflow conditions between the simulations. Both of these changes would have effects on the shock angles inside the isolator, though the bodyside flow composition likely is the dominant effect. The speed of sound changes caused by different quantities of hydrogen and different rates of temperature increase due to combustion heat release would have a significant impact on the speed at which the shock travels through this part of the inlet flow field.

The elliptical recompression shock forms behind the combustor step (BR and CR) as usual, though as indicated in Figure 5.15c, there are additional shock surfaces formed by some of the injectors behind the step. These additional structures form downstream of the injectors along the bodyside of the combustor step, while little evidence of their presence is indicated on the cowlside of the combustor. This is likely due to a difference in the strength of the recompression wave: on the bodyside, the surface is diverging from the direction of incoming flow, and the recompression wave is relatively weak. The injector shock structures can therefore form without being disrupted. On the cowlside, there are additional compression effects caused by the turn in the engine at the step, and the recompression shock is strong enough to disrupt the weaker shock structures generated by the combustor step injectors.

The shock structure of the combined-injection combustor is generally far more complex and compact than that of the unfuelled engine; more shock reflections occur in the constant-area combustor of the combined-fuelling case than occur in the entire unfuelled combustor. While the shocks seem to disappear with downstream distance in the engine, this is an effect of the expansion of the flow in the diverging combustor and nozzle, as well as flow temperature. With a great deal of combustion, the temperature rises high enough across the combustor that the more uniformly high temperature, low density flow further weakens the strength of the gradients across the shock waves.

The pressure rise in the engine is similarly far more severe. Whereas the unfuelled engine shows local pressure rises of up to 75 kPa behind the cowlside shock reflections (Figure 5.15a), the combined-injection combustor shows pressures exceeding 75 kPa across the last half of the constant-area combustor, and most of the diverging combustor. A local peak of 130 kPa occurs behind cowlside reflection C4 in the baseline combined fuelling case. With far stronger pressures than are seen for the inlet-fuelled engine, the engine is able to sustain rapid and robust combustion of the injected fuel.

The presence of robust combustion in the M12 REST engine can also be inferred

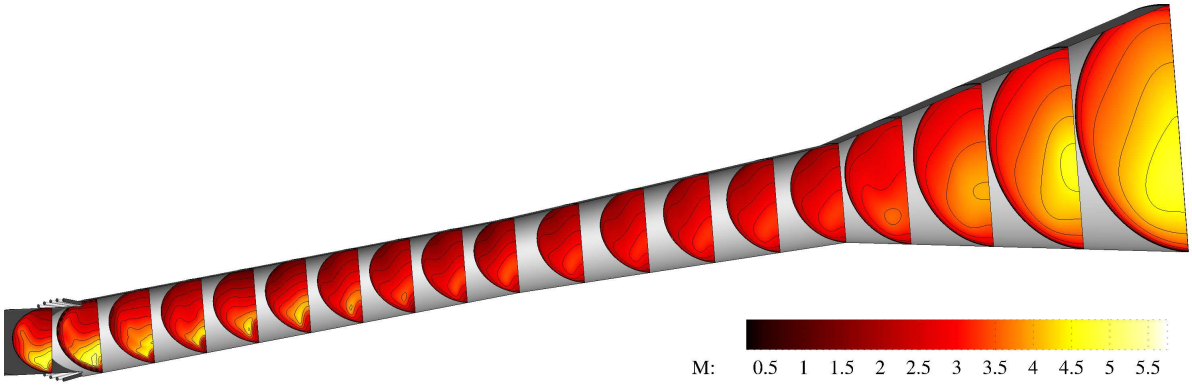

(a) Mach number

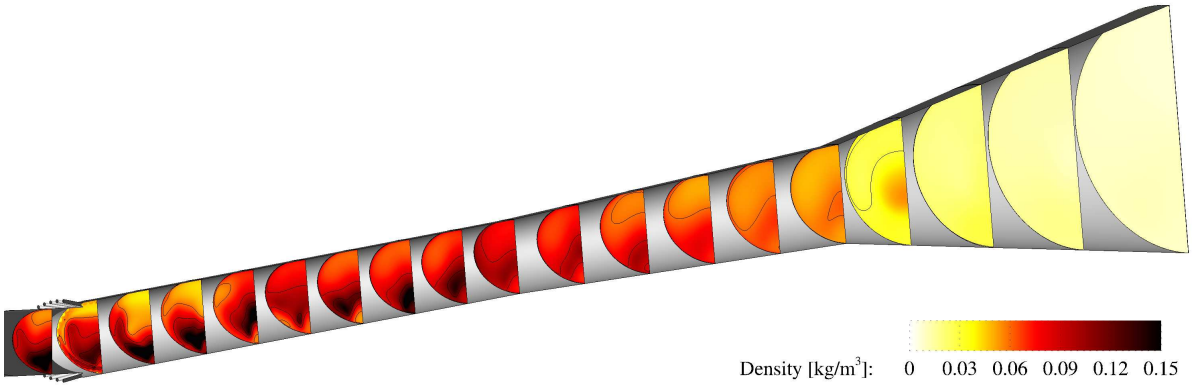

(b) Density

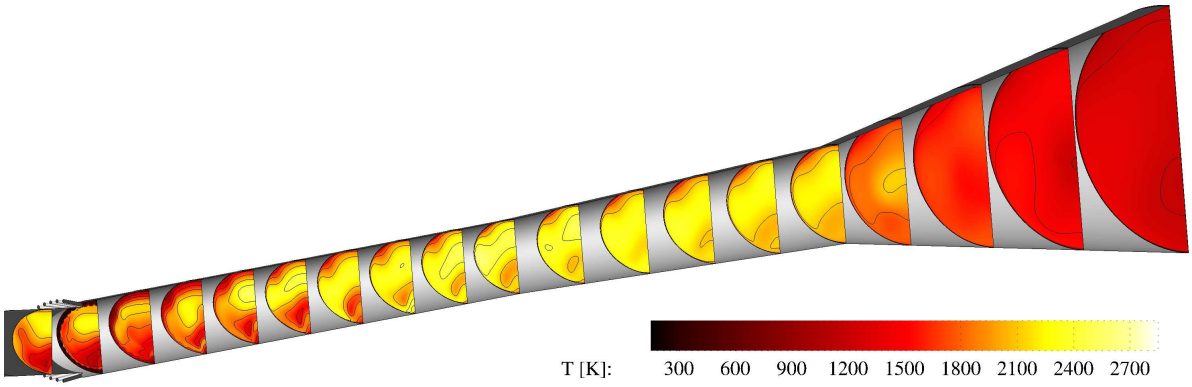

(c) Temperature

Figure 5.16: Mach number, Density, and Temperature Contours in the combustor with combined fuelling

from the local Mach number, density, and temperature contours presented in Figure 5.16. The Mach number contours, for example, show little of the fluctuations observed along the cowlside of the combustor in the inlet-injection case presented in Section 5.2. Instead, the Mach number steadily decreases along the length of the constant-area combustor, and then stays relatively constant through the diverging combustor section. The peak Mach number in the diverging combustor is slightly over Mach 3, which is substantially lower than the Mach number in the inlet-fuelled case. There is still a significant Mach number asymmetry at the nozzle exit which, like the inlet-fuelled case, is due to the higher temperatures and lower densities caused by combustion occurring more vigorously along the bodyside of the engine.

The density flow field in Figure 5.16b shows some evidence of following the pattern of the shock reflections along the cowlside of the engine. This behaviour does not extend to the bodyside portion of the flow, where the low flow density caused by the flow expanding over the bodyside of the combustor step slowly recovers, and then remains relatively stable through the constant-area combustor. This low density locally weakens the density gradients across the faces of shock waves passing through the bodyside region, explaining their apparent absence along the bodyside of the combustor. On average, the density is higher throughout the combined-fuelling combustor when compared to the inlet-injection case, which matches the pressure behaviour observed.

The temperature contours shown in Figure 5.16c clearly demonstrate that the fuel injected behind the step remains close to the walls for at least the first 15 jet diameters downstream, as indicated by the low temperature ring of flow around the circumference of the combustor wall in the first contour slice downstream of the step. This low-temperature layer is somewhat thicker along the bodyside, which is to be expected given the bodyside surface acts as an expansion ramp in this region. Meanwhile, the cowlside flow is already showing signs of additional compression thanks to the turn of the local engine wall. By the time flow reaches the second contour slice downstream, it has already warmed significantly, such that little of the flow has a temperature below 1500 K. Such rapid heating suggests that the stream of hot, partially burned fuel-air mixture entering the combustor along the bodyside wall does indeed pilot ignition of the step-injected fuel; viscous effects alone would not be sufficient to heat the entire flow so quickly.

The temperature continues to rise along the length of the constant-area combustor. Aside from a small pocket of low temperature flow along the top "corner" of the combustor, the flow temperature exceeds 1800 K across the span of the combustor at the

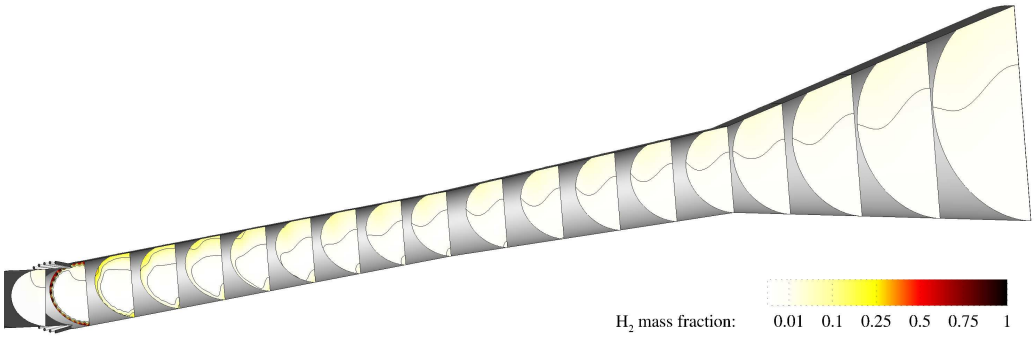

(a) Molecular Hydrogen distribution

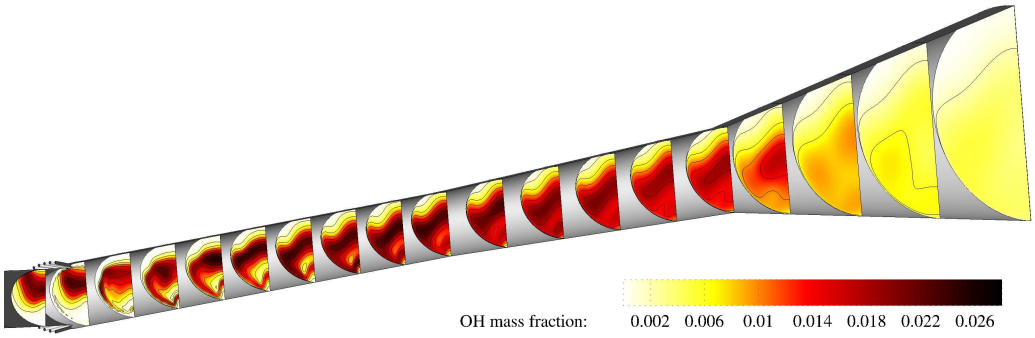

(b) OH radical distribution

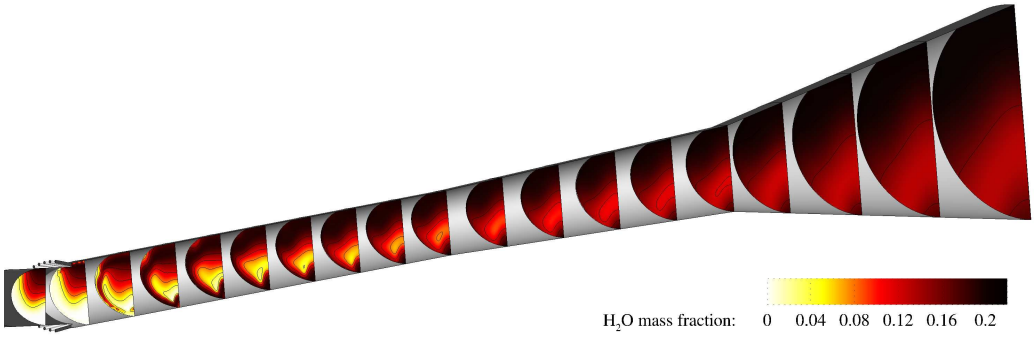

(c) Water vapour distribution

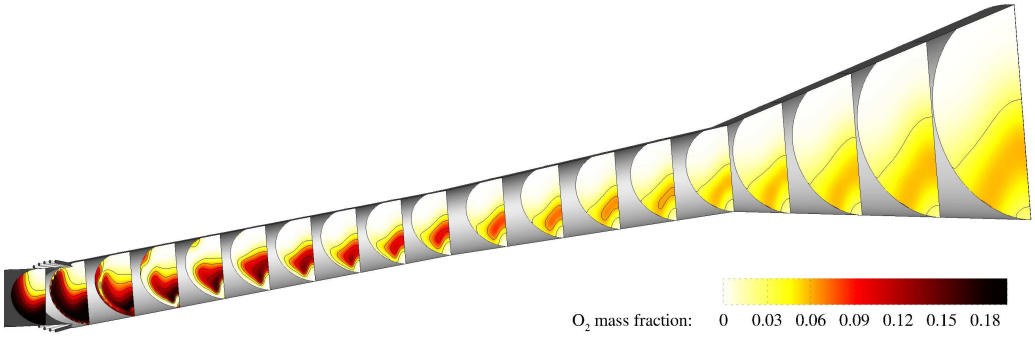

(d) Molecular Oxygen distribution

Figure 5.17: Species mass fraction contours in the M12 REST combustor with inlet fuelling

entrance to the diverging combustor section. Through the first half of the diverging combustor, the flow temperature is greater than 2400 K in a large band of the flow across the bodyside half of the engine. The flow downstream is generally greater than 2100 K, up until it is rapidly expanded and cooled through the engine nozzle.

5.3.3 Flow Chemistry and Combustion

Up to this point, much speculation has been made about the nature of the combustion occurring in the flow, and whether the hydrogen injected from the combustor step remains attached to the wall. In order to check this behaviour directly, the mass fraction distributions of the key combustion species are examined. These distributions are presented together in Figure 5.17.

The distribution of hydrogen confirms that in the first few diameters downstream of the step, the hydrogen remains attached to the walls, though they remain in discrete jets, and do not appear to merge. By the time the flow reaches the location of the first slice downstream of the combustor step, the high-H_2-concentration jet cores are separated by flow regions where the mass fraction of hydrogen is less than 0.1, as shown in Figure 5.17a. This suggests that the portholes are not spaced closely enough to provide a continuous and relatively uniform sheet of fuel along the combustor surface. With a center-to-center spacing of ∼2.5D, these portholes are in fact spaced further apart than the maximum porthole spacing found to give a coherent sheet of fuel in Barth (2012).

There is also considerable variation in the size of the jets in the first combustor slice, which is indicative of another important development in the hydrogen layer injected from the step. As shown in the next slices downstream, the fuel is quickly pushed away from the cowlside wall, leaving the flow along most of the cowlside combustor wall devoid of hydrogen. This is illustrated visually in Figure 5.18.

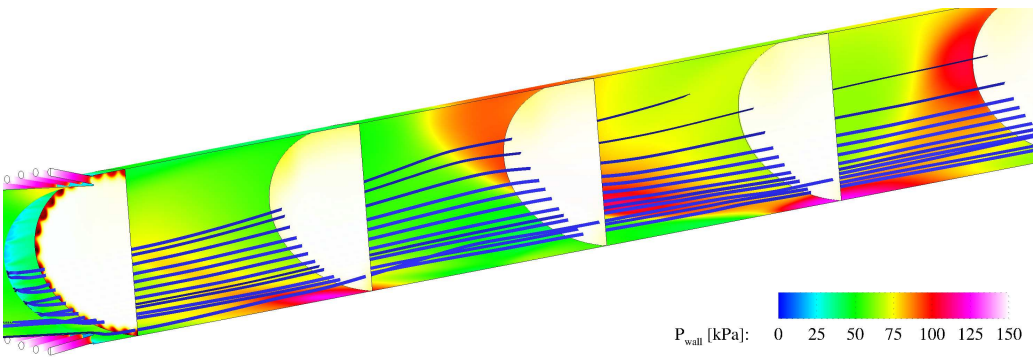

Figure 5.18: Streamlines in the constant area combustor. Wall contours are of local pressure.

The streamlines that pass through the fuel streaks in the first combustor slice in Figure 5.18 almost universally show a lift-off from the cowlside surface that occurs near the high pressure region downstream of the cowlside combustor step. Some fuel is swept up the sidewall, while fuel near the centerline is pushed upward. Flow coming over the step moves into the region immediately adjacent to the wall, indicating that the fuel is either entirely pushed upward, or is well-mixed with the oncoming air. This suggests that any combustion that does occur inside the cowlside boundary layer is short-lived. The fuel concentrated at the centerline of the cowlside wall persists into the diverging combustor before being completely consumed, while a great deal of the fuel injected, or swept up into, the bodyside flow region remains unreacted through the length of the combustor and nozzle.

The OH radical and water vapour distributions in Figure 5.17b and 5.17c both show that the partially burned flow entering the combustor promotes ignition of the step-injected hydrogen, and quite rapidly. A thin strip of OH is already visible along the cowlside of the engine in the first combustor slice (7D downstream of the step injection), and both the OH and water distributions show strong species production across the top of the hydrogen-laden boundary layer in the next slices downstream. By the time the fuel and air has moved a third of the way through the constant area combustor, the cowlside fuel layer appears to be burning.

The oxygen contours shown in Figure 5.17d prove that combustion occurs rapidly inside the constant-area combustor. Any oxygen remaining along the sidewalls as the flow exits the isolator is consumed by the time the flow reaches the middle of the constant area combustor. A thick, oxygen-poor region along the cowlside wall and centerline fuel structures indicates that combustion occurs in these regions. The combustion has evidently ceased along the cowlside wall outside of the centerline region by the time the flow reaches the third combustor slice, 50D downstream of the step. Downstream of this point, further consumption of the central oxygen core only occurs along the its bodyside- and centerline-structure-facing boundaries. Oxygen begins to diffuse downward to the cowlside wall as the flow continues downstream, further indicating the absence of combustion on the cowlside surfaces of the combustor. By the time the flow reaches the nozzle, the oxygen is confined to a chevron-shaped structure anchored on the cowlside wall on either side of the centerplane, and far removed from the remaining hydrogen flowing along the bodyside of the engine.

5.3.4 Engine Performance in the Baseline Injection Case

Having examined the combustion that occurs in the baseline case qualitatively, what the observed reacting flow translates to in terms of combustion efficiency is now examined. Because this simulation is of an engine operating in a fuel-rich mode (i.e. $\phi > 1$), the mixing efficiency needs to be redefined in terms of the limiting reactant, oxygen. Similar to the hydrogen based mixing efficiency, this is defined as

$$\frac{\dot{m}_{O_2,mix}}{\dot{m}_{O_2,total}} = \frac{\int Y_R \rho u dA}{\int Y \rho u dA}, \tag{5.2}$$

where

$$Y_R = \begin{cases} Y_{O_2} & \text{if } Y \leq Y_{stoich}, \\ Y_{stoich}\frac{1-Y_{O_2}}{1-Y_{stoich}} & \text{if } Y > Y_{stoich}. \end{cases} \tag{5.3}$$

The oxygen-based combustion efficiency is defined as the mass of oxygen fully burned and present in the form of water vapour, divided by the total mass of free oxygen captured by the inlet:

$$\frac{0.8881\dot{m}_{H_2O}}{\dot{m}_{O_{2,total}}}. \tag{5.4}$$

This combustion efficiency is equivalent to the hydrogen-based combustion efficiency used previously. The local oxygen-based efficiencies are calculated along the length of the engine, and are presented in Figure 5.19.

The now-familiar mixing and combustion efficiencies seen in the M12 REST inlet give way to a rapid jump in the mixing and combustion efficiencies immediately downstream of the combustor step. This indicates that porthole injection from the combustor step face allows some of the injected hydrogen to mix rapidly with the air flowing over the step from the isolator, due to the large interfacial area between the 36 individual jets and the oncoming air. The previously demonstrated spread of the bodyside flame exiting the isolator around the combustor allows this rapidly mixing fuel to burn quickly, aided in part by the pressure rise induced by the shock reflection (B2) just upstream of the step. By the time the $\phi = 1.33$ flow reaches the end of the constant-area combustor, the local combustion efficiency has jumped from 25.41% to 68.88%, with a mixing efficiency of 83.17% at the same location.

As the flow enters the combustor, the mixing and combustion efficiencies continue to climb, but at a rate indicative of a slower mixing-limited process. Fuel-air mixing is virtually stopped by the time the $\phi = 1.33$ flow enters the nozzle, reaching a final mixing efficiency of 93.17%. This is principally due to the segregation of the remaining

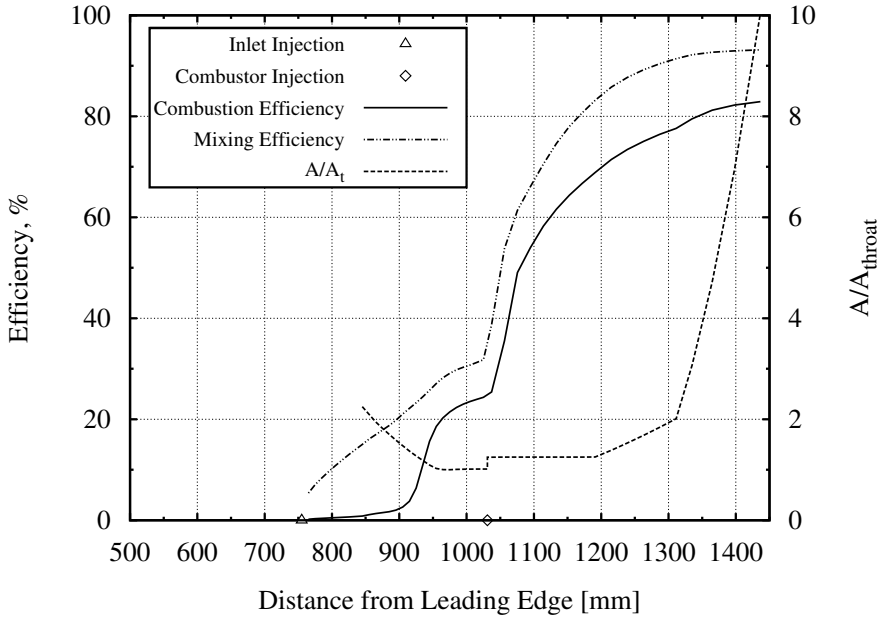

Figure 5.19: Oxygen-based mixing and combustion efficiencies for combined inlet/boundary layer fuelling at a global equivalence ratio of 1.33

hydrogen inside the oxygen-poor bodyside surface of the combustor, further exacerbated by the subsequent expansion of the nozzle. The simulated combustion efficiency at the nozzle exit is 82.9%, which is higher than the 80% threshold suggested to be the point at which net thrust could be achievable (Smart, 2012).

While this combustion performance is impressive, and quite desirable given the amount of fuel injected from the step, it is also very different from the behaviour that has normally been seen in previous boundary-layer combustion studies. In those studies, the layer of hydrogen is continuous, remains confined within the boundary layer, persists for some time downstream of injection, and only burns through a turbulent diffusion process, large-scale mixing does not occur. The mixing efficiency of the $\phi = 1.33$ in Figure 5.19, on the other hand, implies that in the near-field behind the step a nearly reaction-limited combustion of well-mixed fuel takes place. Combined with the lack of a persistent, high-concentration layer of fuel along the cowlside wall of the combustor, it seems likely that the engine does not see the boundary layer combustion the step-injection was intended to provide.

This appears to directly contradict the findings of Chan (2012), who found that a combustor fed by a REST-style inlet should be capable of sustaining boundary layer combustion, thereby reducing the skin friction drag of its combustor. This discrepancy is explained by a key difference in geometry between the engine studied by Chan,

and the M12 REST engine: the engine examined by Chan did not have a turn in the engine at the engine step, and employed injection from a slot, rather than from discrete portholes. The difference in injection likely only has a small effect; the light hydrogen gas will always be subject to the forces of motion in the bulk engine flow. The turn in the engine, on the other hand, likely has a large effect. The "compression corner" effect of the cowlside wall creates a region that either forces fuel away from the wall, or rapidly mixes oncoming air into the fuel layer, ending the opportunity for drag reduction along the portion of the combustor most in need of it.

Aerodynamic forces acting on the engine

The viscous and inviscid forces acting on each internal surface have been obtained, and broken out on a component level. This makes it possible to determine if there was any skin friction drag reduction in the combustor at all. It also makes it possible to see if crossing the 80% combustion efficiency threshold made the shock tunnel engine model come close to producing net thrust. The forces generated in the thrust and drag directions are summarized in Table 5.3.

Table 5.3: Aerodynamic thrust (+) and drag (-) forces acting on the baseline M12 REST

Engine Component	Inviscid Forces [N]	%	Viscous Forces [N]	%	Combined Forces [N]	%
Inlet Internal Surfaces	-12.582	23.57	-12.291	23.03	-24.873	-108.19
Forebody	-5.316	9.96	-3.237	6.06	-8.553	-37.20
Isolator and Combustor Step	0.898	3.30	-1.760	2.65	-0.862	-3.75
Constant Area Combustor	0.265	0.87	-8.305	15.56	-8.040	-34.97
Diverging Combustor	9.335	30.72	-6.181	11.58	3.154	13.72
Nozzle	19.889	65.45	-3.706	6.94	16.183	70.39
Total Drag Forces	-17.898	33.53	-35.480	66.47	-53.378 N	
Total Thrust Forces	30.387	100.0	-	-	30.387 N	
Total Internal Force	17.805 N		-32.243 N		-14.438 N	
Total Force	12.489 N		-35.480 N		-22.991 N	

The internal forces on the tunnel engine model still sum up to a net drag force, despite a large increase in the thrust forces generated by the diverging combustor and nozzle sections. While the internal net (drag) force is two-thirds the magnitude of the inlet-injection case presented in Section 5.2.2, it is still roughly 55% of the net internal force that acts on an unfuelled engine. The inlet remains the single largest source of drag, with its internal surfaces accounting for 46.5% of the total drag generated. While net thrust is up 52% from the inlet-injection case, an additional 47.5% thrust would

need to be generated by the tunnel engine model to overcome the net internal drag. Alternatively, viscous drag would need to decrease by 35% to achieve an uninstalled cruise (zero net force) condition.

It should be stressed at this point that these results are for a tunnel model of the engine. As such, it has 300 K walls (much colder than would be expected in a flight condition) and lacks a well-designed nozzle. As will be shown in Chapter 7, hotter walls would reduce the amount of heat lost from the flow, which would ultimately drive up combustor flow temperature and pressure. Combined with a well-designed thrust nozzle, an engine in flight exhibiting the combustion behaviour seen here could very well produce net thrust where this tunnel model cannot.

The fraction of overall engine drag created by viscous effects *increased* in this baseline combined fuelling case, despite the fuelling ostensibly being designed to reduce the combustor skin friction drag. As described previously, this is caused by the non-uniform compression of the flow that occurred along the cowlside of the combustor, which disrupted the fuel layer within the first 20% of the constant-area combustor's length. If viscous drag reductions did occur, they will have done so in the region just downstream of the combustor step, before the cowlside flow is significantly affected by the effects of the engine turn. The wall shear stress distribution, presented in Figure 5.20, is used to discern how the skin friction distribution changes, and if there is any skin friction reduction benefits from step injection anywhere in the constant-area combustor.

By contrasting the shear stress distribution in the unfuelled combustor to the viscous stress in the combustor when inlet and step injection are combined, it is immediately obvious that the shear along the cowlside wall actually increases significantly in the region 20D downstream of the step. This is due to the combination of a number of effects. First, the hydrogen injected along this surface does not form a coherent sheet, and therefore cannot shield of the wall from the incoming flow, which first reaches the wall in the regions between the jets. This explains the 'streaks' of high shear present in the first third of the constant-area combustor in Figure 5.20b. Second, these regions between the jets also contain significant combustion of well-mixed fuel, rather than the diffusion flame sheet usually seen in boundary layer combustion studies. This kind of burning indicated in the species mass fraction and combustion efficiency plots is not capable of disrupting the turbulent transport of momentum to the combustor wall in the same way the turbulent shear layer mixing and combustion inside of the log-layer of a compressible turbulent boundary layer can (see Barth et al. (2013), which is included as Appendix G).

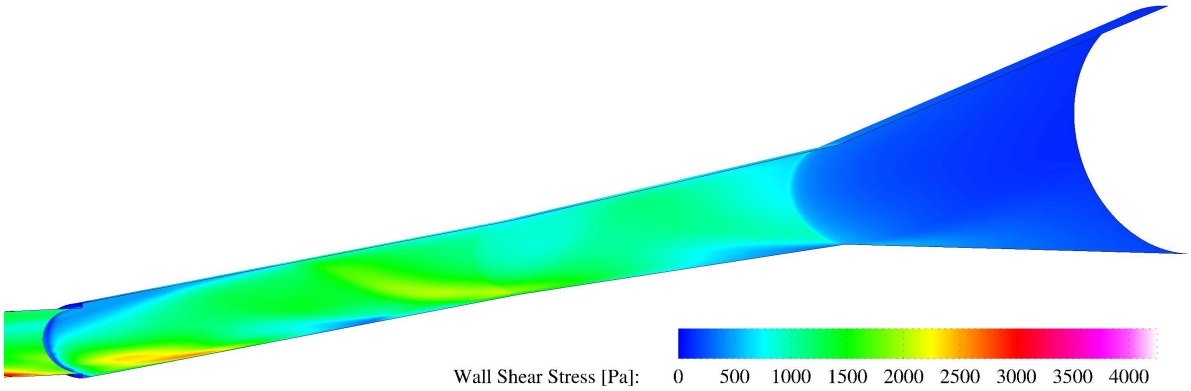

(a) Unfuelled

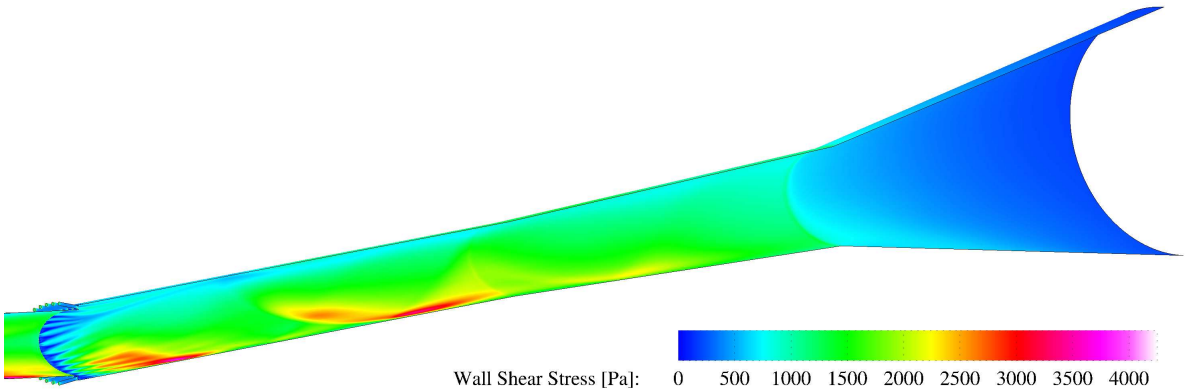

(b) Combined Injection, $\phi = 1.33$

Figure 5.20: Comparison of the unfuelled and combined-injection combustor τ_w distributions

Further downstream, in the region 50 mm upstream of the entrance to the diverging combustor section, a second high shear region is observed to form. This region has the same general form of the similar, but much weaker, shear pattern observed in the unfuelled distribution. This high-drag zone is the result of the stronger shock reflections that occur in the baseline fuelling case, caused by the generally higher pressures and temperatures in the combustor. These conditions lead to a thinner boundary layer than in the unfuelled case, which leaves the wall shear stress more sensitive to the effects of the shock impinging in this region.

By integrating the x- and z- components of the wall viscous shear stress along varying portions of the constant-area combustor, it is possible to determine where inside the constant area combustor that a net reduction in viscous drag is realized. These regions are visualized in the plot included as Figure 5.21. As is visible, the boundary layer injection method used does in fact provide a small reduction in net viscous drag early in the combustor.

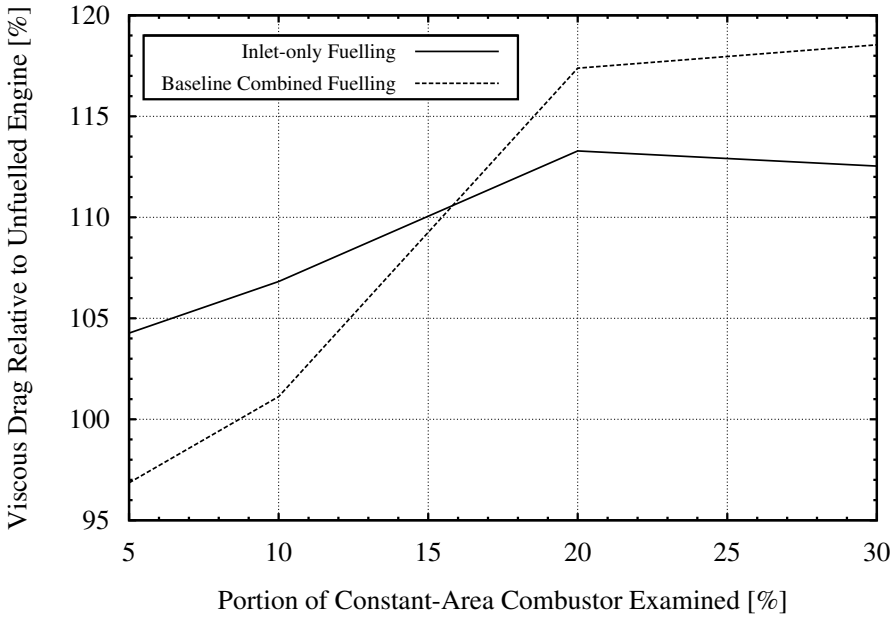

Figure 5.21: Visualization of where step injection provided a skin friction reduction

Relative to the unfuelled internal flow path, the baseline case shows a 4% decrease in drag over the first 5% of the combustor. Downstream of this point, the viscous drag increases sufficiently that any benefits due to step injection are lost within the first 9% of the combustor. The benefits are slightly larger when the baseline case is contrasted with the inlet-only fuelling case: over the first 10% of the combustor length, the baseline case shows 6% less viscous drag than the inlet-injection case examined in Section 5.2. As the cowlside layer of fuel is disrupted, this benefit is lost, and by the time the flow has passed through 16% of the combustor, the baseline case is once again generating greater net viscous drag.

Overall, the viscous shear stress in the combustor is increased by the high pressures and temperatures associated with combustion across much of the engine's cross-section. This increase even extends into the nozzle, where higher shear stress magnitudes are observed downstream of the nozzle expansion corner on the cowlside of the engine. While the present boundary layer injection configuration does not provide any appreciable global benefit to the engine's viscous drag, alternate methods of viscous drag reduction may yet be feasible.

5.3.5 Comparison to experiment

Like the unfuelled and inlet-injection flow cases examined previously in this chapter, the simulation of combined injection in the M12 REST engine has been conducted with

inflow conditions designed to facilitate comparison with experimental data. Pressure data was collected during experimental Shot 11471, and is normalized against a time-offset forebody pressure, as described in Appendix F. The normalized pressure measurements are compared with the distribution extracted from the simulation. This comparison is shown in Figure 5.22.

The simulation and experiment agree quite well along the inlet bodyside surface. The nozzle pressure contours are also generally well-reconstructed, which is quite remarkable given the apparent lack of agreement through the combustor. The measured pressures on the bodyside and cowlside of the combustor all have magnitudes higher than the simulation. Neither side of the engine appears to show the location of strong shock-induced pressure peaks in the experiment, something the other experiments used for comparison so far all show.

This is believed to be due to transient effects in the data. Examining the time-history data for individual pressure transducers during Shot 11471, there is an anomalous signal in the data, indicating that the fuel injected prior to the arrival of the shock burned vigorously during start-up of the engine. This burning in turn leads to the large pressure rise seen in the combustor, and also likely slowed the expulsion of the excess burning fuel from the engine. It is believed that this was responsible for the lack of agreement between the simulated (steady-state) and experimental engine data. More on this phenomenon, and the individual pressure time histories for each engine measurement location can be found in Appendix H.

No wall heat transfer data is available for the experiment in question, but much like the inlet-injection case, an inference can be made based on the values of y^+ along the lines which heat transfer gauges could have been mounted for comparison. The y^+ values are generally between 0.75 and 1 for the majority of the length of the combustor, as shown in Figure 5.23. While this shows that boundary layer viscous sublayers are well-enough resolved to accurately reproduce wall velocity gradients and simulate wall pressure contours, it is likely that the wall heat fluxes in this simulation are an over-prediction. It is likely that the simulation can be considered a reasonable representation of what truly happens during steady operation of the M12 REST at the inflow conditions used, based on both the the good agreement in the inlet and nozzle regions, and the excellent agreement that the CFD simulations have shown with experiments in all previous comparisons.

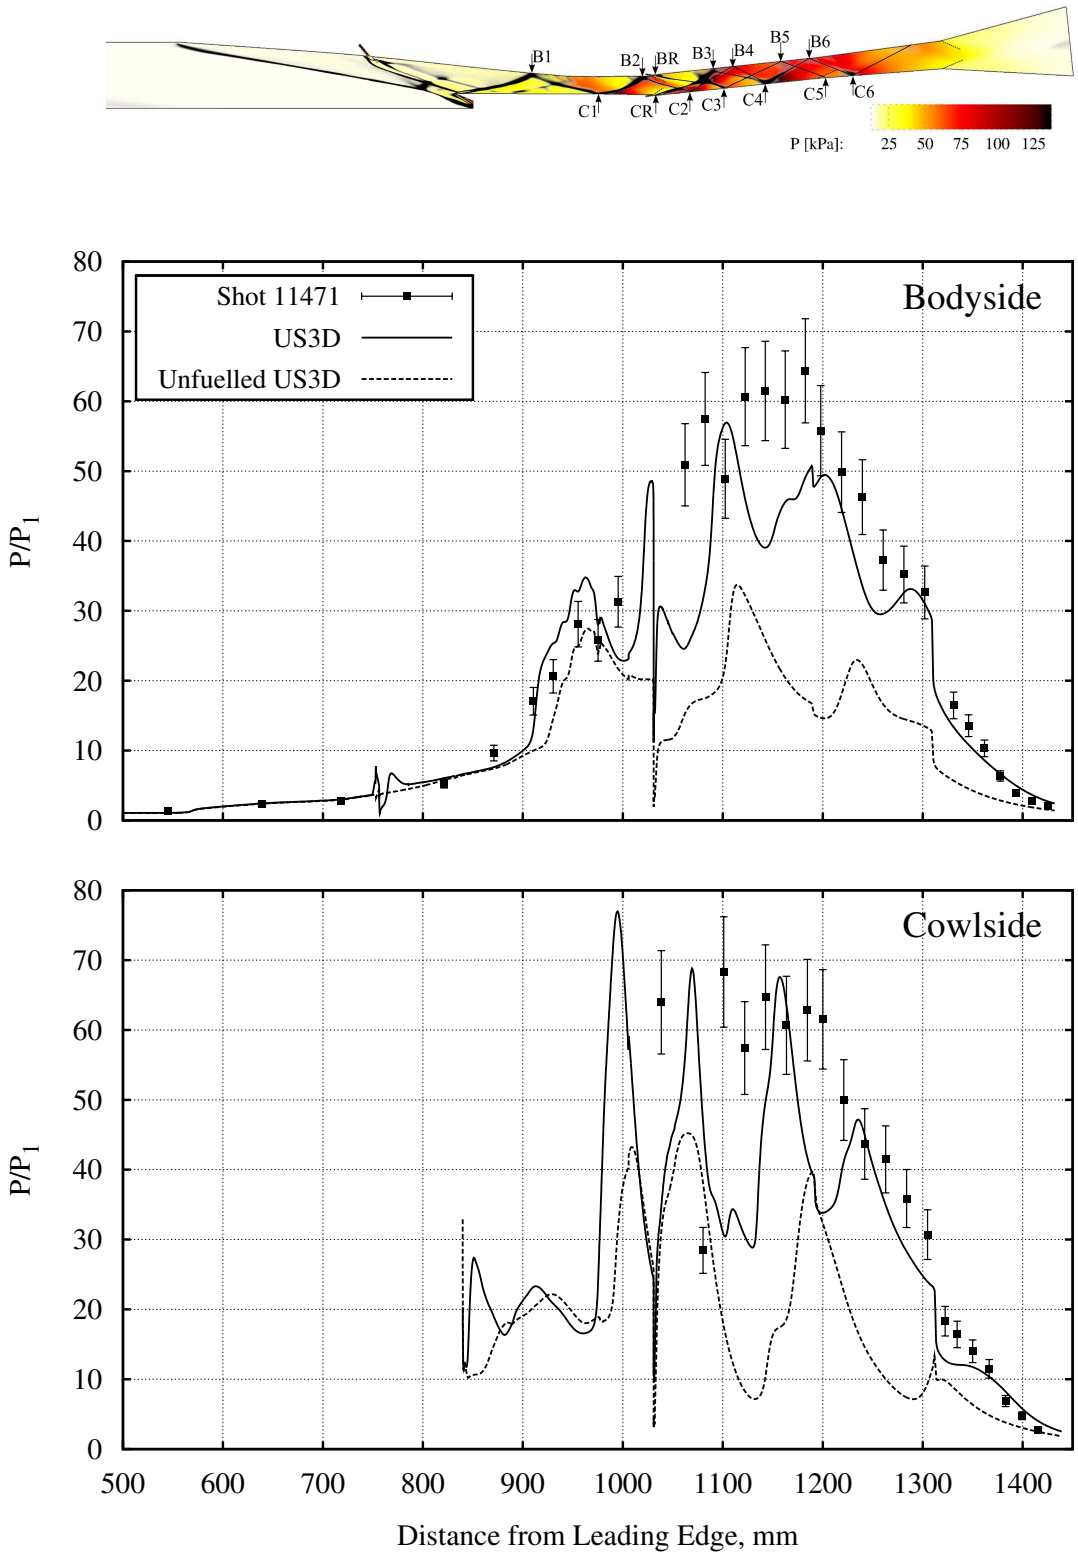

Figure 5.22: Comparison between simulated and experimental pressure data in baseline engine

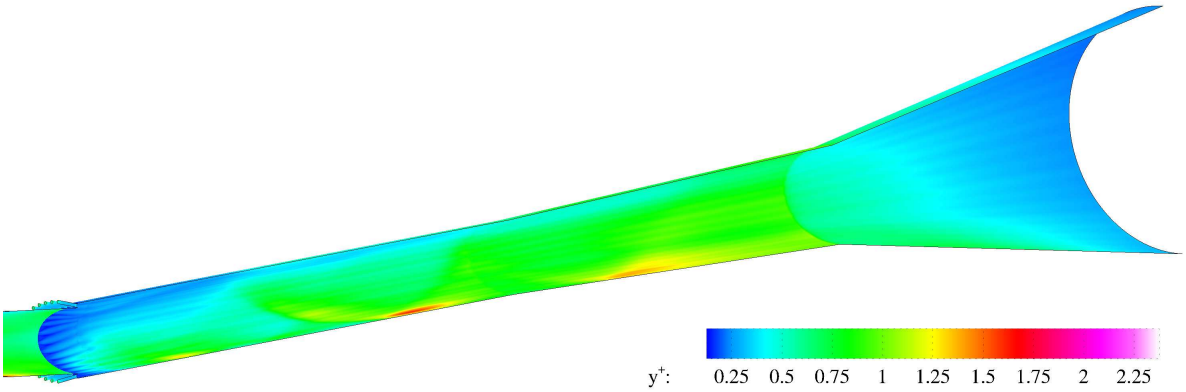

Figure 5.23: The distribution of y^+ in the combustor of the combined injection case

5.4 Summary

Having examined the M12 REST combustor and nozzle flow fields in detail for three separate engine operating scenarios, many important conclusions have been drawn. In general, the simulations agree well with the available experimental data for each case, though the wall-adjacent cell heights are sufficient for velocity and pressure matching only. Further grid refinement along the engine walls will be required if accurate capture the magnitude of wall heat flux in the CFD simulations is desired, though the trends in the heat flux behaviour of the simulation match the experimental trends well. The flow through the combustor is shock-dominated in all cases; any coherent streamwise vortices are weakened and suppressed quickly by passing shock waves. The mixing in the combustor is primarily driven by small-scale turbulent eddies, shear-layer mixing, and by local vorticity generated from the passage of a shock.

Inlet injection at an equivalence ratio of 0.55 is found to give a modest increase in engine thrust, driven by 77.9% complete combustion of the injected hydrogen fuel. Increased inviscid and viscous losses on the inlet brought about by the higher equivalence ratio confirm the postulation of Chapter 4: inlet injection at equivalence ratios higher than 0.33 leads to higher losses in the inlet, without a noticeable improvement to thermal compression at the throat. This case shows poorer mixing efficiency at the throat compared to the lower inlet equivalence ratio in Chapter 4, and based on the comparison of combustion efficiencies at the throat, would cause a greater proportion of the inlet-injected fuel to exit the engine unburned.

Finally, the baseline combined inlet/combustor boundary layer injection case has been examined. Robust burning causes a strong pressure rise throughout the combustor, and further gains in thrust are observed. However, despite more than 80% of

the available oxygen being consumed, net internal thrust has not been achieved in the tunnel model engine. This is likely at least in part due to the fact that the combustor step fuel injection did not provide the appreciable combustor skin-friction reduction intended. The fuel stream from each porthole remains distinct, until the point they are completely disrupted, allowing for much more rapid mixing and combustion of the injected fuel without appreciable film-cooling or boundary layer combustion effects.

The results of this chapter provide a sound understanding of, and baseline for, the engine behaviour and performance at on-design conditions. Efforts to improve the engine's performance can now be undertaken on the basis of this work. These efforts will focus on tailoring fuel injection to the combustor flow field, which is the subject of the next chapter.

6

Tailored Fuel Injection

"Basic research is what I am doing when I don't know what I am doing."

– Wernher von Braun, *in an interview with the NY Times* (1957)

Armed with a complete understanding of the M12 REST engine flow structures with and without fuel injection, it is now feasible to examine how fuel injection might be redesigned to take advantage of the engine's internal flow path. A more in-depth discussion of the flow structures identified in previous chapters is undertaken, followed by case studies of potential injector configurations used as a proof-of-concept for tailored fuel injection. As one of these cases was verified experimentally, a comparison between the tunnel condition with its equivalent flight condition is made to determine how engine performance during tests in the T4 Stalker Tube facility should compare to an engine with an uncontaminated, atmospheric inflow.

Inlet injection is maintained as a constant part of the engine's operation, and so will influence how injection further downstream is placed. As was discussed in Chapter 4, this inlet injection may be used as a source of heat, compression, and chemical radicals to ignite the combustor-injected fuel. Furthermore, the same approximate fuel split used in the baseline configuration of the previous chapter is employed, allowing for a direct comparison to determine the performance gains reaped by moving to a tailored injection configuration in the combustor.

6.1 Targeted Flow Structures

In order to truly tailor fuel injection to the flow path of a scramjet engine, detailed knowledge of the local flow structures in the zone of injection is required. Ideally, this process would be iterative: injectors would be placed near any flow structures considered to have potential for aiding the mixing and combustion of fuel, and moved as

required to maximize engine performance. However, this study is constrained by several geometric considerations that limited the area in which injectors could be placed in the M12 REST engine.

6.1.1 Constraints on Injector Location

Because verification experiments were undertaken jointly with the testing campaign of Wise (2014), there are physical limitations to where fuel injectors can be placed as a part of this tailoring study. Changes to inlet fuel injection are ruled out entirely, as this would require a complete redesign of the experimental model. Similarly, the location of fuel injection in the combustor is constrained to the location of the fuel plenum originally designed for the baseline injection scheme presented in Section 5.3.

This circumferential combustor fuel plenum wraps around the isolator section of the experimental engine model, beginning approximately 25 mm upstream of the entrance to the engine combustor (1006 mm from the leading edge), and is 18 mm long in the streamwise direction. Further constraining the porthole location is the desire to inject the fuel at an oblique angle in order to avoid the total pressure losses and flow blockages associated with normal fuel injection. This constrains the streamwise location of the injectors to an extent imposed by the chosen diameter and angle of the injection portholes.

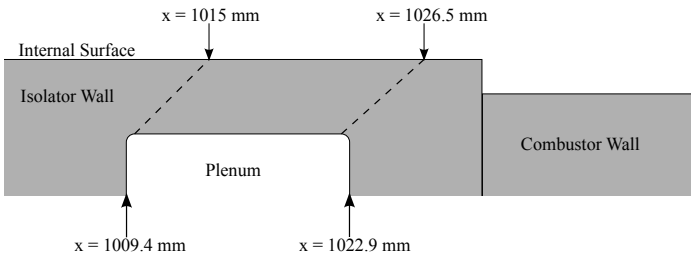

Figure 6.1: Location constraints on a 45°injector fed by the combustor fuel plenum. Streamwise position, x, is defined from the inlet leading edge.

The range of injector locations is shown in Figure 6.1. For simplicity's sake, and to remain consistent with the inlet injectors, the injection angle is held constant at 45°. The furthest upstream location that the leading edge of the injector porthole can exist is 1015 mm from the leading edge, while the furthest distance from the leading edge that the downstream edge of an injector porthole can be located is 1026.5 mm. The actual center of any injector must therefore fall between these lines, at a minimum distance of $D_j \sin(45)$ from either line.

To maximize the fuel residence time inside the engine, it was decided that the injectors would be located as far upstream as possible, with intersection of each injector's centerline with the isolator wall occurring somewhere between 1015-1017 mm downstream of the leading edge, depending on porthole diameter. Armed with this information, it is therefore possible to determine what flow structures should be present in the engine at the point of injection.

6.1.2 Flow Features in the Injection Zone

In previous chapters, effort was made to identify features of the flow that could be of benefit to increasing combustion efficiency by injecting fuel into them. Of primary interest are the regions of the flow that the inlet-injected fuel has been unable to reach. Because this fuel ignites and burns in the isolator, the location of the flame-front in the engine is also of importance, as it provides a heat source for ignition for the combustor-fuel injected fuel. These two features of the flow in the streamwise plane just upstream of the plane where the tailored injectors will be placed are shown in Figure 6.2.

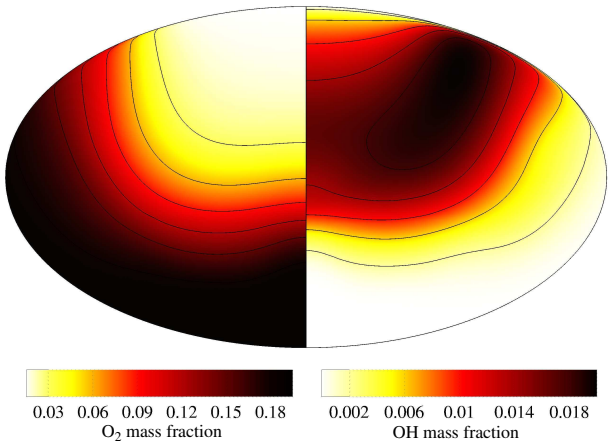

Figure 6.2: Oxygen and OH mass fraction profiles in a plane 1010 mm downstream from the leading edge

Similar to the cross-plane flow structures at the end of the shape-transition of the engine inlet shown in Chapter 4, Figure 6.2 shows that much of the air flowing along the isolator bodyside wall has been mixed and burned with the available oxygen in this region. The OH mass fraction contours show the expected boundary along which a flame-front should exist, as oxygen is transported into the region to react with the hydrogen fuel diffusing outward. This slow arrangement leaves an oxygen-rich zone centered along the cowlside wall at the engine centerplane, and extending up the sidewall.

The location of unreacted oxygen is not the entire story to determining injector location and sizing; the information about flow dynamic pressure, temperature, and the presence of any vortical structures are also of great importance to the ability of a jet to penetrate into, and then mix with, the flow of unburned air. This information is contained within Figure 6.3, and shows that the region along the cowlside wall contains the highest mass flow rate of unreacted molecular oxygen. It also has a dynamic pressure in excess of 800 kPa, with the core flow dynamic pressure exceeding 1000 kPa. The dynamic pressure is a reasonable proxy for mass flux, as its definition includes both flow density and velocity. Thus, the bulk of flow entering the combustor is centered in the regions of highest dynamic pressure. This figure also demonstrates the complexity of the flow through the engine, which is not readily apparent in the diffuse species distribution profiles examined so far.

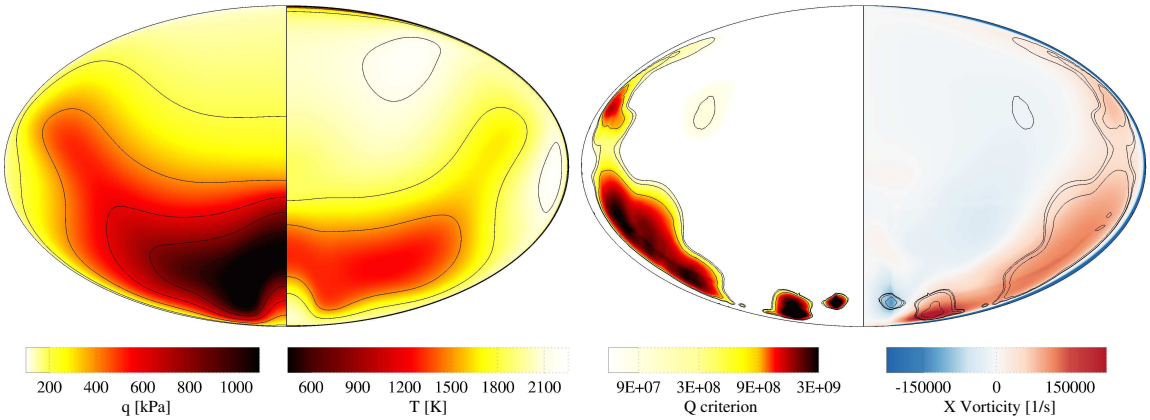

(a) Dynamic Pressure and Temperature Fields **(b)** Q-criterion and Vorticity Fields

Figure 6.3: Other flow properties important to tailored injection in a plane 1010 mm downstream from the leading edge

As was discussed in Section 2.4, the extent to which an injector's jet plume extends into the air crossflow is dependent on both the dynamic pressure ratio between the jet and the crossflow, and on the ratio of jet diameter to boundary layer thickness. The dynamic pressure and diameter of the injector jet are also strongly coupled to the injector mass flow rate. Placing the injector in a high dynamic pressure region creates a challenge in balancing the desire to have the largest diameter and highest dynamic pressure jet possible, with maintaining the mass flow rate balance inside the engine.

Figure 6.3a shows the presence of the hot region along the isolator sidewall that corresponds to the inlet-generated swept separation moving downstream. This weakly rotating flow structure is confirmed by the Q-criterion and streamwise (x) vorticity contours plotted in Figure 6.3b, and may be a good region to place an injector: the

vorticity in the separation could provide extra mixing to a nearby fuel jet, and the high temperature would promote ignition of the fuel. This region may present a tradeoff: having a lower dynamic pressure, getting a jet of fuel well into the mainstream flow along the sidewall would be easier relative to the cowlside flow, but fuel injected into this region would reach less available oxygen than a jet injected out of the cowlside wall.

Also visible in Figure 6.3b are two sets of strong vortex structures moving along the cowlside wall near the engine centerplane. The high Q regions align closely with the strongest regions of vorticity generated, and are propagating downstream. The inner vortex pair is the continuation of the counter-rotating centerline vortex pair generated downstream of the cowl closure, first described in Section 4.1.3. The outer vortices are closer to the isolator wall, and each rotates opposite to its neighbouring centerline vortex. This vortex set was the result of the complex SBLI that occurred when the transmitted cowl closure shock reflects off the cowlside surface just upstream.

The additional vorticity of the fluid in these regions should help promote mixing of the fuel jets with the cowlside core flow of unburned oxygen. As an additional benefit, the centerline counter-rotating vortex pair provides a weaker dynamic pressure region that may be of some benefit to jet penetration.

6.2 Case Study 1: Cowlside Portholes

The first iteration of tailored injection focused on the core of molecular oxygen flowing along the cowlside wall of the isolator into the combustor. Given the high dynamic pressure of the oxygen core, the jets required a similarly high dynamic pressure if they were to penetrated into the flow. In order to keep injector mass flow rates to that required for stoichiometric fuelling with an inlet-to-combustor mass flow ratio of approximately 3:7, total tailored combustor injector porthole area needed to be 2.55 mm^2.

This area was divided amongst four 0.9 mm diameter portholes distributed symmetrically about the engine centerplane to cover the entire high-oxygen mass flow region, as shown in Figure 6.4. The inner injectors were centered 2.25 mm from the centerplane, while the outer injectors were positioned 6.75 mm from the same plane. All injectors were oriented such that their in-plane axis was normal to the local wall, with an injection angle of 45°to the streamwise axis of the isolator. It was determined that the furthest upstream these portholes could be placed was 1015.1 mm from the leading edge, in order to meet the dimensional requirements imposed by the existing

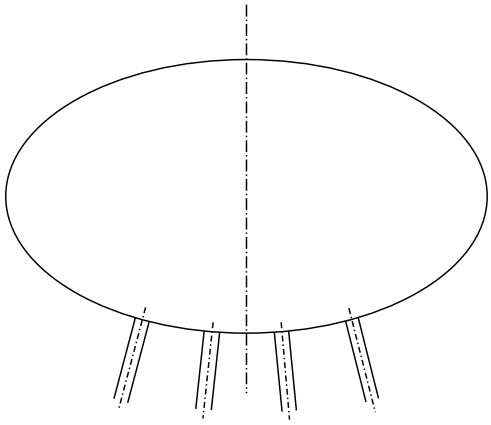

Figure 6.4: Isolator cross-plane position of the tailored injectors in Case Study 1. Streamwise direction is into the page.

combustor fuel plenum. It was calculated that for an isentropic contraction to the injector diameter, that the injector plenum pressure would be approximately 1.1 MPa to achieve a combustor equivalence ratio of 0.7, for a global equivalence ratio of unity. In an experiment, this pressure would need to be higher in order to account for losses in the injector design.

6.2.1 Solution Details

The mesh developed for the tailored flow solution was based on the unfuelled combustor mesh presented in Chapter 3. The injectors were added to the isolator wall and incorporated into the engine mesh using the best practices described in Appendix A. The resulting computational grid in the vicinity of the combustor step is shown in Figure 6.5.

Being integrated into the isolator wall, the injectors had the same wall spacing as the inlet mesh. The high resolution grid downstream of the injectors was coarsened downstream of the combustor step. This allowed for the hydrogen jets to penetrate the freestream and develop for more than 20D before grid coarsening occurred.

This was an exploratory simulation, that is, it was not intended to be compared with experimental data. It was simply to provide a first pass at tailoring fuel injection. As such, this solution was converged first-order accurate in space only. The RMS residual dropped by 3 orders of magnitude, and the mass balance was 6.6×10^{-4} kg/s with a total engine mass capture of 60.47 g/s. Total in-solution time was 0.778 ms, or 1.95 flow-through times. The implicit time step was changing by less than 1×10^{-15} s, which is a good indication that the boundary layers were established. The total fuel mass

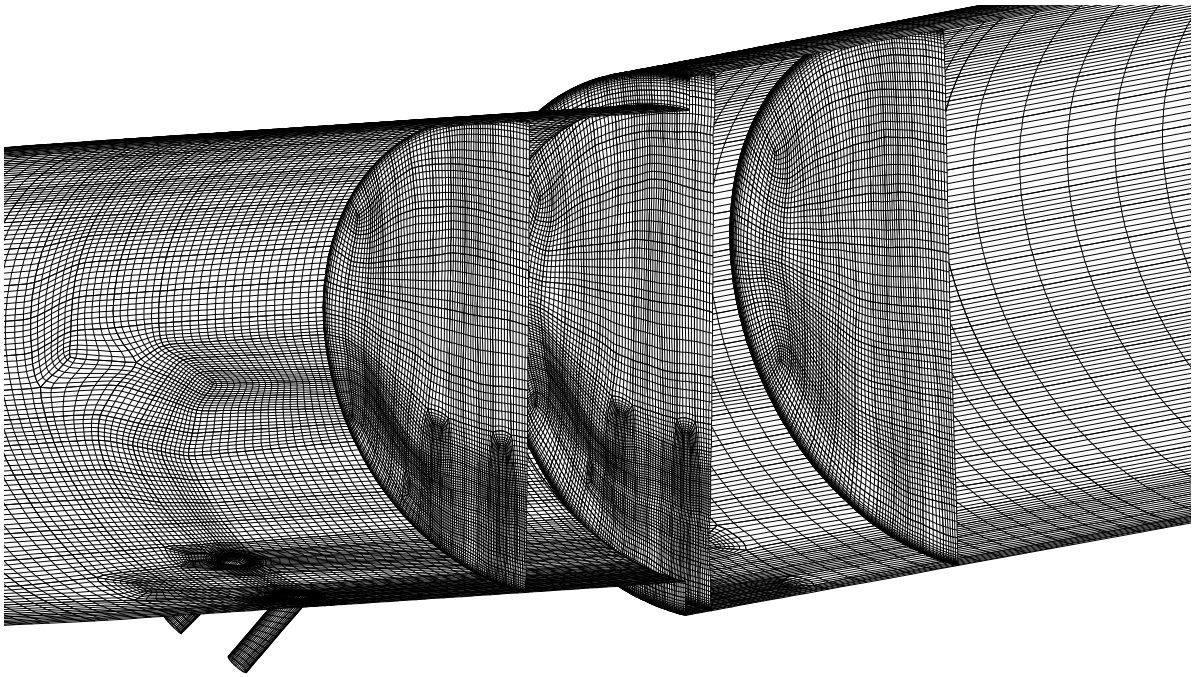

Figure 6.5: The computational mesh used to simulate injection and near-field mixing in Case Study 1

flow rate was 1.556 g/s with an inlet/combustor split of 0.345/0.655, giving a global equivalence ratio of 0.98 based on the mass capture and free oxygen content of the inflow condition used.

6.2.2 The Combustor Flowfield

To examine the effect this injection configuration had on the flow field of the isolator and combustor, the Mach number, density, and static temperature cross-stream contour planes were examined. These variables alone can give some clue of the overall performance of the combustor, and are presented in Figure 6.6.

The Mach number contours shown in Figure 6.6a demonstrate that the flow remains entirely supersonic throughout the combustor. The Mach number of the flow decreases through the constant-area combustor, then remains relatively steady with the cowlside core flow moving at Mach 3, and the rest of the flow moving at or above Mach 2. This is in direct contrast to the Mach 8 REST engine, which demonstrated dual-mode combustion (that is, a mixture of supersonic and subsonic combustion) at on-design conditions (Turner and Smart, 2010).

The density contours in Figure 6.6b suggest that while the complex shock/expansion system downstream of the combustor step largely determines the density field

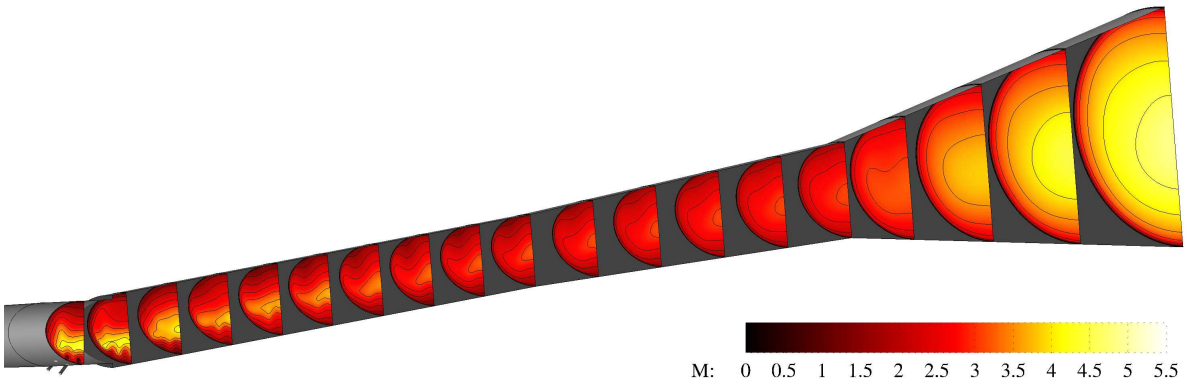

(a) Mach Number

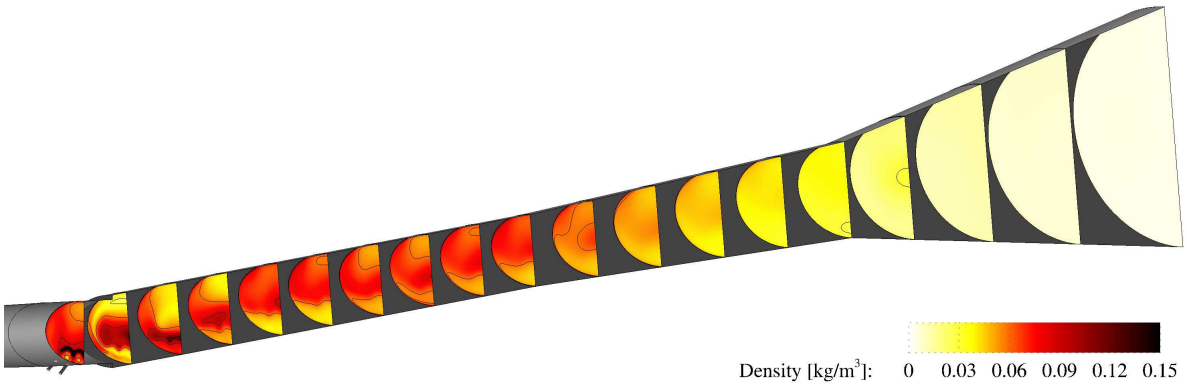

(b) Density

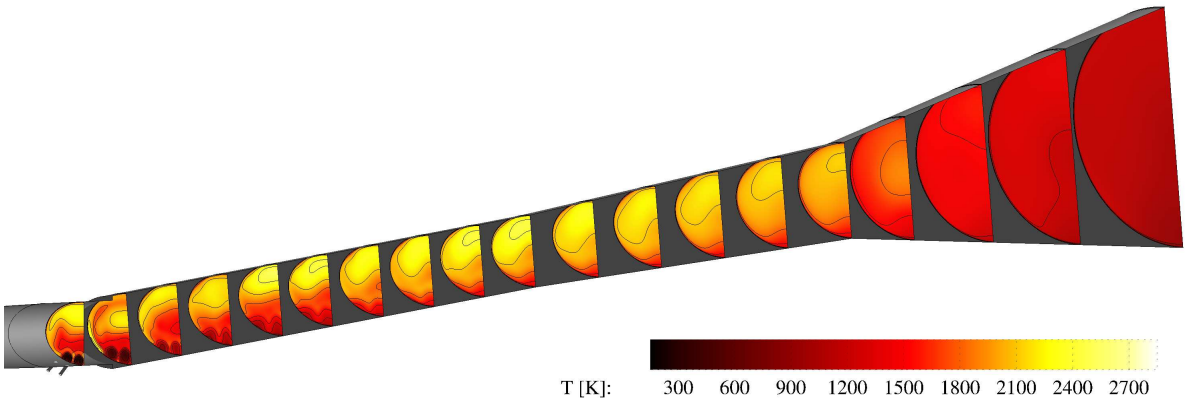

(c) Static Temperature

Figure 6.6: Key variable contours for Case Study 1

in the first 50 mm of the combustor, a consistently low-density region forms along the cowlside wall of the combustor. This region persists downstream, suggesting that much of the injected fuel ends up in this region, and if it is burning, is likely doing so through a mixing-limited process.

The temperature fields appear to support this conclusion, as demonstrated in Figure 6.6c. The low-density region previously identified along the cowlside wall is also at a consistently lower temperature than the rest of the flow. In the first six contour planes downstream from the point of injection. This is primarily due to the over-expanded fuel jets rapidly cooling upon injection, and their temperatures recovering slowly downstream.

The flow along the top half of the combustor is hot, with temperatures exceeding 2400 K for most of the combustor's length, suggesting that combustion of the inlet-injected fuel is continuing. This high temperature region slowly expands across the engine with downstream distance, but even at the end of the diverging combustor, there is a region of flow along the cowlside wall with temperatures at, or below, 1200 K. This suggests that all the heat release due to burning hydrogen is taking place well-away from the cowlside wall and combustion efficiency is likely to be quite low.

6.2.3 Flow Chemistry and Combustion

The actual fate of the hydrogen fuel injected through the tailored injectors can be determined by examining the contours of species mass fractions for key combustion-related molecules and radicals. Of particular interest are the OH radical concentrations which serve as a proxy for flame location, as well as the relative densities of unreacted fuel and oxygen, and the mass fraction of the final product, water vapour. The contours of each of these variables are presented in Figure 6.7.

Many of the features of the flow combustion behaviour predicted by the temperature and density fields are confirmed by the contours presented. The tailor-injected hydrogen only penetrates modestly into the mainstream flow, and clearly stays confined to the cowlside region of the combustor shown in Figure 6.7a. While much of the hydrogen appears to be consumed, a substantial amount is visible at the end of the nozzle, though near the end of the diverging combustor, this relatively hydrogen-rich flow region begins to migrate upward along the side of the combustor duct, as a result of the pressure gradients caused by the turn in the engine at the combustor step.

The contours of unconsumed molecular oxygen in Figure 6.7b mirror the behaviour of the hydrogen. The oxygen near the cowlside wall is rapidly mixed with a fraction

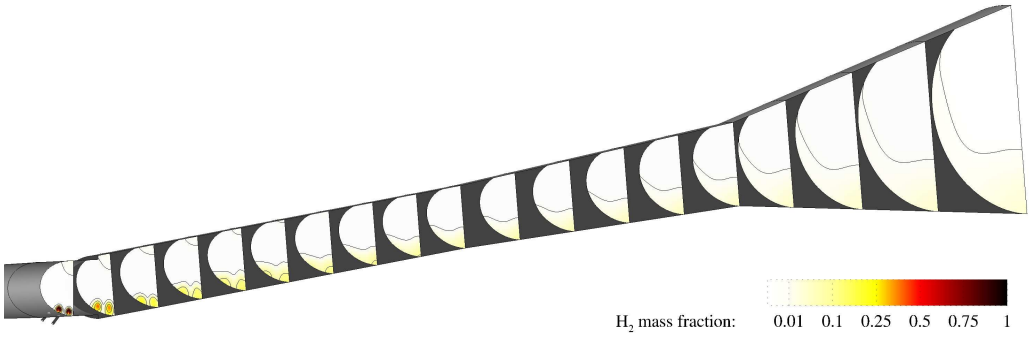

(a) Unreacted Hydrogen

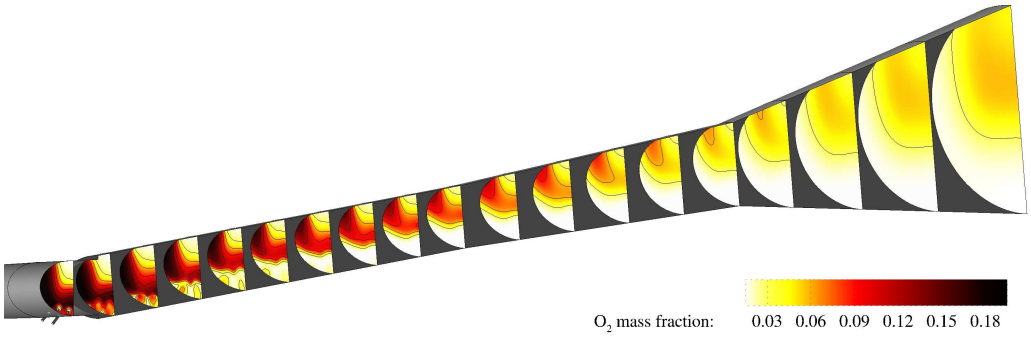

(b) Unreacted Oxygen

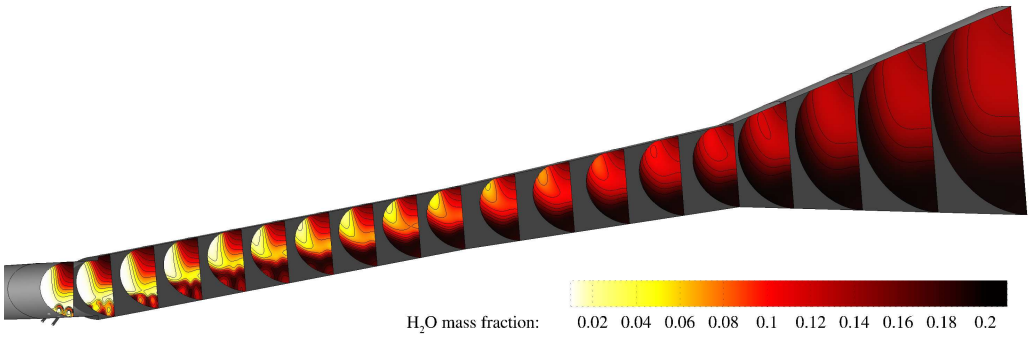

(c) H_2O mass fraction

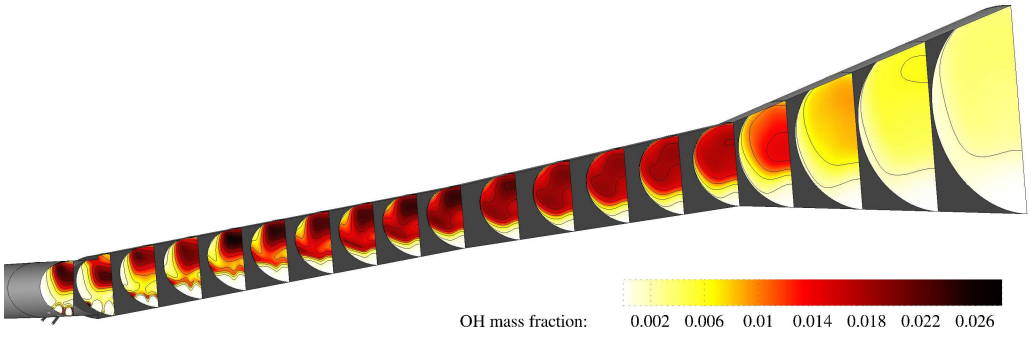

(d) OH mass fraction

Figure 6.7: Chemical indicators of combustion in tailored injection case study 1

of hydrogen downstream of the injectors, and is absent from the cowlside low-density flow region in the last three contour slices of the constant-area combustor. A high-oxygen mass fraction zone persists along the sides of the combustor duct through its entire length, and moves slowly upward with downstream distance. At the end of the nozzle, a large region of unconsumed oxygen exits the engine along the bodyside and central regions of the exit flow area.

The contours for water vapour shown in Figure 6.7 complement the oxygen and hydrogen contours well. The oxygen rich sidewall region shows an absence of water vapour, which slowly diffuses into this area as the flow moves downstream. Similarly, in the oxygen-poor, hydrogen-rich zone along the cowlside wall, the concentration of water vapour increases downstream. They bodyside water-rich region remains relatively unchanged through the constant-area combustor, growing modestly and then mixing across the entire top half of the duct.

Taken together, the three molecular contours suggest that the tailored hydrogen injected is burned in a turbulent diffusion flame sheet that crosses the bottom-third of the combustor. This accounts for much of the behaviour in this region: after the initial oxygen in this region is consumed, the only oxygen available must come from the region above the flame front. Unconsumed hydrogen below the flame font is drawn toward the flame, where it can react with available oxygen, ultimately producing water vapour. This water vapour is primarily pushed downward, taking the place of the upward-diffusing unreacted hydrogen.

This combustion behaviour of the tailor-injected fuel is confirmed by the OH radical contours shown in Figure 6.7d. Just downstream of the combustor step, hydrogen from the outer jet is ignited in the region where the fuel comes in contact with the high-temperature near-wall region created by the interaction of the swept separation with the reflection of the injector bow shock, which extends from the center of the sidewall downward to the fuel plume. Further downstream, other ignition spots occur where the injectors' bow shocks meet between the fuel plumes, and at the top of the central plume as the recompression shock formed from the combustor-step expansion fan passes through the plume-air interface.

These three ignition points rapidly spread across the fuel-air interface, establishing the previously mentioned turbulent diffusion flame sheet, which is distinct from the OH-radical-rich region established along the bodyside surface by inlet-injected fuel. For the entire length of the constant-area combustor, these two regions of burning fuel remain distinct, confirming that different supersonic combustion modes can oc-

cur within the same engine. The two zones appear to merge into a single large reaction zone over the length of the diverging combustor, coinciding with the eventual disappearance of the oxygen-rich, H_2O-poor zone along the combustor sidewall.

While the flow physics of the first tailored case were complex, the global combustion efficiency of the $\phi = 0.98$ flow was estimated to be 68.8% at the exit of the engine nozzle. This is not sufficient to justify replacing the baseline boundary-layer injection configuration presented in Section 5.3.4. As such, a second tailored injection scheme has been developed.

6.3 Case Study 2: Cowl-side and sidewall portholes

The second configuration of tailored injection developed employs fuel injection from the sidewalls of the isolator. This design took into account the fact that the cowlside injection scheme presented in the previous section showed an oxygen-rich zone moving along the sidewall, which never interacted with the fuel injected from the cowlside wall. Injecting along the sidewall of the isolator not only accesses this zone directly, it also has the potential to allow the flame from the inlet-injected fuel to propagate easily around the circumference of the engine from the bodyside combustion zone established upstream. This flame propagation was one of the advantages identified in the baseline injection scheme presented in Chapter 5.

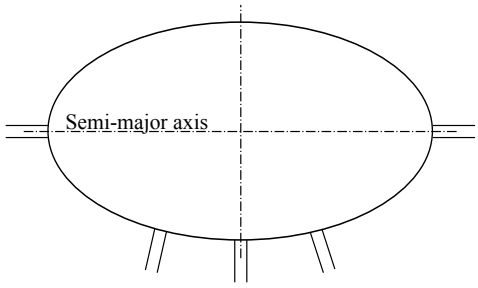

Figure 6.8: In-plane position of injectors in Case Study 2

The injector configuration for this case is shown in Figure 6.8. To take advantage of the counter-rotating vortex pair at the cowl centerline, as was discussed in Section 6.1.2, an injector was positioned at this point. The off-center injectors on the cowlside are in the same position as the outer injectors in the first tailored case, in an attempt to evenly distribute cowlside injected fuel throughout the high-density cowlside air flow. The cowlside injectors were placed such that one was centered upon the engine centreline, while the other two cowlside injectors were placed 6.75 mm (center-to-center) to either side of the central jet. The sidewall injectors were placed on either side of the isolator,

with their cross-plane axes lying on the semi-major axis of the isolator's elliptical cross-section. Moving from four identical injectors equally spaced along the cowlside wall, to five equal-area injectors required that the new injectors have a diameter of 0.8 mm. This keeps the mass flow ratio between the inlet and combustor approximately constant. As with the first tailored case, the injectors were normal to the wall in the cross-stream plane, and angled 45°toward the downstream direction in the streamwise (x-z) plane.

6.3.1 Solution Details

Much like the first tailored injection simulation, the main difference between the computational grid for this simulation and the baseline combustor mesh presented in Chapter 3 was in the regions just upstream and downstream of the combustor step. In this case, the mesh was refined in the last 25 mm of the isolator and the first 40 mm of the combustor, in order to provide sufficient resolution to resolve the injector plumes, as shown in Figure 6.9.

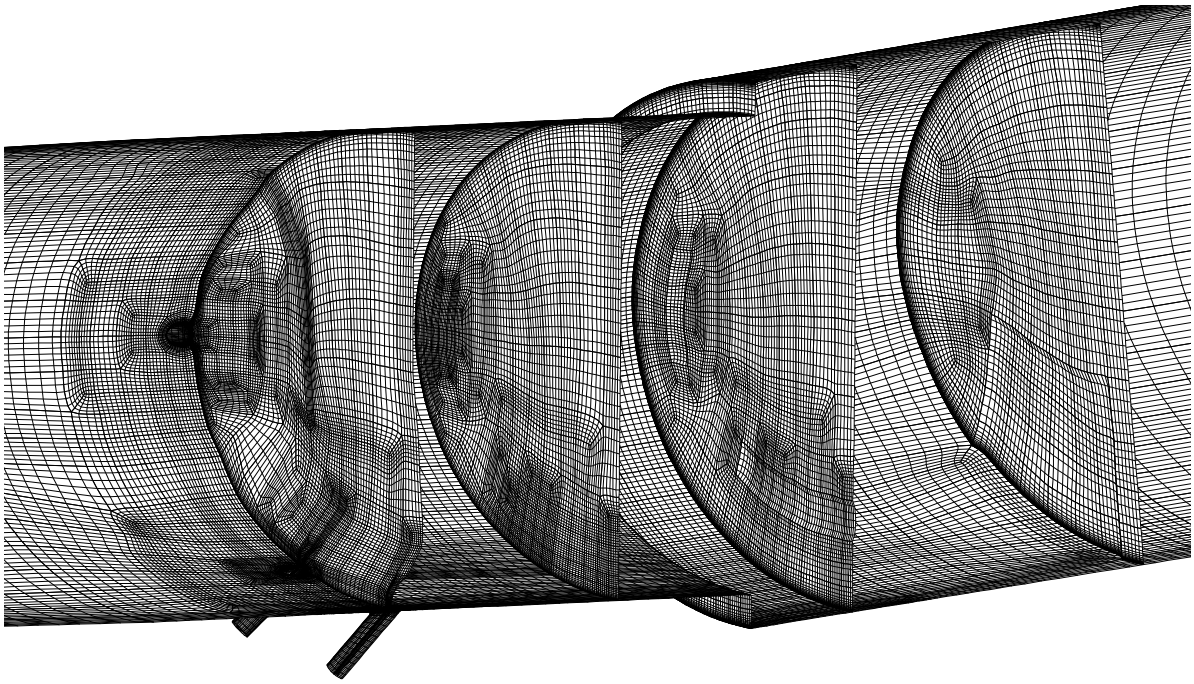

Figure 6.9: The computational mesh surrounding the combustor step in Case Study 2

Similar to previous meshes, the wall spacing in the isolator (and therefore the injectors) remained unchanged at 0.9 μm. The combustor wall spacing profile remained identical to the original combustor mesh, with cell heights decreasing from 0.9 μm to 0.75 μm in the middle of the constant-area combustor. Unlike Case Study 1, this configuration showed significant promise for improvements to the combustion efficiency.

Therefore, the simulation was run to convergence with second-order accuracy, and the injector geometry was built and tested in the T4 Stalker Tube to validate the simulations presented below.

The simulation was run over a total of 84,000 iterations, accounting for an in-simulation time of 1.96×10^{-3} s. When the simulation had converged, the RMS and energy residuals had dropped by 6 orders of magnitude, the solution domain mass flow balance was 0.55×10^{-6} kg/s, and the momentum balance was less than 2.5×10^{-3} N in all three cartesian coordinate directions. The inflow and fuel mass flow rate conditions for this simulation were chosen to match Shot 11491, which can be found in Appendix E. The total mass capture of the engine was 59.82 g/s, which was less than 1% larger than the baseline case's mass capture in Section 5.3. The similar mass captures of the baseline and tailored cases suggest that a comparison between the two should provide an accurate measure of how the performance changed. The total fuel mass flow rate was 1.75002 g/s with an inlet-to-combustor fuel ratio of 29:71; this agreed with the experimental values for fuelling to within 0.02%. In this case study, the engine had a global equivalence ratio of 1.24, based on the inflow mass fractions of molecular and atomic oxygen.

6.3.2 The Combustor Flowfield

The examination of Case Study 2 begins with the shock structures and pressure field produced by the engine with this fuelling configuration, and how they compare with the shock structures and pressure distribution in the baseline case. These structures are plotted as Figure 6.10. While the shock structures in the baseline case shown in Figure 6.10b appear stronger than those in Case Study 2, this was because the shock variable contours in Figure 6.10a were adjusted to make the multiple combustor and isolator shock structures more distinct.

The cowl closure shock generated in the M12 REST inlet in Case Study 2 reflects identically to the shock that was generated in the baseline case. Both reflect off the isolator bodyside wall just upstream of the combustor step (B2), compressing the flow in the isolator just before it enters the combustor. In Case Study 2, the jet bow shocks (JB) and jet wake recompression shocks (JR) appear along the cowlside wall of the isolator between the injector and the combustor step. The jet bow and recompression shocks merge downstream of the step. The combined jet shocks reflect off the bodyside wall somewhere in the vicinity of label BX. Where exactly they reflect is difficult to determine, as the shock density gradients used to create the visualization are weakened

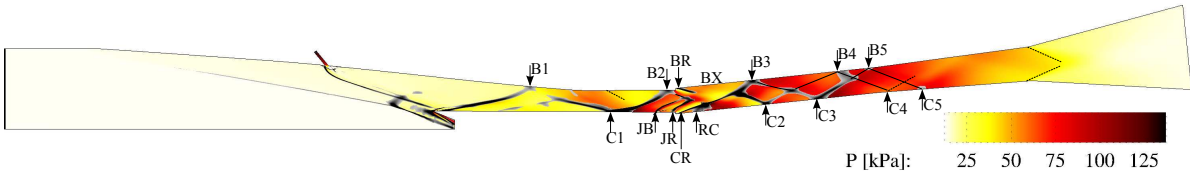

(a) Case Study 2 Symmetry Plane

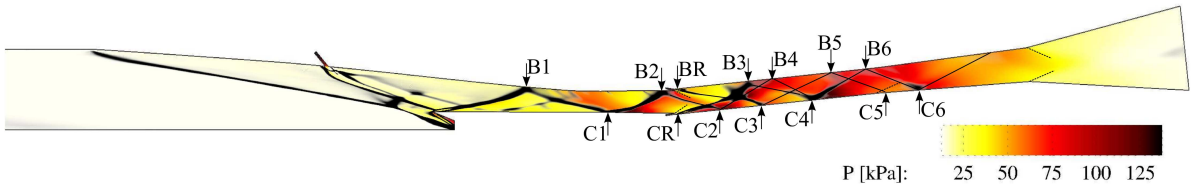

(b) Baseline Symmetry Plane

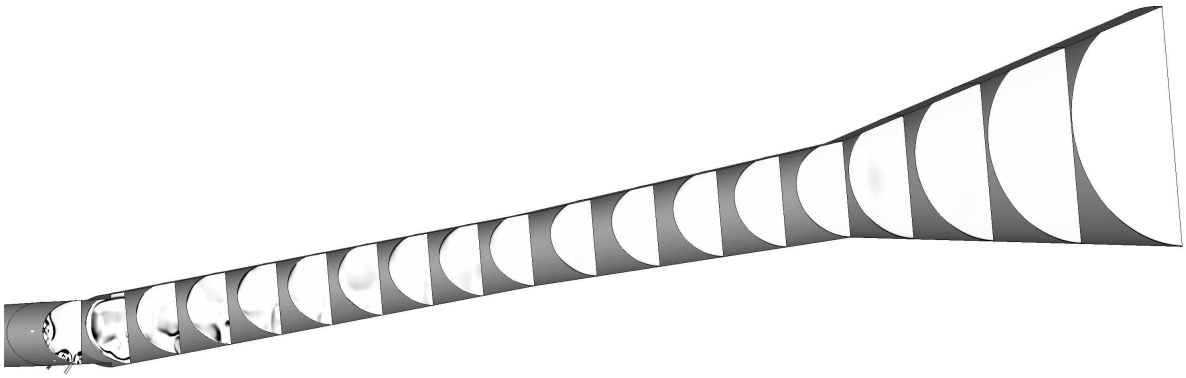

(c) Case Study 2 Cross-plane shock structure

Figure 6.10: Centerline and cross-plane shock structures for case study 2, contrasted with the baseline injection case

as the shocks pass through the low-pressure region along the bodyside wall, to the point where they disappear from view in the figure.

Along the cowlside, an additional shock structure (RC) is observed. This structure is a reflection of the laterally travelling components of the sidewall and off-center cowl-side jet bow shocks. These waves are only weakly visible in the cross-plane contours visible in Figure 6.10c, as the flow in these regions are strongly expanded by the flow passing over the combustor step.

The bodyside shocks that reflect at BX seem to pass through the converged re-compression shock, and merge with the portion of this self-reflecting shock that travels downward toward the cowlside surface, reflecting at C2. Both the bodyside and cowlside components of the recompression and transmitted shocks continue to reflect

through the combustor downstream, and never merge into a single shock structure.

The shock structures become less visible with downstream distance. This is due to the higher temperatures and corresponding lower Mach numbers expected due to the combustion of injected fuel. This weakens the shock gradients sufficiently that the contours eventually disappear in the diverging combustor section. The effects of this fuelling configuration and its associated shock structures on the combustor flow field are confirmed and quantified by examining the contours of Mach number, temperature, and density presented in Figure 6.11.

The Mach number contours in Figure 6.11a show that small areas of the incoming Mach 5 flow entering the combustor along the cowlside of the engine are slowed by the fuel jet plumes, with the sudden expansion of the combustor step allowing some of the high-Mach-number flow to migrate into the center of the combustor. Jet penetration was fairly limited, with the low Mach number region associated with the centerline jet seeming to rise to approximately one third of the constant-area combustor height. Given the small diameter of these jets, the low penetration is not terribly surprising. Despite showing penetration equal to the centerline jet in the region just downstream of injection, the off-center cowlside jets end up staying closer to the cowlside wall. Little can be said about the evolution of the sidewall jets from the Mach number contours alone.

As the flow continues downstream, the high-Mach number region shrinks considerably as the temperature rises across the combustor. By the end of the constant-area combustor, much of the flow has a Mach number ranging from 1 to 3. This distribution in Mach number is far less uniform than that of the first tailored injection case, suggesting a more complex combusting flow field. Outside of the tiny recirculation regions behind the backward step, the entire combustor flow remains supersonic. The flow expansion by the nozzle is not uniform, with the high-speed core exiting the nozzle offset toward the cowlside wall.

In Figure 6.11b, the density contours show a far more complex flow field in the near-field wake of the injectors. The presence of the laterally travelling sidewall injector bow shock is plainly visible in the plane just downstream of injection, and appears to play a role in the formation of a high-density region just above the cowlside fuel plumes downstream of the combustor step. Also visible in the second and third slices of the constant-area combustor are regions of high density indicative of a strong shock interaction with the fuel plumes. This shock is the result of the interaction of the sidewall injector bow shock with the recompression of the combustor step's expansion fan.

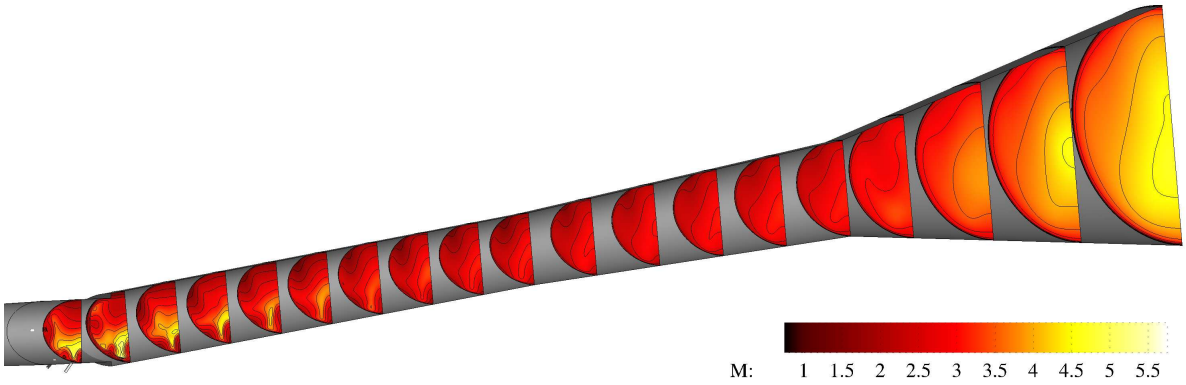

(a) Mach Number

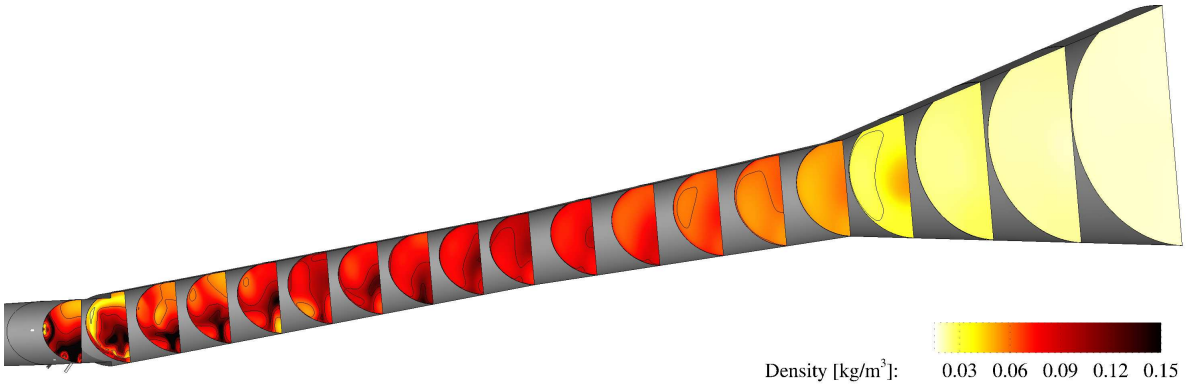

(b) Density

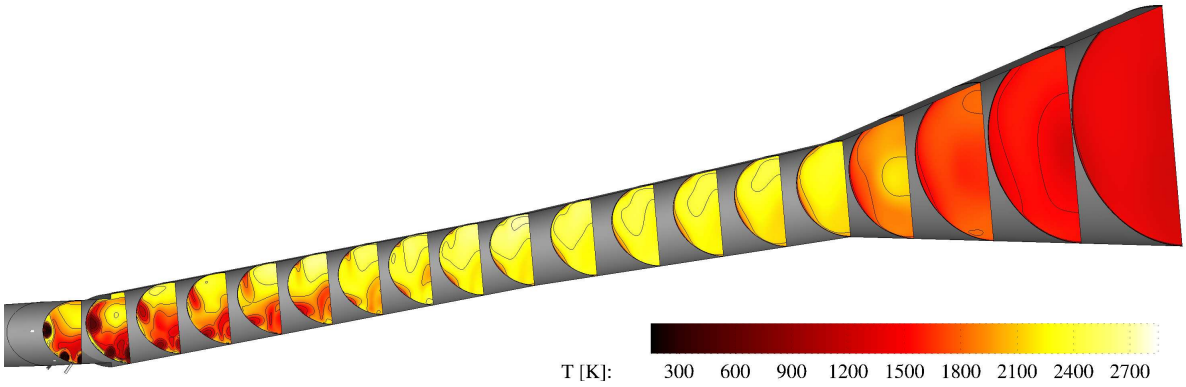

(c) Static Temperature

Figure 6.11: Key variable contours for tailored injection case study 2

The temperature contours in Figure 6.11c show the evolution of the fuel jets in some detail. It is obvious, for instance, that the sidewall jet is injected directly into the middle of the high-temperature swept separation, and the jet plume depicted in the first slice downstream of the combustor step does not cool to the same extent the cowlside jets do. The overall temperature in the cowlside fuel plume region appears to recover faster than that of the injected fuel in Case Study 1. The high-temperature zone near the bodyside wall at the end of the isolator continued along the combustor bodyside wall, and this region of the flow remained the hottest flow region in every contour slice downstream.

The second and third constant-area contour slices show the temperature increasing along the edge of the cool, fuel-rich sidewall jet region, suggesting that ignition of the sidewall injected fuel commenced quickly. This fuel-laden region travelling downstream along the sidewall is also swept upward and inward along the combustor wall, and is essentially indistinguishable from the surrounding flow by the fifth constant-area contour slice.

As the flow progressed further downstream, the temperatures along the cowlside wall rise to 2400 K in the last three constant-area slices, suggesting that robust combustion of the cowlside jets does not commence until this point. Aside from a small band of flow that appears to be associated with the off-centre cowlside jets, the entire combustor flow remains at these high temperatures until the end of the diverging combustor section. This lower-temperature zone still retains a temperature of 2000K, much higher than the cold-temperature zone in Case Study 1, suggesting that heat release from combustion was occurring across nearly the entire diverging combustor cross-section.

6.3.3 Flow Chemistry and Combustion

Both the Mach number and temperature contours give some clue as to the extent of jet penetration into the mainstream and the presence of combustion, but none of the contours presented thus far can shed light on the ultimate fate of the tailored fuel jets. Much as was done in Case Study 1, the attention of the study thus turns to the behaviour of molecular and radical species mass fractions important to the understanding of hydrogen combustion shown in Figure 6.12.

The contours of unreacted molecular oxygen clearly show the extent of each fuel jet's interaction with the mainstream flow. As is visible in Figure 6.12a, each fuel plume remains distinct for the first half of the constant-area combustor. Downstream of this

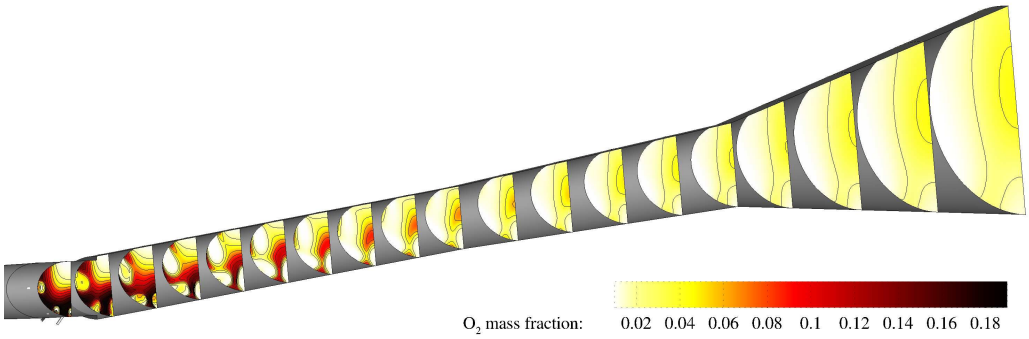

(a) Unreacted Oxygen

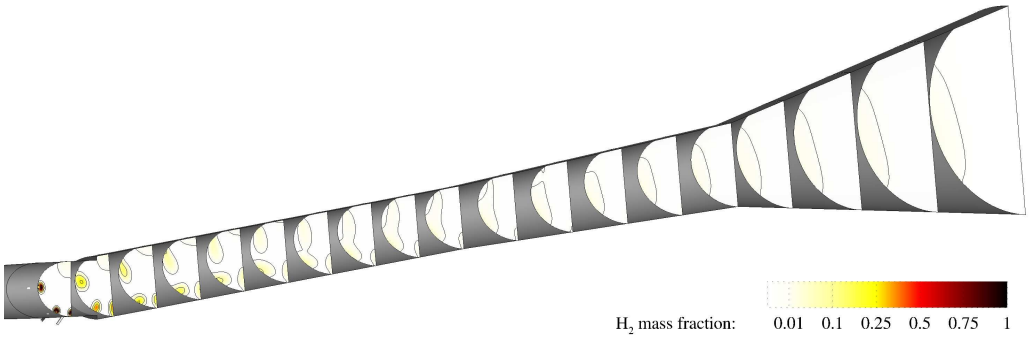

(b) Unreacted Hydrogen

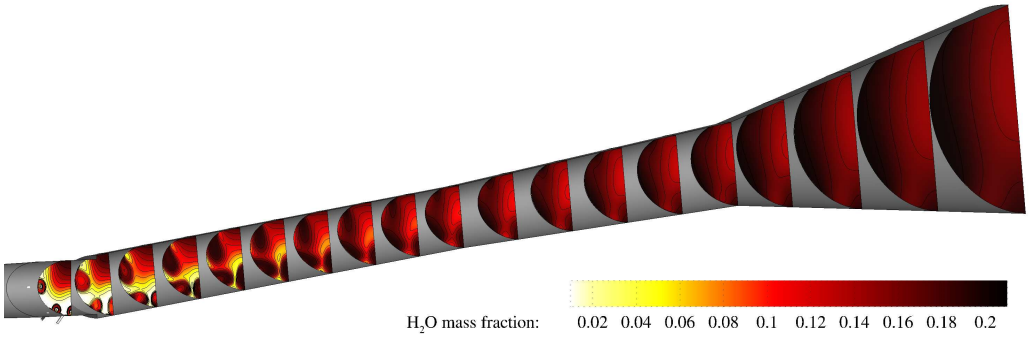

(c) H_2O mass fraction

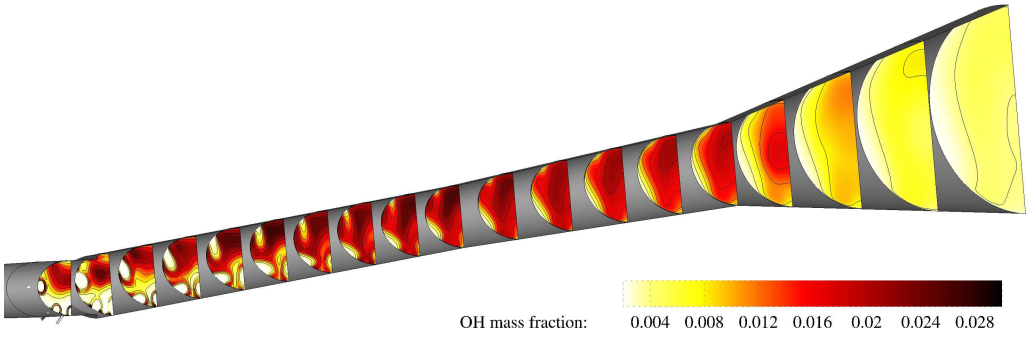

(d) OH mass fraction

Figure 6.12: Species mass fractions indicative of combustion in Case Study 2

point, the sidewall and off-center cowlside fuel structures merge. By this point, neither plume contains any appreciable amount of oxygen, and the combined fuel structure consumes much of the remaining oxygen by the start of the diverging combustor duct. The remaining free oxygen ends up confined to a vertical band running from the centerline of the bodyside surface to a point on the cowlside surface roughly in-line with the off-center injector upstream. This band forms due to a combination of air being forced upward by the cowlside fuel injection process and circulation within the combustor, with the resulting centralized structure linking to zones of under-fuelled air between the off-center and centerline cowlside injector plumes.

The hydrogen contours in Figure 6.12b confirm much of what the oxygen contours suggest. The sidewall jets injected into the swept separation show both higher penetration and lateral spreading than any of the fuel plumes in Case Study 1. The cowlside centerline jet also shows greater penetration, as intended. As such, the overall fuel distribution is far more uniform by the end of the combustor than was achieved by the injection in Case Study 1. These improvements to the delivery of fuel show great promise for improving the combustion efficiency of the engine.

The pressure gradients inside the combustor have a tendency to push the off-center plume toward the sidewall, causing this plume to merge with the sidewall plume in the last third of the constant-area combustor in Figure 6.12b. This merged region of unconsumed hydrogen is eventually confined to a zone moving downstream along the sidewall of the combustor that aligns with the low-temperature zone remarked on in the previous section. This suggests that while fuel distribution was improved overall, the mixing of fuel from the off-center cowlside jets still requires further improvement. Combined with knowledge of where unconsumed oxygen remains, it would appear effort will be needed to keep fuel injected off-center on the cowlside of the engine in the same region downstream. One solution could be to space the off-center injectors closer to the centerline injector. Another solution would be to replace the three cowlside injectors with a single injector at the cowlside centerline. As will be shown in the next chapter, such a jet would get better penetration, and should show better lateral spreading, promoting further mixing. This increased mixing due to improved penetration would likely occur due to the oxygen currently displaced upward by the current jets being pushed sufficiently further upward that it would be forced to turn outward and downward, ultimately coming back into contact with the fuel plume.

The water vapour contours presented in Figure 6.12c suggest that each fuel plume ignites independently from the the others: each plume begins to show a high H_2O mass fraction well before any of the structures appear to come in close contact. Much of

this water vapour ends up migrating toward the combustor sidewall, suggesting that the combustion in the last half of the diverging combustor duct is a strongly mixing-limited combustion process in which each jet plume convects its own combustion products downstream.

These findings are well-supported by the OH mass fraction contours shown in Figure 6.12d. The low-OH mass fraction region associated with the sidewall injector plume quickly comes in contact with the OH-rich, high-temperature bodyside flow, and shows strong indications of burning just downstream of the combustor step. The cowlside fuel plumes, on the other hand, show ignition occuring along the combustor wall at the edges of each plume. These locations correspond to the high-density, high temperature (and therefore high pressure) regions induced by the complex shock structures downstream of the combustor step.

In contrast to the behaviour of Case Study 1, the OH-rich region along the bodyside wall does not remain isolated from the cowlside region for long. By the time flow reaches the third constant-area combustor contour slice, the OH-rich flame regions of the sidewall and off-center cowlside injector plumes have begun to interact. This suggests that the bodyside flame acts to strengthen the combustion of the sidewall and off-center fuel jets.

The cowlside centerline jet's flame front appears to remain independent of the other flames until the last third of the constant area combustor. At this point, a high OH mass fraction region extends across the combustor from the bodyside of the combustor down to the low-temperature, hydrogen-rich structure that has formed along the sidewall. From this point downstream, the mass fraction of OH is high across the majority of the combustor, further indicating that combustion reactions are taking place in a mixing-limited fashion across nearly the entire combustor cross-section.

6.3.4 Comparison with Experiment

The simulated performance of this tailored injection configuration was found to show sufficient performance improvement that the injector configuration was manufactured and tested, as was described in Chapter 3. This permits the simulation results presented above to be validated against experimental data. As mentioned previously, the simulation results presented are for an inflow condition identical to the calculated nozzle exit conditions of shot 11491. The pressure data from this shot are compared to the calculated pressure contours in Figure 6.13.

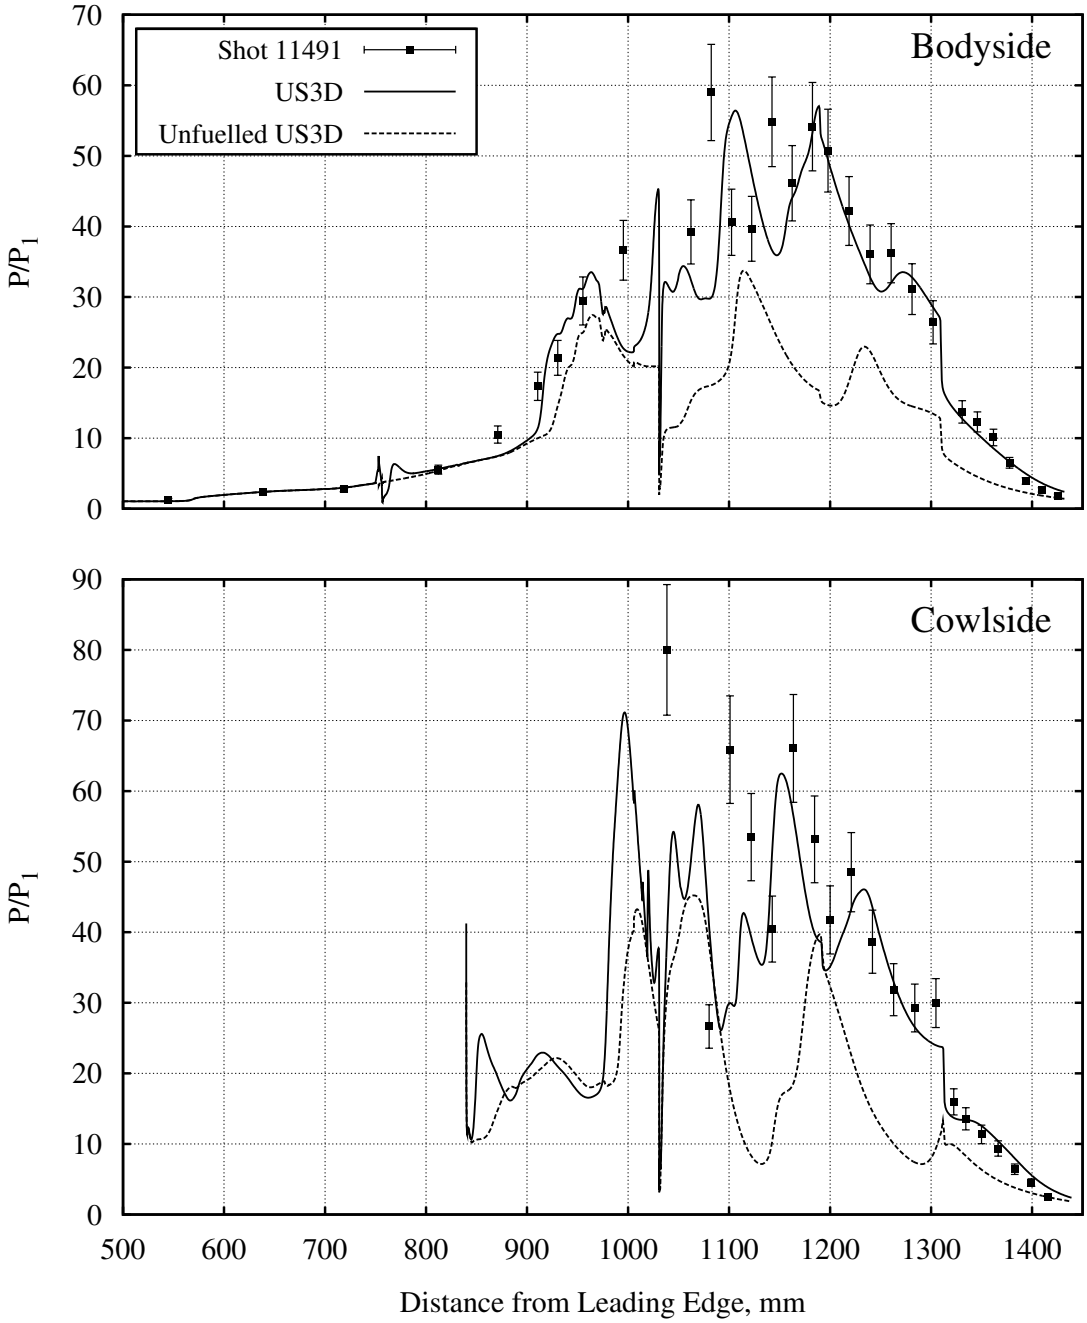

Figure 6.13: Comparison of Simulated and Experimental Centerline Pressure Distributions

It is immediately evident that there is strong agreement between the simulation and the experimental pressure data. As in previous simulations, the pressure rise due to fuel injection and reaction near the throat continue to be captured accurately by the simulation. The CFD pressure distribution matches the data strongly in the nozzle. The agreement in pressure rise location is also generally good on both the bodyside and cowlside surfaces of the engine, particularly in the diverging section of the combustor. Some of this good agreement in shock-related pressure rise (relative to previously made comparisons) may be due to the experimental test-time having been taken slightly later to minimize unsteadiness in the downstream portions of the engine due to engine start-up time.

As with nearly all previous comparisons to experimental data, the simulated peak pressure magnitudes on both the bodyside and cowlside walls generally match the measured pressure peaks. There remain some differences in the peak locations, indicating a difference in the predicted and experimental shock impingement locations. The main areas of divergence between simulation and experiment are confined primarily to the isolator and constant-area combustor, where the greatest divergence is due to the limitations in comparing steady-state RANS modelled flow with the time-averaged data of the experiment. Nevertheless, there is overall good agreement between simulated and experimental pressure distributions in the engine.

While good agreement between the simulated and experimental heat flux distributions is not expected, the comparison is nonetheless made in order to obtain a reasonable determination of the accuracy in wall heat transfer and drag in the simulations when the injection and combustion of fuel is included. The results of this comparison are shown in Figure 6.14. The contours of y^+ were plotted alongside the heat flux in order to quickly determine its effects.

As is the case with the pressure comparison, the first peak in the simulated heat flux along the bodyside of the combustor is offset relative to the peak in the experimental data. Similar behaviour can be seen on the cowlside, though with much higher average values of y^+, the simulated heat flux magnitudes show far less agreement. While the bodyside simulated heat flux matches the general trend of the data, the cowlside heat flux distribution shows less similarity with the experimental data.

At least part of this difference can be attributed to temperature-related effects. Unlike in the non-combusting case examined at the beginning of Chapter 5, several of the cowlside heat flux gauges registered surface temperature rises of over 200 K by the end of the experimental test time. While wall-adjacent cell height still plays a domi-

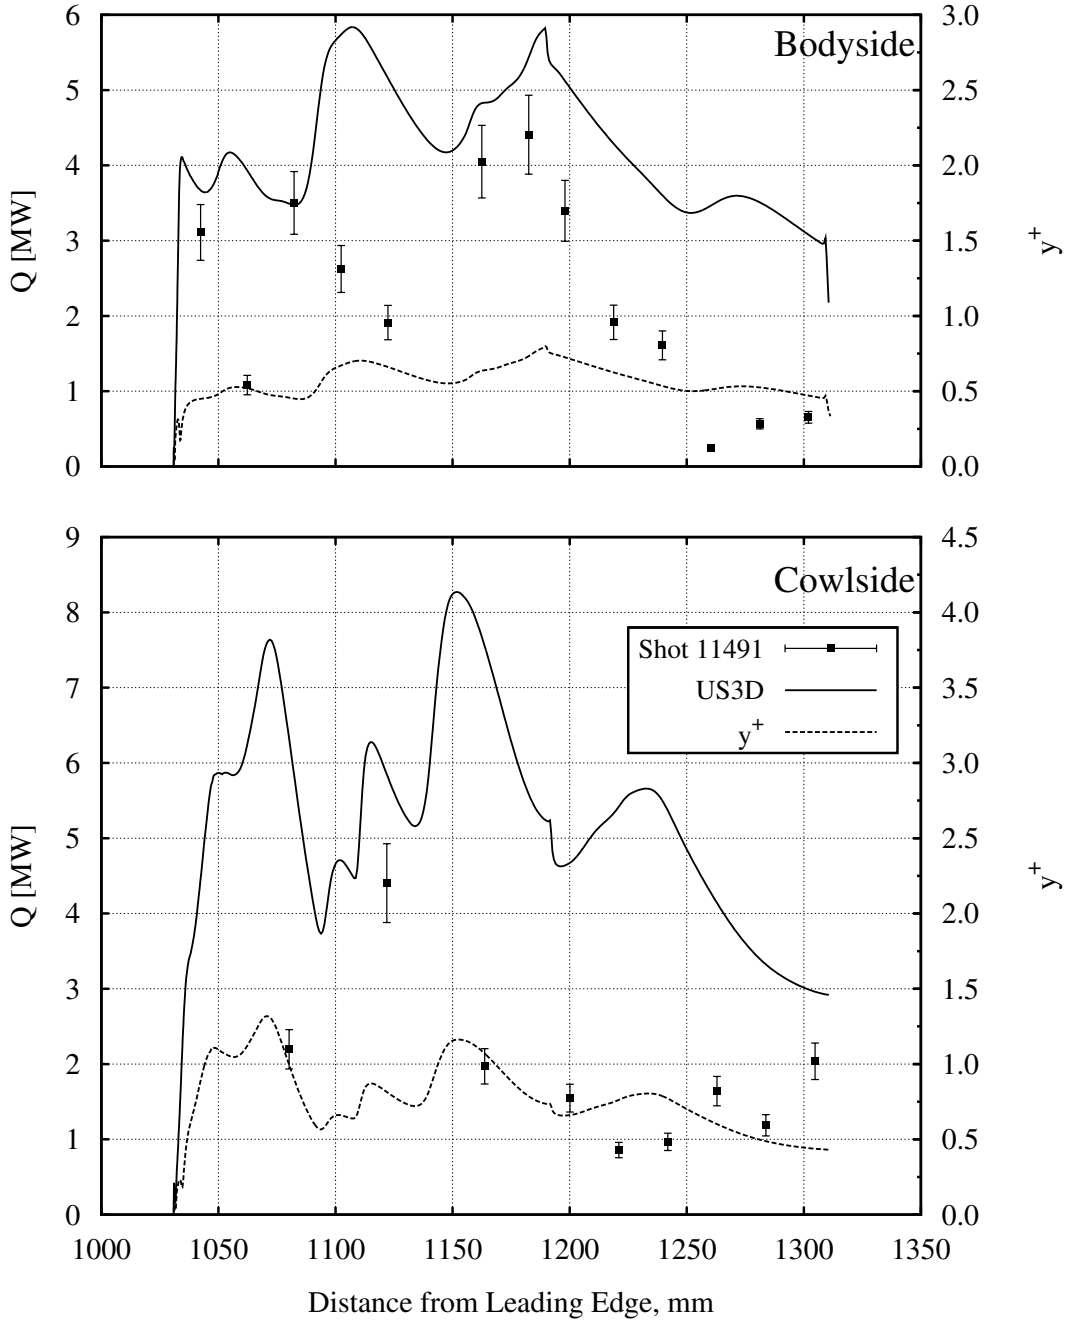

Figure 6.14: Comparison of Simulated and Experimental Combustor Wall Heat Fluxes, with y^+ distributions

nant role in the accuracy of the simulated heat flux, the temperature effect cannot be ignored. Simulating the flow through the engine with a more accurate wall temperature distribution could change the simulated heat flux distribution by as much as 10% from what is currently shown.

Overall, it can be concluded that even with a more accurate wall temperature, the simulated heat flux would still be over-predicted by the CFD simulation. This is likely a result of the inability of RANS turbulence models to accurately capture the viscous sublayer behaviour when wall temperatures are cold relative to the flow, even when y^+ is at a value of one or less. While refining the mesh to decrease y^+ further may improve the model's predictive performance somewhat, doing so could require an order of magnitude more cells in the simulation.

6.3.5 Engine Performance of Case Study 2

Having examined what the simulations indicate is physically happening in the combustor of this tailored injection configuration, the rate at which they predict injected fuel is consumed can be examined by plotting the evolution of the global mixing and combustion efficiencies of the engine with downstream distance, as demonstrated in Figure 6.15. As with the combined injection scheme examined in Chapter 5, these efficiencies were based on the mixing and consumption of available oxygen, which is the limiting reactant in this case.

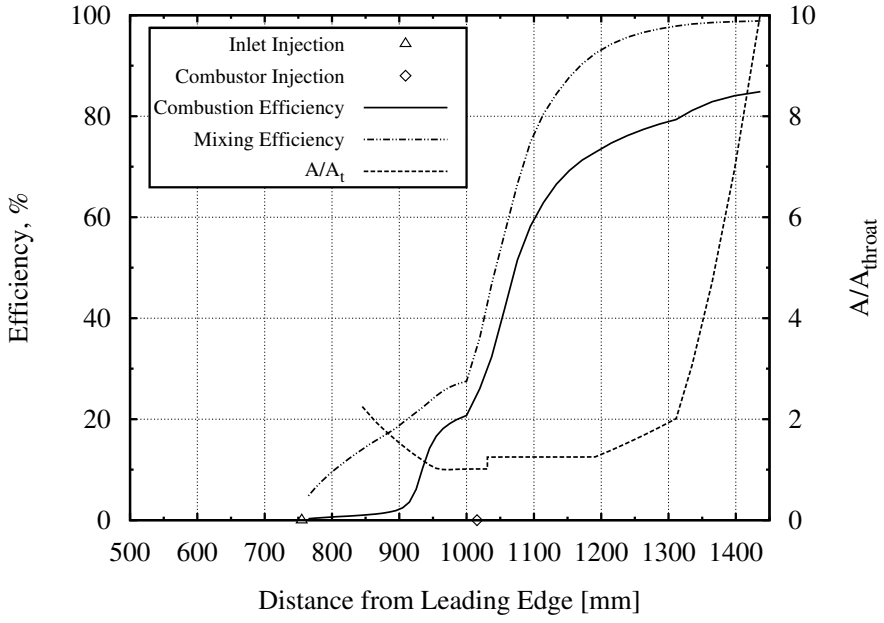

Figure 6.15: Oxygen-based mixing and combustion efficiencies for Case Study 2

The behaviour of the inlet is similar to that of other cases examined, which is of no surprise. Just upstream of the tailored injection point, both mixing and combustion efficiency begin to climb rapidly. This can be explained by the presence of the bow shock and small recirculation zones formed upstream of each injector. Air and fuel entrained in this region mix and react vigorously, thanks to a high temperature and the residence times of fuel and air entrained in the recirculation region. This trend continues into the constant area combustor, with the fuel and air sufficiently mixed to allow reactions to occur very quickly. By the middle of the constant-area combustor (1100 mm downstream of the leading edge) the combustion efficiency has risen to 60%.

Downstream of this point, the rate at which combustion is completed slows dramatically, as the easily-reacted hydrogen is consumed and the combustion process reverts to a mixing-limited process. By the end of the diverging combustor, the mixing efficiency of this $\phi = 1.24$ engine flow field is over 97%, and the combustion efficiency is approximately 79.35%. By the time the flow exits the nozzle, the combustion efficiency is 84.87%, with the mixing efficiency having reached a value of 98.87%.

As with prior calculated efficiencies, these mixing and combustion efficiencies could possibly be higher than those in the experiment due to driver gas contamination and the perfectly-stirred reactor model in the flow solver. However, the experimental data the simulation was compared with in the previous section would suggest that these level of mixing and combustion are realistic estimates. These combustion efficiencies are on the order at which Smart (2012) estimated net engine thrust may be possible for an engine operating in a flight scenario. As with the step-injection case examined in the previous chapter, it must be remembered that this is an engine operating in tunnel conditions, and as such, may have a lower efficiency than an engine in flight, thanks to heat losses to the cold engine walls in the tunnel condition. The magnitude of these heat losses will be discussed in Chapter 7.

As the goal of this study was to see if the engine's performance could be improved from its baseline configuration using tailored fuel injection, a direct comparison to the efficiency curves of the baseline case presented in Section 5.3.4 was undertaken. These comparisons are presented in Figure 6.16. Mixing efficiency in the M12 REST inlet is lower for the tailored injection configuration, as shown in Figure 6.16a. This can be explained by the slight differences in fuel split used in each case in order to match the experiments, where the inexact nature of inlet-fuelling led to shot-to-shot variation in the fuel split, as was described in Chapter 3. The mixing of the tailored case accelerates quickly just upstream of the tailored injectors, and surpasses the mixing of the baseline combustor-step injection case near the start of the combustor. By the time the flow exits

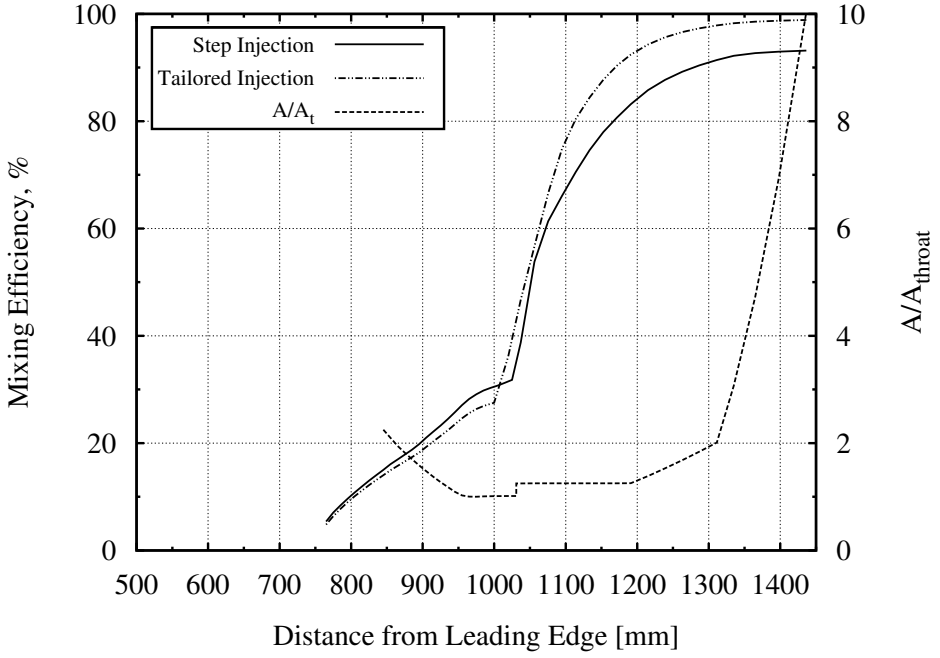

(a) Mixing Efficiency

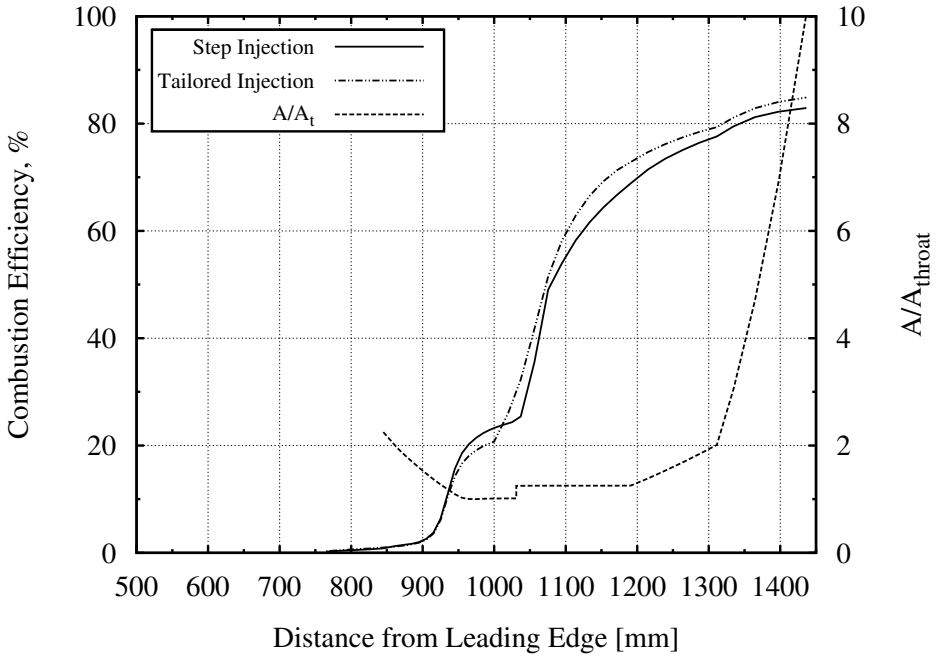

(b) Combustion Efficiency

Figure 6.16: Comparison of the oxygen-based mixing and combustion efficiencies of Case Study 2 with the baseline injection case

the nozzle, the mixing efficiency exceeds the baseline injection case mixing efficiency by nearly 7%.

The combustion efficiency curves presented in Figure 6.16b show that Case Study 2 and the baseline case both have identical combustion efficiencies along the length of the inlet until just before the flow reaches the throat, around x = 940 mm. At this point, the baseline case's combustion efficiency temporarily rises above that of Case Study 2. This is primarily an effect of fuel mass flow rates. With a higher equivalence ratio, and a larger share of fuel injected in the inlet, the baseline case's inlet fuel plume was slightly larger, and able to mix and react with more oxygen than the inlet fuelling in Case Study 2. Not far downstream, at x = 1000 mm, the flow encounters the upstream recirculation zones of the tailored injectors, and the combustion efficiency once again begins to rise. Within 15 mm, sufficient combustion has occurred to cause the combustion efficiency of Case Study 2 to rise above the baseline case, and remain higher through the rest of the engine, exiting the nozzle with a 2% greater combustion efficiency.

This is a somewhat remarkable result. Often, scramjets are over-fuelled in order to improve mixing. Case Study 2 has a significantly lower equivalence ratio than the baseline case: 1.24 versus the 1.33 equivalence ratio of the baseline case. While the baseline case was not designed for high mixing efficiency, the fact remains that when hydrogen consumptions were compared the M12 REST was able to consume 3.5% more hydrogen fuel (by mass) in Case Study 2, despite injecting fuel at a 6% lower mass flow rate.

Aerodynamic forces acting on the flow path in Case Study 2

To examine the net forces acting on the engine surfaces, the pressure and wall shear stress profiles are integrated along the surfaces of the simulated tunnel engine model. As with previous simulations, the primary interest is in the forces acting on the internal flow path, though external forces were quantified where they existed. The thrust and drag contributions found for each component are summarized in Table 6.1.

These numbers are remarkably similar to that of the baseline injection case examined at the end of the previous chapter. The total drag, thrust, and internal force estimates agree with the baseline case to within 0.6% or less. Combined thrust from the nozzle and diverging combustors are up slightly, but the isolator and constant area combustor thrust contributions have both dropped, leaving the total internal drag only 0.6% lower than the baseline configuration. While it may be confusing to see nearly-constant net forces acting in the direction of thrust when combustion efficiency was actually higher, this is actually due in part to the fuelling configuration, and the slight

Table 6.1: Aerodynamic thrust (+) and drag (-) forces acting on the M12 REST in Case Study 2

Engine Component	Inviscid Forces [N]	%	Viscous Forces [N]	%	Combined Forces [N]	%
Inlet Internal Surfaces	-12.548	23.57	-12.347	23.19	-24.895	-107.83
Forebody	-5.320	9.99	-3.238	6.08	-8.558	-37.07
Isolator and Combustor Step	0.662	2.20	-1.593	2.99	-0.931	-4.03
Constant Area Combustor	0.043	0.14	-8.240	15.48	-8.197	-35.50
Diverging Combustor	9.896	32.82	-6.181	11.61	3.715	16.09
Nozzle	19.550	64.84	-3.771	7.08	15.779	68.35
Total Drag Forces	-17.868	33.56	-35.370	66.44	-53.238 N	
Total Thrust Forces	30.151	100.0	-	-	30.151 N	
Total Internal Force	17.603 N		-32.132 N		-14.529 N	
Total Force	12.283 N		-35.37 N		-23.087 N	

differences in engine mass capture. The baseline case injected a significant amount of fuel nearly parallel to the wall, whereas in Case Study 2, a somewhat smaller amount is injected at a 45° angle to the direction of flow. Accounting for the momentum of these fuel streams, the baseline case gets a boost in thrust from the momentum of its hydrogen injectors alone (an effect primarily seen in the Isolator/Combustor step forces). If the thrust forces generated by the fuel jets in each case are neglected, the total force acting in the thrust direction in the tailored case is slightly higher. The rest of the differences in aerodynamic forces come from the differences in the inflow condition: Case Study 2 has a slightly higher mass capture, which in turn means slightly higher net drag forces and lower temperature rises due to combustion heat release.

It is also worth reiterating that there was a difference in fuel equivalence ratio: Case Study 2 had an equivalence ratio of 1.24, while the baseline case was fuelled to $\phi = 1.33$. Near-parity in thrust is achieved despite there being 6% less fuel injected. A comparison of specific impulse, I_{sp} may be more appropriate in this case. Using the net thrust forces generated (i.e. neglecting drag), and knowing that

$$I_{sp} = \frac{F_{thrust}}{\dot{m}_{fuel} g}, \tag{6.1}$$

where g is the acceleration due to gravity, Case Study 2 can be shown to have an uninstalled I_{sp} of 1756 s, 5% higher than the 1667 s calculated for the baseline case. The tailored M12 REST engine simulated in Case Study 2 therefore more efficiently produces thrust relative to the baseline M12 REST engine.

Much like the baseline M12 REST engine, the fact that these results are for an engine

flow path operating at tunnel test conditions should be repeated. As will be discussed in the Chapter 7, the cold (300 K) engine walls in this condition experience high heat fluxes, which would be reduced significantly in an engine in flight, where wall temperatures would likely be held at higher temperatures. This decrease in heat lost to the engine walls would mean a greater increase in mean flow temperature, which in turn would increase combustor and nozzle pressures along with global combustion efficiency. As such, an engine operating at a "true flight condition" would be expected to generate more thrust.

Before moving on, one final remark may be made about viscous (skin friction) drag within the combustor. While the viscous forces in the constant area and diverging combustors of the Case Study 2 engine are virtually identical to those generated by the baseline combustor, the ways in which they were generated are quite dissimilar, as shown in Figure 6.17.

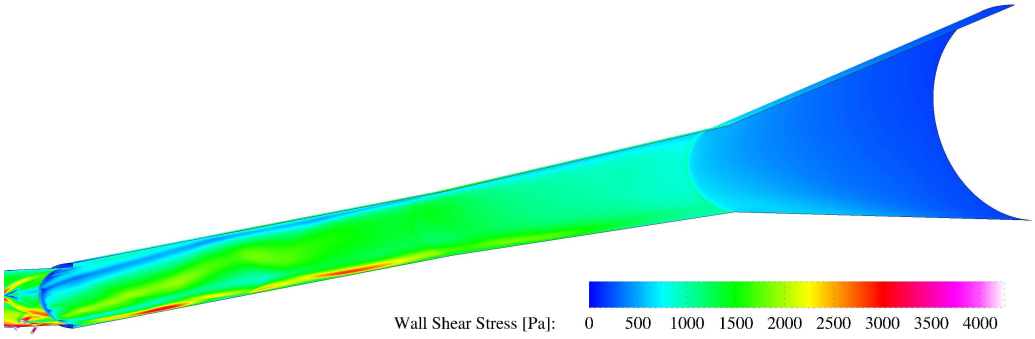

(a) Case Study 2

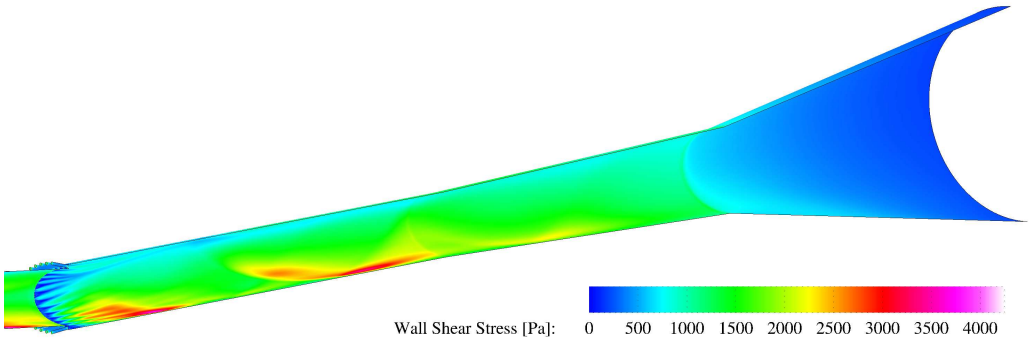

(b)

Figure 6.17: Comparison of the viscous combustor shear stress profiles in Case Study 2 and the baseline engine

While the baseline flow path had far larger high-shear regions along the cowlside of the combustor than are seen in Case Study 2, it also had much lower shear magnitudes along the top of the combustor. This was due to the large amount of cold hydro-

gen fuel injected along the bodyside wall, which provided film-cooling in the baseline case. The Case Study 2 engine, in contrast, has smaller peak shear zones, but a more even distribution of shear around the entire combustor. The only exception to this is the low-shear streak through the combustor in Figure 6.17a, which traces the path of the sidewall injector fuel plume. Clearly, despite improvements in the engine's combustion efficiency, the engine in Case Study 2 would still benefit from further design refinements to reduce the skin friction drag generated by the combustor.

6.4 Tunnel Flow Conditions Versus Flight Conditions

In all the engine flow cases presented in this chapter, and in Chapter 5 before it, there was reference to the equivalent flight conditions that each nozzle exit condition represents. It is generally assumed that by matching the flow stagnation enthalpy and post-forebody-shock static pressure of the flight condition, the measurements taken during an experiment's steady test time are representative of an identical engine configuration in flight (assuming both have cold walls). The excellent agreement of Case Study 2 with its equivalent experiment affords a unique opportunity to test this assumption on a validated computational mesh.

The already validated simulation grid was re-run, with all parameters remaining identical, save one: the inflow condition has been replaced with the equivalent flight conditions for Shot 11491, which are listed in Appendix E. The inflow condition was given the 6° angle of attack assumed in a flight vehicle. While an engine in flight would certainly have hotter walls than the 300 K temperature of the tunnel model, the wall temperature was held constant to minimize the differences between the flight and tunnel simulations being compared. The resulting simulation was run until convergence was achieved after 77000 iterations, for a total in-simulation flow time of 1.84×10^{-3} s. The RMS residual dropped by 4 orders of magnitude, as did the energy residual.

The mass flow rate balance was 1.74×10^{-10} kg/s, and the momentum balance was less than 10^{-5} N in all three coordinate directions. The fuel mass flow rates matched to within 0.005% of those in the tunnel condition simulation, but the engine mass capture of 66.56 g/s was 11% higher than that of the tunnel case. Combined with the fact that with none of the inflow oxygen was contained in NO radicals in the flight simulation, the equivalence ratio for the engine in flight is 0.90, 73% of the equivalence ratio in the tunnel-condition engine simulation.

Several comparisons can be made between the simulations of the flight condition and the tunnel condition presented in Section 6.3. The first, and simplest, comparison

is simply to compare the centerplane Mach number distributions. While the incoming Mach number was much higher for flight, it was expected that the Mach numbers should agree reasonably well once the flow had been processed by the forebody shock. This is indeed the case, as demonstrated in Figure 6.18.

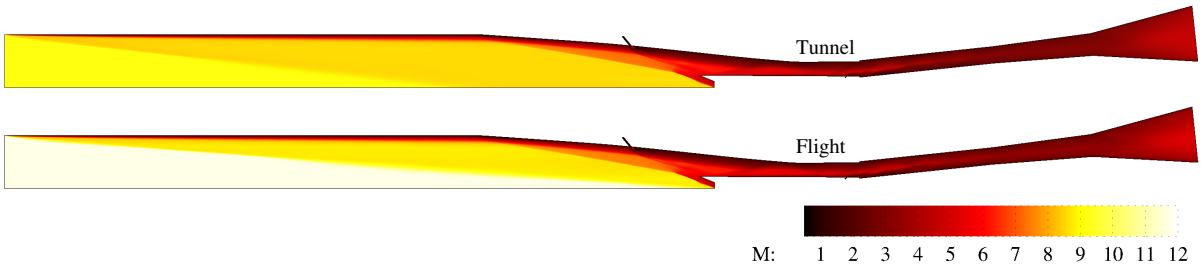

Figure 6.18: Comparison of simulated Tunnel and Flight condition centerplane Mach number contours

With the change in Mach number and angle of attack, the leading edge shock from the lip of the forebody takes on a significantly shallower angle in the Flight condition experiment. Nevertheless, the Mach number fields inside the bounds of the engine are virtually indistinguishable. The bodyside boundary layer entering the engine inlet is also nearly the same thickness in both cases. This encouraging result confirms that the tunnel conditions reconstruct the flow Mach number inside the engine to a reasonable degree. They also help explain the higher inlet mass capture in the flight condition case: the low Mach number, higher density flow approaching the throat is smaller in the flight case, and less flow must be spilled overboard to keep the engine started.

There are some differences in cross-plane Mach number contours of the combustor, as shown in Figure 6.19. The Mach number is slightly higher throughout the length of the combustor and nozzle regions in the flight condition simulation. Given that the behaviour in this region is heavily influenced by fuel combustion, it may be noted that with a higher engine mass capture, the flight condition in effect has a larger thermal sink for the heat released by reaction to be absorbed by. This in turn leads to a lower average temperature across the combustor, assuming similar rates of combustion in each case.

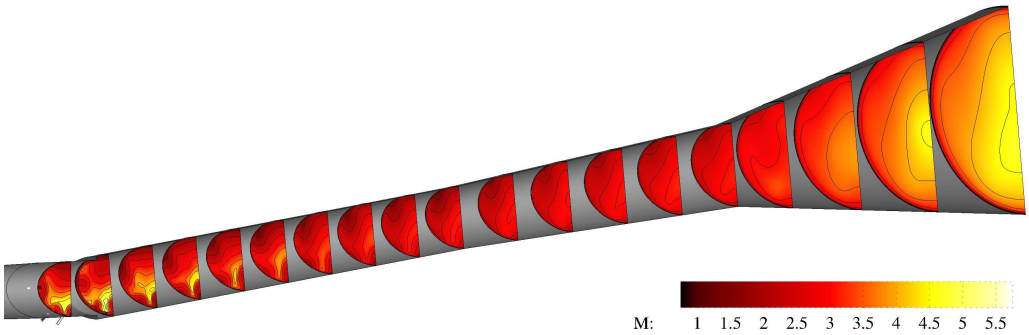

(a) Tunnel Condition

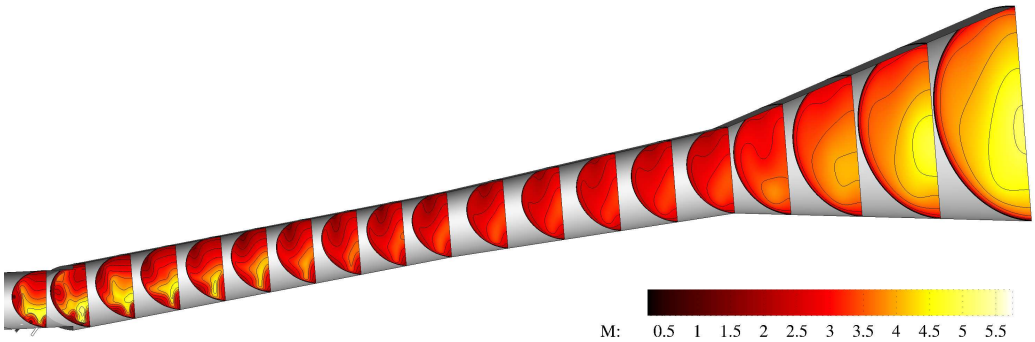

(b) Flight Condition

Figure 6.19: Comparison of Tunnel and Flight simulation combustor Mach number contours

6.4.1 Differences in Tunnel and Flight Shock Structures

Further comparison between flight and tunnel conditions can be made by comparing the shock structures inside the engine for each case. These shock structures are presented in Figure 6.20, and are complemented by the associated pressure distributions on the bodyside and cowlside surfaces, presented in Figure 6.21.

The shock structures in the inlet are virtually identical, with the bodyside reflection of the cowl closure shock (B1) occurring in the same streamwise location in both the flight and tunnel cases. The normalized pressure increases on the bodyside of the inlet show similar agreement, which provides further indication that the tunnel condition does indeed reconstruct flight engine behaviour well.

The simulations begin to show some differences downstream of the engine throat. The cowlside reflection of the cowl closure shock (C1) occurs further downstream in the flight condition than it does in the tunnel condition. This in turn leads to the shock reflecting upstream of the combustor step in the simulation of the tunnel condition, but just downstream of the step in the flight condition simulation. While this appears

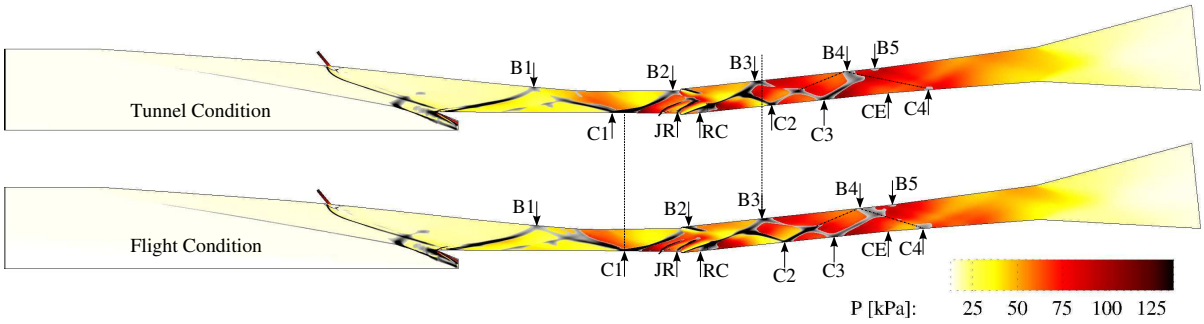

Figure 6.20: Comparison of simulated centerplane shock structure and pressure fields between the Tunnel and Flight conditions of Case Study 2

to have a negligible effect on the injector bow and wake recompression shocks, it has consequences for the pressure field inside the combustor.

A significant mismatch between the shock reflections in the combustor is visible in Figure 6.20, and the pressure peak locations of the bodyside and cowlside distributions in Figure 6.21. This mismatch is not limited to the transmitted cowl closure shock, but is also observed for the reflections of the elliptical combustor recompression shock. The difference in pressures and temperatures entering the combustor is likely the cause for this shift in shock reflection positions.

A degree of asymmetry can also seen at the end of the diverging combustor in the flight condition case, both in the centerplane pressure fields in Figure 6.20, and in the pressure distributions in Figure 6.21. This asymmetry is introduced by a late shock reflection off the bodyside of the engine just upstream of the nozzle that is damped out by combustion pressure rise in the tunnel simulation. This leads to an asymmetric pressure gradient entering the nozzle in the flight simulation.

In the higher pressure region along the bodyside of the combustor, the flow expands normally, and little difference is visible between the bodyside nozzle pressure distributions of the tunnel and flight simulations. Along the cowlside of the engine, the flow expands rapidly as it enters the nozzle, and then appears to stay at a constant pressure before expanding again later in the nozzle. There is no recirculating flow in this region, so it is likely that the cowlside flow downstream of the initial nozzle expansion fan simply experienced an isobaric volume change as the shock reflected off the bodyside wall just upstream of the nozzle passes through the nozzle expansion fan to reach the cowlside wall of the nozzle in the flight simulation.

The normalized combustor peak pressure magnitudes also differ between the flight and tunnel condition simulations, with the tunnel simulation peak pressures 15% higher

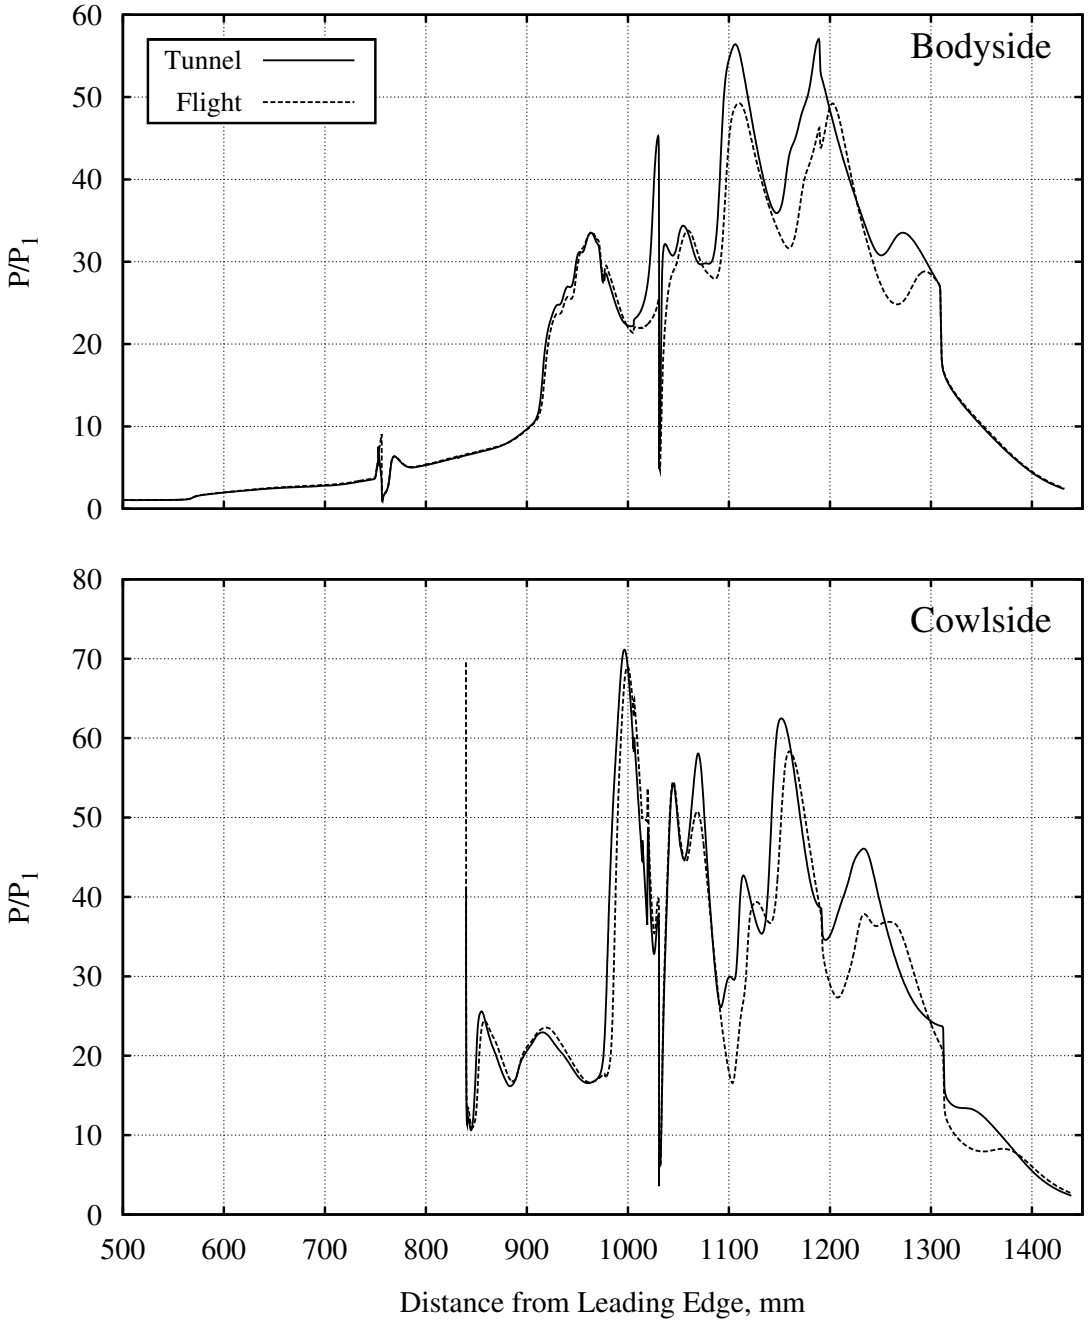

Figure 6.21: Comparison of Tunnel and Flight condition pressure distributions

along the bodyside of the combustor. The cowlside pressure distributions immediately downstream of the step match closely, which is indicative that the tunnel accurately recreates the compression effect caused by the turning of the engine to realign the flow momentum with the direction of assumed vehicle flight. Downstream of this point, however, the tunnel condition once again over-predicts the pressures in the combustor, though not to the same degree observed on the bodyside.

The differences in pressure observed are almost certainly caused by some change in the combustion process. A potential source of this change may be the presence of atomic oxygen in the inflow of the tunnel condition. It has been suggested that this additional free oxygen has the potential to significantly reduce ignition delay times in the engine (Doherty, 2013), an effect that may not be seen in the flight condition.

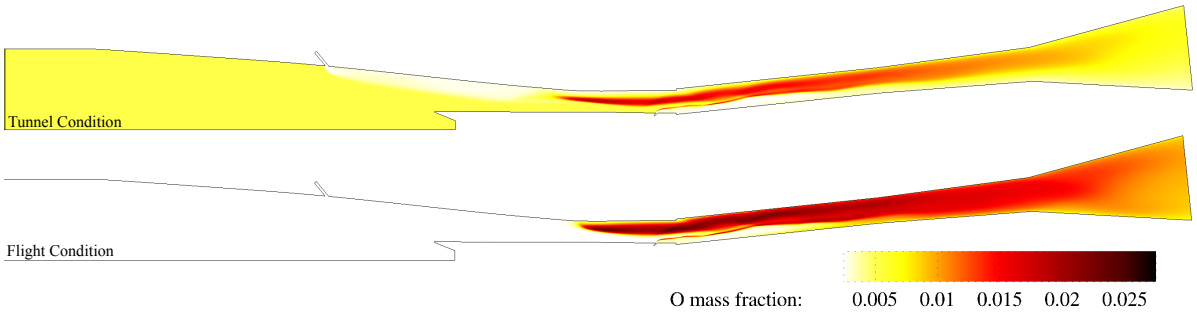

Figure 6.22: Tunnel and Flight condition symmetry plane atomic oxygen mass fractions

The distribution of atomic oxygen along the engine centerplane is compared in Figure 6.22. It is immediately apparent that a reasonable fraction of oxygen is present across the entire simulation domain upstream of the inlet in the tunnel condition simulation. It was expected that the flight condition might show some production of atomic oxygen in the hot bodyside boundary layer, but this is clearly not the case. Instead, both simulations show significant production of atomic oxygen downstream of the point at which the cowl closure shock reflects off the bodyside wall of the inlet, just upstream of the throat. While the increase in oxygen mass fraction in the flight condition simulation occurs slightly downstream of the increase observed for the tunnel condition, both produce significant quantities of atomic oxygen through the isolator and into the combustor. Combined with the similarity in inlet pressure behaviour, this indicates that ignition occurs in virtually the same location in both cases.

The regions of high oxygen production also indicate a slightly different point of initiation in combustion. In the tunnel case, the high atomic mass fraction region is displaced from the bodyside wall, whereas the flight condition shows that production

grows outward from the wall, as a result of the higher pressures and temperatures induced by the interaction of the cowl closure shock with the bodyside boundary layer. The flight condition also shows significantly higher amounts of atomic oxygen produced through the engine length, which may be expected given the much higher mass flow rate of oxygen available to dissociate in this case.

Clearly, atomic oxygen effects alone cannot explain the changes in shock structure and pressure observed so far, and there must be some effect due to the lower equivalence ratio and larger engine mass capture of the flight condition. With more bulk heat capacity to absorb the heat released by hydrogen combustion, if the amount of fuel burned is roughly the same, a lower temperature rise and corresponding pressure rise would be expected. In order to check if the amount of fuel burned in each simulation is indeed similar, it is necessary to calculate the combustion efficiency of the M12 REST engine in the flight simulation.

6.4.2 Effects of Flight Condition on Engine Performance

The examination of the changes in engine performance begins with the mixing efficiencies, presented in Figure 6.23. The efficiencies presented are both fuel-based (as this parameter was held constant between the two simulations) and oxygen-based (which was the limiting reagent in the tunnel case).

The hydrogen-based mixing efficiency is slightly higher through the inlet in the flight condition case, likely due to the higher amount of oxygen available inside the bodyside boundary layer that the inlet-injected fuel enters. The mixing behaviour exhibited in Figure 6.23a indicates that much the same mixing processes occur through the entire engine, but there is simply more oxygen available for the fuel injected in the flight to mix with. In the end, the hydrogen-based mixing efficiency in the flight condition (ϕ = 0.90) simulation is 94.9%, nearly 14% higher than the 81.1% efficiency of the tunnel (ϕ = 1.24) simulation.

The inflection in the mixing efficiency that occurs in the isolator where the incoming flow first encounters the recirculation region of the tailored injectors is further upstream in the tunnel condition simulation. Visible in both efficiency plots in Figure 6.23, this indicates that the injectors' upstream recirculation regions are approximately 5 mm larger in the tunnel simulation case. This may cause some improvement in the tunnel case's combustion efficiency, as these recirculation regions are normally a source of combustion radicals.

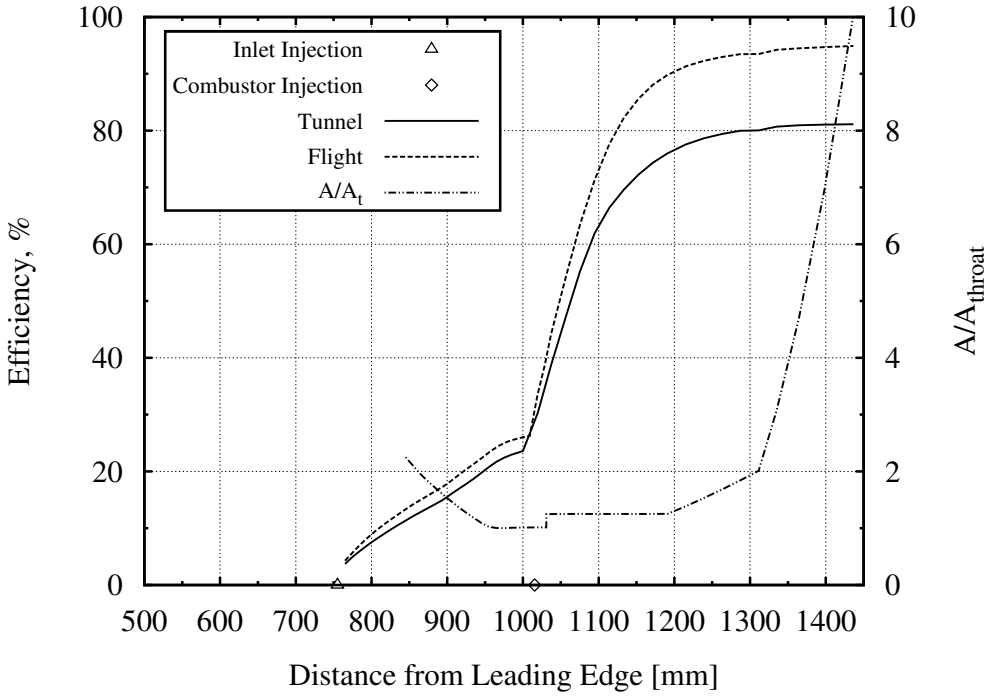

(a) Fuel-Based

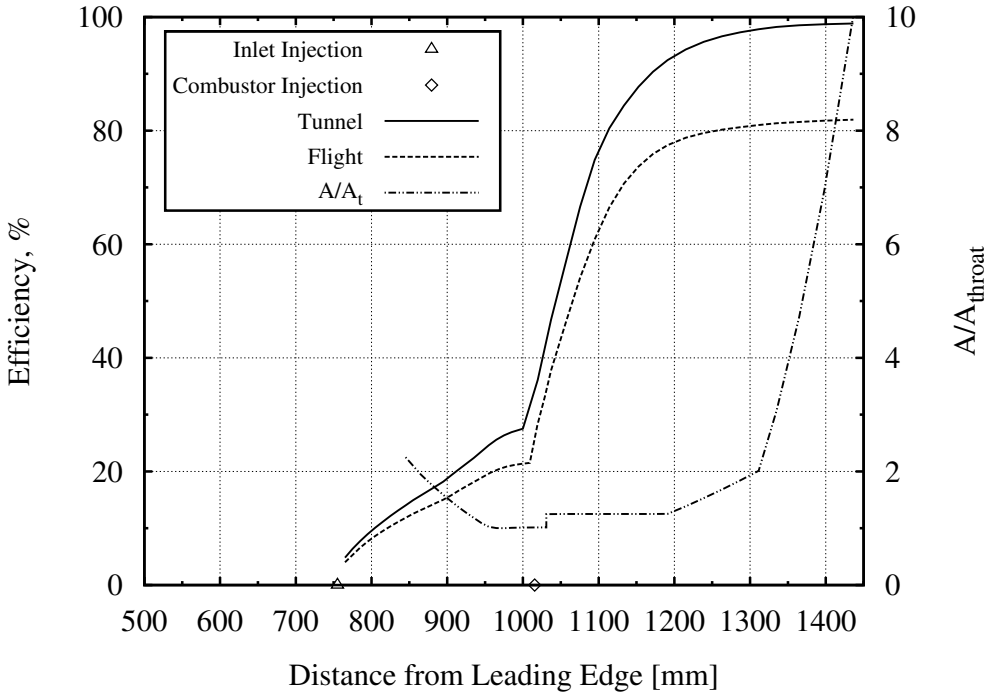

(b) Oxygen-Based

Figure 6.23: Fuel- and Oxygen-Based Mixing Efficiency in the Tunnel and Flight simulations

The oxygen-based mixing efficiency in Figure 6.23b shows the other side of the mixing difference, as the 98.87% mixing efficiency previously presented for the tunnel (ϕ = 1.24) simulation is shown to be 16% higher than the mixing efficiency of oxygen in the flight (ϕ = 1.90)case. This is entirely due to the much greater amount of oxygen available in the flight condition. It is simply not possible for the same fraction of available oxygen to mix and burn with the constant amount of fuel injected in each case.

The combustion efficiencies tell a similar story to that of the mixing efficiency. As presented in Figure 6.24, the overall amount of hydrogen consumed by the engine is higher in the flight simulation, with a fuel-based combustion efficiency of 74.3%. The oxygen-based combustion efficiency is much lower for the flight simulation, on the other hand, with a global combustion efficiency of 66.6% as the flow exits the engine nozzle. Neither result is unexpected: the engine should be able to burn more fuel when there is more oxygen available to react, but an over-supply of oxygen will act to lower the efficiency based on captured air.

Ideally, with it now known that the flight simulation has a much higher mass capture in addition to no strong NO-component of the inflow the simulation would be re-run to match the equivalence ratio of the tunnel, though time constraints prevented this. The higher mass capture has another consequence: a meaningful comparison of net forces acting on the flow path in each case is not possible.

One final comparison between the tunnel and flight condition simulations may yet be made, however. This comparison focuses on their predicted heat fluxes. While the simulations have previously proved under-resolved, and therefore over-estimate the heat flux, it is expected that the two simulations would produce similar over-predictions of heat flux relative to one another. This assumption is based on the y^+ distributions of each simulation, shown in Figure 6.25.

It is evident that the two simulations have similar ability to resolve the near-wall flow in their respective combustors. While the tunnel case may have slightly larger zones in which the y^+ drops below 0.3, the overall differences are negligible. It is therefore assumed that the two simulations have a similar magnitude of heat flux over-prediction, and it is possible to determine the relative ability of a tunnel experiment to predict in-flight combustor heating loads. The format used for this comparison is the same as was previously used to compare the simulated heat flux with experimental measurements. The two predicted heat flux distributions are shown in Figure 6.26.

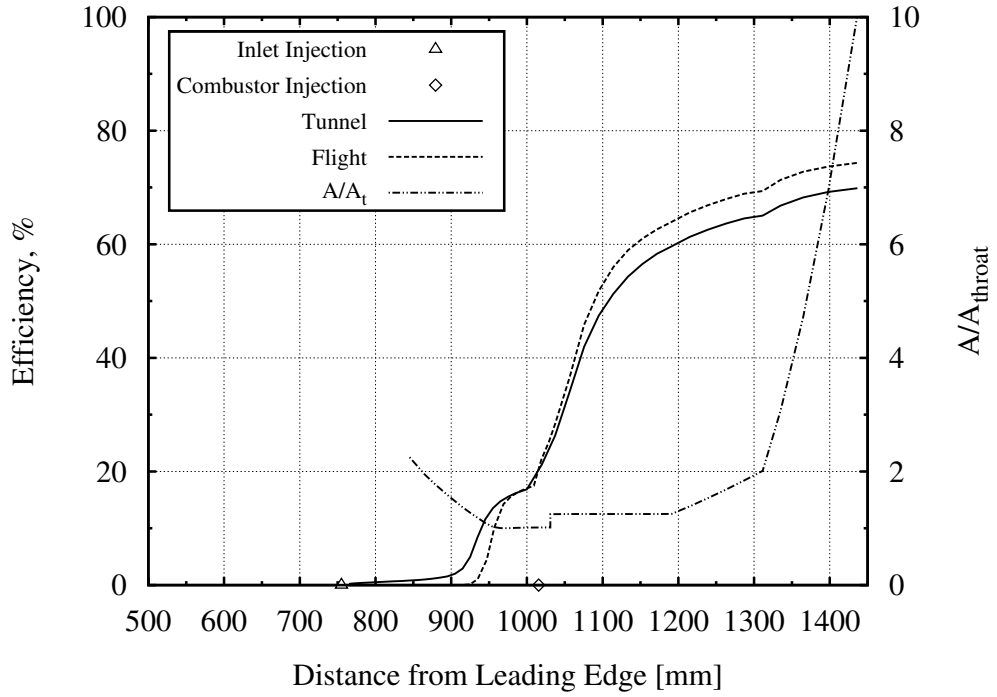

(a) Fuel-Based

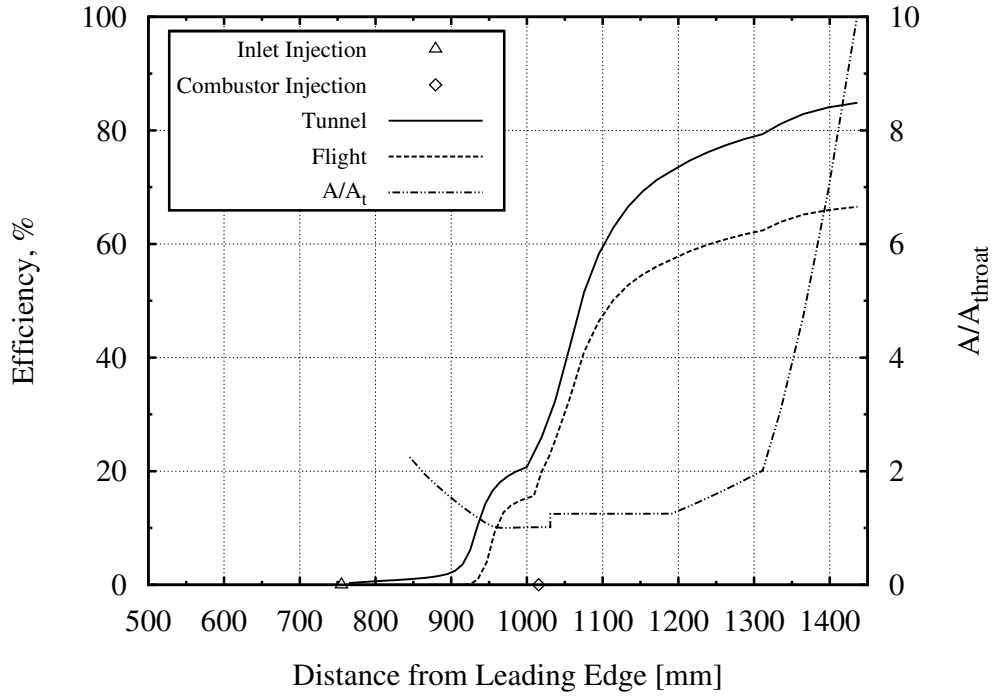

(b) Oxygen-Based

Figure 6.24: Fuel- and Oxygen-Based Combustion Efficiency in the Tunnel and Flight simulations

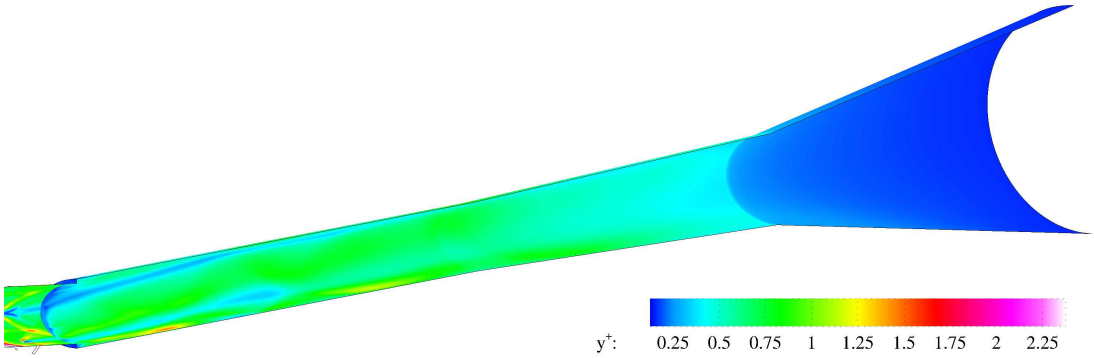

(a) Tunnel

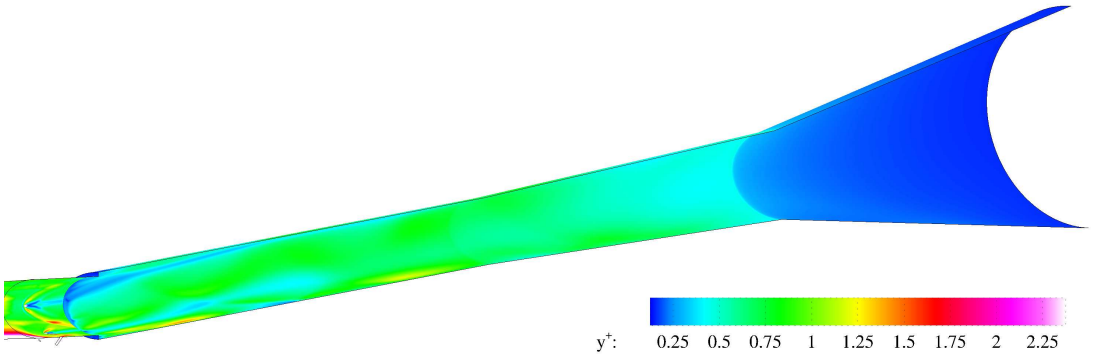

(b) Flight

Figure 6.25: Comparison of the combustor y^+ distributions in the Tunnel and Flight simulations of Case Study 2

It is readily apparent that the tunnel condition produces similar wall heat flux predictions relative to a flight condition with the same engine wall temperatures. Changes in the shock structures mean that in some parts of the combustor, heat fluxes are slightly higher in the flight condition. Conversely, these changes also created zones in which the tunnel simulation predicted a larger heat flux. Overall, it appears as though the tunnel provides a good prediction of in-flight heating loads, assuming identical fuel mass flow rates. This implies, however, that the flight heating loads would rise if the two conditions were operated at the same equivalence ratio. Intuitively, this should be expected: given the relative abundance of available oxygen in the flight case, greater levels of combustion (and therefore heat release) should occur if more fuel was injected to bring the global equivalence ratio up. However, an engine in flight would also likely have higher wall temperatures, which would act to reduce the near-wall temperature gradients, and therefore heat flux. Without running a flight simulation with this higher wall temperature, it is difficult to estimate how the tunnel condition would compare to a flight condition with hot walls and a higher equivalence ratio.

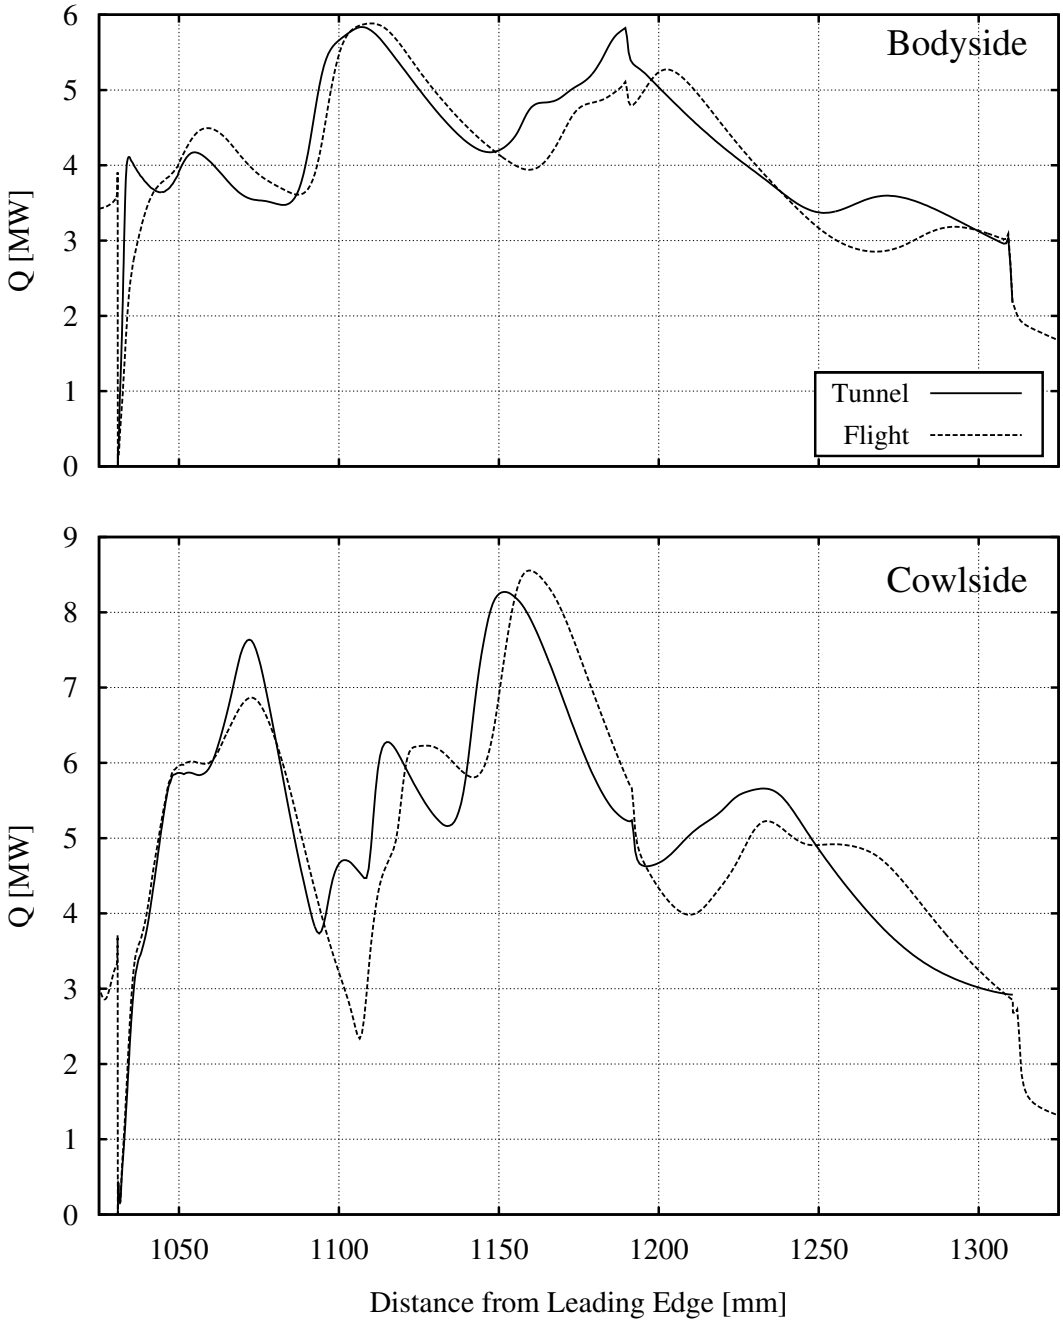

Figure 6.26: Comparison of simulated the Tunnel and Flight heat flux distributions

6.5 Summary

Fuel injection into the M12 REST engine was changed in an attempt to improve its performance. Two separate attempts were made at tailoring the injection to the flow entering the combustor, the second of which was tested experimentally to validate the simulation results. It was found that the second tailored injection case was able to improve the mixing and combustion efficiency of the engine by 5.7% and 2%, respectively. However, due to differences in the equivalence ratio, heat release, and pressure rise, no appreciable increase in thrust was found to occur. Specific impulse, on the other hand, increased by 5%. This suggests that tailored injection has indeed had a strong impact on the efficiency with which the engine produces thrust. Improvements to cowlside injection penetration could further boost engine performance.

The simulation was found to match the experimental pressure data well, providing good validation of the simulation results. Less agreement was seen with the experimental wall heat flux data. As was mentioned in the previous chapter, this was primarily due to y^+ values in the simulation being insufficient to match the magnitudes of the heat transfer. Skin friction remained the dominant source of internal engine drag, and will likely need to be addressed in order to improve the engine performance sufficiently to produce net thrust.

Finally, the tunnel simulation was compared with a simulation in which every parameter was held constant except the inflow condition, which was set to the tunnel inflow's equivalent flight condition. This was done to assess how well tunnel experiments predict the behaviour of an engine in flight. It was found that in addition to the lack of NO radicals in the inflow, the engine mass capture was 11% higher in the flight condition simulation. The free atomic oxygen in the tunnel inflow was found to have little effect on ignition location in the engine relative to the flight condition. The total mass of fuel burned in the flight condition was found to be higher, but with an overall increase in available oxygen and lower equivalence ratio, the oxygen-based combustion efficiency was lower. Wall heat flux distributions were compared, and found to be similar, despite lower relative levels of fuelling.

Having examined the potential for tailoring fuel injection to the engine flow field, there have been several important lessons learned about the relative benefits. Moving forward from this point, further changes must be explored. What some of these changes to the M12 REST might be, and a discussion of their relative benefits, is the subject of the next chapter.

7

Improving Tailored Injection and Engine Performance

"A so-called 'rocket scientist' is really a desperate engineer!"

– Ray Stalker, *Space Travel - Why is it so Expensive?* (1999)

It has become clear from the preceding chapter that tailoring fuel injection can provide some benefit to overall engine performance. However, despite having an estimated combustion efficiency of 84.9%, the tailored-injection M12 REST internal flow path as it is currently designed does not provide net internal thrust under tunnel testing flow conditions. It is therefore worthwhile to consider changes to the both the operation and design of the M12 REST engine that could make it capable of net thrust at Mach 12.

There are several ways in this can be accomplished. One option is to change the properties of the fuel injector jet by increasing plenum temperature, or by employing oxygen enrichment. The injection geometry can be modified as well drawing on the lessons learned from the results of previous chapters. Finally, modifying the geometry or operating parameters of individual engine subcomponents can be investigated, to provide recommendations on how the engine itself might be changed to gain additional performance. All the improvements examined over the course of this chapter are based on the tailored injection configuration of the M12 REST engine presented in Chapter 6.

7.1 Tailored Injection: Fuel Heating

The idea of heating fuel to improve combustion performance in a scramjet engine has arisen in no small part due to the need to actively control vehicle heating during flight

at hypersonic speeds. Using the fuel as a coolant for engine surfaces that experience significant in-flight heating has often been proposed (Segal, 2009). Such a fuel-coolant system was successfully employed by the X-51 experimental vehicle during its successful 210 s in-flight acceleration to a speed above Mach 5 (Rosenberg, 2013). Using the fuel as a heat sink would mean the temperature at which the fuel enters the combustor plenum would increase above the 300 K temperature used in both the simulations and experiments of this study. This higher fuel temperature would also reduce the time it would take for injected fuel to reach a temperature at which auto-ignition is possible, this initiating combustion further upstream in the engine. It would also have the additional benefit of improving jet penetration characteristics. As was mentioned in Section 6.3.3, improving fuel jet penetration could have a number of positive effects on the mixing and combustion of fuel injected from the cowlside wall of the M12 REST.

Recall from Section 2.4.1 of the literature review that the dominant parameter affecting jet penetration was the ratio of jet and freestream dynamic pressures. The jet dynamic pressure is defined as

$$q_{jet} = \frac{1}{2}\gamma P M^2 \tag{7.1}$$

A choked injector always has a Mach number of one, and the jet speed (U_j) is equal to the local speed of sound in the fuel (a). The speed of sound of the fuel jet increases with temperature, as by definition

$$U_j = a_t = \sqrt{\gamma R T}. \tag{7.2}$$

It should be noted that the temperature being referred to is that of the temperature at the injector throat. Assuming isentropic flow, and stagnation conditions in the plenum feeding the injector, the ratio of the injector temperature to its plenum static temperature can be calculated with the isentropic equation

$$\frac{T_j}{T_{plenum}} = \left(1 + \frac{\gamma - 1}{2} M^2\right)^{-1}, \tag{7.3}$$

which is equal to 0.833 at sonic conditions, assuming γ = 1.4. As the plenum fuel temperature rises, so must the jet temperature and speed of sound, despite the Mach number remaining static.

If the injector mass flow rate is held constant, the increase in fuel temperature must be accompanied by a decrease in the jet bulk density. Recalling that mass flow rate

through a duct of constant area (A) is defined as

$$\dot{m} = \rho A U_j = \text{ constant, where } U_j = a = \sqrt{\gamma RT} \text{ at sonic conditions,} \tag{7.4}$$

the density decrease will therefore be proportional to the square root of the temperature increase. From the ideal gas equation,

$$P = \rho RT, \tag{7.5}$$

therefore the static pressure of the jet must increase with the square root of temperature. Recalling equation 7.1, the dynamic pressure of a sonic porthole injector must also increase with fuel temperature. As the jet dynamic pressure increases, the jet penetration also increases.

In order to ascertain how much of a benefit this might have on the M12 REST engine, consider the case of fuel injection into the high-dynamic pressure region along the cowlside of the engine isolator. In the flight-condition simulation, hydrogen was injected at a rate of 0.50253 g/s on the inlet, and 1.24681 g/s from the tailored portholes just upstream of the combustor step, which gave a global equivalence ratio of 0.902. Assuming a true-flight engine would be fuelled with an equivalence ratio of unity, the same inlet-to-combustor fuelling ratio (29:71) would give stoichiometric fuelling rates of 0.55688 g/s on the inlet, and 1.38163 g/s in the combustor. Knowing that this fuel injection rate was for the simulation of half of a full engine, the amount of hydrogen injected through a single porthole in the combustor in stoichiometric fuelling conditions can be determined to be 0.553 g/s.

The jet penetration correlation developed by Portz and Segal (2006), which was presented in the literature review, can be used to determine the change in penetration of the injectors in the combustor. The cross-flow properties for the correlation are taken from the core of the high-speed flow region along the cowlside wall, just upstream of the combustor step. The freestream Mach number is taken to be equal to 5, the ratio of boundary layer thickness to jet diameter is taken to be 1.25, and the freestream dynamic pressure is approximately 800 kPa.

Jet penetration height is calculated 10 jet diameters downstream from the point of injection, for fuel plenum temperatures ranging from 300 K to 1000 K and a constant mass flow rate equal to the stoichiometric rate calculated above. The jet is assumed to be pure hydrogen with a constant γ equal to 1.4. The variation in penetration height, normalized to the penetration of the jet fed by a 300 K plenum is presented in Fig-

ure 7.1.

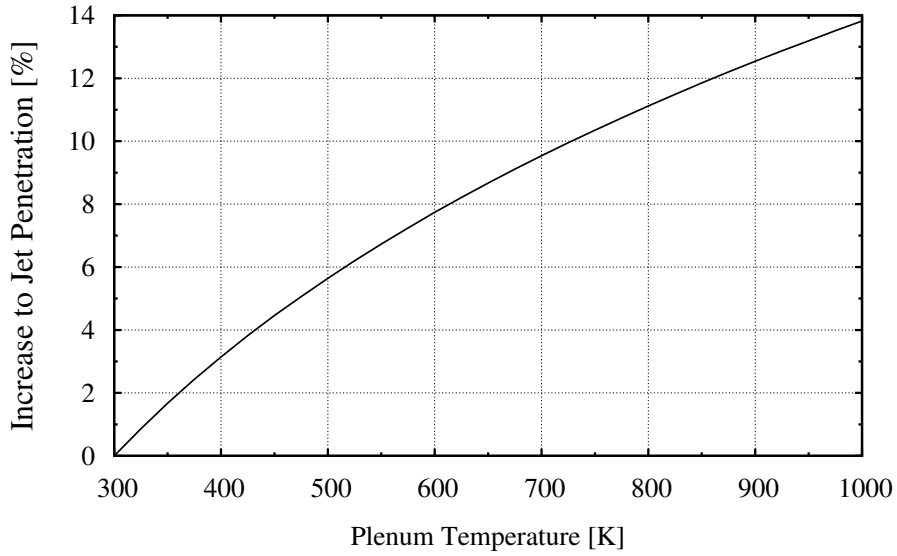

Figure 7.1: Change in jet penetration corresponding to a change in injector plenum temperature

The penetration of the jet increases by a small amount with temperature. If the fuel is heated to 1000 K, the estimated jet penetration increases by 13.8%, which would allow for a modest increase in the mixing and combustion of the jet. At a plenum temperature of 1000 K, the fuel in the injector would be roughly 833 K. The resulting expansion of the jet would cause it to cool, but not to as low a temperature as seen in the fuel plumes of the tailored injection study.

It should be noted that the correlation used was developed for normal injection into the flow, while the injectors were actually inclined 45°to the wall in the streamwise direction. However, given that injection angle is typically modelled using an effective dynamic pressure ratio (McClinton, 1972), it's likely that inclined jets would see similar levels of increased penetration with fuel heating. The tangential portion of a heated fuel jet's momentum would also serve to increase the engine's overall specific impulse, thanks to the jet's higher average speed.

As previously mentioned, heated fuel would also reduce the time it would take injected fuel to ignite and burn. Because the fuel jet would no longer cool to the same extent as it expands, it would also more quickly recover to a temperature at which it can quickly ignite and react. This change in timescale can be estimated using the relations of Pergament (1963), who used a simplified hydrogen-oxygen reaction scheme to develop correlations for ignition and reaction times in flow conditions expected inside

a scramjet combustor. These correlations are defined as

$$\tau_{id} = \frac{8 \times 10^{-9}}{P} e^{\frac{9600}{T}} \tag{7.6}$$

and

$$\tau_{react} = \frac{105 \times 10^{-6}}{P^{1.7}} e^{\frac{1.12T}{1000}}. \tag{7.7}$$

Assuming the mainstream static pressure in the region downstream of injection is constant at the average flow value of approximately 52 kPa, it is possible to use equation 7.6 to show that the fuel ignition delay time falls several orders of magnitude as the fuel warms from 300 K to 900 K. Since a heated fuel jet would need far less time downstream of injection to recover to 900 K, it would be expected that the heated jets would ignite far sooner, which could reduce the length of combustor needed to contain complete combustion.

7.1.1 How much cooling would the M12 REST need?

Knowing what approximate level of jet penetration increase in the combustor would be possible with a given fuel temperature increase, it is useful to determine what the expected heating load on the engine would be at a normal flight condition. It is assumed that only the combustor plenum would experience significant fuel heating under normal operating conditions; the inlet-injected fuel would retain its plenum temperature of 300 K . This is assumed in order to mitigate the risk of heated hydrogen igniting further upstream on the inlet, increasing drag or even potentially causing the inlet to unstart. It is assumed that it is possible to safely heat the hydrogen entering the plenum to a maximum of 1000 K.

Using the equation of state for normal hydrogen (Leachman et al., 2009), it is calculated that heating hydrogen from its liquid form to the 300 K and 1000 K temperatures of each plenum would require 4.29 MJ/kg and 14.53 MJ/kg, respectively. For stoichiometric fuel mass flow rates at the simulated flight conditions, this means that the hydrogen fuel has a full-engine cooling capacity of 44.4 kW. If further cooling capacity is necessary, the engine can be operated in fuel-rich mode. This increase in fuel mass flow rate would further increase jet penetration and fuel mixing, and REST engines operating at equivalence ratios greater than unity have been shown to have higher combustion efficiencies than those operating at stoichiometric conditions (Suraweera and Smart, 2009), (Turner and Smart, 2010), (Moule et al., 2014).

In order to determine what cooling loads would be required by the M12 REST,

some further assumptions must be made about the true operating conditions of the engine. The flight-condition simulation conducted in the previous chapter held the wall temperature of the engine at 300 K in order to facilitate comparison with the tunnel-condition simulation. In a true flight condition, the temperatures would almost certainly be allowed to reach temperatures close to the structural limits of the materials used in engine construction. For a carbon-fibre composite material, it is assumed that the engine wall can be safely held at a temperature of 2000 K.

For simplicity's sake, it is further assumed that the majority of the heat load created by the engine would be generated in the combustor and nozzle. Other potential heating loads are neglected; the heating requirements calculated are therefore the lower limit of what would truly be required. In order to determine the estimated true flight heating loads, the change in heat flux into the combustor and nozzle walls with an increase in wall temperature needs to be estimated. In order to do so, the simulated heat flux over the combustor walls at the current 300 K wall temperature was first integrated, and found to be 129.16 kW.

Having reduced the heat flux to an aggregate value, the heat flux is modelled as

$$Q = hA(T_{aw} - T_w), \tag{7.8}$$

where h is the heat transfer coefficient in W/m^2K (assumed to be a constant), A was the total combustor wall area in m^2, T_{aw} is the adiabatic wall temperature, and T_w is the wall temperature. Using the same relationship for the estimated true flight condition, and assuming that h does not vary with temperature, it is therefore possible to estimate that

$$Q_{flight} = Q_{sim}\frac{T_{aw} - 2000}{T_{aw} - 300}. \tag{7.9}$$

The adiabatic wall temperature can be estimated by taking the average flow condition at a point mid-way along the constant-area combustor. Assuming turbulent flow through the entire combustor, the adiabatic wall temperature is estimated using the reference enthalpy of Meador and Smart (2005):

$$\frac{h_r}{H_\infty} = 0.5\left(1 + \frac{h_w}{H_\infty}\right) + 0.16r\frac{\gamma - 1}{2}M_\infty^2, \tag{7.10}$$

where h_r is the reference static enthalpy at an adiabatic wall ($C_p T_{aw}$), and r is estimated as the cube-root of the flow Prandtl number. The flow mass-flux weighted average Mach number at the combustor mid-point is 3.034, the average flow temperature is 1818 K, and the stagnation enthalpy is taken to be the flow condition stagnation en-

thalpy of 7.01 MJ/kg. The average Prandtl number was assumed to be 0.72 based on air and hydrogen Prandtl numbers at this temperature, and the average value of the specific heat ratio in the flow was calculated to be 1.402. Using equation 7.10 above, the reference enthalpy is calculated to be 5.56 MJ/kg. Using a curve fit for the C_p of air versus temperature (McBride et al., 2002), the adiabatic wall temperature is thus estimated to be 2540 K.

Based on this estimated adiabatic wall temperature, the in-flight bulk combustor heat transfer (with 2000 K walls) can be estimated from equation 7.9 to be 24.1% of the simulated heat transfer, or 31.13 kW. This is roughly 70% of the cooling capacity available if the present mass flow rate of hydrogen was heated from liquid form to 1000 K in the combustor injectors, and 300 K in the inlet injectors. Even if the enthalpy-based method above under-estimated the adiabatic wall temperature, there would still be ample cooling for an adiabatic wall temperature of 2890 K. Similarly, there is potentially spare cooling capacity available in the engine to cool its leading edges if necessary. As mentioned previously, the cooling capacity can also be extended further if necessary, by over-fuelling the engine.

7.2 Tailored Injection: Oxygen Enrichment

As was mentioned in the previous section, REST engines have often been found to show peak performance when the engine is over-fuelled. However, in cases where excess fuel is injected to the engine, while more free stream oxygen may be consumed, more fuel is also exiting the engine unburned. It could be worthwhile to inject supplemental oxygen with the fuel under these circumstances: this additional pre-mixed oxygen would be free to burn inside the engine, further improving the heat release and pressure rise inside the engine and ultimately contributing to net thrust. Oxygen enrichment has previously been investigated as a means for improving engine performance when operating in flight envelopes where engine mass capture might expected to be below that necessary to sustain sufficiently robust combustion (Razzaqi, 2011). As such, oxygen enrichment would have the added benefit of increasing the operational flexibility of the engine.

As with the fuel temperature effects examined in the previous section, oxygen enrichment would also have a positive effect on the penetration and mixing of the injector jet plumes. The addition of oxygen into the jet would do more than change the ratio of the jet's average molecular weight to the freestream average molecular weight; it would also increase the jet-to-freestream dynamic pressure ratio. This is because the

addition of oxygen mass flow rate at constant temperature and hydrogen mass flow rates would drop the speed of sound at the throat, and the additional mass flow rate of the added oxygen would necessarily have to be made up by increasing the mixture flow density, and therefore static pressure, inside the sonic injector.

To investigate the impact of oxygen enrichment on jet penetration, it is assumed that fuelling of the M12 REST at the flight condition examined in the previous chapter is to be fuelled at an equivalence ratio of 1.2, with the same inlet-to-combustor fuelling ratio of 29:71. The oxygen captured by the engine inlet is no longer sufficient to completely burn all the hydrogen injected, but as previously noted, some additional benefit to mixing may already have been achieved: this increase in fuel mass flow rate could give up to a 4% improvement to jet penetration. Additional oxygen could then be mixed into the combustor fuel injectors to burn more fuel. It is assumed that no oxygen enrichment will be applied to the inlet injectors, as additional combustion on the inlet compression surfaces is not desirable. All oxygen enrichment considered will therefore be done from the tailored injectors, where there is the greatest need for improved penetration and mixing.

To measure the amount of oxygen enrichment that can be used in this engine without dropping it below stoichiometric conditions, three new quantities are defined. The first is the freestream addition percentage (FAP), defined as

$$FAP = \frac{\dot{m}_{O_{2,extra}}}{\dot{m}_{O_{2,\infty}}} \times 100\%. \tag{7.11}$$

For this conceptual investigation, the maximum possible value for FAP is 20%. A higher value would push the engine into fuel-lean operation. The second parameter of interest is the enrichment percentage (EP), which is a commonly used parameter in oxygen enrichment studies (Razzaqi, 2011). The enrichment percentage is defined as

$$EP = 0.125\frac{\dot{m}_{O_{2,extra}}}{\dot{m}_{H_2,total}} \times 100\%. \tag{7.12}$$

Given an FAP of 20%, the enrichment percentage required to return the engine to stoichiometric operation is 16.67%.

Finally, because oxygen will be added to the combustor injectors only, a jet mass fraction parameter (JF) is defined. This parameter applies to a single jet only, and is convenient for varying oxygen enrichment in the combustor jets and changing the jet

properties in any penetration correlation used. This parameter is defined as

$$JF = \frac{\dot{m}_{O_2}}{\dot{m}_{total}} \tag{7.13}$$

where the mass flow rates are specific to the injector being examined. Assuming an equal split of enriching oxygen between the five combustor injectors, the maximum value for JF is 0.6525. By varying this parameter from 0 (no enrichment) to its maximum value (stoichiometric engine conditions), the effect of oxygen enrichment on the penetration of a 300 K jet can be determined.

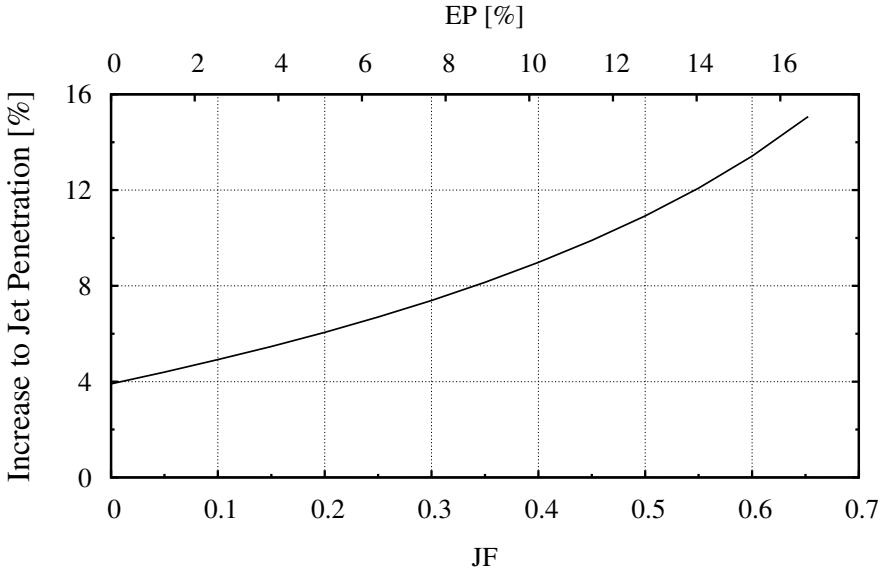

Figure 7.2: Change in jet penetration with oxygen enrichment.

As can be seen in Figure 7.2, the addition of oxygen could increase jet penetration in the combustor by a significant amount. When oxygen enrichment sufficient to return the engine to stoichiometric operation was used, the increase in jet penetration over that from increased fuelling alone is 10.74%, or 15% over an injector in a stoichiometrically fuelled engine. This improvement, combined with the additional oxygen available for combustion could yield a significant performance improvement.

Further performance improvements could also be gained in applying oxygen enrichment to specific jets only. The sidewall jets used in the present tailored injection configuration were found in Chapter 6 to be swept upward into the bodyside region of flow, where most of the oxygen has already been consumed by the combustion of inlet-injected fuel. Shifting additional oxygen away from the cowlside injectors (which will mix with unreacted air flowing in that region of the isolator) to the sidewall injectors could potentially improve combustion efficiency to a greater extent than uniform

enrichment by adding more oxygen to the fuel-rich bodyside flow.

Of course, oxygen enrichment might not be easily combined with a system in which fuel heating is employed. While Razzaqi (2011) successfully tested pre-mixed oxygen and hydrogen injection with mixing chamber pressures as high as 2 MPa, the fuel plenum temperature in those experiments was held at 300 K. Increasing the temperature of the fuel and oxygen during pre-mixing could easily lead to detonation of the mixture prior to injection. While this opens up several interesting new avenues for fuel injection, such as the pulse-detonation injectors studied by Malo-Molina (2014), careful tradeoffs between fuel heating and oxygen enrichment will likely be necessary if they are used in combination.

7.3 Tailored Injection: Geometry, Position, and Mass Flow

All discussion thus far has focused on what can be done to improve the current (tailored) injection design of the M12 REST engine. While constraining the geometry was necessary to quantify the potential benefits of fuel heating and oxygen enrichment alone, there would likely be considerable benefits gained simply from further modifying the physical parameters of the injectors. Whether porthole geometry, position, or simply changing fuel flow rates, there are a myriad of available options.

7.3.1 Non-uniform Injector Size/Mass Flow Rate

Up until this point, the study of tailored injection has focused on a set of identically sized portholes, fed from a common plenum. While this was chosen for simplicity, and gave relatively easy-to-predict porthole mass flow rates and jet penetration, it is unlikely to be the ideal for tailored injection. Instead, varying the porthole sizes could be used to great effect.

One simple example may be found by considering the current tailored injection configuration, in which there are five identical injectors: three spaced along the cowlside of the isolator, and one on either side of the isolator lying along the semi-major axis of the elliptical cross-section. Being equal in size, each porthole has identical mass flow rates, which may limit the engine's combustion performance: as previously mentioned the fuel injected from the sidewall injectors is swept up into the bodyside flow, where any fuel not already reacted finds little oxygen available with which to burn. Some of the fuel from these injectors thus exits the engine unburned.

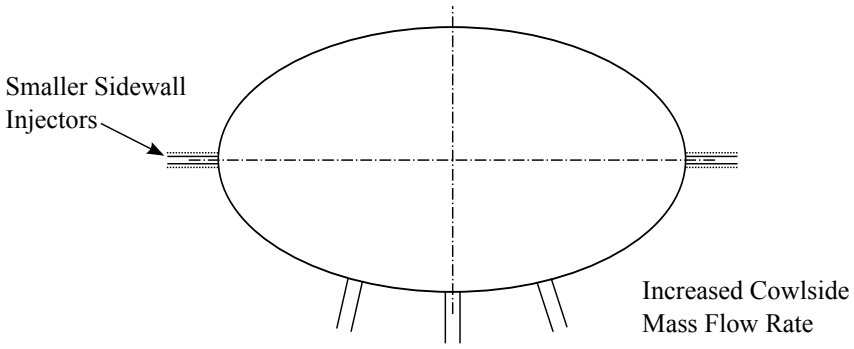

Figure 7.3: Proposed changes to current sidewall injector size (original size in dashed line)

There could be some benefit simply by limiting the mass flow rate through the sidewall injectors, as illustrated in Figure 7.3. This could be done by giving the present injectors a smaller diameter, allowing them to continue to be fed from the same fuel plenum as the cowlside injectors. If the fuel is limited to what would be required to burn the local oxygen along the sides of the combustor, the difference in fuel mass flow rate can be injected through the cowlside portholes to increase their penetration.

As an example, if the mass flow rate through each of the sidewall injectors is halved, this leads to a 33% increase in mass flow rate through each of the cowlside injectors. In order to accept this mass flow rate without changing the plenum pressure, the cowlside injectors' diameters must increase by 15%. Considering that jet penetration is directly proportional to jet diameter, this simple shift to non-uniform injector mass flow rates can yield a significant increase in fuel penetration into, and mixing with, the cowlside core flow of captured air. If this were combined with fuel heating, cowlside penetration could improve by nearly 30%.

7.3.2 Porthole Geometry

Several options for injector geometry were described briefly in Chapter 2. While most alternative porthole shapes showed superior mixing, this was often found to be at the cost of decreased jet penetration height. Nevertheless, some of these may still provide sufficient benefit to fuel jet mixing to be considered viable alternatives to a circular porthole. Of course, if the jet penetration were lowered by moving to a different porthole shape, some effort would have to be put into boosting penetration. This would commonly be achieved by increasing the injector's dynamic pressure, or its diameter, both of which would increase mass flow rates.

There is, however, another aspect to injector geometry which has the potential to improve jet penetration. The angle at which a jet is injected relative to the wall, in the

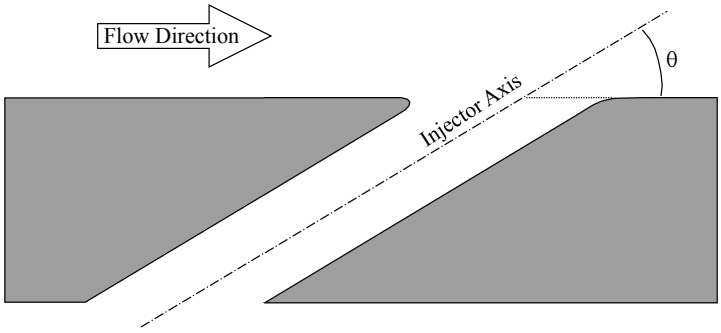

Figure 7.4: Schematic showing streamwise injector angle, θ

streamwise direction (Figure 7.4), can greatly improve the penetration characteristics of a jet. As mentioned in the literature review, McClinton found that shallower angles provide greater jet penetration, and induce lower flow momentum and total pressure losses due to having to turn the jet downstream. Ogawa (2012) performed a multi-objective design optimization study which found the best penetration was achieved by a jet near $\theta = 30°$. McClinton similarly showed that a $\theta = 30°$ injector provided greater penetration than the $\theta = 45°$ injector used presently.

It may be that the best possible injector is a non-circular injector, inclined at a shallow angle to the wall, likely just above or below 30°. This greater degree of penetration does have one additional cost that should be mentioned: it does not occur in the jet near field. Instead, it takes far longer for the maximum jet penetration to be realized; in the study of McClinton this was approximately 60D downstream, after which point the jet plume was pushed back toward the wall by competing flow effects. At the current tailored injector position, a point 60D downstream would be nearly half-way through the constant-area combustor.

7.3.3 Porthole Position

Changing the streamwise injection location such that the cowlside jets are allowed to penetrate deeper into the main engine flow before they enter the combustor and ignite could give significant benefit to the overall engine performance. A well-mixed fuel-air flow tends to ignite and burn quickly, particularly at the high flow stagnation enthalpy of a Mach 12 flight condition. At lower flight speeds, while the flow is less energetic overall, the flow residence times increase as well. Therefore, moving the cowlside injection further upstream could lead to higher combustion efficiencies and would potentially allow for a shorter constant-area combustor, lowering skin friction drag generated in this region.

Given what is known about fuel injection so far, it would seem most prudent that the cowlside tailored injectors be moved upstream such that they are ignited by the burning bodyside flow at the combustor step. Ignition further upstream would lead to increased inlet/isolator pressures, as well as higher skin friction drag along the cowlside wall of the inlet and isolator. If combustion of the tailored jets were to occur too far upstream, the inlet could possibly unstart.

Consider, for the time being, a single injector with a nominal diameter of 1 mm placed on the M12 REST cowlside wall, angled at 30°to the wall. Based on the work of McClinton, the injector fuel plume would be expected to reach its maximum penetration approximately 60 mm downstream. If this point of maximum penetration is set as the combustor step, then the injector would need to be placed at x = 970 mm, just upstream of the inlet throat. McClinton's correlation for jet penetration at an oblique angle is

$$\frac{z_p}{D_j} = 3.385\left(\frac{q_j}{q_\infty}\right)_{eff}^{0.09}\left(\frac{x}{D_j}\right)^{0.18} \tag{7.14}$$

where

$$\left(\frac{q_j}{q_\infty}\right)_{eff} = \frac{q_j \sin\theta}{q_\infty - q_j \cos\theta}. \tag{7.15}$$

Based on the information above, and assuming the jet and freestream dynamic pressures are matched ($q_j = q_\infty$), this relation predicts a penetration of 7.96 mm, or 45% of the isolator's height. The mass flow rate through this injector at the plenum conditions used in the tailored injection study would be approximately 1.04 g/s, or 83% of the hydrogen fuel injected in the isolator of the current tailored injection case. The predicted plume edge is shown in Figure 7.5.

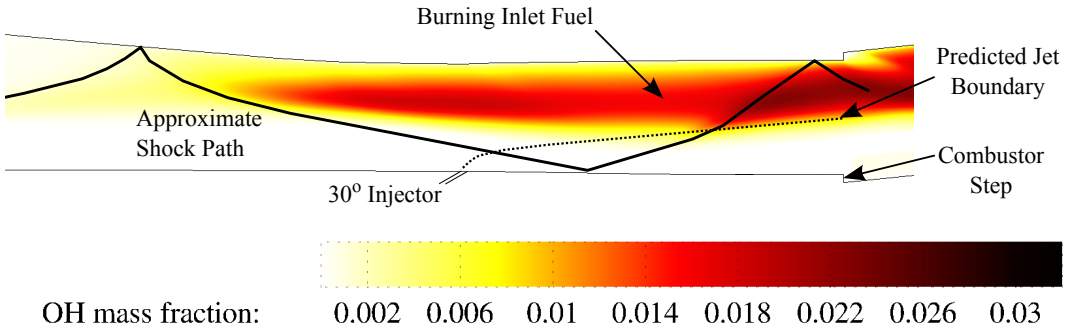

Figure 7.5: Approximate penetration of a 1 mm jet, relative to the OH production of the burning inlet-injected fuel

This predicted penetration height would allow the edge of the cowlside jet plume

to interact with the bodyside plume of fuel approximately 20 mm upstream of the combustor step. The fuel at the edge of the plume would begin to react as it makes contact with the burning bodyside flow, and it is possible that the jet could be burning strongly by the time the flow reaches the combustor step. However, it should be noted that this final plume height would be an optimistic estimate for the actual penetration. The jet would interact with the reflecting cowl closure shock pictured in Figure 7.5, just downstream of the proposed injection point. This shock-jet interaction would increase mixing significantly, but would also reduce the jet's penetration (Schetz et al., 2010). The jet plume boundary would therefore stay nearer to the cowlside wall. Overall, dynamic-pressure-matched cowlside fuel injection at 30°from a streamwise position near the inlet throat could potentially lead to a well mixed flow entering the combustor right next to its ignition source: the burning flow along the bodyside wall of the engine.

If the shock-jet interaction were found to suppress the penetration of the jet so much that ignition at the combustor step were no longer possible, further penetration could feasibly be achieved by using a pair of 'cascading' jets, in which a two jets are closely spaced and aligned in the streamwise direction. As was mentioned in the literature review, the upstream jet would shield the second jet, effectively increasing the downstream jet's dynamic pressure ratio and increasing its penetration. The level of penetration improvement is correlated to jet spacing and could be as much as 50% greater than an equivalent single jet (Lee, 2006).

Assuming this to be a reasonable estimate for penetration improvement in angled injectors as well, the previous example is reconsidered. Rather than taking a single jet of 1 mm diameter, assume the mass flow rate is split equally between two portholes approximately 0.7 mm in diameter (these portholes would have approximately half the cross-sectional area of the 1 mm diameter injector). While the estimated penetration of a single angled jet of this diameter would be only 5.62 mm, the ideal penetration of two cascading jets would be 8.4 mm, or 5% greater than a single 1 mm jet.

When all is said and done, there could be a number of ways to tailor fuel flow in order to take advantage of the large mass of unreacted air that flows along the cowlside of the M12 REST engine. As a note of caution, it may be wise to adopt a more incremental approach to changing injection. A recommended start to further modifying the tailored injection would be to first shrink the sidewall injectors and divert more mass flow to the current cowlside injectors, in order to determine if any performance gains or losses would be recorded from this shift in fuelling, while maintaining the overall fuel mass flow rate. This would help determine what the ideal mass flow rate

from the sidewall injectors should be. The designer would then be free to make modifications to the cowlside injection arrangement, and determine the relative effects on engine performance that those changes had, independently of the sidewall injectors.

7.4 Engine Improvements Beyond Tailored Injection

While there may yet be substantial gains to be made in the combustion efficiency of the M12 REST through further tailoring of the fuel injection, this is not the only avenue open for improving the engine's current performance. There are several aspects of the M12 REST engine's design and operation that, if modified appropriately, could each move the engine closer to achieving net uninstalled thrust Mach 12 flight-equivalent tunnel testing conditions. Potential changes that can be explored include: constraining inlet injection; making improvements to boundary layer injection; and, modifying the geometries of the isolator, combustor, and nozzle.

7.4.1 Improving the Nozzle Design

The present nozzle of the M12 REST engine is a simple uniform expansion of the elliptical combustor geometry. While it generates thrust, the nozzle does not take into account the non-uniformity of the flow, nor does it attempt to minimize entropy losses in the expansion process. While the actual details of an ideal nozzle design are beyond the scope of this study, an estimation of how much additional thrust could be extracted from the flow may be made.

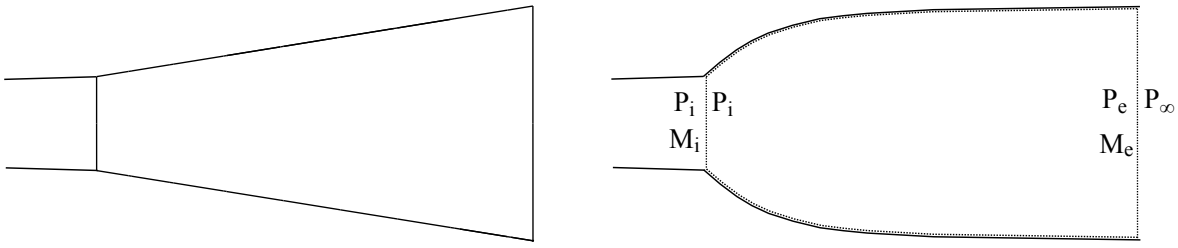

Figure 7.6: Schematic of the current nozzle (left) and ideal nozzle (right), with control volume

As a simplification to the problem, it is assumed that the flow is quasi-one dimensional. The nozzle geometry is then assumed to change from a uniformly expanding surface to an isentropic thrust nozzle with the same exit-to-inlet area ratio of 5, giving the same engine nozzle-to-inlet aerodynamic capture area ratio of 1.35. This is illustrated in Figure 7.6. A control volume is then drawn around the nozzle surface, as pictured. This allows for a stream-thrust analysis of the nozzle to be made. A balance

of the forces on this control volume shows that the nominal thrust force acting on the engine would be equal to

$$F_{thrust} = \dot{m}_e U_e - \dot{m}_i U_i + A_e(P_e - P_\infty) \tag{7.16}$$

where the subscripts e and i denote nozzle exit and inflow conditions, respectively.

It is assumed that the conditions entering this idealized nozzle are equal to the mass-flux-weighted average conditions at the exit of the diverging combustor during the tailored injection case study presented in Section 6.3. This gives the control volume an inflow Mach number of 2.603, an inflow static pressure of 43.32 kPa, and an average specific heat ratio γ = 1.391. The ambient pressure (P_∞) is assumed to be equal to the static pressure of the inflow condition in the same simulation.

Assuming a perfect gas, the nozzle inflow Mach number can be used to calculate the area ratio between the nozzle inlet and the theoretical throat cross-sectional area the flow would need to flow into in order to achieved choked conditions.

$$\frac{A}{A*} = \left(\frac{\gamma+1}{2}\right)^{\frac{1-\gamma}{2\gamma-2}} \frac{\left(1 - \frac{\gamma-1}{2}M^2\right)^{\frac{\gamma+1}{2\gamma-2}}}{M}. \tag{7.17}$$

From this equation, and knowing the nozzle area ratio, the area ratio between the nozzle exit and the hypothetical choked throat can be calculated as

$$\frac{A_e}{A*} = \frac{A_e}{A_i}\frac{A_i}{A*}. \tag{7.18}$$

Employing equations 7.17 and 7.18 shows $A_e/A*$ to be equal to 14.6773. This ratio can in turn be used in equation 7.17 to back-calculate M_e as being equal to 4.317.

Being in possession of the exit Mach number of the nozzle, it is therefore possible to calculate the pressure ratio between the nozzle inlet and exit. This is done by recognizing that, according to the isentropic flow relations,

$$\frac{P_e}{P_i} = \frac{\frac{P_e}{P_0}}{\frac{P_i}{P_0}} = \left(\frac{1 + \frac{\gamma-1}{2}M_e^2}{1 + \frac{\gamma-1}{2}M_i^2}\right)^{\frac{-\gamma}{\gamma-1}}. \tag{7.19}$$

Using this relation, the pressure of the flow exiting the idealized nozzle is found to be P_e = 3696.66 Pa. This pressure is used to calculate the thrust force due to the pressure differential at the nozzle exit. It can also be used to determine the forces due to the

differences in flow momentum at the nozzle inlet and exit. Knowing that

$$\dot{m}U = (\rho UA)U = \rho U^2 A \tag{7.20}$$

the mass formulation of the ideal gas equation can be used to show that

$$\dot{m}u = \frac{P}{RT}U^2 A. \tag{7.21}$$

Knowing that the speed of sound, a, is defined as $\sqrt{\gamma RT}$, the momentum flux can therefore be rewritten as

$$\dot{m}u = \frac{\gamma P}{a^2}U^2 A = \gamma P M^2 A. \tag{7.22}$$

The momentum flux is thus redefined entirely in terms of the two quantities already known at the nozzle inlet and outlet. The formula for the thrust force acting on the nozzle can therefore be rewritten as

$$F_{thrust} = \gamma P_e M_e^2 A_e - \gamma P_i M_i^2 A_i + A_e(P_e - P_\infty). \tag{7.23}$$

Using the inlet and exit conditions already calculated for the nozzle, and assuming a constant specific heat ratio γ = 1.4, the potential thrust this ideal nozzle could generate is 71 N. This number must be halved to 35.5 N to compare with the nozzle in the current simulation, which was for half the engine only.

This ideal nozzle thrust is 81.5% higher than the 19.55 N produced by the current uniform-expansion half-nozzle in the simulation. The stream thrust analysis did not include any viscous (skin friction) drag losses. If it is conservatively assumed that the ideal nozzle generated 20% more skin friction than the current nozzle, this drops to a 58% increase in thrust. Further improvements on generated thrust forces could be achieved if the nozzle to inlet capture area ratio were increased from 1.35 to a larger expansion ratio.

The current aerodynamic forces acting on the engine suggest that the engine would need to increase its thrust by 75% to achieve uninstalled cruise thrust under testing conditions achieved in T4. A well-designed nozzle that provides thrust generation similar to the ideal nozzle examined here would close this thrust gap considerably. It may therefore be concluded that the design of a better nozzle than is presently used in tunnel tests of the M12 REST engine will be required to generate net thrust during experiments.

7.4.2 Changes to Isolator and Combustor Geometry

In addition to a new design for the nozzle, there may be many benefits to modifying the design of the M12 REST isolator and combustor. The constant-area combustor, which was responsible for half of the viscous drag generated downstream of the inlet, was an area of particular concern.

One simple option that may be considered is to shorten the combustor. With combustion efficiency in the M12 REST at nearly 85% for the current tailored case, the combustor could likely be shortened while maintaining reasonably high combustion efficiency. If this length were lost from the constant area combustor, there would be a savings in viscous drag, without greatly impacting the thrust generated. This becomes even more viable if further improvements to injection tailoring are pursued: improvements to combustion efficiency made by moving injection upstream or heating the fuel could potentially shorten the required combustor length even further.

However, a future REST engine designer would do well to recall from previous chapters that the primary source of drag in the constant-area combustor typically comes just downstream of the combustor step, from where the flow is turned by the combustor wall. This high viscous stress region was found in the previous chapter to be magnified by the presence of combustion, and would still exist regardless of combustor length. If this region is responsible for the largest amount of viscous drag created in the combustor, it would seem that attempting to prevent its formation would be a logical choice.

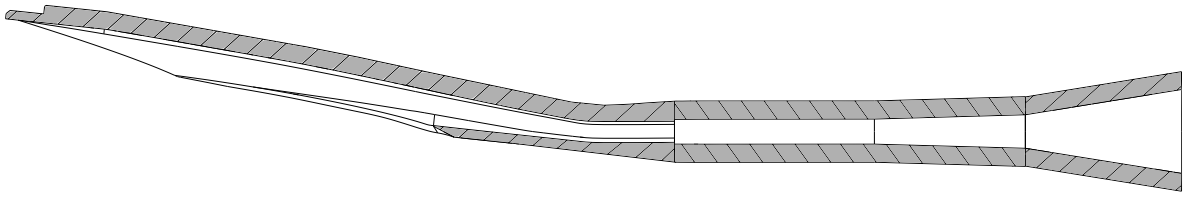

Figure 7.7: Schematic showing realignment of the M12 REST internal flow path in the isolator

It is proposed that this region be dealt with by moving the engine turn upstream into the isolator. Rather than simply moving the problem elsewhere, this could potentially reduce the skin friction generated by the engine turn. Conceptually, putting the turn in the isolator may look similar to the diagram shown in Figure 7.7. Rather than a single, sharp turn, the flow could instead be turned more gradually. This would generate a series of weak compression waves rather than a single strong compression, which would serve to drive down the peak magnitude of compression on the cowlside wall.

Placing the turn in the isolator has also been shown previously to have a negligible effect on the inlet performance (Gollan and Ferlemann, 2011).

Such a turn would also have another benefit beyond reducing combustor shear stress. By turning the internal flow earlier, the engine's external surfaces would also realign with the direction of flow sooner. This would reduce the external pressure drag generated by an integrated engine. As shown by Doherty (2013), the external pressure drag of an integrated engine is an appreciable fraction of its overall drag. Canting the flow path back into alignment with the direction of thrust earlier would shrink the frontal projected area of the engine, thereby reducing pressure drag. Realigning the engine upstream of the combustor also re-opens another avenue for drag reduction that was previously shown to work poorly: boundary layer injection.

7.4.3 Reconsidering Boundary Layer Injection

It was shown in Chapter 5 that the boundary layer injection employed in the baseline M12 REST engine design did not perform as well as expected. This seems to primarily be caused by the portholes being too widely spaced to create a continuous sheet of fuel along the combustor surface, as well as by the injected fuel being quickly disrupted by the high-pressure zone created by the engine turn inside the combustor. If the engine turn were moved upstream, the high pressure zone downstream of the combustor step would be mitigated. This could help to prevent the fuel layer from being disrupted, potentially improving skin friction reduction performance inside the combustor.

As well, the injectors would no longer be constrained by the turn occurring at the combustor, which would allow them to inject fuel completely parallel to each combustor wall, as shown in Figure 7.8.

Moving the injectors on the step face closer together would require either shrinking the porthole size, or lowering the fuel plenum pressure in order to avoid unduly increasing the boundary layer fuel mass flow rate. If it was found that adding more portholes to the step face was not a feasible option, the injectors could also be redesigned to shield the injector portholes from the flow, giving the fuel plumes time to spread into a coherent sheet before they can interact with the main flow. One way this might be done is by recessing the portholes further upstream from the nominal face of the combustor step, and leaving a gap between the combustor and isolator walls, as illustrated in Figure 7.9.

This design could feasibly create a sheet of fuel without the complete structural decoupling of the combustor and isolator that a true slot injector would require. This

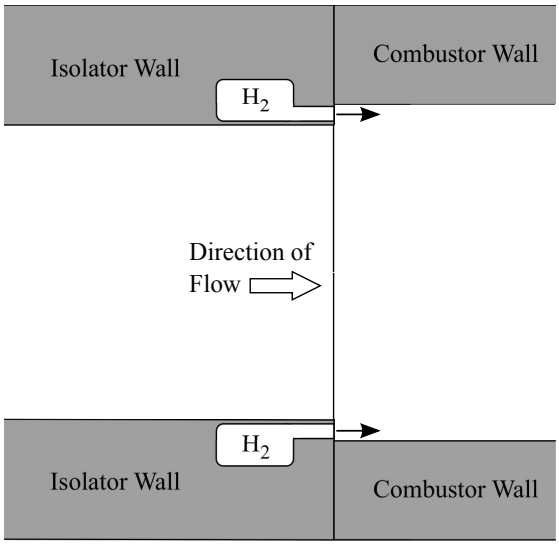

Figure 7.8: Schematic of the proposed change to step injection in the M12 REST combustor

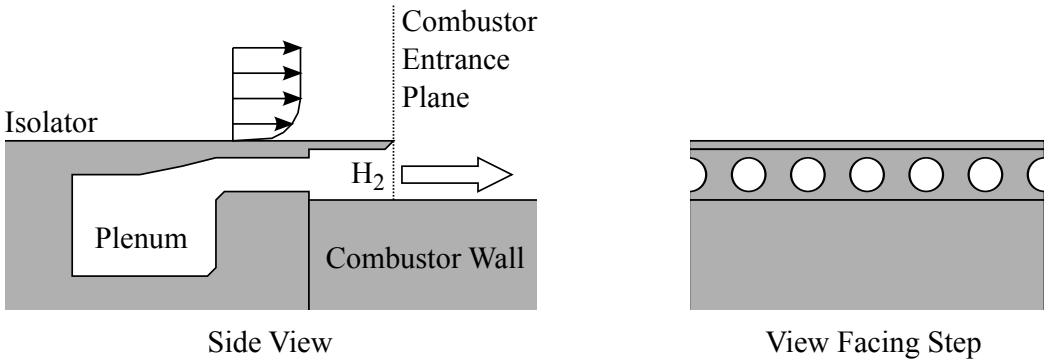

Figure 7.9: Side (left) and step-facing (right) views of possible shielded step injection

design would be more complex to manufacture, and would feature some relatively fine structures, but these should be protected by the cooling flow of hydrogen.

Boundary layer injection could also be employed in a local, rather than global, fashion. Consider the viscous shear stress distribution of the M12 REST engine shown in Figure 7.10. The distribution of shear stress inside the constant-area combustor is heavily concentrated along the bodyside surface.

If only the drag from regions in which the wall shear stress magnitude exceeds 1250 Pa is integrated, the resulting area is found to be on the cowlside wall and part of the sidewall of the combustor. This area is only 44.4% of the total combustor wetted area, but generates 62% of the viscous drag. As such, more focus might be put on boundary layer injection targeting these high-shear regions, rather than on uniformly applying it to the whole combustor

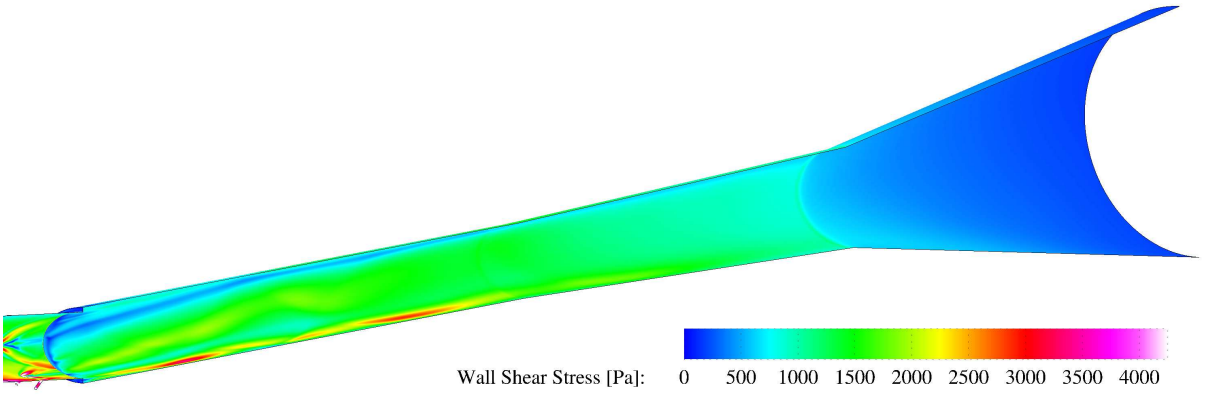

Figure 7.10: Shear Stress in the M12 REST combustor with tailored injection

This targeted skin friction reduction may not be confined to the combustor only. A similar analysis can be done on the inlet, and it can be shown that 8% of the inlet's wetted area (the regions just downstream of the cowl closure, and along the bodyside wall after the inlet-injected fuel ignites) is responsible for generating 26% of the inlet's viscous drag. Implementing targeted viscous drag reduction in these regions of the inlet could be done using the porthole-array technique of Pudsey (2010), or perhaps through some other technique, such as porous wall injection (Schetz and van Overeem, 1975).

Of course, targeted boundary layer injection would require the designer to know how much fuel might be needed in each region targeted for reduction. This may be done on a high level using the analytical method developed in Barth et al. (2013). As a rule of thumb, it is likely that the total fuel devoted to all targeted boundary layer injection employed should be at, or lower than, the $\phi = 0.1$ levels recommended for general combustor boundary layer injection recommended in Appendix G.

7.4.4 Constraining the Inlet-Injection Equivalence Ratio

The final potential engine change to be examined deals with the inlet injection of fuel. Up until this point, the injection of fuel along the inlet bodyside compression surface has been treated as something of a constant: the approximate fuelling split between the inlet and combustor was set, and the actual inlet equivalence ratio was permitted to increase and decrease in line with the global equivalence ratio. While this tended to give good engine performance, putting some constraints on the level of inlet fuelling may be beneficial for a number of reasons.

The first of these reasons is the impact of increasing inlet equivalence ratio on inlet performance. This can be examined in some depth, thanks to the range of inlet

equivalence ratios that were employed over the course of this study. The flight condition simulation was not included, as it lacked an equivalent unfuelled simulation for normalization. It is immediately evident from examining Table 7.1 that increasing the amount of fuel injected on the inlet improves its compression ratio, thanks to thermal compression.

Table 7.1: Effects of varying inlet equivalence ratio. Values are normalized to an unfuelled inlet.

Equivalence Ratio	Normalized Compression Ratio	Normalized P_0 Recovery	Normalized T_0 Recovery	Normalized Inlet Drag	Throat η_{comb}
0.33	1.2320	0.7190	0.9197	1.0484	52.64%
0.356	1.2963	0.6949	0.9204	1.0498	54.21%
0.411	1.3714	0.6568	0.9100	1.0495	52.62%
0.548	1.4767	0.5820	0.8773	1.0742	48.06%

While this increased compression may be a desirable trait, it should also be noted that the total pressure recovery of the inlet (relative to the unfuelled total pressure recovery) drops as more fuel is injected. Total temperature recovery is relatively constant up to a point, but begins to fall at an inlet equivalence ratio between 0.41 and 0.55. Going to an inlet equivalence ratio above 0.5 also appears to have a strong effect on the drag generated by inlet injection. While lower equivalence ratios showed a reasonably constant level of drag, the largest equivalence ratio shows a sudden jump in inlet drag.

The fuel-based combustion efficiency at the throat is also tabulated, and shows that there is a peak in efficiency near $\phi = 0.35$. Higher equivalence ratios show lower efficiency which, based on the behaviour observed in Section 5.2, suggests that more inlet injected fuel would be leave the engine nozzle unreacted. Combined with the considerations above, it would seem prudent to limit inlet equivalence ratio to be no higher than 0.4. Combustion efficiencies at this point are still reasonably close to the peak value.

On the lower limit, it was found in the experimental campaign of Wise (2014) that an inlet equivalence ratio of 0.2 showed no significant pressure rise in the isolator relative to an unfuelled inlet. This suggests that any combustion that did occur was not sufficient to provide thermal compression of the flow. The flight condition simulation of shot 11491 had an equivalence ratio of 0.26, and did show some thermal compression effects. This may be taken as the current lower bound at which inlet fuelling may provide a significant advantage with regard to thermal compression and piloting combustion.

Overall, it is believed that constraining the inlet equivalence ratio to a range between 0.26 and 0.4 would be prudent. If the equivalence ratios are held relatively constant within the range specified above, the thermal compression effects at on-design conditions will have a similarly steady value. Consistent thermal compression could then be assumed, and used to examine a further optimization of the inlet to a lower geometric compression ratio, which could result in a shorter inlet and lower inlet-generated drag.

Finally, constraining the inlet equivalence ratio would mean that when the engine is over-fuelled, the inlet-to-combustor fuelling ratio would no longer be constant, and the relative mass flow rates in the combustor fuel injectors would rise. In the case of mainstream injection such as was used in the tailored injection case studies, this would result in increased fuel jet penetration and mixing, which in turn would improve fuel mixing and combustion efficiency.

7.5 Summary

Having come full circle by once again mentioning changes to tailored fuel injection in the combustor, it may be commented that continuing to improve the performance of the M12 REST will be an iterative process. The engine itself already exhibits reasonably good performance, and other than the nozzle, no major changes to the engine are required. Several small changes, on the other hand, could result in reasonable gains in engine performance.

An update to the nozzle design has the potential to improve thrust considerably, as do improvements in combustion efficiency brought about by further tailoring the fuel injection. These further improvements in combustion efficiency could also permit a reduction in viscous drag, by shortening the engine combustor. Realigning the engine with the nominal axis of vehicle travel further upstream could reduce overall drag in several ways, and could possibly make boundary layer injection and combustion a more viable technique for reducing viscous losses in the M12 REST combustor. Taken in aggregate, it is reasonable to believe that some combination of the incremental improvements discussed in this chapter could ultimately result in net thrust at Mach 12.

8

Conclusions and Recommendations

"Everything looks nonsensical before it works."

– Burt Rutan, *Remarks at Oshkosh AirVenture* (2011)

The purpose of this study was to determine what flow physics are at work inside of a streamtraced hypervelocity scramjet engine, as well as to determine the potential of tailored fuel injection as a method for improving scramjet performance. To determine the answers to these problems, an extended campaign of CFD simulations was carried out, focusing on the M12 REST scramjet. The simulations were conducted using the research CFD code US3D, and included a full finite-rate chemistry model to conduct the first combusting simulations of a fuelled REST scramjet engine operating at Mach 12. These simulations were supplemented by an experimental campaign, in order to validate the flow fields obtained from the simulations. The experiments were carried out in the T4 Stalker Tube facility at The University of Queensland, using an engine model identical in size to that which was simulated. Several engine operational modes were examined, including unfuelled operation; inlet-fuelled operation; combined inlet and boundary layer fuelled operation; and, two attempts at replacing boundary layer injection in the combustor with tailored fuel injection.

The examination of the M12 REST engine's flow physics was conducted in two parts: first the inlet flow physics alone were investigated, followed by the combustor flow physics. In both cases, performance metrics were extracted in order to determine the changes in engine performance as the engine's fuelling was modified. The relative differences between an engine operating in the flow generated by the T4 Stalker Tube was compared with the predicted performance of an engine operating in a true flight condition. Wherever possible, areas for potential improvement in the engine design, fuelling, and operation were noted, and later examined.

8.1 Conclusions

Based on the simulations and experimental work undertaken in this study, several conclusions can be drawn. These conclusions cover both the scientific findings of the investigation, as well as the validity of the methodologies employed. It is believed that all the conclusions made should be generally applicable to any scramjet with a shape-transitioning and/or three-dimensional inlet.

The primary flow structures generated in the M12 REST inlet are: a thick, hot, bubble-shaped boundary layer along the bodyside surface of the engine, a symmetric pair of swept separations, and the strong cowl closure shock wave. The majority of captured air ends up flowing along the cowlside wall of the engine. The bodyside boundary layer shape is a result of the thick incoming vehicle forebody boundary layer being compressed inward by the contraction of the inlet sidewalls. This induces a vortex in the top corner of the inlet, which conducts low-momentum boundary layer flow inward along the bodyside surface. This inward flow meets at the engine centerplane, and ultimately pushes the bodyside boundary layer outward from the wall. This boundary layer "bubble" in turn directs most of captured air downward, such that the majority of the air mass flow is along the engine cowlside wall. The swept separations are symmetric regions of slowly rotating, high temperature flow, generated by the interaction of the cowlside boundary layer with a laterally travelling shock inside the inlet. This lateral shock is generated upstream by the formation of the cowlside leading edge, and reflects off its symmetric pair upstream of the inlet closure, after which it travels away from the engine centerplane along the internal cowlside inlet surface. The lambda foot formed by this shock causes two separate regions of the cowlside boundary layer to roll up, and eventually coalesce to form the swept separation further downstream. The cowl closure shock stands off from the cowl notch, owing to the small portion of the leading edge that is perpendicular to the incoming flow. This shock travels upward along and outward from the engine centerplane, reflecting off the inlet bodyside compression surface upstream of the throat. The resulting shock train compresses the flow further, and has implications for the downstream flow physics summarized in conclusions that follow.

Fuel injected from the inlet bodyside compression surface at low equivalence ratios remains trapped within the thick bodyside boundary layer. The fuel ignites at the inlet throat, providing thermal compression of the flow inside the isolator. Owing to both the large thickness of the bodyside boundary layer and the relatively low momentum of the fuel jet, the injector plumes are turned downstream before they

can escape the boundary layer. The boundary layer thickness grows slightly due to the additional fuel mass inside of it, and boundary layer turbulence quickly mixes the fuel. The hot, thin, boundary layer being swept by the inlet sidewalls inward toward the engine centerplane interacts with the outside edge of the fuel-laden boundary layer, creating a local hot spot in which heat and radicals are produced. Low pressure in the boundary layer prevents this zone from igniting the entire fuel-laden bodyside boundary layer, until after it interacts with, and is compressed by, the cowl closure shock. Downstream of this shock-boundary layer interaction (SBLI), the production of radicals and heat release rapidly spreads across the boundary layer, and combustion of the inlet-injected fuel is established. By the time the flow reaches the inlet throat, as much as half of the fuel may be completely consumed by combustion. The heat release due to combustion of inlet-injected fuel is sufficient to provide thermal compression of the flow inside the isolator, making the flow conditions entering the combustor more ideal for the rapid ignition and reaction of fuel injected downstream.

Flow inside the combustor is strongly shock-dominated. Both the reflecting cowl closure shock train transmitted from the inlet and recompression shocks downstream of the combustor step propagate through the engine combustor. These shocks rapidly compress and weaken any coherent streamwise vortices in the flow, leaving the mixing inside the combustor to be driven by lateral vorticity and shear flows; local vorticity in the vicinity of shock waves; and SBLI effects. The presence of these shock trains drives up local combustor pressure, allowing combustion to occur rapidly. This leads to local regions of high chemical reaction rates, particularly on the bodyside of the engine, where SLBI effects are dominant.

Fuel injected from the face of the combustor step for the purposes of boundary layer combustion is ignited by the burning bodyside flow created by inlet-fuelling. The discrete fuel jets never coalesce into a single fuel structure, and are quickly disrupted. As such, no appreciable skin friction reduction is achieved. The portholes in the face of the combustor step are too widely spaced, causing the core each injector plume to flow downstream in discrete fuel-rich streaks, joined to one another by regions of well-mixed fuel and air. Well-mixed regions of fuel and air along the edges of the jets allow the bodyside flame to ignite the boundary-layer-injected fuel, which then propagates circumferentially around the engine. The realignment of the engine walls with the nominal direction of vehicle travel creates a compression ramp effect along the cowlside of the combustor. The resulting pressure gradient just downstream of the combustor step forces unconsumed hydrogen either inward toward the cowlside centerline, or outward and upward along the combustor sidewall. As such, the fuel is

unable to burn within the cowlside boundary layer, and no appreciable skin friction reduction is achieved. This unintended mixing does, however, have a positive effect on the combustion of step-injected fuel with the mainstream flow, achieving a final combustion efficiency of 82.9% with an equivalence ratio of 1.33.

Fuel injection tailored to the internal flow structure of a three-dimensional scramjet increases the mixing and combustion performance of the engine. By placing mainstream porthole injection preferentially along the cowlside of the engine, it is possible to ensure that fuel injected into the engine isolator has access to the majority of the inlet-captured air. Fuel injected from the isolator sidewalls into the swept separation ignites faster than fuel along the cowlside wall, though this sidewall-injected fuel is ultimately swept upward into the oxygen-poor bodyside flow region. Overall, the M12 REST with tailored fuel injection at $\phi = 1.24$ has a global combustion efficiency of 84.9%, a 2% increase from the boundary layer injection case. The tailored injection case also produces nearly the same thrust as the baseline case, despite using 6% less fuel. The tailored engine can therefore be concluded to operate more efficiently, having a 5% higher uninstalled specific impulse.

RANS simulations are sufficient to accurately capture the pressure behaviour of experiments. Every simulation with an inflow condition taken from a T4 Stalker Tube nozzle exit flow condition has been compared with its corresponding experiment. In nearly all cases, good agreement exists between the experimental pressure measurements and the simulated distribution of pressure in the engine. The simulation of the boundary layer injection configuration did not agree well with the experimental data, though in this case, there is thought to have been unusually long transient effects in the experimental data caused by early fuel timing. These transient effects were not smoothed out by the end of the engine's experimental test time. The simulations tend to over-predict the magnitude of heat flux inside the engine combustor. They do, however, often manage to capture the trend in the heat flux. This is most likely a turbulence-model-related effect, which may possibly be improved on with a finer mesh near the wall, or a change in the management or values of the turbulent Prandtl and Schmidt numbers.

Engines tested at hypervelocity conditions in the T4 Stalker Tube accurately capture the normalized pressure and heat flux distributions of a cold-walled scramjet in flight, but have lower inlet mass capture. The tunnel-validated simulation of the best tailored injection case has been compared with the operation of an identical engine at its pressure- and enthalpy-equivalent flight condition. The largest difference recorded was in engine mass capture, which was 11% higher in flight than it was in the tunnel

condition simulation. Combined with the lack of NO in the flight inflow condition and identical fuel mass flow rates, the flight equivalence ratio was 0.90. This is 27% lower than the tunnel equivalence ratio, and corresponding changes in combustion efficiency. Despite these differences, the overall normalized pressure and heat flux distributions predicted in each simulation are in close agreement. A match of equivalence ratios would likely provide a more appropriate comparison for combustion efficiency between the two cases. The presence of atomic oxygen in the tunnel condition simulation was not found to have an appreciable effect on the streamwise location of ignition in the scramjet.

While tailored injection improves engine performance, it is not sufficient to generate net thrust in tunnel testing. While an engine in flight will no doubt have better performance, further design improvements will likely still be required. Despite achieving a high combustion efficiency, the viscous drag generated by the engine's internal surfaces still exceed the thrust generated by the engine. An engine in flight would lose less heat to the combustor walls, and therefore see greater flow compression, but would still benefit from improvements in engine design. Several options for modifying the engine have been examined, including injector placement, size and geometry; modifications to the engine isolator and combustor; a retooling of boundary layer combustion; and, the potential impact of redesigning the engine nozzle. Many of the options examined have potential benefits, and form some of the recommendations for future work.

8.2 Recommendations for Future Study

Based on the conclusions of this study, several recommendations for future work have been formulated. These are in part based on issues identified during the course of this work that could not be solved within the scope of this project. They also include logical building blocks upon which future experimental work, simulations, and, design changes may be based.

8.2.1 Methodological

The first set of recommendations relate to the methodology by which the M12 REST is studied, and ways in which future studies may be changed to increase their efficacy. Both experimental and computational methods are considered. The first of these recommendations is simple: engage in greater collaboration between experimental and computation studies. This study in particular benefited greatly from collaboration

with another study on a similar engine. The validation experiments carried out in this study would not have been possible if there was not another PhD candidate working on a similar engine, and such data is incredibly valuable for validating simulations.

The next recommendation is to conduct any future experimental investigations of tailored injection utilizing individual Ludweig tubes for each fuel plenum. Using a needle valve to moderate plenum pressure and mass flow rate has been demonstrated to reduce the responsiveness of the moderated plenum. This leads to a lack of accurate repeatability in the fuel equivalence ratio, and the fuel split between two plena. Moving to a system in which each plenum is fed from its own Ludweig tube would not only solve this problem, but would give future experimenters greater control over fuel splitting: the ratio of fuel injected from each plenum would no longer be mostly controlled by the relative area of the injectors fed by each plenum.

A third recommendation regarding methodology is more specific: the REST inlet with fuel injection should be studied using a higher-fidelity computational tool, such as the hybrid RANS/LES methods available in US3D. While such an inlet simulation would undoubtedly have a higher computational cost than the complete engine simulations run in this study, a more detailed simulation of the inlet flow could unravel the causes of the minor differences noted between the simulations and experiments in Chapter 4. A better understanding of where the RANS model is unable to fully resolve the detailed inlet flow physics would also aid the interpretation of future RANS-based studies of the M12 REST engine, and any other streamtraced scramjet inlet geometries.

8.2.2 Topical

The other recommendations to be made are purely based on issues related to the design of the engine, its operation, and its optimization. These recommendations are believed to offer potential avenues for further research, through which hypervelocity scramjets (and the M12 REST specifically) may see significant improvements in performance.

The first of these recommendations is to continue investigating options for tailored injection, which has already shown the ability to improve engine combustion efficiency. It was shown in Chapter 7 that moving the cowlside injectors upstream, injecting at a shallower angle of 30°, and/or combining them into a single jet has the potential to greatly improve the penetration and mixing of fuel into the core mass of air flowing along the cowlside wall. The sidewall injectors should likely be kept in their present location, but have their fuel mass flow rates reduced to prevent excess fuel from being swept into the bodyside flow, where it is unlikely to find any oxygen with which to

burn. Exploring fuel heating could yield several benefits, including increased jet penetration and shorter ignition lengths, both of which would lead to improved engine performance. Oxygen enrichment may similarly improve jet penetration and, in the case of the sidewall jet, may allow otherwise unburned fuel to be consumed.

It is further recommended that an optimization study should be carried out for the REST inlet and isolator geometry when fuel injection is included. The thermal compression produced by the combustion of inlet fuel at the inlet throat, and through the isolator, leads to an opportunity to make inlet injection a more integral part of the inlet geometry. By fixing inlet equivalence ratio and relying on inlet fuel combustion for some predictable level of flow compression, the inlet could be redesigned with a lower geometric compression ratio, which has the potential to significantly reduce the size and drag of the inlet.

Outside of this inlet optimization, all REST engines presently employ a turn somewhere in the engine geometry to realign the internal flow with the nominal axis of vehicle motion. Moving this turn upstream into the isolator and making it a gradual direction change, rather than a sharp turn, could improve the internal drag characteristics of the engine, and make combustor boundary layer injection more likely to work as intended. It would also reduce the projected frontal area of the engine, reducing the external drag generated by an engine in flight.

It is recommended that the topic of viscous drag reduction through boundary layer combustion be investigated further. Particularly, the porthole spacings in the combustor step should be made closer, and/or have their geometries modified to ensure the fuel forms a continuous sheet. Combined with moving the canting of the engine flow path into the isolator, this may make boundary layer combustion a more viable technique.

Targeted boundary layer combustion should also be considered. It was shown in the previous chapter that the majority of the viscous drag generated by a single component of the engine flow path may be generated by only a fraction of its wetted area. Employing some amount of limited boundary layer fuel injection and combustion in these zones could reduce viscous drag without overly impacting the operation of the engine.

The final recommendation is that a new nozzle design should be investigated for REST engines. As a general class of engine, a single nozzle design concept may be applicable for REST engines with a wide variety of design flight conditions. Given that the engine begins with a shape transition, a shape transitioning nozzle may be a

strong choice, and has been the subject of recent investigations in literature (Mo et al., 2014). Alternatively, a given nozzle design has relatively few key variables that define its geometry. Multi-objective design or evolutionary optimizations may be viable paths for determining a nozzle design that achieves optimal thrust with minimal losses.

As a final remark, streamtraced scramjet engines appear to be a strong contender for the design of a scramjet that can operate at hypervelocity conditions. While their geometries often appear unconventional and complex, it should be reiterated that the current issues preventing the 3D streamtraced scramjets from achieving net thrust are not insurmountable. A great deal of performance may yet be achieved by making small changes, starting in the engine inlet, and progressively modifying the engine components downstream as optimal designs of upstream components are locked in. Once the optimal injection locations in the inlet and isolator are determined, for example, an in-depth examination and re-design of the engine combustor and nozzle may be more fruitfully undertaken. Of course, even if flight at Mach 12 proves to be not viable, the lessons learned from these efforts will no doubt make the sustained flight of an scramjet at higher Mach numbers than those achieved by the X-51 a real possibility.

References

Abdel-Salam, T.; **Tiwari, S.** and **Mohieldin, T.** (**2003**). "Effects of Ramp Side Angle in Supersonic Mixing". *AIAA Journal*, **vol. 41, no. 6**, pp. 1199–1201.

Ali, M. and **Sadrul Islam, A.** (**2006**). "Study on main flow and fuel injector configurations for Scramjet applications". *International Journal of Heat and Mass Transfer*, **vol. 49**, pp. 3634–3644.

Anderson, G.; **McClinton, C.R.** and **Weidner, J.** (**2000**). *Scramjet Propulsion*, chap. Scramjet Performance, (pp. 369–446). American Institute of Aeronautics and Astronautics, Inc.

Anderson, J.D. (**2003**). *Modern Compressible Flow: with Historical Perspective*. McGraw-Hill, Boston, MA, USA, 3rd ed.

Anderson, J.D. (**2006**). *Hypersonic and High-Temperature Gas Dynamics*. American Institute of Aeronautics and Astronautics, Inc., 2nd ed.

Axdahl, E.; **Kumar, A.** and **Wilhite, A.** (**2012**). "Study of Forebody Injection and Mixing with Application to Hypervelocity Airbreathing Propulsion". In *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. AIAA 2012-3924, Atlanta GA.

Banks, J.; **Aslam, T.** and **Rider, W.** (**2008**). "On sub-linear convergence for linearly degenerate waves in capturing schemes". *Journal of Computational Physics*, **vol. 227**, doi:10.1016/j.jcp.2008.04.002.

Barth, J.E.; **Wheatley, V.** and **Smart, M.K.** (**2012**). "Streamwise Porthole Fuel Injection for Boundary-Layer Combustion Inside a Scramjet Engine". In *Poster presented to the 28th International Congress of the Aeronautical Sciences*. Brisbane, Australia.

Barth, J.E.; **Wheatley, V.** and **Smart, M.K.** (**2013**). "Hypersonic Turbulent Boundary-Layer Fuel Injection and Combustion: Skin-Friction Reduction Mechanisms". *AIAA Journal*, **vol. 51, no. 9**, pp. 2147–2157.

Baurle, R. and **Edwards, J.** (**2010**). "Hybrid Reynolds-Averaged/Large-Eddy Simulations of a Coaxial Supersonic Freejet Experiment". *AIAA Journal*, **vol. 48, no. 3**, pp. 551–571.

Ben-Yakar, A. and **Hanson, R.** (**2001**). "Cavity Flame-Holders for Ignition and Flame Stabilization in Scramjets: An Overview". *Journal of Propulsion and Power*, **vol. 17, no. 4**, pp. 869–877.

Ben-Yakar, A.; **Mungal, M.** and **Hanson, R.** (**2006**). "Time evolution and mixing characteristics of hydrogen and ethylene transverse jets in supersonic crossflows". *Physics of Fluids*, **vol. 18, no. 026101**, pp. 1–16.

Beresh, S.J.; **Henfling, J.F.**; **Erven, R.J.** and **Spillers, R.W.** (**2007**). "Vortex Structure Produced by a Laterally Inclined Supersonic Jet in Transonic Crossflow". *Journal of Propulsion and Power*, **vol. 23, no. 2**, pp. 353–363.

Berglund, M.; **Fedina, E.**; **Fureby, C.**; **Tegner, J.** and **Sabel'nikov, V.** (**2010**). "Finite Rate Chemistry Large-Eddy Simulation of Self-Ignition in a Supersonic Combustion Ramjet". *AIAA Journal*, **vol. 48, no. 3**, pp. 540–550.

Berman, H.A.; **Anderson, J.D.J.** and **Drummond, J.P.** (**1983**). "Supersonic Flow over a Rearward Facing Step with Transverse Nonreacting Hydrogen Injection". *AIAA Journal*, **vol. 21, no. 12**, pp. 1707–1713.

Billig, F. (**1992**). "Research on Supersonic Combustion". In *30th Aerospace Sciences Meeting and Exhibit*, (pp. 1–20). No. AIAA-92-0001 in 30th Aerospace Sciences Meeting & Exhibit, AIAA, AIAA.

Billig, F. and **Jacobsen, L.** (**2003**). "Comparison of Planar and Axisymmetric Flowpaths for Hydrogen Fueled Space Access Vehicles (Invited)". In *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, (pp. 1–12). AIAA 2003-4407.

Billig, F.; **Orth, R.** and **Lasky, M.** (**1971**). "A Unified Analysis of Gaseous Jet Penetration". *AIAA Journal*, **vol. 9, no. 6**, pp. 1048–1058.

Blake, L. (**1970**). "Approximate Transport Calculations for High-Temperature Air". *AIAA Journal*, **vol. 8, no. 9**, pp. 1698–1701.

Bonanos, A.M.; **Bergthorson, J.M.** and **Dimotakis, P.E.** (**2009**). "Molecular Mixing and Flowfield Measurements in a Recirculating Shear Flow. Part II: Supersonic Flow". *Flow, Turbulence and Combustion*, **vol. 83**, pp. 251–268.

Boyce, R.; **Takahashi, M.** and **Stalker, R.** (**2005**). "Mass Spectrometric Measurements of Driver Gas Arrival in the T4 Free-Piston Shock-Tunnel". *Shock Waves*, **vol. 14, no. 5/6**, pp. 371–378.

Brieschenk, S.; **Kleine, H.** and **O'Byrne, S.** (**2013**). "Laser ignition of hypersonic air-hydrogen flow". *Shock Waves*, **vol. 23, no. 5**, pp. 439–452.

Burrows, M. and **Kurkov, A.** (**1973**). "Analytical and experimental stufy of supersonic combustion of hydrogen in a vitiated airstream". Tech. Rep. Technical Report TM X-2828, NASA.

Burtschell, Y. and **Zeitoun, D.** (**2004**). "Numerical investigation of H2 injection in Mach 5 air flow with a strong shock/boundary layer interaction". *Shock Waves*, **vol. 13**, pp. 465–472.

Bushnell, D.M. and **Hefner, J.** (**1990**). *Viscous Drag Reduction in Boundary Layers*, vol. 123, Progress in Aeronautics and Astronautics. American Institute of Aeronautics and Astronautics, Inc., Washington DC, USA.

Cary, A. and **Hefner, J.** (**1972**). "Film-Cooling Effectiveness and Skin Friction in Hypersonic Turbulent Flow". *AIAA Journal*, **vol. 10, no. 9**, pp. 1188–1193.

Catris, S. and **Aupoix, B.** (**2000**). "Density corrections for turbulence models". *Aerospace Science and Technology*, **vol. 4**, pp. 1–11.

Chan, W.Y.K. (**2012**). *Effects of flow non-uniformities on the drag reduction by boundary layer combustion*. Ph.D. thesis, The University of Queensland.

Chapuis, M.; **Fedina, E.**; **Fureby, C.**; **Hannemann, K.**; **Karl, S.** and **Martinez Schramm, J.** (**2013**). "A computational study of the HyShot II combustor performance". *Proceedings of the Combustion Institute*, **vol. 34**, pp. 2101–2109.

Croker, B. (**2007**). "On the design of hypersonic inward-turing inlets". Tech. Rep. AFRL-RB-WP-TP-2009-3016, Air Force Research Laboratory, Air Force Research Laboratory Air Vehicles Directorate Wright-Patterson Air Force Base, OH 45433-7542 Air Force Materiel Command United States Air Force.

Curran, E. (**2001**). "Scramjet Engines: The First Forty Years". *Journal of Propulsion and Power*, **vol. 17, no. 6**, pp. 1138–1148.

Curran, E.; **Heiser, W.** and **Pratt, D.** (**1996**). "Fluid Phenomena in Scramjet Combustion Systems". *Aunnal Review in Fluid Mechanics*, **vol. 28**, pp. 323–360.

Czysz, P.A. and **Bruno, C.** (**2009**). *Future Spacecraft Propulsion Systems: Enabling Technologies for Space Exploration*. Berlin, Heidelberg : Springer Berlin Heidelberg, 2 ed.

Denman, A. (**2007**). *Large-Eddy Simulation of Compressible Turbulent Boundary Layers with Heat Addition*. Ph.D. thesis, The University of Queensland.

Dimotakis, P.E. (**1991**). *High-Speed Flight Propulsion Systems, Process in Astronautics and Aeronautics Series*, (pp. 265–340). American Institute of Aeronautics and Astronautics, Inc., Washington DC.

Doherty, L.J. (**2013**). *Experimental Investigation of an Airframe Integrated 3-D Scramjet at a Mach 10 Flight Condition*. Ph.D. thesis, The University of Queensland.

Doherty, L.J. and **Wise, D.J.** (**2013**). "Pitot Survey Data for the T4 Mach 10 Nozzle". Tech. Rep. 2013/07, The Universit of Queensland School of Mechanical and Mining Engineering.

Doherty, L.J.; **Zander, F.**; **Jacobs, P.A.**; **Gollan, R.J.**; **Chan, W.Y.K.** and **Kirchhartz, R.M.** (**2012**a). "NENZF-r: Non-Equilibrium Nozzle Flow, Reloaded. A User Guide." Tech. Rep. 2012/08, The University of Queensland School of Mechanical and Mining Engineering.

Doherty, L.J.; **Smart, M.** and **Mee, D.J.** (**2012**b). "Design of an Airframe Integrated 3-D Scramjet and Experimental Results at a Mach 10 Flight Condition". In *Proceedings of the 18th AIAA/3AF International Space Planes and Hypersonic Systems and Technologies Conference, Tours, France, 24-28 September*. AIAA Paper 2012-5910.

Donohue, J.M.; **McDaniel, J.C.** and **Haj-Hariri, H.** (**1994**). "Experimental and Numerical Injection into a Supersonic Flowfield". *AIAA Journal*, **vol. 32, no. 9**, pp. 1860–1867.

Dove, J. and **Teitelbaum, H.** (**1974**). "The vibrational relaxation of H2. I. Experimental measurements of the rate of relaxation by H2, He, Ne, Ar, and Kr". *Chemical Physics*, **vol. 6, no. 3**, pp. 431–444.

Drayna, T.W.; **Nompelis, I.** and **Candler, G.V.** (**2006**). "Hypersonic Inward Turning Inlets: Design and Optimization". In *44th AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, USA. 9-12 January*.

Drummond, J.P.; **Diskin, G.S.** and **Cutler, A.D.** (**2002**). "Fuel-Air Mixing and Combustion in Scramjets". In *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. AIAA-2002-3878, Indianapolis USA.

Dunlap, R.; **Brehm, R.** and **Nicholls, J.** (**1958**). "A Prelminary Study of the Application of Steady-State Detonative Combustion to a Reaction Engine". *Journal of Jet Propulsion*, **vol. 28, no. 7**, pp. 451–456.

Edwards, J. and **Chandra, S.** (**1996**). "Comparison of Eddy Viscosity-Transport Turbulence Models for Three-Dimensional, Shock-Separated Flowfields". *AIAA Journal*, **vol. 34, no. 4**, pp. 756–763.

Eklund, D.R.; **Stouffer, S.D.** and **Northam, G.B.** (**1997**). "Study of a Supersonic Combustor Employing Swept Ramp Fuel Injectors". *Journal of Propulsion and Power*, **vol. 13, no. 6**, pp. 697–704.

Ferrante, A.; **Matheou, G.** and **Dimotakis, P.E.** (**2011**). "LES of an inclined sonic jet into a turbulent crossflow at Mach 3.6". *Journal of Turbulence*, **vol. 12, no. 2**, pp. 1–32.

Ferri, A. (**1973**). "Mixing-Controlled Supersonic Combustion". *Annual Review of Fluid Mechanics*, **vol. 5**, pp. 301–338.

Fuller, R.; **Wu, P.**; **Nejad, A.** and **Schetz, J.** (**1998**). "Comparison of Physical and Aerodynamic Ramps as Fuel Injectors in Supersonic Flow". *Journal of Propulsion and Power*, **vol. 14, no. 2**, pp. 135–145.

Gaitonde, D.; **Shang, J.** and **Visbal, M.** (**1995**). "Structure of a Double-Fin Turbulent Interaction at High Speed". *AIAA Journal*, **vol. 33, no. 2**, pp. 193–200.

Gaitonde, D.; **Malo-Molina, F.J.**; **Croker, B.** and **Ebrahimi, H.B.** (**2007**). "The Flowfield Structure in Generic Inward-Turning-Based Propulsion Flowpath Components". In *43rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. AIAA 2007-5744, AIAA.

Gardner, A.; **Paull, A.** and **McIntyre, T.** (**2002**). "Upstream Porthole Injection in a 2-D scramjet model". *Shock Waves*, **vol. 11**, pp. 369–375.

Gehre, R.; **Wheatley, V.** and **Boyce, R.** (**2012**). "Revised model coefficients for vibrational relaxation in a nitrogenâĂŞoxygen gas mixture". *Shock Waves*, **vol. 22, no. 6**, pp. 647–651.

Genin, F. and **Menon, S.** (**2010**). "Simulation of Turbulent Mixing Behind a Strut Injector in Supersonic Flow". *AIAA Journal*, **vol. 48, no. 3**, pp. 526–538.

Gerlinger, P.; **Nold, K.** and **Aigner, M.** (**2010**). "Influence of reaction mechanisms, grid spacing, and inflow conditions on the numerical simulation of lifted supersonic flames". *International Journal for Numerical Methods in Fluids*, **vol. 62**, pp. 1357–1380.

Gerlinger, P.; **Stoll, P.**; **Kindler, M.**; **Schneider, F.** and **Aigner, M.** (**2008**). "Numerical Investigation of Mixing and Combustion Enhancement in Supersonic Combustors by Strut Induced Streamwise Vorticity". *Aerospace Science and Technology*, **vol. 12**, pp. 159–168.

Gollan, R.J. and **Ferlemann, P.G.** (**2011**). "Investigation of REST-class Hypersonic Inlet Designs". In *17th AIAA International Space Planes and Hypersonic Systems and Technologies Conference April 11-14, 2011, San Francisco, CA*.

Gollan, R.J. and **Ferlemann, P.G.** (**2009**). "Parametric Geometry, Structured Grid Generation, and Initial Design Study for REST-Class Hypersonic Inlets". In *JANNAF 43rd Combustion, 31st Airbreathing Propulsion, 25th Propulsion Systems Hazards Joint Subcommittee Meeting, La Jolla CA*, (pp. 1–26).

Gomez, T.; **Flutet, V.** and **Sagaut, P.** (**2009**). "Contribution of Reynolds stress distribution to the skin friction in compressible turbulent channel flows". *Physical Review E*, **vol. 79**, pp. 035301–1 – 035301–4.

Goyne, C. (**1999**). *Skin Friction Measurements in High Enthalpy Flows at High Mach Number*. Ph.D. thesis, The University of Queensland.

Goyne, C.; **Stalker, R.**; **Paull, A.** and **Brescianini, C.** (**2000**). "Hypervelocity Skin-Friction Reduction by Boundary-Layer Combustion of Hydrogen". *Journal of Spacecraft and Rockets*, **vol. 37, no. 6**, pp. 740–746.

GridPro (**2014**). "GridPro Homepage". Software Package, Ver. 5.1.
URL `http://www.gridpro.com`

Gruber, M.; **Nejad, A.**; **Chen, T.** and **Dutton, J.** (**2000**). "Transverse Injection from Circular and Elliptic Nozzles into a Supersonic Crossflow". *Journal of Propulsion and Power*, **vol. 16, no. 3**, pp. 449–457.

Gruber, M.; **Nejad, A.**; **Chen, T.** and **Dutton, J.** (**1995**). "Mixing and Penetration Studies of Sonic Jets in a Mach 2 Freestream". *Journal of Propulsion and Power*, **vol. 11, no. 2**, pp. 315–323.

Gruber, M.; **Nejad, A.**; **Chen, T.** and **Dutton, J.** (**1997**). "Compressibility effects in supersonic transverse injection flowfields". *Physics of Fluids*, **vol. 9, no. 5**, doi:10.1063/1.869257.

Gruhn, P. and **Gülhan, A.M.** (**2011**). "Experimental Investigation of a Hypersonic Inlet with and Without Sidewall Compression". *Journal of Propulsion and Power*, **vol. 27, no. 3**, pp. 718–729.

Guoskov, O.; **Kopchenov, V.**; **Lomkov, K.**; **Vinogradov, V.** and **Waltrup, P.** (**2001**). "Numerical Research of Gaseous Fuel Preinjection in Hypersonic Three-Dimensional Inlet". *Journal of Propulsion and Power*, **vol. 17, no. 6**, pp. 1162–1169.

Häberle, J. and **Gülhan, A.M.** (**2008**). "Experimental Investigation of a Two-Dimensional and a Three-Dimensional Scramjet Inlet at Mach 7". *Journal of Propulsion and Power*, **vol. 24, no. 5**, pp. 1023–1034.

Hariharan, A.R. and **Babu, V.** (**2014**). "Transverse Injection Into a Supersonic Cross Flow Through a Circular Injector With Chevrons". *Journal of Fluids Engineering*, **vol. 136**, pp. 021204–1 to 021204–8.

Heiser, W.H. and **Pratt, D.T.** (**1994**). *Hypersonic Airbreathing Propulsion*. American Institute of Aeronautics and Astronautics, Inc.

Holland, S.D. (**1995**). "Internal Aerodynamics of a Generic Three-Dimensional Scramjet Inlet at Mach 10". Tech. Rep. NASA TP-3476, NASA.

Hornung, H. (**1988**). "28th Lanchester Memorial Lecture - Experimental real-gas hypersonics". *Aeronautical Journal*, **vol. 92(920)**, pp. 379–389.

Huh, H. and **Driscoll, J.** (**1996**). "Shock-Wave-Enhancement of the Mixing and the Stability Limits of Supersonic Hydrogen-Air Jet Flames". In *Twenty-Sixth Symposium (International) on Combustion*, (pp. 2933–2939). The Combustion Institute.

Hunt, J.; **Wray, A.** and **Moin, P.** (**1988**). "Eddies, stream, and convergence zones in turbulent flows". Proceedings of the 1988 Summer Program CTR-S88, Stanford, NASA Centre for Turbulence Research.

Im, H.; **Bechtold, J.** and **Law, C.** (**1993**). "Analysis of thermal ignition in supersonic flat-plate boundary layers". *Journal of Fluid Mechanics*, **vol. 249**, pp. 99–120.

Im, H.; **Helenbrook, B.**; **Lee, S.** and **Law, C.** (**1996**). "Ignition in the supersonic hydrogen/air mixing layer with reduced reaction mechanisms". *Journal of Fluid Mechanics*, **vol. 322**, pp. 275–296.

Ingenito, A. and **Bruno, C.** (**2010**). "Physics and Regimes of Supersonic Combustion". *AIAA Journal*, **vol. 48, no. 3**, pp. 515–525.

Jachimowski, C.J. (**1992**). "An Analysis of Combustion Studies in Shock Expansion Tunnels and Reflectect Shock Tunnels". Tech. Rep. NASA Technical Paper 3224, NASA.

Jacobs, P.A.; **Rogers, R.**; **Weidner, E.** and **Bittner, R.** (**1992**). "Flow Establishment in a Generic Scramjet Combustor". *Journal of Propulsion and Power*, **vol. 8, no. 4**, pp. 890–899.

Kim, J.H.; **Yoon, Y.**; **Jeung, I.S.**; **Huh, H.** and **Choi, J.Y.** (**2003**). "Numerical Study of Mixing Enhancement by Shock Waves in Model Scramjet Engine". *AIAA Journal*, **vol. 41, no. 6**, pp. 1074–1080.

Kirchhartz, R.M. (**2009**). *Upstream Wall Layer Effects on Drag Reduction with Boundary Layer Combustion*. Ph.D. thesis, The University of Queensland.

Kirchhartz, R.M.; **Mee, D.J.**; **Stalker, R.**; **Jacobs, P.A.** and **M.K., S.** (**2010**). "Supersonic Boundary-Layer Combustion: Effects of Upstream Entropy and Shear-Layer Thickness". *Journal of Propulsion and Power*, **vol. 26, no. 1**, pp. 57–66.

Kouchi, T.; **Masuya, G.**; **Hirano, K.**; **Matsuo, A.** and **Tomioka, S.** (**2013**). "Supersonic Combustion Using a Stinger-Shaped Fuel Injector". *Journal of Propulsion and Power*, **vol. 29, no. 3**, pp. 639–647.

Kovachevich, A. (**2010**). *Investigation of an Intake Injected Scramjet with a Hot Wall*. Ph.D. thesis, The University of Queensland.

Kovachevich, A.; **Hajek, K.**; **McIntyre, T.**; **Paull, A.** and **Abdel-jawad, M.** (**2006**). "Imaging of hydrogen fuel injection on the intake of a heated wall scramjet". In *Proceedings of the 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, Sacramento, CA*. AIAA 2006-5039.

Kroo, I. (**1997**). "Standard Atmosphere Computations".
URL `aero.stanford.edu/stdatm.html`

Kulgein, N. (**1962**). "Transport Processes in a Combustible Turbulent Boundary Layer". *Journal of Fluid Mechanics*, **vol. 12**, pp. 417–437.

Kutschenreuter, P. (**2000**). *Scramjet Propulsion*, chap. Supersonic Flow Combustors, (pp. 513–598). American Institute of Aeronautics and Astronautics, Inc.

Leachman, J.; **Jacobsen, R.**; **Penoncello, S.** and **Lemmon, E.** (**2009**). "Fundamental equations of state for parahydrogen, normal hydrogen, and orthohydrogen". *Journal of Physical and Chemical Reference Data*, **vol. 38, no. 3**, pp. 721–748.

Lee, S. (**2006**). "Characteristics of Dual Transverse Injection in Scramjet Combustor". *Journal of Propulsion and Power*, **vol. 22**, pp. 1012–1019.

Levin, V. and **Larin, O.** (**2003**). "Skin Friction Reduction by Energy Addition into a Turbulent Boundary Layer". In *41st Aerospace Sciences Meeting and Exhibit*. AIAA paper 2003-0036.

Lorrain, P.; **Brieschenk, S.**; **Capra, B.** and **Boyce, R.** (**2012**). "A detailed investigation of nominally 2-D radical-farming scramjet combustion". In *18th AIAA/3AF International Space Planes and Hypersonic Systems and Technologies Conference*. Tours, France.

Ma, K.L.; **van Rosendal, J.** and **Vermeer, W.** (**1996**). "3D shock wave visualization on unstructured grids". In *Proceedings of the 1996 symposium on Volume Visualization*. IEEE, IEEE Press, 87-94,104.

Malo-Molina, F.J. (**2014**). "Three Dimensional Analysis of a Single Hydrocarbon-Air Pulse Detonation in a Supersonic Combustor". In *Proceedings of the 52nd Aerospace Sciences Meeting, AIAA SciTech*. AIAA 2014-0626, National Harbor, MD, USA.

Malo-Molina, F.J.; **Gaitonde, D.V.**; **Ebrahimi, H.B.** and **Ruffin, S.M.** (**2010**). "Three-Dimensional Analysis of a Supersonic Combustor Coupled to Innovative Inward-Turning Inlets". *AIAA Journal*, **vol. 48, no. 3**, pp. 572–582.

Marshall, L.A.; **Bahm, C.**; **Corpening, G.P.** and **Sherrill, R.** (**2005**). "Overview With Results and Lessons Learned of the X-43A Mach 10 Flight". In *AIAA/CIRA 13th International Space Planes and Hypersonics Systems and Technologies Symposium*. AIAA 2005-3336.

McBride, B.; **Gordon, S.**; **Svehla, R.**; **Reno, M.** and **Zehe, M.** (**2002**). "Chemical Equilibrium with Applications".
URL `http://www.grc.nasa.gov/WWW/CEAWeb/ceaWhat.htm`

McClinton, C.R. (**1972**). "The effect of Injection Angle on the Interaction Between Sonic Secondary Jets and a Supersonic Freestream". NASA Technical Note TN D-6669, NASA.

McGilvray, M. (**2008**). *Scramjet Testing at High Enthalpies in Expansion Tube Facilities*. Ph.D. thesis, University of Queensland.

McGilvray, M. and **Morgan, R.** (**2009**). "Effects of Upstream Injection on Scramjet Performance Using an Entropy-Based Method". *Journal of Propulsion and Power*, **vol. 25, no. 2**, doi:10.2514/1.38325.

McGuire, J. (**2007**). *Ignition Enhancement for Scramjet Combustion*. Ph.D. thesis, University of New South Wales.

McMurtry, P.; **Riley, J.** and **Metcalfe, R.** (**1989**). "Effects of heat release on the large-scale structure in turbulent mixing layers". *Journal of Fluid Mechanics*, **vol. 199**, pp. 297–332.

Meador, W. and **Smart, M.** (**2005**). "Reference Enthalpy Method Developed from Solutions of the Boundary-Layer Equations". *AIAA Journal*, **vol. 43, no. 1**, pp. 135–139.

Mee, D.J. (**1993**). "Uncertainty Analysis of Conditions in the Test Section of the T4 Shock Tunnel". Tech. Rep. Report No. 4/93, Department of Mechanical Engineering, The University of Queensland.

Mehta, U. (**1996**). "Strategy for Developing Air-Breathing Aerospace Planes". *Journal of Aircraft*, **vol. 33, no. 2**, pp. 377–385.

Millikan, R. and **White, D.** (**1963**). "Systematics of Vibrational Relaxation". *Journal of Chemical Physics*, **vol. 39, no. 12**, pp. 3209–3213.

Mo, J.; **Xu, J.**; **Gu, R.** and **Fan, Z.** (**2014**). "Design of an Asymmetric Scramjet Nozzle with Circular to Rectangular Shape Transition". *Journal of Propulsion and Power*, **vol. 30, no. 3**, pp. 812–819.

Mölder, S. and **Szprio, E.** (**1966**). "Busemann Inlet for Hypersonic Speeds". *Journal of Spacecraft*, **vol. 3, no. 8**, pp. 1303–1304.

Moule, Y.; **Sabel'nikov, V.**; **Mura, A.** and **M.K., S.** (**2014**). "Computational Fluid Dynamics Investigation of a Mach 12 Scramjet Engine". *Journal of Propulsion and Power*, **vol. 30, no. 2**, pp. 461–473.

Moule, Y. and **Smart, M.** (**2013**). "Performance Analysis of a Mach 12 Scramjet at Off-Design Conditions". *Journal of Propulsion and Power*, **vol. 1, no. 1**, pp. 282–285.

NOAA (**1976**). "The U.S. Standard Atmosphere". Tech. rep., U.S. Government Printing Office, Washington, DC, USA.

Nompelis, I. (**2004**). *Computational Study of Hypersonic Double-Cone Experiments for Code Validation*. Ph.D. thesis, The University of Minnesota.

Nompelis, I.; **Candler, G.V.** and **Drayna, T.W.** (**2004**). "Development of a Hybrid Unstructured Implicit Solver for the Simulation of Reacting Flows Over Complex Geometries". In *34th AIAA Fluid Dynamics Conference and Exhibit, Portland, Oregon*.

Oberkampf, W. and **Roy, C.J.** (**2010**). *Verification and Validation in Scientific Computing*, chap. 8, (pp. 286–342). Cambridge University Press.

Ogawa, H. and **Boyce, R.** (**2012**). "Multi-Objective Design Optimization of Fuel Injection for Mixing Enhancement in Scramjets by Using Surrogate-Assisted Evolutionary Algorithms". In *Proceedings of the 18th AIAA International Space Planes and Hypersonic Systems and Technologies Conference*. AIAA 2012-5815.

Park, C.; **Howe, J.**; **Jaffe, R.** and **Candler, G.V.** (**1994**). "Review of Chemical-Kinetic Problems of Future NASA Missions, II: Mars Entries". *Journal of Thermophysics and Heat Transfer*, **vol. 8, no. 1**, pp. 9–23.

Parthasarathy, K. and **Zakkay, V.** (**1970**). "An Experimental Investigation of Turbulent Slot Injection at Mach 6". *AIAA Journal*, **vol. 8, no. 7**, pp. 1302–1307.

Paull, A.; **Stalker, R.** and **Mee, D.J.** (**1995**). "Experiments on supersonic ramjet propulsion in a shock tunnel". *Journal of Fluid Mechanics*, **vol. 296**, pp. 159–183.

PDC (**1999**). *GridPro v4.1 User's Guide and Reference Manual for TIL, Ggrid And Other Utilities*. Program Development Corporation, White Plains, NY.

Pergament, H. (**1963**). "Theoretical analysis of non-equilibrium hydrogen-air reactions in flow systems". In *AIAA-ASME Hypersonic Ramjet Conference*, vol. 63113.

Peterson, D.M.; **Boyce, R.** and **Wheatley, V.** (**2013**). "Simulations of Mixing in an Inlet-Fueled Axisymmetric Scramjet". *AIAA Journal*, **vol. 51, no. 12**, pp. 2823–2832.

Peterson, D.M. and **Candler, G.V.** (**2011**). "Simulations of Mixing for Normal and Low-Angled into Supersonic Crossflow". *AIAA Journal*, **vol. 49, no. 12**, pp. 2792–2804.

Pečnik, R.; **Terrapon, V.**; **Ham, F.**; **Iaccarino, G.** and **Pitsch, H.** (**2012**). "Reynolds-Averaged Navier-Stokes Simulations of the HyShot II Scramjet". *AIAA Journal*, **vol. 50, no. 8**, pp. 1717–1732.

Pointwise (**2014**). "Pointwise Homepage". Software Package, V17.0.
URL `http://www.pointwise.com`

Portz, R. and **Segal, C.** (**2006**). "Penetration of Gaseous Jets In Supersonic Flows". *AIAA Journal*, **vol. 44, no. 10**, pp. 2426–2429.

Pudsey, A. and **Boyce, R.** (**2010**). "Numerical Investigation of Transverse Jets Through Multiport Injector Arrays in Supersonic Crossflow". *Journal of Propulsion and Power*, **vol. 26, no. 6**, pp. 1225–1236.

Pulsonetti, M. (**1995**). "Experimental Methods for a Scramjet Scaling Study". Tech. Rep. Report No. 7/95, Department of Mechanical Engineering, The University of Queensland.

Pulsonetti, M.; **Erdos, J.** and **Early, K.** (**1988**). "An engineering model for analysis of scramjet combustor performance with finite rate chemistry". In *Presented at the 24th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. AIAA Paper 88-3258, Boston MA.

Razzaqi, S. (**2011**). *Oxygen Enrichment in a Hydrogen Fuelled Scramjet*. Ph.D. thesis, The University of Queensland.

Roache, P. (**1994**). "Perspective: A Method for Uniform Reporting of Grid Refinement Studies". *Journal of Fluids Engineering*, **vol. 116**, pp. 405–413.

Rosenberg, Z. (**2013**). "Hypersonic X-51 programme ends in success". Flight International. Accessed online http://www.flightglobal.com/news/articles/hypersonic-x-51-programme-ends-in-success-385481/.

Rowan, S.A. (**2003**). *Viscous drag reduction in a scramjet combustor*. Ph.D. thesis, The University of Queensland.

Rowan, S.A. and **Paull, A.** (**2006**). "Performance of a Scramjet Combustor with Combined Normal and Tangential Fuel Injection". *Journal of Propulsion and Power*, **vol. 22, no. 6**, pp. 1334–1338.

Roy, C.J. (**2003**). "Grid Convergence Error Analysis for Mixed-Order Numerical Schemes". *AIAA Journal*, **vol. 41, no. 4**, pp. 595–604.

Schetz, J. and **Billig, F.** (**1966**). "Penetration of Gaseous Jets Injected into a Supersonic Stream". *Journal of Spacecraft and Rockets*, **vol. 3, no. 11**, pp. 1658–1665.

Schetz, J.; **Maddalena, L.** and **Burger, S.K.** (**2010**). "Molecular Weight and Shock-Wave Effects on Transverse Injection in Supersonic Flow". *Journal of Propulsion and Power*, **vol. 26, no. 5**, pp. 1102–1113.

Schetz, J. and **van Overeem, J.** (**1975**). "Skin Friction Reduction by Injection Through Combinations of Slots and Porous Sections". *AIAA Journal*, **vol. 13, no. 8**, pp. 971–972.

Schumacher, J. (**2000**). *Numerical Simulation of Cantilevered Ramp Injector Flow Fields for Hypervelocity Fuel/Air Mixing Enhancement*. Ph.D. thesis, University of Toronto.

Schwartzentruber, T.; **Sislian, J.** and **Parent, B.** (**2005**). "Suppression of Premature Ignition in the Premixed Inlet Flow of a Shcramjet". *Journal of Propulsion and Power*, **vol. 21, no. 1**, pp. 87–94.

Segal, C. (**2009**). *The Scramjet Engine: Processes and Characteristics*. Cambridge University Press.

Segal, C.; **McDaniel, J.**; **Whitehurst, R.** and **Krauss, R.** (**1995**). "Mixing and Chemical Kinetics Interactions in a Mach 2 Reacting Flow". *Journal of Propulsion and Power*, **vol. 11, no. 2**, pp. 308–314.

Seiner, J.; **Dash, S.** and **Kenzakowski, D.** (**2001**). "Historical Survey on Enhanced Mixing in Scramjet Engines". *Journal of Propulsion and Power*, **vol. 17, no. 6**, pp. 1273–1285.

Sislian, J. (**2000**). *Scramjet Propulsion: Detonation-Wave Ramjets*, vol. 189 of *Progress in Astronautics and Aeronautics*, chap. 13, (pp. 823–889). AIAA, Reston, VA.

Smart, M. (**1999**). "Design of Three-Dimensional Hypersonic Inlets with Rectangular-to-Elliptical Shape Transition". *Journal of Propulsion and Power*, **vol. 15, no. 3**, pp. 408–416.

Smart, M. (**2001**). "Experimental Testing of a Hypersonic Inlet with Rectangular-to-Elliptical Shape Transition". *Journal of Propulsion and Power*, **vol. 17, no. 2**, pp. 276–283.

Smart, M. (**2007**). "Scramjets". *The Aeronautical Journal of the Royal Aeronautical Society*, **vol. 111**, pp. 605–619.

Smart, M. (**2012**). "How Much Compression Should a Scramjet Inlet Do?" *AIAA Journal*, **vol. 50, no. 3**, pp. 610–619.

Smart, M. and **Ruf, E.** (**2006**). "Free-jet Testing of a REST Scramjet at Off-Design Conditions". In *25th AIAA Aerodynamic Measurement Technology and Ground Testing Conference, San Francisco CA*. AIAA 2006-2955.

Smart, M. and **Tetlow, M.** (**2009**). "Orbital Delivery of Small Payloads Using Hypersonic Airbreathing Propulsion". *Journal of Spacecraft and Rockets*, **vol. 46, no. 1**, pp. 117–125.

Smart, M. and **Trexler, C.** (**2004**). "Mach 4 Performance of Hypersonic Inlet with Rectangular-to-Elliptical Shape Transition". *Journal of Propulsion and Power*, **vol. 20, no. 2**, pp. 288–293.

Smart, M. and **White, J.A.** (**2002**). "Computational Investigation of the Performance and Back-Pressure Limits of a Hypersonic Inlet". In *40th AIAA Aerospace Sciences Meeting & Exhibit, Reno NV*. AIAA 2002-0508.

Spalart, P. and **Allmaras, S.** (**1992**). "A One-Equation Turbulence Model for Aerodynamic Flows". In *30th Aerospace Sciences Meeting & Exhibit*, (pp. 1–22). AIAA-92-0439, AIAA, AIAA.

Srinivasan, R.; **Stoll, P.**; **Kindler, M.**; **Schneider, F.** and **Aigner, M.** (**2008**). "Transverse Injection Through Diamond and Circular Ports into a Mach 5.0 Freestream". *AIAA Journal*, **vol. 46, no. 8**, pp. 1944–1962.

Stalker, R. (**2005**). "Control of Hypersonic Turbulent Skin Friction by Boundary-Layer Combustion of Hydrogen". *Journal of Spacecraft and Rockets*, **vol. 42, no. 4**, pp. 577–587.

Stalker, R.; **Paull, A.**; **Mee, D.J.**; **Morgan, R.** and **Jacobs, P.A.** (**2005**). "Scramjets and Shock Tunnels: The Queensland Experience". *Progress in Aerospace Sciences*, **vol. 41**, pp. 471–513.

Stern, F.; **Wilson, R.**; **Coleman, H.** and **Paterson, E.** (**2001**). "Comprehensive Approach to Verification and Validation of CFD Simulations - Part 1: Methodology and Procedures". *Journal of Fluids Engineering*, **vol. 123**, pp. 793–802.

Suraweera, M.V. (**2006**). *Reduction of Skin Friction Drag in Hypersonic Flow by Boundary Layer Combustion*. Ph.D. thesis, The University of Queensland.

Suraweera, M.V. and **Smart, M.** (**2009**). "Shock-Tunnel Experiments with a Mach 12 Rectangular-to-Elliptical Shape-Transition Scramjet at Offdesign Conditions". *Journal of Propulsion and Power*, **vol. 25, no. 3**, pp. 555–564.

Tanimizu, K. (**2008**). *Nozzle optimization study and measurements for a qasi-axisymmetric scramjet model*. Ph.D. thesis, The University of Queensland.

Tanimizu, K.; **Mee, D.J.**; **Stalker, R.** and **Jacobs, P.A.** (**2009**). "Drag force on quasi-axisymmetric scramjets at various flight Mach numbers: theory and experiment". *Shock Waves*, **vol. 19**, pp. 83–93.

Townend, L. (**2001**). "Domain of the Scramjet". *Journal of Propulsion and Power*, **vol. 17, no. 6**, pp. 1205–1213.

Turner, J.C. (**2010**). *An Experimental Investigation of Inlet Fuel Injection in a Three-Dimensional Scramjet Engine*. Ph.D. thesis, The University of Queensland.

Turner, J.C. and **Smart, M.** (**2013**). "Mode Change Characteristics of a Three-Dimensional Scramjet at Mach 8". *Journal of Propulsion and Power*, **vol. 29, no. 4**, doi:10.2514/1.B34569.

Turner, J.C. and **Smart, M.** (**2010**). "Application of Inlet Injection to a Three-Dimensional Scramjet at Mach 8". *AIAA Journal*, **vol. 48, no. 4**, pp. 829–837.

Van Driest, E. (**1951**). "Turbulent Boundary Layer in Compressible Fluids". *Journal of the Aeronautical Sciences*, **vol. 18, no. 3**, pp. 145–160.

Van Driest, E. (**1956**). "The Problem of Aerodynamic Heating". *Aeronautical Engineering Review*, **vol. 15, no. 10**, pp. 26–41.

Van Wie, D.M. (**2000**). *Scramjet Propulsion*, chap. 7 - Scramjet Inlets, (pp. 447–511). American Institute of Aeronautics and Astronautics, Inc.

Viti, V.; **Neel, R.** and **Schetz, J.** (**2009**). "Detailed flow physics of the supersonic jet interaction flow field". *Physics of Fluids*, **vol. 21, no. 046101**, pp. 1–16.

Warwick, G. (**26 May 2010**). "First X-51A hypersonic flight deemed success." Aviation Week.

Watanabe, J.; **Kouchi, T.**; **Takita, K.** and **Masuya, G.** (**2011**). "Numerical Study on the Turbulent Structure of Transverse Jet into Supersonic Flow". *AIAA Journal*, **vol. 49, no. 9**, pp. 2057–2067.

Weber, R. and **MacKay, J.** (**1958**). "An Analysis of Ramjet Engines Using Supersonic Combustion". Tech. Rep. TN 4386, NACA, Lewis Flight Propulsion Laboratory, Cleveland, OH.

Wheatley, V. (**2014**). "Schematic of a typical scramjet". Figure.

Wheatley, V. and **Jacobs, P.A.** (**2010**). "Fuel injection via rectangular cross-section injectors for mixing enhancement in scramjets". In *17th Australiasian Fluid Mechanics Conference, Auckland NZ*.

Wise, D.J. (**2014**). *Experimental Investigation of a 3D Scramjet Engine at Hypervelocity Conditions*. Ph.D. thesis, The University of Queensland.

Wooldridge, C. and **Muzzy, R.** (**1965**). "Measurements in a Turbulent Boundary Layer with Porous Wall Injection and Combustion". In *Proceedings of the Tenth International Symposium on Combustion, Pittsburgh, PA*, (pp. 1351–1362). Combustion Institute.

Xiao, X.; **Hassan, H.** and **Baurle** (**2007**). "Modeling scramjet flows with variable turbulent Prandtl and Schmidt numbers". *AIAA Journal*, **vol. 45, no. 6**, pp. 1415–1423.

Zeldovich, Y. (**1951**). "On the Theory and Combustion of Initially Unmixed Gases". NACA Technical Memorandum NACA TM1296, National Advisory Committee for Aeronautics.

Zheltovodov, A. (**1996**). "Shock waves/turbulent boundary-layer interactions - Fundamental studies and applications". In *27th AIAA Fluid Dynamics Conference*. New Orleans LA.

Appendix A

Mesh Generation for a REST engine

Creating a mesh of sufficient quality to accurately simulate a REST-class engine at the Mach 12 flow conditions encountered is a non-trivial problem. The highly swept leading edges, the internal transition from a rectangular cross-section to an elliptical shape, and the inclusion of fuel injector geometry are individually complicated problems; combined they pose an even greater challenge.

A set of best-practice guidelines for meshing such a geometry has been developed over the course of this research, and is presented for the aid of anyone who might wish to simulate a REST-class engine of their own design. These guidelines assume grid generation is being done using the GridPro software package (GridPro, 2014), and that the end-user requires a wall-resolved viscous mesh for their solution (i.e. no wall functions will be used). As such, some basic familiarity with the operation of GridPro is also assumed.

A.1 Surface Preparation

It is worthwhile taking time to appropriately prepare the engine surfaces to ensure adequate smoothness and rounded features. GridPro typically takes in an .stl formatted geometry file for user-defined surfaces, which in turn provides a reference for all the generated blocks that border the loaded surface. Block corners linked to the surface will often 'lock' to regions where there is a sudden change in size and/or aspect ratio of the triangular mesh that defines the surface contour. This in turn can lead to bad blocks of topology with negative cell volumes at worst, or a highly skewed mesh topology at best. While sharp surface features require a high density of triangular elements to resolve properly, CAD or other geometry generation methods may require some manipulation to avoid a sudden change surface in element size as a smooth surface gives way to a detailed feature.

One simple method for generating a good quality mesh is to create a surface mesh in another meshing utility, such as Pointwise (2014). Routines that place more emphasis on user action may make the generation of a complex 3D grid onerous, but it has been found that a Pointwise generated surface mesh can quickly create a high-quality surface file which, when exported as a Nastran-format file, can easily be read into GridPro as a high-quality surface.

A.1.1 Smoothing sharp features

Of course, a well defined surface file can still cause difficulties in mesh generation if the surface involves highly discontinuous features such as: sharp leading edges on an inlet, the lip of a porthole injector, or the lip of a backward-facing step; all of which are assumed in the initial design of a REST flow-path employing both inlet and combustor boundary-layer injection. While it is not impossible to generate a good-quality mesh in GridPro with these sort of sharp surface features, doing so tends to require a significant level of complexity in terms of additional construction surface constraints, and typically a large increase in the number of required topology blocks. In the absence of these more extreme measures, the sharp feature will tend to take on a non- uniform curvature, typically in staggered one- or two-cell increments.

The introduction of a small amount of curvature can significantly ease the level of gridding complexity required, while still achieving a high-quality mesh. A good 'rule of thumb' for such rounding of sharp features is to keep the radius of curvature well below 1mm. Applying a 1mm curvature has been shown to have a negligible impact on inlet performance for streamtraced inlets of a similar scale, parameters such as total pressure recovery and mass capture are nearly identical to that of the the sharp-edge geometry(Drayna et al., 2006). As a lower bound, the surface need not employ curvatures less than that of a typical tool-nose radius: smaller curvatures would not only require a prohibitively large number of cells to adequately model, they would also be impossible to manufacture. As such, there is little to be gained in going to curvatures less than 0.1mm, and no advantages in curvatures less than approximately 25-50 μm.

A.1.2 M12 REST Surface Modifications

To illustrate how such guidelines may be applied to REST engine surface specification, the M12 REST engine studied can provided examples of surface modifications. One

simple example worth mentioning is applying a 0.25 mm radius to the internal corners of the inlet: rather than blending from a sharp 90° corner to the smooth elliptical surface, the issue becomes a matter of changing the radius of curvature as the distinct body, side, and cowl walls coalesce into a single elliptical bounding surface.

The leading edges are the most significant regions of the engine where smoothing must be applied to a sharp edge. The original CAD edge topology featured a "half-blunt" leading edge profile, in which the internal surface terminated in a sharp corner, while the external surface blended outward along a radius of approximately 1mm, as shown in the dashed lines in Figure A.1.

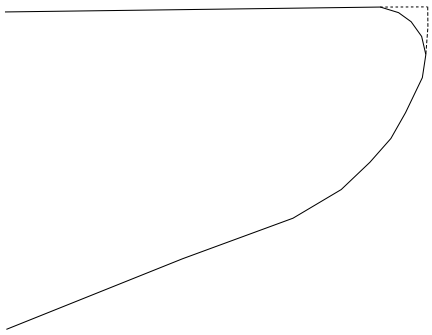

Figure A.1: Profile of M12 REST rounded leading edges (original half-blunt profile shown in dashed lines)

The intent of this design was to minimize the loss in inlet performance while providing a finite leading edge radius to avoid the grid generation difficulties and thermal protection issues that a sharp feature often causes. However, attempting to mesh this topology with GridPro and achieve a satisfactory mesh proved difficult. The solution was simple: by applying a 0.25 mm radius to this sharp edge, the surface mesh would be easily able to wrap around the topology, requiring no special handling. There is no reason a REST inlet couldn't use a symmetric radius on the leading edge: while it would slightly increase the effective capture area of the inlet (and push the location of the leading edges slightly downstream), following the guidelines above should not result in any discernible performance loss in the inlet.

Of course, simply applying this radius will not completely solve gridding issues. The Side-Cowl Corner and the Cowl Notch both are simple to construct with a sharp leading edge, but require further design when moving to a rounded radius. Examining the Cowl Notch first, it becomes obvious that special blending must be done to a rounded leading edge, as simply rounding the corner leads to a recessed sharp feature that would be just as likely to experience severe heating loads as a sharp edge would be. Furthermore, when constructing topology in GridPro, a block placed in the sharp

corner shown in dashed line in Figure A.2 would lead to a topological singularity that would result in a highly skewed mesh in this region.

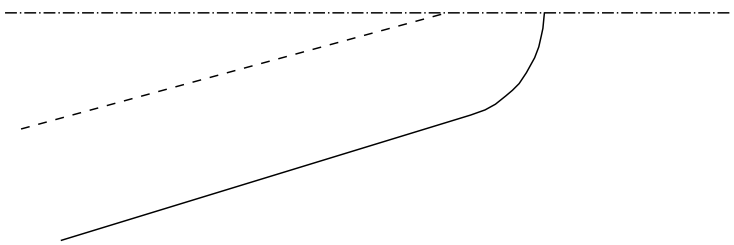

Figure A.2: Difference between sharp (dashed) and rounded cowl notch (solid line) geometries. Flow direction is left to right.

As such, introducing a curvature to the notch itself is desirable. Shown in Figure A.2, this curvature allows for an appropriately blunted leading edge at the point of cowl closure, as well as allowing for the grid block topology to terminate on the center plane in a smooth fashion. The radius of curvature shown of the rounded notch is 30 μm. Much like rounding on the leading edge, the cowl closure point is pushed slightly back, but once again the impact is likely to be minimal with an appropriate selection in leading edge radius.

Turning attention to the Side-Cowl corner, it becomes immediately obvious that a number of competing factors must be accounted for: not only are two leading edges intersecting at an oblique angle, they must also blend to the internal curvature applied to the corner that was mentioned at the start of this section, as illustrated in Figure A.3. A few guidelines may provide the designer with a relatively easy path for generating this feature.

The first suggestion is that blending begin at the point where the internal curvature becomes tangent to the cowl leading edge curvature, as shown in Figure A.3b above. This serves to preserve the internal grid block topology for as long as possible. Upstream of this point, the internal and leading-edge curvatures slowly blend and turn outward to match the side wall leading edge curvature. Care must be taken in creating this blended surface: an early attempt at achieving this blending unintentionally created a small obstruction on the external wall that hindered the spilling of the boundary layer outward away from the inlet. This in turn allowed for the formation of a subsonic recirculation on the external corner, which quickly grew to unstart the inlet!

There may be other sharp features needing some smoothing as well, depending on the engine design. As will be shown in Section A.6, it is often desirable to round the corner of an injector exit, allowing the boundary layer cells to wrap into the injector.

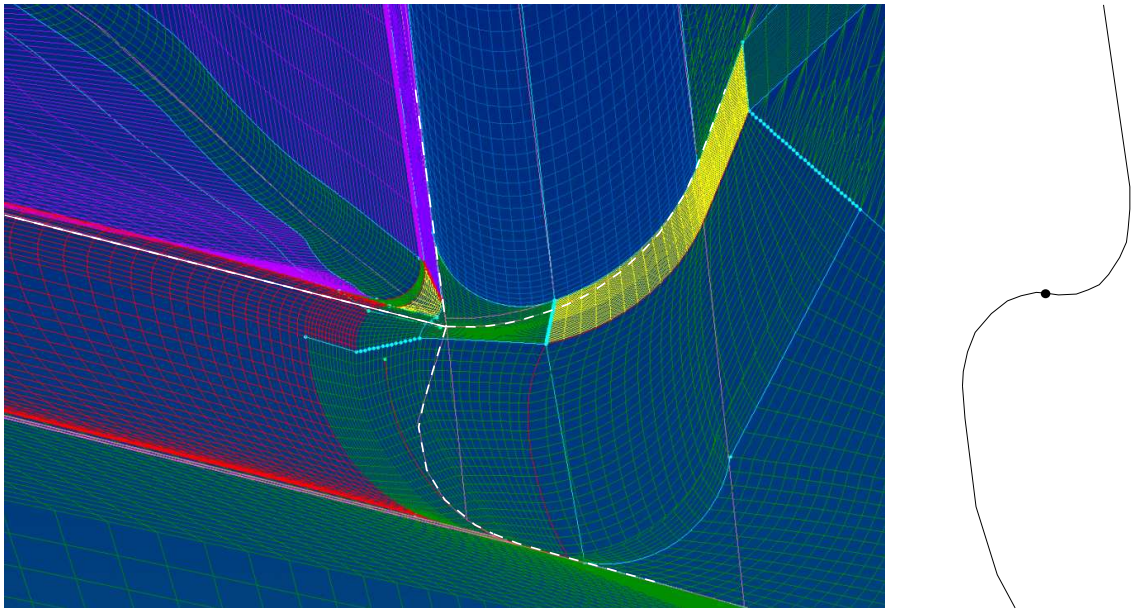

(a) Corner Surface **(b)** Inflection

Figure A.3: Wiremesh rounded corner surface (original sharp feature in white dashes) and schematic of inflection point

In the case of a backward facing step, such as that found in the M12 REST engine, a radius of 50μm was found to allow for a well-structured mesh at the step, without unnecessary extra cells or an irregular block topology.

A.2 Inlet Meshing

Having discussed at length the considerations for surface topology that ease the grid-generation process, it must be remarked that actual topology creation requires some significant thought as well. GridPro's solving algorithms will always attempt to achieve a grid that has smooth transitions in cell aspect ratio, as well as a high degree of orthogonality to any surfaces, be they a physical surface, a symmetry plane, or a useful virtual surface for block construction and anchoring that will have no effect on simulations of the mesh.

Of course, the usual best-practices in grid generation still apply: little cell density is required in regions where the flow is uniform and surfaces have little curvature. Similarly, in complex geometrical regions, or regions where the flow contains shocks or other discontinuities, the cell density must be high. But how should block topology be constructed when the blocks used must be able to provide a high quality mesh both in a sharp 90° corner, which then blends into a smooth elliptical surface downstream.

While the blocks may stretch and bend, the overall topology will be maintained in the streamwise direction. The solution used in the grids generated for this thesis is the "fan" of blocks illustrated in Figure A.4.

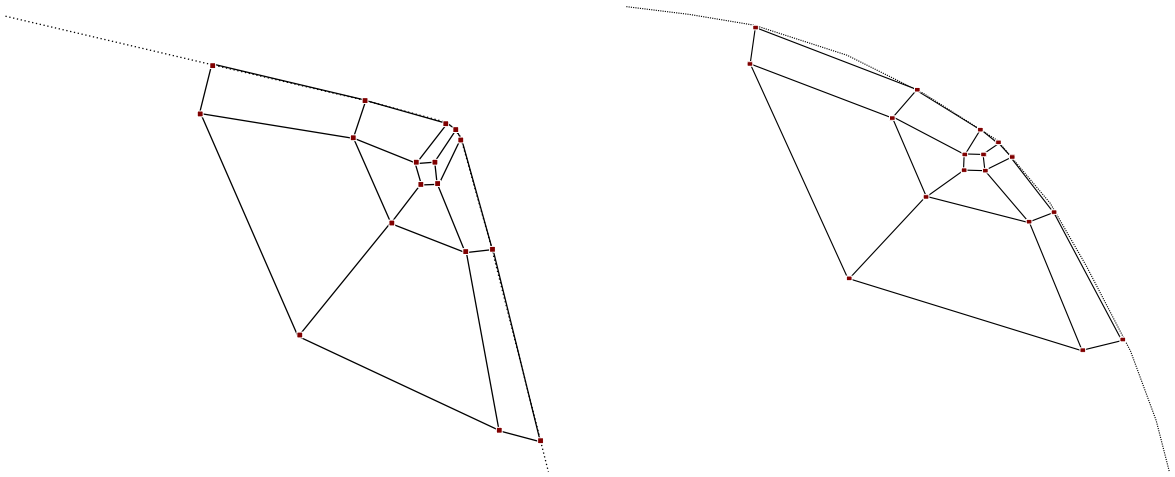

(a) Sharp Corner **(b)** Smooth Curve

Figure A.4: "Fan" block topology for sharp corner and smooth curvature surfaces (dashed line)

As demonstrated in Figure A.4a, in regions where the surface normal direction changes rapidly, such as the internal corners at the beginning of the inlet, the fan block structure allows the grid to curve cleanly around the geometry while maintaining good block orthogonality to the surface. Meanwhile, further downstream the blocks are free to spread to accommodate the uniformly curved surface without any tortuous warping or turning of the blocks, as Figure A.4b shows. A tube surface of the same radius as the inner quarter-circle 'edge' of the fan, running the length of the shape-transitioning portion of the inlet, is of some utility in constructing the fan geometry, but is unlikely to improve final mesh quality. On the contrary, block corners associated with the tube surface are more likely to cause a topological singularity that diminishes cell quality in this region of the mesh. The second "outer" quarter-circle of blocks allows easy integration of such fan structures into both the top and bottom corners of the inlet block topology, leading to a typical inlet cross-plane block structure like that shown in Figure A.5.

The fans integrate cleanly into the rest of an otherwise simple block topology with a minor 5-block singularity, allowing a great degree of copy and paste block generation along the streamwise length of the inlet surface between the side-cowl corner and the cowl notch. There are, of course, several special cases to be considered, each of which will be discussed in turn.

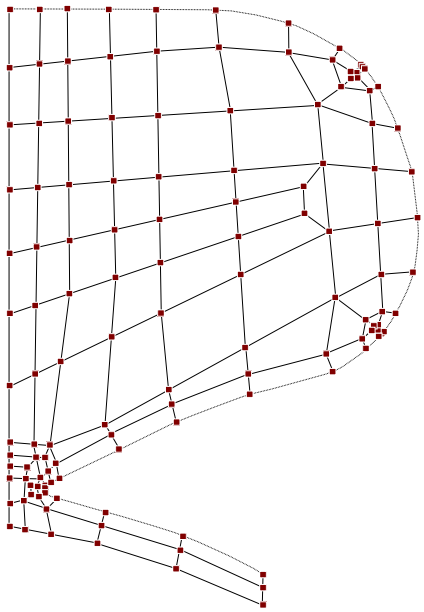

Figure A.5: Typical cross-plane block structure. Dashed line denotes inlet surface in-plane

A.3 Block creation and destruction

A high-fidelity mesh will necessarily mean a high cell count. However, this should not be thought to imply a uniform mesh density. Rather, regions of the flow that lack any sharp gradients or flow features (e.g. freestream flow along the inlet forebody) should have a minimal amount of grid devoted to them, while complex flow regions (e.g. injectors) or regions with strong gradients (e.g. the cowl closure notch) should have a higher than average cell density. The challenge becomes to smoothly transition between regions of differing grid density.

One simple option is to apply a "patch" region of cells. Some flow solvers have routines to specifically translate between two different grids that share a common block face. Others may allow for a hybrid mesh in which regions of structured hexahedral cells are connected via a block of unstructured tetrahedral cells, as is the case with US3D. In the latter case, it is important that any unstructured zones be located sufficiently far from any wall surfaces to avoid the poor solution quality that arises in most flow solvers when attempting to resolve flow to the wall using an unstructured mesh.

Another solution, and that which is recommended for this application, is to use some simple techniques that allow for the creation and destruction of extra blocks around regions where extra grid resolution is required. The first, and simplest, of these techniques to create new blocks is illustrated in Figure A.6. An additional connector is added along a boundary surface, and the resulting block terminates at a 45° angle to

the same surface. This allows these additional cells to be propagated in the streamwise direction without disrupting the propagation of cells already moving parallel to the surface; the block of cells closest to this surface is simply compressed into a smaller cross-section. This method is of particular use when creating complex block topologies like the corner 'fan' structure used for transitioning from the inlet corners into the smooth isolator curvature.

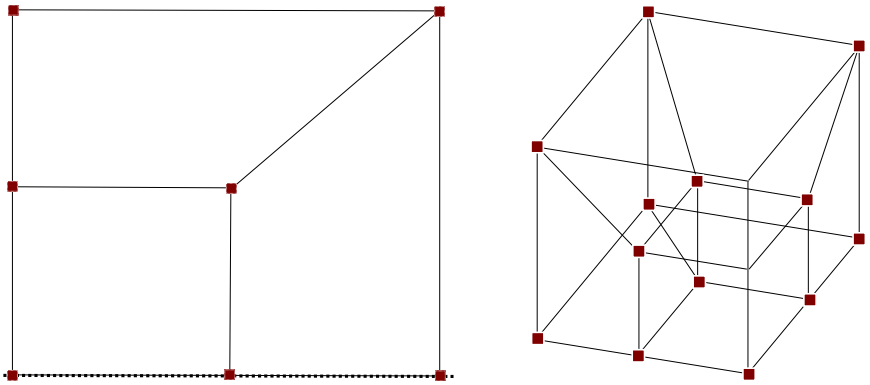

(a) 2D **(b)** 3D

Figure A.6: Boundary Block Creation in Two and Three Dimensions

Such a block creation scheme need not even propagate across the entire domain. The same technique can be used in three dimensions to avoid unnecessary cell refinement in regions where it is unnecessary, as shown in Figure A.6b. Block destruction at a boundary can be accomplished in the same manner, simply by reversing the procedure: blocks can be turned through angled boundaries such that they'll terminate at the wall.

In the case where grid refinement is necessary in regions where pulling cells off a boundary is impractical, and/or the level of refinement required necessitates the creation of multiple blocks, a trapezoidal block structure can be employed. As illustrated in Figure A.7, the narrowing of a single block on one end allows for the creation of two new rows of identical-cell-count blocks, effectively tripling (or even greater) the number of cells.

As a note of caution, this technique cannot be used immediately adjacent to any boundary surfaces, as this will result in a topological singularity that will degrade the local mesh quality significantly. As with the boundary block method, this refinement method can be easily extended to a three-dimensional topology, and can be reversed to provide rapid grid coarsening. This method can also be used multiple times in parallel, provided each instance of the trapezoidal block geometry is separated by a block layer

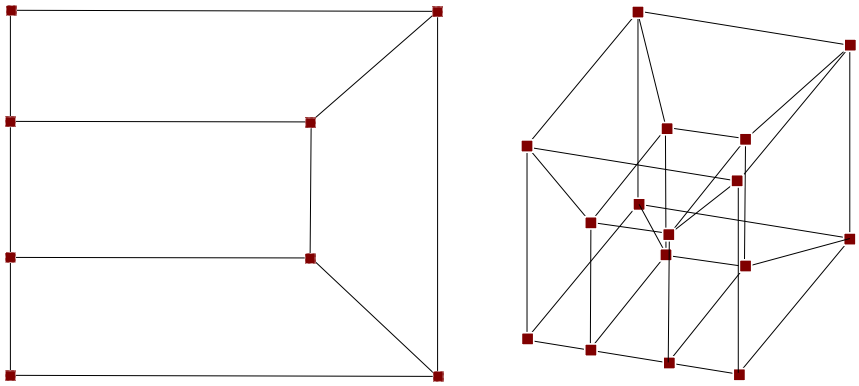

(a) 2D **(b)** 3D

Figure A.7: Trapezoidal Block Creation in Two and Three Dimensions

in order to avoid the creation of a six-connector singularity, as recommended by the GridPro manual (PDC, 1999).

The technique also need not use only a single block in between the two newly generated block layers; any number of blocks can be contained by this technique, provided the edge blocks narrow accordingly. A good example of this is the destruction of the "fan" block structure introduced in Section A.2. After destroying the diamond block at the center of the fan structure with the boundary block technique, the remainder of the fan can be destroyed quickly using a nested trapezoidal block structure. Destroying this fan structure greatly simplifies the grid topology downstream of the point where the shape transition is completed, and removes unnecessary complexity from the mesh entering the combustor section of the engine.

A.4 Cowl Notch Block Structure

The closure of the cowl notch is a specific application of the boundary block destruction method, but one that is complicated by the presence of the blocks required to adequately wrap the mesh around the leading edge. The keystone of this blocking structure is a rectangular block located one layer of blocks away from the inlet surface, and bounded by the engine symmetry plane. As demonstrated in Figure A.8, the corners of this block can be used as an anchor to wrap the leading-edge blocks into the symmetry plane surface.

The effect of using this topology is that the centerplane blocks are continuous through the entire length of the mesh topology. The three point singularity along the inner corner of this block also allows for the mesh to remain a high level of orthogonality at the wall.

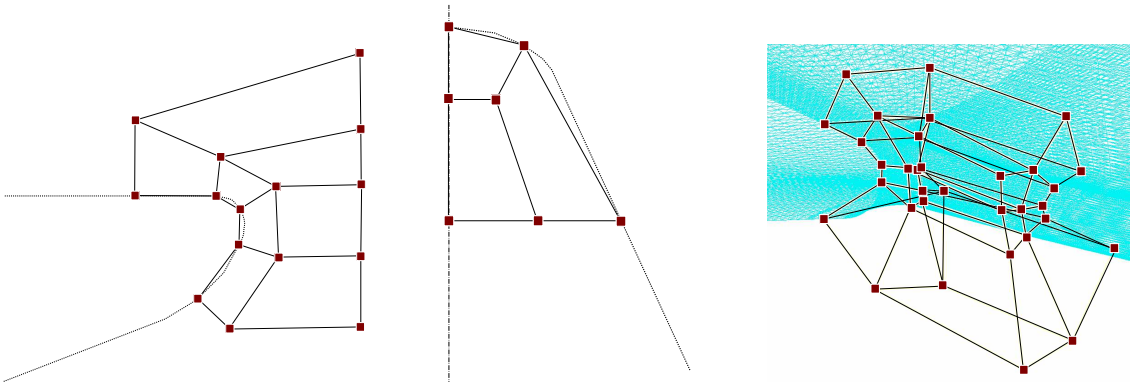

(a) Side View **(b)** Top View **(c)** Iso View

Figure A.8: Top, Side, and Isometric views of Cowl Closure Notch Block Geometry

The main deficiency of the topology as constructed is that this particular point of the inlet is a stagnation region. As such, the high temperature and density of the flow in these blocks makes achieving a wall-resolved mesh (i.e. $y^+ < 1$) problematic. Because wall clustering requires a uniform number of cells to be created in adjacent blocks, achieving a y^+ of less than one in this region would lead to an untenable amount of cells close to the wall.

However, in most simulation applications, the practical effects of under-resolving this region are likely small enough to ignore: as mentioned in Chapter 3, the regions in which the wall-spacing is not sufficient are physically small, and should have a negligible effect on the overall domain. An LES simulation would require this region to be fully resolved, in which case block creation would be necessary in order to achieve the nanometer-scale cell heights in this zone, and could be coupled with the '-Ov' flag in the grid clustering tool in order to locally increase wall resolution in the vicinity of the cowl notch.

A.5 Side-Cowl Corner Mesh

No other point in the REST topology is as geometrically challenging, and requires as much effort to mesh, as the Cowl-Leading-Edge/Side-Wall corner. In many ways, this region of the mesh encourages the most artistic license, as there are any number of combinations of block topology, creation, and destruction that can generate a viable mesh. Since there is no one "right" way to geometrically construct this corner with blunt leading edges, the mesh builder will have to adjust the advice given in this section to suit their own needs. The surface block topology used for the present study is shown in Figure A.9.

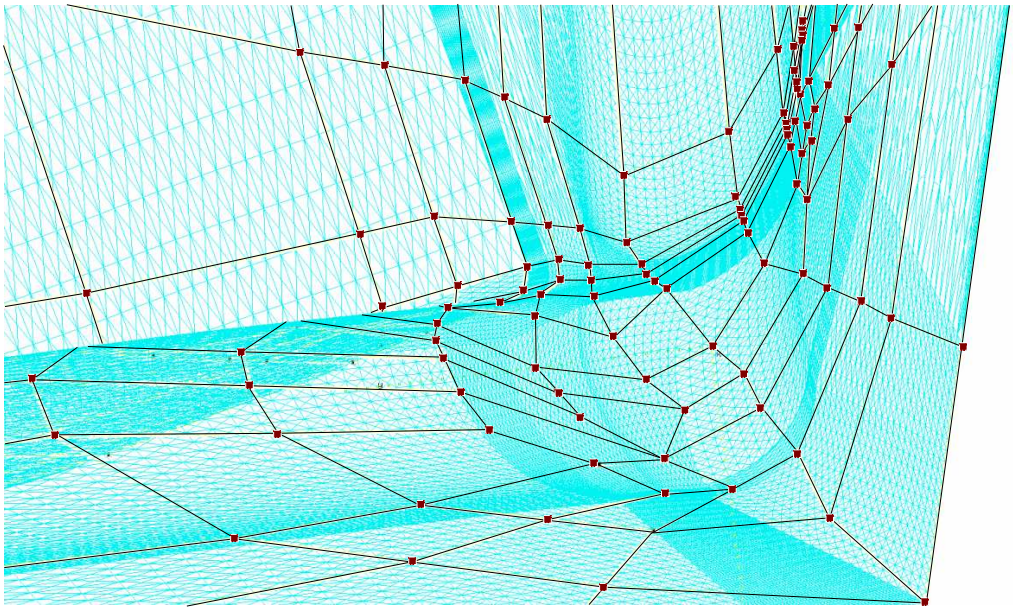

Figure A.9: Typical surface block structure for the Side-Cowl Corner

As a rule, it is recommended that blocks travelling along the cowl surfaces be wrapped around to the external surfaces of the inlet whenever possible. This allows cells propagating along the sidewall leading edge to move downward along the external surface, and cells propagating in the streamwise direction to enter the engine capture area in the smoothest way possible. Where necessary, block destruction should be used, though it is best to do this in regions of low curvature whenever possible. This block destruction will avoid unnecessary over-resolution of the flow around this corner, which is simpler than the block topology may suggest. In all cases, the builder is encouraged to remember that GridPro always seeks to build meshes that are orthogonal to local wall surfaces, and to attempt to construct blocks accordingly.

A.6 Injector Meshing

There are two competing philosophies for constructing the block topology of porthole fuel injectors. The first design involves constructing the blocks along the axis of the injector, and turning them downstream after the blocks have been extruded the approximate penetration height that the jets are expected to achieve. The other design is to create a uniform "block of blocks" which successively refines inward toward the injector. Immediately at the junction of the injector tube and the surface, the block structures would look roughly identical, while their large-scale structures differ significantly, as is shown in Figure A.10.

The advantage to the extruded block method is that the cell axes will roughly align

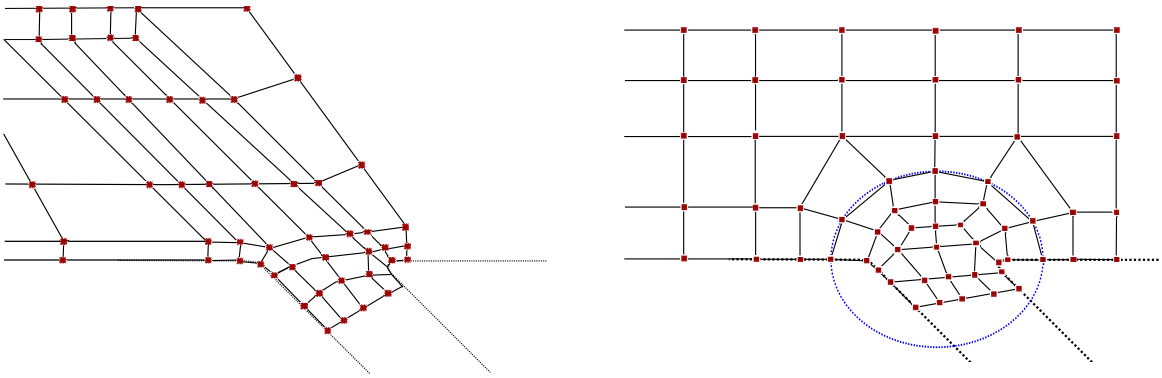

(a) Extruded Injector **(b)** "Block of Blocks"

Figure A.10: Comparison of extruded and refined "Block of Blocks" topologies

with the dominant flow features in the near-field of the jet. This allows for high resolution of the jet shock and vortical structures with the bare minimum number of blocks and cells. This comes at a significant cost in geometric complexity. While a single jet can make use of this construction method effectively, multiple injectors in close proximity to one another will make integration into the inlet mesh difficult at best. In cases where injectors may be located radially along a curved surfaces, as is the case in some combustor meshes, the extruded block method will result in regions of low cell quality and/or extraneous high-cell-density regions in otherwise uninteresting regions of the flow field.

In contrast, the "block of blocks" method involves the construction of a regular rectangular block structure that is easily integrable with the surrounding block topology, and with other nearby injectors. While this typically requires a larger number of blocks and a higher number of cells to achieve good resolution of the jet, the payoff in time savings and mesh quality are worth the slight increase in computational cost. For this reason, the remainder of this section will focus on the construction of these blocks. In general, all images will be of one half of the injector to facilitate visualization of the block structure. Because the injectors have a plane of symmetry, it is a simple matter to mirror such a topology to create a full injector block structure.

For convenience's sake, all examples will be shown for a 45° injection angle, as this strikes a compromise between the uniform mesh of perpendicular injection, and the highly asymmetric block topology of a shallower angle. It should only be a minor exercise for the reader to apply the principles illustrated to any injection angle.

The block structure inside and in the immediate vicinity (roughly 0.5D in all directions from the jet) is the most difficult piece of the injector topology to generate a high

quality mesh for. The key to this structure is to understand that the boundary layer cells from the surrounding wall surfaces ought to wrap around and into the injector tube. This can be accomplished in a controlled fashion for any injector angle using four blocks, as demonstrated in Figure A.11.

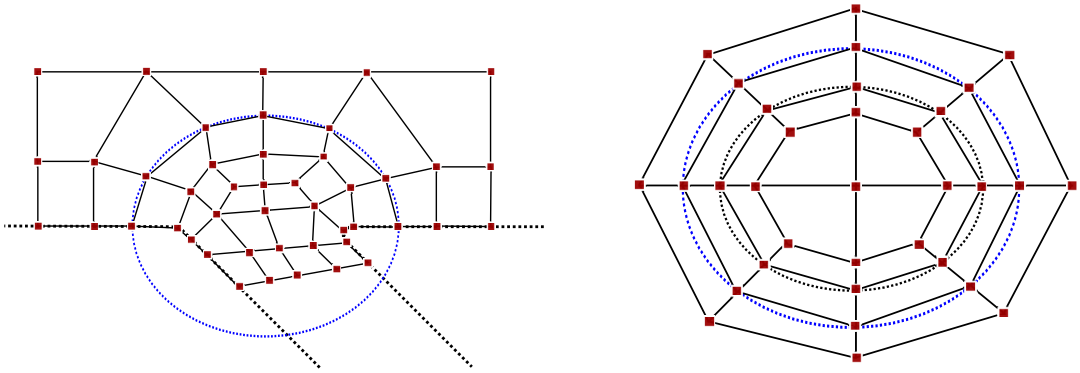

(a) Side View **(b)** Top View (at surface)

Figure A.11: Recommended injector block topology

These blocks allow the wall-bounded grid turn into the injector tube with a minimum in non-orthogonality in the cells at the wall surface. The use of these structures leads to the shown singularities, though their presence does not cause any undue grid distortion, in large part thanks to the elliptical surface built to prevent the grid from moving or stretching near the injector orifice. The top view of the injector topology in Figure A.11b shows how this profile can be rotated around the axis of the injector orifice: by utilizing a quartet of square core blocks in the injector, the blocks that wrap around the injector lip can use the core block edges as anchors to sweep around the circumference of the orifice.

Outside the ellipsoid surface that contains the injector blocks in the regions of high curvature, a layer of blocks is required to avoid a singularity being attached to the surface, which would result in poor grid quality. As is readily visible in Figure A.11a above, an additional layer of blocks laterally along the wall surface allows for the construction of a rectangular subset of blocks, hereafter referred to as the inner injector block group.

Because the inner injector block group is rectangular in shape, it becomes a relatively trivial exercise in block creation and destruction in order to integrate the injector mesh into the surrounding inlet mesh. The methodology for doing so is demonstrated in Figure A.12. Starting from the inlet block, it is first prudent to determine the expected approximate jet penetration height. This will allow the overall inlet grid

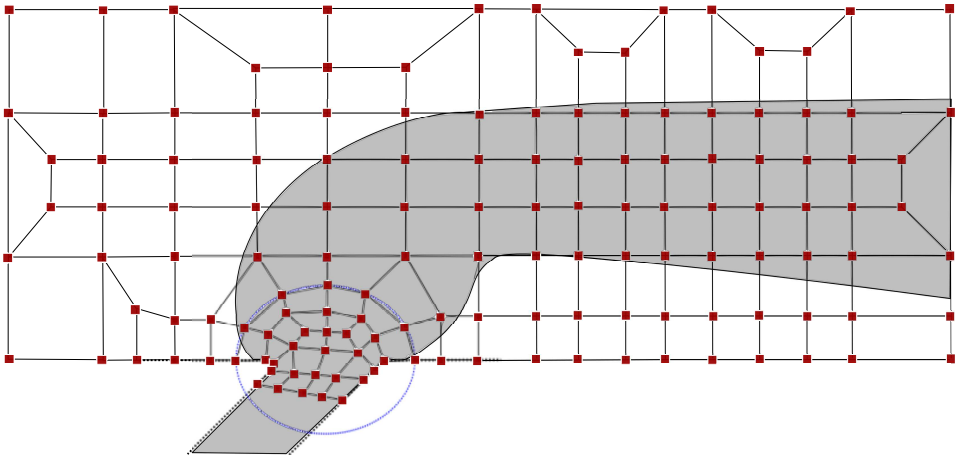

Figure A.12: Block creation and destruction to integrate injector mesh into surrounding topology. Shaded region shows expected injector plume.

structure to retain adequate resolution into the mainstream flow region in order to accurately capture jet development and mixing. Upstream and to the sides of the injector orifice it is prudent to retain fine mesh resolution for approximately 1D of length, particularly near the wall, in order to accurately capture the upstream recirculation zone and the development of the injector bow shock. Downstream of the injector orifice, the mesh should retain a reasonably high density for a length of at least 10D. Longer lengths may be desirable if doing so does not push the overall mesh size too large to be computed by available resources, but it has been shown that coherent jet structures tend to break down 8-10D downstream of injection, so lower grid resolution is not likely to cause a loss in physical flow information downstream of this point.

As mentioned previously, this "block of blocks" topology is (relatively) easy to integrate into complex geometries, and especially in the presence of other nearby injectors. An example of this is the tailored injection case shown in Figure A.13.

The simple shape of the injector blocks makes integrating injector topologies at right angles to one another on a curved surface a rather painless procedure. The injector topology shown quickly blends with the coarser block structure upstream, allowing for rapid generation of the overall grid topology. The same injector topology was attempted with the extruded block method, but quickly abandoned given the complexity and low quality of the blocks created.

A.7 Grid Generation and Wall Clustering

Up until this point, there has been no discussion of how to choose grid spacing across the inlet, nor how the grid should be refined as it approaches the wall. This is, to some

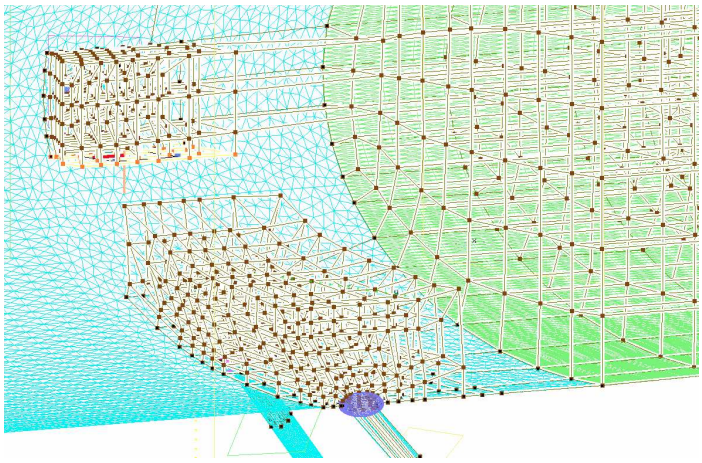

Figure A.13: Example of integrating multiple injector blocks.

extent, a matter of preference and of experience, and the grid maker is encouraged to experiment with grid spacing to meet their own requirements. However, some general recommendations may be made that have been found through experience to give good resolution of flow features at a reasonable cost.

For general mainstream grid spacing within the bounds of the inlet walls, choosing a block cell count that allows for 0.25 mm grid spacing will provide a reasonably high degree of grid fidelity, without overly inflating the global cell count. This cell size can obviously be grown considerably in the mainstream flow regions beneath the inlet forebody, where no complicated flow features are expected. Such grid spacing was used in the inlet of the M12 REST engine, and the total inlet cell count (including the 500 mm forebody) was approximately 20 Million cells after grid refinement at the inlet wall. In all cases, the mesh maker is encouraged to continually be asking themselves what they will need to do up- or downstream of the topology they are currently creating; since block cell counts propagate along shared edges, a requirement to refine cell count in one small region of the flow could lead to a large growth in global cell counts.

Taking steps to ensure the blocks are constructed with due care, and with reasonable block density along leading edges, running the GridPro *Ggrid* command to smooth the grid should not be a time consuming procedure. As a point of reference, the block edges tangent to the leading edge surface for the blocks shown in Figure A.8a were usually given a dimension of 8 to 12 cells. In this example, the grid was typically smoothed to a point where block folds were eliminated within 150-200 sweeps of the smoothing algorithm. For a grid constructed with less care, additional construction surfaces may be required in order to constrain sufficient numbers of cells to the inlet leading edges.

For grid refinement approaching the inlet walls, it is recommended that the GridPro

clu command be run, using a first-cell height of 0.9 microns or less and a cell growth factor of 1.05. This was found to provide a y^+ value less than unity across nearly the entire inlet flow domain, with the exception of the previously mentioned small region near the cowl closure notch, where grid requirements would be untenable for nearly any simulation given current computational resources.

A.8 Concluding Remarks

The preceding documentation completely describes the procedure required to reproduce the mesh topology used in the study of the M12 REST engine inlet. While no discussion has been made regarding the engine combustor, the comparatively trivial mesh topology required downstream of the inlet should be well within the capabilities of the grid maker once they have created a mesh of similar complexity to the one demonstrated here.

Appendix B

The Jachimowski (1992) Hydrogen-Air Reaction Model

This appendix contains the coefficients for the forward reaction rates of the hydrogen-air reaction mechanism detailed in Jachimowski (1992). The forward reaction rate is expressed using an expanded form of the Arrhenius equation:

$$k_f = AT^B e^{-E/RT}, \tag{B.1}$$

where R is the universal gas constant, T is the bulk temperatures, A is a pre-exponential factor, B is a temperature exponent, and E is the activation energy of the reaction. These coefficients are listed in Table B.1. The values of these coefficients differ from those listed in the reference, as US3D requires all quantities to be in SI units.

In cases where the reaction requires the collision of a third body, M, a third body efficiency multiplier is required. All third body efficiencies (which act as a multiplier to the calculated value of k_f) are assumed to be equal to one, with the following exceptions: H_2 has a third body efficiency of 2.5, and H_2O has a third body efficiency of 16.

Table B.1: Forward Rate Coefficients for the Jachimowski (1992) Model

Reaction	*A*	*B*	*E*
$H + OH + M \rightarrow H_2O + M$	8.620E+15	-2.00E+00	0.00000E+00
$H + H + M \rightarrow H_2 + M$	7.300E+11	-1.00E+00	0.00000E+00
$H + O + M \rightarrow OH + M$	2.600E+10	-0.60E+00	0.00000E+00
$O + O + M \rightarrow O_2 + M$	1.100E+11	-1.00E+00	0.00000E+00
$H + O_2 + M \rightarrow HO_2 + M$	2.300E+12	-1.00E+00	0.00000E+00
$OH + OH + M \rightarrow H_2O_2 + M$	1.600E+16	-2.00E+00	0.00000E+00
$N + N + M \rightarrow N_2 + M$	2.000E+11	-0.80E+00	0.00000E+00
$H + NO + M \rightarrow HNO + M$	5.400E+09	0.00E+00	-3.02133E+02
$NO_2 + M \rightarrow N + O + M$	1.160E+13	0.00E+00	3.32347E+04
$H_2 + O_2 \rightarrow HO_2 + H$	7.000E+10	0.00E+00	2.86020E+04
$H + O_2 \rightarrow OH + O$	2.200E+11	0.00E+00	8.45974E+03
$O + H_2 \rightarrow OH + H$	5.060E+01	2.67E+00	3.16737E+03
$OH + H_2 \rightarrow H_2O + H$	2.160E+05	1.51E+00	1.72720E+03
$OH + O_2 \rightarrow H_2O + O$	1.500E+06	1.14E+00	0.00000E+00
$HO_2 + H \rightarrow OH + OH$	1.500E+11	0.00E+00	5.03556E+02
$HO_2 + O \rightarrow O_2 + OH$	2.000E+10	0.00E+00	0.00000E+00
$HO_2 + OH \rightarrow H_2O + O_2$	2.000E+10	0.00E+00	0.00000E+00
$HO_2 + HO_2 \rightarrow H_2O_2 + O_2$	2.000E+09	0.00E+00	0.00000E+00
$H + H_2O_2 \rightarrow H_2 + HO_2$	1.700E+09	0.00E+00	1.90344E+03
$H + H_2O_2 \rightarrow OH + H_2O$	1.000E+10	0.00E+00	1.80273E+03
$O + H_2O_2 \rightarrow OH + HO_2$	2.800E+10	0.00E+00	3.22276E+03
$OH + H_2O_2 \rightarrow H_2O + HO_2$	7.000E+09	0.00E+00	7.22602E+02
$N + O_2 \rightarrow NO + O$	6.400E+06	1.00E+00	3.17240E+03
$N + NO \rightarrow N_2 + O$	1.600E+10	0.00E+00	0.00000E+00
$N + OH \rightarrow NO + H$	6.300E+08	0.50E+00	0.00000E+00
$H + HNO \rightarrow NO + H_2$	4.800E+09	0.00E+00	0.00000E+00
$O + HNO \rightarrow NO + OH$	5.000E+08	0.50E+00	0.00000E+00
$OH + HNO \rightarrow NO + H_2O$	3.600E+10	0.00E+00	0.00000E+00
$HO_2 + HNO \rightarrow NO + H_2O_2$	2.000E+09	0.00E+00	0.00000E+00
$HO_2 + NO \rightarrow NO_2 + OH$	3.400E+09	0.00E+00	-1.30924E+02
$HO_2 + NO \rightarrow HNO + O_2$	2.000E+08	0.00E+00	5.03556E+02
$H + NO_2 \rightarrow NO + OH$	3.500E+11	0.00E+00	7.55334E+02
$O + NO_2 \rightarrow NO + O_2$	1.000E+10	0.00E+00	3.02133E+02

Appendix C

Record of Experiments

The operating conditions, and nozzle supply conditions for each experimental shot of the validation experiments are listed in Table C.1, and nozzle exit conditions are listed in Table C.2. The fuel system pressures and mass flow rates for each shot are recorded in Table C.3. The variable $t_{Dia.}$ in Table C.1 is the primary diaphragm thickness, and the variable PP in Table C.2 is the pitot pressure.

In Table C.3, the pressures listed are for the inital Ludweig Tube pressure ($p_{LT,i}$), and the pressures of the inlet and combustor plena during test time.

Table C.1: Validation Experiment Shot Condition Summary

	Res.	Comp. Tube			Shock Tube					Nozzle supply cond.		
Shot No.	p_{fill} MPa	p_{fill} kPa	Fill Gas %Ar	%He	p_{fill} kPa	T_{fill} K	Test Gas	$t_{Dia.}$ mm	u_{shock} m/s	p_s MPa	T_s K	H_0 MJ/kg
11486	12.0	156	35	65	160	300	Air	6	2605	81.58	4980	6.951
11487	12.0	156	35	65	160	300	Air	6	2653	82.79	5082	7.178
11488	12.0	156	35	65	160	300	Air	6	2666	83.52	5113	7.241
11489	12.0	156	35	65	160	300	Air	6	2617	84.50	5032	7.062
11490	12.0	156	35	65	160	300	Air	6	2605	84.18	5006	7.008
11491	12.0	156	35	65	160	300	Air	6	2605	84.90	5014	7.010
11492	12.0	156	35	65	160	300	Air	6	2617	84.55	5032	7.063
11493	12.0	156	35	65	160	300	Air	6	2605	84.29	5008	7.010

Table C.2: Validation Experiment Nozzle Exit Condition Summary

	Nozzle exit conditions							
Shot No.	p_∞ Pa	T_∞ K	ρ_∞ g/m^3	u_∞ m/s	M_∞ -	$Re_{u,\infty}$ 10^6 1/m	$p\gamma_\infty$ -	PP kPa
11486	1131	382.5	10.22	3615	9.195	1.655	1.399	123.4
11487	1119	394.2	9.85	3674	9.228	1.599	1.397	123.4
11488	1134	397.0	9.92	3689	9.238	1.609	1.397	124.9
11489	1186	391.5	10.53	3645	9.194	1.704	1.397	129.8
11490	1172	385.8	10.56	3633	9.234	1.722	1.398	129.5
11491	1177	386.8	10.51	3630	9.183	1.695	1.398	118.6
11492	1047	377.9	9.63	3649	9.367	1.601	1.398	118.6
11493	1064	375.4	9.85	3637	9.368	1.640	1.398	120.8

Table C.3: Validation Experiment Fuelling Summary

Shot No.	$p_{LT,i}$ MPa	p_{inlet} kPa	p_{comb} kPa	$\dot{m}_{inlet}$ g/s	$\dot{m}_{comb}$ g/s	$\dot{m}_{total}$ g/s
11486	-	-	-	-	-	-
11487	2.100	295.1	1914.4	0.957	1.601	2.558
11488	2.105	240.0	1927.5	0.802	1.611	2.413
11489	2.105	228.5	1928.7	0.769	1.612	2.381
11490	2.496	254.4	2278.4	0.864	1.905	2.769
11491	3.332	289.0	2972.9	1.005	2.494	3.499
11492	2.890	289.2	2639.3	0.985	2.206	3.192
11493	2.801	0.000	2599.6	0.000	2.168	2.168

Appendix D

Calibrations and Uncertainty Analysis

This appendix details the measured and calculated uncertainties for data collected over the course of the joint experimental campaign. The procedure by which the fuel system was calibrated is also presented. In all cases, the error estimates are derived using the Root-Sum-Square (RSS) method, in which the total uncertainty in any quantity, Y, is calculated using the formula

$$\mathrm{Y}_{\psi} = \sqrt{\psi 1^2 + ... + \psi_i^2 + ... + \psi_N^2} \tag{D.1}$$

where ψ_i is the uncertainty in each fundamental measurement that was used to calculate the quantity (Y) in question. Many of the mathematical details and derivations employed in the uncertainty analysis have been omitted, and may be found by consulting the summary of Doherty (2013).

D.1 Nozzle Exit Flow Properties

The nozzle exit flow conditions for the average tunnel condition, and for each experiment simulated and validated in this study, are calculated using the space-marching nozzle flow code NENZFr, as mentioned in Chapter 3. This code requires a total of nine inputs to solve for the nozzle exit flow properties, four of which have been shown to have a negligible effect on the uncertainty of the values calculated (Doherty, 2013). The remaining five variables consist of four inputs from the experimental conditions recorded, and one purely numerical value. The experimental values consist of the shock tube fill pressure, P_{ST}; the shock tube fill temperature, T_{ST}; the measured shock speed, V_{shock}; and the nozzle supply pressure, P_{NS}. The fifth input is the estimated location at which the nozzle boundary layer transitions from laminar to turbulent, x_{tr}.

Table D.1: Relative uncertainties in variables used by NENZFr

Input Variable	Nominal Value	Relative Uncertainty
P_{ST}	160 kPa	1.5%
T_{ST}	300 K	2.0%
V_{shock}	2642.9 m/s	5.0%
P_{NS}	82.5 MPa	3.6%
x_{tr}	75 mm	100%

The uncertainties for each of these variables are summarized in Table D.1. The shock tube fill pressure uncertainty is based on the accuracy of the fill pressure gauge used. The gauge face has 2 kPa increments, and the gauge itself has an accuracy of 0.5%. It was assumed that a tunnel operator could accurately read the gauge in 1 kPa increments (half-way between two gauge markings), and that the ambient laboratory pressure is accurate to within 2% of the calibrated value. Based on these assumptions, the RSS relative uncertainty in P_{ST} is 1.5%. The temperature of the test gas in the shock tube is never measured during an experiment, but is allowed to settle for between 10 and 15 minutes from the time it is filled to the time the experiment is conducted. The standard assumption that is used in T4 is that the shock tube temperature is equal to a nominal laboratory temperature of 300 K, with an uncertainty of 5%. The published uncertainties in the measured shock speed and nozzle supply pressure are 5% (Mee, 1993) and 3.6% (Doherty, 2013), respectively.

The nominal boundary layer transition location is chosen to be 75 mm. This location was determined by matching the calculated pressure of the flow after it had passed through the forebody shock, to the pressure measured on the forebody of fuel-off experiment Shot 11486. While this pressure match did not hold across all experiments, it was not expected to. The boundary layer may transition nearly anywhere in the region immediately downstream of the throat during an experiment. As such, the uncertainty associated with this transition location is set at 100%.

The uncertainty in the exit flow properties provided by NENZFr is estimated by first running a perturbation study of these five variables on the nominal experimental tunnel conditions. In this procedure, a simulation is run to perturb each of the five variables off its nominal value an increment of $\pm 1.5\%$. The resulting response surface is then used to perform a sensitivity analysis, resulting in the uncertainties for each freestream flow variables summarized in Table D.2. For more details on this sensitivity study, refer to Doherty et al. (2012b).

Table D.2: Relative uncertainties of NENZFr-calculated nozzle exit core flow properties

Variable (Input uncertainty)	P_{ST} 1.5%	T_{ST} 2.0%	V_{shock} 5.0%	P_{NS} 3.6%	x_{tr} 100%	Total RSS Uncertainty
ρ_{NS}	0.25%	-0.41%	-5.81%	2.95%	0.00%	6.53%
T_{NS}	-0.20%	0.34%	4.83%	0.60%	0.00%	4.88%
H_{NS}	-0.31%	0.52%	7.43%	0.76%	0.00%	7.50%
ρ_{∞}	0.30%	-0.51%	-7.21%	2.62%	4.63%	8.98%
U_{∞}	-0.14%	0.24%	3.43%	0.38%	0.00%	3.46%
P_{∞}	-0.12%	0.18%	2.77%	3.68%	6.08%	7.62%
a	-0.21%	0.34%	4.93%	0.52%	0.83%	5.04%
μ	-0.31%	0.50%	7.31%	0.77%	1.17%	7.47%
μ_t	-0.17%	0.29%	4.16%	-0.71%	58.10%	58.26%
M_{local}	0.06%	-0.10%	-1.50%	-0.14%	-0.84%	1.73%
P_{pitot}	0.01%	-0.02%	-0.31%	3.37%	4.55%	5.67%
P_0	0.23%	-0.35%	-5.41%	2.98%	0.94%	6.26%
H_0	-0.31%	0.52%	7.45%	0.76%	0.00%	7.52%
q	0.02%	-0.03%	-0.35%	3.37%	4.63%	5.74%
Re_u	0.47%	-0.77%	-11.10%	2.22%	3.46%	11.87%
Y_{N_2}	0.00%	-0.00%	-0.03%	0.00%	0.00%	0.03%
Y_{O_2}	0.05%	-0.09%	-1.24%	0.00%	0.00%	1.24%
Y_O	-1.53%	2.55%	36.32%	-0.21%	0.00%	36.44%
Y_{NO}	-0.02%	0.03%	0.49%	0.01%	0.00%	0.49%

D.2 Thin-Film Heat Transfer Gauges

The thin-film heat transfer gauges (HT gauges) used to measure surface heat flux in the validation experiments were manufactured in The University of Queensland's Instrumentation Laboratory. The manufacturing method used, the gauges' calibration, and the resulting sources of uncertainties their measurements are described in detail in Wise (2014). These uncertainty estimates are summarized briefly below.

The first source of uncertainty in the measured heat flux is due to the thermal inertia of the nickel film coating the surface of the gauge. This source of uncertainty has been demonstrated to be negligible, if there is at least 2 ms of model startup flow prior to the start of the experimental test envelope (Wise, 2014). A thin Silicon dioxide film has been applied to the gauges to improve their durability, and prevent the possibility of a short circuit caused by a gauge being exposed to ionized flow. This coating contributes an uncertainty of 5% to the measured heat flux.

The substrate of each heat transfer gauge is made from quartz, which is a relatively poor thermal conductor. As such, the heat transfer gauges mounted in the engine will

see their surfaces warm significantly during the start-up flow time. This temperature increase has been estimated to be as high as 100 K, with some gauge surface temperatures inside the M12 REST combustor having been recorded as increasing by more than 200 K on occasion. Taking the 100 K value as a nominal value, there is a 10.5% measurement uncertainty contribution due to this effect (Goyne, 1999). Calibration errors and voltage measurement errors have been shown to be 2.5% and 0.5%, respectively (Wise, 2014).

Using these individual contributions to uncertainty, summarized in Table D.3, an overall RSS uncertainty was estimated. This resulted in a total uncertainty of 11.9%, assumed to be constant across all of the heat transfer gauges.

Table D.3: Uncertainty in Heat Flux Measurements using HT Gauges

Source	Uncertainty
Nickel Film	N/A
SiO_2 coating	5%
Quartz Substrate	10.5%
Calibration	2.5%
Voltage	0.5%
Total	11.9%

As was mentioned in Chapter 3, there were a total of 14 heat transfer gauges mounted on the bodyside compression surface of the inlet, and 28 mounted in the combustor. The streamwise locations of the inlet heat transfer gauges are listed in Table D.4. The combustor heat transfer gauges were located along the bodyside and cowlside of the combustor at the same streamwise locations as the combustor pressure transducers. These are listed in tables in the next section, and are not repeated here for the sake of brevity.

Table D.4: Streamwise location of HT Gauges

Gauge ID	Streamwise Location [mm]	Gauge ID	Streamwise Location	Gauge ID	Streamwise Location
IBHT1	530.0	IBHT6	658.6	IBHT11	791.6
IBHT2	560.0	IBHT7	673.5	IBHT12	806.5
IBHT3	579.0	IBHT8	688.5	IBHT13	836.3
IBHT4	598.9	IBHT9	703.4	IBHT14	851.2
IBHT5	618.8	IBHT10	776.8		

D.3 Experimental Pressure Measurements

The experimental pressures measured in this study are acquired using Kulite pressure transducers. These transducers have a 1% experimental uncertainty in their measured output voltage. This uncertainty covers both the mounting of the transducers, and any possible thermal effects (Doherty, 2013). In order to obtain a pressure from each gauge's voltage signal, the output is divided by the sensitivity of each sensor. The resulting sensitivity and calibration uncertainty of each pressure transducer is shown in Tables D.6 and D.5.

Table D.5: Engine Cowlside Pressure Transducer Details

Sensor ID	Serial Number	Range [psi/kPa]	Sensitivity [V/kPa]	Calibration Uncertainty [%]	Streamwise Position [mm]
CC1	Y97-6	100/689.48	$1.4675x10^{-4}$	$\pm$0.13	1038.5
CC2	Y97-1	100/689.48	N/A		1059.3
CC3	Y97-16	100/689.48	$1.4719x10^{-4}$	$\pm$0.19	1080.2
CC4	Y97-3	100/689.48	$1.4720x10^{-4}$	$\pm$0.14	1101.0
CC5	Y97-12	100/689.48	$1.4766x10^{-4}$	$\pm$0.17	1121.9
CC6	S94-12	50/344.74	$2.9494x10^{-4}$	$\pm$0.12	1142.8
CC7	X97-64	50/344.74	$2.9460x10^{-4}$	$\pm$0.06	1163.7
CC8	X97-68	50/344.74	$2.9517x10^{-4}$	$\pm$0.08	1184.6
CC9	X97-59	50/344.74	$2.9575x10^{-4}$	$\pm$0.08	1200.1
CC10	X97-67	50/344.74	$2.9464x10^{-4}$	$\pm$0.04	1221.0
CC11	S94-13	50/344.74	$2.9584x10^{-4}$	$\pm$0.05	1241.9
CC12	X97-62	50/344.74	$2.9238x10^{-4}$	$\pm$0.16	1262.9
CC13	X97-63	50/344.74	$2.9573x10^{-4}$	$\pm$0.17	1283.8
CC14	Y92-30	50/344.74	$2.9317x10^{-4}$	$\pm$0.08	1304.8
NC1	X97-65	25/172.37	$2.9355x10^{-4}$	$\pm$1.11	1322.4
NC2	UU8-69	25/172.37	$5.8252x10^{-4}$	$\pm$0.23	1334.4
NC3	UU8-70	25/172.37	$5.9074x10^{-4}$	$\pm$0.30	1349.9
NC4	UU8-67	25/172.37	$5.8777x10^{-4}$	$\pm$0.14	1366.4
NC5	UU8-68	25/172.37	$5.7831x10^{-4}$	$\pm$0.18	1382.8
NC6	UU8-71	25/172.37	$5.0742x10^{-4}$	$\pm$1.11	1399.3
NC7	P97-79	25/172.37	$5.8249x10^{-4}$	$\pm$0.06	1415.7

As is plain from the tables presented, the largest calibration uncertainty for these transducers' sensitivity coefficient is 2.0% (Wise, 2014). This leads to an RSS uncertainty of 2.2% for the pressure measurements. In the cases where a normalized pressure is plotted, the pressures are normalized by the nominal forebody pressure, and the ratio of the measured and nominal nozzle supply pressures (see Appendix F for the exact form of this relation). The nominal forebody pressure is calculated from the nozzle exit static pressure using isentropic flow relations:

$$\frac{P_{1,nom}}{P_\infty} = 1 + \frac{2\gamma}{\gamma - 1}(M_\infty^2 \sin^2\beta - 1), \tag{D.2}$$

where the shock angle, β, is calculated from the upstream Mach number (M_∞) and the deflection angle (θ) using the θ-β-M relation:

$$\tan\theta = 2\cot\beta \frac{M_\infty^2 \sin^2\beta - 1}{M_\infty^2(\gamma + \cos 2\beta) + 2}. \tag{D.3}$$

Knowing that the measurement uncertainty in the 1.6° forebody angle of attack is 0.1°, the RSS uncertainty in β is 6.49%, giving nominal forebody pressure uncertainty of 10.15%. The normalized experimental pressure uncertainty can then be calculated knowing the uncertainties in measured and nominal P_0, and in measured static pressure. The RSS uncertainty of P/P_1 is thus found to be 11.57%.

D.4 Fuel System

Unlike previous sections, where calibrations were completed by Wise (2014), the fuel system must be calibrated specifically for the validation experiments used in this study. As such, before describing the uncertainty in the measured fuel mass flow rates, it is first necessary to describe how the fuel system is calibrated.

D.4.1 Calibration Procedure

The technique of Pulsonetti (1995) is used to calibrate the mass flow rate of fuel through each injector during any experiment's test time. It is assumed that the fuel stored in the Ludweig tube prior to the test is allowed to expand isentropically, that flow through the tubing to each injector plenum is adiabatic, and that the flow through the injector orifices is choked. This allows the mass flow rate of hydrogen fuel through the injectors to become a function of the stagnation pressure in the plenum chamber Turner and

Table D.6: Engine Bodyside Pressure Transducer Details

Sensor ID	Serial Number	Range [psi/kPa]	Sensitivity [V/kPa]	Calibration Uncertainty [%]	Streamwise Position [mm]
FB1	P97-74	25/172.37	$5.9266\text{x}10^{-4}$	$\pm$0.05	162.5
FB2	P97-73	25/172.37	$5.8860\text{x}10^{-4}$	$\pm$0.06	162.5
IB1	UU8-75	25/172.37	$5.9052\text{x}10^{-4}$	$\pm$2.00	545.0
IB2	UU8-72	25/172.37	$5.8287\text{x}10^{-4}$	$\pm$0.19	638.7
IB3	P97-71	25/172.37	$5.8610\text{x}10^{-4}$	$\pm$0.08	718.3
IB4	X97-97	100/689.48	$1.4671\text{x}10^{-4}$	$\pm$0.09	821.3
IB5	X97-98	100/689.48	$1.4579\text{x}10^{-4}$	$\pm$0.27	871.0
IB6	Y97-7	100/689.48	$1.2132\text{x}10^{-4}$	$\pm$0.27	890.8
IB7	Y97-2	100/689.48	$1.4702\text{x}10^{-4}$	$\pm$0.19	910.6
IB8	Y97-9	100/689.48	$1.4630\text{x}10^{-4}$	$\pm$0.19	930.4
IB9	Y97-4	100/689.48	$1.4656\text{x}10^{-4}$	$\pm$0.25	955.3
IB10	Y97-99	100/689.48	$1.4639\text{x}10^{-4}$	$\pm$0.18	975.3
IB11	Y97-13	100/689.48	$1.4678\text{x}10^{-4}$	$\pm$0.18	995.3
CB1	CC9-30	50/344.74	$2.9251\text{x}10^{-4}$	$\pm$0.73	1042.4
CB2	CC9-25	50/344.74	$3.0452\text{x}10^{-4}$	$\pm$0.25	1062.3
CB3	CC9-26	50/344.74	$2.9263\text{x}10^{-4}$	$\pm$0.46	1082.3
CB4	CC9-31	50/344.74	$2.9300\text{x}10^{-4}$	$\pm$0.14	1102.4
CB5	CC9-29	50/344.74	$2.9242\text{x}10^{-4}$	$\pm$0.20	1122.4
CB6	CC9-32	50/344.74	$2.9326\text{x}10^{-4}$	$\pm$0.30	1142.5
CB7	CC9-28	50/344.74	$2.9400\text{x}10^{-4}$	$\pm$0.28	1162.5
CB8	CC9-27	50/344.74	$3.1685\text{x}10^{-4}$	$\pm$0.21	1182.5
CB9	CC9-24	50/344.74	$3.1866\text{x}10^{-4}$	$\pm$0.21	1198.0
CB10	X97-55	50/344.74	$2.9440\text{x}10^{-4}$	$\pm$0.07	1218.8
CB11	X97-61	50/344.74	$2.9271\text{x}10^{-4}$	$\pm$0.05	1239.6
CB12	X97-54	50/344.74	$2.9402\text{x}10^{-4}$	$\pm$0.12	1260.4
CB13	X97-69	50/344.74	$2.9370\text{x}10^{-4}$	$\pm$0.09	1281.2
CB14	X97-58	50/344.74	$2.9586\text{x}10^{-4}$	$\pm$0.03	1302.0
NB1	X97-60	50/344.74	N/A		1319.3
NB2	X97-53	50/344.74	$2.9422\text{x}10^{-4}$	$\pm$0.11	1330.8
NB3	P97-80	25/172.37	$5.8551\text{x}10^{-4}$	$\pm$0.07	1345.8
NB4	P97-75	25/172.37	$5.8720\text{x}10^{-4}$	$\pm$0.07	1361.7
NB5	P97-77	25/172.37	$5.8879\text{x}10^{-4}$	$\pm$0.09	1377.7
NB6	P97-78	25/172.37	$5.8692\text{x}10^{-4}$	$\pm$0.06	1393.6
NB7	P97-76	25/172.37	$5.8353\text{x}10^{-4}$	$\pm$0.04	1409.5
NB8	P97-72	25/172.37	$5.8513\text{x}10^{-4}$	$\pm$0.06	1425.4

Smart (2010). This mass flow rate can be expressed in the form

$$\dot{m}_f = \alpha_j P_{L,i}^{\frac{\gamma-1}{2\gamma}} P_0^{\frac{\gamma+1}{2\gamma}}, \tag{D.4}$$

where $P_{L,i}$ is the initial Ludwieg tube pressure, P_0 is the plenum stagnation pressure, and α_j is a calibration constant that is a function of the injector discharge coefficient.

Since the Ludweig tube volume is constant, the magnitude of the drop in Ludweig tube pressure during the time the solenoid valve is open can be used to calculated mass flow rate directly. The injection coefficient can therefore be calculated for any initial Ludweig tube pressure using the expression

$$\alpha_j = \frac{V_L}{RT_L} \frac{\Delta P_L}{P_{L,i}^{\frac{\gamma-1}{2\gamma}} \overline{P_0}^{\frac{\gamma+1}{2\gamma}} \Delta t} \tag{D.5}$$

where V_L is the volume of the Ludweig tube, $\overline{P_0}$ is the average plenum stagnation pressure over the time the solenoid valve is open, Δt (Doherty, 2013).

This calibration constant is calculated for a wide range of initial Ludwieg tube pressures, by measuring the average steady pressure in each plenum in each case. As shown in Figure D.1, alpha is quite insensitive to a wide range of plenum (and therefore Ludweig tube) pressures, and can therefore be assumed constant for each injector plenum. This in turn yields the observed relationships between the measured plenum pressure and injector mass flow rate.

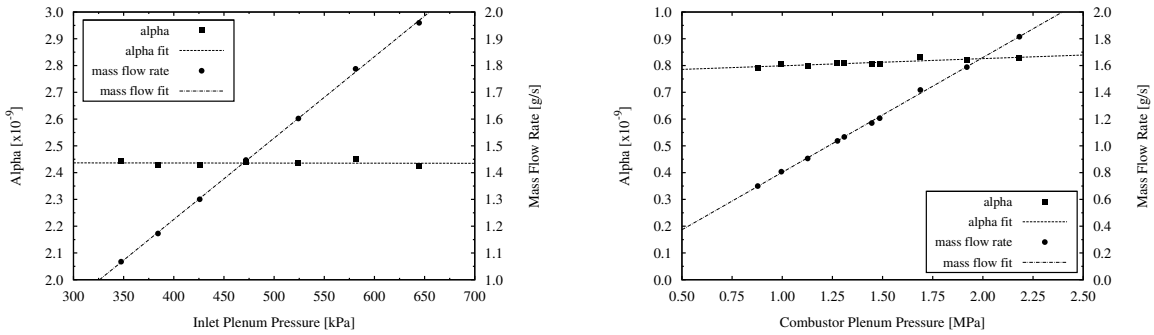

(a) Inlet Plenum **(b)** Combustor Plenum

Figure D.1: Calibration constant and mass flow rate pressure variance for each fuel plenum

Subject to experimental uncertainty, the instantaneous mass flow rate through a plenum at any moment during the test time can be calculated using equation D.4. As was mentioned in Chapter 3, the inlet plenum pressure is actually still rising during

the test time, but does so very slowly relative to the test time duration. The average steady mass flow rate can therefore be obtained by integrating the plenum pressure trace (Turner and Smart, 2010), and is considered accurate to within the calculated uncertainty below.

D.4.2 Uncertainty Analysis

Based on the equations introduced in section D.4.1, it is possible to use the RSS method to calculate an experimental uncertainty in the fuel mass flow rate. To begin, consider the formula for the calibration constant α_j, shown in equation D.5. It can be assumed that the universal gas constant, R, the specific heat ratio γ, and the pressure measurement interval Δt all have a negligible uncertainty. The Ludweig tube volume was measured by Chan (2012) to within 3% accuracy. The Ludweig tube was allowed to equilibrate with the surrounding atmosphere, and can be assumed (consistent with the previous assumption regarding the shock tube initial temperature) to have an internal temperature of 300 K $\pm$ 2%. The initial and final pressures of the Ludwieg tube can be measured to within 2% (Chan, 2012). The plenum pressures were measured using Kulite pressure transducers which have the previously calculated measurement uncertainty of 2.2%.

Applying these values to an RSS error calculation for equation D.5, the total uncertainty in α_j was found to be 4.7%. Armed with this uncertainty in the calibration constant, it was therefore possible to apply the RSS method to equation D.4, yielding a total uncertainty in the calculated injector mass flow rate of 5.1%.

Appendix E

Shot-Specific Inflow Conditions

This appendix contains the inflow conditions used to match specific experiments, as well as their pressure- and stagnation-enthalpy-matched flight conditions. The nozzle exit conditions were calculated using NENZFr, as detailed in Section 3.4. The flight conditions were calculated for the 1976 standard atmosphere, using an online calculator provided by Stanford University's Aircraft Aerodynamics and Design Group (Kroo, 1997).

Fuelling conditions are included where appropriate, and are listed as the mass flow of half the engine, to allow easy comparison between simulation and experiment.

Table E.1: Shot 11464 Tunnel and Equivalent Flight Conditions

Variable	Nozzle Exit Condition	Equivalent Flight Condition
P, Pa	1092.7	372.0
ρ, kg/m^3	0.0095253	0.00528
U, m/s	3658.8	3713.29
T, K	395.87	245.11
M, K	9.145	11.83
Re_u, $1/m$	1.52 x 10^6	1.25 x 10^6
q, kPa	63.76	36.46
H_0, MJ/kg	7.141	7.141
Y_{N_2}	0.72837	0.767
Y_{O_2}	0.18258	0.233
Y_{NO}	0.082921	0.0
Y_O	0.006129	0.0
Altitude, km	-	37.87
$\dot{m}_{H_2,inlet}$, g/s	0.70663	-
$\dot{m}_{H_2,comb}$, g/s	0.00000	-

Table E.2: Shot 11471 Tunnel and Equivalent Flight Conditions

Variable	Nozzle Exit Condition	Equivalent Flight Condition
P, Pa	1154.4	390.97
ρ, kg/m^3	0.010337	0.00557
U, m/s	3626.4	3674.53
T, K	385.73	244.1
M, K	9.185	11.73
Re_u, $1/m$	1.67 x 10^6	1.31 x 10^6
q, kPa	67.97	37.67
H_0, MJ/kg	6.996	6.996
Y_{N_2}	0.72845	0.767
Y_{O_2}	0.18365	0.233
Y_{NO}	0.082736	0.0
Y_O	0.005164	0.0
Altitude, km	-	37.52
$\dot{m}_{H_2,inlet}$, g/s	0.57717	-
$\dot{m}_{H_2,comb}$, g/s	1.28059	-

Table E.3: Shot 11486 Tunnel and Equivalent Flight Conditions

Variable	Nozzle Exit Condition	Equivalent Flight Condition
P, Pa	1131.3	3982.64
ρ, kg/m^3	0.010217	0.00545
U, m/s	3615.0	3662.0
T, K	382.46	244.5
M, K	9.195	11.68
Re_u, $1/m$	1.66 x 10^6	1.27 x 10^6
q, kPa	66.76	36.55
H_0, MJ/kg	6.951	6.951
Y_{N_2}	0.72846	0.767
Y_{O_2}	0.18372	0.233
Y_{NO}	0.082714	0.0
Y_O	0.005106	0.0
Altitude, km	-	37.67

Table E.4: Shot 11491 Tunnel and Equivalent Flight Conditions

Variable	Nozzle Exit Condition	Equivalent Flight Condition
P, Pa	1176.6	398.66
$\rho, kg/m^3$	0.010507	0.00569
U, m/s	3630.2	3678.41
T, K	386.79	243.71
M, K	9.183	11.75
$\mathrm{Re}_u, 1/m$	1.695 x 10^6	1.34 x 10^7
q, kPa	69.23	38.55
$\mathrm{H}_0, MJ/kg$	7.01	7.01
Y_{N_2}	0.72845	0.767
Y_{O_2}	0.18372	0.233
Y_{NO}	0.082728	0.0
Y_O	0.005102	0.0
Altitude, km	-	37.38
$\dot{m}_{H_2,inlet}, g/s$	0.502635	-
$\dot{m}_{H_2,comb}, g/s$	1.246795	-

Appendix F

The Effect of Nozzle Supply Pressure on Experimental Data

In general, the agreement between the simulations presented in this study and the experimental data they were intended to match is determined by comparing the average pressures measured along the M12 REST engine's length during the test time, with the simulated pressure distributions along a line passing through the measurement locations on the bodyside and cowlside surfaces of the engine. In general, the absolute pressure distributions in the simulations show reasonable, though not perfect, agreement. An example of this weak agreement is the absolute pressure comparison between the simulation of the unfuelled M12 REST engine and the experiment it was meant to replicate (Shot 11486) shown in Figure F.1a.

While the general trend of the experimental data is similar to the simulation in every case, at first it appears as though there are significant discrepancies between the magnitudes of peak pressures across both the bodyside and cowlside surfaces of the engine. However, if the comparison is made between the normalized experimental and simulated pressure distributions, the simulation shows far better agreement, as shown in Figure F.1b. The labelled zones in both parts of Figure F.1 show that these improvements in agreement are found across the engine's entire length, as well as on both the bodyside and cowlside of the combustor and nozzle.

This better agreement between normalized pressure distributions is universal to all the validation shots compared with experimental data. The simple conclusion to be drawn is that in order to achieve a proper comparison between the experiment and the simulation, the appropriate method for analyzing the experimental data for comparison to CFD is in its normalized form. Why the absolute pressures are not the appropriate measure for comparison is the subject of this appendix.

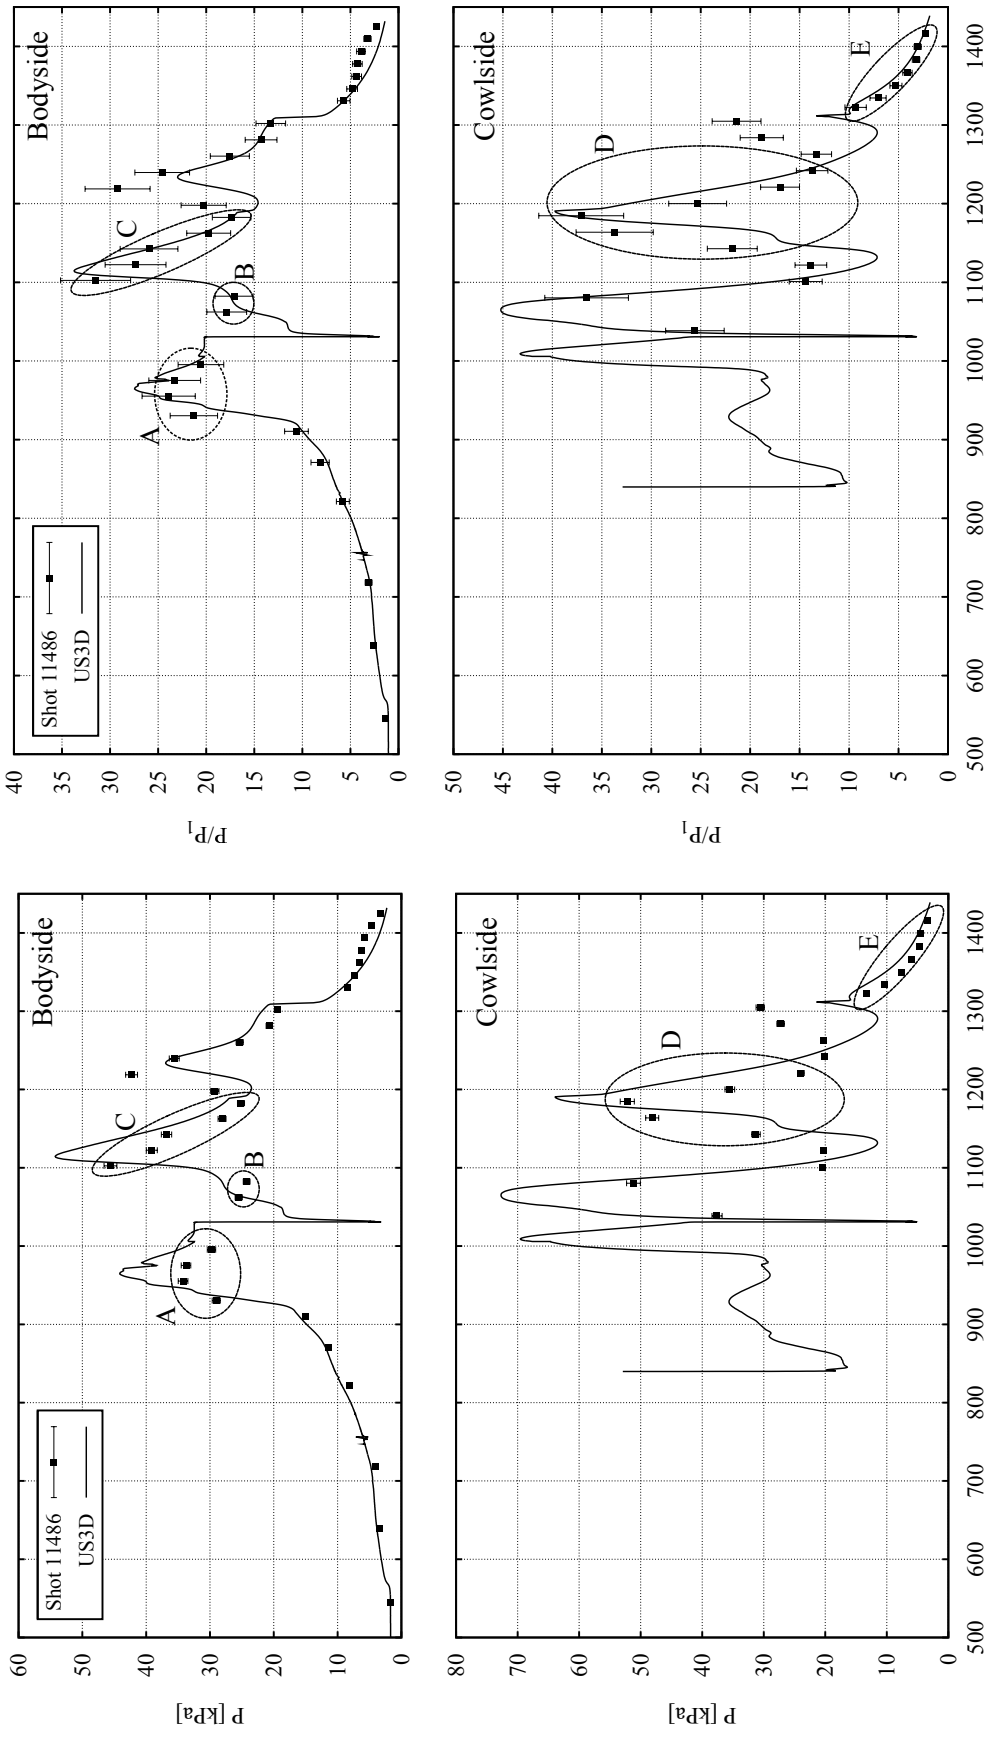

(a) Absolute Pressure Comparison

(b) Normalized Pressure Comparison

Figure F.1: Differences between the Absolute and Normalized Pressure Data for Shot 11486

The analysis begins by defining the pressure, P_1 that the simulated and experimental pressure distributions are normalized against. This pressure is defined as the nominal static pressure of the nozzle exit flow, after it has been processed by the forebody leading edge shock. This pressure is calculated by using the nominal nozzle exit static pressure calculated with NENZFr (this value is listed in Appendix E) and the oblique shock relations. As such, this pressure will be referred to as $P_{1,nom}$ for the rest of this appendix.

The absolute pressure distribution from the CFD simulation is normalized against this nominal forebody pressure, using the relation

$$\frac{P}{P_1} = \frac{P(x)}{P_{1,nom}} \tag{F.1}$$

where $P(x)$ is the pressure at any streamwise location along the line from which pressure values are extracted from the simulation. It should be emphasized at this point that $P(x)$ is an inherently time-averaged value, since it is extracted from a RANS simulation which was used to calculate the steady-state flow field of the M12 REST engine.

Before the normalized experimental pressure calculation can be presented, the absolute pressure measurement being normalized needs to be understood. This absolute measured pressure is calculated for each pressure transducer using the equation

$$\overline{P(t)} = \frac{\sum_{i=0}^{N} P(t_i)}{N} \tag{F.2}$$

where $P(t_i)$ is an instantaneous pressure measurement, i is the index number of each measurement included in the average, and N is the total number of measurements taken over the test time. As mentioned in Chapter 3, the test time over which the instantaneous measurements are averaged was determined by the onset of steady flow conditions in the engine, after the establishment of steady nozzle flow. All pressure measurements ie averaged over the same test time interval, which usually lasted 0.5 ms. The nozzle supply pressure trace for shot 11486, showing the interval over which test time was taken, is reproduced in Figure F.2.

The nozzle supply pressure is not steady (on average) during the test time, instead it slowly decreases with time. This is due to the under-tailored nature of the Mach 12 enthalpy tunnel condition, which causes the supply pressure (which is also the stagnation pressure of the flow) to decay slowly both before and during the test time. This pressure decline is accompanied by an associated decrease in the flow stagnation enthalpy.

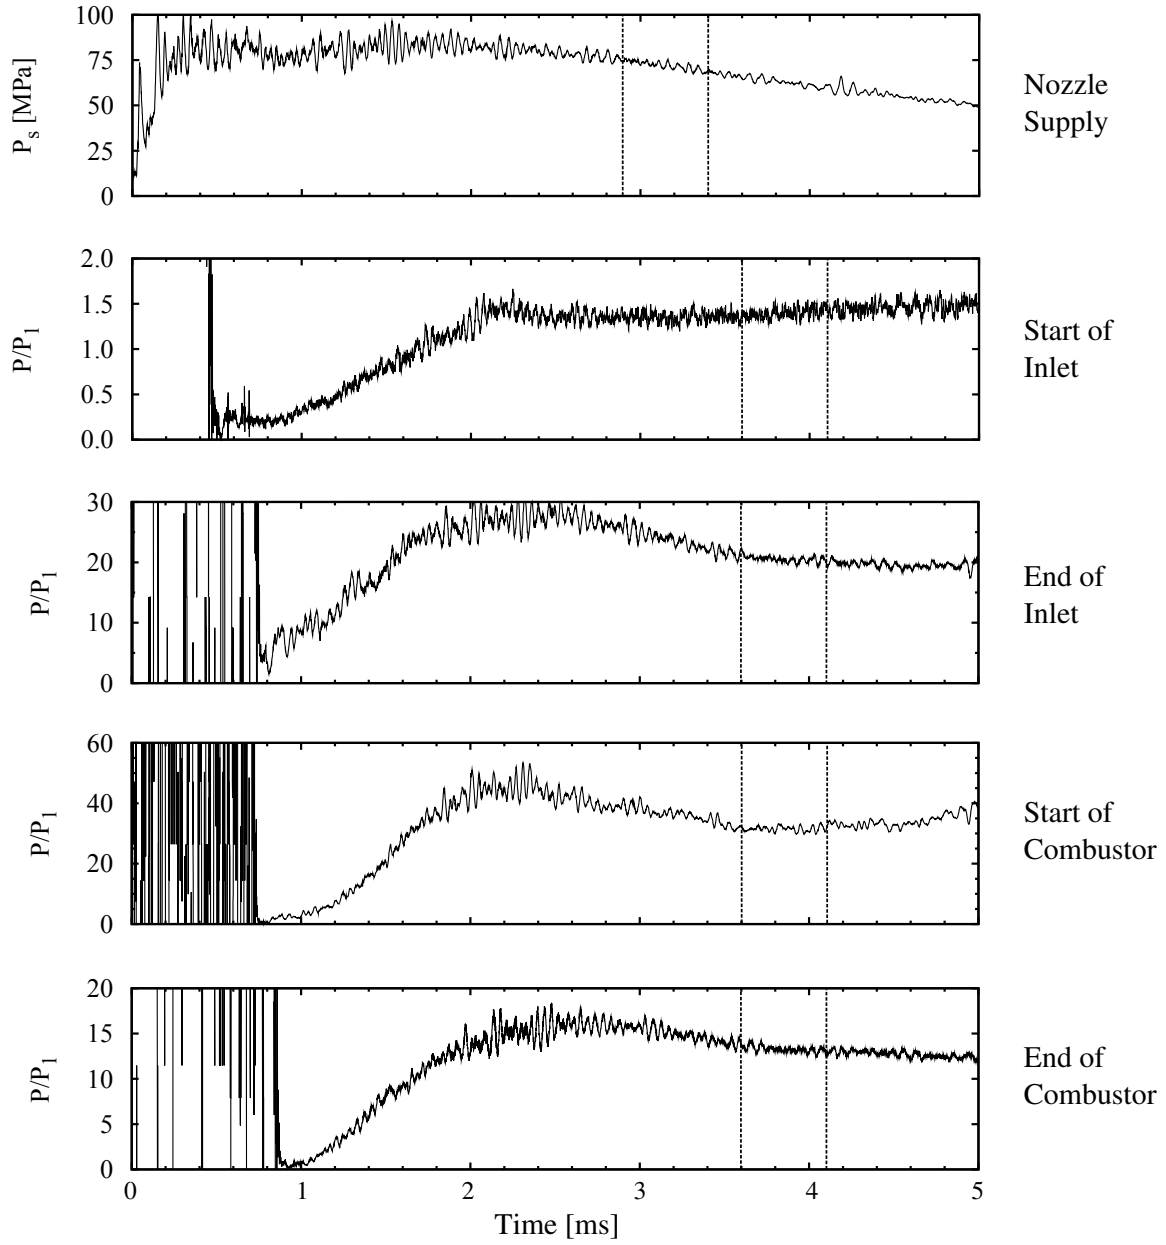

Figure F.2: Supply pressure and selected normalized engine pressure traces, showing test time interval

The pressure drop is only exacerbated by the late test time, caused by a longer-than-expected flow establishment time in the engine. This delay in test time, shown clearly in the individual pressure traces for select individual pressure transducers, leads to the averaging interval occurring outside the time when the nozzle supply pressure curve is relatively flat.

For the experimental normalization to be done, another fact needs to be accounted for: the pressure signal from the nozzle takes a finite time to reach the model, and indeed each pressure transducer along the model's length. This temporal delay in the pressure signal is large enough that, following the rupture of the secondary diaphragm at the nozzle throat, the start of the sudden pressure rise recorded by each transducer during flow establishment has a measurable offset from each transducer upstream of it. Effectively, this means that if the freestream conditions are changing upstream of the model, then the effects of this change are delayed by a measurable amount of time for each downstream pressure transducer. The pressure transducers in the engine nozzle are being affected by an average pressure signal that occurred earlier in time than the pressure signal being recorded by transducers along the engine's inlet.

The temporal lag in the nozzle pressure signal is normally accounted for in the normalization of the mean experimental pressure data, by using the following relation:

$$\frac{P}{P_1} = \overline{\frac{P(t)}{P_0(t-\tau_o)}} \frac{P_{0,nom}}{P_{1,nom}}, \tag{F.3}$$

where P_0 is the nozzle supply (stagnation) pressure. This relation accounts for the unique apparent nozzle pressure signal at each pressure transducer by first normalizing the absolute pressure signal against a time-shifted nozzle supply pressure signal. The offset in time, τ_o, is unique to each pressure measurement and is calculated as the time lag between the rupture of the secondary diaphragm, and the start of the sudden pressure rise recorded by each pressure transducer. The averaging process effectively shifts the averaging window for the nozzle supply region backward in time by the increment τ_o.

This time-offset normalized pressure is then renormalized to the engine forebody pressure, by multiplying it by the ratio of the nominal average nozzle supple pressure and the nominal forebody pressure that were defined previously. The pressure distribution that results is assumed to be equivalent to an engine operating under a steady flow condition, at the nominal supply pressure. For more details on this normalization procedure, refer to Chapter 3 of Doherty (2013).

Knowing that the normalized pressures are taking into account the unsteady nature of the nozzle supply pressure, the differences in the observed behaviour of the two interpretations of the pressure data are readily explained. In places where the two datasets differ significantly, it can be inferred that the absolute pressure measurement is reacting to a nozzle supply pressure that is different from the nominal supply pressure. In shots where the nozzle supply pressures has smaller variations between the test time and time-offset averaging periods, there is a more uniform pressure signal through the engine and less variation between the normalized and absolute measured pressure distributions.

To demonstrate this effect, the ratio $\frac{\overline{P_0(t-\tau_o)}}{P_{0,nom}}$ is calculated for the unfuelled and tailored injection validation experiments used in this study by dividing the ratio of $\overline{P(t)}$ and $P_{1,nom}$ by the normalized experimental pressure. This ratio can be thought of as the factor that each average pressure has to be divided by in order to account for the unsteadiness in the nozzle supply pressure. The distribution of this ratio is plotted in Figure F.3.

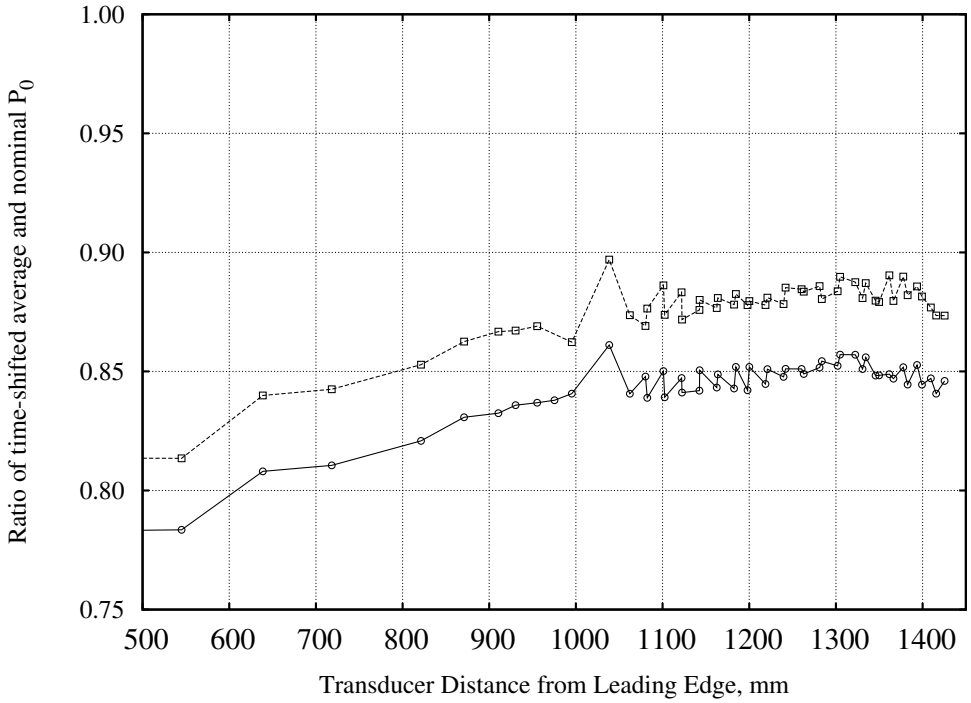

Figure F.3: The ratio of $\overline{P_0(t-\tau_o)}$ and $P_{0,nom}$ for each pressure transducer

While the absolute pressure traces of both shots match their respective simulations reasonably well, the absolute pressure magnitudes are typically low compared to the absolute pressure trace from the CFD simulations. This is due to the measured pressures being determined by a nozzle supply pressure that is lower than the nominal

supply pressure calculated from NENZFr, something that the normalization corrects for. The ratio plotted in Figure F.3 is therefore less than unity through the whole engine, which raises the normalized pressures to match the pressures at the nominal stagnation pressure.

The ratio of time-offset and nominal supply pressures is relatively constant, but does rise slightly with downstream distance. This is to be expected, since the average nozzle supply pressure decreases with time. As the averaging window for nozzle supply pressure is offset further back in time with increased downstream distance, the average stagnation pressure signal seen by downstream transducers is therefore closer to the nominal stagnation pressure. Therefore, the ratio the averaged and nominal supply pressures should be closer to unity at the exit of the nozzle than at the start of the inlet.

In conclusion, the CFD simulations examined in this study were conducted using a time-averaged steady state RANS flow solver, assuming inflow conditions based on the nominal nozzle supply pressure. As such, the most appropriate way to compare the simulations with experimental data is to compare them with the normalized pressure traces, as this normalization process accounts for the transient effects inherent to the test condition. It should be noted that this particular mismatch between absolute and normalized pressure trends is unlikely to be a common feature of tests at other flight Mach numbers in the T4 Stalker Tube. These experiments were limited both by the relatively long flow establishment times of the M12 REST engine, and the relatively short time in which the under-tailored nozzle supply pressure curve is relatively flat. In any situation where the nozzle supply pressure is comparatively more stable than the one used in the joint campaign in which these data were collected, the differences between the absolute and normalized pressure trends should be minimal.

Appendix G

Hypersonic Boundary Layer Combustion

This appendix is a reformatted reproduction of the analytical paper published by the author on the topic of fuel injection and combustion inside of a hypersonic turbulent boundary layer (Barth et al., 2013). The copyright of this paper rests with the author and has been included to supplement the discussion of this technique in Chapter 5.

Abstract

Drag reduction is important to improving the performance of scramjet engines operating at high Mach numbers. One demonstrated method for reducing skin friction drag on a surface exposed to hypersonic flow is the injection and combustion of hydrogen fuel in the boundary layer. However, there are other fuels of interest in scramjet applications, and the underlying mechanisms that drive the reduction of skin friction are not well understood. An existing analytical model for boundary layer combustion of hydrogen is re-derived for a general fuelling condition, and then extended to allow investigation of the underlying flow physics in this model. Applying this theory to ethylene fuelling indicates that skin friction reduction through boundary layer combustion is possible with fuels other than hydrogen. Analysis of the modelled boundary layer profiles demonstrates that skin friction reduction is accomplished through several coupled mechanisms: a change in near-wall viscosity, density changes and combustion act to reduce Reynolds stresses, and the low-momentum fuel stream thickens the boundary layer and changes the wall-normal velocity gradient. Finally, the theory is used to estimate the maximum fraction of fuel that should be used for skin friction reduction in a typical scramjet engine.

Nomenclature

A^* = $KU_e/[v^*(R^*)^{1/2}]$
C_{pi} = constant pressure specific heat of species i, J/(kg K)
C_{ps} = constant pressure specific heat of gas mixture, $=\sum_i c_i C_{pi}$, J/(kg K)
c_f = skin friction coefficient, $2\tau_w/(\rho_e U_e{}^2)$
c_i = mass fraction of species i
D = mass diffusivity, m^2/s
f = relative stoichiometric mass fraction of fuel
$fc_{oe}\Delta Q$ = heat release due to combustion of fuel in air, J/kg
g = streamwise velocity ratio, u/U at flame front
H = stagnation enthalpy, $C_{ps}T + u^2/2$, J/kg
h = static enthalpy, J/kg
h_{io} = enthalpy of formation of species i, J/kg
K = von Kármán constant, 0.41
k = molecular thermal conductivity, W/(m·K)
Le = Lewis number, Sc/Pr
M = Mach number
$\dot{m}$ = mass flow rate of fuel injected into boundary layer, kg/(s·m)
Pr = Prandtl number, $\mu C_{ps}/k$
R^* = gas constant and specific heat ratio, $R\cdot C_{psw}/R_w\cdot C_{ps}$
R_{ij} = Reynolds shear stress, $\rho\overline{u_i u_j}$, N/m^2
R_s = mixture gas constant
Re_x = Reynolds number, $\rho_e U_e x/\mu_e$
Sc = Schmidt number, ν/D
T = temperature, K
U_e = freestream velocity, m/s
u = streamwise component of velocity along the surface, m/s
u' = u/U_e
u_{inj} = fuel injection velocity, m/s
$\overline{u_i u_j}$ = averaged product of velocity fluctuations,m^2/s^2

v = surface-normal component of velocity, m/s
x = streamwise distance along surface, m
y = distance normal to surface, m
Z = Shvab-Zeldovich coupling coefficient
ΔQ = heat of combustion of fuel, J/kg
γ = ratio of specific heats
δ = boundary layer thickness, m
θ = momentum thickness, m
μ = molecular viscosity, kg/(m·s)
ν = kinematic viscosity, m^2/s
ρ = density, kg/m^3
τ = shear stress, N/m^2
v^* = friction velocity, $\sqrt{\tau_w/\rho_w}$

Subscripts

CO2 = carbon dioxide
e = mainstream
F = fuel
H2O = water
n = no fuel injection
o = oxygen
P = products
T = temperature
t = turbulent
w = wall

G.1 Introduction

Recently there has been a push for more economical access to space, which has spurred interest in air breathing hypersonic vehicles as a means to achieve this goal. One major challenge in bringing such a transportation system to fruition is the development of a supersonic combusting ramjet (scramjet) capable of operating at high altitudes and Mach numbers. At these flight conditions, predominantly turbulent boundary layers can be expected, which exacerbate an already sizable performance penalty due to drag. Skin friction drag losses inside the engine combustor are responsible for as much as 25% of the drag force generated by an engine (Tanimizu, 2008). In theory, this drag force can be reduced by controlling the hypersonic boundary layer attached to the vehicle. However, with the apparent dominance of the scramjet engine in generating this drag, substantial benefit could be achieved through boundary layer control inside the scramjet combustor alone. As such, a viable scramjet need not only see substantial improvements in combustion efficiency; arguably, minimizing engine drag losses is of nearly equal importance (Czysz and Bruno, 2009).

One well-investigated method for controlling turbulent boundary layers is the phenomenon of film cooling. Originally designed to reduce heat transfer rates from the hot boundary layer to the surface it is flowing over, film cooling relies on the injection of cold gas into the boundary layer, forming a buffer layer over the surface. This has an additional benefit: the skin friction drag acting on the surface is also reduced (Parthasarathy and Zakkay, 1970). Experiments have demonstrated higher skin friction reductions as flow speed increases (Cary and Hefner, 1972). Much of this effect is attributed to the injected gas reducing the near-wall momentum of the flow, thickening the boundary layer. This thicker boundary layer in turn will have a smaller wall-normal streamwise velocity gradient at the wall, and therefore reduce the shear stress (and therefore skin friction) at the wall, per equation G.1 below (Suraweera, 2006).

$$\tau_w = \mu_w \frac{du}{dy}_w \tag{G.1}$$

Significant work has also been done examining the effect of boundary layer combustion on skin friction. Kulgein (1962) examined the effects of combustion of methane injected into the turbulent boundary layer of a subsonic flow through a porous wall. Similar experiments with hydrogen injection were conducted by Wooldridge and Muzzy (1965). Kulgein's study found there to be little effect due to combustion, though any benefits may have been counteracted by an increase in skin friction due to the rough-

ness of the porous wall (Schetz and van Overeem, 1975). Woodridge and Muzzy's study showed far more promising results, including that the turbulent Prandtl and Schmidt numbers could be assumed to have a value of unity across the boundary layer (Wooldridge and Muzzy, 1965). This particular assumption will be used in the analysis to come.

There have also been several experiments involving the tangential slot injection of fuel along a wall surface. Burrows and Kurkov (1973) is a well-studied example of this injection geometry, which showed that a film of injected hydrogen will remain close to the wall despite combustion occurring in the outer edge of the boundary layer. It was concluded that both film cooling and combustion effects work together to provide a reduction in skin friction drag.

More recently, skin friction reduction through boundary layer combustion has been experimentally demonstrated (e.g. Goyne et al. (2000); Kirchhartz et al. (2010); Suraweera (2006)) in shock tunnel experiments involving hypersonic boundary layers. The study conducted by Goyne et al. (2000) on hydrogen combustion in a hypersonic turbulent boundary layer showed that the local skin friction drag could be reduced by as much as 70%. Computational work has shown similar results: Levin and Larin showed that energy addition via electrical discharge in a turbulent boundary layer could provide significant skin friction reductions (Levin and Larin, 2003). Heat addition analogous to the heat release from hydrogen boundary layer combustion has the same effect (Denman, 2007).

While experiments and simulations both have clearly demonstrated the potential utility of boundary layer combustion, a simple method for predicting this reduction in skin friction drag was not available until 2005, when Stalker published an analytical theory for hydrogen combustion in a hypersonic turbulent boundary layer (Stalker, 2005). The theory agress well with available experimental skin friction data for hydrogen combustion in hypersonic turbulent boundary layers, but is not without its limitations. While hydrogen is certainly an attractive fuel for scramjet applications, there are many potential cases for which hydrocarbon fuels may be more appropriate (Townend, 2001).

The accuracy of Stalker's model indicates that the theory captures at least some of the flow physics driving skin friction reduction in the presence of combustion, but its present form does little to illuminate how a choice in operating parameters is physically affecting boundary layer: the predicted skin friction coefficient is merely presented as a data point for design. While this is certainly a valuable engineering tool, a deeper

understanding of the flow physics may aid the designer further when considering the integration of a boundary layer combustion system into an engine or vehicle.

Thus, the motivation for the study that follows is clear: to re-derive the theory presented by Stalker (2005) for a more general fuelling condition, and then extend its reach to elucidate the underlying mechanisms that reduce skin friction. After a brief review of the flow model used, the paper continues with a derivation of the generalized forms for the skin friction and fuel mass flow rate equations. Equations for the variation in primitive variables and eddy (turbulent) viscosity across the boundary layer are developed. The updated theory is then used to contrast the performance of hydrocarbon fuel injection with hydrogen injection under identical conditions, as a starting point upon which to base an examination of the physical flow mechanisms that lead to reductions in skin friction drag. Finally, the analysis is used to predict the amount of fuel that is necessary for effective boundary layer combustion in a scramjet engine of the size typically examined in current experimental work.

G.2 Flow Model

In order to analytically solve combusting flow along the wall of a hypersonic vehicle or scramjet combustor, a simplified model is necessary. The flow case considered in this analysis is a two-dimensional, unbounded zero-pressure-gradient compressible hypersonic turbulent boundary layer flowing along a flat plate. Gaseous fuel is injected tangentially along the surface in the streamwise direction, and diffuses through the boundary layer. In reality, this fuel injection is most often done through a slot, as shown in Figure G.1a. Downstream of this injection point, the flow develops to yield a constant pressure across the boundary layer, and ultimately boundary layer profiles that exhibit local similarity. In this region between injection and locally similar flow, a layer of fuel next to the surface is separated from the freestream by a mixing layer and the fuel mass fraction at the wall, c_{Fw}, equals 1. It is assumed that for low fuel mass flow rates this mixing region is very short, and no combustion takes place within it. When $c_{Fw} = 0$, combustion ceases, and the flow contains only air and combustion products. The flow solution in this region is also not modeled.

The flow model is thus only derived in the region where the boundary layer is self-similar, and therefore simplifies to the case shown in Figure G.1b: a flat plate with fuel travelling along the wall from the leading edge.The model is solved by assuming fuel is injected at the leading edge at a constant mass flow rate, then allowing c_{Fw} to vary from 1 to 0 as fuel diffuses away from the wall downstream of the injection point.

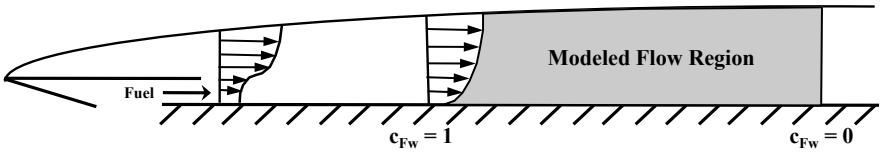

(a) Physical Injection Geometry

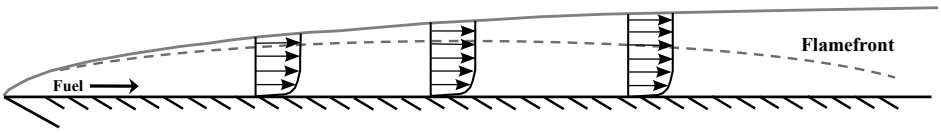

(b) Modeled Flow Geometry

Figure G.1: Comparison of Physical and Modeled Flow Geometry

The diffusing fuel may eventually react with oxygen from the freestream to form a flamefront inside the boundary layer. More about the combustion reaction itself is presented in the next subsection.

In choosing a boundary layer model upon which to base the original analysis, Stalker (2005) chose the compressible turbulent boundary layer model of Van Driest (1951), (1956). The model developed in Van Driest (1951) allows calculation of the effect on skin friction caused by temperature-induced density changes in a turbulent boundary layer, and has successfully predicted skin friction on flat plates in hypersonic flows. This model lends itself well to the problem of combustion in a turbulent boundary layer: temperature change due to combustion will alter the density, having a strong effect on the Reynolds stresses that dominate the value of shear stress inside a turbulent boundary layer. Assuming that $Pr_t = 1$ across the boundary layer, an expression for density can be obtained from the variation in flow stagnation enthalpy with velocity through the Crocco relation:

$$\rho_w/\rho = 1 + bu' - a^2u'^2 \tag{G.2}$$

where $u' = u/U_e$, $a^2 = U_e^2/2H_w$ and $b = (H_e - H_w)/H_w$.

From Van Driest II (1956), the turbulent shear stress in the boundary layer is given by the Prandtl mixing length equation:

$$\tau = \rho l^2 \left(\frac{du}{dy}\right)^2 \tag{G.3}$$

where in this case, l is taken as the von Kármán mixing length:

$$l = \left[K\left(\frac{du}{dy}\right) \Big/ \left(\frac{d^2u}{dy^2}\right) \right] \tag{G.4}$$

The resulting shear stress τ is assumed constant near the wall, and therefore equal to its surface value τ_w. This equation forms the basis for the detailed analysis presented in the following section.

Combustion Model

The assumed model for fuel combustion inside the boundary layer requires attention in order for the analysis to proceed. As stated previously and illustrated in Figure G.1b, the fuel begins flowing along the surface, then diffuses through the boundary layer toward the freestream. When the fuel encounters oxygen, it is assumed to combust instantaneously to form an equilibrium non-premixed diffusion flamefront. This yields the result that oxygen and fuel will never coexist in their molecular forms; beneath the flamefront all oxygen is consumed, and above it no fuel is present, as demonstrated for hydrogen combustion in Figure G.2. Similar curves can be obtained for any fuel.

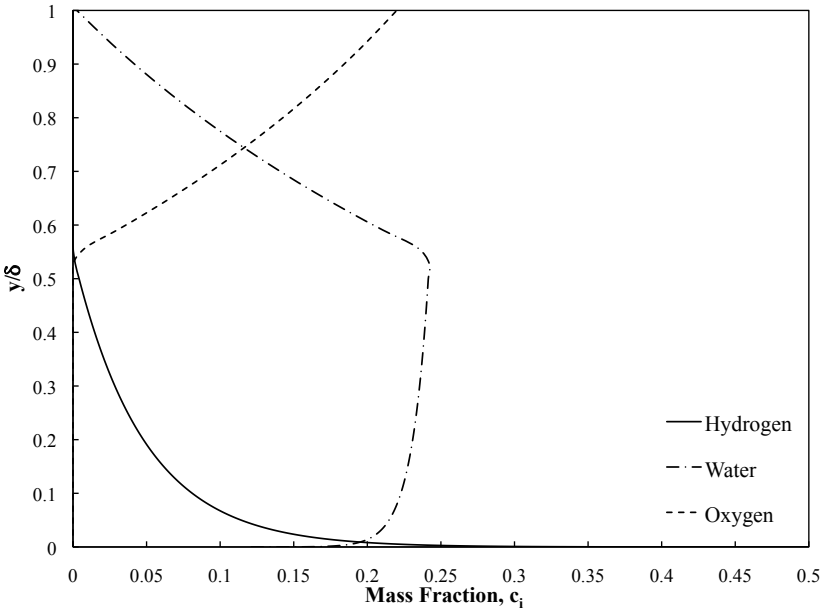

Figure G.2: Typical variation in species mass fractions for hydrogen combustion in a boundary layer ($c_{Fw} = 0.5$)

The complete combustion reaction is assumed to occur in a single step, taking the

form:

$$O_2 + \chi_F(Fuel) \longrightarrow \sum_i \chi_i P_i \tag{G.5}$$

where χ is the stoichiometric fraction of each species relative to oxygen, and P_i is a placeholder for each reaction product. This reaction is of the same form used in Stalker (2005) (hereafter referred to as "S2005"), and indeed simplifies easily to the single-step reaction for hydrogen combustion.

This combustion model is obviously a simplification of the real process, and would be limited by some minimum ignition temperature. For hypersonic flows, the temperature across most of the boundary layer is expected to be higher than the auto-ignition temperature of the fuels typically used in scramjet applications. On this basis, it is assumed that the reaction scheme above is sufficent, though it will overestimate the heat release of combustion, and therefore exaggerate its effects. It can therefore be presumed to give an upper limit for the action of combustion on the flow physics inside the boundary layer.

G.3 Analysis

In re-deriving expressions for the skin friction coefficient for a combusting turbulent hypersonic boundary layer, the analysis follows the same procedure as outlined in S2005, and starts by considering the boundary layer density profile. The flow is assumed to have Prandtl and Lewis numbers of unity, which was previously shown to be reasonable by Wooldridge and Muzzy (1965). Using Shvab-Zeldovich coefficients (Zeldovich, 1951) to couple species and thermal flow properties, the two-dimensional boundary layer momentum-, species-, and energy-transfer equations can be written in a common form:

$$\rho u \frac{\partial Z}{\partial x} + \rho v \frac{\partial Z}{\partial y} = \frac{\partial}{\partial y}\left[(\mu + \mu_t)\frac{\partial Z}{\partial y}\right] \tag{G.6}$$

where μ_t is the eddy viscosity, and Z is any of the Shvab-Zeldovich coefficients $Z_{Fo} = c_o - \frac{c_F}{f}$; $Z_{oT} = c_o + \frac{H}{f\Delta Q}$; and $Z_{FT} = c_F + \frac{H}{\Delta Q}$. The typical variation in any of these coefficients is shown in Figure G.3.

With the exception of Z_{Fo}, the variation in coupling coefficient across the boundary layer takes a form closely resembling that of the related species mass fraction. The different profile for Z_{Fo} can be largely explained by recognizing it is a weighted difference of the species profiles for fuel and oxygen.

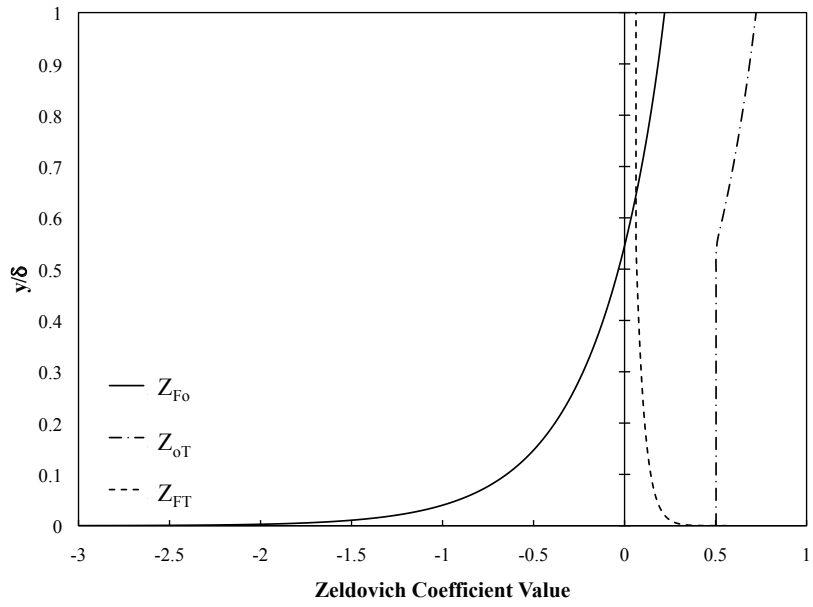

Figure G.3: Typical variation in coupling coefficients across a combusting boundary layer ($c_{Fw} = 0.5$)

The general form in equation G.6 can be shown to have the solution

$$Z = d_1 + d_2 u \tag{G.7}$$

where d_1 and d_2 are constants that can be determined from the boundary conditions for the coupling coefficients. The fuel species mass fraction, c_F, is defined to be $c_F = c_{Fw}$ at the wall, and $c_F = 0$ at the flamefront and boundary layer edge. Similarly, oxygen fractions at the wall and flamefront are $c_o = 0$, and $c_o = c_{oe}$ at the boundary layer edge. The velocity boundary conditions enforce the no-slip condition at the wall, and $u = U_e$ at the boundary layer edge. The normalized velocity at the flamefront inside the boundary layer, g, can be found by evaluating equation G.7 for Z_{Fo}:

$$g = \frac{c_{Fw}}{c_{Fw} + f c_{oe}} \tag{G.8}$$

Evaluating the solutions for Z_{FT} and Z_{oT}, expressions for H on both sides of the flamefront are obtained. For $0 < u' < g$

$$\frac{C_{ps} T}{H_w} = 1 + b_1 u' - a^2 u'^2 \tag{G.9}$$

and for $g < u' < 1$

$$\frac{C_{ps} T}{H_w} = 1 + \alpha + b_2 u' - a^2 u'^2 \tag{G.10}$$

where $H_w = C_{psw}T_w$; $b_1 = (H_e - H_w + fc_{oe}\Delta Q)/H_w$; $b_2 = (H_e - H_w + c_{Fw}\Delta Q)/H_w$, and $\alpha = c_{Fw}\Delta Q/H_w$.

Pressure remains constant across the boundary layer. So, if it is assumed that the gas constant and constant-pressure specific heats for each species remain invariant with temperature, then the following relationship for density is obtained:

$$\frac{\rho_w}{\rho} = \frac{R_s C_{psw}}{R_{sw} C_{ps}} \frac{C_{ps} T}{H_w} = R^* \frac{C_{ps} T}{H_w} \tag{G.11}$$

Gas Property Ratio

Up until this point in the analysis, all the changes necessary to generalize the fuelling condition have been implemented by generalizing the chemistry model, and changes to the Shvab-Zeldovich coupling coefficients. The gas property ratio R^*, on the other hand, requires some change in the analysis. In all fuelling cases, the values for C_p for each gas species are taken at their constant value for an intermediate temperature between the wall and freestream temperatures examined in S2005. (The effect of this assumed temperatue for property determination is examined later.) R^* can be therefore expressed entirely in terms of species mass fractions, and through the Shvab-Zeldovich coefficients can be represented by a function of c_{Fw} and one other species. This expression for R^* as a function of the fuel species mass fractions is plotted for various values of c_{Fw} in Figure G.4.

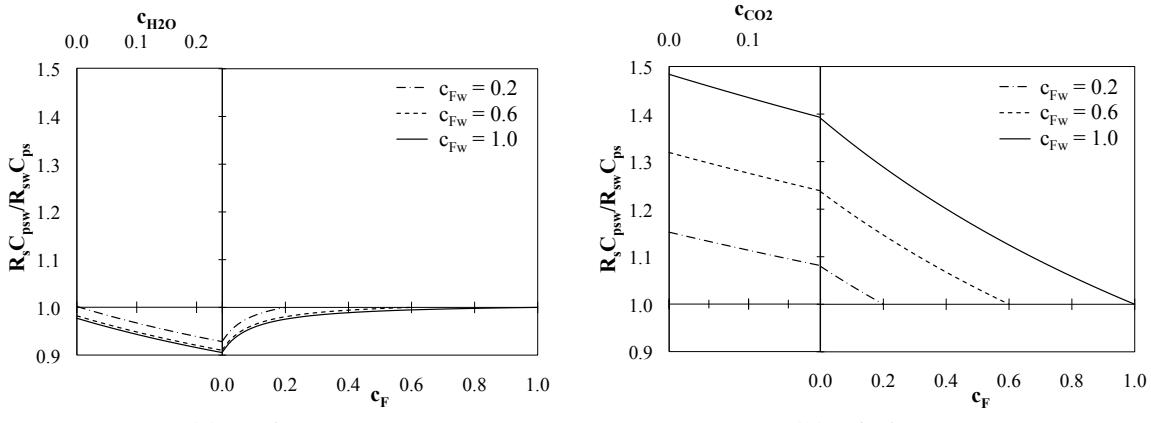

(a) Hydrogen **(b)** Ethylene

Figure G.4: Behavior of R^* across boundary layers of varying c_{Fw} for (a) Hydrogen, and (b) Ethylene.

Note that because $c_F = 0$ when $u' \geq g$, R^* is plotted as a function of one of the product species mass fractions in left side of each plot, where the right hand side of the

plot corresponds to the boundary layer edge, and the central axis is the location of the flamefront. On the right side of the vertical axis, the relative location of the wall can be inferred from the point where $c_F = c_{Fw}$.

On the basis of the behavior of R^* for hydrogen plotted in Figure G.4a, Stalker assumed it could be approximated as equal to one. This would yield a maximum deviation of ∼10% from the exact value for R^* when $c_{Fw} = 1$. As clearly demonstrated for a hydrocarbon fuel, in this case ethylene, in Figure G.4b, the same assumption would yield a maximum error of nearly 45%. Thus, it is prudent to consider an alternative simplification for R^*.

One attractive choice is to simply model R^* as a local constant with a value greater than one; i.e. R^* varies only as a function of c_{Fw}. This has several advantages, not the least of which is that the exact form for R^*, or even a linear approximation to it, quickly renders the analysis that follows intractable. If a constant value for R^* is assumed, what can be said is being assumed about the boundary layer, exactly? The implicit definition given for R^* in equation G.11 can be simplified using the definition of the gas constant $R = C_p - C_v$ to yield

$$R^* = \frac{\gamma - 1}{\gamma} \frac{\gamma_w}{\gamma_w - 1} \tag{G.12}$$

Thus, assuming a constant value for R^* is equivalent to assuming that away from the wall, the ratio of specific heats is constant across the boundary layer. Note that this does not imply constant C_{ps} across the boundary layer; full variation in specific heats across the boundary layer is included in the model.

For hydrogen fuel, this assumption is quite accurate since both the hydrogen and air have $\gamma \approx 7/5$, so significant variations in γ occur only in the vicinity of the flame front where the product species concentrations are high. For hydrocarbon fuels, the ratio of specific heats is significantly less than $7/5$, leading to significant differences in γ between the wall and freestream for high values of c_{FW}.

The question remains: what constant value should R^* take? Choosing the integrated mean value of R^* (and therefore of γ) across the boundary layer as its constant value is a logical choice: the maximum deviation from the exact value of R^* will be approximately 20% (for $c_{FW} = 1$). Since the value of c_{Fw} drops rapidly as the flow moves downstream from its starting point, the deviation from the exact R^* drops to the same order as S2005 in short order. As such, this assumption can be considered as adequate for the purposes of carrying forward with the analysis with R^* held constant in equation G.11.

Skin Friction and Fuel Mass Flow rate

Having settled on a fixed form for the variation in density across the boundary layer, the skin friction may be solved for using the boundary layer momentum integral equation, employing the boundary layer momentum thickness. Before proceeding, the momentum thickness must be derived. Combining equations G.3 and G.4, following Van Driest (1951) and S2005 the resulting expression is integrated to give

$$v^* \ln \frac{dy}{du} + constant = K \int \sqrt{\frac{\rho}{\rho_w}} du \tag{G.13}$$

where the friction velocity $v^* = \sqrt{\tau_w/\rho_w}$. From this equation, the rate of change in distance from the wall with velocity below and above the flamefront can be obtained. This is done by combining equation G.11 with equations G.9 and G.10 respectively, and substituting the result in place of ρ/ρ_w in equation G.13. The equation obtained for $0 \leq u' \leq g$ is:

$$\frac{dy}{du'} = U_e \frac{\nu_w}{(v^*)^2} \exp\left[B_0 - \frac{A^*}{a}\Big(\arcsin\frac{b_1 - 2a^2u'}{Q_1} - \arcsin\frac{b_1}{Q_1}\Big)\right] \tag{G.14}$$

where $A^* = KU_e/[v^*(R^*)^{1/2}]$, $Q_1 = \sqrt{b_1^2 + 4a^2}$ and B_0 is a constant. Similarly, the expression obtained for $g \leq u' \leq 1$ is

$$\frac{dy}{du'} = U_e \frac{\nu_w}{(v^*)^2} \exp\left[B_1 - \frac{A^*}{a}\Big(\arcsin\frac{b_2 - 2a^2}{Q_2} - \arcsin\frac{b_2 - 2a^2u'}{Q_2}\Big)\right] \tag{G.15}$$

where B_1 is also a constant, and $Q_2 = \sqrt{b_2^2 + 4a^2(1+\alpha)}$. Since the velocity must be continuous across the flamefront, B_1 can be expressed in terms of B_0 by setting equations G.14 and G.15 to be equal when $u' = g$.

Next, the expression for momentum thickness, which is required in the boundary layer momentum integral equation, is split into portions above and beneath the flamefront as follows:

$$\theta = \Big(\frac{\rho_w}{\rho_e}\Big)\left[\int_0^g \frac{\rho}{\rho_w}(1-u')u'dy + \int_g^1 \frac{\rho}{\rho_w}(1-u')u'dy\right] \tag{G.16}$$

Both integrals can be expanded into a series in A^{*-1}, and employing the fact that $A^* \gg 1$, terms with negative exponents greater than A^{*-2} can be discarded. This assumption follows from the analysis of S2005, and can be shown to hold true for both

hydrogen and at least one hydrocarbon fuel: ethylene. By adding the resulting expansions of each integral together, the final form of the boundary layer momentum thickness is found to be

$$\theta = \frac{\rho_w \nu_w}{\rho_e (v^*)^2} \frac{U_e}{R^* (A^*)^2} \exp\left[B_0 + \left(\frac{A^*}{a}\right) F_1\right] \left(1 + G \exp\left[\left(\frac{A^*}{a}\right)(F_2 - F_1)\right]\right) \tag{G.17}$$

where

$$G = \frac{0.5(b_1 - b_2) g (1 - g)}{1 + b_1 g - a^2 g^2}, \tag{G.18}$$

$$F_1 = \arcsin \frac{b_1}{Q_1} + \arcsin \frac{2a^2 - b_2}{Q_2} + \arcsin \frac{2a^2 g - b_1}{Q_1} - \arcsin \frac{2a^2 g - b_2}{Q_2}, \tag{G.19}$$

and

$$F_2 = \arcsin \frac{b_1}{Q_1} + \arcsin \frac{2a^2 g - b_1}{Q_1}. \tag{G.20}$$

At this point, one final piece of information is required to solve the boundary layer momentum balance: an expression for the total mass flow of fuel in the boundary layer. As shown in S2005, by using the Shvab-Zeldovich coefficient for fuel and product species, $Z_{FP} = c_{Fw} + (f \sum c_P)/(1 + f)$, the total mass flowrate of fuel can be shown to be

$$\dot{m} = \frac{c_{Fw} + f c_{oe}}{1 + f c_{oe}} \rho_e U_e \theta. \tag{G.21}$$

Remembering now that this analysis is based in part on the assumption that fuel is continuously injected along the wall surface, the wall shear stress can be related to the momentum thickness and fuel mass flow by differentiating the momentum balance equation of the boundary layer with respect to x, obtaining

$$\tau_w + (U_e - u_{inj}) \frac{d\dot{m}}{dx} = \frac{\rho_e U_e^2 d\theta}{dx} \tag{G.22}$$

which can be solved for c_f by substituting in equation G.17 and integrating with respect to A^*, remembering that as $x \to 0$, $A^* \to 0$ and $c_f \to \infty$. The resulting implicit equation in c_f should collapse to the same form as developed by Van Driest, and therefore, B_0 is

set to -3.462. The final form of the relation for the coefficient of skin friction is

$$4.15\log\left(c_f Re_x \mu_e/\mu_w\right) + 1.7 = \frac{F_1\sqrt{2}}{M_e\sqrt{c_f(\gamma_e - 1)}} + 3.97\log\left(\frac{1 - c_{Fw} + (u_{inj}/U_e)(c_{Fw} + fc_{oe}}{1 + fc_{oe}}\right) + 3.97\log\left(1 + G\exp\left[\frac{2(F_2 - F_1)K}{M_e\sqrt{c_f(\gamma_e - 1)}}\right]\right) \quad \text{(G.23)}$$

where the logarithm terms have been changed to base 10, and the left hand side leading term has been raised from 3.97 to 4.15, in order to agree with Van Driest's formulation for c_f in a compressible turbulent boundary layer when $c_{Fw} = 0$. This value in Van Driest's model was in turn set in order to match incompressible boundary layer theory as $M_e \to 0$.

By substituting equation G.17 into equation G.21, along with $B_0 = -3.462$, a final form for the fuel mass flow rate in a combusting boundary layer is also obtained:

$$\dot{m} = 0.187\mu_w\left[\frac{c_{Fw} + fc_{oe}}{1 + fc_{oe}}\right] exp\left(\frac{2F_1K}{M_e\sqrt{c_f(\gamma_e - 1)}}\right)\left[1 + G\exp\left(\frac{2(F_2 - F_1)K}{\sqrt{c_f(\gamma_e - 1)}}\right)\right] \quad \text{(G.24)}$$

With equations G.23 and G.24, the local skin friction coefficient can be determined for any rate of fuel mass injection into the boundary layer. By assuming a constant value for $\dot{m}$, equation G.24 can be solved iteratively for c_f for successive values of c_{Fw} between 1 and 0. These related values for c_f and c_{Fw} may then be used to obtain a value for Re_x (as an analogue to streamwise position) from equation G.23.

Since the viscosities at the wall and freestream boundaries must be known to obtain a solution, they must also be calculated. The Sutherland formula for each species can be used to obtain their individual viscosities at a given temperature, from which a mixture viscosity may in turn be obtained by applying Wilke's method as detailed in Blake (1970).

One characteristic of note in both equations G.23 and G.24 is that c_f and $\dot{m}$ both appear to be independent of R^*. While the latter immediately makes physical sense (there is no reason why fuel mass flow rate should depend on properties of the flow), it is interesting that the formulation for c_f shows no explicit dependence on R^*. R^* certainly has some effect on the underlying physical mechanisms that lead to skin friction reduction, the simplifications that also lead to the elimination of A^* in equation G.23

have simply made dependence on R* implicit. Given its relation to C_{ps} (which is itself related to c_{Fw} and ρ demonstrated in equation G.11, the effects of R^* are integrated into the final expression through a variety of terms, including the numerators of F_1 and F_2.

Skin Friction with Suppressed Combustion

It is possible to discern the relative magnitude of the effects of combustion alone if the skin friction reduction of an identical amount of fuel used for film cooling is known. The derivation of these equations is quite similar to the analysis presented in the previous subsections. The principal difference lies in the definition of the coupling coefficients and the value chosen for R^*. If no combustion takes place, the boundary layer will be a mixture of air and fuel only. No flamefront will be present, no heat release occurs, and many of the defined variables simplify to forms identical to the Van Driest II analysis. Making the definitions $Q = (b^2 + 4a^2)^{1/2}$ and $F_0 = \arcsin\left[(2a^2 - b)/Q\right] + \arcsin\left(b/Q\right)$, it is possible to define the skin friction and mass flow relations as

$$4.15 \log\left(c_f Re_x \mu_e/\mu_w\right) + 1.7 = \frac{F_0\sqrt{2}}{M_e\sqrt{c_f(\gamma_e - 1)}} + 3.97 \log\left(1 - c_{Fw} + (u_{inj}/U_e)c_{Fw}\right) \tag{G.25}$$

and

$$\dot{m} = 0.187 \mu_w c_{Fw} \exp\left(\frac{2F_0 K}{M_e\sqrt{c_f(\gamma_e - 1)}}\right) \tag{G.26}$$

respectively. Setting $c_{Fw} = 0$ in the above equations will naturally set $\dot{m}$ equal to zero, and retrieve the Van Driest skin friction coefficient for a boundary layer, allowing for a direct evaluation of the efficacy of film-cooling flows in reducing skin friction relative to an unmodified boundary layer.

Boundary Layer Profiles

While equations G.23-G.26 are useful in determining the net effect of fuel injection and combustion on skin friction generation in a turbulent boundary layer, they do not provide any explanation of the underlying physical mechanisms that produce the observed reductions in c_f. So, it is worth some extra effort to derive expressions for the boundary layer profiles of velocity, density, temperature, and eddy viscosity. Recall that density and temperature are already defined in terms of the velocity of the flow at any point inside the boundary layer through equations G.9-G.11. As such, determining the boundary layer profiles of the primitive flow variables requires only that the velocity variation with wall distance, y, be defined. This can be obtained in implicit form by

integrating equations G.14 and G.15, as well as their non-combusting-flow equivalent, to obtain the wall-normal distance, y, as a function of u'. For the combusting boundary layer, for $0 \leq u' \leq g$:

$$y = Pe^{B_0}\Bigg[(A^*\sqrt{Q_1^2 - (2a^2u' - b_1)} + 2a^2u' - b_1) \times \exp\left(A^*\left[\arcsin\left(\frac{b_1}{Q_1}\right) + \arcsin\left(\frac{2a^2u' - b_1}{Q_1}\right)\right]\right) - (A^*\sqrt{Q_1^2 - b_1^2} - b_1)\Bigg], \quad \text{(G.27)}$$

for $g \leq u' \leq 1$:

$$y = Pe^{B_0}\Bigg[(A^*\sqrt{Q_2^2 - (2a^2u' - b_2)} + 2a^2u' - b_2) \times \exp\left(A^*\left[F_2 + \arcsin\left(\frac{2a^2u' - b_2}{Q_2}\right) - \arcsin\left(\frac{2a^2g - b_2}{Q_2}\right)\right]\right) + (b_2 - b_1)e^{A^*F_2} - (A^*\sqrt{Q_1^2 - b_1^2} - b_1)\Bigg], \quad \text{(G.28)}$$

and for the non-combusting case, on the range $0 \leq u' \leq 1$:

$$y = Pe^{B}\Bigg[(A^*\sqrt{Q^2 - (2a^2u' - b)} + 2a^2u' - b) \times \exp\left(A^*\arcsin\left(\frac{2a^2u' - b}{Q} + \arcsin\left(\frac{b_1}{Q_1}\right)\right)\right) - (A^*\sqrt{Q^2 - b^2} - b)\Bigg], \quad \text{(G.29)}$$

where $B = B_0 = -3.462$, and P is defined as

$$\frac{\mu_w M_e^2(\gamma_e - 1)}{\rho_e U_e a^2[4K^2 + c_f M_e^2(\gamma_e - 1)]}. \quad \text{(G.30)}$$

These equations, combined with equations G.9-G.11 allow for the profiles of velocity, density, and temperature to be constructed analytically. The only other major equation requiring effort to derive is that of the variation in eddy viscosity μ_t.

To begin, recall the definition for τ in equation G.3. With the previous assumption that shear stress is dominated by turbulent transport of momentum to the walls, the turbulent Reynolds stress R_{ij} can be assumed equivalent to $\tau = \mu_t(du/dy)$ away from

the wall, to render an equation of the form

$$\mu_t = \rho U_e K^2 \frac{\left(\frac{dy}{du'}\right)^3}{\left(\frac{d^2y}{du'^2}\right)^2} \tag{G.31}$$

A solution to this equation allows for a direct calculation of the turbulent viscosity at any point in the boundary layer. Substituting in the appropriate equation for dy/du', along with its derivative, a set of compact equations can be found for the eddy viscosity in combusting and non-combusting boundary layers with fuel injection. For the combusting cases, the final form of the equations for are presented below. For $0 \leq u' \leq g$:

$$\mu_t = \mu_w \exp\left[B_0 + A^*\Big(\arcsin\Big(\frac{2a^2u' - b_1}{Q_1}\Big) + \arcsin\Big(\frac{b_1}{Q_1}\Big)\Big)\right], \tag{G.32}$$

and for $g \leq u' \leq 1$:

$$\mu_t = \mu_w \exp\left[B_0 + A^*\Big(F_2 + \arcsin\Big(\frac{2a^2u' - b_2}{Q_2}\Big) - \arcsin\Big(\frac{2a^2g - b_2}{Q_2}\Big)\Big)\right]. \tag{G.33}$$

Similarly, for the case where fuel is injected but combustion is suppressed, the eddy viscosity across the range $0 \leq u' \leq 1$ is given by

$$\mu_t = \mu_w \exp\left[B + A^*\Big(\arcsin\Big(\frac{2a^2u' - b}{Q}\Big) + \arcsin\Big(\frac{b}{Q}\Big)\Big)\right]. \tag{G.34}$$

Since B and B_0 retain their previously defined values, the eddy viscosity will decay to a very small value as u' approaches 0, which is the desired behavior. Immediately above the wall surface is the viscous sublayer, inside which turbulent eddy structures are smaller than the Kolmogorov scale. These small-scale eddies rapidly dissipate within the viscous sublayer due to the viscous action of the fluid. As such, in the limit as $y \to 0$, μ_t must also go to zero. At the wall, the eddy viscosity given by equations G.32-G.34 is on the order of 3% of the laminar viscosity, which is an acceptably small value. With these equations defined, a thorough analysis of boundary layer combustion can be undertaken.

G.4 Results and Discussion of Analysis

The equations presented over the course of the previous section can provide a daunting amount of information about the scale, and mechanisms by which, combustion of

a general fuel in the boundary layer of a hypersonic flow reduces skin friction. Fortunately, a robust understanding of these phenomena and their causes can be gained from an examination of a few key cases. To facilitate comparison with the original theory for hydrogen presented in S2005, the same standard freestream flow condition is adopted. This condition is for the equivalent to a flight Mach number of 12.7, and has a flow stagnation enthalpy of 7.54 MJ/kg. The freestream velocity is 3380 m/s, with a temperature of 1500 K and density of 0.159 kg/m^3. These conditions correspond to the experimental data of Goyne (2000), and are in the mid-to-high range of flight conditions a scramjet may one day operate at.

An isothermal wall with a temperature of 300 K is assumed. This temperature is chosen in part for comparison with the original model, and also because it is similar to the wall temperature of models in most short-duration hypersonic experiments. Increased wall temperature is demonstrated clearly in S2005 to have little effect on the magnitude of skin friction reduction. As such, no further remarks on the effects of T_w are made here.

Hydrogen and Ethylene Performance Comparison

At this point, it is prudent to first test the model derived to ensure it behaves as expected. Ethylene (C_2H_4) is a popular propellant choice for experimental hydrocarbon-fuelled scramjets, and will be the principal fuel examined and contrasted with the original hydrogen combustion model. Some care must be taken in making a comparison with the data presented in S2005, where the hydrogen flow was defined in terms of its injection velocity and mass flow rate. A direct comparison of the same velocity and mass flow for ethylene leads to a wildly different injection condition.

A more appropriate comparison is made by matching fuel injection Mach number, and equivalence ratio. While matching the injection Mach number is straight-forward, care should be taken when speaking about equivalence ratios. Recall that this model is for an unbounded surface flow. As such, any oxygen mass flow rate (and therefore equivalence ratio) will have a completely arbitrary value. A simple way of ensuring fuel flow conditions match is to recognize that if the hydrogen and ethylene equivalence ratios can be assumed equal and have an abundance of oxygen to react with, the standard definition for equivalence ratio can be solved to find:

$$\frac{\dot{m}_{hydrogen}}{f_{hydrogen}} = \frac{\dot{m}_{ethylene}}{f_{ethylene}} \tag{G.35}$$

So, the hydrogen mass flow rate of 0.2 kg/s·m, and injection velocity of 1550 m/s given in S2005 yield an injection Mach number condition of 1.17. The same injection Mach number and equivalence ratio yields an ethylene injection velocity of 413.78 m/s and mass flow rate of 0.473 kg/s·m.

To determine how local values of c_f evolve in a boundary layer with fuel injection, the fuel mass flow rate is set to a constant value and is assumed to be injected along the wall at $x = 0$. Recognizing that Re_x is analogous to streamwise distance when freestream properties are held constant, setting $c_{Fw} = 1$ allows the values of Re_x and local c_f to be solved for at the point where equations G.23-G.26 become valid. Due to the boundary layer being assumed self-similar, marching the value of c_{Fw} from $1 \rightarrow 0$ will yield a corresponding set of values for c_f (using equation G.24 or G.26) and then Re_x (from equation G.23 or G.25) for the entire region downstream of the first solved point in which the model is valid. Comparing these skin friction coefficients to those for a boundary layer with no fuel injection at the same Re_x yields skin friction reduction curves for combusting and non-combusting cases, such as those presented in Figure G.5 for the conditions defined previously.

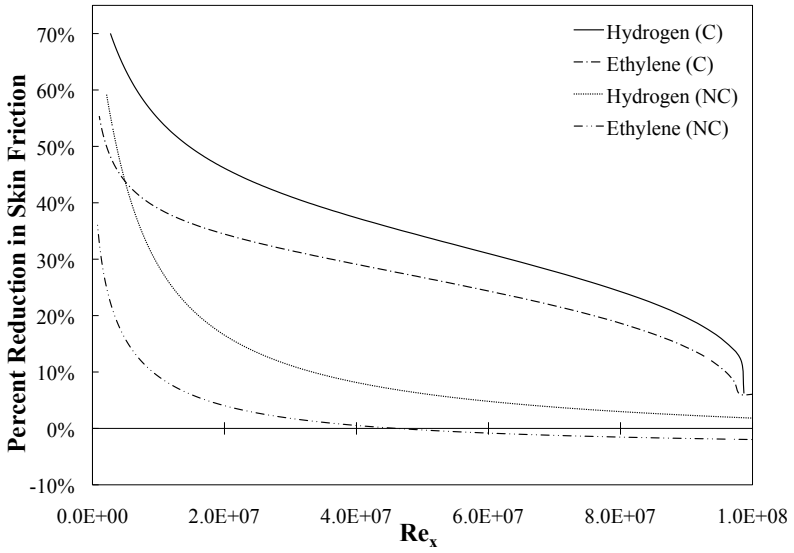

Figure G.5: Reduction in local c_f for fuel injection, with (C) and without (NC) combustion

Immediately noticeable is the fact that ethylene provides a reduction in local c_f when injected into the boundary layer. Hydrogen outperforms ethylene in both combusting and non-combusting cases. That hydrogen outperforms ethylene is not an especially surprising result: hydrogen has both lower viscosity and a higher heat of combustion. It may be noted at this point that when the properties for hydrogen are in-

serted into equations G.23-G.26, they take on the exact form presented in S2005, and as S2005 presented, will still provide good agreement with available experimental data.

Figure G.5 also shows that film cooling effects are strong immediately downstream of the injection point, but the non-combusting curves show that the effect drops off rapidly as the fuel diffuses away from the wall further downstream. It should be remarked that if the ethylene does not ignite, film cooling flow will eventually lead to an increase in local c_f. The value of c_f in boundary layers with fuel combustion is much less sensitive to the local value of c_{Fw}, and only decays rapidly when Re_x is very large and the fuel is almost depleted.

Interestingly, the effect of combustion appears to be stronger for ethylene than for hydrogen as $c_{Fw} \to 1$ near the leading edge of the plate. When $c_{Fw} = 1$, the reduction in c_f for combusting hydrogen is roughly 16% greater than the film cooling case, whereas for ethylene the combusting case gives a 50% increase in c_f reduction over film-cooling alone. This suggests that the dominance of the various underlying physical reasons for reduction in skin friction differ for ethylene and hydrogen at different points along the wall surface. What these mechanisms are, and which are dominant for each fuel, is the next point for discussion.

Understanding the Flow Physics

One popular understanding of how the injection and burning of fuel in a boundary layer reduces skin friction is as follows: the combustion reaction releases vast quantities of heat, raising the local temperature of the fluid next to the wall. This temperature rise in turn causes a reduction in density, and therefore a local dilatation of the fluid that increases the boundary layer thickness. A thicker boundary layer implies smaller velocity gradients next to the wall which, recalling that $\tau_w = \mu_w (du/dy)|_w$, means a smaller wall shear stress and therefore lowered skin friction drag (Suraweera and Smart, 2009). The validity of this explanation is challenged by the boundary layer velocity profiles generated using equations G.27-G.29, normalized to the boundary layer thickness of a boundary layer without fuel injection at the same Reynolds number, as plotted in Figure G.6.

Examining the normalized velocity profiles over $0 \leq y \leq \delta$, it is readily apparent that when fuel is injected into the boundary layer δ increases significantly. The magnitude of the increase appears to be strongly tied to the momentum of the injected fuel. The heavier, slower-moving ethylene fuel stream causes a 66% increase in δ compared with the 4% increase in the equivalent hydrogen injection case. While Figure G.6 does

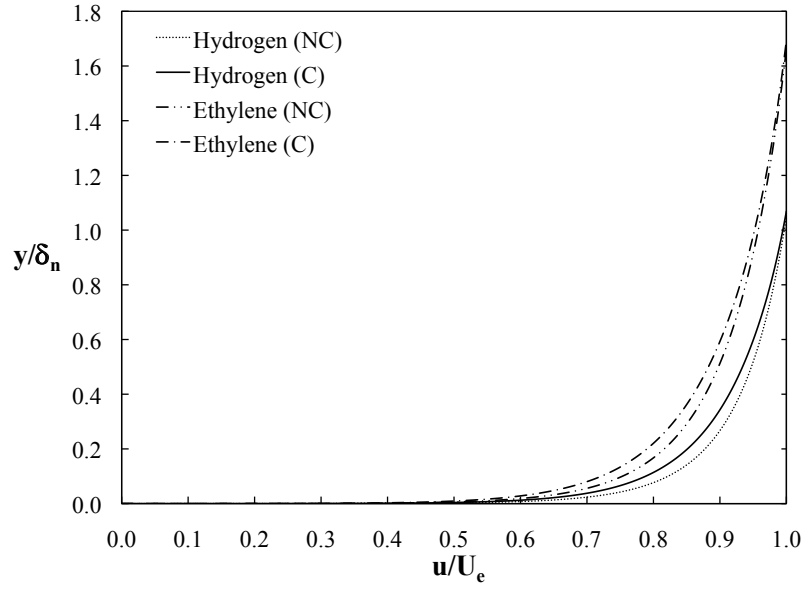

Figure G.6: Velocity profiles with fuel injection at $c_{Fw} = 0.5$, with and without combustion

show some slight increase in boundary layer thickness for both fuel injection cases when combustion occurs, this increase is modest compared with the increase due to fuel injection alone and does not scale with the observed increases in reduction of c_f.

While there may be some merit to the argument that an increase in boundary layer thickness lowers the streamwise velocity gradient at the wall, the increase in δ alone is not indicative of reductions in c_f when combustion occurs; combustion effects must be predominantly driven by other physical phenomena inside the boundary layer. This assertion is further supported by the change in the "fullness" of the combusting velocity profiles. The velocity profiles for the combusting cases are less "full" than that of the non-combusting and non-injection cases. A fuller velocity profile implies a greater level of turbulent momentum transport from the freestream to the wall. Thus, it can be inferred that combustion acts, at least in part, to reduce the rate of turbulent transport across the boundary layer. How combustion does this is examined in a later section.

The behavior of the other classical component of τ, the total flow viscosity, is found by adding the local laminar viscosity to its complementary eddy viscosity. This total viscosity in the near-wall region is plotted in Figure G.7. As previously demonstrated through physical reasoning and equations G.32-G.34, the total viscosity is dominated by the value of the laminar viscosity at the wall. This leads to the asymptotic behavior observed as $y/\delta \to 0$. The injection of fuel into the boundary layer lowers the value of the total viscosity at the wall, with an additional reduction in combusting cases due the addition of combustion products (and lack of oxygen) in the gas mixture beneath

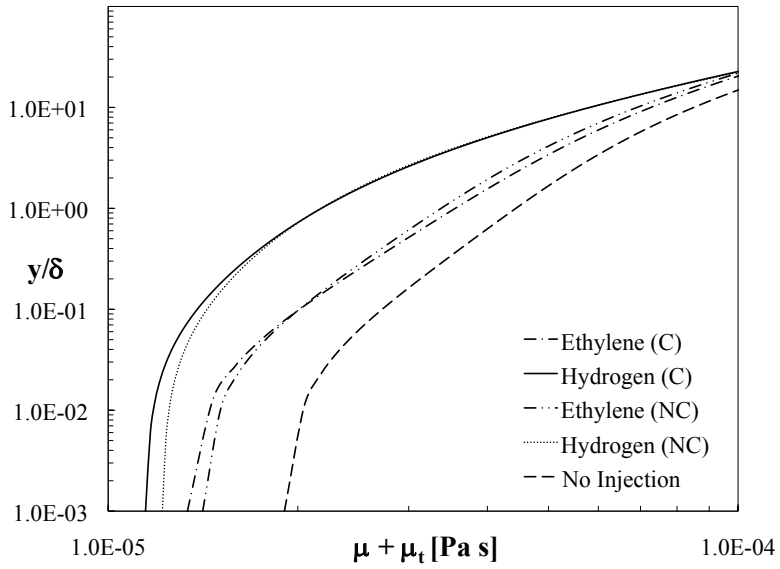

Figure G.7: Near-wall total viscosity profiles for $c_{Fw} = 0.5$, with and without combustion

the flamefront. The reduction observed is greater for hydrogen than for ethylene, as expected: with a viscosity half that of the ethylene mixture, in pure film cooling mode hydrogen injection will cause greater reduction in c_f, despite the much larger increase in δ that implies a larger reduction in $[du/dy]_w$ for ethylene injection. It is worth recalling at this point that the model assumed an isothermal wall boundary condition. Thus, regardless of the magnitude of heat released through combustion, there will be no temperature-driven increase in viscosity at the wall in combusting cases.

It can be remarked that, due to the analytical requirement that dy/du' be continuous at g, the boundary layer velocity profile in Figure G.6 shows no evidence of the existence of a flamefront in the combusting cases. The flamefront is, however, immediately visible when examining the boundary layer profiles for temperature and density shown in Figure G.8. These plots, taken for identical local values of Re_x, illustrate further the difference in dominant mechanisms for reduction in c_f for hydrogen and ethylene fuels. Despite the higher heat release of combustion for hydrogen, the overall maximum temperature in the boundary layer is somewhat higher for ethylene. The flamefront is also closer to the wall for ethylene combustion compared with hydrogen at any value of Re_x. Also noteworthy is the fact that for an equivalent injection condition, the consumption of ethylene will proceed at a different rate than the hydrogen, as evidenced by the difference in c_{Fw} between hydrogen and ethylene.

Considering the link between temperature and density, the behaviour of the density profiles of the combusting cases (Figures G.8a and G.8b) are strikingly dissimilar from the temperature profiles. While the density is very high next to the wall due to

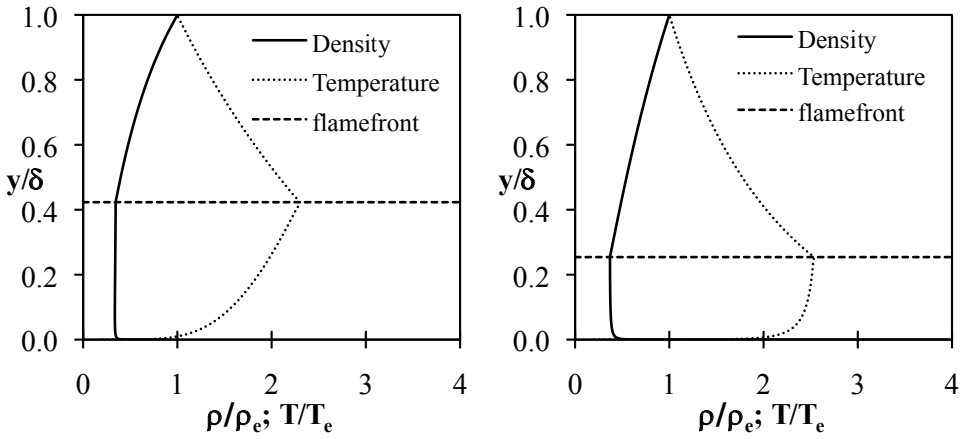

(a) Hydrogen, c_{Fw}=0.34 **(b)** Ethylene, c_{Fw}=0.43

Figure G.8: Combusting boundary layer density and temperature profiles at $Re_x \approx 10^7$

the low value of T_w, it rapidly drops to a nearly constant value that is a fraction of the freestream density. This drop is certainly linked to the heat release from combustion, but seems insensitive to the temperature variation beneath the flamefront. Given that the pressure is constant across the boundary layer, this seeming lack of coupling between ρ and T must be due to the changes in mass fraction of the constituent gas species, and hence the mixture gas constant, between the wall and the flamefront.

The behavior of the mean density is important due to its strong coupling with the boundary layer Reynolds (turbulent) stress,

$$R_{ij} = \overline{\rho u_i' u_j'} \tag{G.36}$$

which drives turbulent momentum transport across the boundary layer. As density decreases, the overall rate of momentum transfer also decreases, leading to a less "full" velocity profile and a reduction in c_f. By assuming that the Reynolds stress can be represented by the product of the turbulent viscosity calculated with equations G.32-G.34 as appropriate, and velocity gradients within the boundary layer, it can be shown that

$$\overline{u_i' u_j'} = \frac{\mu_t}{\rho} \frac{du}{dy}_w \tag{G.37}$$

The resulting curves are contrasted against the density profiles in Figure G.9 for non-combusting injection of fuel, and in a later figure for combusting hydrogen.

When hydrogen is injected, Figure G.9a shows a large deviation in the density profile from the non-injecting case. This is expected, given gaseous hydrogen's low den-

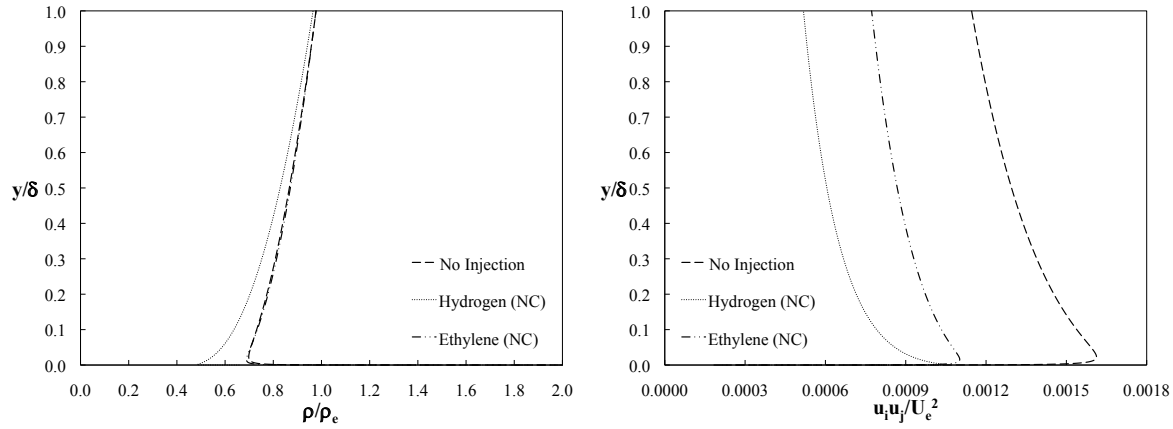

(a) Density **(b)** Mean Turbulent Fluctuations

Figure G.9: Density and velocity fluctuation profiles for non-combusting boundary layers at $c_{Fw} = 0.5$

sity, whereas the density of ethylene is nearly identical to that of dry air. This does much to explain the film-cooling performance of hydrogen: its low density will have a strong influence on the Reynolds stress profile.

There is, however, another effect obviously present: Figure G.9b clearly demonstrates that ethylene injection reduces the magnitude of velocity fluctuations across the boundary layer, despite its mean boundary layer density profile being virtually identical to that of air. While ethylene's density is similar to air, its viscosity is approximately 50% lower, as previously demonstrated in Figure G.7. Low viscosity is also a feature of hydrogen, which together with its previously discussed effect on density explains why hydrogen film-cooling is much more effective than ethylene.

While the lower viscosity of the fluid is one source of skin friction reduction for ethylene film-cooling, the significant increase in boundary layer thickness for ethylene injection is indicative of the more dominant effect: the injection of a low-velocity, low-momentum fluid next to the wall causes a reduction in $[du/dy]_w$. This can be easily discerned by modelling a case where the ethylene viscosity is artificially set as equal to the viscosity for air, and is included as Figure G.10.

The exercise illustrated in Figure G.10 serves to isolate the effects of viscosity. With ethylene's viscosity changed, but other properties remaining constant, the effects of the low-momentum injection of fuel are kept constant (including constant boundary layer thickness) while the viscosity effects undergo a dramatic change. Closest to the point of injection, where $c_{Fw} = 1$, there is only a 15% smaller reduction in c_f due to injection effects alone. As more and more of this fuel diffuses away from the wall,

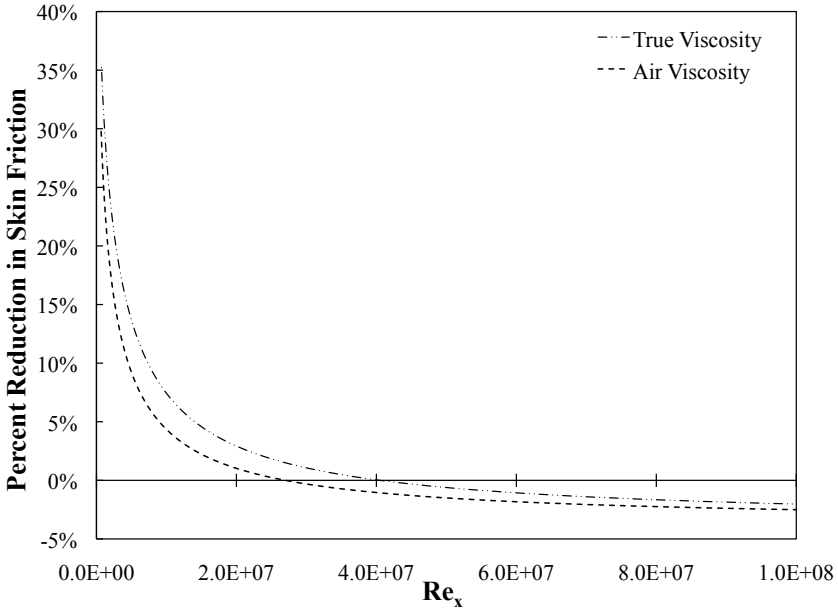

Figure G.10: Effect of fuel viscosity on skin friction reduction in an ethylene film-cooled boundary layer

the difference in between the true and artificial viscosity cases asymptotically drops to zero. It can thus be concluded that the dominant effect for ethylene film cooling is, by far, the reduction in $[du/dy]_w$ brought about by the low-momentum injection of fuel next to the wall.

This effect is also certainly present in a hydrogen-film-cooled boundary layer. However, at the same stoichiometry, a hydrogen fuel stream would have lower mass flow rate than that of an ethylene stream. This results in a much smaller increase in δ, and a greater dependence on reduced viscosity in decreasing skin friction. As hydrogen has a significantly lower viscosity than ethylene, it still provides a superior c_f reduction overall.

Figure G.11 clearly demonstrates that combustion reduces the overall shear stress across the boundary layer, as the factor by which the density is reduced beneath the flamefront is greater than the intensification factor of the velocity fluctuations. The heat release due to combustion affects the physical phenomena responsible for the transport of momentum from the freestream to the wall. While it is plainly visible that turbulence is intensified at the flame front, the turbulent fluctuations near the wall actually decrease relative to the non-combusting injection case, thus further driving down momentum transport to the wall. This agrees with the conclusions of Denman (2007), who found that there are reductions in streamwise turbulent velocity fluctuations and near-wall density due to combustion. These mechanisms combine to reduce

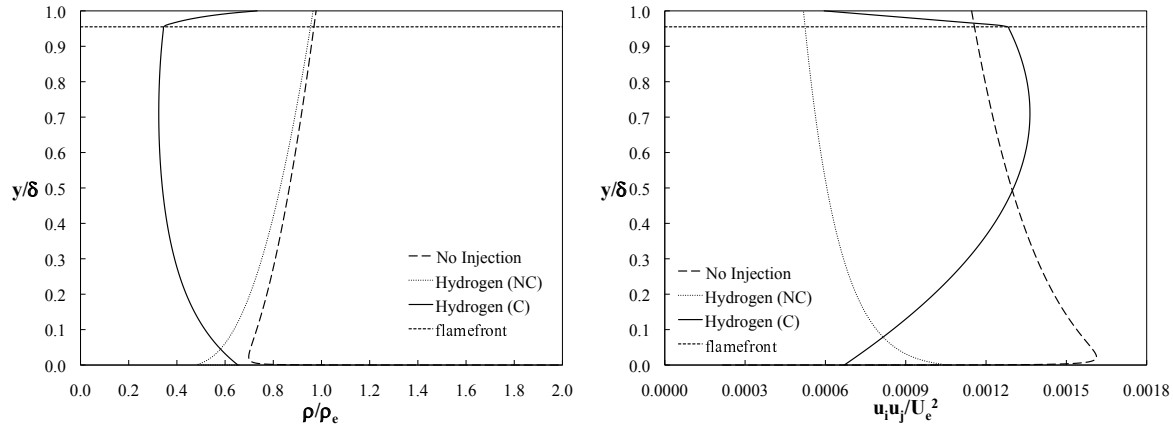

(a) Density **(b)** Mean Turbulent Fluctuations

Figure G.11: Density and velocity fluctuation profiles for hydrogen combustion at $c_{Fw} = 0.5$

shear stress at the wall, and therefore skin friction. It should also be noted that while the results presented in Figure G.11 are for hydrogen injection, boundary layer combustion of ethylene exhibits the same behaviour.

Overall, there are several linked mechanisms that work to reduce skin friction when fuel is injected into a compressible turbulent boundary layer. First, injection will affect the total viscosity at the wall, to which τ_w is directly proportional. Second, the density of the injected fuel can have a strong influence on the Reynolds stresses, and therefore on c_f. The low speed and low momentum of the injected fuel relative to the freestream also has a strong effect on c_f through the wall-normal velocity gradient $[du/dy]_w$ (and also on boundary layer thickness). This appears to be especially strong in the case of ethylene injection. Finally, if the injected fuel burns, a deep layer of hot, low-density fluid beneath the flamefront interferes with the transmission of freestream momentum to the wall, reducing the Reynolds stresses despite the amplification of turbulent velocity fluctuations across much of the boundary layer. The Reynolds stress is the dominant source of shear stress in, and momentum transport across, the boundary layer. While all three appear to have a common link to flow temperature and density across the boundary layer, the combustion mechanism has the strongest influence on skin friction.

Model Sensitivity to Assumed Constant Parameters

Having examined the flow physics predicted by the analytical model, it is worth evaluating the impact of some of the assumptions made in the model. Some assumptions,

such as $Le = Pr = 1$, have been previously shown to have a sound basis in experiment. Others, like assuming a diffusion flame model or a constant R^* (equivalent to a constant γ), were necessary in order to obtain a tractable analytical solution. Many of these assumptions will introduce some deviation from the real processes and effects of boundary layer combustion on the magnitude of local skin friction. How much of a deviation such assumptions introduce is generally difficult to quantify. Despite all these potential errors, S2005 demonstrated good agreement with the experiments of Goyne et al. (2000) when applying the model to hydrogen combustion. As the model presented in this paper collapses to a form identical to S2005 when properties for hydrogen are used, the new model can be said to show the same agreement. Taking these assumptions as a whole, while not necessarily an accurate representation of the true processes inside a combusting turbulent boundary layer, at least reasonably quantify the effects of these processes on c_f.

One assumption that can be directly checked for its effect is the assumption of constant C_p for each species. These values of C_{pi} are normalized to C_p for air for ease of calculation. The assumed values for C_{pi} were originally taken for each gas at a temperature between that of the freestream and the wall. However, the temperature in the boundary layer can be two to three times higher than the freestream when combustion is present. As such, setting C_p at a higher average temperature may be desirable. Using temperature of 1500 K (as compared with ~700K for the original C_{pi} in S2005) for evaluating C_{pi} causes the predicted values for local c_f reduction to change in the manner illustrated in Figure G.12.

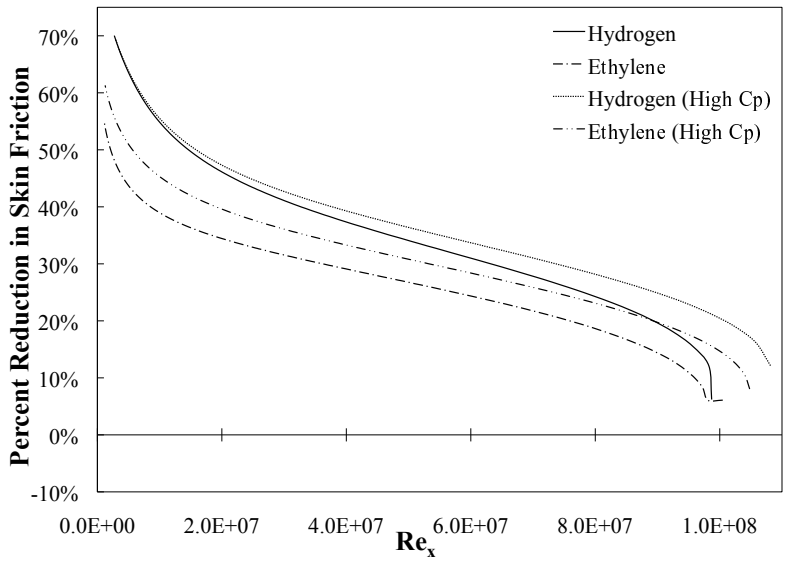

Figure G.12: Changes in reduction to local c_f with a change in assumed values for C_{pi}

Compared with the original curves for combusting skin friction reduction replotted from Figure G.5, constant values for C_p taken at a higher temperature yield a small increase in skin friction reduction that lasts further downstream from the point of injection. Thus, a constant C_p calculated at a lower average temperature gives a more conservative estimate of skin friction reduction.

Application to a Scramjet Engine

In previous sections, the analysis has been focused on local reductions in skin friction. However, the curves from Figure G.5 can be integrated over the range of Re_x to obtain a net reduction in skin friction across any distance downstream of the injection location. In doing so, it is possible to obtain a first estimate of the optimum boundary layer fuel mass flow rate for a scramjet combustor.

Having assumed a hypersonic vehicle flying at Mach 8 conditions, with a 6 degree forebody, and a 1m long combustor fed by a 60 mm diameter circular duct, the total reduction in internal skin friction drag versus boundary layer equivalence ratio takes the form shown in Figure G.13. This calculation further assumes that the boundary layer remains small relative to the height of the combustor, and there is no mainstream fuel injection. Thus, corrections to local mainstream stagnation enthalpy to account for pressure rise due to mainstream combustion are not necessary, and the analysis is greatly simplified.

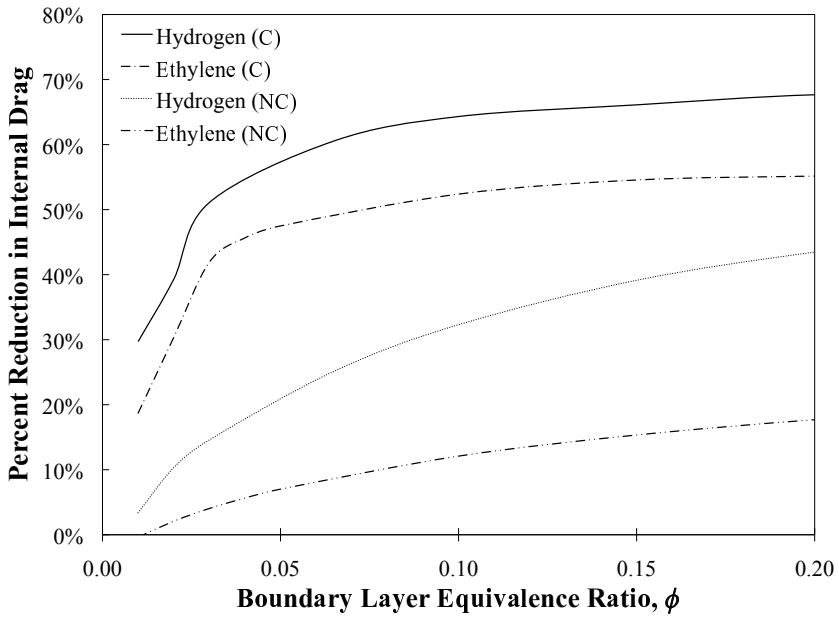

Figure G.13: Total skin friction drag reduction versus boundary layer fuel equivalence ratio

As Figure G.13 demonstrates, above $\phi \approx 0.1$, there is little benefit to further increases in fuel mass flow rate in the boundary layer; the remaining 90% of available fuel (assuming a global equivalence ratio of unity) should be injected into the mainstream flow of the engine. For higher equivalence ratios, significant amounts of fuel would be unburned at the combustor outlet, and the total reduction in c_f levels off. It is worth repeating that it is critical that ethylene burn if it is injected into a boundary layer. The net effect from ethylene film-cooling is small enough that, if no boundary layer combustion occurs, there would be more benefit from improving engine mainstream mixing and combustion to generate extra thrust. Overall, the result is promising: devoting a small fraction of available fuel to boundary layer combustion inside a scramjet may give a large overall reduction in skin friction drag. Recall that, as previously stated, the combustor skin friction drag may account for as much as 25% of total engine drag (Tanimizu, 2008). Thus, combined with boundary layer combustion's contribution to thrust generation, and the strong effect tangential injection can have on engine specific impulse, this large reduction in this drag could yield a significant increase in engine net thrust.

G.5 Conclusions and Recommendations

An analytical model for the injection and combustion of hydrogen in a hypersonic turbulent boundary layer has been extended to apply to a general fuelling condition. An analysis of the extended model shows that ethylene boundary layer injection is a viable candidate for skin friction reduction provided that the fuel ignites; film-cooling ethylene injection can lead to an increase in skin friction at downstream locations. The analysis is sensitive to the choice in constant pressure specific heats for each species. Choosing a low-temperature value for constant C_p is a conservative choice, as higher constant values of C_p predict larger, longer lasting skin friction reduction for the same flow conditions. The results generated with this model are consistent with available literature, demonstrating the utility of this model for predicting skin friction in hypersonic turbulent boundary layers.

An examination of the underlying flow physics predicted by the model suggests that several coupled mechanisms are responsible for the reduction of skin friction: a local reduction in wall mixture viscosity, a reduction in flow density and the injection of low-speed, low-momentum fuel decreasing the near-wall velocity gradient, and heat release due to combustion inside the boundary layer. All these effects are coupled with the variation of primitive variables across the boundary layer. Heat release due to combustion is the strongest mechanism for reducing skin friction. This heat release

affects the density profile and turbulent velocity fluctuations such that the turbulent transport of momentum from the freestream to the wall is decreased. This in turn acts to reduce the wall skin friction drag.

An application of the model to a typical scramjet combustor at Mach 8 flight conditions shows that skin friction reduction does not require a prohibitively large fraction of carried fuel to be injected into the boundary layer. At most, the global stoichiometric fuel fraction contained the boundary layer should be limited to 0.1. Combined with the other results, it can be inferred that many scramjet designs could benefit from inclusion of boundary layer fuel injection and combustion.

Finally, it should be noted that this study remains uncompared with experimental data for boundary layer combustion for ethylene. While the results for hydrogen injection have previously been validated, the authors are not aware of any currently available literature that can validate the results for ethylene. Undertaking such a study would have great value not only in validating this more general model, but would undoubtedly give further insight into the physical phenomena that reduce skin friction during boundary layer combustion.

Appendix H

Time Histories for Pressure Transducers in Shot 11471

This appendix contains the time histories for experimental shot 11471, from the experiments of Wise (2014). These traces were used in the discussion of the comparison between experimental and CFD pressure distributions in the baseline engine case, discussed in Chapter 5. This appendix begins with the normalized pressure distributions from Shot 11471 and Shot 11491, with select data points labelled to assist with the interpretation of the time history plots presented later on. This normalized distribution is presented as Figure H.1.

In these two shots, the engine was fuelled at similar equivalence ratios: the engine had an equivalence ratio of 1.33 in Shot 11471, and 1.24 in Shot 11491. Despite this, each exhibits remarkably different agreement with their respective equivalent CFD simulations. The natural question is, why do the experimental data in Shot 11471 appear to show such different combustion pressure behaviour in the combustor relative to the simulation intended to match it?

It is believed that the behaviour of shot 11471's experimental pressure distribution was caused by long-time-scale transient effects, that were a result of both the timing of fuel injection, and the amount of fuel injected. It is a common practice in experiments in the T4 Stalker Tube for fuel injection to commence prior to the arrival of the nozzle start-up flow at the engine. However, it has also been found that REST engine performance can be sensitive to the amount of time that the fuel is flowing prior to engine startup, and at what mass flow rate (Doherty, 2013).

Consider, for example, the time histories of transducers IB7 and IB11 shown in Figure H.2. The time histories of pressure at each of these measurement locations are compared for the two shots. The time histories for each pressure transducer are also

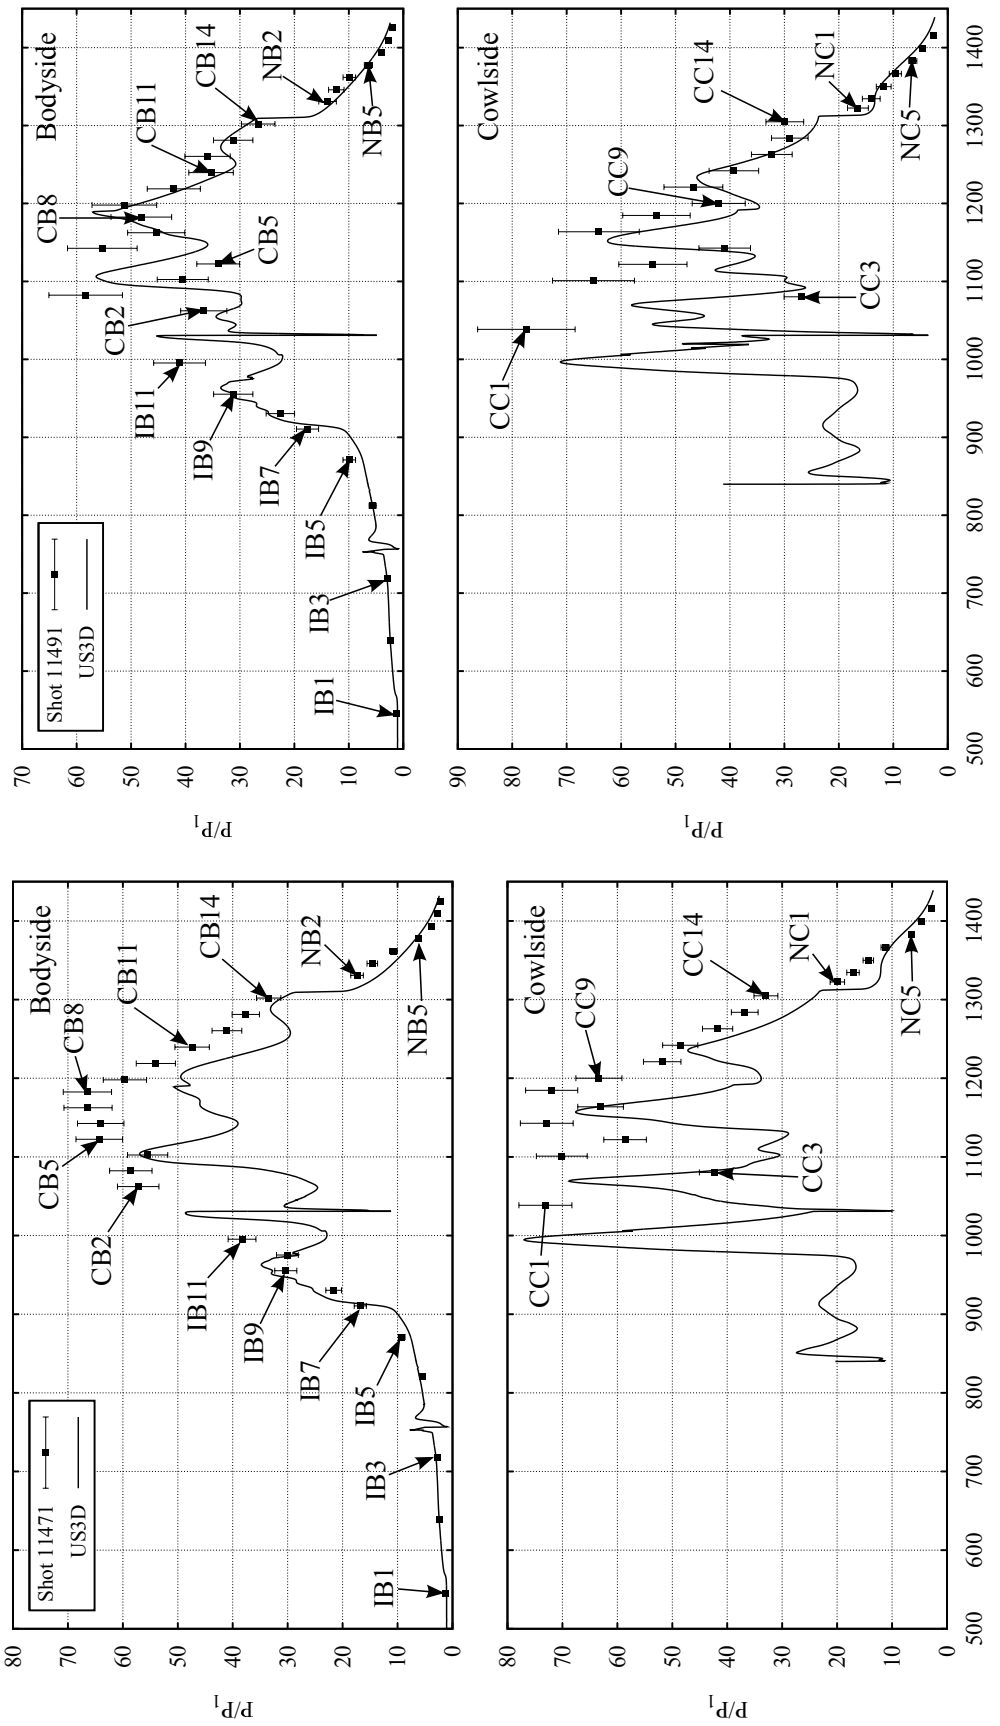

(a) Shot 11471

(b) Shot 11491

Figure H.1: Differences in the Normalized M12 REST Pressure Distributions for Shots 11471 and 11491

presented in their normalized form. This normalization accounts for the shot flow tailoring effects, allowing for a steady pressure signal to be extracted. In each individual pressure trace, the two vertical red lines represent the start and finish of the test time. This test time is the interval over which the instantaneous pressures were averaged. The pressure traces for IB7 show at once some similarity, and some difference in the behaviour of the inlet flow fields during start-up and test time.

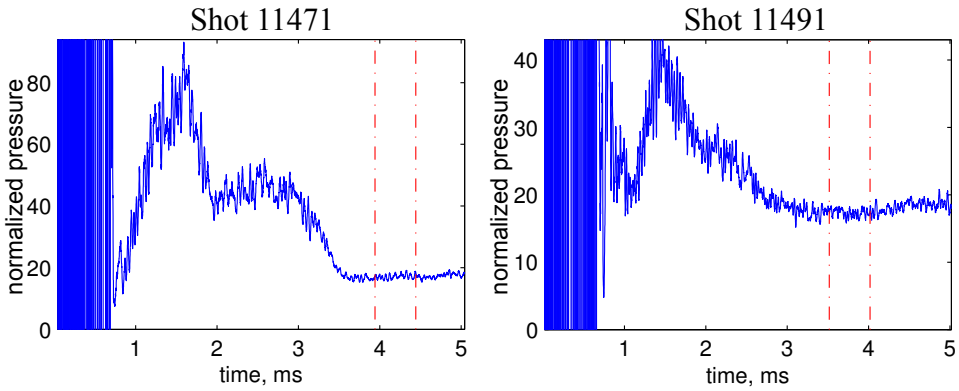

(a) IB7

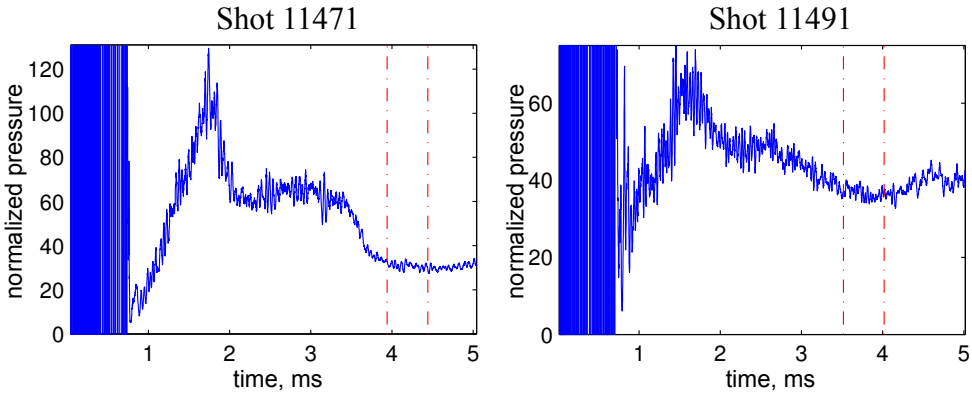

(b) IB11

Figure H.2: Differences Normalized Pressure Traces for Shots 11471 and 11491

As Figure H.2a shows, both shots show an initial spike in pressure between 1 and 2 ms after the start of the nozzle pressure rise, corresponding to the arrival of the nozzle startup shock. However, the initial pressure rise in Shot 11471 is nearly double that of Shot 11491, despite their nominal nozzle inflow conditions agreeing to within 2% (see Appendix E for both nozzle exit conditions). Downstream of this spike, the IB7 pressure trace for Shot 11491 drops to roughly half its initial peak magnitude, and plateaus, remaining nearly constant for 1 ms. At 3 ms, the pressure drops again by half, and then remains roughly constant at a normalized pressure of 20 times the nominal forebody condition. A similar plateau is visible in the trace of Shot 11491, but lasted

for nearly a quarter the time of the plateau in Shot 11471, before dropping to a similar final steady pressure.

Nearly identical behaviour is seen in Figure H.2b. The plateaus in IB11 last longer in both shots, but the plateau in Shot 11491 is once again several times shorter than that observed in Shot 11471. In both cases the pressures eventually settle to an apparently steady flow, though Shot 11491 does so sooner. There is also some evidence of pressure rise in the last millisecond of the recorded pressure, likely due to the now significant drop in nozzle supply pressure in this under-tailored nozzle flow condition.

This behaviour also extends into the engine combustor in shot 11471. However, as Figure H.3 shows, the behaviour in Shot 11491 behaves as is typically expected during an experiment: The pressure trace begins with a peak caused by the passage of the nozzle start-up shock, then drops briefly before recovering to a steady flow pressure determined by the shocks established within the combustor. The start-up peak pressure remains higher in Shot 11471 than it does in Shot 11491.

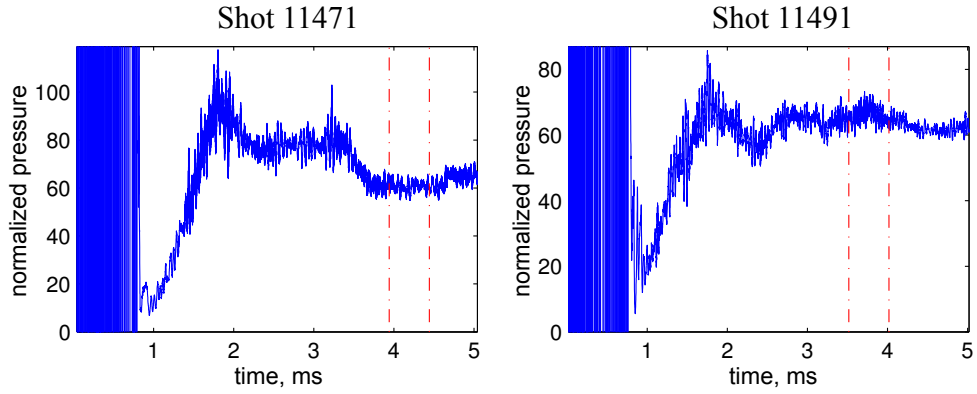

(a) CC7

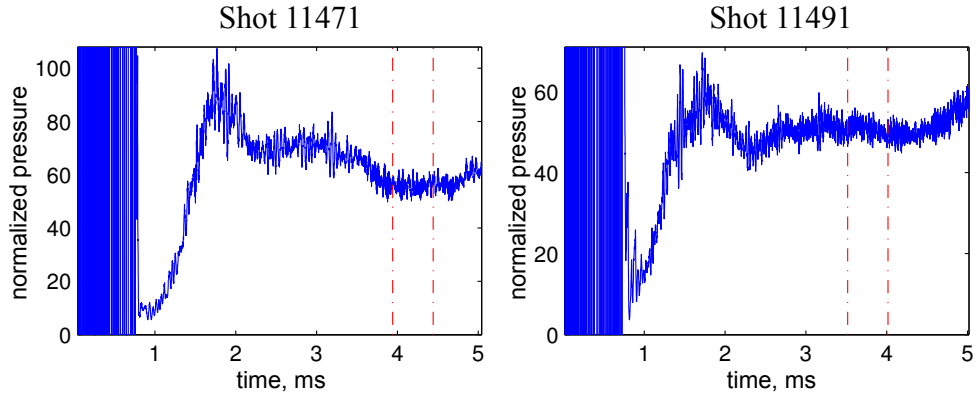

(b) CB9

Figure H.3: Differences Normalized Pressure Traces for Shots 11471 and 11491

This inter-shot difference in the behaviour of the pressures was common to both the

bodyside and cowlside of the combustor. It is believed that this plateau was caused by the early fuelling time used during the experiments. A large mass of fuel is present in the inlet and combustor prior to the arrival of flow from the nozzle. When the start-up flow arrives at the engine, the strong shock locally increases the temperature, and the fuel is able to mix and burn. This causes the much higher initial spikes in pressure seen in Shot 11471, as well as the sustained pressure plateau effect.

This behaviour is exacerbated by increases in the fuel mass flow rate, which explains why Shot 11491, with its lower equivalence ratio and smaller proportion of fuel injection into the inlet, shows a far shorter plateau effect. The larger mass of fuel takes longer to be pushed through the engine. It is also likely that the flow has either not completely reached its final operating condition (in which case the high pressure region inside the combustor would eventually be blown out through the nozzle), or else has reached an alternative steady state made possible by the slow, high pressure start-up caused by early fuelling.

It may therefore be concluded that the lack of agreement between simulation and experiment demonstrated here, and in Chapter 5, is the result of early fuelling in the experiment. While the inlet is able to return to the expected pressure distribution by the time Shot 11471 runs out of test flow, the combustor remains in a significantly different state. The behaviour of the nozzle, and every transducer in Shot 11471, are included as the end of this appendix.

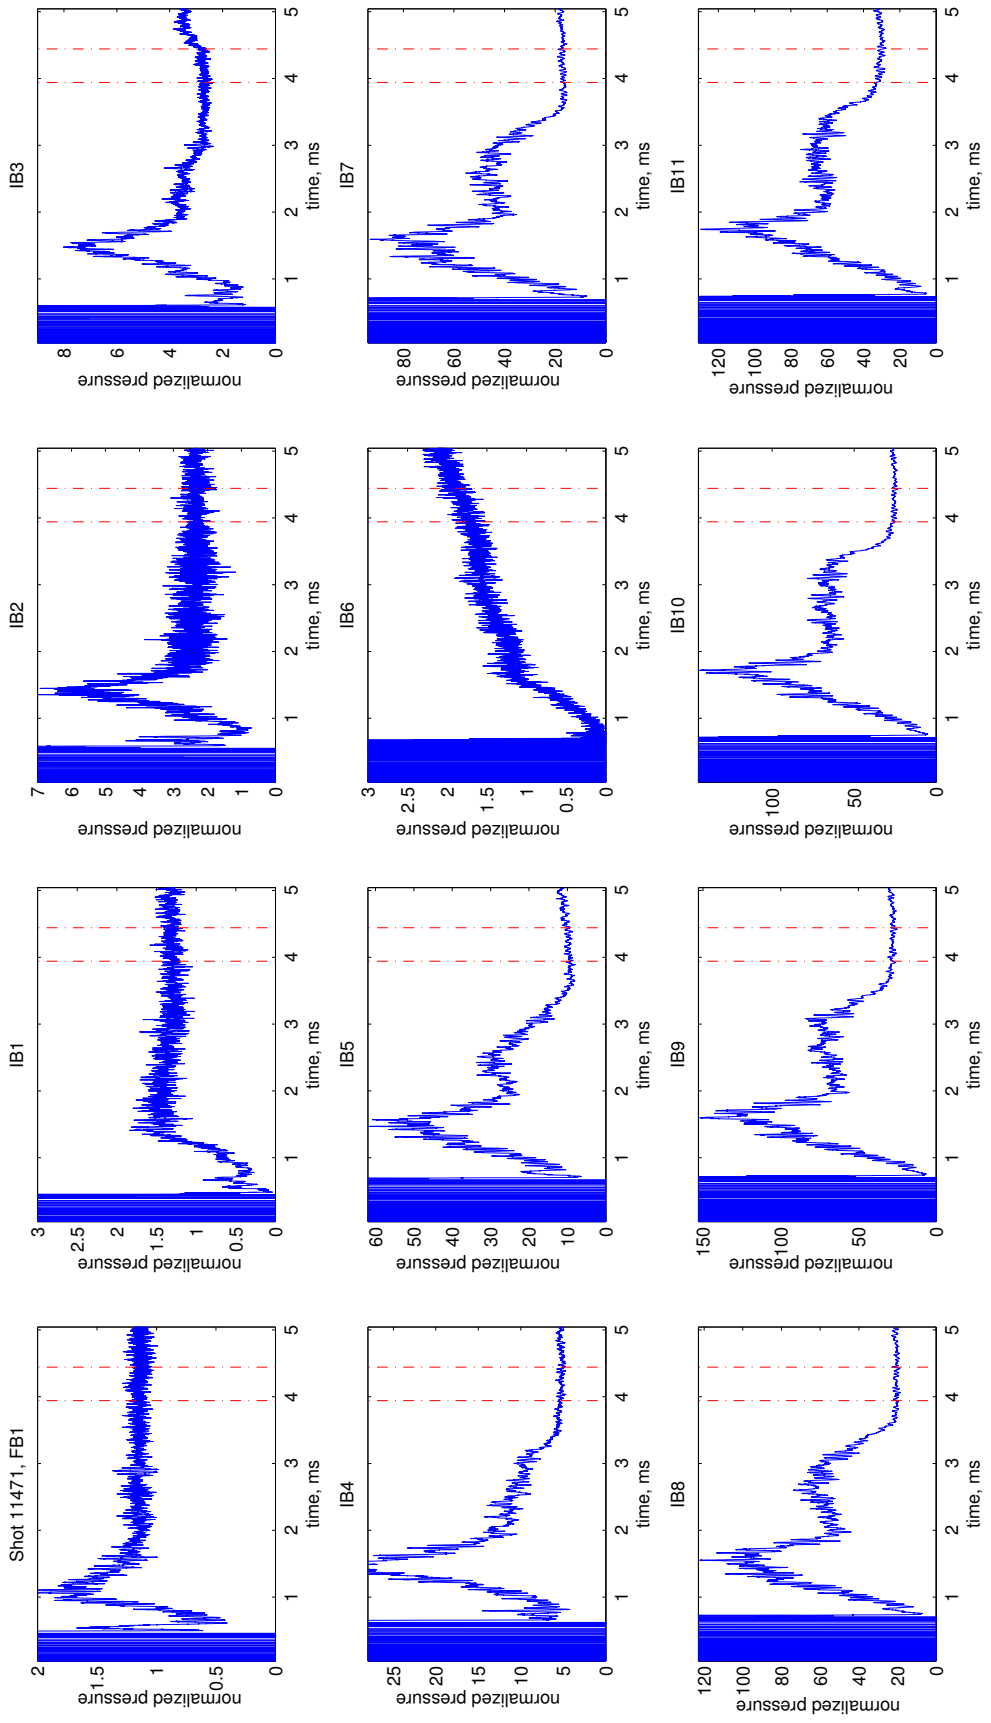

Figure H.4: Shot 11471 time histories of bodyside pressure transducers in the M12 REST inlet

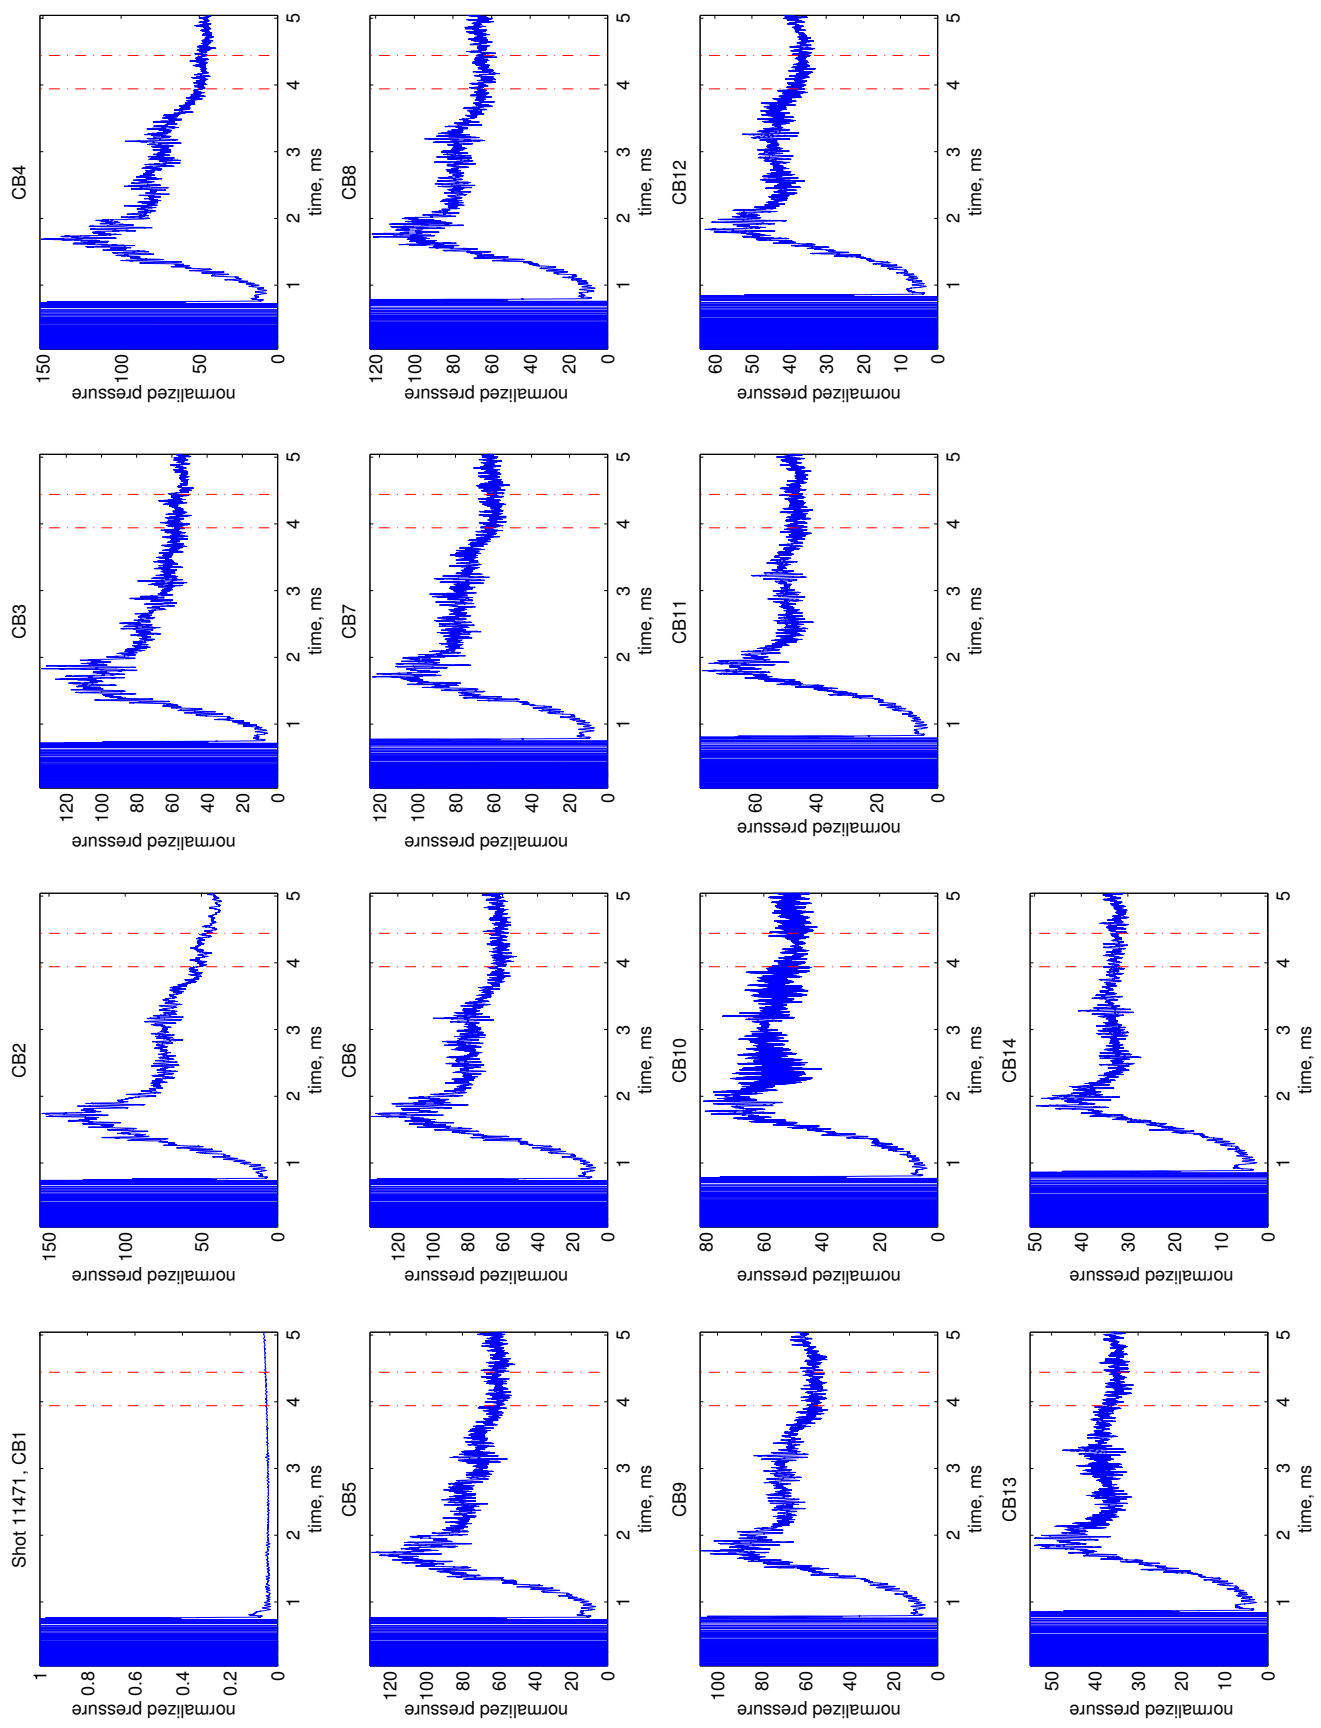

Figure H.5: Shot 11471 time histories of bodyside pressure transducers in the M12 REST combustor

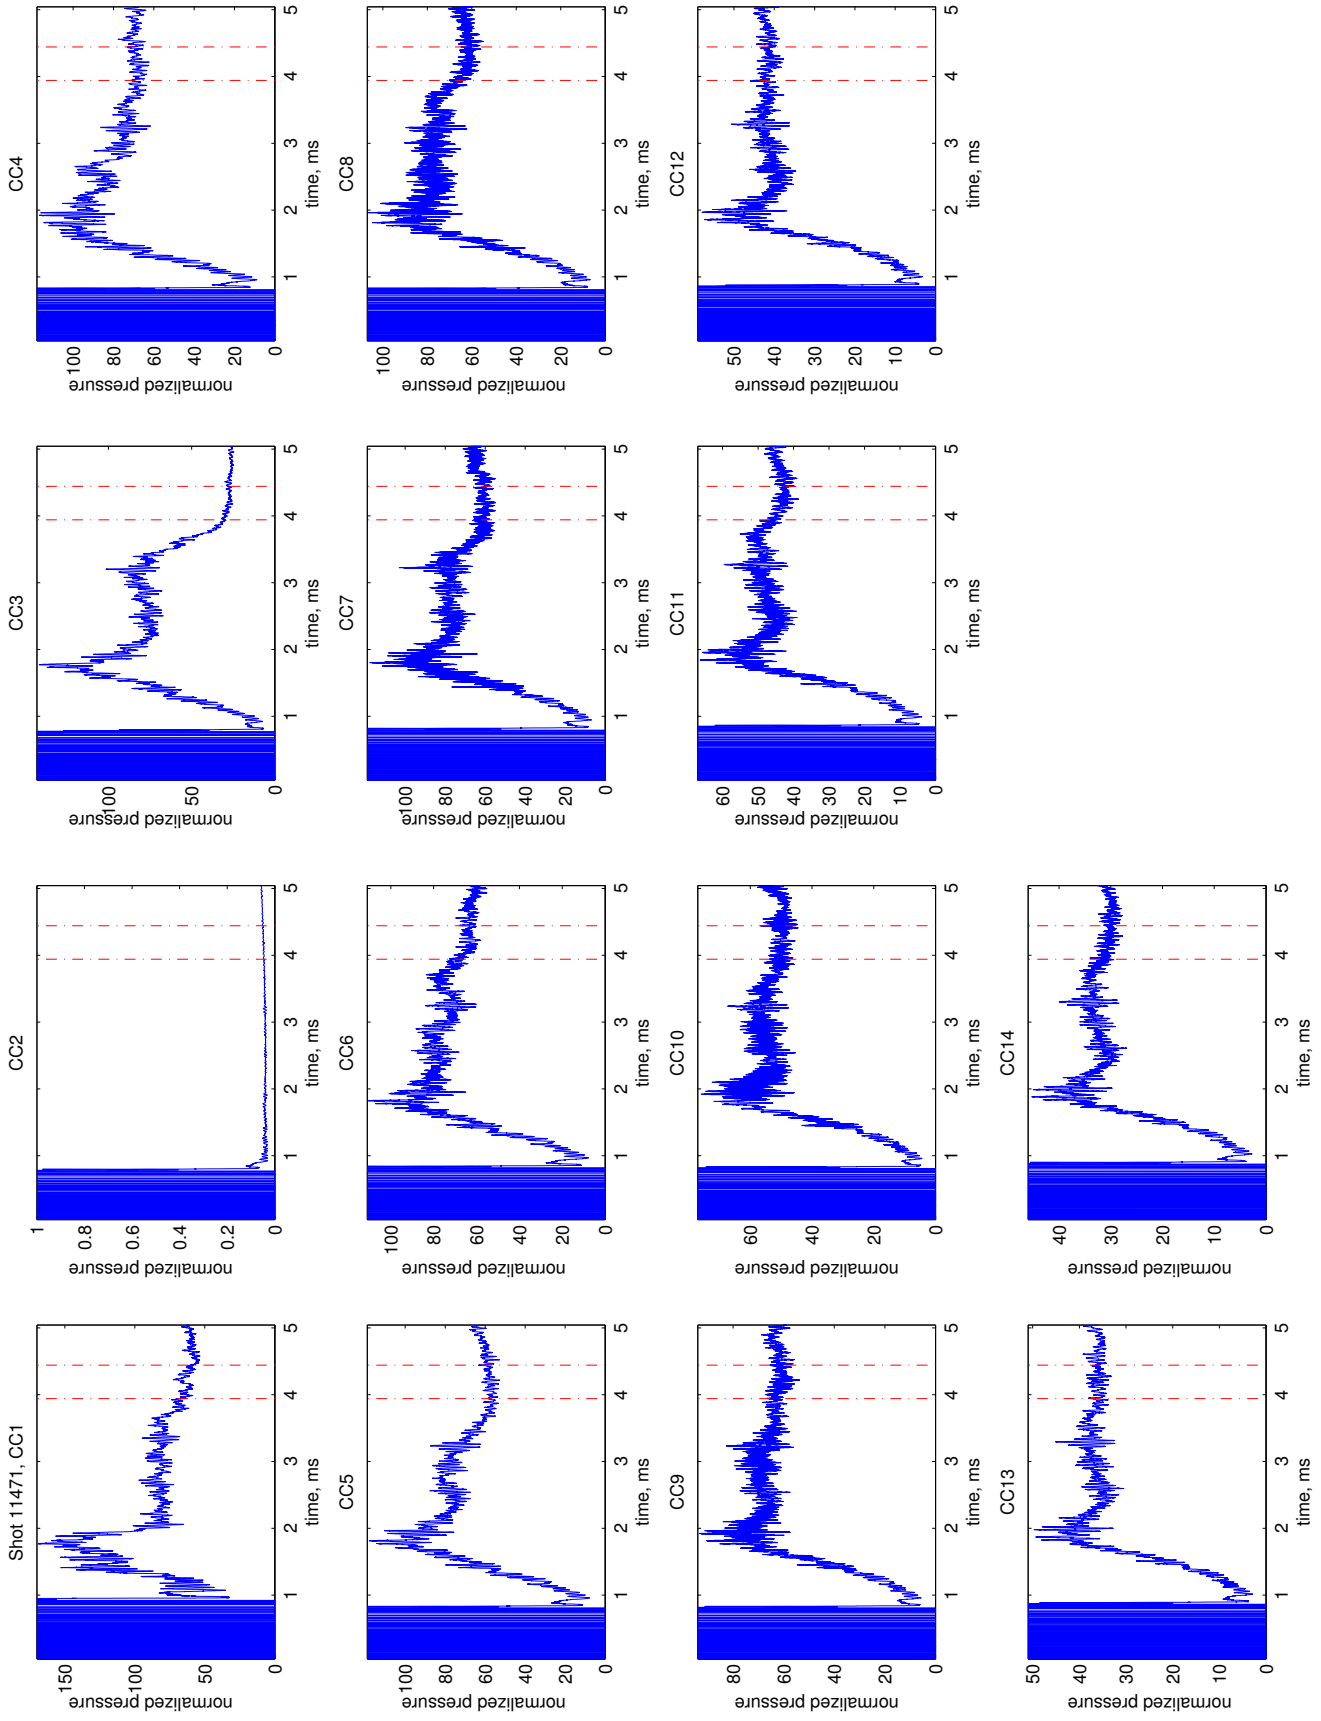

Figure H.6: Shot 11471 time histories of cowlside pressure transducers in the M12 REST combustor

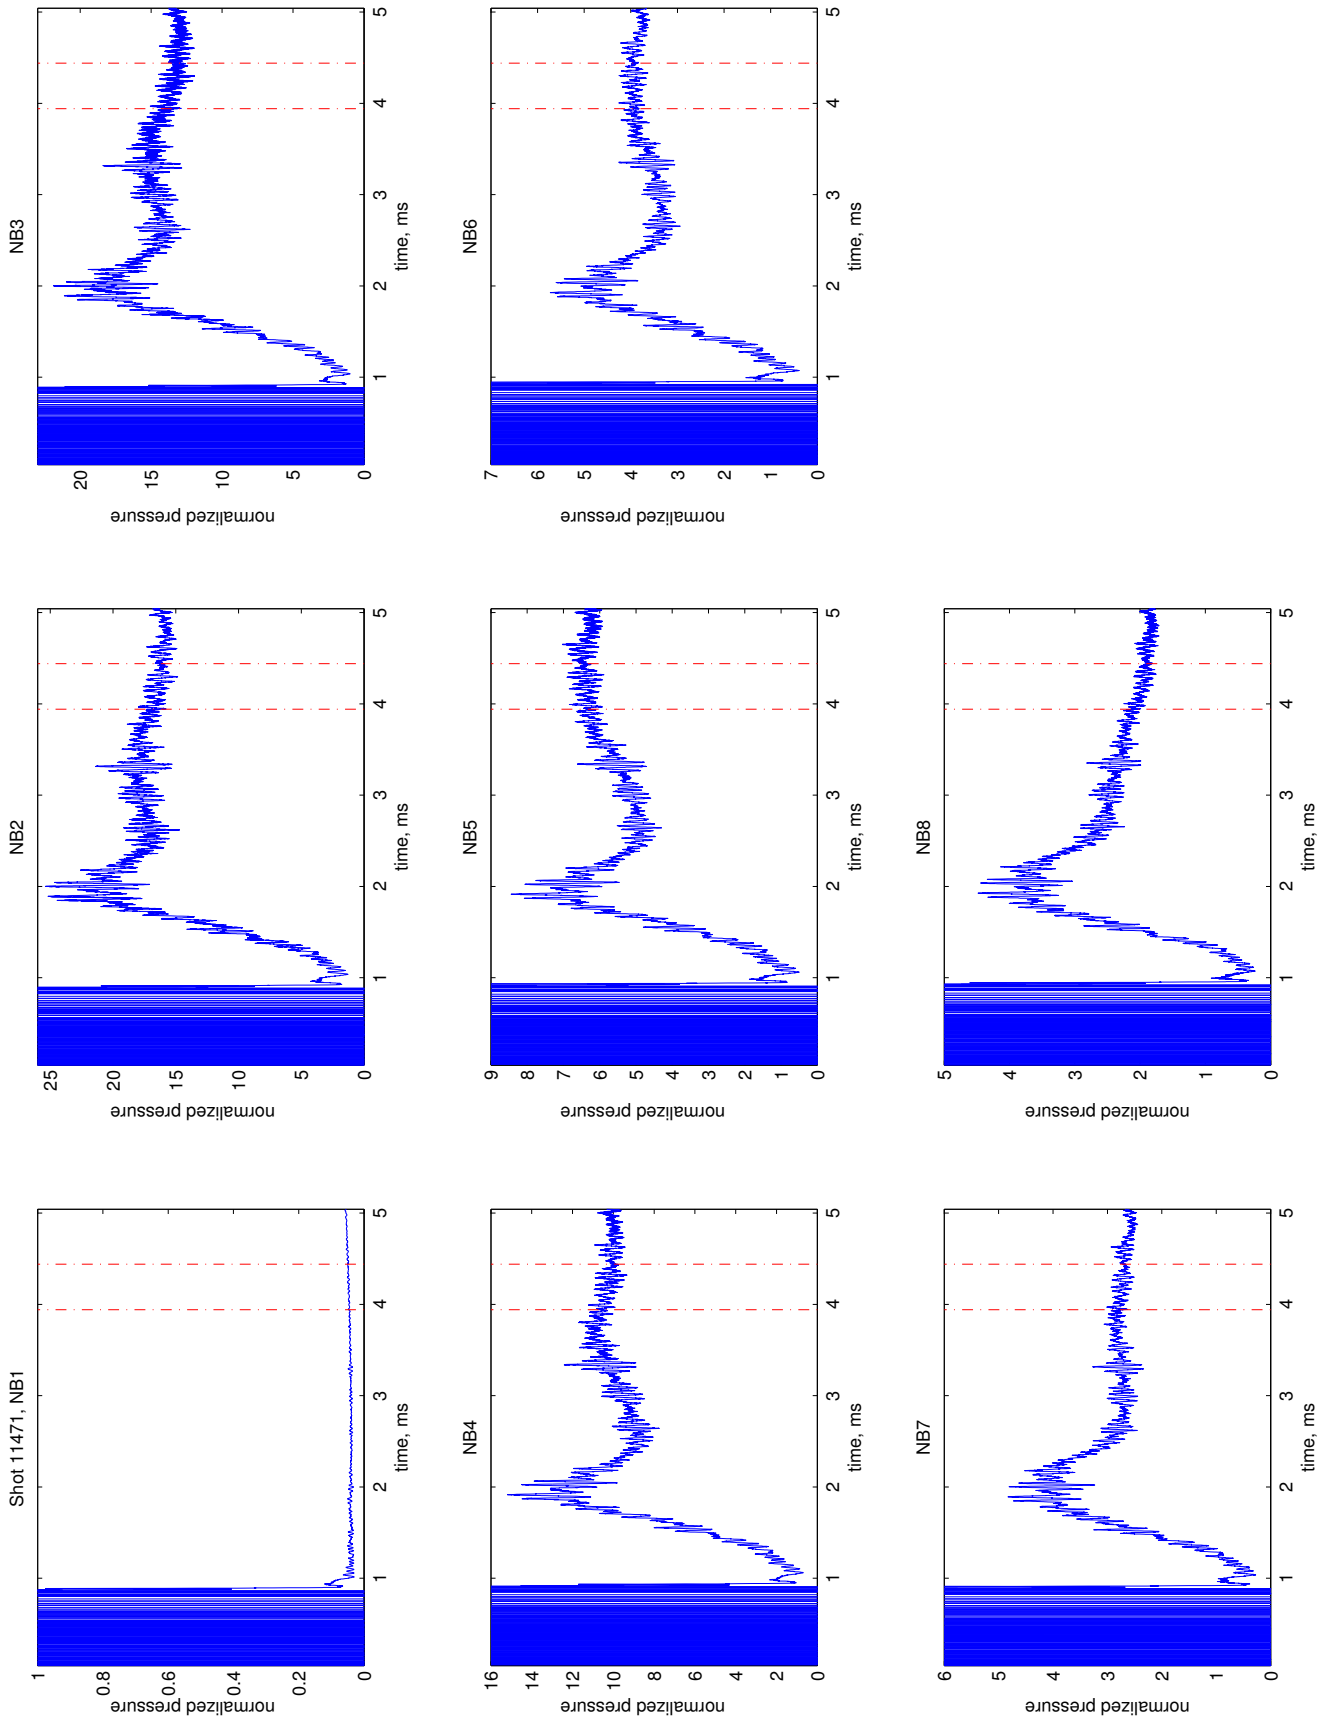

Figure H.7: Shot 11471 time histories of bodyside pressure transducers in the M12 REST nozzle

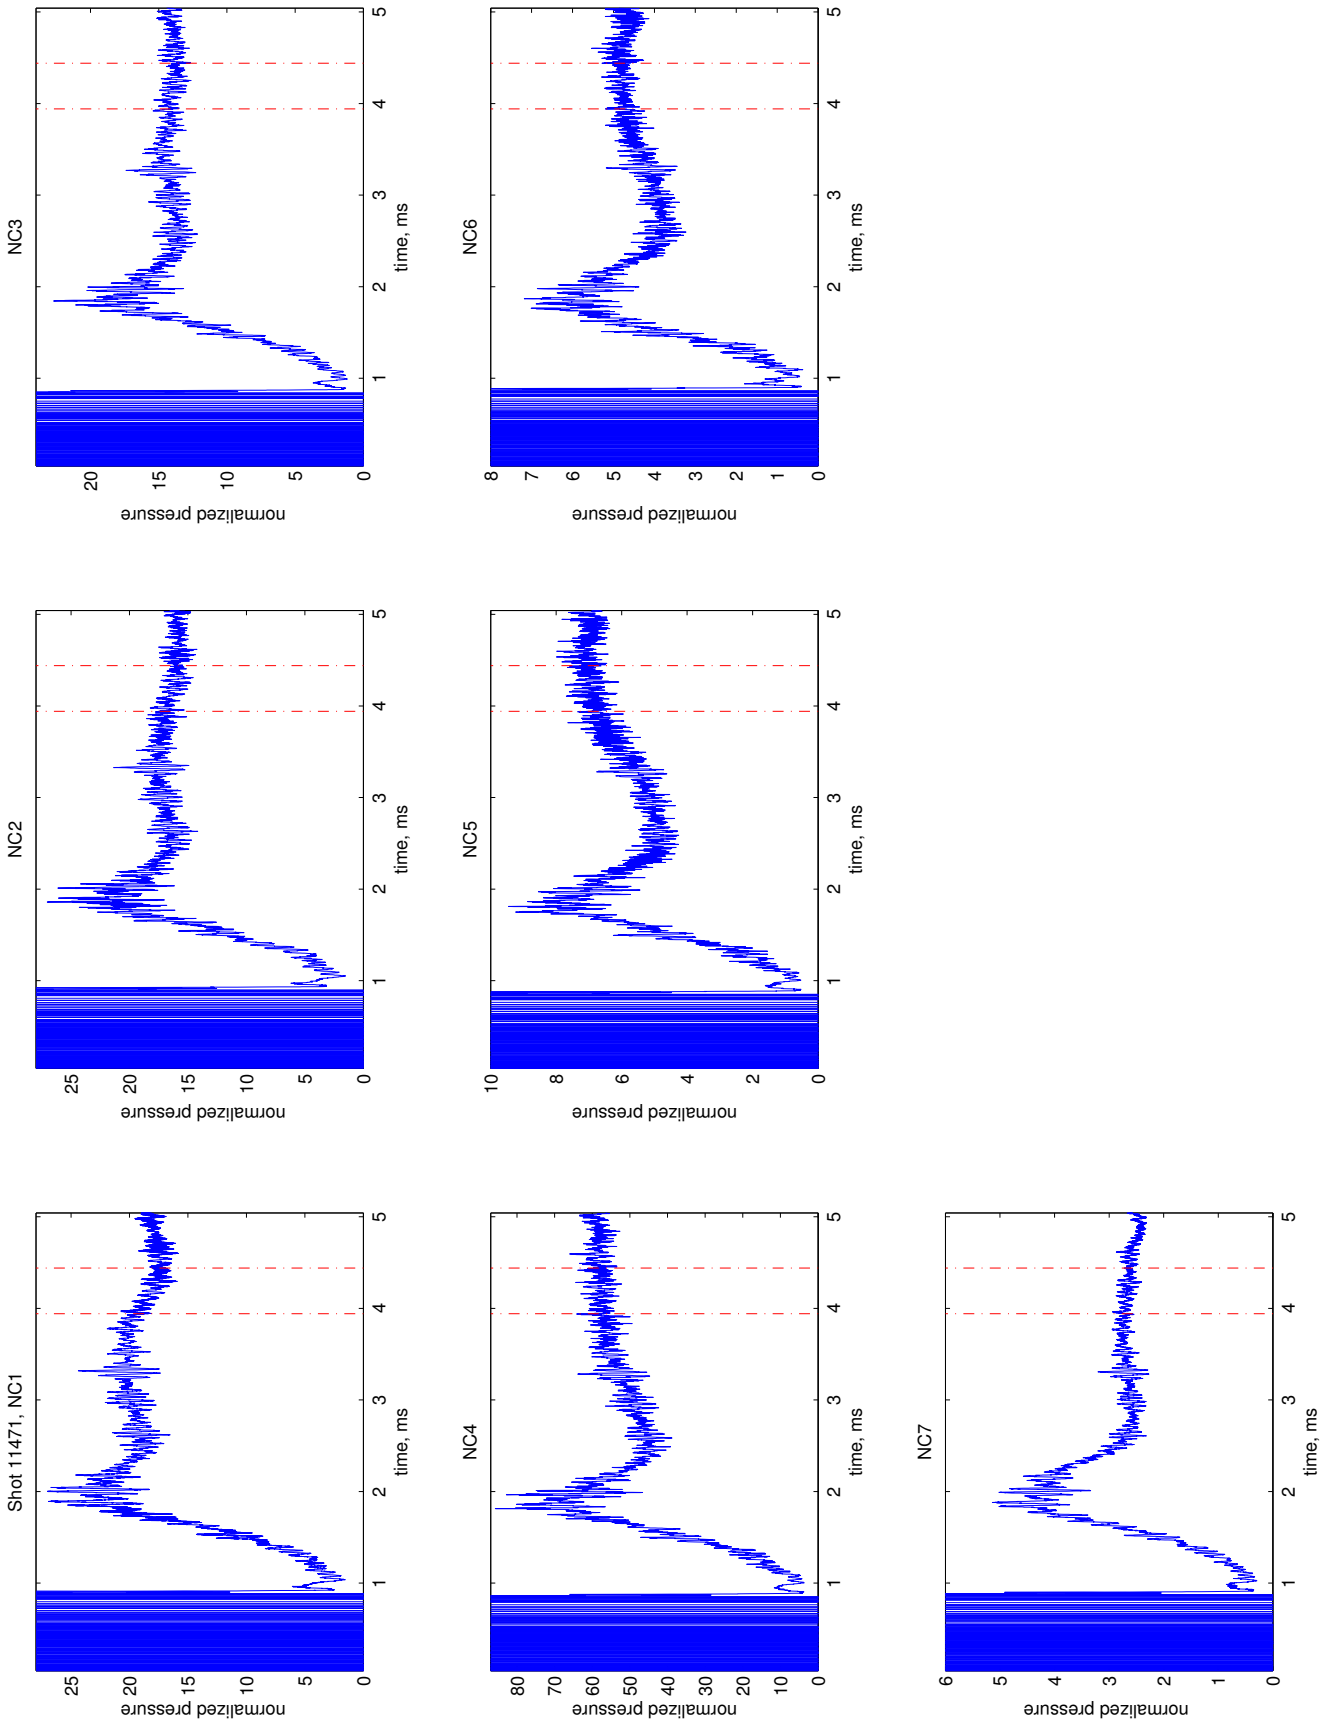

Figure H.8: Shot 11471 time histories of cowlside pressure transducers in the M12 REST nozzle