The numerical prediction of the acoustic response of liquid fuelled, swirl stabilised flames

by

Nicholas C. W. Treleaven

A Doctoral Thesis

Submitted in partial fulfilment of the requirements for the award of
Doctor of Philosophy of Loughborough University

January 2019

Abstract

This thesis sets out and tests several different approaches to predicting and understanding the acoustic response of an industrially representative liquid fuelled, swirl-stabilised, lean-burn fuel injector using numerical simulations. This work is important as it contributes to the design of fuel injectors with a low susceptibility for thermoacoustic instabilities or ‘rumble’.

The flame transfer function (FTF), a transfer function relating the mass flow rate through the fuel injector and heat release rate of the combustor, has been chosen as the best way to describe the flame response as it can be used in conjunction with a acoustic field solver to predict the stability of a combustion system. The FTF of a chosen injector geometry has been predicted using conventional compressible methods and a novel incompressible method which has been shown to be consistent with the compressible method at two frequencies of forcing. This is in contrast with mass flow forced incompressible simulations that fail to reproduce the correct downstream flow field fluctuations. The single phase flow field and acoustic response of the injector has also been predicted and compared to experiments with good agreement.

The injector hydrodynamic response has also been investigated along with how hydrodynamics, acoustics, the fuel spray and heat release are related. Acoustic forcing can be seen to actively alter the strength of large scale fluid structures, the mean pressure field and the mass flow rates through the different injector passages. The fuel spray may also couple with these structures causing additional local changes to the mixture fraction field and heat release rates. The effects of fuel spray SMD (Sauter Mean Diamter) and fuel spray injection velocity have on the flame have also been tested showing that the fuel spray atomisation, which can also be affected by acoustic forcing, may play a significant role in combustion instabilities.

Several novel numerical methods have been developed and are also discussed in detail including methods relating to the reproduction of acoustic forcing in incompressible simulations and the reproduction of turbulent fields at inlets. Several innovative post-processing techniques have been employed to identify the relationship between large scale flow structures, the fuel spray and combustion.

Modifications of the original injector geometry have been proposed to reduce the sensitivity of the injector to instabilities. These include better atomisation and mixing, better placement of swirl vanes, better aerodynamic design and improved hydrodynamic stability.

Acknowledgements

Firstly a big thank you to my supervisors at Loughborough University: Andrew Garmory and Gary Page, and at Rolls-Royce: Marco Zedda. I would also like to thank my unofficial supervisors at Rolls-Royce Deutschland: Max Staufer, André Fischer and Claus Lahiri.

Thank you to Rolls-Royce and the EPSRC for funding my project and for paying my extortionately high fees as an international student. I am also indebted to those I have met in France, Britain and Germany who have made me feel welcome over the last few years.

I must thank my ‘co-authors’, Jialin Su and Mark Cassel from Loughborough University and Matthew Juniper and Andrea Giusti from Cambridge University. I must thank my other colleagues from Loughborough University and Rolls-Royce for their support and insight. The support staff at Loughborough deserve a medal for putting up with my constant travel related badgering. Next come my colleagues from Cambridge and Oxford who I completed my second masters degree with, the *CDT elite* and the newer students who I have met through the program. After that comes *Team Turbulence* a fantastic bunch of people who became good friends and helped me through a life-changing masters program in France. I should also acknowledge the help and support I received during the beginning of my academic journey from my colleagues at LTRAC, Monash University.

A big thank you goes out to those I've met along the way who have inspired me and helped me understand the greatest unsolved problem in classical physics. They are, among others: Michel Stanislas from LML, Bill George from Imperial, Vassili Kitsios from LTRAC, Rob Miller from the Whittle lab, Li He from Osney, Sébastien Candel from EM2C, Thierry Poinsot from Cerfacs and lastly but definitely not least Laurent Cordier and Peter Jordan from PPRIME Poitiers.

I will forever be indebted to my family and my wonderful partner Tegan for their ongoing support and understanding. I am a very lucky guy!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Thermoacoustic instabilties

Thermoacoustic instabilities occur as a result of a feedback mechanism between unsteady heat release, fluctuations in pressure and perturbations of flow and mixture within the combustion chamber. This phenomenon is responsible for the non-linear growth of pressure and velocity fluctuations within the combustion system to magnitudes that are significant enough to cause the emission of sound at levels audible from within the aircraft. Furthermore these fluctuations may become so large that the associated fluctuation of stress within the engine structure leads to fatigue of engine components or outright failure [74].

The sound produced by these fluctuations is often termed ‘rumble’, ‘growl’ or ‘howl’ depending on the fundamental frequency of instability [73]. Thermoacoustic instabilities are of fundamental interest due to the difficulties involved with their prediction and the cost associated with failing to predict them early enough in the design of a new component or system.

An early example of these kind of instabilities is the F1 rocket engine used for the Saturn V rocket used on the Apollo missions to the moon. A combustion instability problem led to the use of 1332 full scale tests to iteratively design a set of baffle plates to prevent the oscillation. Despite the expansion of expertise and the tools available to the modern combustion system design engineer, there can still be problems [113].

In the late 2000’s, an aircraft engine manufacturer produced and installed roughly 200 new helicopter engines before discovering that a thermoacoustic instability existed of such a magnitude that each engine would have had to be completely rebuilt after a very small number of flight hours. At a large cost to the enterprise, due to both the delay in airframe delivery as well as the direct cost of the research, a fully annular, compressible, large eddy simulation (LES) was undertaken of the entire combustion system [163]. This geometry was found to have an unsteady fluctuation of pressure and temperature which spun around the annulus of the combustion chamber that could only be observed in a fully annular test or at least using a method that could reproduce the acoustic field of the full annulus [163].

The difficulties involved with predicting this phenomenon are driven by the complexity and number of interactions within the combustion system. An important aspect involves the prediction of how sound waves propagate within the combustion system. This requires that any simulation or experiment must be able to correctly resolve the speed of sound, temperature, velocity, geometry and boundary conditions of the entire combustion system, including in some cases, the interfaces with the compressor and turbine. Furthermore the assumption of incompressible flow, a common simplification used in low-speed flows cannot be used in isolation due to the need to resolve the acoustic field. Finally, the use of liquid fuel in jet engines requires that atomisation, convection and evaporation of the fuel is modelled in a sufficiently accurate fashion.

The aim of this thesis is to show how expensive fully annular and multi-phase compressible simulations can be replaced by a set of smaller, smarter simulations that are hopefully cheap enough such that thermoacoustic instabilities can be identified early in the design process so that expensive late-stage modifications of combustion systems can be avoided in the future. In performing these simulations, this thesis also identifies a number of interaction mechanisms between acoustics, the flow field, the fuel spray and combustion that can be targeted in future fuel injector designs to limit the prevalence of thermoacoustic instabilities in gas turbines.

1.2 Modern aero-engine combustion chambers

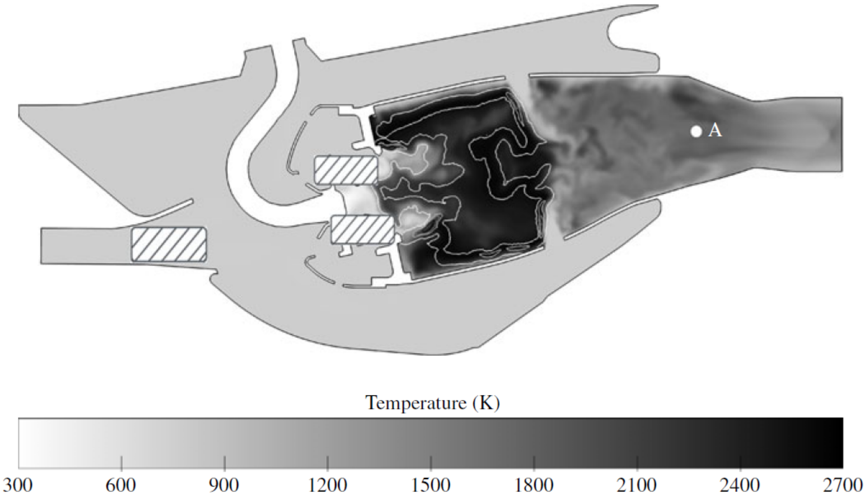

Figure 1.1: An example of a conventional liquid fuelled, swirl stabilised aero-engine combustion system of the RQL (rich, quench, lean) type. Reproduced from Motheau et al. [95].

An example of a modern aero-engine combustor is visualised in Figure 1.1 with some of the features covered for reasons of geometric sensitivity. The modern combustion chamber is fully annular in shape meaning that it is continuous in the azimuthal direction. This annular chamber is divided into a series of *sectors* which contain at the upstream end a fuel injector. This differs from the canular combustion systems preferred in land based gas turbines for power generation (see Figure 1.2) which comprise of a set of individual *cans* arranged around the centreline of the engine. These cans are often extremely large as this allows the fuel and air to mix thoroughly enabling very efficient, clean combustion. Aero-engine combustion chambers must be as small as possible to minimise the size and weight of the engine. Aero-engines also operate on kerosene as a fuel due to its desirable energy to volume ratio as compared to gaseous fuels. The liquid fuel must then be introduced through the fuel injector and then atomised into small droplets that convect downstream, evaporate and mix into a combustible mixture. Gas turbines on land are often run on gaseous fuels removing the complicating effects of atomisation.

The thermodynamic efficiency of an engine is largely dependent on the minimisation of pressure losses across the combustor. Unfortunately the need to minimise the size of the combustion chamber requires an extremely high power density provided by a highly turbulent flame. The turbulence present within the combustion chamber is driven by the pressure drop across it. In order to minimise this pressure drop, swirl is introduced to air passing through the fuel injector. This creates a central recirculation zone (CRZ) that promotes mixing, turbulent production and combustion. This also ensures that the hot combustion products are continously brought back into contact with the freshly injected reactant, stabilising the flame [73].

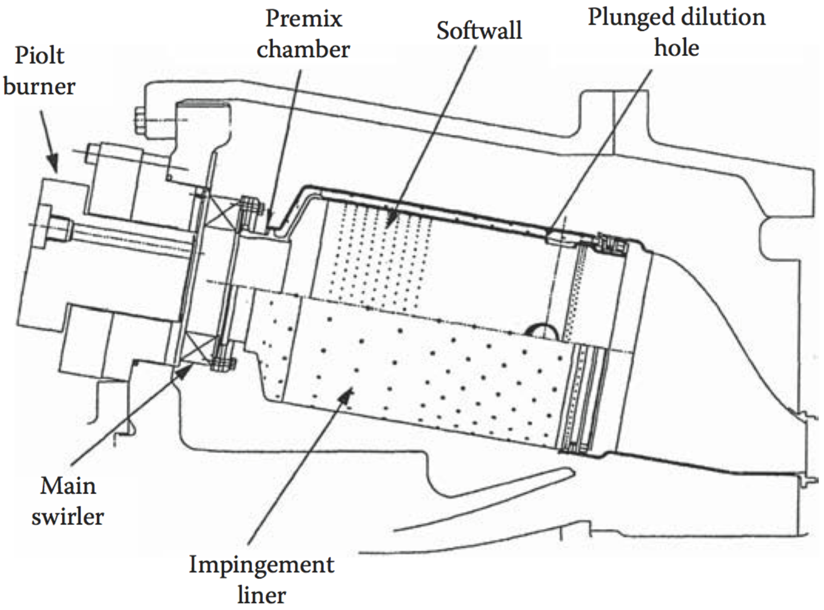

Figure 1.2: An example of a premixed lean burn system. Reproduced from Gallimore et al. [40].

1.2.1 Lean-burn combustors

Combustion instabilities have been a recurring problem but have drawn more significant interest recently due to the emergence of (Dry) Lean-Burn combustion systems. Combustor design has long been influenced by the production of Nitrous oxides (NO_x) that contribute to smog, acid rain and climate change [134]. In order to limit the chances of this occurring, regulation has been introduced and is being progressively tightened to force manufacturers and operators of gas turbines to reduce NO_x production [34]. The first generation of technology designed to limit the production of NO_x was the RQL burner, so called this because of the two different combustion zones separated by a series of jets. RQL (Rich, Quench, Lean) technology (see Figure 1.1) prevents the production of NO_x by avoiding the conditions which favour their formation such as high temperatures with an abundance of oxygen. In this case the fuel is first burnt richly to limit both the temperature and more importantly the abundance of oxygen and then undergoes a quick quenching process through the injection of large quantities of fresh air that switches the combustion quickly from being rich to being lean and lowering the temperature to a level that restricts the production of NO_x.

Whilst regulation is imposed by government, it is industry that advises government on what is possible in the short, medium and longer term. In the case of land based gas turbines the next technology was referred to as Wet Low Emissions (WLE) [3]. As nitrous oxide is produced by high temperatures, steam or water can be injected into the engine to limit the flame temperature and reduce emissions.

This technology was never adopted for aircraft engines as it would have required very heavy sub-system addition and the inclusion of large water tanks, reducing payload and increasing fuel burn for even short haul operations. Due to this fundamental difference in feasible technology, the rate at which regulations have tightened for each sector has been different with land based turbines being required of continually higher reductions of NO_x emissions. This difference led to the introduction of the current state of the art system for land based turbines named Dry Low Emissions (DLE). Dry low emission systems (see Figure 1.2) reduce the flame temperature by ensuring that the mixture fraction is relatively low. The equivalent system for jet engines is called 'Lean-burn' due to the combustion always remaining lean.

In the case of land based systems, the size of the combustion chambers has been increased significantly, as has the size of the fuel injector or air swirler. Dilutions ports have been eliminated leaving only small cooling holes. Furthermore there is a neck between the air swirler and main combustion system, the purpose of which is to keep the velocity of the swirling flow high enough to promote mixing without allowing for combustion. This kind of system is referred to as being 'premixed' and, providing there is adequate space, is extremely effective. Lean-Burn systems also have the advantage that they significantly reduce soot production and Unburnt Hydrocarbons (UHC's) from leaving the combustor, increasing efficiency [73].

In the case of aero-engines, several different lean-burn approaches have been tested however none of them take a truly premixed approach to fuel injection. This fundamentally comes down to the size restraints imposed in the case of propulsion systems. Length cannot be increased as it leads to excessive weight and drag whilst diametric increases get in the way of the bypass flow.

1.3 Rayleigh criterion

The fundamental equation that defines whether or not a thermoacoustic instability will grow is the Rayleigh Criterion, it can be derived by considering the definition of specific pressure volume work over a cyclic oscillation [120]:

$$\int_{\mathrm{T}} p'(t)\dot{Q}'(t)dt \geq \int_{V}\int_{\mathrm{T}}\sum_{i} L_i(\mathbf{x},t)dtdV \tag{1.1}$$

where p' is the pressure fluctuation in the region of the flame, $\dot{Q}'$ is the integral heat release rate fluctuation, L_i is the acoustic damping but the ith loss mechanisms and the integrals are computed over the acoustic cycle period and the combustor volume respectively.

The heat release rate does not affect itself or the pressure directly, but through a series of more complicated phenomena. However it is customary to assume that $\dot{Q}'$ is related to p' through a transfer function such that [74]:

$$\frac{\hat{\dot{Q}}'(\omega)/\overline{\dot{Q}}(\omega)}{\hat{p}'(\omega)/\overline{p}(\omega)} = \mathcal{G}_{p,Q}(\omega) \tag{1.2}$$

$\mathcal{G}_{p,Q}$ is the pressure to flame transfer function with the circumflex designating a Fourier transform. The transfer function can be decomposed into a gain (A) and phase (ϕ). The phase will determine the sign of the left hand side of equation 1.1 and is related to the delay between the pressure fluctuation and the heat release rate fluctuation through:

$$\tau = \phi/\omega \tag{1.3}$$

The job of the combustion engineer is to design the combustor such that τ results in the left side of equation 1.1 being smaller than the right side for the frequencies of interest (ω) which corresponds to resonant frequencies of the combustion system.

The delay τ is actually the sum of a number of processes including acoustic propagation time, fuel injector acoustic impedance, fuel atomisation, fuel transport, fuel evaporation, fuel mixing and fuel combustion. Each of these processes has its own characteristic time and process gain and as such τ can be re written as:

$$\tau = \tau_a + \tau_z + \tau_{sa} + \tau_{st} + \tau_{se} + \tau_{fm} + \tau_f \tag{1.4}$$

τ_z relates to the acoustic impedance; the relationship between changes of pressure close to the flame and the mass flow rate of air entering the combustion chamber. To avoid having to model the acoustic waves directly during computation of the heat release, the pressure to flame transfer function ($\mathcal{G}_{p,Q}$) can be split into the acoustic impedance and the flame transfer function, defined as:

$$\mathcal{F} = \frac{\hat{\dot{Q}}'(\omega)/\overline{\dot{Q}}(\omega)}{\hat{\dot{m}}'(\omega)/\overline{\dot{m}}(\omega)} \tag{1.5}$$

where $\dot{m}'$ is the mass flow rate fluctuations at the injector.

Studies using fully reacting compressible Large Eddy Simulation (LES) have proven useful and accurate in the determination of the flow both within and outside to the combustor primary zone [24, 37, 43, 64, 163]. The problem with this single approach is that its usefulness alone is almost as limited as a full scale experimental test. This is for two reasons, the first is that the acoustic field is sensitive to the entire combustion system and can interact with the compressor and turbine systems therefore requiring that the entire system must be explicitly meshed and modelled. The other limitation is due to the sensitivities of combustion instabilities to operating conditions. Instabilities may only develop at particular operating points, power settings or atmospheric conditions requiring many different environmental parameters to be individually altered and tested. The combined cost of all of this, from meshing to computation to operating conditions, would be enormous and almost certainly negate any advantages of computational methods against full scale experimental testing. In contrast this study follows methods to divide the thermoacoustic system into two distinct parts: the large wavelength acoustic field and short wavelength convective (combustion) field as shown in Figure 1.3 [123].

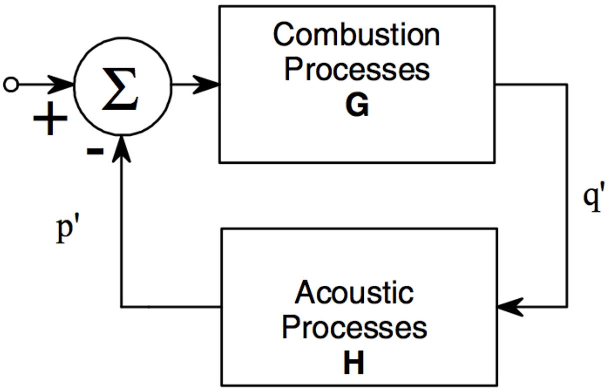

Figure 1.3: Splitting the thermoacoustic cycle into acoustic and convective (combustion) processes. Reproduced from [123].

This acoustic field, in a practical application of the work discussed in this thesis, would be calculated using a low-order acoustic network model which is briefly described in Section 1.5.1. The acoustic model would therefore generate the favourable acoustic modes of the entire combustion system (H in Figure 1.3) and provide an input fluctuation of pressure (see Figure 1.3) or its corresponding mass flow through the fuel injector. It is then assumed that all of the geometric sensitivity outside of the fuel injector is captured in this acoustic network model.

The flame is therefore assumed to be sensitive to changes in mass flow through the fuel injector only and independent of the rest of the combustion system. This allows the flame to be studied in isolation using either numerical or experimental methods. In this thesis the aim was to focus on numerical methods that enable the testing of future injector designs without the requirement of manufacturing them. The outcome of these simulations is therefore a flame transfer function (G in Figure 1.3) that gives the heat release rate of the combustor as a function of the injector mass flow. The acoustic field is then assumed to be perturbed by this heat release according to Lighthill's equation [75]. The key processes that must be reproduced and key transfer functions are shown in Figure 1.4 and will be explained in the following sections: pressure fluctuations in the combustion chamber lead to mass flow rate fluctuations through the fuel injector which then leads to direct modifications of the heat release rate through interaction with the flame while also perturbing the atomisation process leading to fluctuations of mixture fraction that also alters the heat release rate. Finally heat release fluctuations then contribute to the pressure fluctuations in the combustion chamber.

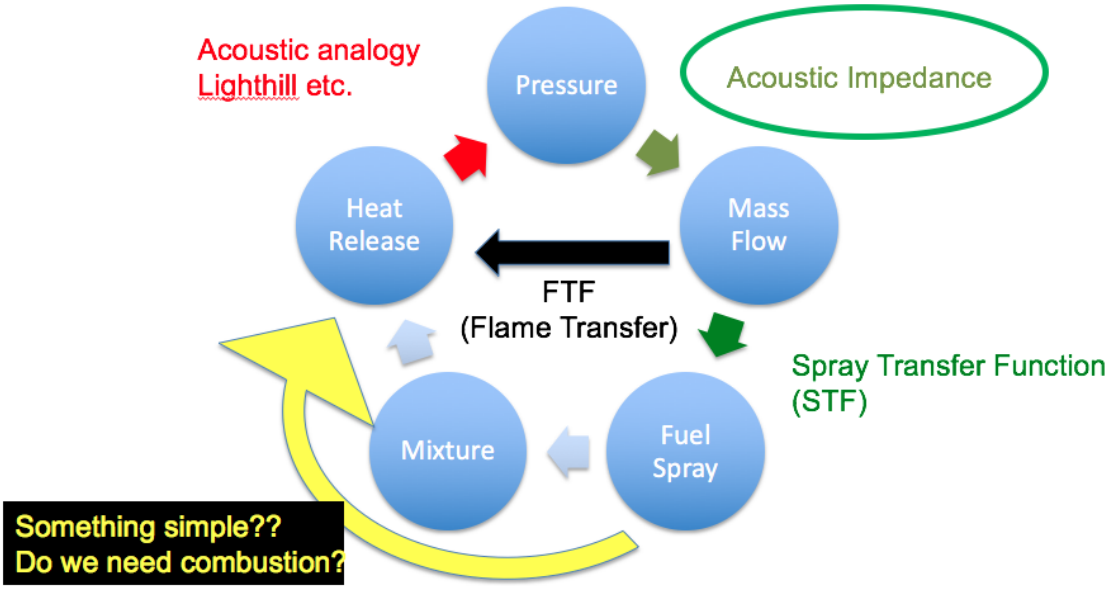

Figure 1.4: The thermo-acoustic cycle and the various transfer functions that exist between processes.

1.4 Neglected mechanisms

There are several possible mechanisms for how the flame may be affected by acoustics that have not been taken into account within the methodology presented here. They are the amplitude dependence of flame response, assumptions of flame compactness, fluctuations in fuel flow rate, low frequency rumble and the effect pressure can have on the immediate heat release.

1.4.1 Non-linear effects

Non-linear effects are related to the amplitude of the forcing of the flame. In order to increase the signal-to-noise ratio of either an experimentally or numerically determined flame transfer function the amplitude of the imposed pressure or mass flow rate fluctuations must be high, however, when a combustion system switches from stable to unstable operation the fluctuation magnitude at the characteristic frequency will be low and grow in time. In order to predict this growth the gain in the system must be positive for this low amplitude [89]. Therefore in order for the flame to respond similarly in both the real engine and the FTF test, the flame response must be sufficiently linear. As can be seen from the FDF (Flame Describing Function) in Figure 1.5, the FDF gain does change as a function of amplitude however it is hoped that in the case of the injector studied here that the system remains sufficiently linear to enable prediction of instability even if the growth rate is calculated incorrectly. If the FTF is not sufficiently accurate it can be replaced by an FDF that accounts for the effect of forcing amplitude. The disadvantage of this method it that the flame must now be forced at various different amplitudes to obtain the flame response making testing significantly more expensive. Linearity also extends to the frequency response, if the flame response is linear then forcing at one frequency will cause the flame to respond at the same frequency. If the flame response is not linear it is possible for phenomena at different frequencies to interact. Linearity allows multiple frequencies to be tested at once, non-linearity implies that each frequency must be tested individually making testing expensive and time consuming. The acoustic field should remain linear at the amplitudes of forcing applied however, the flame is dependent on the turbulent flame speed and turbulence is a fundamentally non-linear phenomenon. By forcing the flame at single and multiple frequencies, the linearity of the flame transfer function has been discussed in Chapter 8.

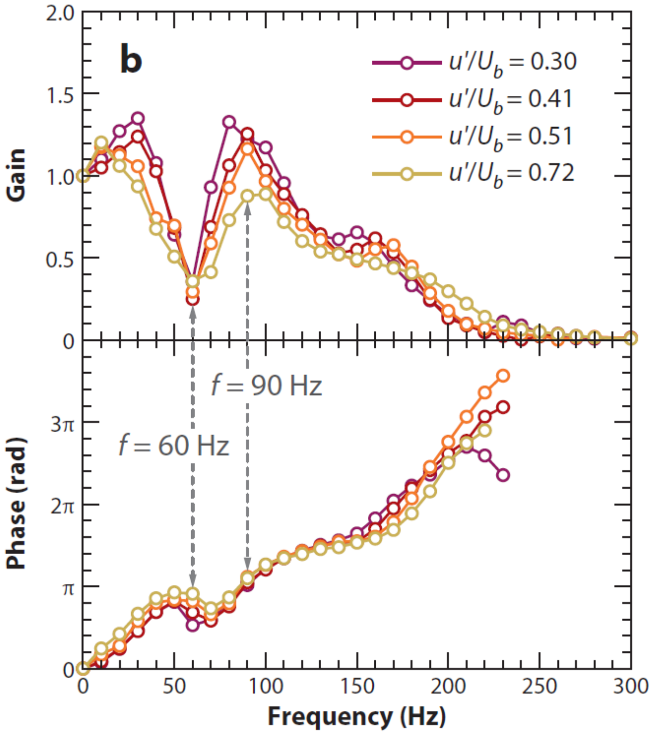

Figure 1.5: The FDF for a premixed swirl stabilised flame (reproduced from [16]). U_b is the mean flow velocity measured by the HWA probe.

1.4.2 Non-compactness of the flame

Similarly to other acoustic components in the combustion system the use of a transfer function implies that the flame is compact with respect to the acoustic wave. It is true that the flame reaction zone is extremely thin compared to the acoustic wave however according to the mechanisms of flame response discussed in Section 1.5.5 the flame response is dependent on the time delay between the acoustic wave and the convection of the azimuthal fluctuations. Acoustic waves travel both upstream and downstream meaning that this time delay will depend on the direction of the acoustic wave. Because the FTF only relates the fluctuation of mass flow and heat release at a point it makes no distinction as to the direction of the acoustic wave. In this case it may not be a sufficient measure of flame response and the FTF should be replaced with a flame transfer matrix.

1.4.3 Fuel flow rate fluctuations

It has been assumed in this study, as it is often [37, 43, 64, 65, 163], that the mass flow rate of fuel remains constant. Obviously if fluctuations of air mass flow rate can alter the flame response through changes in mixture fraction, so can fluctuations in fuel flow rate. It is assumed that the acoustic impedance of the fuel system is extremely high and as the fuel is quite incompressible, acts to damp out any pressure driven fluctuations in mass flow rate. Studies at LUTC have shown that fluctuations in mass flow rate can also occur through the interaction of fuel system components with vibrational modes of the engine casing [19]. When a structural mechanical oscillation is established the fuel manifold may be found to vibrate. This vibration sets up acoustic waves inside the fuel manifold leading to fluctuations of fuel flow rate through the injection nozzles. This fluctuation of fuel flow rate already corresponds to a structurally unstable vibrational mode and if this also corresponds to a favourable acoustic mode could lead to extremely large oscillation of pressure and stress.

1.4.4 Low frequency rumble

Low frequency rumble is a term given to combustion instabilities that occur due to fluctuations in flame response that lead to compositional or temperature inhomogeneities (entropy waves) in the combustion chamber [95]. These inhomogeneities produce acoustic waves when they are accelerated through the nozzle guide vanes (NGV's). The methodology presented here could be used to identify these instabilities with an additional model that identifies these inhomogeneities leaving the flame and adds a delay filter corresponding to the convection of these structures to the NGV's.

1.4.5 Pressure effects on heat release

The combustion model used in this study is the Presumed Probability Density Function (PPDF) Flamelet Generated Manifold (FGM) method [157, 161]. Such a methodology describes combustion as being a function of a small number of control variables. In the case of this study, those control parameters are the progress variable, a measure of how reacted the mixture is, the fuel mixture fraction and the variances of the progress variable and mixture fraction. Pressure fluctuations are assumed not to affect the reaction rates. In the case of gas turbine combustion, fluctuations in pressure are assumed to be small and of such a long wavelength that local pressure gradients are also negligible. For this reason, pressure fluctuation dependence on reaction is ignored. In the case of internal combustion engines pressure fluctuations over a cycle are extremely large and the flow may be of sufficiently high Mach number such that shock waves and hence, large gradients of pressure can exist. These do effect the reaction rates of certain reactions between fuel, air and intermediate species. Mittal and Pitsch [88] have suggested an extension of the FGM method to include the pressure as an additional control parameter. This means that the neglected physics associated with pressure fluctuations could be better represented through modification of the combustion model. The expansion of the model to include an additional control parameter does however increase the integration time of the FGM table and the FGM table size, making simulations more expensive. For this reason, direct pressure effects on combustion are ignored in this study but offer a reasonably easy extension of the method that may improve the accuracy of the results.

1.5 Factors affecting thermoacoustics

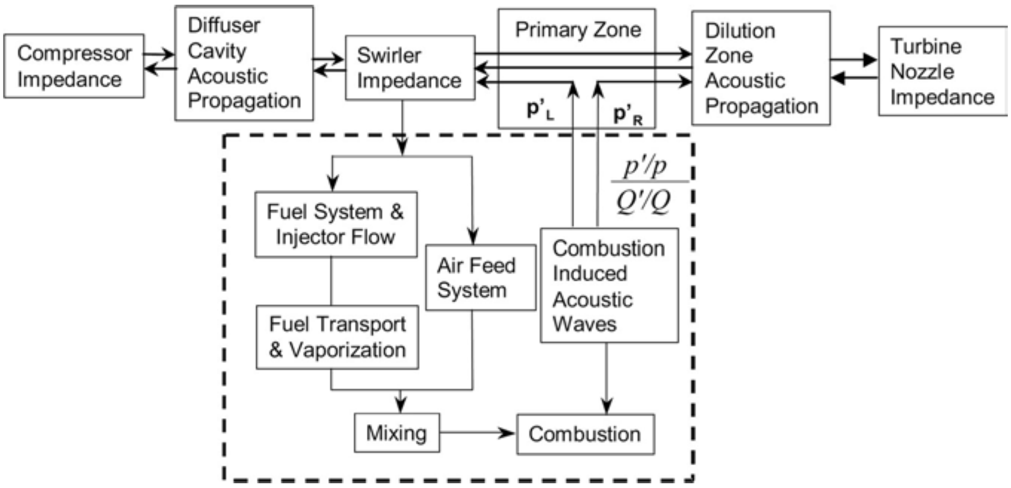

Figure 1.6: The acoustic and convectively active processes in the case of a liquid fuelled engine. Reproduced from Mongia et al. [93].

Figure 1.6 shows how combustion can be affected by variations of mixture fraction due to the processes effecting the injection and mixing of the fuel and air. All of these convectively driven processes are contained within the dashed line and are affected by the acoustic field determined by the components outside the dashed area. Each of these processes are described in Section 1.5.3. The combustion process then feeds energy back to the acoustic field according to Rayleigh's criterion.

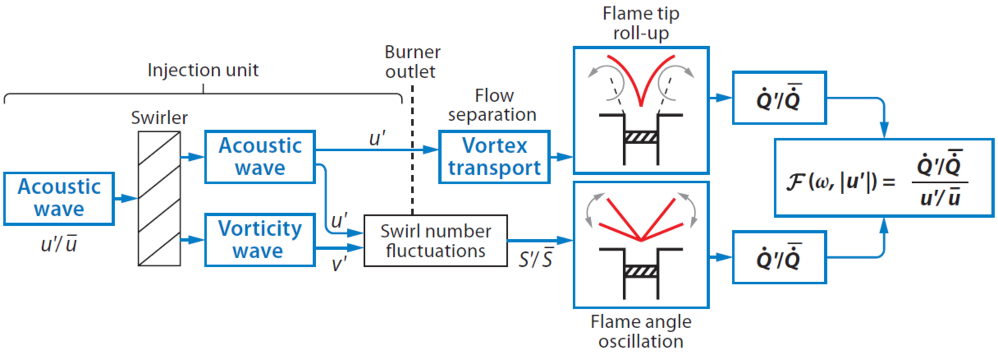

Figure 1.7: The effect of fluctuations in axial velocity and swirl number on the flame transfer or flame describing function (reproduced from [16]).

Variations of mixture fraction are not the only way the heat release can be modified. Studies by Palies *et. al.* [105] on premixed flames have highlighted two independent mechanisms driving fluctuations in swirl stabilised flames where the equivalence ratio or mixture fraction is constant.

When approaching the fuel injector, the acoustic wave induces an acceleration of the air within the tube shaped injector body. This modified velocity induces vorticity at the exit of the injector leading to a vortex ring that is convected towards the flame. The flame in this case can be considered to be analogous to a laminar 'V' flame with a modified turbulent flame speed S_T. The approaching vortex ring causes the arms of the V to roll up, modifying the flame area and therefore the heat release rate.

The same axial velocity fluctuations that cause the production of the vortex ring also interact with the swirl vanes within the injector creating an azimuthal velocity fluctuation. The azimuthal fluctuation will then cause a variation of swirl number within the combustor. The swirl number also contributes to heat release through its contribution to turbulent flame speed and therefore the angle of the 'V' flame [51]. The individual effects of axial velocity fluctuations and swirl number fluctuations on the flame and the FTF are shown in Figure 1.7 and further described in Section 1.5.5.

1.5.1 The acoustic field

The first delay (τ_a) is associated with the acoustic field and assuming that the Mach number is low will be approximately equal to the distance from the flame to the injector divided by the speed of sound. The acoustic field is important as it determines the resonant frequencies of the combustion system. Just as a flute resonates at particular eigenfrequencies so does a jet engine combustion chamber. The flute player alters the resonance frequency by changing the boundary conditions. When a new hole along the flute is uncovered, the point along the flute corresponding to the newly opened hole becomes a pressure node, altering the resonant waveshape and hence the eigenfrequency. Frequencies corresponding to non-resonant modes will be naturally damped by the system and hence any instability will only occur at a system eigenfrequency.

Numerous methods exist for the determination of the system eigenfrequencies, these include non-linear methods such as the solution of the Navier-Stokes equations, or linearised methods such as the solution of the linearised Euler equations [107], Helmholtz equation [98] or linearised Navier-Stokes equations [17]. These computational methods are aided by the wavelengths for frequencies of interest being quite large requiring only modest mesh cell counts. Linearised methods have a further benefit in that they can be solved in Fourier space, significantly reducing the time to convergence. Linearisation of the relevant equations for acoustics is common and justified on the basis that even large sound pressure levels are caused by relatively small fluctuations in pressure. The associated fluctuations in velocity are commonly two orders of magnitude lower than fluctuations in pressure and therefore assumed to vary linearly.

The large characteristic acoustic wavelengths within the combustion chamber mean that many of the small scale features inside the combustion chamber are irrelevant and may be simplified. The low order acoustic network model (LOTAN) method [35] assumes that the combustion chamber and surrounding components may be represented as a series of ducts and acoustic elements. The ducts are of sufficiently simple geometry that the solution of their eigenmodes may be found analytically. The additional acoustic elements such as orifices, dampers and fuel injectors may be represented using simple analytical models, acoustic impedances or acoustic transfer matrices.

A simple analytical model may consider the fluid inside an orifice to be incompressible and that the mass and energy fluxes must be unchanged while the momentum flux is altered by the change in duct area only (see Section 5.1 for an example). The acoustic impedance is a transfer function that relates the fluctuation of acoustic particle velocity to the local fluctuation of pressure:

$$\mathcal{Z}(\omega) = \frac{\hat{p}'(\omega)}{\hat{u}'(\omega)} \tag{1.6}$$

where the circumflex designates a Fourier transform. In the case where an acoustic element cannot be assumed to be much smaller than the wavelength, the acoustic impedance should be replaced by the acoustic transfer matrix ($\mathbf{T}(\omega)$) [89]. The acoustic transfer matrix is a 2 × 2 matrix that describes the relationship between upstream and downstream travelling waves and the acoustic element. It can either be expressed in terms of the upstream ($\hat{p}_1$,$\hat{u}_1$) and downstream ($\hat{p}_2$,$\hat{u}_2$) pressure and velocity:

$$\begin{bmatrix} \hat{p}_2 \\ \hat{u}_2 \end{bmatrix} = \mathbf{T}(\omega) \begin{bmatrix} \hat{p}_1 \\ \hat{u}_1 \end{bmatrix} \tag{1.7}$$

or it can be represented as an acoustic scattering matrix which relates the upstream and downstream Riemann invariants $\mathcal{B}$ and $\mathcal{C}$ [15]:

$$\begin{bmatrix} \mathcal{B}_2 \\ \mathcal{C}_2 \end{bmatrix} = \mathbf{N}(\omega) \begin{bmatrix} \mathcal{B}_1 \\ \mathcal{C}_1 \end{bmatrix} \tag{1.8}$$

where $\mathcal{B}$ and $\mathcal{C}$ are representative of acoustic waves propagating in the downstream and upstream direction:

$$\mathcal{B} = \frac{1}{2}\left(\frac{p'}{\rho c} + u'\right) \tag{1.9}$$

$$\mathcal{C} = \frac{1}{2}\left(\frac{p'}{\rho c} - u'\right) \tag{1.10}$$

The subscripts 1 and 2 represent the upstream and downstream locations. The acoustic impedance or acoustic transfer matrix may be obtained using either analytical, experimental or numerical methods.

The important contributions to the acoustic field are shown by the objects outside of the dashed line in Figure 1.6. These include the impedances of the compressor, fuel injector (swirler) and turbine as well as the acoustic propagation through the upstream diffuser and dump cavity and downstream dilution zone.

1.5.2 From acoustic field to flow field

The delay (τ_z) and gain between the arrival of an acoustic wave at the fuel injector and its corresponding mass or velocity fluctuation is governed by the injector acoustic impedance [149]. The subsequent delays associated with the fuel spray and combustion are also strongly related to the acoustic response of the injector [148].

Acoustic Impedance ($\mathcal{Z}$)

The acoustic impedance of the fuel injector is a critical part of the determination of the system stability and has been measured extensively at LUTC (Loughborough University Technology Centre) using the methodology developed by Rupp [126] which involves using loudspeakers to excite the flow acoustically whilst measuring the system response using the multi-microphone technique [136]. This method requires the injector be physically built and therefore remains an expensive design tool. Gunasekaran and McGuirk [45] first showed that a URANS (Unsteady Reynolds-Averaged Navier Stokes) equations solver could be used to predict the reflection coefficient of a duct-fed orifice by assuming mildly compressible pressure. The acoustic boundary conditions developed for that study were then implemented in the open source compressible CFD code OpenFOAM. Following this Su et al. [149] have shown that this methodology can be extended to the far more complicated geometries typical of jet engine fuel injectors. The importance of this work stems from the use of the URANS methodology instead of the more costly LES or from experimental testing. Acoustic impedance calculations were validated against data acquired using the multi-microphone method.

1.5.3 From flow field to mixture fraction

In a lean burn system, the majority of air is introduced through the fuel spray nozzle or fuel injector and from here the flow field within the combustion chamber is established. In an aero-engine the fuel (usually kerosene) is introduced as a liquid into the combustion chamber through the fuel injector where it is atomised, convected by the flow, evaporated, mixed with the surrounding air and burnt. The time delays relating to the first four of these mechanisms are discussed in this section ($\tau_{sa}, \tau_{st}, \tau_{se}, \tau_{fm}$). This section corresponds to the left hand branch of the processes outlined within the dashed line in Figure 1.6.

Atomisation and fuel spray transport

Lefebvre [73] provides the following estimate of the largest droplet of fuel atomised as a function of the relative flow velocity between fuel and air (U_R), flow density (ρ) and fuel surface tension (σ).

$$D_{max} = \frac{12\sigma}{\rho U_R^2} \tag{1.11}$$

As the droplet diameter is a function of air velocity it is likely that fluctuations of air velocity induced by the acoustic field will cause fluctuations in the fuel droplet diameter.

Measurements of the atomisation of sprays under high amplitude pressure fluctuations have been performed using PDA (Phase Doppler Anemometry) at Loughborough University using a method similar to that originally described by Qiu and Sommerfield [118]. These experiments provide data on the distribution, velocity, and size of fuel droplets at various locations around the fuel injector [18, 148].

Because the fuel spray is composed of a very large number of droplets of different sizes, the spray droplet diameter is usually expressed in terms of a statistical quantity called the Sauter Mean Diameter (SMD) defined as:

$$D_{32} = \frac{\sum_{i=1}^{N} D_{s,i}^3}{\sum_{i=1}^{N} D_{s,i}^2} \tag{1.12}$$

where $D_{s,i}$ is the diameter of the ith droplet and N are the total number of droplets in the interrogation volume. The PDA experiments show that the SMD of the droplet spray will change due to the acoustic forcing. Figure 1.8 shows the variation of SMD at a point downstream of the injector as a function of the air velocity close to the atomiser. This plot shows the characteristic phase portrait that is the result of the delay between the velocity increasing and the subsequent reduction of SMD of the fuel spray.

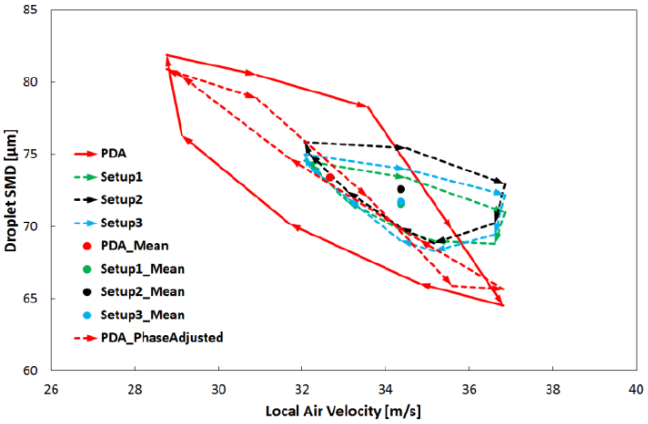

Figure 1.8: The variation of droplet SMD with local air velocity when the chosen injector is forced at $St = 0.5$. Reproduced from Su et al. [148].

Once the fuel has been atomised, the position and evolution of the fuel droplets must be taken into account. The critical value of interest is the Stokes number of each droplet which relates the particle response time to a characteristic flow time:

$$Stk = \frac{\tau_d}{\tau_f} \tag{1.13}$$

where the particle response time can be related to the particle diameter, density and air viscosity via the relation [83]:

$$\tau_d = \frac{\rho_d D^2}{18\mu} \tag{1.14}$$

The characteristic flow time may be defined as either being the turbulent eddy turnover time, the period of rotation of large scale eddies, the period associated with acoustic forcing or quite simply L/U_0 where L is a characteristic length and U_0 a characteristic velocity. The choice between which characteristic time to use is driven by what mechanism dominates the flow at a particular point. Regardless of the definition of τ_f, the fuel particle response can be described by the Stokes number. Those particles with $Stk << 1$ will respond quickly to any changes to the flow and disperse evenly amongst the flow structures (Figure 1.9a) whereas particles with $Stk >> 1$, will not respond to small high-frequency oscillations in air flow velocity. This will cause these particles to be attracted to the boundaries of large flow structures, see Figure (1.9d). For particles with a Stokes number of anywhere in between these ranges, the particle response will be a mixture of the previous two behaviours (Figures 1.9b and 1.9c). For droplets with a large Stokes number this effect may cause congregation of droplets in particular regions of the combustor hampering evaporation, leading to localised hot-spots or incomplete combustion.

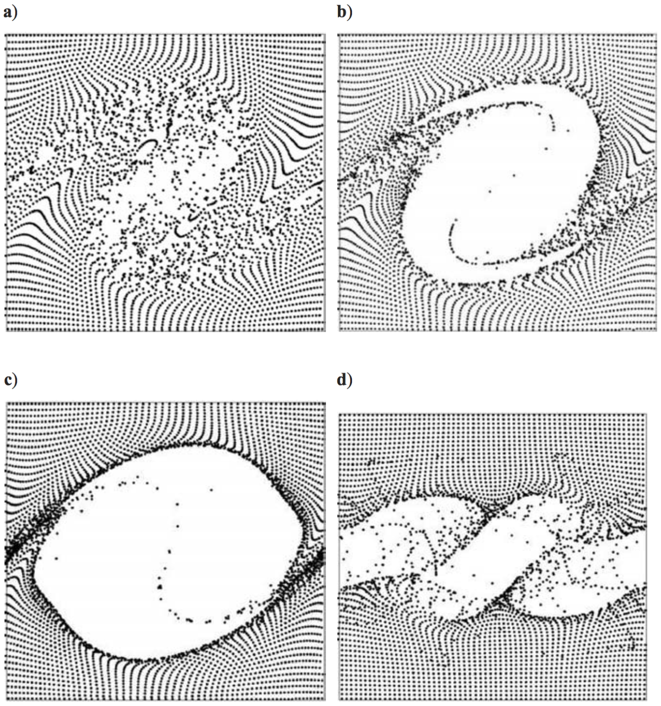

Figure 1.9: The droplet dispersion in a mixing layer for droplets with a Stokes number of a: 0.1, b: 1.0, c: 4.0 and d: 100. Reproduced from Menon [83].

Menon [83] performed a DNS (Direct Numerical Simulation) using a pseudo-spectral code of a mixing layer in a 64x64x64 box to look at how particles of different sizes would respond to to the flow features. Menon [83] defined a droplet dispersion function that measures the dispersion of droplets across the mixing layer:

$$D_y(t) = \sqrt{\frac{1}{N_p}\sum_{i=1}^{N_p}[Y_i(t) - Y_m(t)]^2} \tag{1.15}$$

where N_p is the total number of droplets, $Y_m(t)$ is the mean value of droplet displacement across the mixing layer at time t and Y_i is the displacement of the ith droplet across the mixing layer. He found that there was a preferential dispersion for droplets with $1.0 < Stk < 10$ implying that these are the droplets that will respond most favourably to velocity fluctuations.

It is convenient to combine atomisation and spray convection effects together such that:

$$\tau_s = \tau_{sa} + \tau_{st} \tag{1.16}$$

De la Cruz Garcia et al. [30] defined a transfer function relating to the response of the fuel spray to fluctuations of velocity or mass flow through the injector:

$$\mathcal{S}(\omega) = \frac{\hat{M}'(\omega)}{\hat{u}'(\omega)} \tag{1.17}$$

where M' is the fluctuations of light intensity Mie scattered by the fuel droplets. $\mathcal{S}(\omega)$, the spray transfer function (STF), remains somewhat of a theoretical concept. The most important property of the STF is that it can be related to the FTF [30]. For this reason M' could be replaced by any number of variables such as fuel spray mixture fraction, fuel spray velocity, fuel spray mass etc. or a combination of any of the above quantities depending upon which quantities have the greatest effect on the flame response.

The determination of a spray transfer function that closely resembles a flame transfer function opens a pathway to the prediction of spray driven combustion instabilities without the need to resolve combustion. If such a transfer function could be obtained, it implies that the combustion instability is driven by changes in the atomisation or convection of the fuel spray. In this case, combustion does not necessarily need to be included in the analysis as it can be concluded that the transfer functions relating fluctuations of heat release rate and any non-fuel spray related phenomena are relatively weak. The STF and FTF would converge and could be linked by a simple delay type model that relates the spray and the heat release such as that proposed by Crocco and Cheng [26]. This would significantly decrease the cost of both simulations and experiments as the problems associated with simulating and measuring high pressure, high temperature, reacting and compressible flows may be avoided. Such a simplification is shown graphically in Figure 1.4 by the yellow arrow.

Fuel evaporation and mixing

The time till evaporation (τ_{se}) is described in detail in Chapter 7. It is a function of pressure, temperature, droplet diameter, velocity, and the properties of the fuel [73]. Table 1.1 shows the approximate evaporation times of three different sized droplets based on the correlations shown in [73]. The three droplet sizes are representative of some of the smallest, largest and mid-sized particles. The approximations make use of the effective evaporation constant (λ_{eff}) defined as:

$$\lambda_{\text{eff}} = D_0^2/\tau_{se} \tag{1.18}$$

where D_0 is the initial diameter. Table 1.1 shows that the evaporation time may vary from anywhere between 1 μs to 1.7×10^4 μs. Also shown in the table are the characteristic times scales of the droplets and the corresponding characteristic frequency. Assuming that a stokes number of one will result is the largest particle response, the characteristic frequencies of droplets somewhere between one and eighty microns are aligned most closely to frequencies corresponding to rumble in gas turbines (100-1000 Hz). The mass of one droplet of fuel for each droplet size is also shown in Table 1.1. Whilst the smaller droplets might show the largest sensitivity to acoustic forcing they may contain proportionally less of the fuel. For this reason the effects of acoustic forcing on spray driven heat release should show more significance at lower frequencies.

Table 1.1: Approximate evaporation times at 2000 K and ambient pressure, droplet masses, characteristic timescales and characteristic frequencies ($f_d = 1/\tau_d$) for fuel droplets of different sizes.

Diameter (μm)	1	80	150
λ_{eff} (mm^2/s)	0.9	1.1	1.3
τ_{se} (μs)	1.1	5.8×10^3	1.7×10^4
τ_d (μs)	2.6	1.7×10^4	5.8×10^4
m_d (kg)	4.4×10^{-16}	2.3×10^{-10}	1.5×10^{-9}
f_d (Hz)	3.9×10^5	6.1×10^1	1.7×10^1

1.5.4 From mixture fraction to heat release

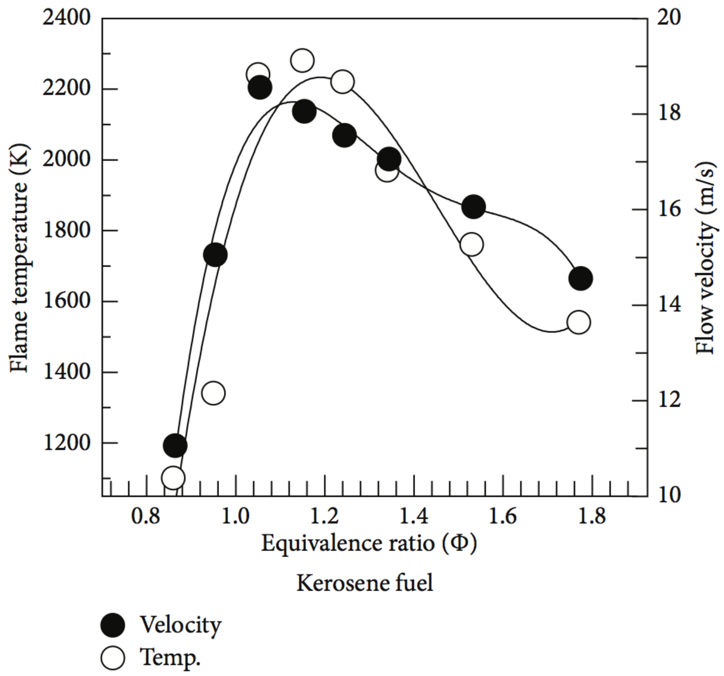

Figure 1.10: Flame temperature and speed for premixed kerosene at various mixture fractions. Reproduced from Shehata and ElKotb [137].

Figure 1.10 shows the adiabatic flame temperature of kerosene as a function of equivalence ratio. The gradient of flame temperature with respect to equivalence ratio is much higher in a lean flame than when the flame is close to stoichiometric or when the flame is rich. Because the flame is lean, it is much more likely to be perturbed by instability as the heat release rate is more sensitive to fluctuations of mixture fraction. This section describes the branch from mixture fraction to heat release in Figure 1.4 and from mixing to combustion in Figure 1.6.

Because there is a very real chance of the flame extinguishing in the case of a Lean Burn system, especially at low power conditions, the combustion system is split into two zones, the pilot zone, designed to always operate at a safe level and the main zone, only in operation when the engine is at a higher power setting. This zonal approach has resulted in a set of staged combustion systems, either axially or radially.

Rolls-Royce has preferred to go with an alternative option, the so called LDI (Lean Direct Injection) approach. The LDI system aims to produce two combustion zones but staged one inside the other. The pilot zone is kept close to the centreline and front face of the swirler while the main zone extends further downstream and further outwards radially. Whilst this approach is extremely compact it does significantly increase the complexity of the flame and airflow system.

1.5.5 From flow field to heat release

The time delay until combustion (τ_f) includes the effects of reaction chemistry, flame surface area and flame stretch [114]. The chemical reaction is extremely fast in most cases and flame stretch effects are minimal in the case where the flame front is accurately resolved. The vast majority of heat release in the case of premixed flames, where the equivalence ratio is constant, is therefore due to wrinkling of the flame by turbulence and large scale flow structures [114]. Turbulence has the effect of increasing the effective (turbulent) flame speed by increasing the surface area of the flame. Large scale flow structures such as vortices also distort the shape and surface area of the flame in the same manner [145].

Experiments and simulations with premixed flames

This section describes in more detail the processes shown in Figure 1.7 and how the mass flow fluctuations in the fuel injector may lead to heat release rate fluctuations that are unrelated to changes in mixture fraction. Studies by Palies et al. [105] of premixed flames have shown that there are two separate mechanisms that can affect the heat release rate without requiring a change in mixture fraction. In these studies the flame was fed by an upstream plenum with a loudspeaker embedded in the wall leading to a long premixing chamber with a single axial flow swirler. The velocity of air through the swirler was measured using a single hot-wire anemometry (HWA) probe upstream and the heat release was measured using OH^* chemiluminescence. The velocity to heat release flame describing function (FDF) obtained using this methodology is shown in Figure 1.5. The FDF is defined as [16]:

$$\mathcal{F}_{df}(\omega, |\hat{u}'|) = \frac{\hat{q}'(\omega)/\overline{q}}{\hat{u}'(\omega)/\overline{u}} \tag{1.19}$$

The FDF is the non-linear extension of the flame transfer function (FTF) defined as:

$$\mathcal{F}(\omega) = \frac{\hat{q}'(\omega)/\overline{q}}{\hat{u}'(\omega)/\overline{u}} \tag{1.20}$$

In this study mass flow rate will be frequently interchanged with the area averaged velocity fluctuation as the injector flow surface area does not change and it is assumed that the air density is constant.

For laminar flames the imaginary part of the FDF, related to the time delay, has been shown to be relatively invariant with respect to fluctuation amplitude [102] and Figure 1.5 shows that the same is true for premixed swirling flames. The prediction of instability relies on the correct prediction of the delay between pressure and heat release which can be resolved through correct resolution of the convective and reactive time delays together with the acoustic characterisation of the combustion system [30]. The advantage of using a non-linear FDF is that the transient, non-linear behaviour of the system can be simulated including growth rates, triggering and limit cycles [102].

The magnitude (or gain) of the FDF (see Figure 1.5) can be seen to exhibit a characteristic 'V' shape with a maximum gain at around 90 Hz and a minimum gain at around 60 Hz. The swirl number and its fluctuations, first defined by Beer and Chigier [10] was also approximated at the exit of the burner. The swirl number is a non-dimensional number that is linked to the likelihood of a central recirculation zone (CRZ) appearing in the combustion chamber and is defined as the ratio between tangential and axial thrust and is often approximated by [73]:

$$S = \frac{\int_0^{R_o} u_\theta r \rho u_x r dr}{R_o \int_0^{R_o} u_x^2 \rho r dr} \tag{1.21}$$

where R_o is the outer radius of the fuel injector, r is the radial coordinate, u_θ is the azimuthal velocity and u_x is the axial velocity. The overall phase averaged flame shape was obtained by recording the the flame using an ICCD cammera and then performing an Abel transform of the resulting image to get a 2D slice of the flame. By dividing the viewing window into a top and bottom component, Pallies *et. al.* [105] were able to show that the variation of heat release was due to two independant mechanisms: Swirl number fluctuations and vorticity at the exit of the burner.

The observed minimum and maximum flame response was attributed to the destructive and constructive interference respectively of the effects due to axial velocity fluctuations and swirl number fluctuations. The dependence on frequency of this phenomenon is related to the difference between the convection speed of the azimuthal and the axial velocity perturbations. The axial velocity fluctuations move from the swirler to the injector exit plane at the speed of sound while the azimuthal fluctuations convect at the mean flow velocity [106]. This velocity difference when coupled with the frequency of the incoming wave and distance between the swirler and injector exit corresponds to a phase shift between the two velocity components and a subsequent fluctuation of the swirl number.

The instantaneous swirl number fluctuations normalised by the mean swirl number can be expressed to first order as [106]:

$$\frac{\hat{S}'}{\overline{S}} = \frac{\hat{u}_\theta'}{\overline{u_\theta}} - \frac{\hat{u}_x'}{\overline{u_x}} \tag{1.22}$$

It was shown by Palies *et. al.* [105] that fluctuations in turbulent burning velocity could be linked to the axial and azimuthal velocity fluctuations as:

$$\frac{\hat{S}_T'}{\overline{S_T}} = -0.4\frac{\hat{u}_\theta'}{\overline{u_\theta}} + 0.4\frac{\hat{u}_x'}{\overline{u_x}} \tag{1.23}$$

Attempts to model the swirl flame using a perturbed G-equation by Palies et. al [106] showed that the turbulent flame speed had a similar effect of the flame to the axial velocity perturbations leading to a relation between the swirl flame heat release, axial velocity fluctuations and turbulent flame speed:

$$\frac{\hat{\dot{Q}}'}{\overline{\dot{Q}}} = \mathcal{F}_V(\omega)\left[\frac{\hat{u}'_x}{\overline{u_x}} - \frac{\hat{S}'_T}{\overline{S_T}}\right] \tag{1.24}$$

where $\mathcal{F}_V$ is the flame transfer function for a laminar V-flame. Combining Equations (1.22), (1.23) and (1.24) gives a relation for the heat release fluctuations as a function of the axial velocity perturbations, swirl number and the laminar V-flame transfer function:

$$\frac{\hat{\dot{Q}}'}{\overline{\dot{Q}}} = \mathcal{F}_V(\omega)\left[\frac{\hat{u}'_x}{\overline{u_x}} + 0.4\frac{\hat{S}'}{\overline{S}}\right] \tag{1.25}$$

1.6 Aims and thesis layout

The aim of this thesis is to produce a numerical methodology that can be used to accurately predict the FTF of any lean burn fuel injector. The objectives are to show how each element of the interaction between acoustics and the flame can be accurately and efficiently captured using numerical methods. This includes analysis of the mechanisms of interaction and suggestions as to how designs can be made to be less sensitive to acoustics. Chapter 2 discusses the generation of the computational mesh for each of the methods discussed, while Chapter 3 deals with questions of how the single phase flow should be calculated such that convection and mixing in the combustor are accurately modelled. Chapter 4 deals with the concepts of flow stability and structure. Structures, brought about by flow instabilities can account for a large part of the fluctuating kinetic energy of the flow and therefore can dominate the processes that depend upon fluctuations in velocity. Furthermore these structures can sometimes be predicted using linearised methods which are, at least in theory, significantly cheaper than full-scale simulation. Chapter 5 discusses the effects of the acoustic field on the flow field and discusses the many different mechanisms of possible interaction between them. The effect forcing has on flow structure is also discussed. In order to reduce the costs and increase the accuracy of reacting simulations, in chapter 6 methods that can reproduce the effects of acoustic forcing are discussed for incompressible simulations. Chapter 7 focusses on the fuel spray, atomisation, evaporation and convection processes. It also attempts to draw links between flow structure, spray and combustion. Combustion is discussed in chapter 8, including a comparison of different methods for resolving the flame transfer function. Finally, chapter 9 gives some general conclusions of the work and suggests some future directions. Although this thesis aims to resolve the FTF of a single fuel injector type, each chapter attempts to present a blueprint for how each design feature of the fuel injector can be used to change the flame sensitivity to forcing.

Rather than considering each interaction separately in each subsequent chapter, this thesis is structured such that each chapter builds on the previous one to reinforce understanding. Techniques and methods are recycled as appropriate from one chapter to the next.

Chapter 2

Non-dimensionalisation, geometry and meshing

The computational mesh is a critical component within the development of numerical methodologies for the resolution of thermoacoustic instabilities. The accurate resolution of all time and spatial scales of all of the variables of interest across the entire combustor system is well beyond the capability of today's computers and will remain prohibitively expensive for decades to come. As such the accurate resolution of such quantities of interest requires the clever use of different numerical methodologies to resolve different parts of the problem such that the bare minimum but still sufficient amount of information is resolved in each case. This means that as part of this study, numerical methods tested and used will cover a broad range of turbulence closure methods and different assumptions will be used to limit the complexity of the governing equations. Each of these different methodologies will place an associated requirement on the mesh as the size of characteristic waves or quantities of interest change. This chapter discusses the requirements of the mesh and the different meshing strategies used as well as describing the physical geometries represented. In this thesis, RANS, URANS and LES turbulence closures are used and as such the mesh requirements are described for each closure used.

2.1 Non-dimensionalisation

As this thesis contains operating points and geometries of industrial interest, values have been non-dimensionalised in the following way: Velocities have been non-dimensionalised using the reference velocity defined as:

$$U_{ref} = \sqrt{\frac{2(p_{0,in} - p_{out})}{\rho}} \tag{2.1}$$

where $p_{0,in}$ is the upstream total pressure and p_{out} is the downstream static pressure. Lengths have been non-dimensionalised by the diameter of the fuel injector D. For all other values, quantities are represented as a proportion of the maximum value of the mean. Using pressure as an example:

$$p^* = \frac{p}{max(\overline{p})} \tag{2.2}$$

Times have been in some cases non-dimensionalised using t^* such that:

$$t^* = \frac{D}{U_{ref}} \tag{2.3}$$

2.2 Coordinate axis

The simulation results will be presented in cartesian or polar cylindrical coordinate systems. In the case of cartesian coordinates, the x-direction runs in the direction of the outlet duct from the injector. The y-direction is parallel to the fuel feed arm and z is defined according to the right hand rule. In cylindrical coordinates, the axial coordinate is equivalent to the x-direction in the cartesian case, r and θ are defined according to the right hand rule. The datum is at the centreline of the injector exit, or the point where the injector body meets the downstream duct.

2.3 Geometry

2.3.1 The LU acoustic duct

The acoustic test facility at Loughborough consists of a $1.54D \times 1.54D$ square duct attached to a large plenum at either end with the fuel injector mounted at the upper end of the duct. Air is drawn into the large upper plenum of approximately $105000D^3$ from the atmosphere before passing through the fuel injector, along the duct and out to the lower plenum where the air is returned to the atmosphere via a large centrifugal fan. The fan is set to achieve the desired 3% pressure drop across the injector and duct. Two JBL AL615 600 Watt loudspeakers are mounted towards the lower end of the duct and connected to a Chevin Research A3000 amplifier that can generate acoustic waves over a wide range of frequencies. The amplifier receives its input from a sine wave signal generator that is tuned to the required frequency.

Four Kulite dynamic pressure miniature transducers were installed at 4 different axial locations along the duct at $2.21D$, $5.87D$, $7.01D$ and $9.42D$ downstream relative to the outlet of the injector. Static calibration of these transducers was performed by applying a known reference pressure and measuring the corresponding voltage output. Dynamic calibration was performed by placing each transducer at the same location in the duct and applying the same pressure signal while the amplitude and phases of each transducer was monitored [126].

The output signals from the transducers were digitalised through use of a National Instruments LabView data acquisition system with a 16 bit data acquisition card [7, 127]. The output signals were sampled at 40 kHz in eight blocks of 32,678 samples each. Each block was Fourier transformed using a Hamming window and the corresponding Welch periodogram (see Section A.3.2) constructed by averaging over each of these blocks. This was then used to calculate the auto power spectrum of each signal and the phase angle relative to the Kulite furthest downstream via a cross

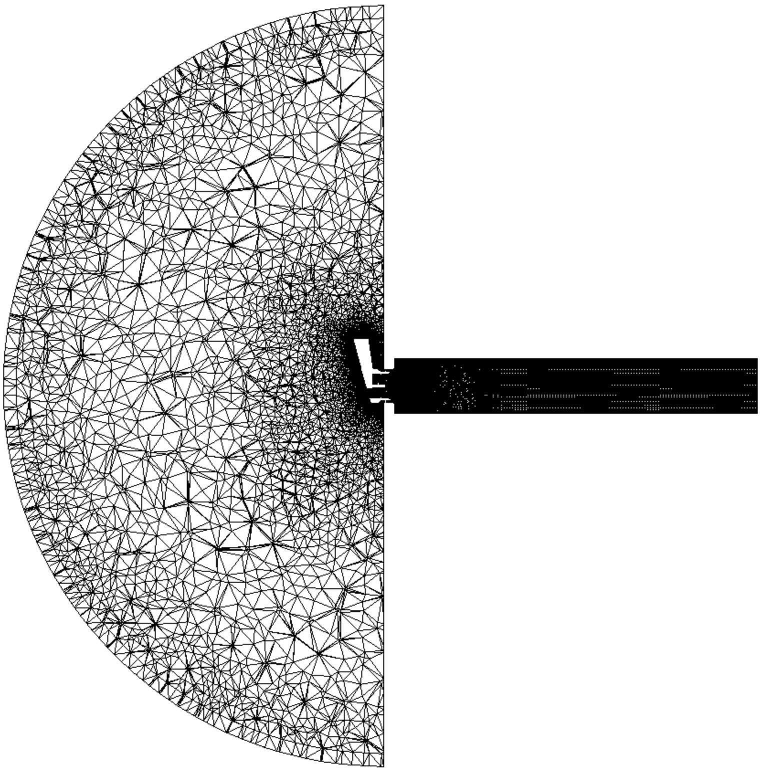

Figure 2.1: Cross section of the LU acoustic duct rig computational domain. The mesh has been generated using the LUBPAI method.

correlation. The multi-microphone technique was then used to derive the acoustic impedance of the injector (see Section 5.2).

In the computational domain the upstream plenum is not modelled exactly and instead replaced with a large hemispherical dome of diameter $20D$. The downstream plenum is not modelled at all and the duct is shortened to $10.26D$. The duct must be long enough to ensure that the approaching acoustic waves are sufficiently planar and that the velocity field is not affected by the imposition of a single static pressure across the exit. The first $0.23D$ of the mesh in the duct is generated directly while the rest is extruded. The cells in the upstream plenum are large such that they act to dissipate acoustic waves and the curved surface of the dome becomes the inlet boundary for the simulation. The cells lying on the inlet boundary are extruded from the surface mesh and are smaller than the other plenum cells. This is to ensure that the gradient of velocity and pressure is calculated accurately for the non-reflective boundary conditions (see Section 5.3). This configuration prevents non-physical standing waves from forming upstream of the fuel injector. Standing waves may form in the downstream duct, however Su [147] calculated a cut-on frequency of well over 1000 Hz which is much higher than the frequencies of interest. An example mesh of the domain is shown in Figure 2.1.

2.3.2 The Rolls-Royce lean-burn fuel Injector

The injector used in this study is a combined (non-radially staged) dual zone fuel injection system for lean burn combustors. The geometry is representative of Rolls-Royce designed modern lean-

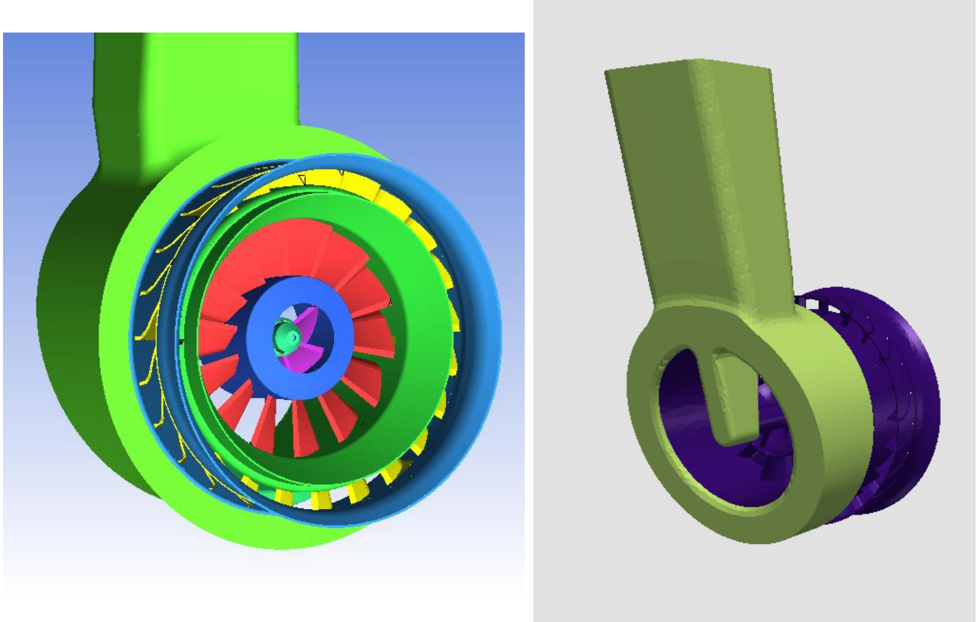

(a) Injector front

(b) Injector back

Figure 2.2: The lean burn injector showing passage A (purple), passage B (red) and passage C (yellow)

burn combustors used in engine demonstrator tests. It provides air to two separate combustion zones (designated pilot and main zones) via three flow passages. Fuel is injected into the pilot zone via a pressure swirl atomiser placed on the rearmost tip of a central bullet within the innermost flow passage designated passage A. The main fuel is supplied via an airblast atomiser located between the two outer flow passages designated as passage B and passage C. Each flow passage includes a set of swirler vanes designed to induce a large enough swirl number in the combustion chamber to create a stable recirculation zone that enhances the burning of the fuel. In the case of the fuel injector used in this study, swirlers in passages B and C rotate in the opposite direction to passage A, this was chosen to improve mixing and increase turbulence levels within the combustor. The injector model used in simulations and experiment includes the fuel feed arm that acts as a connection between the engine casing and the fuel injector in a real engine. The addition of the fuel feed arm breaks the symmetry of the model. This injector was found to exhibit low levels of NO_x production in full-scale tests but unacceptable combustion instabilities. This injector configuration is referred to as *injector I1* and was used within all of the simulations of the LU acoustic duct rig.

2.3.3 SCARLET rig

The SCARLET rig is designed to measure the FTF of real engine fuel injectors. It consists of a round combustion chamber sandwiched between two long ducts. The ducts are fitted with a series of high speed pressure transducers. At the far end of each of the ducts is a set of acoustic sirens. The acoustic sirens consist of a perforated plate that rotates at the desired frequency. When the plate perforation lines up with the inlet to the rig, a jet of high pressure air is injected, resulting in a sound wave being generated. These sirens are chosen over loudspeakers as they are capable

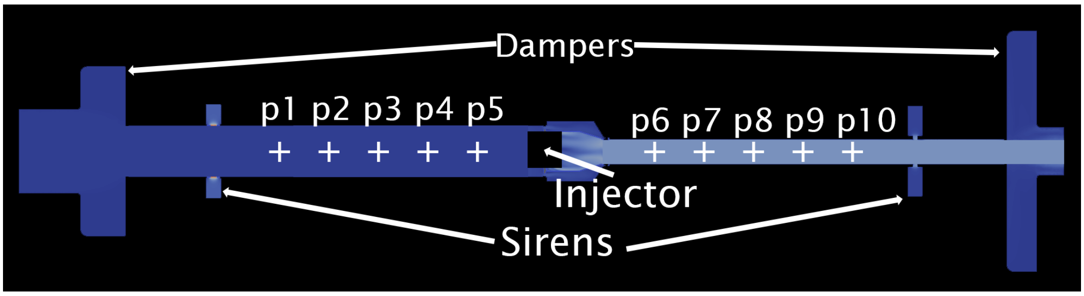

Figure 2.3: The SCARLET rig showing the positions of the injector, dampers, sirens and pressure probes used to analyse the acoustic waves. Flow is from left to right.

of producing extremely high sound pressure levels however they can only operate at one frequency at a time. Beyond the sirens are a set of dampers designed to absorb the sound being produced by the sirens and flame. The inlet flow can be heated to engine representative temperatures and the rig can be pressurised up to 20 Bar. The combustion chamber includes an outer annulus that feeds cool air to the combustor walls. By operating the sirens and recording the pressure fluctuations along the ducts a plane wave decomposition can be performed to analyse the strength of the acoustic waves travelling upstream and downstream (see Section 5.2). This then allows for calculation of the acoustic transfer matrix of the combustion chamber. By comparing the results of this decomposition under reacting and non-reacting conditions a flame transfer matrix and FTF may be obtained. The SCARLET rig simulations and data are achieved at an engine representative pressure, temperature and fuel flow rate. The scarlet rig pressure drop is set to 3.8% but unlike the acoustic duct simulations the mass flow rate is metered and therefore known. To ensure that the correct air to fuel ratio is achieved in the combustion chamber, the mass flow rate of air is set instead of the pressure drop as the upstream boundary condition.

2.3.4 Tip on a stick (TOS) injector

The I1 injector geometry presented above was designed to be used in engine tests. In the case of the SCARLET rig, the combustion chamber consists of a single axisymmetric can and the upstream duct is not large enough to fit a fuel injector with the fuel feed arm attached. The fuel injector used in engine tests is manufactured by welding or brazing sections of sheet metal subcomponents making the design rugged but difficult to modify. The TOS series of injectors are injectors where the fuel arm has been removed for rig testing purposes. The fuel is fed from directly upstream in these injectors, and as such, they remain axisymmetrical. These injectors are also composed of a series of screw-on and slide-on components that allow various subcomponents to be modified quickly such as the angle of the swirl vanes. In this study a TOS injector that will be referred to as *Injector I1A* was used for all of the experiments and simulations of the SCARLET rig. The two injectors are quite similar but contain a number of important design changes, the most important are listed below:

1. The location of the main fuel pre-filmer has been brought inboard slightly, reducing the outer radius of passage B slightly.

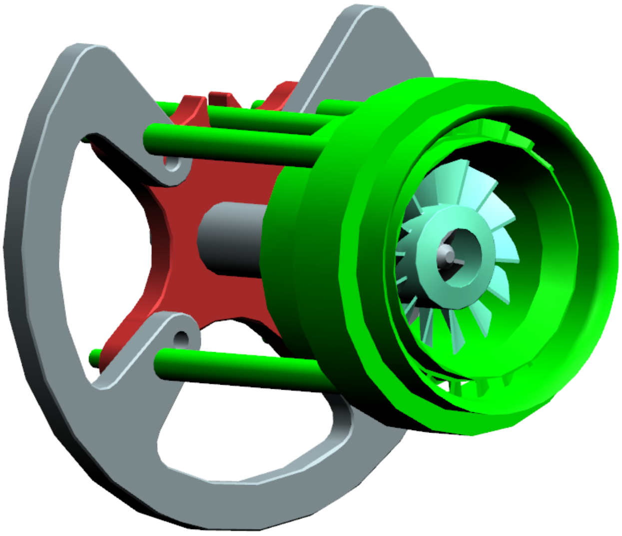

Figure 2.4: The TOS injector I1A.

2. The size of passage C has been reduced along with the outer diameter of the fuel injector.
3. The TOS injector is mounted in the rig with the aid of a large upstream mounting plate (see Figure 2.4). This plate likely affects the passage of acoustic waves in the upstream passage and hence it was decided that it needed to be included in the computational domain.

2.4 Mesh requirements

Meshes were produced using the *ANSYS ICEMCFD* software package. The meshing strategies are based on the following 3 principles:

1. Algorithm stability and accuracy.
2. Accurate resolution of the scales of interest.
3. Minimising mesh size.

The first principle of mesh generation relates primarily to mesh quality, a measure of how close the mesh represents a perfect regular structured mesh with uniform hexahedral elements. This type of cell facilitates the easy and accurate resolution of fluxes and field gradients at the cell boundaries reducing errors and improving stability. The uniformity of cells is also important as very large changes in cell size can also cause instability issues. If this was the only guiding principle, then the mesh would be full of very fine regular cells, however this will lead to extremely large cell counts.

The second principle requires the size of the mesh cells to be small enough to resolve the scales of interest and tends to push the mesh cell count upwards while the third principle aims to lower the cell count. Algorithm stability is method dependent and can act to either increase or decrease the cell count.

The accurate resolution of scales depends on a number of factors however the four most important are:

1. Local Reynolds number.
2. Proximity to unsteady flow structures.
3. Proximity to walls.
4. Turbulence closure used.

The following subsections deal with each of these four points.

2.4.1 RANS

The URANS and RANS methodology use the Boussinesq approximation to add additional numerical viscosity to the flow to produce the same net effect as the unresolved turbulence. In a RANS simulation a steady solution of the ensemble average is sought. All of the unsteady flow structures are assumed to be turbulent eddies and as such can be represented by the turbulence model. RANS tends to fail in one of two situations:

1. The mesh is poorly refined so that the mean gradient of the flow cannot be resolved or
2. That there is inherent unsteadiness in the flow caused by large scale flow structures that are not turbulent eddies. This includes wakes, separations and jets.

These two failures will manifest themselves in different ways: A poorly resolved mesh will converge to the wrong solution, however a mesh that resolves unsteady flow structures accurately will never converge due to the inherent non-turbulent unsteadiness of the flow. The effects of large scale unsteady flow structures on the accuracy of RANS simulations is discussed in Section 4.2.1.

2.4.2 Walls

Close to the wall, a boundary layer will be formed. The law of the wall [117] states that the relationship between streamwise velocity (u_x) and distance from the wall (y) is (in terms of the modified turbulent viscous wall units (y^+ and u^+)):

$$u^+ = y^+ \text{ for } 0 < y^+ < 5 \tag{2.4}$$

and

$$u^+ = \frac{1}{\kappa} ln(y^+) + C \text{ for } 30 < y^+ \text{ and } y << \delta \tag{2.5}$$

where δ is the boundary layer thickness and κ and C are constants. Provided that the first mesh point is placed within this region a wall function can be used that takes into account the effect of the unresolved boundary layer [81]. In the case of attached boundary layers these methods work well, however if the boundary layer separates then the model will not capture the true mean field because of the large unsteady flow structures present close to the wall [20]. In order to capture these effects the unsteady solution must be found.

2.4.3 Incompressible URANS

In this case, the turbulent eddies are not resolved and as such the mesh size is limited by the smallest unsteady flow structures. Unsteady flow structures exist primarily in regions of flow separation such as behind swirler vanes, in the main zone of the combustor and close to any separated wall bounded flows. In these regions the size of the largest structures is set by the geometry, such as the width of a wake or jet. The size of the smallest eddies is inversely proportional to the Reynolds number requiring ever finer meshes to capture them, however at some point these structures begin to resemble turbulent flow and can accurately be taken into account by the turbulence model. Higher Reynolds number flows can become highly turbulent and due to their high kinetic energy can delay or prevent separation allowing larger mesh sizes to be used. Inaccurate capture of these large unsteady structures can lead to an incorrect mean field being calculated as described in Section 4.2.1. The mean flow is important as the quantity of air entering the combustion chamber will affect the air-to-fuel ratio and the temperature of the combustion chamber. The large unsteady structures and their interaction with the acoustic field will also impact on the fluctuating surface area of the flame and rate of fuel air mixing, both of which will alter the heat release rate.

2.4.4 Compressible URANS

Solving the compressible URANS equations results in the generation and propagation of sound waves. In the case of incompressible flow the size of the mesh cells is driven by the need to resolve small flow structures, or pertubations of velocity. Far from the region of interest the mesh can be coarsened as only the slowly varying gradient of pressure must be captured by the mesh. Acoustic waves exist everywhere in the domain and in places where the acoustic waves must be resolved, the mesh must be sufficiently well refined. While acoustic waves have characteristic lengths several orders of magnitude higher than the velocity perturbations, they also travel through the fluid, in most cases, isentropically. This is in contrast to velocity perturbations that decay in time due to viscous forces. When solving the Navier-Stokes equations the (laminar) viscosity of the fluid (ν) is replaced by the effective viscosity (ν_e) which includes the effects of unresolved turbulent fluctuations on diffusion of the mean velocity gradients. For the RANS closure it is [39]:

$$\nu_e = \nu + \nu_t \tag{2.6}$$

where ν_t is the turbulent viscosity. Although not explicitly added, the effective viscosity of the fluid in the simulation includes a term relating to the ratio of the acoustic or velocity perturbation wavelength to the cell size and the order of the numerical scheme [44]. This additional viscosity does not cause great problems for the velocity field as the turbulent viscosity is relatively high, which tends to mask this numerical viscosity term. Acoustic waves should not be damped by the mesh and as such considerably more cells per wavelength must be used, increasing the size of the mesh. The effect of mesh size and the order of the numerical scheme is discussed in Appendix B. In places where the propagation of acoustic waves should be avoided, the size of the mesh cells can be increased to damp out the acoustic waves using the numerical viscosity. The addition of compressibility also makes the solvers used in this thesis more sensitive to large jumps in cell size. For this reason, meshes in the compressible domain are usually generated such that they are tetrahedral cell dominated with less variation between largest and smallest cell size. This increases the cost of compressible simulation. Su [147] generated the downstream mesh by extruding cells of no more than $0.046D$ and found this to be sufficient to prevent unwanted dissipation of sound waves over the frequencies of interest.

2.4.5 Incompressible LES

The Large Eddy Simulation (LES) method used in this thesis replaces the RANS transport equations for turbulent quantities with a simple equation relating the local strain rate of velocity to the local size of the mesh (see Section 3.2.8). LES assumes therefore that the unresolved scales of turbulence act locally only. A common estimate at quantifying this requirement has been developed by Pope [117] which states that a well-resolved LES should resolve more than 80% of the turbulent kinetic energy (k) or TKE. Dianat et al. [33] used an approximation of the turbulent dissipation rate (ϵ) developed by Moeng and Wyngaard [92] to approximate the integral length scale (l_c):

$$l_c = \frac{k^{3/2}}{\epsilon} \tag{2.7}$$

where,

$$k = k_{res} + k_{sgs} \tag{2.8}$$

and,

$$\epsilon = 0.93 \frac{k_{sgs}^{3/2}}{\Delta} \tag{2.9}$$

where k_{sgs} is the unresolved TKE, k_{res} is the mesh resolved TKE and Δ is the mesh spacing. These relations can be combined to show that if $k_{sgs} \leq 0.2k$ and in order for the mesh to be well resolved the following inequality should stand:

$$\frac{l_c}{\Delta} > 12 \tag{2.10}$$

Practically, the length scale can be approximated from RANS calculation. For RANS based on the $k-\epsilon$ model, k and ϵ can be used directly in equation 2.7 while simulations using the $k-\omega$ model must make use of the relation $\epsilon = k\omega$. For unsteady simulations the resolved TKE must be calculated from the fluctuating velocity fields, $k_{res} = 1/2R_{ii}$ where R_{ij} is the ij component of the Reynolds stress tensor:

$$R_{ij} = \overline{u'_i u'_j} \tag{2.11}$$

and u' is the fluctuating component of the Reynolds decomposition:

$$u = \overline{u} + u' \tag{2.12}$$

For LES, an estimation of the dissipation rate (ϵ) can be obtained by computing the energy spectral function ($E(\kappa)$)(see Section 3.2.8) and computing the integral [117]:

$$\epsilon = 2\nu_t \int_0^{\kappa_c} \kappa^2 E(\kappa) d\kappa \tag{2.13}$$

An estimation of k_{sgs} can then be obtained using Equation (2.9).

RANS simulations described in Chapter 3 show that within the main combustion chamber the length scales are around 0.1D to 0.5D, this requires therefore, a mesh size of no more than $0.008D$. The swirler passages require a much higher resolution, there are length scales down to around $1 \times 10^{-5} D$ which would require a mesh spacing of around $8.3 \times 10^{-7} D$. This number is very low due to the presence of boundary layers. LES is known to perform badly in close proximity to walls [81, 117], for this reason it is recommended that the grid spacing should be no larger than $\Delta x^+ < 150, \Delta y^+ < 1, \Delta z^+ < 20$ expressed in viscous wall units [112]. The y^+ values in the injector calculated at one instant of a URANS simulation shows y^+ values of between 7 and 25. At the lower end of the range, the cell spacing has been set to $0.0032D$ while at the upper end it is $0.0064D$. In order to guarantee a constant $y^+ < 1$ the mesh would need to be refined by around a factor of 10 or prism layers could be employed. The practical use of prism layers for such complicated industrial geometries is very rare and as refining the mesh further resulted in prohibitively large mesh sizes, the mesh was not refined beyond this level. Hopefully with improved meshing algorithms and more powerful computers, fully refined LES for complex geometries will become more common. The failure of LES methods to capture boundary layers correctly [63] has led to a series of hybrid RANS-LES methodologies [9] such as the Scale Adaptive Simulation (SAS) method [84] or the Detached Eddy Simulation (DES) method [141].

2.5 Meshing Strategies

2.5.1 Rolls-Royce best practice for single phase lean burn combustor geometries (RRBP)

The starting point for the development of a computational meshing strategy was the Rolls-Royce best practice methodology (RRBP) for lean burn single phase combustor geometries [6]. Density boxes are applied to limit the cell size in the regions of interest. The mesh is generated using tetrahedral elements which are then converted to hexahedral cells after smoothing.

By varying the cell size within each density region, a number of different but related meshes can be generated of varying sizes. Table 2.1 shows the densities used and total cell size for each mesh generated. The density boxes used are placed as follows:

1. Around the passage A path.
2. Around the injector body.
3. Around the main body of the injector and fuel feeder arm.
4. Within the primary zone of the combustion chamber or duct.

5. Within the downstream section of the duct.

Several of meshes generated by this method can be seen in Figures 2.5, 2.6 and 2.7.

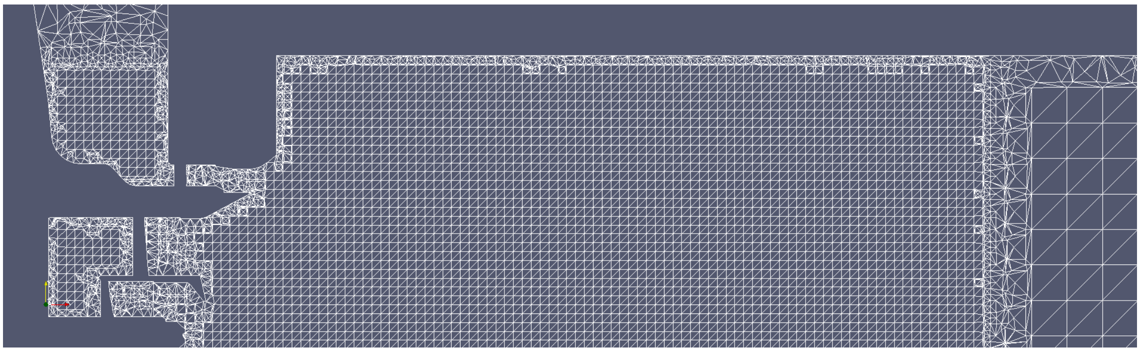

Figure 2.5: A cross-section of mesh 4.

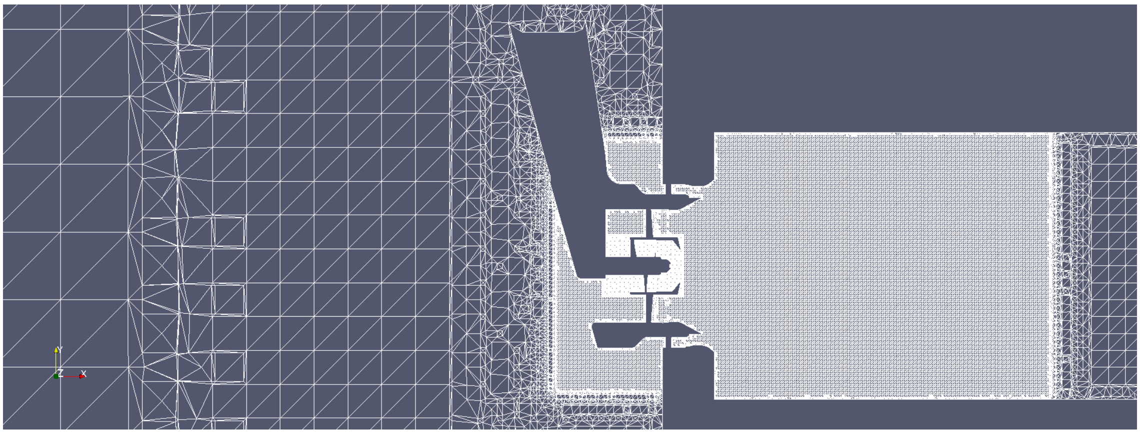

Figure 2.6: A cross-section of mesh 15.

2.5.2 Loughborough University best practice for acoustic impedance (LUBPAI)

In previous studies conducted by Su [147] the generation of the mesh was performed within ICEM-CFD using a very different methodology. The Rolls-Royce methodology tends to produce extremely efficient meshes in the sense that the use of density boxes and and conversion to hexahedral elements produces a mesh with a small number of cells of high quality with small cells located only where they are required. However this efficiency leads to large changes in cell size at the extremity of the density boxes. It has been observed that this is especially detrimental to the stability of the unsteady compressible flow solver, leading to large, high frequency oscillations in the pressure field. Because of this, the Loughborough University best practice for the measurement of acoustic impedance (LUBPAI), as developed by Su et al. [149], was used for the resolution of the acoustic impedance. Firstly, a triangular surface mesh is generated and smoothed and then a Delauney tetrahedral mesh grown from the surface mesh. Prism layers are then added at the duct walls and the mesh smoothed. This methodology produces meshes that are extremely continuous

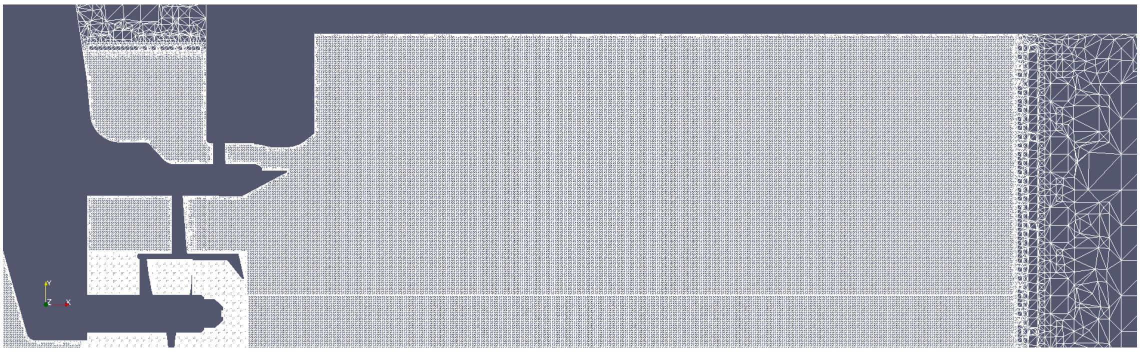

Figure 2.7: A cross-section of mesh 16.

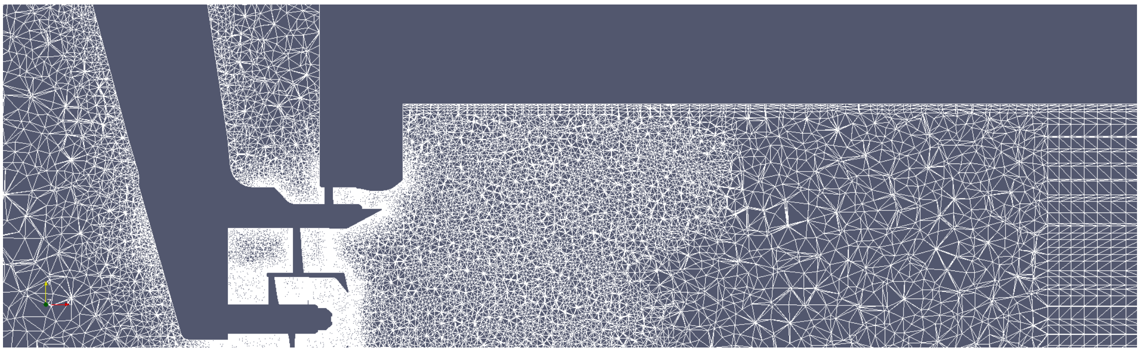

Figure 2.8: A cross-section of the 9 million cell mesh generated using the LUBPAI method.

Table 2.1: Meshes generated using the Rolls-Royce best practice methodology

Mesh number	Global size/D	Pilot density/D	Main swirler density/D	Primary zone density/D	Outlet density/D	Number of cells
2	1.3	.0032	0.0064	0.064	0.26	10×10^6
3	2.6	0.0064	0.013	0.13	0.51	2×10^6
4	1.3	0.013	0.064	0.064	0.13	1×10^6
8	0.12	0.0064	0.0064	0.013	0.026	7×10^6
9	0.12	.0032	0.0064	0.0064	0.026	40×10^6
11	1.3	.0032	0.0064	0.013	0.077	19×10^6
12	1.3	0.013	0.064	0.064	0.13	4×10^6
13	1.3	0.0064	0.013	0.064	0.13	5×10^6
15	1.3	0.0064	0.013	0.064	0.13	7×10^6
16	1.3	.0032	0.0064	0.0064	0.13	41×10^6
23	0.82	0.0064	0.0064	0.0064	0.013	69×10^6

with smooth variations of cell size across the domain (see Figure 2.1) however these meshes are considerably larger in cell count as compared to RRBP methodology.

2.5.3 Rolls-Royce Deutschland methodology for reacting flow

Figure 2.9: Diagram highlighting the main features of meshes generated using the RRBP and LUBPAI meshing strategies and how the *worst aspect ratio* parameter can be used to move the RRD meshes between these two extremes.

The RRD methodology is preferred by some of the engineers at the Rolls-Royce site in Dahlewitz in Germany. The methodology is similar to the RRBP methodology except that mesh coarsening is employed during the generation of the octree mesh where the number of iterations of mesh smoothing and grid coarsening can be selected in ICEMCFD. The mesh coarsening works by absorbing small cells into neighbouring ones. As the number of coarsening and smoothing operations are increased, the uniformity of the octree mesh is degraded and the mesh begins to resemble a smooth delaunay mesh similar to those produced using the LUBPAI meshing technique. The number of coarsening steps is controlled by the user but the effectiveness of each coarsening iteration is controlled by the *worst aspect ratio* parameter. A large value will only allow cells of very regular proportions to be formed in the coarsening process while a very small value will allow the cell quality to be degraded during the process, allowing more coarsening to take place. Using this parameter the user can effectively control how similar the mesh is to the RRBP mesh type or how similar it is to the LUBPAI mesh type. The advantage of this method over the LUBPAI method is that it utilises the significantly more robust octree meshing method. The flexibility of the method is illustrated with two meshes of the SCARLET geometry. The first is a mesh of the entire combustion chamber and upstream and downstream ducts. In this case, a smooth mesh was important to ensure stability of the compressible algorithm, and as such, the *worst aspect ratio* was set quite low to ensure that plenty of coarsening was performed. The mesh density was set to between $0.0064D$ and $0.072D$ resulting in a mesh of 5 M cells. The mesh is visualised in Figure 2.10a. The next mesh (Figure 2.10b) is of the same geometry but with the fuel injector and upstream duct removed. The mesh density is reduced to $0.0064D$ with some places down to $0.0032D$. This mesh is for high fidelity LES calculations assuming incompressible flow. Due to the much simpler geometry and the less sensitive numerical method the cells can be more uniform and a higher value of *worst aspect ratio* was employed to reduce the amount of coarsening. The tetra-to-hexa cell conversion reduced the cell count from ≈ 120 M cells to ≈ 25 M cells resulting in a very high quality and uniform mesh.

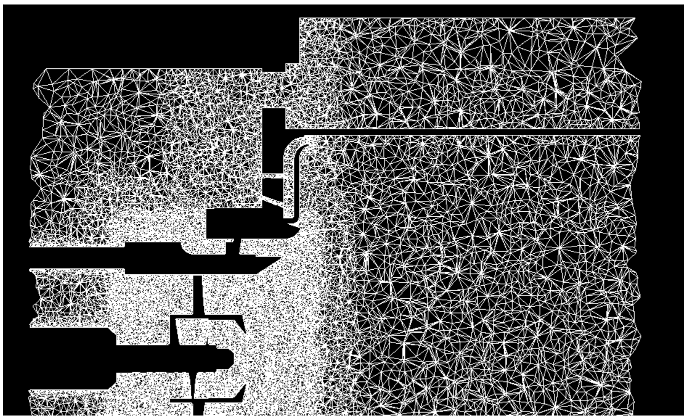

(a) small *worst aspect ratio*

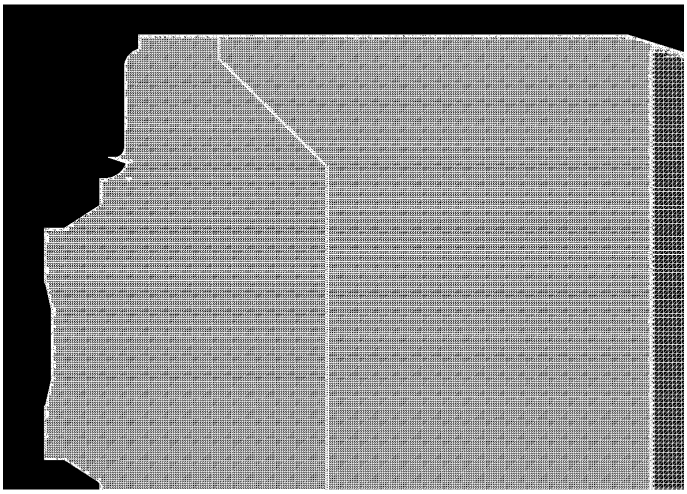

(b) large *worst aspect ratio*

Figure 2.10: The effect of *worst aspect ratio* on the mesh.

Chapter 3

Computations of single phase flow

Before the acoustic response of the flame can be measured it is prudent to ensure that the numerical tools chosen are capable of resolving the simple non-reacting flow fields produced by the chosen injector geometry. This chapter explores the differences in the mean field statistics between unsteady and steady simulations, LES and RANS, and the effects of compressibility.

3.1 Numerical Methodology

The governing equations are solved using the finite volume method. The equations and methods are described in detail in Appendix A.1. Second order accuracy in space is provided by using second order estimates of the flux and gradients at the boundaries of each cell. In the case of the pressure, a symmetric stencil is used, however for velocity, a symmetric stencil may lead to the generation of instabilities and as such a linear upwind stencil is used. Upwinding computes the flux and flux gradient by only considering the points that are upstream, as defined by the local velocity. In certain cases even this approximation can become unstable in which case 1st order upwinding is used instead, or a blending factor is used that blends the two methods together to achieve an order of accuracy and stability between the two. TVD schemes [47] use a blending function unique to each cell allowing for the order and stability of the scheme to vary across the domain.

3.1.1 OpenFOAM

OpenFOAM [104] (Open source Field Operation and Manipulation) is a collection of tools written in C++ to solve for the solution of continuum mechanics problems, primarily CFD. The code is free, open-source and based on the finite volume method. It contains far too many features to list here but several of the solvers have been used in this body of work and the relevant functions are described in this thesis.

OpenFOAM simulations have been run using the PISO algorithm or PIMPLE algorithm which is a combination of SIMPLE and PISO. These algorithms are described in Appendix A.1 and by Ferziger and Peric [39]. The OpenFOAM solver used for the PIMPLE algorithm is *PimpleFoam* or *rhoPimpleFoam* for compressible flows and can be used for steady and unsteady flows, while in the case where the number of outer iterations are restricted to one with no under-relaxation, the method is refered to as *rhoPisoFoam* and can only be used for unsteady flows. It was has been found by Su [147] that the *rhoPisoFoam* solver is more accurate for acoustically forced simulations than the *rhoPimpleFoam* solver due to the lack of under-relaxation of either velocity or pressure. As such the *rhoPimpleFoam* solver was only used to find the unsteady swirling bias flow of the injector while the acoustically forced simulations were computed using *rhoPisoFoam*.

3.1.2 PRECISE-UNS

PRECISE-UNS (Predictive system for Real Engine Combustors with Improved Sub-models and Efficiency - UNStructured) is Rolls-Royce's in-house combustor analysis software. Originally formulated to be used with block structured grids and incompressible flows, it has been given the capability to perform fluid simulations on unstructured grids in incompressible and compressible flows. Coupled with the base solver are various different kinds of turbulence closures such as k-epsilon, k-omega, SAS and LES methods.

The code is a cell-centred finite volume method that has up to 2nd order accuracy in space and time. Six convection schemes are available including the first or second order upwind method, central differencing and two total variation diminishing (TVD) schemes. Temporal schemes are restricted to the 1st order backward euler scheme and a 2nd order backwards quadratic method. Further information on the code can be found in [4]. PRECISE-UNS is written primarily in FORTRAN 90 and is based on the *dolfyn* solver [28]. PRECISE-UNS uses a method equivalent to the PIMPLE method and may be forced to conduct extra pressure correction steps or as many iterations as required by the user.

It has been noted in [39] that the advantage of the PISO method is that the additional pressure correction steps means that the pressure equation does not need to be under relaxed and equivalently in the simulations run using OpenFOAM it has been found by Su [147] that the velocity and pressure need not be under-relaxed between timesteps either. This allows the simulations to be performed with only one iteration per timestep. Despite the similarity of method, simulations using PRECISE-UNS had to be run with more that one iteration per time step. This reduces the amount of time required to perform simulations using OpenFOAM versus PRECISE-UNS despite a slower per iteration computation time due to the use of C++ instead of FORTRAN.

3.1.3 Compressibility

In order to resolve a steady or unsteady compressible or reacting flow, an additional equation describing the evolution of the flow total energy (E_h) or an alternative form must be solved. OpenFOAM uses as default the equation for total energy, however the source term due to viscous heating is neglected due to the assumption that the Mach number is low (see Appendix A.1.5). Garnier et al. [41] have identified that the total energy equation must be used if the equations are to remain conservative. This is important in the case of high Mach number flows where discontinuities exist. As the Mach number in the combustion chamber is significantly less than one, the total energy equation may be replaced with the equation for specific enthalpy h which is useful from the point of view of the integration of pre-existing combustion models in PRECISE-UNS.

Su [147] has modified the PISO algorithm in OpenFOAM [39] to be similar to the algorithm proposed by Issa [53] where the energy equation is solved within the PISO pressure correction loop. The same approach was attempted within PRECISE-UNS, however initial results did not show any benefit when compared to the SIMPLE approach where the energy equation (specific enthalpy) is solved only once per iteration.

3.1.4 Unsteadiness

PRECISE-UNS and OpenFOAM have a number of different methods that can be used for resolving the unsteady flow equations. A second order quadratic method (see Appendix A.1.6) was preferred and employed in all simulations except where additional stability was required, in which case, the backwards Euler method was used to damp out some of the high frequency components of the pressure oscillations leading to smoother results and better convergence.

3.1.5 Turbulence Closure

In order to close the Reynolds averaged or filtered Navier-Stokes equations a turbulence closure method must be chosen. For RANS and URANS simulations two models were used. The first is the $k - \epsilon$ RNG model [170] as it is a commonly used model in industry and the prefered model used in Rolls-Royce. The second model tested was the $k - \omega$ SST model which has been used by Su et al. [147, 149, 150] in simulations to derive the acoustic impedance of orifices and simplified injector geometries.

It has been assumed that URANS is not sufficiently accurate to be able to efficiently capture the interactions between the acoustic field and unsteady combustion based on the work of Dunham et al. [36] and therefore an LES closure method is also used and compared to the URANS results. The LES sub-grid scale model used is the constant Smagorinski model chosen for its simplicity and robustness. Turbulence closure methods are further described in Appendix A.1.7.

3.2 Results and Discussion

3.2.1 PIV

PIV (particle image velocimetry) was used to characterise the flow of Injector I1 by Abo-Serie *et al.* [1]. Very little detail was provided in this report on the experimental procedure and the results are not easily understandable, however the mean flow field and Reynolds stresses in a slightly different flow geometry and at a 2% pressure drop are included here as a reference.

Experimental methodology

The PIV was conducted in a square section duct with a width and height of $1.79D$, which is slightly wider than the $1.54D$ of the Loughborough acoustic duct rig. The boundary conditions were also different, with a pressure drop of 2% being imposed. Both of these changes limit the usefulness of the results for direct comparison with the CFD however despite the alterations the normalised results are surprisingly similar, at least for the mean flow.

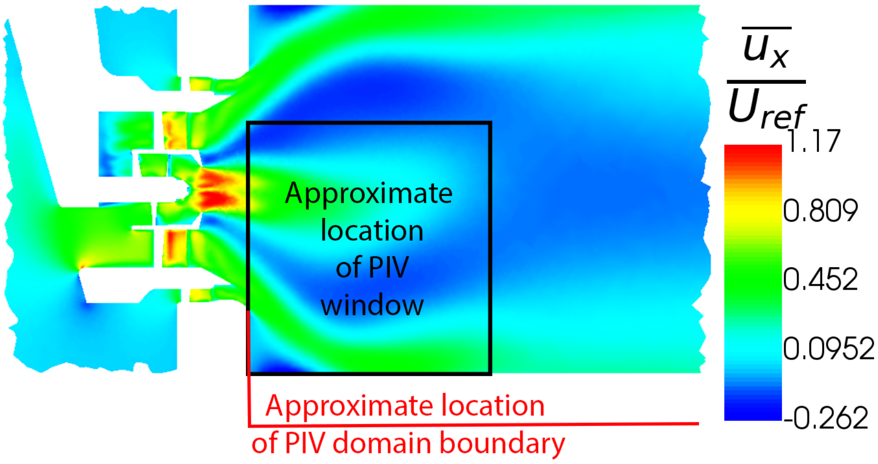

Figure 3.1: The location of the PIV window superimposed upon the velocity field for the compressible OpenFOAM simulation.

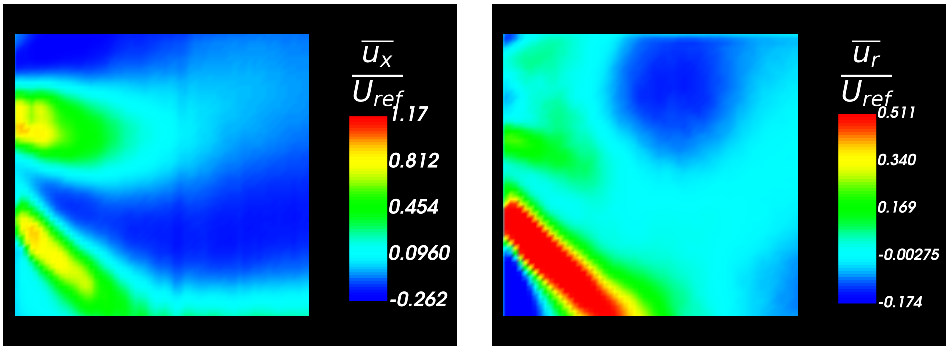

(a) Mean velocity in the axial direction from PIV.

(b) Mean velocity in the radial direction from PIV.

Figure 3.2: The PIV results for mean velocity adapted from Abo-Serie et al. [1].

PIV works through the injection of a passive additive to the flow that forms small particles that when illuminated with a laser sheet can be viewed using a digital camera. Two exposures are taken using either a fast shutter or a pulsed laser to freeze the flow at two different times. The resultant image pair is then broken into series of interrogation windows that are usually a square block of pixels with a resolution that corresponds to a factor of 2 (16x16 pixels, 32x32 pixels, etc.) Each interrogation window then undergoes a 2-dimensional cross-correlation between the two images [119]. This is usually undertaken using a fast-Fourier transform, for which, the speed is maximised for windows of the afore-mentioned sizes. The cross-correlations will usually contain a peak value that corresponds to the distance the passive particles have moved between the two image exposures, therefore allowing the velocity of the particles to be calculated. Providing the Stokes number of the particles is sufficiently low and other body forces are minimised (buoyancy, Coriolis etc.), the particles are a good representation of the flow itself [119]. Basic PIV gives a two-dimensional field of velocity vectors with two velocity components. This can be improved with the addition of a second camera which can resolve the third velocity component (stereo PIV) or many additional cameras that can then resolve a three-dimensional interrogation volume (tomo-PIV). The approximate location of the PIV field of view can be seen in Figure 3.1.

Whilst the concepts are simple, there exists many nuances to the experimental process, pre-processing and post-processing of results, none of which have been made available to the author. The information available is that the results were taken at a frequency of 15 Hz and 1000 images were used to form the averaged statistics presented here. The resolution of the PIV system resulted in $.017D$ per vector. The mean axial and radial velocities are shown in Figure 3.2.

3.2.2 Incompressible RANS

It seems likely that in some cases the flow field within the combustion chamber is too complex for a steady simulation to be able to resolve the flow accurately. Fundamentally, the steady simulation forces all of the natural flow unsteadiness to be taken account by the turbulence model. Turbulence models are calibrated against simple known solutions where the turbulence and flow can be assumed to be homogeneous and if not isentropic, then axisymmetric [81]. Swirling flows have large unsteady structures that are far from homogeneous turbulence [146]. Simulations using the steady, incompressible methodology typically overestimate the pilot penetration length in combustor simulations as evidenced in Figure 3.3a as compared to Figures 3.1 and 3.2. Despite the limitations of the method, it is still used extensively for preliminary design of fuel injectors in Rolls-Royce where URANS or LES simulations remain too expensive. Although inaccuracies are large in absolute terms, RANS simulations appear to be able to capture general trends in aerodynamic and emissions performance which can then be verified using experiments [143]. As experimental validation is required even in the case of an LES simulation, a RANS simulation can provide the required information at a substantially lower cost.

Whilst the accuracy of the steady flow simulations is questionable, they also serve the useful purpose of getting the flow to a near convergence point before unsteady simulations are started. The PIMPLE methodology used with sufficient under relaxation allows for rapid convergence of the RANS equations to a point where the URANS simulations can be started when the flow statistics are almost converged. All of the steady simulations presented here are iterated 4000 times before

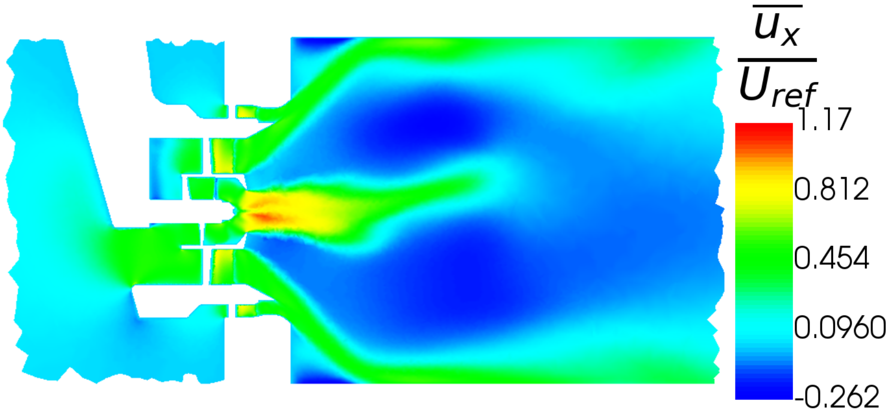

(a) The LUBPAI mesh.

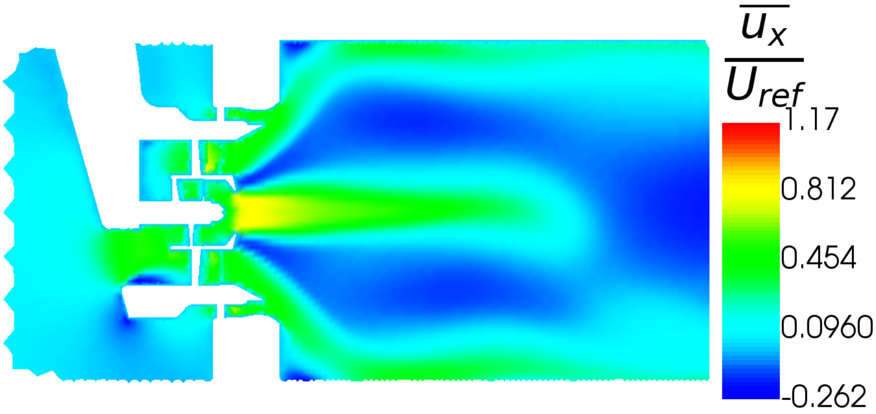

(b) Mesh 4.

Figure 3.3: The axial velocity calculated with two different meshes assuming incompressible steady flow and the PRECISE-UNS solver.

the simulation is allowed to become unsteady.

All of the PRECISE-UNS simulations attempted to use the PRECISE-UNS best practices for isothermal CFD modelling of lean burn combustors [6] as provided by Rolls-Royce. Relaxation factors were set to 0.5 for velocity and 0.3 for the other variables. The gradient of the velocity variables were limited using a Van Albada limiter [155]. The convection scheme was set to being a second order upwind method for velocity and first order upwind for the other variables. The velocity scheme was blended slightly towards a first order upwind method to maintain stability.

Steady incompressible PRECISE-UNS simulation using the LUBPAI mesh and the $\mathbf{k-\epsilon}$ RNG model

This simulation was performed using the same mesh as that used in Chapter 5 which deals with the acoustic response of the injector and therefore serves as a useful comparison between the methods used. In comparison to the unsteady computations the RANS method overestimates the pilot penetration length and underestimates the peak axial velocity as can be seen in Figure 3.3a. The contour can also be seen to exhibit a large degree of asymmetry, a symptom of the highly unsteady flow being restricted to a steady flow solution.

Steady incompressible PRECISE-UNS simulation using mesh 4 and the $\mathbf{k-\epsilon}$ RNG model

This mesh was developed in order to test the numerical methods employed in this study. Its relatively small cell count made the solution quickly obtainable. In comparison to the above simulation that uses over twice as many cells, the simulation shows a even more over-predicted pilot jet length and under-predicted peak velocity (see Figure 3.3b). The CRZ is predicted to have the same magnitude of velocity but it is also predicted to take up a smaller volume. The cone angle is predicted to be wider and therefore the ORZ is predicted to be smaller. The simulation looks to be quite well converged and steady, the large cell sizes work well in damping out any unsteadiness with a penalty for accuracy.

3.2.3 Compressible RANS

The compressible solver in PRECISE-UNS was still in development at the beginning of this project. Previous to this project, the code was extremely sensitive to mesh quality with all but the simplest and highest quality meshes leading to divergence of the flow solution. Even with the improvements to the code it must be noted that compressible methods appear to be considerably more mesh sensitive than incompressible. This includes large changes in cell size that can occur with the Rolls-Royce best practice method. The majority of meshes built with this method were hoped to be used in forced compressible unsteady simulations but failed to converge in either compressible steady or compressible unsteady simulations. Just one mesh of this type, mesh 4, converged for both compressible steady and unsteady flow. The LUBPAI mesh generated by Jialin Su fared much better due to the Delaunay algorithm generating a very homogenous mesh. Whilst mesh 4 was capable of providing a converged solution with PRECISE-UNS under unforced compressible conditions, it failed to converge in acoustically forced OpenFOAM simulations. This highlights the

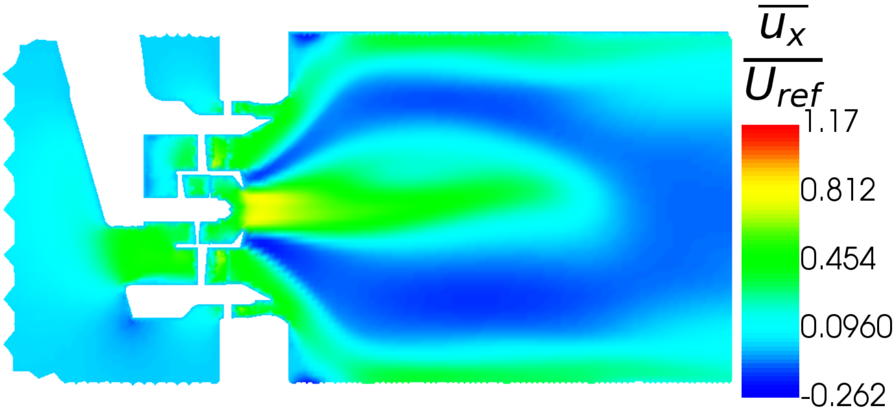

Figure 3.4: The axial velocity for the compressible simulation using Mesh 4 and PRECISE-UNS.

importance of the mesh quality for the compressible regime even across different computational codes.

Steady compressible PRECISE-UNS calculation with Mesh 4 and the $k - \epsilon$ RNG model

This simulation was run to provide a starting point for unsteady compressible simulations with the low resolution mesh 4. The compressible flow result for axial velocity visualised in Figure 3.4 can be seen to be very similar to the incompressible results in Figure 3.3b, except that the maximum flow velocity has been suppressed further in comparison to the unsteady flow results and the PIV.

3.2.4 Incompressible URANS

All of the incompressible URANS equations were started from the steady state flow result and run for 14000 timesteps according to Rolls-Royce best practice [6]. The time step was set to 2.0×10^{-5} s and the second order backwards quadratic method for temporal integration. The maximum number of iterations per timestep was set to 20 and the target residual set to 10^{-5}. The flow was then averaged in time using every 20th data field. Due to computational resource restrictions the average was sometimes computed by linearly interpolating the field onto the low resolution mesh, mesh 4, at each time step and then averaging in time. This was achieved through use of the *vtk* package [133]. Figure 3.5 shows the axial velocity computed using three different meshes.

Unsteady incompressible PRECISE-UNS calculation with mesh 4 and the $k - \epsilon$ RNG model

The low resolution mesh 4 was used primarily for testing. In comparison to higher resolution meshes and the PIV data it shows lower peak velocities and a longer pilot jet.

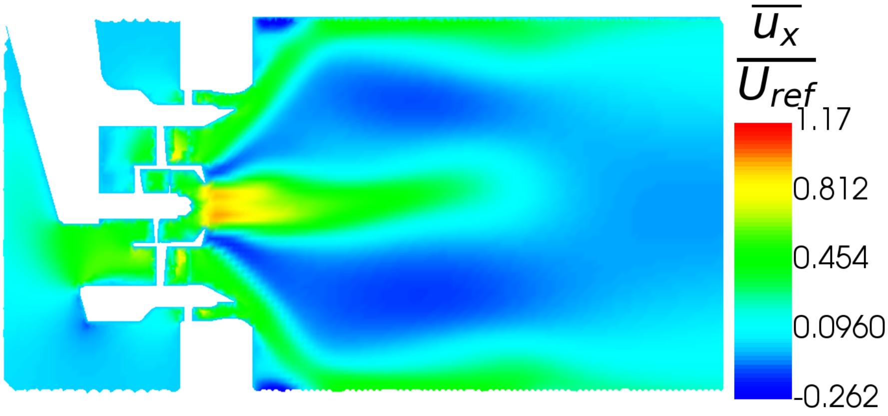

(a) Mesh 4.

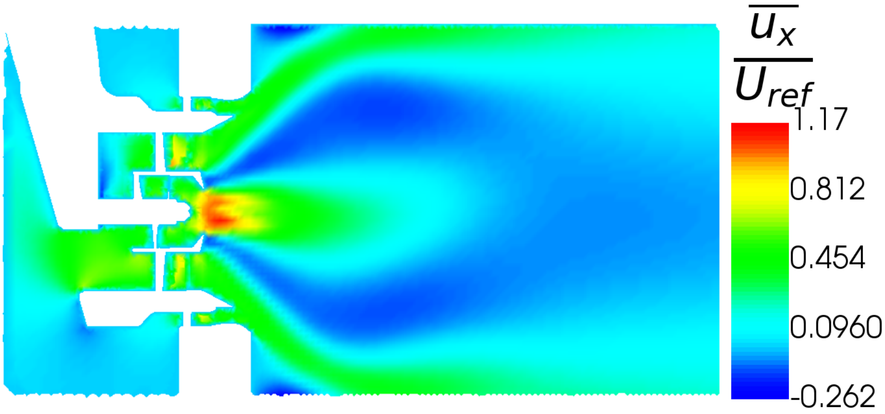

(b) The LUBPAI mesh.

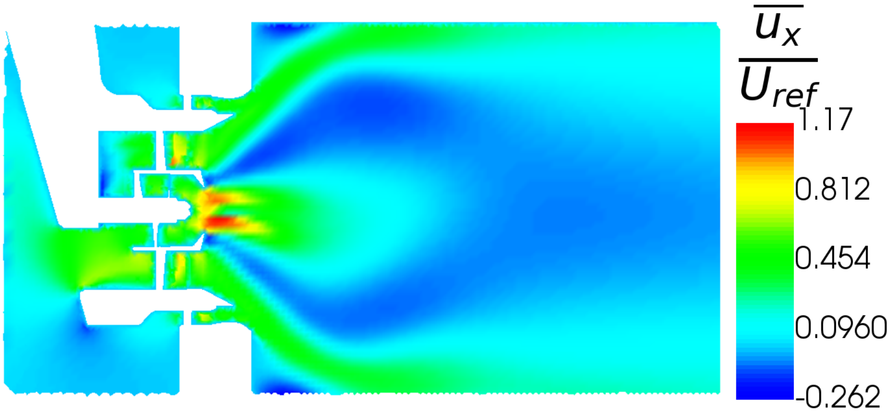

(c) Mesh 16.

Figure 3.5: The axial velocity produced using 3 different meshes assuming unsteady, incompressible flow and PRECISE-UNS.

Unsteady incompressible PRECISE-UNS calculation with the LUBPAI mesh and the $\mathbf{k}-\epsilon$ RNG model

This simulation was run to compare the results obtained using incompressible PRECISE-UNS with those obtained using compressible OpenFOAM. The peak velocity seems to be around 10 percent lower in this case. This may be due to the numerical method, compressible effects, turbulence model or a considerably larger timestep as the compressible simulations were run with a time step of 10^{-6} s.

Unsteady incompressible PRECISE-UNS calculation with mesh 16 and the $\mathbf{k}-\epsilon$ RNG model

Mesh 16 was an attempt to produce a mesh of a resolution similar to that proposed in [6] but without the large jumps in cell size that caused problems with previous meshes. The results show a small improvement with regards to the LUBPAI mesh results with a higher peak velocity and shorter pilot jet. The mesh has over 40 million cells which is over four times more than the LUBPAI mesh making simulations with this mesh very expensive in comparison. The results were in this case interpolated onto the coarser mesh for averaging and therefore some of the benefit of this mesh may be understated due to the effective spatial filter that has been applied.

3.2.5 Compressible URANS

The compressible unsteady OpenFOAM simulations using the $\mathbf{k}-\omega$-SST model

This simulation was performed by Jialin Su using OpenFOAM version 2.3.1. The (pseudo)-steady solution was first obtained using the *rhoPimpleFoam* solver with a timestep of 1×10^{-4} s using a first order backwards scheme. The timestep was then reduced to 1×10^{-6} s and run for a total of 45000 time steps using the PISO algorithm of Issa [53] with a second order backwards scheme. The full flow field was saved every 200th timestep and used to form the average flow field as visualised in Figure 3.1. Spatial discretisation was achieved through use of a second order TVD [47] scheme for velocity and a first order upwind method for the turbulence and energy variables. The peak axial velocity is around 10% higher than in the incompressible case using PRECISE-UNS.

The mean velocity profiles show good similarity for the mean velocity. The peak axial velocity is almost the same for the PIV and simulation located in the outlet of passage A. This pilot flow jet appears to be slightly extended in the case of the simulation. The simulation mains jet is pushed inwards by the wall resulting in a lower velocity in the region and a smaller inner recirculation region. The outer recirculation region is not apparent in the experimental case. The velocity in the core of the mains jet is significantly higher in the PIV measurements. The radial velocity profile is slightly lower in magnitude in the case of the simulation, this appears to be primarily due to the cone angle of the flow being higher in the experimental case. How much of this is due to the difference in domains is hard to judge however the change in pressure drop is likely to change the cone angle observed. The simulation shows that the main flow is being drawn back into the CRZ, whereas the PIV seems to suggest that this is happening further downstream, again this may be a consequence of the domain geometry.

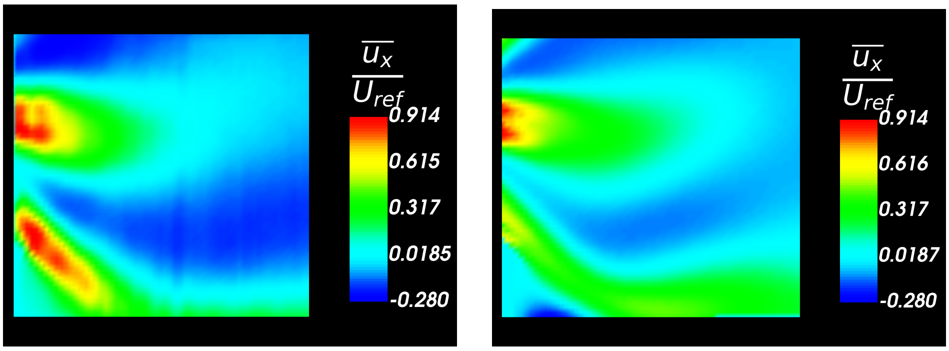

(a) Velocity in the axial direction from PIV.

(b) Velocity in the axial direction from the compressible OpenFOAM simulation.

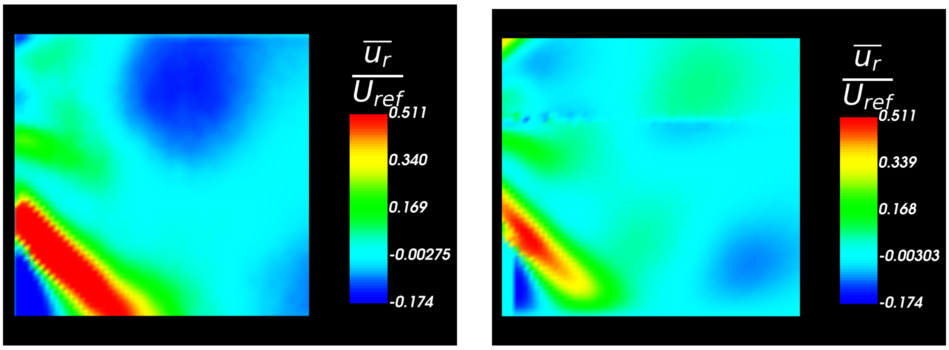

(c) Velocity in the radial direction from the PIV.

(d) Velocity in the radial direction from the compressible OpenFOAM simulation.

Figure 3.6: The PIV results for mean velocity adapted from Abo-Serie et al. [1] compared to the Compressible OpenFOAM simulation.

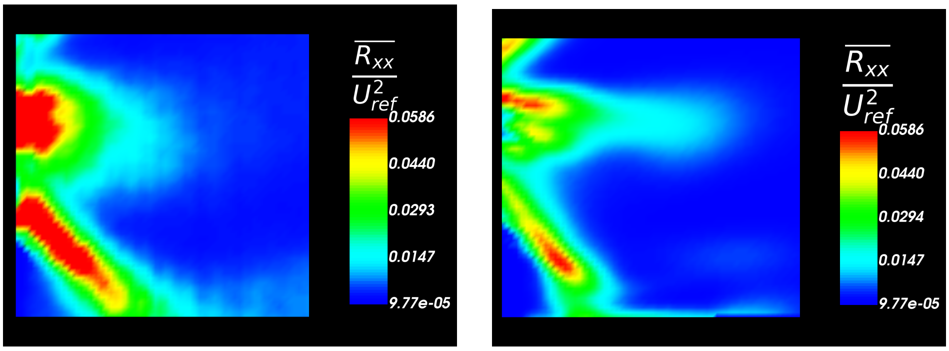

(a) R_{xx} from PIV.

(b) R_{xx} from the compressible OpenFOAM simulation.

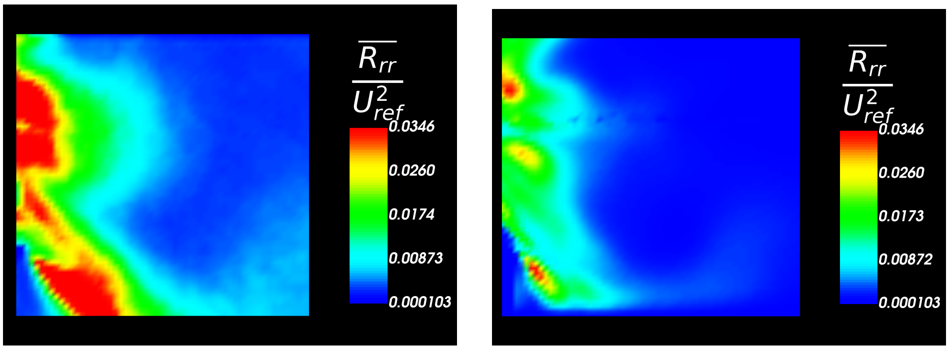

(c) R_{rr} from PIV.

(d) R_{rr} from the compressible OpenFOAM simulation.

Figure 3.7: The PIV results for Reynolds stress adapted from Abo-Serie et al. [1] compared to the compressible OpenFOAM simulation.

The Reynolds stress plots show that the Reynolds stresses calculated by the simulation are roughly 40% lower in the case of the simulation. This is not surprising as a significant amount of the instantaneous flow is being modelled by the turbulence model through the k equation. It appears that the pilot jet penetration is over-predicted by the simulation. The peak axial Reynolds stress in the mains jet is predicted to be located further radially outwards in the case of the simulation. The peak of the radial Reynolds stress appears to be concentrated at the edge of the ORZ in both cases.

Hot-wire measurements

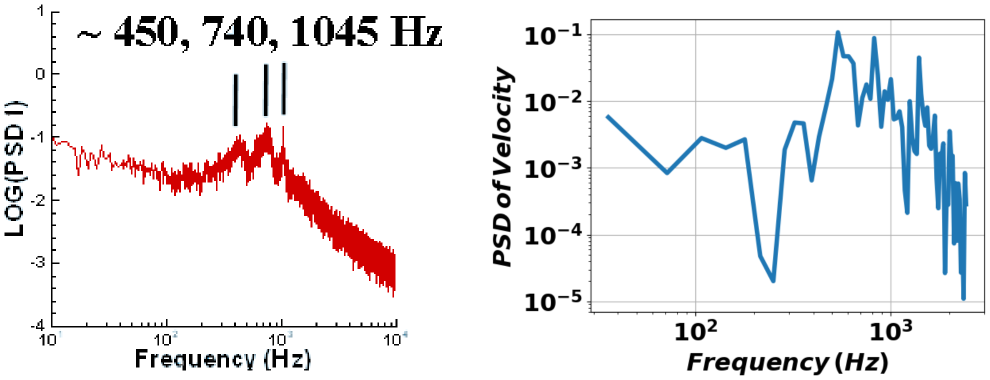

(a) Measured with HWA from [1].

(b) From an axial velocity probe in the compressible OpenFOAM simulation.

Figure 3.8: The PSD of a velocity probe located $0.026D$ downstream of the injector between the exit of passages B and C in the HWA experiment and in the compressible OpenFOAM simulation.

Abo-serie et al. [1] also conducted Hotwire Anemometry (HWA) measurements of velocity at a location $0.026D$ from the injector exit between passages B and C with a pressure drop of 3%. The PSD (power spectral density) of the hot-wire measurement is shown in Figure 3.8a. It shows three peaks at 450 Hz, 740 Hz, and 1045 Hz. An axial velocity probe was used to extract the axial velocity in the same place in the OpenFOAM simulation from the saved fields. A raw PSD of the data is presented in Figure 3.8b, it also shows three distinct peaks at 500 Hz, 800 Hz and 1400 Hz. Su [147] identified the 1400 Hz peak as being caused by a transverse acoustic mode in the duct. As the duct in the experiment is of a different size this peak should be shifted. There is an additional small peak in the numerical spectrum at 1000 Hz but it is difficult to ascertain whether this is real or simply noise. Welche's method [159] can be used to improve the signal-to-noise ratio but at the cost of reducing the resolution of the system. Using this method has the effect of blending the 800 and 1000 Hz peaks.

3.2.6 Axial velocity, pressure and azimuthal velocity

The most obvious difference between the steady and transient solutions are the maximum axial velocities reached in the pilot and the subsequent length of the pilot jet. The maximum axial

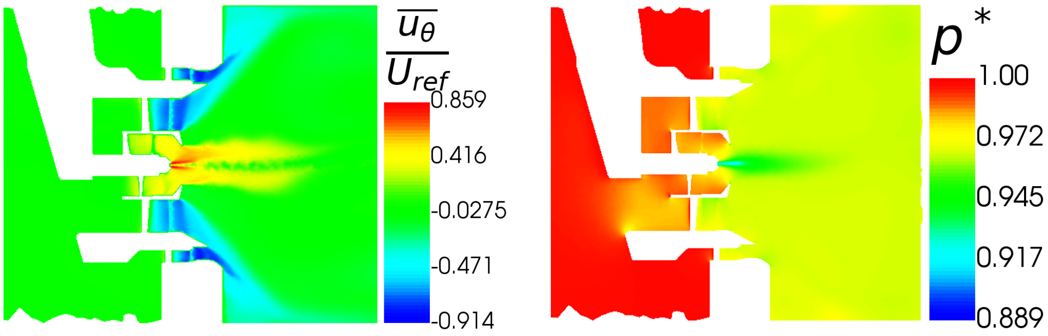

(a) Mean azimuthal velocity from the incompressible steady PRECISE-UNS simulation.

(b) Mean pressure from the incompressible steady PRECISE-UNS simulation.

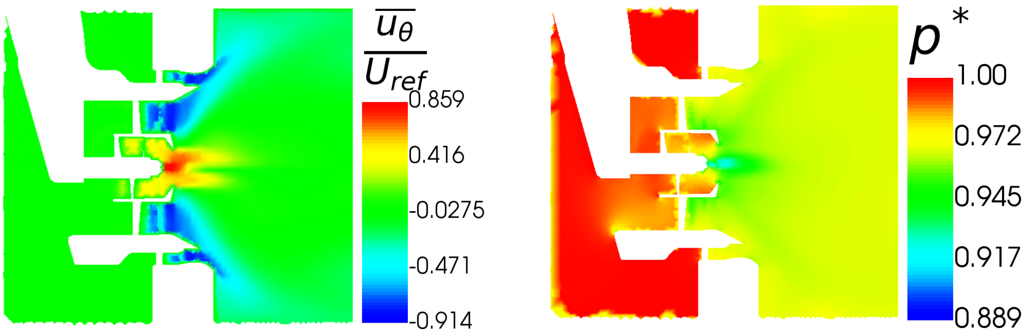

(c) Mean azimuthal velocity from the incompressible unsteady PRECISE-UNS simulation.

(d) Mean pressure from the incompressible unsteady PRECISE-UNS simulation.

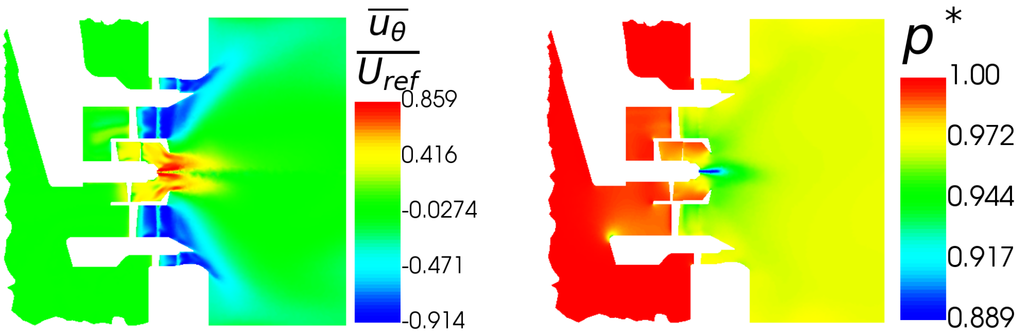

(e) Mean azimuthal velocity from the compressible unsteady OpenFOAM simulation.

(f) Mean pressure from the compressible unsteady OpenFOAM simulation.

Figure 3.9: The mean azimuthal velocity (left) and mean pressure normalised by the maximum (right) for three different simulation strategies using the LUBPAI mesh.

velocity is around 10-20% higher than the reference velocity, the theoretical maximum obtainable velocity of the flow. The driving force in the flow is the three percent pressure drop applied across the the injector. Referring to the right hand column of Figure 3.9 it can be seen that the minimum pressure inside the exhaust duct is between 9 and 11 percent lower than the upstream. This extreme low pressure region is concentrated in the core of the pilot jet, where the axial velocity is maximum, and is due to the vortex that forms at the centre of of the pilot jet. This vortex is driven by the azimuthal velocity of the flow through passage A. It can be seen that there is a good correlation between higher azimuthal velocities and lower pressures that then drive a higher axial velocity. The steady incompressible simulation has the lowest maximum axial velocity, lowest maximum azimuthal velocity and highest minimum pressure, while the unsteady compressible simulation has the highest maximum axial velocity, highest maximum azimuthal velocity and the lowest minimum pressure. The incompressible unsteady simulation has a maximum azimuthal velocity lying between the two but with the highest minimum pressure. This is most likely due to the average being calculated over the coarser mesh which tends to blur out the maximum values. This low pressure region may also help to shorten the pilot jet penetration length as it effectively provides a backwards pointing force that slows down the pilot. The pilot jet is also dispersed more effectively by large unsteady structures that convect downstream from the splitter plate between passages A and B (see Chapter 4). The fact that the inner vortex may have such a large role in the mean velocity field of the combustion chamber implies that the mesh should be refined along the centreline such that the vortex is accurately resolved and the turbulence model does not become too active in this region, potentially damping out the swirling flow. Compressibility may also play a role as the reference velocity as calculated assuming incompressible flow is around one percent lower than in the case of calculating it assuming isentropic flow. In the case where the local pressure drop is reduced by the vortex at the centre of the pilot jet, this discrepancy increases.

3.2.7 The effect of turbulence model

Figure 3.10 shows the principle components of resolved Reynolds stress for the transient incompressible and compressible simulations using the LUBPAI mesh. These are calculated by considering the fluctuating velocities as calculated and excluding the Reynolds stress that can be derived from the turbulent kinetic energy calculated by the turbulence model. In this case large numbers are better as it shows that the turbulence model is not damping out too much unsteady information. In this study the fluctuating components are important as they are influenced by the acoustic field and will affect the structure of the flame. Ideally the turbulence model would be doing very little and the simulation would be approaching a DNS. Of course the limitation of this is that the mean flow should be correct and for that the turbulence model must add additional dissipation given that turbulent scales will not be resolved by the simulation.

The Reynolds stresses are quite similar across the two simulations with a slightly higher magnitude in the case of the $k-\omega$ model although some of this will be due to the higher values in the $k-\epsilon$ simulation being interpolated onto the coarser mesh. In the case of the $k-\epsilon$ simulation the axial velocity fluctuations in the mains jet are considerably weaker. These fluctuations are caused by large velocity structures emanating from a region behind the splitter plate between passages A and B (see Chapter 4). These structures appear to be dissipated more quickly by the $k-\epsilon$ model.

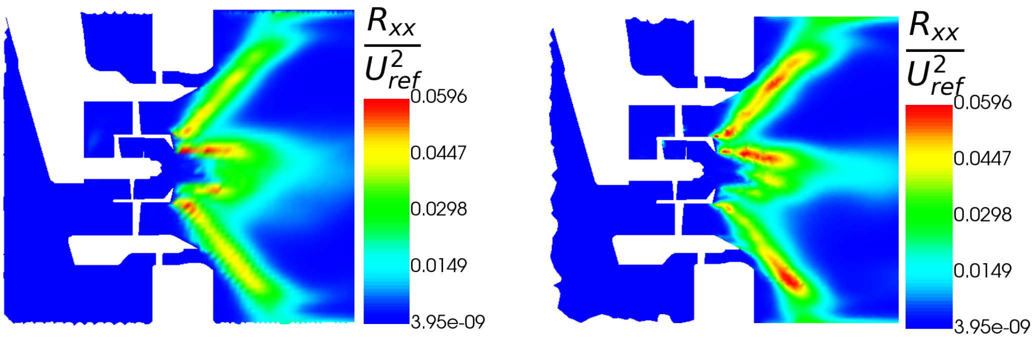

(a) R_{xx} from the $k-\epsilon$ model simulation.

(b) R_{xx} from the $k-\omega$ model simulation.

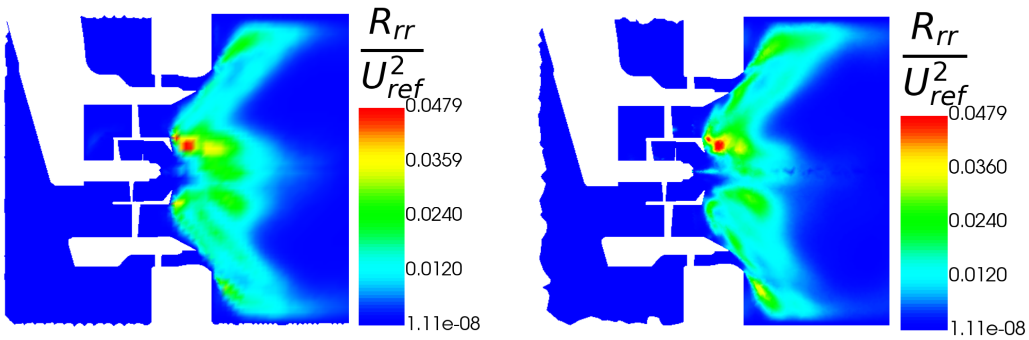

(c) R_{rr} from the $k-\epsilon$ model simulation.

(d) R_{rr} from the $k-\omega$ model simulation.

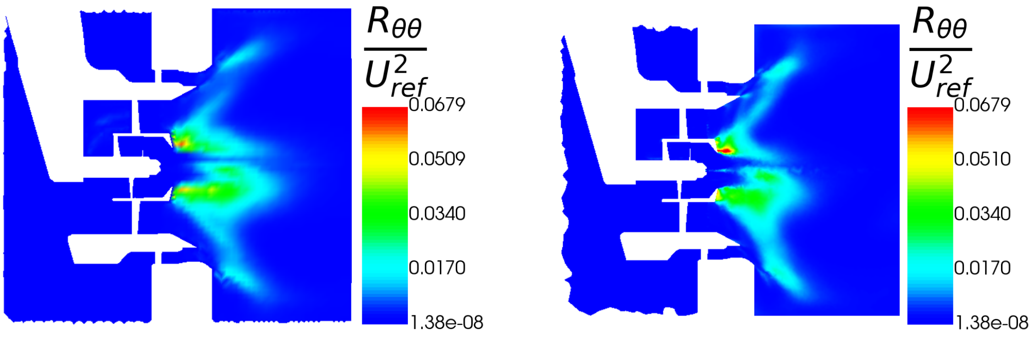

(e) $R_{\theta\theta}$ from $k-\epsilon$ model simulation.

(f) $R_{\theta\theta}$ from the $k-\omega$ model simulation.

Figure 3.10: Three components of resolved Reynolds stress from the incompressible, PRECISE-UNS simulation using the $k-\epsilon$-RNG model compared to the compressible OpenFOAM simulation using the $k-\omega$-SST model.

The contour plots of Reynolds stress show that the $k-\omega$ model does a better job at representing the fluctuations of velocity than the $k-\epsilon$ model, albeit at a greatly reduced magnitude as compared to the LES calculation as expected.

3.2.8 Incompressible LES

The PRECISE-UNS LES solver can show some sensitivity to the location of the upstream and downstream boundaries when using a total pressure upstream boundary condition. For this reason the large hemispherical plenum used with other simulations was replaced by a long upstream duct of diameter $3.9D$ and length $9.0D$. The first $2.5D$ was generated along with the rest of the mesh with the remainder extruded with $0.25D$ layers. The simulation was run for 100,000 time steps with a time step of $1 \times 10^{-3} t^*$ and the mean fields calculated over the last 50000 timesteps using mesh 23 from Table 2.1. Figure 3.11 shows the mean pressure field, two components of mean velocity and the principle components of Reynolds stress. The peak axial velocity is lower than in the unsteady compressible simulation and quite consistent with the other unsteady incompressible simulations, however the cone angle appears to be significantly shallower. This may indicate the compressibility does play an important role even in non-reacting simulations. The mean pressure field shows a lower minimum value despite a lower peak azimuthal velocity, this is probably due to the better resolution and lower dissipation of the central vortex. The largest observable difference between the LES and the other simulations is the greatly increased magnitude of resolved velocity fluctuations which should help to better predict the heat release in reacting simulations.

The energy spectral function

The quality of the LES in terms of how well resolved it is can be measured by considering the energy spectral function $E(\kappa)$ defined by [117]:

$$E(\kappa) = \oint \frac{1}{2} F_{ii}(\kappa) dS(\kappa) \tag{3.1}$$

$E(\kappa)$ is the integral of the velocity-spectrum tensor ($F_{ij}(\kappa)$) around spherical shells of radius $\kappa = |\kappa|$ where κ is the wavenumber vector. This function is important as it has been shown that at infinitely high Reynolds number, the slope of this function approaches $(-5/3)$ providing that the flow is homogeneous and isotropic [62].

Often the velocity spectrum tensor is approximated by the one dimensional tensor in the stream wise direction defined as two times the Fourier transform of the two-point velocity correlation tensor $\mathbf{B}_{ij}$ integrated in the stream wise direction only:

$$F^{(x)}_{i,j}(k_x) = \frac{1}{\pi} \int_{-\infty}^{\infty} \mathbf{B}_{ij}(r_x) e^{-i\kappa_x r_x} dr_x \tag{3.2}$$

where $\mathbf{B}_{ij}(\mathbf{r})$ is:

$$\mathbf{B}_{ij}(\mathbf{r}) = \overline{\mathbf{u}_i(\mathbf{x})\mathbf{u}_j(\mathbf{x}+\mathbf{r})} \tag{3.3}$$

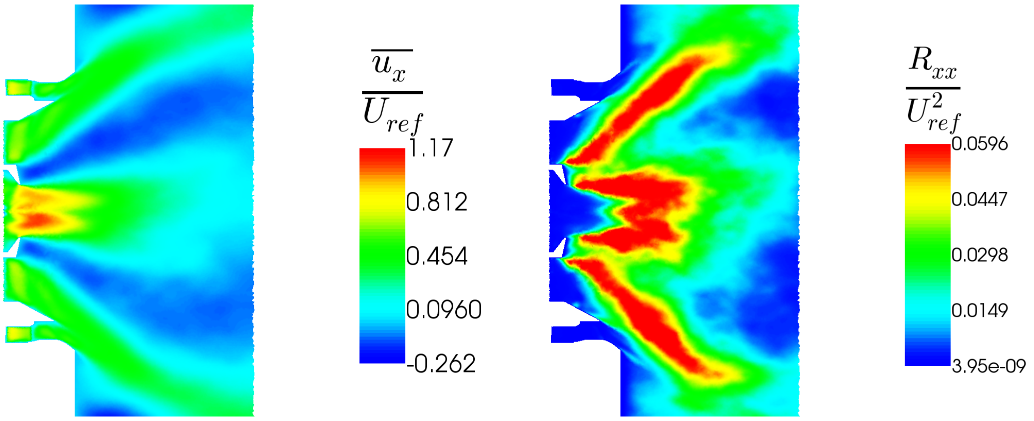

(a) u_x (b) R_{xx}

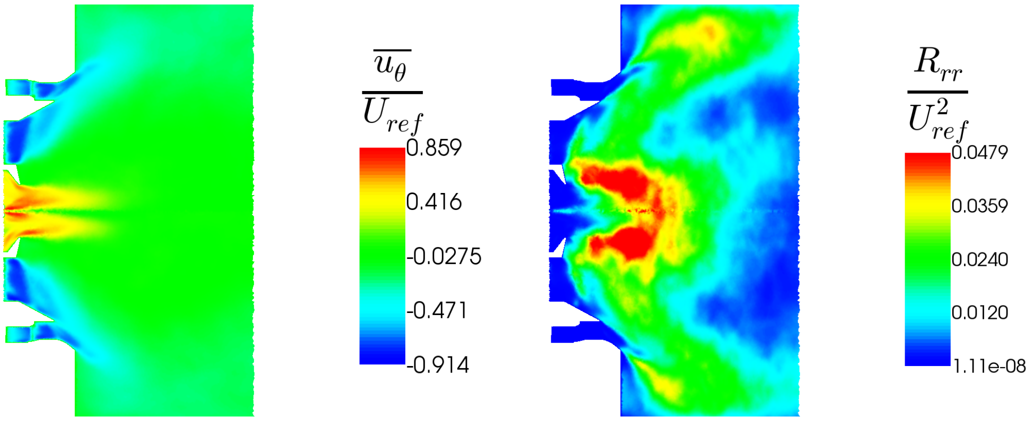

(c) u_θ (d) R_{rr}

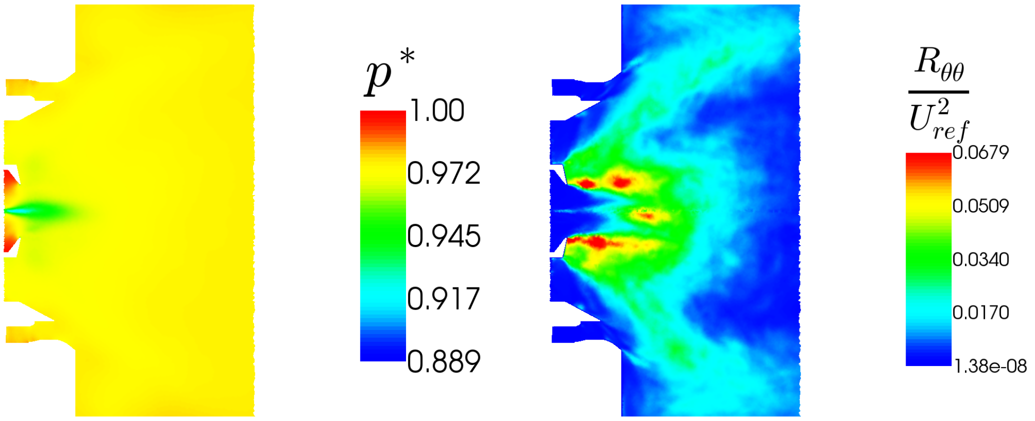

(e) p (f) $R_{\theta\theta}$

Figure 3.11: The results from the incompressible LES simulation.

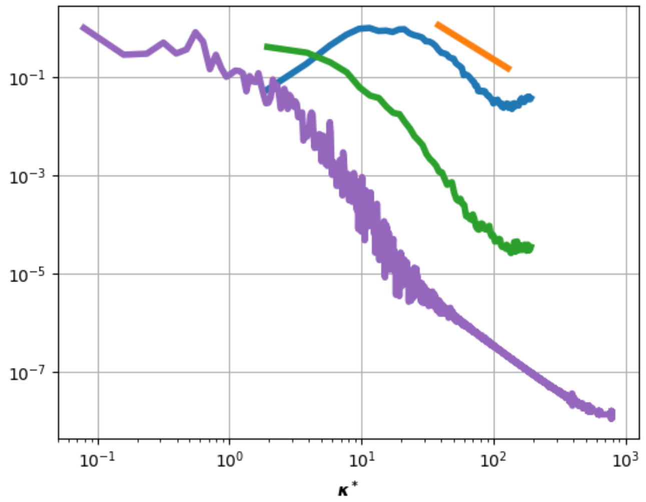

Figure 3.12: The different kinds of wavenumber spectrum for a point on the centreline of the injector at the injector exit. $E^*(\kappa^*)$ is in blue, $F_{11}^{(x)}(\kappa_x^*)$ calculated from the line probe (green) and approximated using Taylor's hypothesis (purple). $E^*(\kappa) \propto \kappa^{*\frac{-5}{3}}$ is in orange.

By using Taylor's hypothesis this quantity can be approximated by using a single high speed probe of velocity such as a hot-wire by making the transformation $r = |\mathbf{u}|\delta t$ or $\kappa = f/|\mathbf{u}|$. This is the preferred method for checking the turbulent spectrum in experiments and simulations as it requires significantly fewer velocity probes than by direct measurement of $\mathbf{B}_{ij}$. The full velocity spectrum tensor has also been resolved by extracting the velocity from 3-dimensional CFD velocity fields using line probes lying in the x, y and z directions centred at six different points. The points are located such that they lie:

1. Along the centreline, in the centre of the pilot jet at the exit plane of the injector.
2. In the CRZ at the exit plane of the injector.
3. In the centre of the mains exit flow at the exit plane of the injector.
4. Similar to 1, but $0.13D$ downstream.
5. Similar to 2, but $0.13D$ downstream.
6. Similar to 3, but $0.13D$ downstream.

Three line probes of length $0.64D$ and resolution 0.0064 collect the velocities about these points every 200th timestep. The mean velocities are removed so that $\mathbf{B}_{ij}$ can be calculated for each point. A fast Fourier transform (FFT) is then used to obtain the wavenumber spectrum for each timestep. Finally the wavenumber spectrum and two point velocity tensor is averaged in time and $E(\kappa)$ calculated. Figure 3.12 shows the non-dimensionalised energy spectral function plotted against the one-dimensional spectral tensor in the x-direction and the same calculated using Taylor's hypothesis for the first point. The spectra are plotted against the non-dimensional wavenumber $\kappa^* = \kappa D$. The plots show the marked difference between the one dimensional wavenumber spectrum (green) and the spectral energy function (blue) which is calculated from the former by multiplying it with

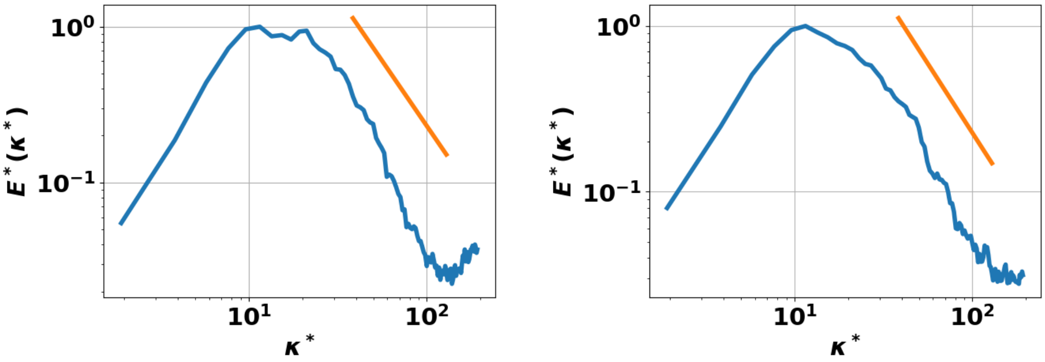

(a) Centre line, injector exit plane.

(b) CRZ, injector exit plane.

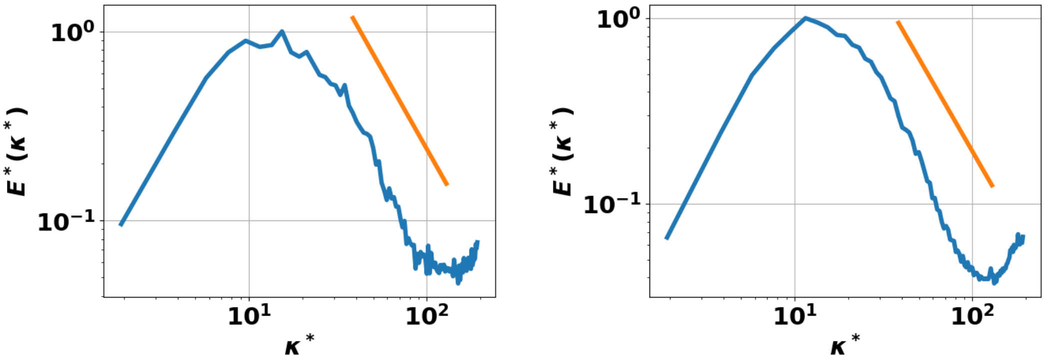

(c) Mains jet, injector exit plane.

(d) Centreline, $0.13D$ downstream.

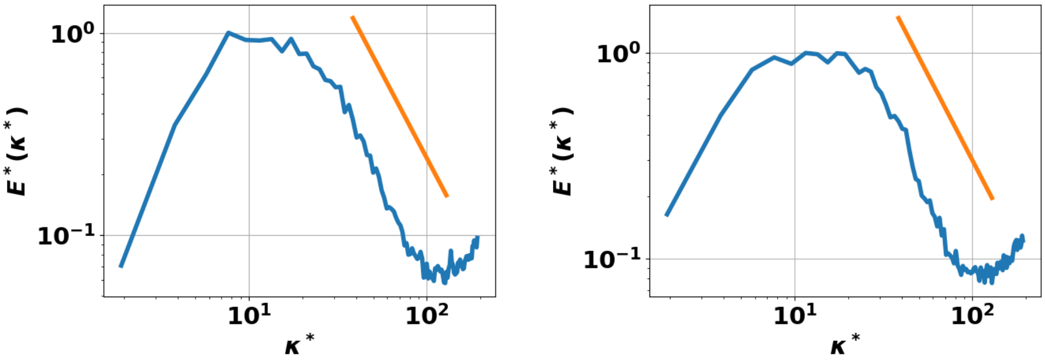

(e) CRZ, $0.13D$ downstream.

(f) Mains jet, $0.13D$ downstream.

Figure 3.13: The normalised energy spectral function $E^*(\kappa^*)$ (blue), $E^*(\kappa) \propto \kappa^{*\frac{-5}{3}}$ is in orange .

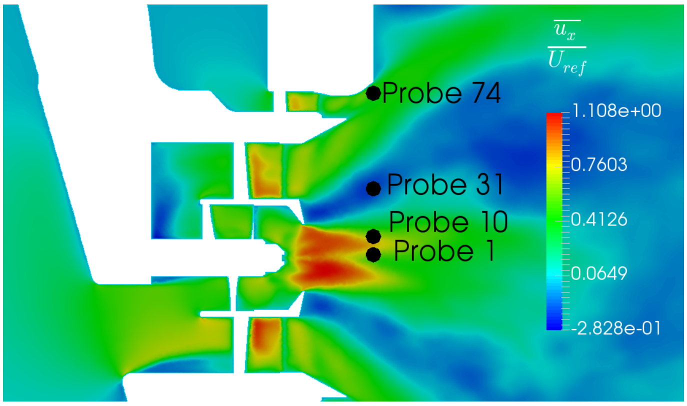

Figure 3.14: The location of the velocity probes used to generate the PSD's of axial velocity.

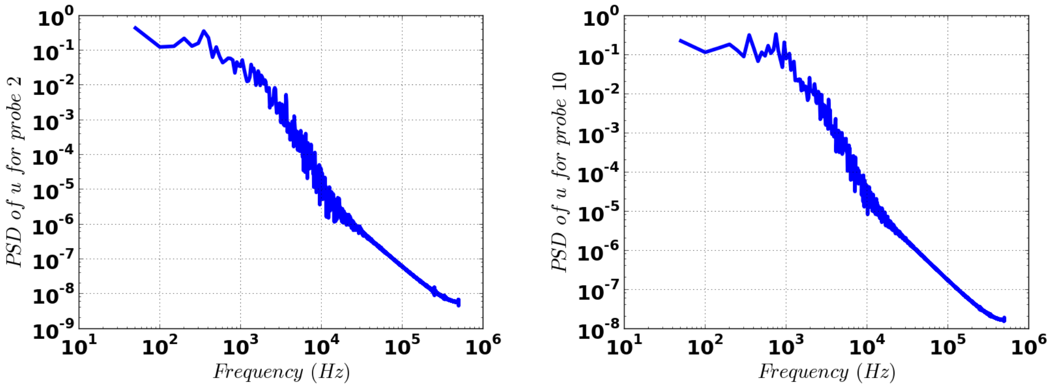

(a) Probe 1.

(b) Probe 10.

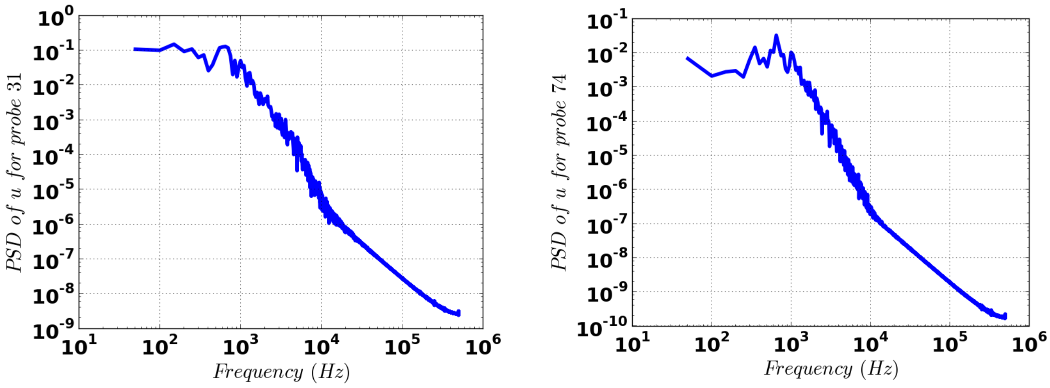

(c) Probe 31.

(d) Probe 74.

Figure 3.15: The normalised PSD of axial velocity at the locations shown in Figure 3.14.

the area of the spherical shell $dS = \kappa^2 d\phi d\theta$. This multiplication by κ^2 boosts the effect of the high wavenumber structures and reduces the gradient of the curve. The one-dimensional spectrum calculated using Taylor's hypothesis shows a similar shape to spectrum calculated from the line probe but will a discernible shift in the observed wavenumber. The two curves can be made to collapse by simply adjusting the assumed mean velocity magnitude to around a quarter of the value obtained from the simulation. This discrepancy is likely due to the large turbulence levels in the combustion chamber.

The non-dimensional energy spectral functions for each of the six points are shown in Figure 3.13. In every case the inertial range of turbulent fluctuations shows very good agreement with the -5/3 law. In all of the plots except for Figure 3.13b, the spectrum approaches the minimum before then increasing at very high wavenumbers. This is the wavenumber range over which the LES sub-grid scale model is working to remove kinetic energy from the flow. An ideal LES spectrum would show the inertial range continuing to follow the -5/3 rule until the very highest wavenumbers. The high wavenumber bump is representative of kinetic energy piling up at the very smallest scales. This implies that either not enough energy is being removed from the flow by the model or that there is insufficient grid resolution. It would therefore be useful to try to refine the mesh further or change the sub-grid scale model and see how these spectra change. This would also prove interesting in comparison with high quality experimental data.

This conclusion is balanced by the requirement that LES simulations are inexpensive enough to be useful. Simulations using the Taylor hypothesis to check the spectra may not show the high wavenumber bump seen in this study however some of these simulations may require additional refining of the mesh and yet appear in peer-reviewed literature. The reality is that academic guidelines on LES simulations are frequently flouted and yet these simulations remain useful. In an industrial setting, combustion engineers typically believe that a bad LES is better than no LES. It can be self-defeating to attempt the perfect LES simulation in industry as it may become unviably expensive.

Around 300 point probes were added to the simulation to record the local velocity at every time step. Four probes have been identified as being of particular interest the locations of which are shown in Figure 3.14. They are all located at the injector exit plane. The first probe (probe 1) is located in the centre of the pilot jet, the second is located on the outside edge of the pilot jet, the third is located in the CRZ and the fourth is located towards the outside of the mains jet. The PSD's of the axial velocity fluctuations at these four locations are shown in Figure 3.15. Three peaks can be identified as in Figure 3.8a at 350 Hz, 650 Hz and 1000 Hz.

3.3 Conclusions

It is difficult to be truly objective about the accuracy of each of the methods presented here without better experimental data. The results achieved using different unsteady methods show reasonable similarity of mean flow field solution. The steady RANS methods fail to capture the length of the pilot jet, but industrial experience says that these methods will remain popular for early stage design well into the future. As turbulence models have a set of constants that are tuned to best represent certain fundamental fluid flows, the question should be asked as to whether turbulence models should be tuned to better represent injector flows. LES methods show significant advantages when it comes to unsteady structures downstream of the injector but may suffer from poor near wall performance in the swirl passages. Hybrid methods and better meshing algorithms should be explored further. There is also a distinct lack of good quality time resolved statistics using either HWA or high speed PIV. Having this data available would also improve the validation of methods for the calculation of single phase flow.

The simulations performed in this chapter were of the geometry under a known pressure drop. This may be physically realistic, but the resulting mass flow rate is then dependent on the accuracy of the method. Because the mass flow rate of air feeds directly into the air-to-fuel (AFR) and hence the heat release, any error in the velocity field calculation will likely cause significant error in the temperature field. If the mass flow rate is set, using either experimental data or assuming a certain effective area of the fuel injector, the AFR can be maintained and any error will only appear in the pressure field which the heat release rate is only loosely dependent on.

Chapter 4

The hydrodynamic stability of the combustor

The hydrodynamic stability of the combustion system may at first not seem all that important for the understanding of combustion instabilities however the hydrodynamically unstable modes of the combustor can interact with the burner flame in several different ways. Hydrodynamic stability refers to how the flow may develop in time or space as a linear or nearly linear combination of modes.

It can be seen in Chapter 3 that the steady simulations results were not the same as the time averaged results from unsteady simulations. This discrepancy is related to the assumptions made about what the unsteady fluid is doing in the steady simulations and how this unsteadiness affects the mean flow. It is assumed in steady simulations that any unsteadiness in the flow is related to turbulent fluctuations, however this is not always the case and the neglected large scale structures will influence the mean flow field and fluctuating heat release rate differently to turbulence [91, 145]. Large scale flow structures may be present in the combustion chamber which begin as hydrodynamic instabilities. The shape and strength of the most important large scale structures in the flow are first identified using POD (Proper Orthogonal Decomposition) [77], and as hydrodynamical instabilities can be identified using linear stability analysis [58, 146, 151], an attempt is made to predict these structures using only the mean flow field.

For a simple 2D flow it might be assumed that a flow instability mode has the form [22]:

$$\mathbf{u}'(x,r,t) = \mathrm{Re}(\phi_\kappa(r)e^{\imath(\kappa x-\omega t)}) \tag{4.1}$$

where ϕ, κ and ω are the complex valued eigenvector, wavenumber and frequency respectively. The wavenumber can therefore be decomposed into $\kappa = \kappa_r + \kappa_i \imath$ where κ_r designates the spatial wavenumber and κ_i is the spatial growth rate. Similarly the frequency is decomposed into $\omega = \omega_r + \omega_i$, where the real part designates the oscillation frequency while the imaginary part designates the negative of the temporal growth rate of the mode. A positive growth rate designates the mode as being unstable, while a negative growth rate shows that a mode is stable.

In the case of an injector flow an instability mode may take the form [151]:

$$\mathbf{u}'(x,r,t) = \mathrm{Re}(\phi_m(r,x)e^{\imath(m\theta-\omega t)}) \tag{4.2}$$

In the case of $m = 1$, this is representative of a radially and axially varying flow structure with a single velocity peak in the azimuthal (θ) direction that is rotating in the positive theta direction with time. If the mean flow is rotating about the x-axis then this flow will have a vortex core extending downstream from the fuel injector. This instability mode will act to push the vortex core out from the centre line and force it to rotate about the x-axis. This leads to the well observed phenomenon of the precessing vortex core (PVC). The PVC has an azimuthal wavenumber of $m = 1$ as there is one velocity peak in the azimuthal direction per revolution. Whilst $m = 1$ structures are most common, structures of any natural number are possible. $m = 0$ structures, such as vortex rings, are possible and usually the result of longitudinal forcing. Collectively, structures with $m \neq 0$ are referred to as *helical modes* [103].

The instability expressed by equation 4.1 is referred to as a *local* instability while instabilities of the form expressed in equation 4.2 are referred to as being *global*. Local instabilities are dependant on axial location while global instability modes are calculated over the entire field. If the local instability modes lose dependence on the axial direction then the flow is referred to as being *parallel* [22]. Parallel flows rarely exist in practical applications however some engineering flows can sometimes be described *weakly non-parallel* allowing some results from parallel theory to be extended to such applications.

Unstable local instabilities can be further divided into *absolutely* and *convectively* unstable modes [22]. An absolute local instability is one that, given the correct forcing, will grow infinitely in time at any fixed location while a convective local instability is locally unstable and grows to infinite size in at least one Galilean frame while decaying in time at the fixed position in the laboratory frame of reference.

A locally convectively unstable flow can be globally stable as unstable pockets of flow can convect downstream to a point where the local conditions are no longer favourable to that instability. These flows are of interest as they act as amplifiers to external forcing such as when an injector flow is disturbed by a sound wave. Such a flow within a combustion chamber may be stable in the absence of acoustic forcing but may then transition to a violent instability when triggered by external forcing [82]. In contrast a locally absolutely unstable flow will be self excited and will lead to dynamics that are superimposed upon and will often supersede the external forcing. Such flows will typically decompose into turbulence further downstream if not already turbulent.

4.1 Hydrodynamics and combustion instabilities in premixed burners

The PVC has been observed to vanish in the case of reacting flow [65, 124, 130, 132, 135], Mannoharan et al. ([82]) show that this may be caused by the density gradient imposed by the flame. However it has also been shown to perturb the flame structure in studies by Moeck et al. [90] and Paschereit et al. [108]. Moeck [91] used tomographic reconstruction of chemiluminescence images to show that the heat release fluctuation field is helical in shape. Planar laser-induced fluorescence of OH combined with PIV was used by Stöhr [145] to show that the PVC enhances local mixing, increasing the size of the flame and enhancing heat release. Because the PVC is asymmetric about the centreline any heat fluctuations caused by it tend to cancel each other out when they are integrated in the radial and azimuthal directions. For this reason the acoustic signature of the PVC in the far-field is negligible in the case of axisymmetric burners [14, 90, 108]. This means that the PVC is unlikely to be responsible for longitudinal combustion instabilities. However, in the case of a fully-annular combustion chamber, azimuthal instabilities are possible and synchronisation of the PVCs between multiple burners could lead to combustion instabilities.

Acoustics and the PVC may also interact directly. It has been observed by many studies that the presence of a planar acoustic wave ($m = 0$) tends to damp out the PVC ($m = 1$) or higher order helical flow instabilities. This is most efficiently done at low frequencies [66, 90, 108] but has been shown to lead to amplification of the PVC at frequencies twice the natural frequency [54]. There can be non-linear interaction where the acoustic forcing results in a third frequency of oscillation being observed that is the sum of the forcing frequency and the natural oscillation frequency [14, 90, 144]. The interaction between acoustics and hydrodynamics are discussed more fully in Chapter 5.

One of the more interesting interactions between hydrodynamics and combustion instabilities was captured by Stöhr et al. [146]. A single sector premixed lab scale burner was set to run at a bi-stable operating point that saw an alternating pattern of axial ($m = 0$) combustion instability followed by the promotion of a PVC ($m = 1$) that subsequently damped the original axial combustion instability. This alternating cycle was shown to be easily predictable by simply conducting a linear stability analysis around the 'mean' flow at various times in the cycle. It was shown that the PVC generated additional turbulence which modified the mean flow, dampening the PVC and promoting the axial mode instead. The 'mean' in this case is not a true mean but instead shows that a linear stability analysis can be performed around a temporally changing flow solution.

4.2 The effects of turbulence

The effects of turbulence and unsteady fluctuations on the mean flow and flow hydrodynamics are summarised in Figure 4.1. The first is through modification of effective viscosity of the flow, which, for the purposes of this thesis will be referred to as the *mean field modification.* The increased effective viscosity also alters the stability of particular hydrodynamic modes. The turbulence itself is of course composed of fluctuations of velocity that can at least in theory interact with the hydrodynamic modes in a non-linear way however this *non-linear* or *weakly-non-linear* behaviour is considered to be of secondary importance compared to the other two. These three effects are described in this section.

4.2.1 Mean flow modification

Which mean flow?

There is a difference between solving the RANS equations and temporally averaging a flow solved with the URANS equations. The RANS equations are often used to simulate the mean flow around

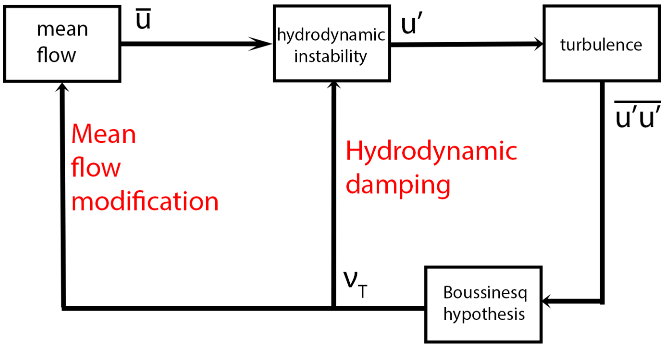

Figure 4.1: The interaction pathways between the mean flow and hydrodynamic instabilities.

or within an object and it is often assumed that the averaging process is conducted over many samples in time. However the RANS average is an *ensemble* average. In this thesis it is not always explicitly stated as to whether an average is conducted over time or as an ensemble because often it is a matter of semantics. However when it comes to conducting a stability analysis of the flow, the averaging process becomes important [52].

The ensemble average

The ensemble average is the average conducted over a number of statistically independent samples. One method for collecting these samples might be to conduct an experiment in a wind tunnel where the tunnel is started and a sample of velocity taken 5 minutes later. The tunnel is then stopped and the process repeated N_s times. The ensemble average is therefore [52]:

$$\overline{u}_s(N_s) = \lim_{N_s \to \infty} \frac{1}{N_S} \sum_{i=0}^{N_s} u_i \tag{4.3}$$

This would not be a very efficient way to conduct the experiment as it would be very time consuming so one might be tempted to calculate the temporal average instead.

The temporal average

The temporal average is completed by taking many samples of velocity within the same experiment [52]:

$$\overline{u}_t(T) = \lim_{T \to \infty} \frac{1}{T} \int_{t=0}^{T} u dt \tag{4.4}$$

Because turbulence is a *random deterministic* process [42], there is some correlation between the velocity measured at one time and the next, however the correlation drops to zero in the limit as $\Delta T \to \infty$ where ΔT is the time between two samples. The temporal average and ensemble average for a *physical experiment* will therefore converge in this same limit.

$$\overline{u}_s = \lim_{\Delta T \to \infty} \overline{u}_t \tag{4.5}$$

with

$$N_s = N_t = \frac{T}{\Delta T} \tag{4.6}$$

Luckily for the experimentalist, the correlation of turbulent fluctuations drops very quickly to zero so ΔT only has to be fractions of a second for the above relation to be valid. A good experimentalist will however check the correlation between samples to ensure that they are reasonably statistically independent.

The difficulty in being a CFD user

Although the ensemble and temporal average converge for the experimentalist, in CFD, the RANS equations solve for the steady solution of a turbulent flow and the URANS equations solve for the unsteady turbulent flow. Whilst a real turbulent flow is unsteady in time, the RANS equations solve for the stochastic mean flow which does not require time as a variable [52]. The addition of the unsteady term in the URANS equations then allows for the flow to be unfrozen in time. In order to solve for the ensemble average the integral effect of the turbulence on the mean flow must be taken into account by a turbulence model, the so called *turbulence closure* problem. The effect of turbulence is manifested through the Reynolds stress term in the RANS momentum equation and by using the Boussinesq hypothesis this is equivalent to augmenting the (effective) viscosity of the fluid (ν_e). By increasing the viscosity, the effective Reynolds number drops ($Re_e = UL/\nu_e$), stabilising the flow and allowing for a steady solution where none existed before [113].

The calculation of just the right amount of additional viscosity such that the solution converges to a steady state that resembles the ensemble averaged flow solution depends on the turbulence model used. Turbulence models are usually calibrated against particular reference flows such as homogeneous decaying turbulence and cannot be expected to reproduce the same effective viscosity as is produced by a very different flow such as the large periodic shedding of vortices behind a cylinder [81]. It is for this reason that a RANS simulation of a cylinder will be quite different (if it converges at all) to the time average of a URANS simulation of a cylinder where the large unsteady fluctuations are resolved and the turbulence modelled. Continuing this argument to its conclusion results in the timestep and grid spacing reduced sufficiently such that the turbulence closure does not need to do anything, resulting in a DNS simulation.

The shift mode ($\mathbf{u}_\triangle$)

The shift mode was first coined by Noack [101], to describe the difference between an ensemble averaged mean flow (such as from a RANS simulation) and a temporally averaged mean flow (such as from a URANS). Noack was attempting to reproduce the dynamics of the flow around a cylinder using a Reduced-Order Model (ROM). The ROM was constructed from a handful of Kahunen-Loeve (see Section A.3) modes, that were obtained from the temporally evolving flow obtained via a DNS, as a basis for a finite element approximation of the flow. The reduced order model was capable of reproducing the Von-Karmen vortex sheet behind the cylinder, however there were two problems: the first was that the flow could become suddenly unstable and shoot off to a non-physical solution. The second problem was that if the flow was started from a stationary solution the growth rate of the model was incorrect compared to the Navier-Stokes solution. It was discovered that because the Kahunen-Loeve decomposition is computed at the stable limit cycle, the modes contain no information about how the flow evolved to this solution. The shift mode is the modification of the mean flow by the unsteady dynamics and provides the model with two bits of information. Firstly it ensures that the initial flow is correct and improves the growth rate of the unsteady modes. Secondly it stabilises the limit cycle and pulls the simulation back to the limit cycle if it begins to diverge. Noack computed the shift mode ($\mathbf{u}_\Delta$) from the time average and the stationary solution of the Navier-Stokes equations of the flow around a cylinder (visualised in Figure 4.2). The flow in this case was laminar so the stationary solution is equivalent to the ensemble average (or RANS) solution in a turbulent flow. Whilst Noack performed an additional normalisation, the operation is essentially described by:

$$\mathbf{u}_\Delta = \overline{\mathbf{u}}_t - \overline{\mathbf{u}}_s \tag{4.7}$$

The shift mode can therefore be considered as the physical manifestation of the *mean flow modification* mechanism, or how the large scale unsteady dynamics modify the mean field away from the statistically steady mean field and towards the temporal mean. This concept is important to understand the *mean field correction* mechanism [152], which describes one mechanism for how acoustic forcing can alter the mean flow field and the corresponding *acoustic shift mode* described in Section 5.6.

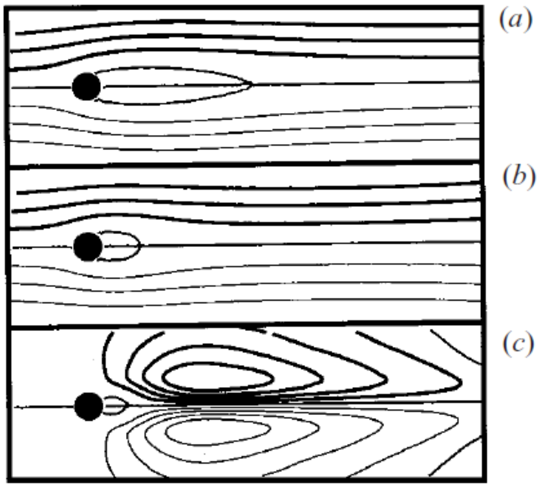

Figure 4.2: The steady flow streamlines (a), time average flow streamlines (b) and the shift mode streamlines (c). Reproduced from [101].

4.2.2 Hydrodynamic damping

Classical linear stability theory is applied to a steady solution of the Navier-Stokes equations that is devoid of fluctuating components that are not of the assumed mode form [22]. This implies that the base flow is laminar upstream of the initial perturbation. In most practical situations including swirl stabilised fuel injectors the base flow is not a steady solution of the Navier-Stokes equations as there is no such thing as a steady solution to this problem. Instead the base flow is a time-average of an unsteady simulation or a RANS average of a steady simulation. This means that the base flow and the perturbations assumed by the analysis are superimposed on the turbulent flow whether or not the turbulent flow is computed or simply modelled by the turbulence modelling strategy. The analysis is therefore only valid if the presence of the turbulence does not invalidate the assumption that perturbations grow linearly. This requires that the interaction between linear instabilities and non-linear turbulence can be safely neglected. Tammisola et al. [151], following the analysis of Turton et al. [154], have chosen to model the turbulent fluctuations using an eddy viscosity hypothesis:

$$\mu_t(\mathbf{s}_{l,1}) = \nabla \cdot \widetilde{\mathbf{u}'\mathbf{u}'} \tag{4.8}$$

Implicit in this assumptions is the idea that the turbulent fluctuations act only to damp the coherent periodic motions. Tammisola et al. [151] admit that these assumptions can only be tested *posteriori* from a fully non-linear simulation, however if the a linear global stability mode has an amplitude close to zero and shares a similar spatial structure and frequency as the leading POD (see Appendix A.3) mode then this serves as a good check as to whether the above assumptions are valid. A more complete explanation of this section is included in Appendix A.2

4.3 The calculation of unstable modes

The determination of globally or locally unstable modes begins with the linearised incompressible Navier-Stokes equations (in cylindrical polar co-ordinates):

$$\frac{\partial r u_r'}{\partial r} + \frac{\partial u_\theta'}{\partial \theta} + r\frac{\partial u_x'}{\partial x} = 0 \tag{4.9}$$

$$\begin{aligned}\rho\Big[\frac{\partial u_r'}{\partial t} + \overline{u}_r\frac{\partial u_r'}{\partial r} + \frac{\overline{u}_\theta}{r}\frac{\partial u_r'}{\partial \theta} + \overline{u}_x\frac{\partial u_r'}{\partial x} + u_r'\frac{\partial \overline{u}_r}{\partial r} + \frac{u_\theta'}{r}\frac{\partial \overline{u}_r}{\partial \theta} + u_x'\frac{\partial \overline{u}_r}{\partial x} - \frac{2u_\theta'\overline{u}_\theta}{r}\Big] \\ = -\frac{\partial p}{\partial r} + \mu_e\Big[\frac{1}{r}\frac{\partial}{\partial r}\Big(r\frac{\partial u_r'}{\partial r}\Big) + \frac{\partial^2 u_r'}{\partial \theta^2} + \frac{\partial^2 u_r'}{\partial x^2} - \frac{2}{r^2}\frac{\partial u_\theta'}{\partial \theta} - \frac{u_r'}{r^2}\Big]\end{aligned} \tag{4.10}$$

$$\begin{aligned}\rho\Big[\frac{\partial u_\theta'}{\partial t} + \overline{u}_r\frac{\partial u_\theta'}{\partial r} + \frac{\overline{u}_\theta}{r}\frac{\partial u_\theta'}{\partial \theta} + \overline{u}_x\frac{\partial u_\theta'}{\partial x} + u_r'\frac{\partial \overline{u}_\theta}{\partial r} + \frac{u_\theta'}{r}\frac{\partial \overline{u}_\theta}{\partial \theta} + u_x'\frac{\partial \overline{u}_\theta}{\partial x} - \frac{u_\theta'\overline{u}_r + u_r'\overline{u}_\theta}{r}\Big] \\ = -\frac{1}{r}\frac{\partial p}{\partial \theta} + \mu_e\Big[\frac{1}{r}\frac{\partial}{\partial r}\Big(r\frac{\partial u_\theta'}{\partial r}\Big) + \frac{\partial^2 u_\theta'}{\partial \theta^2} + \frac{\partial^2 u_\theta'}{\partial x^2} + \frac{2}{r^2}\frac{\partial u_r'}{\partial \theta} - \frac{u_\theta'}{r^2}\Big]\end{aligned} \tag{4.11}$$

$$\begin{aligned}\rho\Big[\frac{\partial u_x'}{\partial t} + \overline{u}_r\frac{\partial u_x'}{\partial r} + \frac{\overline{u}_\theta}{r}\frac{\partial u_x'}{\partial \theta} + \overline{u}_x\frac{\partial u_x'}{\partial x} + u_r'\frac{\partial \overline{u}_x}{\partial r} + \frac{u_\theta'}{r}\frac{\partial \overline{u}_x}{\partial \theta} + u_x'\frac{\partial \overline{u}_x}{\partial x}\Big] \\ = -\frac{\partial p}{\partial x} + \mu_e\Big[\frac{1}{r}\frac{\partial}{\partial r}\Big(r\frac{\partial u_x'}{\partial r}\Big) + \frac{\partial^2 u_x'}{\partial \theta^2} + \frac{\partial^2 u_x'}{\partial x^2}\Big]\end{aligned} \tag{4.12}$$

By assuming that the mean flow is homogeneous in the θ-direction the above equations can be simplified and the (global) instability and fluctuating quantities are then assumed to exist in the form:

$$\mathbf{u}'(x,r,\theta,t) = \tilde{\mathbf{u}}(x,r)e^{(\imath(-\omega t+m\theta))}, \tag{4.13}$$

Substituting these expressions into equations 4.9, 4.10, 4.11 and 4.12 results in the following system of equations:

$$-\omega_{gm}\mathbf{B}\tilde{q}_{gm} + \mathbf{L}(\tilde{q}_{gm} : \overline{q}) = 0 \tag{4.14}$$

where $q = [\mathbf{u}, p]$, $\mathbf{B}$ has diagonal elements $[\rho, 0]$ and $\mathbf{L}$ is the linear Navier-Stokes operator. This equation can now be rearranged into the form:

$$D(m, \omega, R) = 0 \tag{4.15}$$

where D is referred to as the dispersion relation and R represents the control parameters such as the Reynolds number of the flow or boundary conditions. This equation can then be solved to find the complex growth rate ω given an azimuthal wave number m with the applied control parameters R.

For the calculation of locally unstable flows, the dispersion relation becomes [23]:

$$D(m, \kappa, \omega, R) = 0 \tag{4.16}$$

with:

$$q'(x,r,\theta,t) = \tilde{q}(r)e^{(\imath(\kappa x+m\theta-\omega t))} \tag{4.17}$$

where the relation is now in terms of the axial wave number. Such problems are solved by iterating through different values of m, κ, and R and calculating the corresponding value of ω. This is also the method that is used to decide whether a local instability is convectively or absolutely unstable. The values of κ are altered until $d\omega/d\kappa = 0$ which corresponds to a saddle point of ω_i. This condition is where there is a zero group velocity. The value of ω at this point is the spatio-temporal growth rate and is designated as ω_0. If $\omega_0 < 0$ then the instability is absolute.

A region of absolute instability is a prerequisite for a globally unstable mode. Global mode identification can be achieved through at least two methods. The first is termed the *linear selection criterion* and assumes that the global mode remains quasi-linear over the domain. In this case the global mode frequency (ω_g) will correspond to the local frequency at the saddle point of the local mode frequency (ω_i) with respect to the axial location or,

$$\omega_g = \omega_i \text{ where } \frac{\partial \omega_i}{\partial x} = 0 \tag{4.18}$$

The second identification criterion is termed the *non-linear selection criterion* [110], it assumes that once the instability reaches a certain amplitude locally it will evolve in a non-linear fashion downstream [94]. In this case the global mode frequency will correspond to the frequency of the local mode furthest upstream where $Im(\omega_i) < 0$ and $d\omega/d\kappa = 0$.

The computational effort required to isolate local instabilities is significantly less than that required to compute the global modes directly and as such the local analysis is often computed first and then used to find the globally unstable modes [58].

4.4 POD analysis

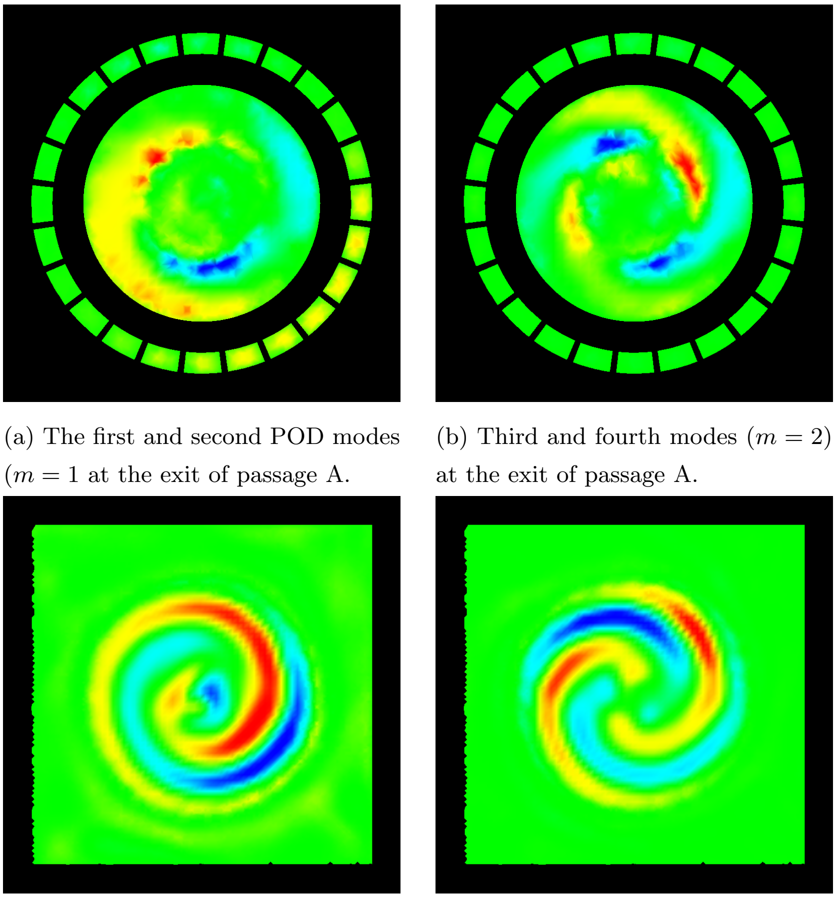

(a) The first and second POD modes ($m = 1$ at the exit of passage A.

(b) Third and fourth modes ($m = 2$) at the exit of passage A.

(c) The first and second POD modes ($m = 1$) 0.032D downstream from the injector exit plane

(d) Third and fourth modes ($m = 2$) 0.032D downstream from the injector exit plane.

Figure 4.3: The reconstructed axial velocity field from POD calculated from the compressible OpenFOAM calculation. Regions highlighted in red are high velocity regions and regions coloured in blue are low velocity regions.

POD analysis of the velocity field for the compressible OpenFOAM simulation was undertaken and showed that the first four spatial modes had a helical structure. Reconstruction of the velocity using these two modes revealed a swirling helical mode emanating from the splitter plate between passages A and B (Figure 4.3a). Reconstruction of the velocity from the third and fourth modes revealed a higher order swirling helical mode emanating from the same region (Figure 4.3b). The

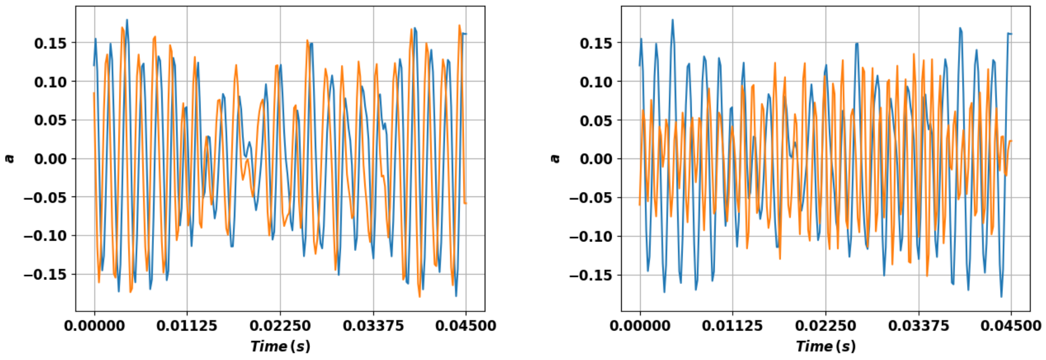

(a) The first (blue) and second (orange) POD modes plotted against time.

(b) The first (blue) and third (orange) temporal POD modes plotted against time.

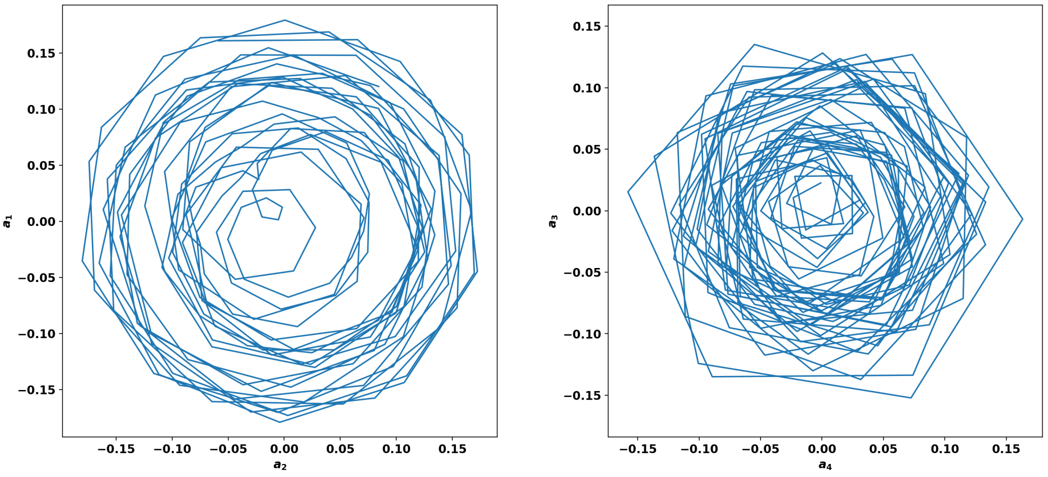

(c) Modes plotted against each other.

(d) The third and fourth modes plotted against each other.

Figure 4.4: Some of the temporal POD mode amplitudes for the unforced compressible OpenFOAM simulation.

Table 4.1: Dominant azimuthal wavenumbers for the first 10 POD modes for various simulations

Simulation	OpenFOAM	Mesh 16 URANS	Mesh 23 LES
mode 1	1	2	2
mode 2	1	2	2
mode 3	2	1	1
mode 4	2	1	1
mode 5	0	3	0
mode 6	0	0	1
mode 7	0	3	0
mode 8	0	1	0
mode 9	0	1	0
mode10	0	1	1

Table 4.2: Energy contained in the two highest energy modes of each azimuthal wavenumber expressed as a percentage of the compressible OpenFOAM simulation energy for various simulations. The resolved kinetic energy for each simulation is also expressed as a percentage of the energy resolved in the compressible OpenFOAM simulation.

Simulation	OpenFOAM	Mesh 16 URANS	Mesh 23 LES
m = 1	22.10	8.78	18.31
m = 2	12.16	16.90	20.28
m = 3	0	7.07	8.71
resolved kinetic energy	100	103	240

Table 4.3: Peak frequency of each azimuthal wave number for various simulations in Hz.

Simulation	OpenFOAM	Mesh 16 URANS	Mesh 23 LES
m = 1	450	550	400
m = 2	900	700	600
m = 3	0	1250	900

azimuthal Fourier transform of each of these modes showed that the first two modes had a dominant wavenumber of $m = 1$ and the third and fourth modes had a wavenumber of $m = 2$. The dominant azimuthal wavenumber (as defined by the wavenumber with the largest value of E_m) for each POD mode is detailed in table 4.1. The same modes are shown slightly further downstream in Figure 4.3c ($m = 1$) and Figure 4.3d ($m = 2$). These *helical modes* are seen to rotate in a clockwise direction when viewed from downstream in the same direction as the flow coming from passages B and C.

Computation of the PSD using blocks of 80 sample points using Welch's method resulted in identifying peaks in the temporal mode spectrum of around 450 Hz to 500 Hz for the first two modes and peaks of 800 Hz to 1000 Hz for the third and fourth modes. These peaks correspond quite well to the HWA and velocity probes presented in Figure 3.8. The temporal modes themselves are plotted in Figure 4.4. Figure 4.4a shows the first two modes plotted against time. The second mode (orange) appears to lead the first mode by 90 deg. The same modes are plotted against each other in Figure 4.4c showing the characteristic circular shape. The two modes can be seen to be circulating around an *attractor*, or stationary point. The modes appear to be *bi-stable*, where there appear to be two *limit cycles* at amplitudes of around 0.1 and 0.15. There also appears to be one cycle where the amplitude of both modes significantly reduce to being very close to the attractor. Referring to Figure 4.4b the third mode (orange) is plotted against the first (blue). The amplitude of the third mode is also variable and appears to be inversely proportional to the amplitude of the first mode. From left to right the amplitude of the first mode is maximum before decreasing to a minimum. The third mode does the opposite, increasing until the first mode amplitude reduces. Both modes continue at a similar magnitude before the first mode amplitude re-reduces and the third mode grows. The modes can therefore be seen to be in competition. The amplitude of the $m = 1$ helical mode is intermittently replaced by the amplitude of the $m = 2$ mode. The

$m = 2$ mode dynamics are represented in Figure 4.4d. This shows that the 3rd and 4th modes are also rotating about an attractor with two stable limit cycles. The mode is oscillating at a faster frequency so the orbits appear less circular due to the temporal discretisation and the increased number of periods.

The same analysis has been repeated for two other simulations, the PRECISE-UNS unsteady incompressible URANS simulation with mesh 16 from Table 2.1 and the incompressible LES simulation. The first is computed on a larger mesh using the PRECISE-UNS solver and the $k - \epsilon$ turbulence model. Analysis of the wavenumber of the first ten POD modes shows that the largest energy containing modes are $m = 2$ helical modes followed by $m = 1$ modes, $m = 3$ modes also appear. The resolved kinetic energy is slightly increased due to the larger mesh. The energy contained in the ($m = 1$) modes is significantly reduced and the ($m = 2$) increased slightly. If the total energy contained in all the helical modes are added together it is constant across both simulations at around 32%. Dominant mode frequencies are shifted slightly higher in this case for the $m = 1$ and lower for the $m = 3$ mode.

The LES is computed on a much larger mesh (mesh 23 from Table 2.1) using the LES solver of PRECISE-UNS. The resolved kinetic energy is almost 2.4 times larger in this case due to the higher resolution mesh and much lower effective viscosity. The energy contained in the $m = 1$ modes is reduced slightly while the energy in the $m = 2$ mode is increased significantly. $m = 3$ modes appear as modes number 11 and 16 and have around the same amount of energy as compared to the PRECISE-UNS URANS simulation. The peak frequency of the $m = 1$ mode is slightly lower in this case while the $m = 2$ peak frequency is significantly lower.

Once the higher resolved energy content of the LES is taken into account, the energy content per mode across the simulations shows reasonable consistency despite the changing meshes and numerical methods which shows that even relatively low cost URANS simulations are capable of reproducing the largest scale flow features. The relative magnitude of these helical modes with respect to the others does show quite a lot of sensitivity to method and given that each of these modes are in competition with each other even small changes to the mesh could easily push one mode to being more dominant. The preference of the $m = 2$ mode for the incompressible simulations instead of the $m = 1$ may indicate that the stability of these modes is dependent on local changes to density as well as the velocity. A higher resolution mesh seems to be necessary to capture $m = 3$ modes.

The peak frequency of these structures is also sensitive to method although part of the inconsistency observed is due to the difficulty is obtaining good frequency spectra with CFD data due to the relatively short running times (0.045 s). The dominant azimuthal wave numbers for each simulation are shown in table 4.1, the mode energies in table 4.2 and the peak frequencies in table 4.3.

4.5 Local Stability Analysis

A local stability analysis of the mean, unforced URANS flow computed with OpenFOAM was completed at Cambridge University by Matthew Juniper. The aim of this investigation was to see if the linear stability analysis was capable of finding the $m = 1$ and $m = 2$ global modes in line

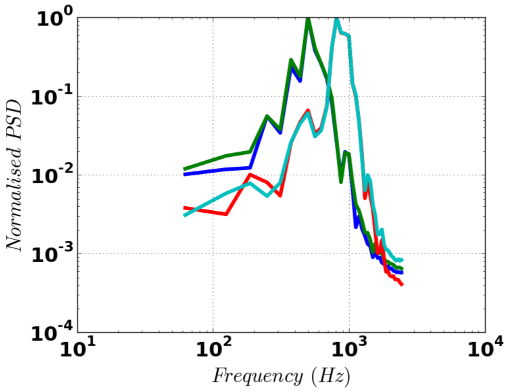

Figure 4.5: The PSD of the first (blue), second (green), third (red) and fourth (cyan) temporal POD modes for the compressible OpenFOAM simulation

with those found using the POD method. An additional benefit of finding these globally unstable modes is that their interaction with acoustics might be able to be predicted using the theory of Hemchandra [60] as discussed in Section 5.6.3. Pre- and post-processing of the data was completed by the author.

4.5.1 Methodology

The fully 3D flow field was first averaged in time over 225 snapshots and then circumferentially averaged to create a 2D flow field. This averaged field was then analysed at Cambridge University using a program called InstaFlow [58] to determine the locally unstable modes in the injector flow field.

Structured cylindrical mesh generation, interpolation and averaging

In order to create the 2D flow field required by the linear stability analysis, the flow must circumferentially averaged. Because the data is originally computed over an unstructured mesh, the flow must be interpolated onto a structured cylindrical polar mesh.

The mesh was generated using the *vtk* package [133]. The points of the mesh were first generated by moving one point in the x-direction followed by one move in the radial direction followed by N movements in the circumferential direction before moving radially outwards once more. Once the outer radius was achieved the mesh was continued in the x-direction. The cells were then generated in a similar order by defining the mesh connectivity as a set of tetrahedral elements along the centreline and hexahedral elements elsewhere. The mesh generation program was made in such a way that the number of radial, azimuthal and axial segments could be easily altered. The mesh dimensions are obtained automatically from the the donor mesh however the user has the ability to truncate the extent of the mesh as required. Once the mesh has been generated each physical quantity is interpolated linearly from the surrounding points in the unstructured mesh to the structured polar mesh using the *vtkProbeFilter* function.

The cylindrical mesh generated for the linear stability analysis was chosen to have 1000x400x500 points in the x, r, θ directions resulting in a resolution of $0.005D \times 0.004D \times 0.013$ radians or around

$0.001D$ at the most extreme radius. This resolution is similar to the original mesh however slightly courser in parts due the the limitations of the computing resources used. Azimuthal averaging was achieved by simply summing up all of the values at a given x and r location and dividing by the number of azimuthal segments. The mean axial velocity field used in the linear stability analysis is visualised in Figure 4.6.

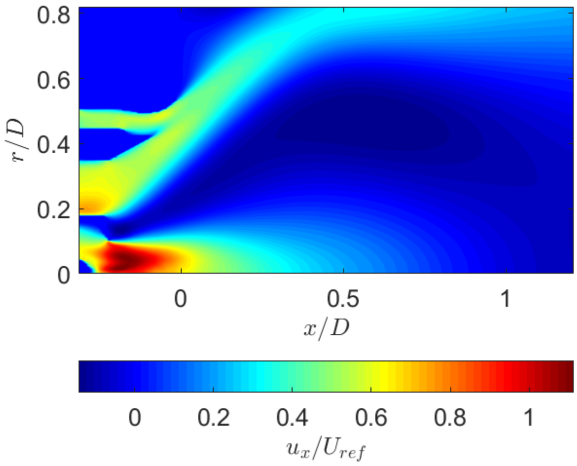

Figure 4.6: The mean axial velocity field for the unforced simulation interpolated onto a cylindrical grid and averaged in the θ direction as used in the linear stability analysis.

InstaFlow

InstaFlow [58] is software developed by Matthew Juniper designed to calculate the locally and globally unstable modes from a 2D velocity and effective viscosity field. The local dispersion relation is reformulated into a generalised eigenvalue problem:

$$\mathbf{A}\hat{q} = \omega \mathbf{B}\hat{q} \tag{4.19}$$

where $\mathbf{A}$ and $\mathbf{B}$ are matrices. The maximum temporal growth rate, ω, and the maximum spatio-temporal growth rate, ω_0 , are then calculated by iterating through the values of κ, where κ is a real number until the highest value of the imaginary part of ω is found (ω_i) given an azimuthal wave number (m). The maximum spatio-temporal growth rate is found by starting from the maximum temporal growth rate (ω) and iterating through values of κ in the complex plane until $d\omega/d\kappa = 0$, the corresponding maximum spatio-temporal growth rate (ω_{0i}) corresponds to the growth rate of the wave with zero group velocity, the wave that stays at the point of impulse. If $\omega_{0i} > 0$ and $\omega_i > 0$ then the instability is absolute however if $\omega_{0i} < 0$ and $\omega > 0$ then the flow is absolutely stable and convectively unstable. *Instaflow* is then capable of finding the globally unstable modes following the local analysis however the local analysis relies on the flow being sufficiently *parallel*. From the results found in this study it appears that this flow is not sufficiently parallel to render the local analysis suitable and as such a global analysis must be undertaken separately. It was decided that the global stability analysis was not of sufficient priority in this study to warrant further effort.

4.5.2 Results and discussion

The local stability analysis identified an unstable helical mode with azimuthal wavenumber $m = -2$ corresponding to a clockwise rotating double helix travelling downstream at a radii of around 0.27D. The instability is absolutely unstable over frequencies between 400 Hz (upstream) and 1000 Hz (downstream) with wavelengths equivalent to 1.08D to 1.44D. The instability is driven by the azimuthal and axial shear between passages A and B. The growth rate of the instability is maximum where the two streams come into contact and decays thereafter. The absolute growth rate and frequency as a function of axial location is presented in Figure 4.7.

The frequency of the $m = 2$ instability identified by POD shows a peak frequency of around 1000 Hz in agreement with the frequency observed in [1] using a hot wire. This implies that the instability grows extremely quickly from when the streams come into contact and becomes sufficiently strong such that the instability becomes non-linear and turns into a vortex that convects downstream. This is analogous to the example given in section 5.1 of [94] of a helium jet discharging into a 'semi-infinite domain' where the linear global mode is given, to first approximation, by the conditions where the streams first come into contact.

A clockwise rotating helix ($m = -1$) mode was also identified in the local stability analysis with a wavelength of 0.59D and an oscillation frequency of 574 Hz. This mode has a noticeably higher frequency than the $m = 1$ mode isolated using POD and measured using a hot wire [1]. Whilst the approximation of the flow as being axisymmetric is one source of error, it is the opinion of Matthew Juniper that this discrepancy is more likely due to the flow being insufficiently parallel. This is evidenced by the large change in growth rates and oscillation frequencies as a function of axial location in the case of the $m = -2$ helical mode. For this reason the identification of the global mode from the local is problematic and requires the use of alternative techniques.

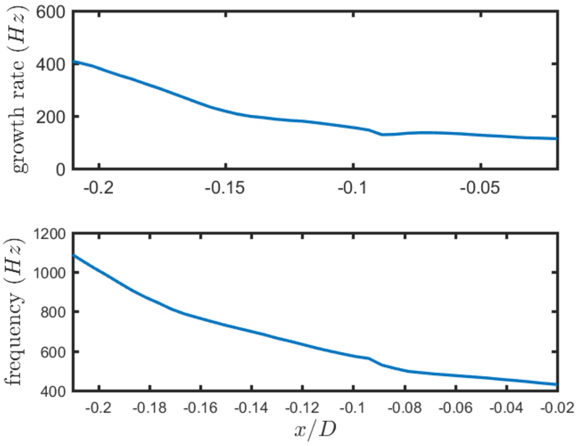

Figure 4.7: The growth rate and frequency of the $m = -2$ mode found using linear stability analysis as a function of axial distance at a radial distance of 0.27D.

4.6 Conclusion

The motivation for undertaking the stability analysis is more fully explored in chapter 5 however the observation that such a large part of the flow field consists of a relatively small number of large scale flow structures is encouraging. These structures have been identified here using POD however at least in theory these structures should also be observable using global linear stability analysis methods. Provided that the mean flow field is obtainable using computationally efficient methods then these structures can be resolved much more quickly than having to wait for LES computations. Due to the very limited resources available to the design engineer, this opens the possibility of analysing some of the unsteady flow field as part of the early design process. Observations made by others [145, 146, 152] that helical modes do interact with acoustic modes and influence heat release fluctuations adds weight to the argument that the consideration of hydrodynamic stability should become one of the tools available to industry. The problem is that the current state of the art for linear stability analysis appears to be insufficiently developed for it to be used in a predictive capability. The local stability analysis competed here does find evidence of the $m = 1$ and $m = 2$ modes however it is only after the modes have been identified using POD that any confidence in the findings may be obtained. Due to difficulties in identifying these local modes, the global modes could not be computed. If linear stability analysis is to become useful in the design of new fuel injectors it must become more reliable and robust. The identification of these structures using POD is still useful as their large kinetic energy are bound to influence spray convection, atomisation and evaporation as discussed in Chapter 7 and mixing and combustion as discussed in Chapter 8. Chapter 5 discusses how these structures may interact with the acoustic field and the injector acoustic response.

Chapter 5

The acoustic response

The acoustic impedance is defined as being the complex transfer function that links the incoming pressure fluctuation and the corresponding velocity [59]. From the point of view of the linear acoustic field, the acoustic impedance is often considered enough to define the response of a small component such as a fuel injector. A large component might require an acoustic transfer matrix instead because there is sufficient delay between waves entering and exiting the component. For the flow field however, the fluctuation of mass that the acoustic impedance describes is made up of a series of complicated non-linear interactions that have a localised effect on the flow field and therefore the spray convection and the flame.

After describing in Chapter 1 how the flame may be perturbed by acoustics through interaction between the acoustic field and the fuel injector, this chapter attempts to test the ability of a compressible URANS method to reproduce the acoustic impedance and acoustic transfer matrix of the chosen injector geometries. The acoustically forced flow fields are also analysed to identify the effects that acoustic forcing have on the downstream flow field and to describe the possible interaction mechanisms between the acoustics, mean flow field and large scale structures identified in Chapter 4.

5.1 A simple example

As an example, consider a simple orifice in a wall, subjected to an incoming pressure wave. A fuel injector is not an orifice, however whilst it is significantly more complicated, it is fundamentally a hole in a wall separating a pressure differential. As such, one expects similar mechanisms of interaction between the injector and the sound wave. As the pressure wave approaches, the pressure differential across the orifice will change and then the wave either passes through or is reflected by the wall. As the pressure differential decreases, the fluid residing inside the orifice will accelerate in the negative direction and then as the wave passes, it will accelerate in the positive direction. At the limit of infinite wavelength or zero frequency limit the flow through the orifice will act in a quasi steady manner with any induced flow rate through it being equivalent to that induced when a pressure differential is applied across the orifice and left for an infinite amount of time.

As the frequency of incoming wave increases, the pressure differential is applied for progressively smaller amounts of time. The system can be modelled as a solid mass being subjected to an oscillating force. If the pressure is applied over a uniform area and the fluid inside the orifice is considered incompressible then [59]:

$$\delta p = \rho L \frac{du}{dt} \tag{5.1}$$

where L is the length of the orifice, u is the acoustic particle velocity and ρ is the fluid density. If δp (the pressure fluctuation) is a simple sinusoidal signal ($\delta p = cos(\omega t)$) then this can be represented as the real part of the complex pressure signal:

$$\delta p = \mathrm{Re}(\tilde{p} e^{j\omega t}) \tag{5.2}$$

leading to the differential equation:

$$\rho L \frac{du}{dt} = \tilde{p} e^{j\omega t} \tag{5.3}$$

by assuming a solution of the form $u = Ae^{j\omega t}$ [59], the particular solution can be found to be:

$$u = -\frac{\tilde{p} j}{\omega \rho L} e^{j\omega t} \tag{5.4}$$

The specific acoustic impedance is formally defined as:

$$\mathcal{Z}(\omega) = R + jX = \frac{\hat{p}}{\hat{u}} \tag{5.5}$$

and for the simple example given can be shown to be equal to:

$$\mathcal{Z}(\omega) = L\omega j \tag{5.6}$$

The real part of this transfer function (R) is referred to as the resistance and describes the in-phase response of the system. The imaginary part (X) is referred to as the reactance and describes the out of phase response of the mass of air to to the pressure fluctuation. Together they describe both the system gain and the phase shift or time lag between the peak incoming pressure wave and the peak disturbance to mass flow rate. As mentioned in the introduction, the time lag is especially important for the prediction of acoustic instabilities. The resistance, reactance, gain and phase shift can be linked through:

$$\phi_Z = \tan(X/R) \tag{5.7}$$

and:

$$Z = |Z| e^{j\phi_Z} \tag{5.8}$$

If the gain ($|Z|$) is calculated incorrectly then the sensitivity of the system will be affected, if the phase shift is significantly different then the time it takes for a pressure wave to disturb the heat release rate will be calculated incorrectly and hence frequencies identified as those being most responsible for instability will be incorrect.

In the above analysis it is assumed that the fluid is inviscid and that there is no bias flow. The result therefore shows a viable value of reactance but no resistance. in order to find a theoretical expression for the resistance, the effect of viscosity must be included in the analysis.

Kinsler et. al. [59] approximate the effect of the laminar boundary layer on the acoustic impedance leading to:

$$\mathcal{Z}(\omega) = \frac{1}{\pi}\Big[\frac{\omega^2}{c} + \frac{L'}{a^3}\sqrt{2\nu\omega} + \frac{L'}{a^2}\big[\omega + \frac{1}{a}\sqrt{2\nu\omega}\big]j\Big] \tag{5.9}$$

where the first term is the effect of radiation, the second term represents the resistance due to viscosity, the third term represents the phase lag due to the mass inside the orifice and the 4th term is representative of the phase lag induced by viscosity. a is the radius of the orifice and the length of the orifice (L) has been replaced with a modified length L' that includes the effects of the orifice shape on how much fluid is moved by the acoustic wave.

Up until this point the bias flow, induced by the mean pressure drop across the orifice, has been neglected. In order to take this effect into account it is best to consider the modified Cummings [27] equation developed by Luong et. al. [78]. This leads to an impedance of:

$$\mathcal{Z}(\omega) = \frac{1}{S}\left(\frac{U}{C_d^2} + \omega L' j\right) \tag{5.10}$$

where C_d is the discharge coefficient of the orifice and U is the mean velocity. An alternative approach was developed by Bellucci et al. [11] that casts the fundamental relationship between acoustic pressure and velocity as being composed of a mass component, viscous component and a component related to the discharge coefficient:

$$\mathcal{Z}(\omega) = \frac{1}{\pi}\Big[\frac{U}{C_d a^2} + \frac{L'}{a^3}\sqrt{2\nu\omega} + \frac{L'}{a^2}\big[\omega + \frac{1}{a}\sqrt{2\nu\omega}\big]j\Big] \tag{5.11}$$

Comparing this equation to Equation (5.9) that neglects bias flow, it can be seen that this equation does not take into account the loss of acoustic energy due to the radiation of sound. The following equation is also aided by noticing that the effect of bias flow will almost certainly lead to the flow becoming turbulent, at least in the boundary layer or across the orifice face, therefore the viscosity ν might be replaced by ν_e the effective viscosity, a combination of the fluid viscosity and the apparent turbulent viscosity that approximates the net effect of turbulence on the flow:

$$\mathcal{Z}(\omega) = \frac{1}{\pi}\Big[\frac{U}{C_d a^2} + \frac{\omega^2}{c} + \frac{L'}{a^3}\sqrt{2\nu_e\omega} + \frac{L'}{a^2}\big[\omega + \frac{1}{a}\sqrt{2\nu_e\omega}\big]j\Big] \tag{5.12}$$

In this equation each of the terms represent: the bias flow effect on resistance, radiation effect on resistance, viscosity effect on resistance, mass effect on reactance and the viscosity effect on reactance. Several steps in the analysis have been left out for reasons of brevity, for a more complete analysis, please refer to Section A.4.

5.2 Multi-microphone technique

The multi-microphone technique is a method for ascertaining the acoustic impedance of an acoustic element. It is a development of the two-microphone technique described by Seybert and Ross [136] and has been used extensively to measure the acoustic impedance of injectors and orifices [127]. Experiments of this kind are conducted by placing the acoustic element of interest at the end of a long duct. At the other end of the duct is placed a noise generation unit, most commonly a set of loudspeakers. In the case of the linear regime testing, a bias flow and pressure drop must be present and as such the duct needs to be open ended at both ends with the noise generation unit(s) being placed around the edges of a central duct that acts and an exhaust. A differential pressure is applied across the duct using a fan to induce a bias flow.

Close to the acoustic element, two or more high speed pressure transducers are placed to record the fluctuations of pressure generated by acoustic waves in the duct. Providing the duct is long enough and the frequencies of interest are low it is safe to assume that only plane acoustic waves can exist in the duct. The pressure recorded at the wall of the duct is therefore equal to the pressure throughout the cross section of the duct at that location and is the sum of the pressure contributions from a wave travelling upstream from the noise source to the acoustic element and a second reflected wave. Mathematically this is represented as [150]:

$$\tilde{p}(x) = \hat{p}e^{j\omega t} = \hat{p}_i e^{j\omega t + j\kappa_- x} + \hat{p}_r e^{j\omega t - j\kappa_+ x} \tag{5.13}$$

where the subscript i refers to the upstream, incident acoustic wave and the subscript r refers to the reflected acoustic wave and $\kappa_\pm = \omega/(U \pm c)$. By noting that the acoustic velocity $\tilde{u}_p$ can be related to the acoustic pressure by using the linearised Euler equations as:

$$\tilde{u}_i = \frac{-\tilde{p}_i}{\rho c} \text{ and } \tilde{u}_r = \frac{\tilde{p}_r}{\rho c} \tag{5.14}$$

where c is the local speed of sound, the equation for the local acoustic velocity can be written as:

$$\tilde{u}_p(x) = \hat{u}_p e^{j\omega t} = \frac{1}{\rho c}\Big(- \hat{p}_i e^{j\omega t + j\kappa_- x} + \hat{p}_r e^{j\omega t - j\kappa_+ x}\Big) \tag{5.15}$$

If the pressure is measured simultaneously at n locations, the following matrix equation can be set up relating the pressure signal observed at the n locations and the magnitude of the incident and reflected waves:

$$\begin{bmatrix} e^{j\kappa_- x_1} & e^{-j\kappa_+ x_1} \\ \vdots & \vdots \\ e^{j\kappa_- x_n} & e^{-j\kappa_+ x_n} \end{bmatrix} \begin{bmatrix} \hat{p}_i \\ \hat{p}_r \end{bmatrix} = \begin{bmatrix} \hat{p}_1(x_1) \\ \vdots \\ \hat{p}_n(x_n) \end{bmatrix} \tag{5.16}$$

This system has a unique solution for the case where $n = 2$, however for the case of $n > 2$ the system is over-defined and can be solved by using the method of least square fit. Two pressure transducers are therefore sufficient for determining the magnitude of the incident and reflected pressure waves, however one must be careful to ensure that the location of the pressure transducers does not correspond to or is too close to the location of a pressure node otherwise the transducer may not observe fluctuations of pressure. This can become a problem if the wave frequency and therefore wavelength is altered over a large range as the chance of a pressure node falling onto the location of a pressure transducer is greatly increased. One solution is to move the location of the pressure transducers for differing test frequencies however this increases the time for experiments to be undertaken. Sufficient experience at Loughborough university [126] has shown that by using four pressure transducers the likelihood of pressure nodes falling onto the location of the pressure transducers is reduced without overly increasing the cost and complexity of the experimental setup.

Once the incident and reflected wave amplitudes are found, the acoustic velocity and pressure at the acoustic element can be ascertained using Equations 5.13 and 5.15 respectively. Finally the acoustic velocity within the acoustic element is found by assuming that the acoustically driven mass flow rate through the acoustic element and the duct are the same and therefore:

$$\tilde{u}_{p,AE} A_{\text{eff}} = \tilde{u}_{p,\text{duct}} A_{\text{duct}} \tag{5.17}$$

where A is the area and the subscripts AE, eff and duct refer to the acoustic element, effective (area) of the acoustic element and duct respectively. Error analysis by Rupp [126] has shown that pressure and velocity wave amplitudes were measured to within 3% while the phase angles were measured to be within 0.5° with a repeatability of 1%.

5.3 Characteristic boundary conditions

When the switch is made between solving the incompressible to compressible flows the flow now supports the creation, convection and destruction of sound waves [39]. Sound waves propagate through the flow at $\mathbf{u} \pm c$ meaning that if the flow leaving the domain is supersonic then information cannot travel back upstream and hence no boundary conditions are required whereas if the flow is subsonic, then at least the pressure must be set at outflow boundary similar to incompressible flows [115]. A sound wave however is not just a wave propagating in the pressure field, but also in the velocity field as well, therefore in the case of subsonic flows, the pressure and velocity should be set at the outflow boundary. the problem is that the velocity boundary condition also applies to the non-acoustic velocity field and therefore this boundary condition must be set while taking into account the pressure and velocity fields together. If the amplitude of the acoustic waves are small, the effect of viscosity is negligible and the acoustic field can be approximated by the Euler equations. If each of the boundaries are carefully placed such that viscous effects are negligible (far from walls, low unsteadiness, low turbulence levels) then, in the vicinity of the boundary, the full Navier-Stokes equations can also be replaced by the Euler equations [100]. By diagonising these equations and reorienting the axis such that on of the principal directions corresponds to the boudnary normal direction $\mathbf{n}$. The original flow variables ($\mathbf{V} = [\rho, u, v, w, P]^T$) are replaced with new wave (or characteristic) variables $\mathbf{W}_n$ orientated in the $\mathbf{n}, \mathbf{s}_1, \mathbf{s}_2$ directions [100]:

$$\delta \mathbf{W}_n = \begin{bmatrix} \delta\rho - \frac{1}{c^2}\delta P \\ \mathbf{s}_1 \cdot \delta\mathbf{u} \\ \mathbf{s}_2 \cdot \delta\mathbf{u} \\ \mathbf{n} \cdot \delta\mathbf{u} + \frac{1}{\rho c}\delta P \\ -\mathbf{n} \cdot \delta\mathbf{u} + \frac{1}{\rho c}\delta P \end{bmatrix} \tag{5.18}$$

where the $\mathbf{s}_1$ and $\mathbf{s}_2$ directions lie in the plane of the boundary. Each component of this vector is a different wave with the first line representative of an entropy wave, the second and third lines representative of velocity waves parallel to the boundary surface and the 4th and 5th components representative of acoustic waves travelling normal to the surface. W_n^5 is the acoustic wave entering the domain while W_n^4 is the acoustic wave leaving the domain [115]. Each of these waves has a convection velocity of $[u_n, u_n, u_n, u_n + c, u_n - c]$.

These two acoustic waves can be related back to the pressure and velocity at the boundary by [100]:

$$\frac{\partial u_n}{\partial t} = \frac{1}{2}\Big(\frac{\partial W_n^4}{\partial t} - \frac{\partial W_n^5}{\partial t}\Big) \tag{5.19}$$

$$\frac{\partial p}{\partial t} = \frac{\rho^n c^n}{2}\Big(\frac{\partial W_n^4}{\partial t} + \frac{\partial W_n^5}{\partial t}\Big) \tag{5.20}$$

If the 1st order Euler finite difference method is applied to these equations, the updated velocity and pressure at the new timestep can be found using:

$$u_n^{n+1} = u_n^n + \frac{1}{2}(\delta W_n^4 - \delta W_n^5) \tag{5.21}$$

$$p^{n+1} = p^n + \frac{\rho^n c^n}{2}(\delta W_n^4 + \delta W_n^5) \tag{5.22}$$

In order to set this boundary value, the magnitude of the change of both the incoming and outgoing characteristic waves must be known. The outgoing characteristic wave magnitude change can be found by [147]:

$$\delta W_n^4 = -\delta t(u_{n,P}^n + c_{n,P}^n)\Big(\frac{u_n^n - u_{n,P}^n}{\delta x_n} + \frac{1}{\rho_P^n c_P^n}\frac{p^n - p_P^n}{\delta x_n}\Big) \tag{5.23}$$

Only first order accurate schemes in time are considered here. The primary reason for this is that the timestep is limited in this explicit methodology by the CFL number driven by the acoustic field, rather than the convective velocity. In the case of incompressible flows the efficiency of explicit methods are often limited by the restriction that:

$$CFL = \frac{u_n \delta t}{\delta x} \leq 1 \tag{5.24}$$

This is because velocity waves cannot move further than one grid space per timestep in order to be correctly represented and to help prevent numerical instability. Once the flow becomes compressible the explicit solver must now track the evolution of velocity, entropy and sound waves. As sound waves move at the speed of sound the CFL condition becomes;

$$CFL_c = \frac{(c+u)\delta t}{\delta x} \leq 1 \tag{5.25}$$

As in the case of the injector acoustic measurement, the main duct flow has an average velocity of the order of 1 m/s and the speed of sound is of the order of 300 m/s, the timestep must be two orders of magnitude smaller in the compressible case compared to the incompressible. The size of the allowable time step is also linked to the grid spacing, with larger grid spacing leading to a larger timestep. The spatial resolution that is required to represent the acoustic waves is dependent on the spatial order of the numerical scheme. A high order scheme will allow for fewer grid points to be used per wavelength while a lower order scheme will require a denser mesh. The maximum allowable timestep can now be related to the number of mesh points per wavelength by:

$$\delta t \leq \frac{1}{f}\frac{\delta x}{\lambda} \tag{5.26}$$

This requirement shows that the maximum timestep can be linked to the maximum frequency of the sound wave and the mesh spacing compared to the size of the wavelength which is in turn linked to the spatial order of the numerical scheme. The acoustic wave is however, not the wave with the shortest wavelength, the shortest wavelength within the flow is the wavelength associated with unsteady flow features such as vortices or if they are resolved, the turbulent eddies. These flow structures are much smaller than the acoustic wavelengths of the range of frequencies of interest and hence restrict δx much more greatly than the presence of sound waves. Because δx is restricted by the velocity fluctuations, the corresponding maximum timestep is significantly smaller than that which is required to accurately resolve the acoustic field. As such a higher order temporal scheme is likely to improve the accuracy of the non-acoustic velocity field but will have limited effect on the accuracy of the acoustic field without an associated improvement of spatial order of accuracy. Typical timesteps used for non-acoustic flows in this study are $\delta t \approx 10^{-5}\ s$ while for acoustic flows $\delta t \approx 10^{-6}\ s$.

Despite the apparent lack of sensitivity to the order of the temporal discretisation scheme, some benefit has been identified in using a second order scheme in space. By assuming that the gradient of velocity and pressure is calculated to second order within the cell, a second order upwind approximation of the gradient at the boundary can be derived by writing the Taylor series expanded in the x-normal direction at both the boundary and the centre of the boundary cell. The terms of order δx_n^2 can be eliminated resulting in the following second order approximation of the gradient:

$$\frac{\partial p}{\partial x_n} = \frac{3}{2}\frac{p_n - p_{n,P}}{\delta x_n} - \frac{1}{2}\frac{\partial p}{\partial x_n}\bigg|_{x=x_P} \tag{5.27}$$

Such an approximation does improve the accuracy of the scheme (see appendix C) however it significantly reduces the robustness of the method. The fact that the gradient is calculated as the difference of two other gradient approximations makes it susceptible to short wavelength instabilities. If there is a large variation in pressure between the cell boundary and centre then either the upwind or cell centred gradient may have the incorrect sign and drive the pressure or velocity in the wrong direction.

Curvature correction

In cases where the boundary is curved, the characteristic boundary conditions are corrected to account for the boundary curvature. The equations for the outgoing and incoming waves at the boundary become:

$$\frac{\partial W_n^4}{\partial t} + (u_n + c)\frac{\partial W_n^4}{\partial n} + \frac{2u_n c}{R_b} = 0 \tag{5.28}$$

$$\frac{\partial W_n^5}{\partial t} + (u_n - c)\frac{\partial W_n^5}{\partial n} + \frac{2u_n c}{R_b} = 0 \tag{5.29}$$

where R_b is the radius of the boundary.

5.3.1 Subsonic Total Pressure Inlet

In the case of a subsonic inlet boundary condition the characteristic equation has four positive eigenvalues and one negative. This implies that the value of four boundary conditions must be imposed in order to obtain the flow solution: the two components of velocity in the plane of the boundary, and two thermodynamic variables, chosen in this case to be pressure and temperature. The four characteristic waves entering the domain are representative of an entropy wave, an acoustic wave and the velocity components acting in the plane of the boundary. The final characteristic wave is the acoustic wave leaving the domain. By assuming that the velocity is zero in the plane of the boundary, only the velocity normal to the inlet will be modified by the acoustic waves. The boundary temperature is assumed to be constant but a simple mathematical analysis shows that the effect of the acoustic waves on temperature is minimal. Therefore in order to calculate the required pressure and velocity at each boundary, the magnitude change of the characteristic wave entering the domain (δW_n^5) must be derived and combined with the magnitude change of the outgoing characteristic using equations 5.20, 5.19 and 5.23.

Non-Reflective Inlet

In most situations it is desirable to set the inlet boundary condition to be acoustically non-reflective. The total pressure inlet usually sets the inlet velocity by calculating the magnitude of velocity from the difference between the calculated static and desired total pressures, while the direction of velocity is always imposed normal to the boundary surface. This method of imposing velocity is equivalent to the imposition of a fixed velocity and is similar acoustically to the interaction of a sound wave with a wall, a velocity node. In this case the pressure will fluctuate at the inlet as a wave hits it causing the wave to reflect back into the domain. The two-microphone technique requires that only one incident and one reflected wave can be within the duct at once, a requirement that would be violated if the acoustic waves were allowed to reflect off the inlet and back towards the injector. The requirement for a non-reflective boundary is that the incoming characteristic wave magnitude is set to zero:

$$\delta W_n^5 = 0 \tag{5.30}$$

In this case the static pressure and hence total pressure is not fixed and the static pressure and velocity are allowed to fluctuate to allow any acoustic waves to pass unhindered through the boundary and out of the domain. This can lead to pressure drift where the static pressure at the boundary changes over time altering the original required boundary condition [50]. One solution is to use the feedback approach describe by Hirsch [50] as:

$$\delta W_n^5 = -K_p(p^n - p_\infty) \tag{5.31}$$

The problem with this methodology is that it is only non-reflective in the case when the static pressure at the boundary is equal to the desired value and also that the inlet total pressure is prescribed, the desired static pressure value is dependent on the mass flow rate through the injector, a quantity that may change as acoustic waves are applied to the domain. In order to take this into account the local velocity is used to replace the upstream static pressure with the upstream total pressure:

$$\delta W_n^5 = -K_p(p^n - p_{0,\infty} + \frac{1}{2}\rho|\mathbf{U}|^2) \tag{5.32}$$

The value of K_p should be chosen such that it is sufficiently close to zero to prevent the reflection of acoustic waves. Typical values are 10^{-6}–10^{-5}.

Pressure Node

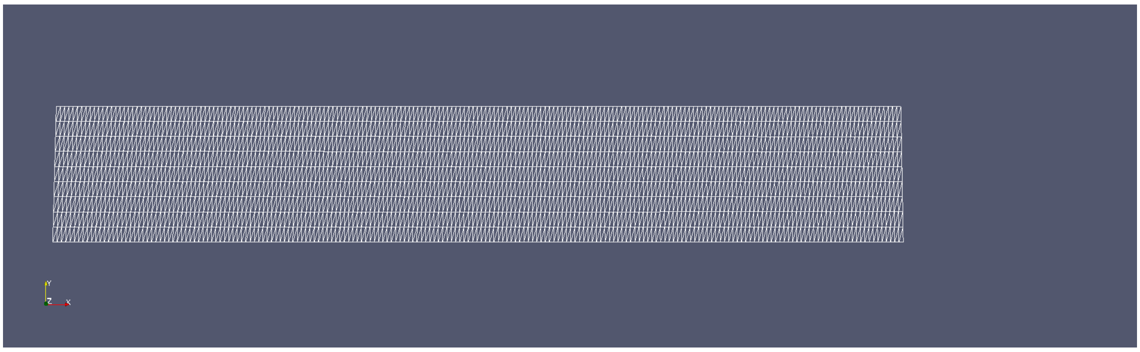

Figure 5.1: Simple duct mesh used for the testing of the characteristic boundary conditions

Whilst not used for calculations of acoustic impedance, the inlet was converted into a reflective type inlet to test the ability of the numerical method to accurately convect acoustic waves across a domain. The upstream boundary condition for a pressure node implies (from Equation 5.22) that the magnitude change of the incoming characteristic wave should be [147]:

$$\delta W_n^5 = -\delta W_n^4 \tag{5.33}$$

This will then result in a standing wave forming along the duct.

5.3.1.1 Velocity Node

Similar to the pressure node simulation, a velocity node can be induced at the upstream boundary (see Equation 5.19) by setting [147]:

$$\delta W_n^5 = \delta W_n^4 \tag{5.34}$$

5.3.2 Subsonic Forced Pressure Outlet

The subsonic outlet has four negative eigenvalues and one positive. In this case only the incident acoustic wave magnitude must be set.

Sinusoidal Forcing

In order to generate acoustic waves, the downstream pressure boundary condition was converted to a forced pressure outlet whereupon recognition that $\delta u = -\delta p/(\rho c)$, the incoming characteristic wave is set to [147]:

$$\delta W_n^5 = \frac{-2|\hat{p}|(\sin(\omega(n+1)\delta t) - \sin(\omega n \delta t))}{\rho^n c^n} \tag{5.35}$$

The generation of sinusoidal waves at the pressure outlet allows the testing of acoustic impedance of an injector using a numerical version of the multi-microphone method.

A simple square section duct was chosen as a test case for the characteristic boundary conditions in PRECISE-UNS. The duct dimensions are 6.25 m × 1 m × 1 m. A structured mesh was generated with 200×10×10 equally space nodes (see Figure 5.1). This resulted in a mesh spacing of 31.25 mm, 100 mm and 100 mm in the x, y, z directions respectively. A steady flow field was generated after 4000 iterations of the compressible PIMPLE algorithm within the PRECISE-UNS solver running with the k-epsilon RNG turbulence model and simple wall functions. The upstream boundary condition was a total pressure and total temperature of 101600 Pa and 287 K, while the downstream boundary condition was a pressure outlet set at 101325 Pa. After 4000 iterations the solver was switched into unsteady compressible mode and a 63 Pa 54.33 Hz acoustic wave generated at the downstream boundary using the method described in the following section. The upstream boundary was set to be a pressure node, similar to that found at a wall, except that the wall was allowed to be porous to maintain a bias flow.

Figure 5.2 shows the resulting pressure plot at different points along the duct as a function of time while Figure 5.3 shows a contour of pressure at a time corresponding to the wave's maximum magnitude. The same simulation was repeated except with using the velocity node upstream boudnary condition. Figure 5.4 shows the pressure field along the duct at one instant in time. It shows that the maximum pressure is achieved at the inlet where the acoustic particle velocity is constrained to be zero.

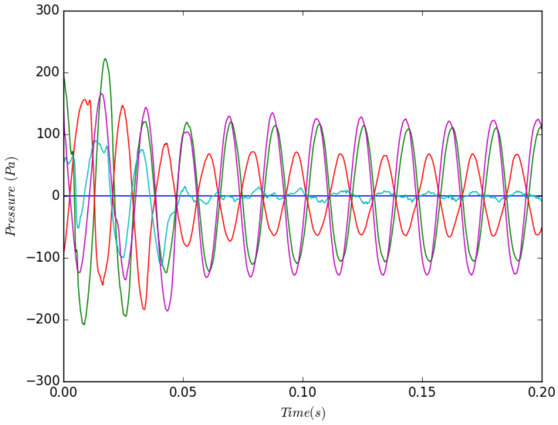

Figure 5.2: Trace of pressure measured at different lengths along duct assuming a pressure node upstream. Blue: x/L=0, Green: x/L = 0.25, Red: x/L = 0.5, Cyan: x/L = 0.75, Purple: x/L = 1.0.

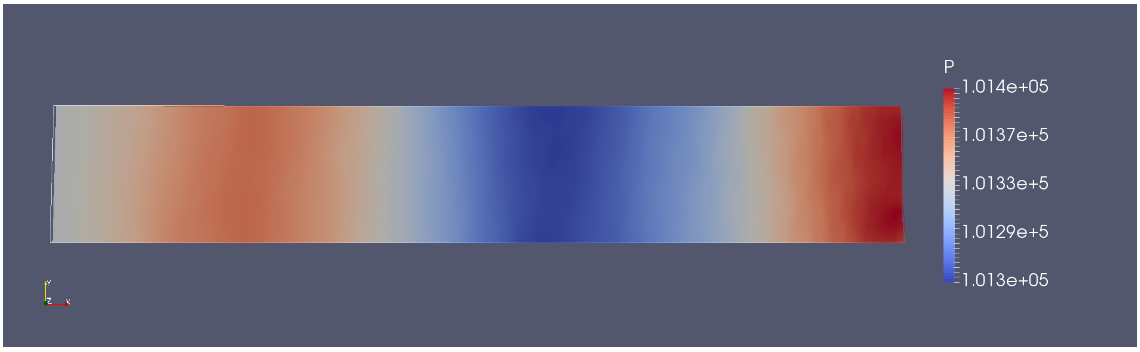

Figure 5.3: Snapshot of pressure field assuming a pressure node at the upstream (left hand side) boundary.

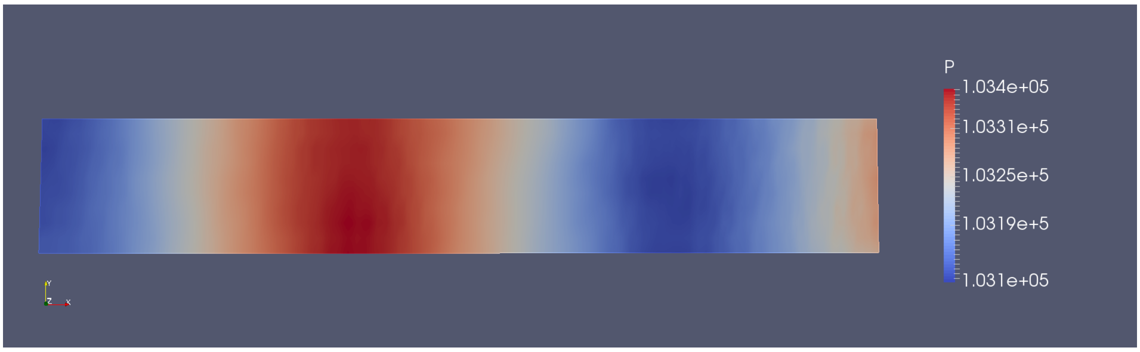

Figure 5.4: Snapshot of pressure field assuming a velocity node at the upstream (left hand side) boundary.

Multiple frequency forcing

The numerical computation of the acoustic impedance methodology, developed by Su et al. [149], has been shown to reproduce results from the multi-microphone method for a number of geometries of interest including simplified fuel injector geometries and orifices. The length of the time required to complete the simulations necessary to compute the acoustic impedance of a fuel injector is dependant on the frequencies(s) of interest but a typical geometry for impedance measurements between 50 and 600 Hz at a 50 Hz interval will require roughly 30000 CPUh. This cost is not prohibitive when compared to experimental methods, especially in the case of prospective injector designs that would need to be manufactured through a rapid prototyping method before they could be tested.

The efficiency of the method can be improved significantly through the recognition that the acoustic field is linear. This linear nature of the governing equations allows the interaction of the injector with the acoustic field to be deconvoluted from one signal into its fundamental frequencies. Several methods within system identification exist using this principle where a suitable input signal with a broad band frequency spectrum is applied to the system and the response of the system measured using a suitable method [76]. Such methods include impulse methods, where an impulse function such as a step change or Dirac delta function is applied, or noise methods, where a noisy broad band signal is applied over time or multiple sine waves are applied simultaneously.

White Noise Forcing

One possibility to decrease the cost of forced simulations by testing multiple frequencies of interest is to consider the application of sinusoidal waves of all the frequencies of interest at once. In this case he magnitude of the incoming characteristic wave becomes:

$$\delta W_n^5 = \frac{-2|\hat{p}|(\delta_{n+1} - \delta_n)}{\rho^n c^n} \tag{5.36}$$

where δ is a random number generated between -1 and 1. In the limit of infinitely small timesteps, this is exactly equivalent to a white noise signal, however as the timestep is finite the highest fundamental frequency present is equal to:

$$f_{max} = \frac{1}{\delta t} \tag{5.37}$$

where δt is the timestep. In the case of the simulations performed in this study, the timestep is so small that the highest frequencies injected into the domain are several order of magnitude higher than the frequencies of interest or what the mesh and solver are capable of resolving. Furthermore the presence of high frequency oscillations may induce numerical instabilities. As such it is prudent to low-pass filter the incoming characteristic waves. A true and perfect low pass filter cannot ever applied to a signal of finite length however a sufficiently accurate filter could be applied to the noise signal if the amplitude of the signal is known far into the future and past when the filter is applied. The problem with this methodology is that the amplitude of the incoming signal must be generated sufficiently in advance and saved for a sufficiently long time [46]. This requirement complicates the algorithm, increases computational cost and computational memory requirements.

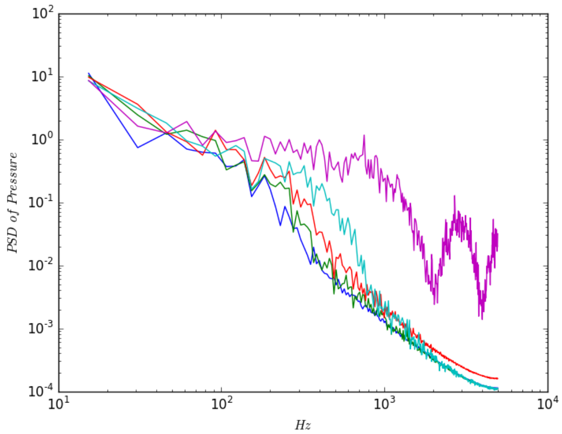

Figure 5.5: PSD of pressure measured at different lengths along duct for the case of 1000 Hz low pass white noise. Blue: x/L=0, Green: x/L = 0.25, Red: x/L = 0.5, Cyan: x/L = 0.75, Purple: x/L = 1.0.

If it is assumed that spurious high frequency components of the input signal will not adversely affect the stability or accuracy of the result it becomes attractive to apply a simple quasi low pass filter. The quasi low-pass filter simply increases the number of timesteps between the calculation of δ_{n+1} and δ_n such that Equation 5.36 becomes:

$$\delta W_n^5 = \frac{-2|\hat{p}|(\delta_{n+o} - \delta_n)}{\rho^n c^n} \tag{5.38}$$

where o is the number of timesteps chosen to skip and the maximum frequency present in the duct, neglecting spurious high frequencies induced by the imperfect quasi-low pass filter, is now:

$$f_{max} = \frac{1}{o\delta t} \tag{5.39}$$

Figure 5.5 shows the PSD of pressure taps along the length of the low resolution duct simulation in the case where the simulation is forced with white noise using the quasi-low pass filter method set at 1000 Hz. It can be seen that the coarseness of the mesh severely attenuates the signal frequencies over 300 Hz.

Frequency array forcing

The white noise method is attractive for reasons of maximising the amount of information that can be extracted from one simulation. However, the fact that the method excites every frequency between the inverse of the simulation length and the cut-off means that the Fourier coefficient of any particular frequency is very low. In order to avoid an unacceptable signal to noise ratio this requires therefore an extremely high excitation amplitude or extended run time. The maximum amplitude of the excitation is limited by the point at which the acoustics can not longer be considered linear. Therefore the extended run time may then increase the cost of the simulation such that the white noise forcing method becomes comparable with forcing at discrete frequencies. The experiments carried out at both Loughborough and Rolls-Royce are done at discrete frequencies and in some cases multiple frequency excitation of the injector is impossible due to the nature of the experimental apparatus (acoustic sirens). For this reason the FTF or acoustic impedance is not measured as a continuous function but instead as a set of discreet points in frequency space. Due to the assumption of linearity, one should be able to reproduce this FTF by exciting the flow at all frequencies simultaneously. Practically, this is done in exactly the same method as the *sinusoidal forcing* method however each additional frequency is added as an additional term to the calculation of δW_n^5 similar to Equation (5.35).

5.3.3 Subsonic constant mass flow inlet

Similarly to the subsonic total pressure inlet, the subsonic constant mass flow inlet must provide four boundary conditions. In this case the three components of velocity, and temperature are chosen. Again, the components of velocity in the plane of the boundary are set to zero while the velocity normal to the boundary cannot be set explicitly as it would correspond to a velocity node, reflecting the outgoing acoustic waves. Therefore the normal velocity and pressure must be modified according to the magnitude of the outgoing acoustic wave.

This boundary condition is of significant importance in reacting flow simulations or in simulations that will be used to generate boundary conditions for reacting flow simulations as the heat release is a strong function of mixture fraction while only weakly linked to pressure. As the URANS simulation is unlikely to capture the effective area of the injector exactly, it is more important that the mass flow rate through the injector is correct rather than the pressure drop.

Relaxed mass flow inlet

A truly non-reflective inlet is set in exactly the same way as in the total pressure case:

$$\delta W_5 = 0 \tag{5.40}$$

However, this boundary condition does not fix the mass flow rate at the boundary as only the gradient of velocity is set. As such the velocity must be relaxed towards the target value:

$$\delta W_5 = K_v(u_n - u_{in}) \tag{5.41}$$

The value of K_v needs to be set such that it is sufficiently close to zero but as the acoustic particle velocity is much lower than the acoustic pressure, typical values of K_v are higher, in the range of 10^{-4}–10^{-3}.

Inlet forcing

The inlet boundary condition can also be used to inject acoustic waves into the computational domain by adjusting the amplitude of the incoming acoustic wave according to Equation 5.42:

$$\delta W_5 = K_v(u_n - u_{in}) - \frac{2|\hat{p}|(sin(\omega(n+1)\delta t) - sin(\omega n\delta t))}{\rho^n c^n} \tag{5.42}$$

Smoothing

One of the disadvantages of using the characteristic and non-reflective boundary conditions described in Sections 5.3 and 5.3.1, is that it is assumed that locally the only fluctuations associated with an acoustic wave propagate in the face normal direction. This can cause problems if a local pressure perturbation moves upstream towards the inlet. The non-reflective boundary conditions sees this perturbation as an acoustic wave and adjusts the pressure and velocity to compensate. This then induces a gradient across the inlet that then causes fluid and acoustic waves to propagate along the inlet surface. The velocity and pressure perturbations that this induces are perceived as being acoustic waves impacting on the wall in the normal direction promoting a further localised pressure and velocity fluctuation. This effect can amplify to the point where the inlet becomes supersonic or the velocity negative and the boundary condition fails. In order to avoid this problem it is necessary to apply a small amount of damping at the inlet, especially in the case of complex geometries. The damping applied in this work is applied by taking the average velocity and pressure across the inlet and relaxing each inlet face and inlet cell towards this value:

$$u_n = (1 - \delta_s)u_n + \delta_s \langle u_n \rangle \tag{5.43}$$

$$u_{n,P} = (1 - \delta_s)u_{n,P} + \delta_s \langle u_n \rangle \tag{5.44}$$

where δ_s should be as small as possible and has been tested successfully down to $\delta_s = 0.05$. Strictly speaking this smoothing should only need to be added to the inlet face values to remove any transverse gradients. The problem is that large transverse gradients can create a disconnect between the boundary cell values and boundary inlet faces which are then propagated to the other cells across the inlet. Once the inlet cells and faces become sufficiently disconnected, the gradient calculation used in Equation 5.23 is no longer valid. Unfortunately such a smoothing procedure reduces the gradients between the boundary cells and the boundary faces, also invalidating the same measure of the gradient. For this reason, the gradient must be calculated differently.

Gradient blending

Due to the reasons highlighted above the upwind gradient at the boundary can be corrupted in the case of insufficient or over-enthusiastic smoothing. As such it is safer to average the gradients at the boundary with the gradient within the cell:

$$\frac{\partial p}{\partial x_n} = \frac{1}{2}\frac{p - p_P}{\delta x_n} + \frac{1}{2}\frac{\partial p}{\partial x_n}\bigg|_{x=x_P} \tag{5.45}$$

This is a 1.5 order approximation of the gradient at the point halfway between the boundary cell centre and inlet face. It therefore lowers both the order and accuracy of the gradient calculation at the boundary. It is however extremely robust as provided that the mean of the two gradients has the correct sign, the pressure and velocity at the boundary will converge.

5.4 Test Case 1: The LU acoustic Duct

Experiments and simulations were performed of the injector being forced by an upstream propagating wave as generated by loudspeakers downstream of the I1 injector. In the numerical case, the acoustic waves were introduced from the downstream boundary.

5.4.1 Experimental methodology

Experiments were undertaken from 50 Hz to 650 Hz at a 25 Hz interval where the amplifier gain was set such that at each measurement frequency the amplitude of the pressure signal measured by the Kulite closest to the injector face measured 135 dB. In some cases this amplitude was unreachable due to resonance within the system or because this Kulite was close to a pressure node, in these cases, the amplitude was set to the maximum allowable by the loudspeaker system. The lowest amplitude recorded was 125.8 dB. The maximum amplitude corresponds to a relative pressure fluctuation ratio to mean pressure drop of less than 4% and hence, no back flow is generated.

5.4.2 Numerical methodology

The simulations of the acoustic duct rig were performed in OpenFOAM by Jialin Su at Loughborough university. The approach follows the methodology developed by Su et al. [149, 150]. The LUBPAI mesh described in Chapter 2 is used with the $k - \omega$ SST model providing turbulence closure. The mean flow field was found first by applying a 3% pressure drop across the injector and iterating the flow field using the PIMPLE solver until convergence. The timestep was chosen as 1×10^{-4} s using the backwards Euler scheme to ensure rapid convergence and to dissipate any acoustic waves. The velocity field was solved using a second order TVD scheme with a first order upwind scheme chosen for the other variables.

Once the field was converged, the simulation was run using the PISO algorithm of Issa [53] with a timestep of 1×10^{-6} s for 5 ms plus 12 acoustic cycles using a second order backwards scheme. This small timestep ensures that there is negligible dissipation of the acoustic waves and that the CFL at the boundary remains less than one. This was repeated for each of the frequencies of interest. The frequencies tested were from 250 Hz to 550 Hz at 50 Hz intervals with the acoustic wave magnitude set to 300 Pa. The 300 Hz and 450 Hz simulations were extended to 30 acoustic cycles to allow for a more in-depth analysis of the unsteady flow field. The acoustic impedance of the injector was calculated using the multi-microphone technique by monitoring the pressure at the same locations as in the case of the experiment.

5.4.3 Results and discussion

Acoustic Impedance

Figure 5.6 shows the resistance and reactance calculated using the experimental methodology for the representative lean burn geometry chosen (labelled as I1) and for comparison, a more modern fuel injector (labelled as I2) as measured by Mark Cassell at Loughborough University.

Results from the acoustic impedance experiments show that the acoustic impedance of the injector geometry used in this study is similar to that found with a more modern fuel injector. The

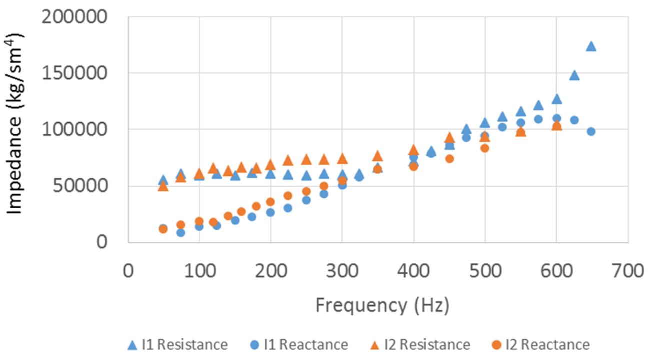

Figure 5.6: Acoustic impedance measurements made of the I1 and I2 injectors.

most noticeable deviation between the injectors occurs at medium frequency. The newer injector is significantly more complicated in shape and is likely to have stronger viscous forces acting on the flow within the injector body.

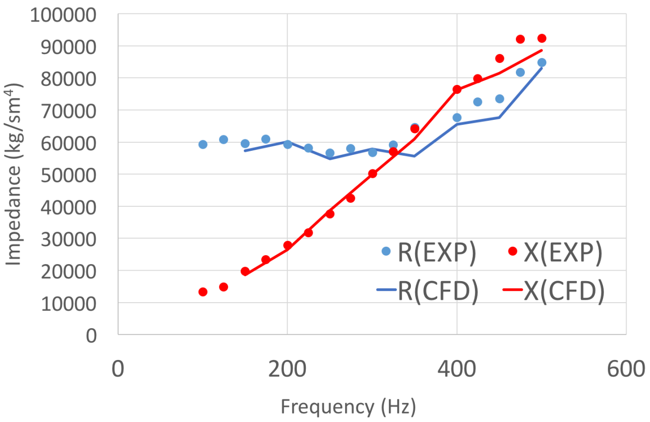

Figure 5.7: The acoustic impedance of the I1 injector as measured by experiment and simulation.

Figure 5.7 shows the impedance of the I1 injector as calculated using the experimental and numerical methodologies. It shows excellent agreement between the numerical prediction and the experimental data with only a small deviation at higher frequencies. This result is important as it validates the use of the URANS methodology for the calculation of the acoustic impedance for complex fuel injector geometries. Some analytical models exist for simpler geometries (see for example [11] or [78]) however they are not capable of reproducing the *roll up* of the resistance of *roll down* of the reactance at high frequencies. The model of Yang and Morgans [171] does reproduce some of these features however it cannot reproduce the complex interactions between each passage nor the fluctuations of swirl number as explained in the following section.

In the case of the numerical results the velocity field and pressure is known at the exit of the injector. For this reason the acoustic impedance of the injector can be calculated directly. At every time step the mass flow rate and average pressure was calculated at the exit plane of the injector and saved. After the simulation was run the acoustic impedance was calculated as:

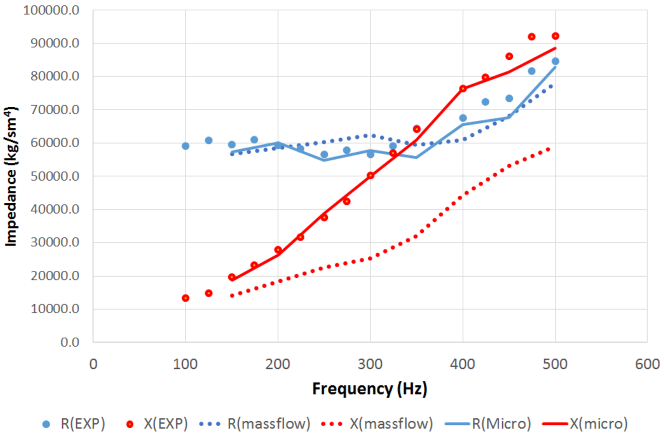

Figure 5.8: Acoustic impedance measurements made of the I1 injectors. Using the multi-microphone technique using data taken from simulations and experiments and compared with the impedance calculated directly from the area averaged velocity and pressure fields.

$$Z = \frac{\hat{p}^* \hat{\dot{m}}}{A_{\text{eff}} \hat{\dot{m}}^* \hat{\dot{m}}} \tag{5.46}$$

where the hat designates the Fourier transform of the signal and the asterisk is the complex conjugate. The result is plotted in Figure 5.8. It shows that the multi-microphone technique does an excellent job of predicting the resistance however the reactance is significantly lower in the case of this direct method. This means that the out of phase mass flow fluctuations are larger than expected at the injector exit. There are a few possibilities for why this is the case, firstly the Mach number close to the injector exit passages is significantly higher than in the main duct and therefore may invalidate the assumption that the group velocity of the wave remains constant close to the injector face. There also exists the possibility that the fluctuations in velocity induced by the incoming acoustic wave feed back to the acoustic energy by themselves generating an acoustic wave. It has been suggested by Jialin Su at Loughborough University that the assumption of linearity in the acoustic field is violated in the region close to the injector face due to the mean velocity in the recirculation zones being very low and therefore the ratio of acoustic fluctuations of velocity to the mean velocity being no longer small.

5.4.4 Fluctuations of pressure, mass flow rate, angular velocity and swirl number

The injector consists of three flow passages which all contribute to the fluctuation of mass flow rate observed. Furthermore, the fluctuating axial velocity leads to a corresponding fluctuation of angular velocity and swirl number through each passage and the injector itself. The pressure, axial and angular velocity are averaged over the planes shown in Figure 5.9, while the swirl number [10] is calculated over each of the planes according to:

$$S_i = \frac{\int_{R_{i,i}}^{R_{o,i}} u_\theta r \rho u_x r dr}{R_o \int_{R_{i,i}}^{R_{o,i}} u_x^2 \rho r dr} \tag{5.47}$$

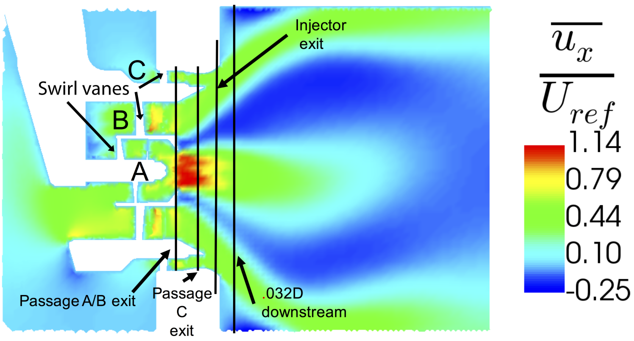

Figure 5.9: The mean velocity from the unforced compressible OpenFOAM simulation showing the locations of the planes where the pressure, mass flow rate, angular velocity and swirl number are calculated for each passage and the injector total.

where $R_{i,i}$ is the inner radius of the ith passage, $R_{o,i}$ is the outer radius of the ith passage and R_o is the outer radius of the injector. The magnitude of each of these fluctuations normalised by the mean value is shown in Figure 5.10a for the 300 Hz forced simulation. The fluctuation magnitude of pressure is less than one percent but results in a fluctuation of mass flow through each passage of around eight percent which leads to another eight percent fluctuation of angular velocity and around a 3 percent fluctuation in swirl number. The fluctuation of mass flow rate and pressure across the passages and the injector as a whole is very similar while in the case of the angular velocity and swirl number the fluctuation magnitudes are significantly higher in the case of the injector exit plane than the individual passages. Figure 5.10b shows the injector exit plane fluctuations plotted next to the average fluctuation values across the passages in the case of pressure and mass flow rate and plotted next to the sum of the fluctuation values in the case of the angular velocity and the swirl number. It suggests that the pressure and mass flow rate through the injector tend to average across the passages while the swirl number and angular velocities tend to sum.

Given that the mass flow fluctuation magnitude is similar across the passages, it follows that the swirl number fluctuations are proportional to the angular velocity fluctuations. Conversely, passage B has a large variation of angular velocity but a small variation of swirl number whereas passage A has a large fluctuation in swirl number and a smaller fluctuation in angular velocity.

Figure 5.11 shows the fluctuations of each quantity normalised by their magnitude for passages A and B. In the case of passage A there appears to be a comparatively large phase shift between the mass flow fluctuations and the angular velocity fluctuations. Because the swirl number is a balance between axial and tangential momentum, this phase shift results in a large fluctuation in swirl number as the axial velocity is reducing when the azimuthal velocity has reached a maximum. This phase shift is caused by the axial velocity fluctuation travelling through the injector at the

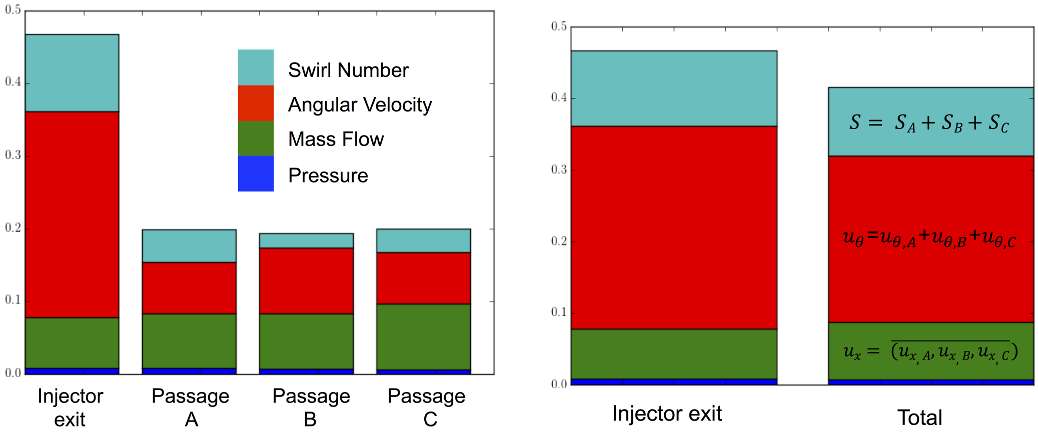

(a) The normalised fluctuation amplitudes for each passage and the injector total.

(b) The fluctuation amplitudes at the injector exit and the sum of swirl number and angular velocity and average of pressure and mass flow rate across the passages.

Figure 5.10: The fluctuation magnitudes of pressure, mass flow, angular velocity and swirl number normalised by their mean quantities for the compressible OpenFOAM simulation forced at 300 Hz.

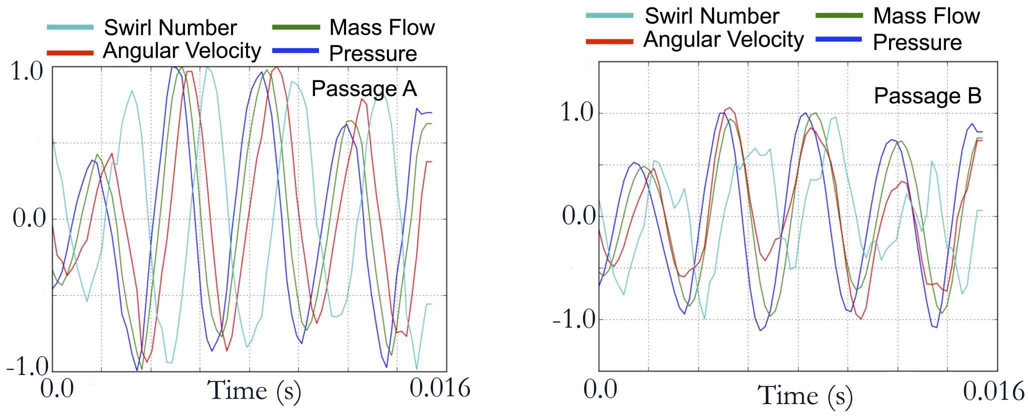

(a) Passage A.

(b) Passage B.

Figure 5.11: The normalised fluctuations of pressure, mass flow, angular velocity and swirl number through passage A and B for the compressible OpenFOAM simulation forced at 300 Hz.

speed of sound while the azimuthal fluctuations are convected at the mean flow axial velocity. The axial fluctuations are instigated by the incoming pressure wave that travels upstream to the flow swirler where the azimuthal fluctuations are generated. The fluctuations then travel downstream to the injector exit where they combine with the axial fluctuations to become swirl number fluctuations. This process is further described in Section 1.5.5. Looking now at the injector geometry (see Figure 5.9), the fluctuations in swirl number can be related to the distance between the exit of the flow passage and the swirl vanes. Passage A has a large distance, a large phase shift and a large swirl number response while passage B has a short distance, a small phase shift and a large swirl number response. As described in Section 1.5.5 the fluctuations in swirl number are capable of affecting the shape of the flame and the heat release. The flame transfer function for a swirl stabilised premixed flame shows a dip in the gain of the flame transfer function which corresponds to deconstructive interference of the effects of axial and azimuthal fluctuations. As the strength and phase of these fluctuations can be related to the swirler location, the location in the frequency range where this dip in the FTF occurs may be altered by moving the swirler locations backwards or forwards inside the injector body or by modifying the swirler response.

5.4.5 POD analysis

Similar to the unforced simulation, the simulations forced at 300 Hz and 450 Hz were analysed using POD. The helical structures seen in the unforced simulation were evident in both simulations. The velocity reconstruction of the the fist two modes for the 300 Hz forced simulation are visible in Figure 5.12b and the third and fourth modes in Figure 5.12a. These modes appear to have wavenumbers of $m = 1$ and $m = 2$ respectively however they appear to be scrambled somehow by the application of the forcing.

The first four modes were identified using the method described in Section A.3.2 to be $m = 1, 2, 2$ and 2 respectively. The PSD of these four modes were then considered where it can be seen that the first two modes and third and fourth modes are clearly paired (see Figure 5.14a) furthermore the third and fourth modes have a higher peak in the region of 800-1000 Hz corresponding to the PSD of the $m = 2$ modes in the unforced simulation. For this reason the first two modes are assumed to be more closely aligned with $m = 1$ modes and the third and fourth more representative of $m = 2$ modes.

The difficulty in calculating the dominant azimuthal wavenumber is that the operation depends both on space and component of velocity. The quantity used to characterise the wavenumber is ψ_m which is a function of both radius and axial location and contains three components $\psi_m(x, r) = [\psi_{m,x}, \psi_{m,r}, \psi_{m,\theta}]$. The helical modes are created due to shear instabilities leading to linearly unstable regions of the flow with clearly defined azimuthal wavenumbers and axial wavelengths [58, 152]. However as the flow develops downstream and the magnitude of the instability grows, the flow becomes progressively more non-linear until the structure decays into chaos. This means that the coherence of the structure will decrease in space and time, therefore the dominant wavenumber of a structure may change leading to regions with different dominant wavenumbers. The structures may also be separated by components where the structure may comprise of very strong fluctuations in the axial direction but much weaker ones in the radial or azimuthal directions. This was apparent in the calculation of azimuthal wavenumber where different components of velocity often

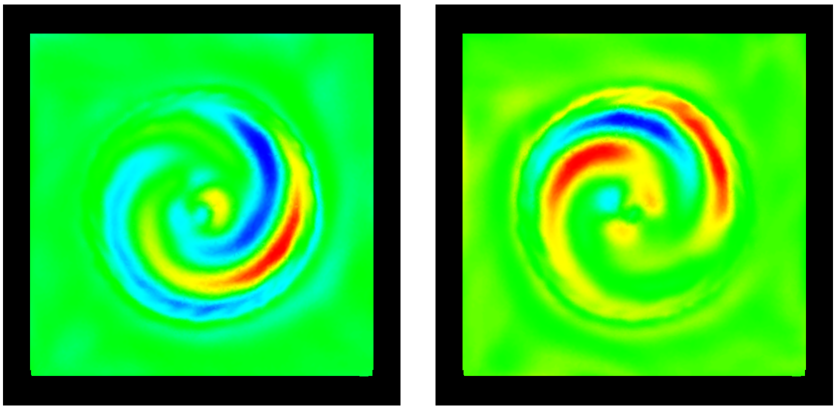

(a) The $m = 1$ swirling mode generated from the two most energetic POD modes.

(b) The $m = 2$ swirling mode generated from the 3rd and 4th most energetic POD modes.

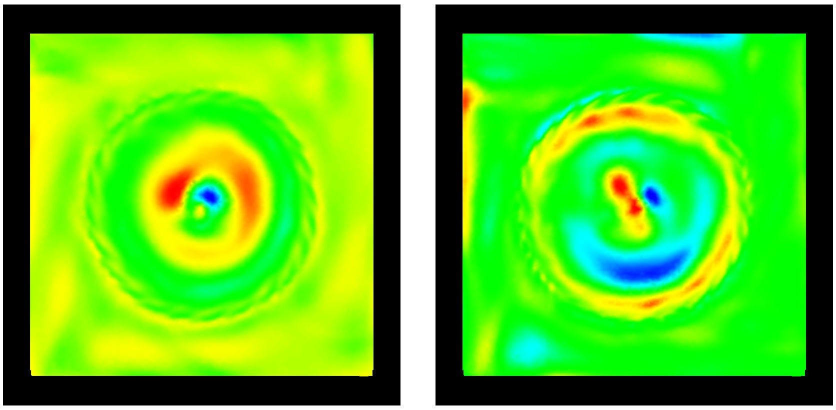

(c) The first of two symmetrical forced ($m = 0$) modes generated from the 6th and 7th most energetic POD modes.

(d) The second of two forced modes ($m = 0$) generated from the 9th and 10th most energetic POD modes.

Figure 5.12: The reconstructed axial velocity field 2.5 mm downstream of the injector exit calculated from various POD modes from the 300 Pa, 300 Hz OpenFOAM calculation of the Loughborough acoustic duct rig. High velocity regions highlighted in red and low velocity regions highlighted in blue.

had different dominant wavenumbers. The integration of the wavenumber magnitude in space is shown in Table 5.1 however these values have sometimes been ignored due the possibility of errors as mentioned and given the appearance of the PSD of the corresponding temporal modes. Given this, the scrambling of the modes that hinders their identification would likely be reduced with more snapshots.

Critics of this analysis have pointed out a number of methods to assist in cleaning up the data and making the mode identification easier. The first would be to perform the azimuthal Fourier transform on the raw data and then compute the POD on the filtered velocity fields corresponding to each azimuthal wavenumber. This would have the effect of forcing the structures identified to have a particular wavenumber. The second method would be through the use of Sieber et. al.'s spectral POD method (see [138]) or the third method would be through the use of DMD [131] (see Section 5.6.4). Both of these methods (DMD more so) assume that each structure identified will fluctuate at a given frequency. The imposition of a set azimuthal wavenumber or fluctuation frequency across the domain seems to violate the idea that the flow is non-linear and because of this structures should be expected to fluctuate in shape and frequency.

Table 5.1: Dominant azimuthal wavenumbers for the first 10 POD modes for unforced and forced simulations

Forcing Frequency	Unforced	300 Hz	450 Hz
mode 1	1	1	1
mode 2	1	2	1
mode 3	2	2	2
mode 4	2	2	1
mode 5	0	1	1
mode 6	0	1	1
mode 7	0	0	0
mode 8	0	0	0
mode 9	0	0	0
mode10	0	0	0

In the case of the 450 Hz simulation the first two modes were identified using the PSD as being more representative of $m = 2$ modes and the third and fourth as being representative of $m = 1$ modes. Table 5.2 shows the proportion of kinetic energy contained within each mode type for the unforced, 300 Hz forced simulation and the 450 Hz forced simulation. It shows that the proportion of kinetic energy in the $m = 1$ helical mode has been augmented in the case of being forced at 300 Hz and reduced in the case of being forced at 450 Hz, the opposite is true for the $m = 2$ mode. The mechanisms by which the strength of these modes may be altered by the acoustic forcing is discussed in Section 5.6.

The temporal components of the helical modes are plotted against each other in Figure 5.13 for the unforced calculation and the simulations forced at 300 Hz. The shape of the orbits between the forced and unforced calculations seem to be similar, however the stationary point or attractor appears to be moving diagonally from top right to bottom left in the case of the forced $m = 1$

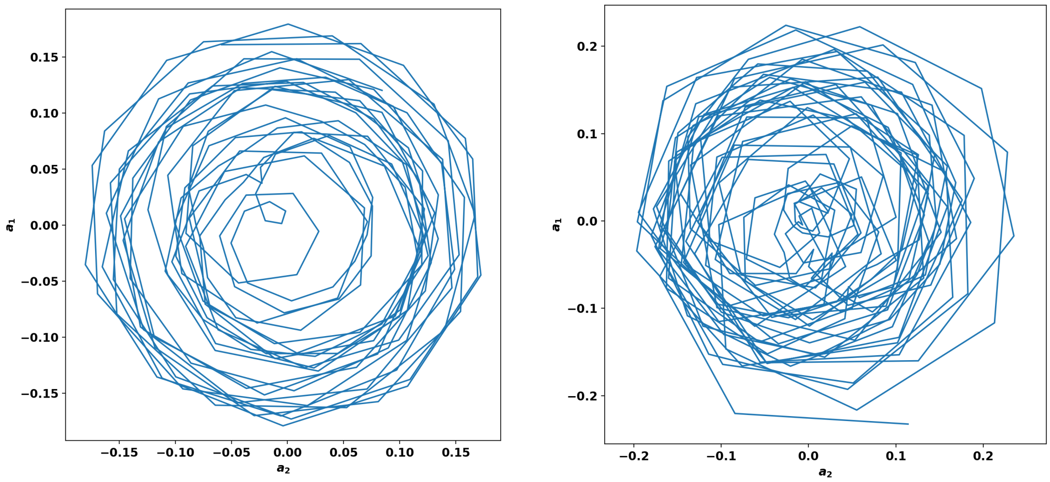

(a) 1st and 2nd modes plotted against each other for the unforced simulation.

(b) 1st and 2nd modes plotted against each other for the 300 Hz forced simulation.

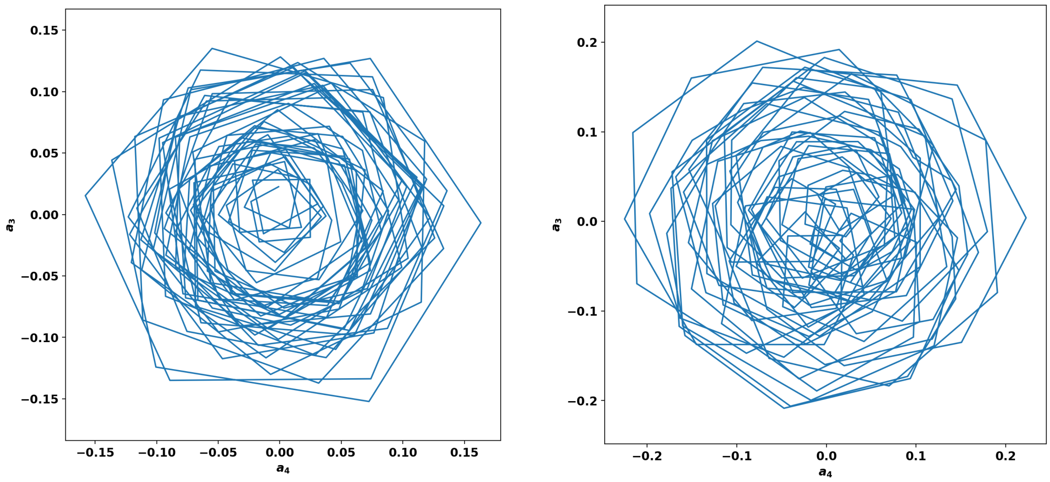

(c) 3rd and 4th modes plotted against each other for the unforced simulation.

(d) The third and fourth modes plotted against each other for the 300 Hz forced simulation.

Figure 5.13: Some of the temporal POD mode amplitudes for the forced and unforced cases

mode (Figure 5.13b) and vertically up and down in the case of the forced $m = 2$ mode (see Figure 5.13d). The attractor remains close to stationary in the unforced case for both mode types.

Table 5.2: Energy contained in the two highest energy modes of each azimuthal wavenumber and the total resolved kinetic energy for each simulation expressed as a percentage of the energy resolved in the unforced simulation.

Forcing frequency	Unforced	300 Hz	450 Hz
m = 1	22.10	27.49	12.61
m = 2	12.16	11.17	16.83
Forced modes	N/A	11.05	7.95
resolved kinetic energy	100	199	105

In the case of the 300 Hz simulation the 7th, 8th, 9th and 10th most energetic modes contained a strong peak in their PSD at the forcing frequency. The azimuthal wavenumber analysis identified these modes as being axisymmetric and the reconstructed velocity field of the 7th and 8th or 9th and 10th modes shows that these modes are representative of the fluid being pulled backwards and forwards through the injector. It is expected that if the injector was axisymmetric that these four modes would collapse into two. However, due to the scarfing on the back face of the injector, that assists in the removal and installation of the injector in the engine, the length of the injector is shorter on the side opposite the fuel feed arm. This explains the asymmetry of the mode shown in Figure 5.12d. In the case of the 450 Hz forced simulation, only three modes were identified by their PSD as being forced modes. It is likely that there is a fourth mode that has not been identified due to there not being a sufficient number of snapshots in the analysis. This is supported by Table 5.2 that shows that the energy contained in these three modes is close to three quarters of the energy contained in the four forced modes identified in the 300 Hz simulation.

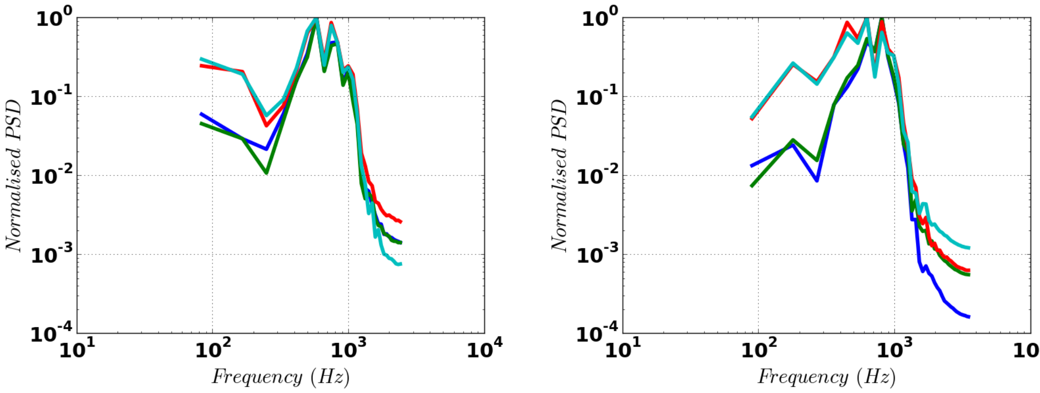

(a) Forced at 300 Hz.

(b) Forced at 450 Hz.

Figure 5.14: The PSD of the first (blue), second (green), third (red) and fourth (cyan) most energetic POD modes.

5.5 Test Case 2: The SCARLET rig

5.5.1 Experimental methodology

The SCARLET rig has four acoustic sirens with two placed upstream and one downstream of the combustion chamber. This allows measurement of the full acoustic transfer matrix of the injector combustion chamber system under reacting and non-reacting conditions. The acoustic sirens are comprised of a spinning perforated disc that separates the interior of the rig from a high pressure source. They can produce extremely high magnitudes of sound waves, however unlike loudspeakers they can produce sound at only one frequency. The rig upstream pressure and temperature were elevated to engine representative conditions and the frequency of the sirens stepped through a range of frequencies from $St = 0.05$ to $St = 0.60$. Due to the assumption of system linearity the combustion chamber was excited simultaneously from upstream and downstream but at different frequencies. It was ensured that the frequencies chosen were not harmonics of each other to ensure the signal could be accurately decomposed. The pressure signal was monitored at the four microphones located upstream and the four microphones located downstream. The multi-microphone technique was then used to calculate the magnitude and phase of the acoustic waves in each of the ducts and from that, the acoustic transfer matrix. The experiments were completed at DLR (German Aerospace Centre) Cologne by Claus Lahiri and André Fischer from RRD together with technicians from the DLR.

5.5.2 Numerical Methodology

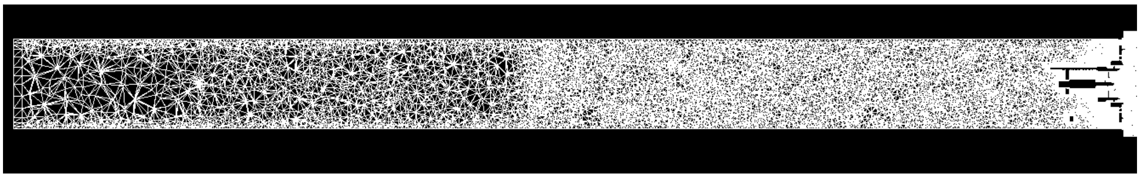

Figure 5.15: The upstream sponge mesh from inlet until injector.

Being that PRECISE-UNS was still under development, and in order to reproduce the acoustic transfer matrix using a numerical methodology, a mesh was created that was small enough that it could be run relatively quickly but with sufficient resolution in the injector and ducts to provide a reasonable representation of the acoustic and velocity fields. The mesh was chosen to have similar resolution around the injector as in the LUBPAI mesh and the same resolution in the upstream and downstream ducts relative to the speed of sound. Furthermore, the mesh was extruded downstream similarly to the LUBPAI mesh but the upstream duct was generated using the octree method provided in ICEMCFD. In order to facilitate good absorption of the acoustic waves travelling upstream, the upstream mesh was extended significantly upstream and a set of density boxes added to progressively coarsen the mesh in the upstream direction (see Figure 5.15). Finally a single layer of cells was extruded to ensure that the gradient would be accurately captured at the inlet. The mesh was generated using the RRD methodology with a small *worst aspect ratio* to ensure smooth transition between the very largest and smallest cells.

The simulation was iterated at a constant mass flow rate corresponding to a mean pressure drop of 3.8% using the PRECISE-UNS unsteady compressible solver for 2000 iterations with a timestep of 1×10^{-4} s to ensure rapid convergence and to dissipate any acoustic waves. A backwards Euler method was used for the temporal integration and a first order upwind spatial scheme was used for all variables.

Four simulations were run, one with upstream acoustic forcing at $St = 0.14$ and one with downstream acoustic forcing at $St = 0.14$ and then these two simulations were run again at multiple frequencies using a frequency array method where the frequencies were the same as those used in the experiment up to $St = 0.60$. The simulations were run for 50,000 timesteps plus however many steps were needed to overcome any transients. A second order backwards method was used for the temporal integration with a timestep of 1×10^{-6} s. A constant mass flow inlet was imposed with a non-reflecting boundary condition, the relaxation factor K_v was set to 0.005 and the inlet boundary smoothing factor δ_s was set to 0.05. The downstream boundary condition was set to a non-reflective pressure outlet with no relaxation. Enforcing the mass flow rate is important for reacting simulations as it ensures the correct air-to-fuel ratio in the combustion chamber however it also prevents the inlet from being truly non-reflective and very long transients where the pressure in the upstream duct equalises. Transients may take 100,000 steps or even more, however once the mean pressure field is established the forcing frequency, magnitude or direction can be changed with little additional transient calculation.

The simulation was run with a first order upwind convection scheme to maintain numerical stability. This poses problems when considering the characteristic waves inside the combustion chamber. Four out of five characteristic waves convect in the direction of the mean flow and hence the upwind differencing will capture their progress correctly. The fifth characteristic wave is an acoustic wave that travels upstream. While the pressure gradient in calculated symmetrically, the upwinding of the velocity components will likely lead to increased diffusion of the acoustic wave.

5.5.3 Results and Discussion

The wave fits for the upstream and downstream forcing (shown in Figure 5.16) show that with sufficient forcing amplitude and sufficient running time a good wave fit is possible within the ducts. Figure 5.17 shows the acoustic transfer matrix as calculated in experiments and using two different sets of simulations. The first simulation set was forced at $St = 0.14$ only while the second set was forced using a frequency array spanning from $St = 0.05$ to $St = 0.60$. The results show significant errors for both types of simulation. Errors appear to be smaller from zero to $St = 0.3$ showing that for higher frequencies the mesh should be further refined. The gain of the transfer matrix is significantly better captured than the phase, especially for the T_{11} and T_{12} components.

In the case of the T_{22} component, which is the component used to derive the FTF, both the single frequency simulations and frequency array simulations show the lowest level of error. The errors have a number of potential sources: The first is the low quality of the wave fit, a better quality fit would produce better results. This suggests that additional timesteps would help in eliminating noise and reducing the effects of transients. Another problem is that the numerical geometry is different to the experimental in that the cooling holes that surround the combustion chamber are simply modelled by assuming a constant mass flow from the outer annulus into the chamber. In

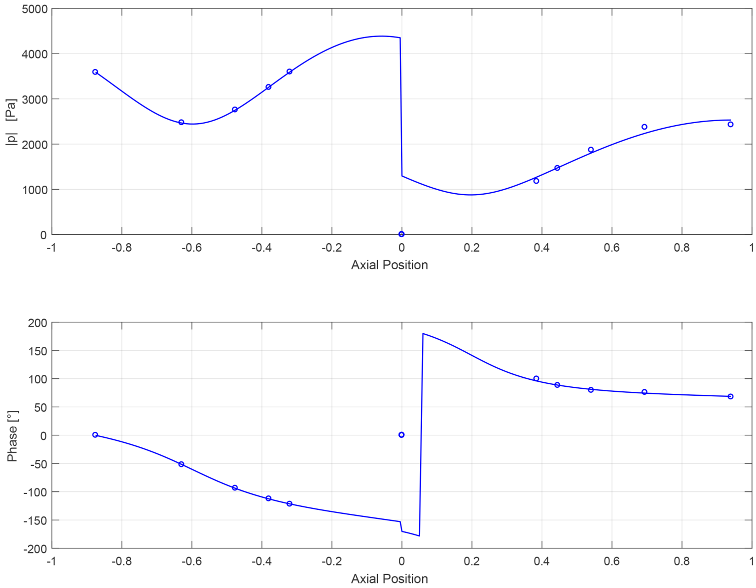

(a) Upstream forcing.

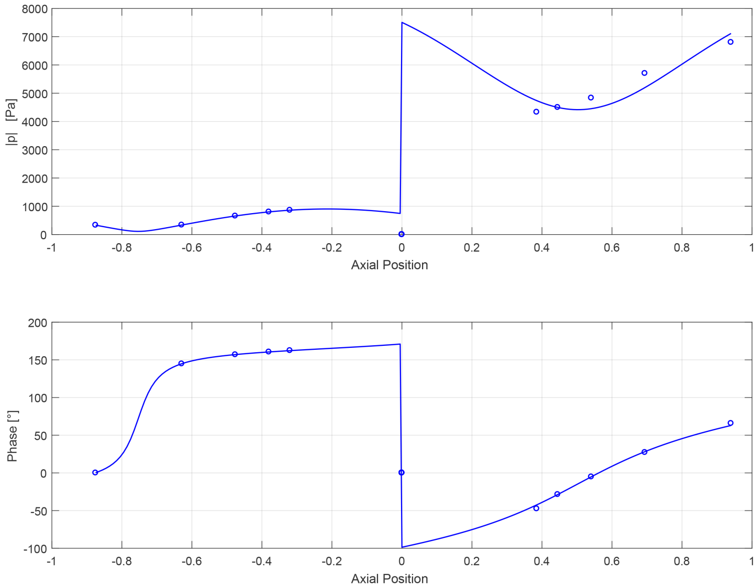

(b) Downstream forcing.

Figure 5.16: Wave fits from the non-reactive SCARLET simulation forced at $St = 0.14$. Circles are the magnitudes/phases of the individual pressure measurements and the lines are the standing wave fits.

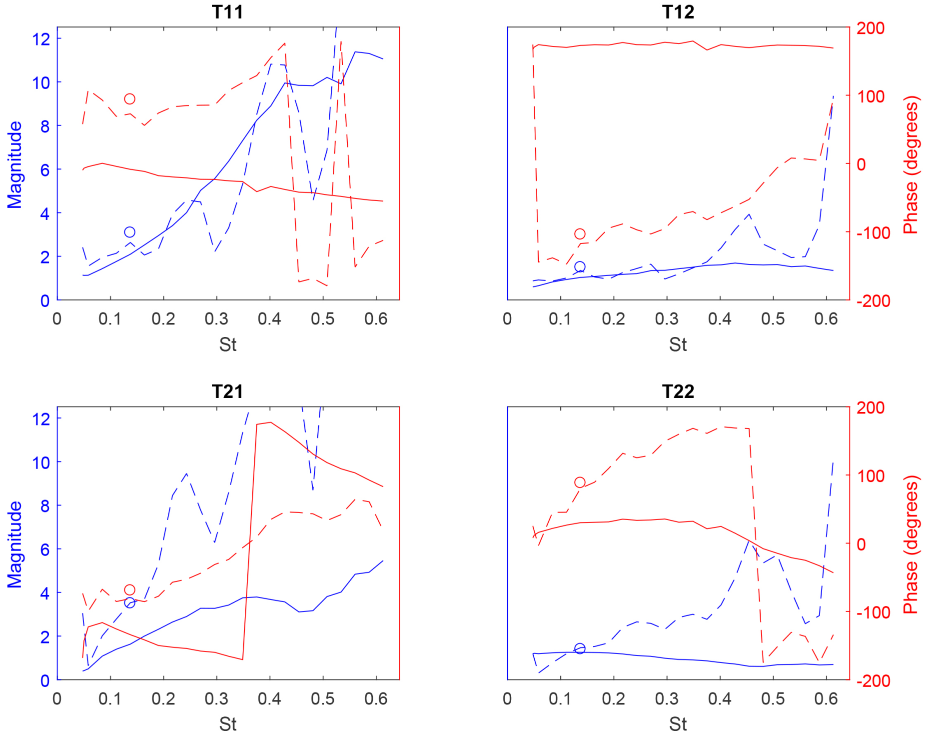

Figure 5.17: The experimentally derived transfer matrix for the non-reacting case (solid lines). The numerical results for 260 Hz are shown as circles while the frequency array results are shown with the dashed lines.

reality these holes will have their own acoustic impedance which would contribute to the result. The single frequency simulation results and frequency array results show very good consistency which shows that at this frequency errors are most likely due to the geometrical differences than the numerical or post-processing method. Furthermore, this also shows that the acoustic field is linear which means that for non-reacting simulations at least, many frequencies can be considered at once.

The difficulty in resolving the pressure field accurately, the quality of the wave fit and the use of first order gradients is strongly related to the mesh and geometry. The quality of the mesh could be improved and the removal of the sponge upstream will help, however the use of first order gradients is due to the solver becoming unstable in regions with a very large CFL number, primarily in the cooling holes for the heatshield around the outer radius of the injector. These holes can be replaced with an effusion boundary condition which simply remove the mass from one side and add it to the other but this increases the number of potential sources of error.

It would be useful to develop an effusion boundary condition within the code with an acoustic impedance that could be set using empirical values. This would then allow for small features to be removed and a smoother, more uniform mesh, significantly increasing the code performance. The code itself would also benefit from some more rigorous validation on simple geometries such as the orifice work of Su et al. [150].

5.6 Interaction pathways between acoustics and hydrodynamics

As previously mentioned in Chapter 4, it has been observed in both experimental and numerical studies that there can be interaction between the acoustic forcing and the naturally unstable modes of the combustor. It as been observed that there can be natural competition between the $m = 0$ forcing mode that seeks to quench the $m = 1$ or higher order helical modes [152]. The forcing modes and the helical modes may also interact producing an additional frequency peak in the measured pressure spectrum corresponding to the sum or difference between the peak frequencies of the helical and forced modes. Finally it has been observed that forcing may in fact amplify helical modes or trigger helical modes that would otherwise be suppressed. A set of theories about which mechanisms are responsible have been developed over the last few years and can be broadly classed as being either through linear [152] or non-linear interaction [60]. A pictorial of the different interaction pathways is offered in Figure 5.18. The figure shows how the fuel injector (in blue) acts as a device, with a given effective area and acoustic impedance, that acts to convert a given mean and phase varying pressure field into mean, phase varying and turbulent downstream velocity field. The different possible pathways for this conversion are described in the following subsections.

5.6.1 Mean flow correction

This is the first of two mechanisms described in [152]. The mechanism describes the diminution or amplification of the strength of the helical modes as being caused by an alteration of the local effective viscosity by the coherent velocity fluctuations excited by the acoustics, similar to the

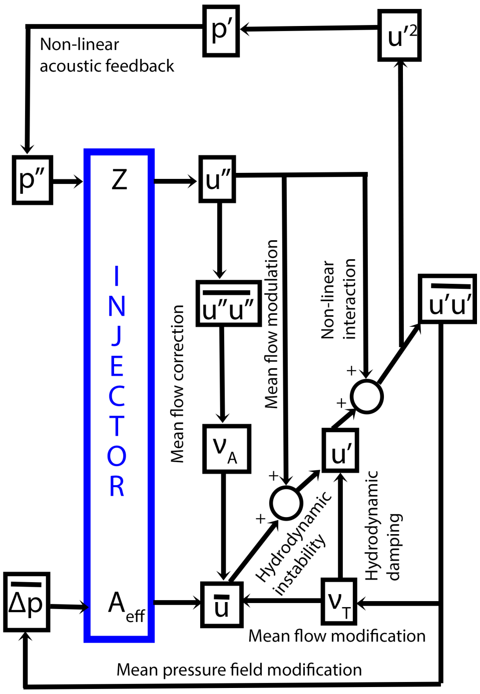

Figure 5.18: Acoustic interaction pathways

mean flow modification (Section 4.2.1). This modification of the effective viscosity also has an effect on the mean flow which then alters the stability of particular modes. The theory assumes no interaction between the acoustically driven fluctuations and the turbulence. This becomes troublesome if the hydrodynamic modes become significantly amplified due to the modified mean field. Such a large increase in unsteadiness would also increase the turbulence levels and feedback to the mean field through the *mean flow modification* mechanism described in Section 4.2.1. For this reason Terhaar et al. [152] refer to this mechanism in the case where the unforced flow is hydrodynamically stable (for helical modes) where the addition of forcing then causes the flow to become unstable.

The relative total resolved kinetic energy of the flow fields is shown for the unforced and two of the forced simulations in Table 5.2. In the case of the 300 Hz forced simulation the strength of all the unsteady non-turbulent fluctuations is almost double that of the unforced simulation. In the case of the 450 Hz forced simulation, the strength of these structures is only slightly higher than in the unforced case. It can be seen that the unresolved turbulent kinetic energy in the 450 Hz simulation is increased compared to the other simulations (see [153]). This might be explained by this mean field modification mechanism where the helical structures are damped by the increased effective viscosity generated by the acoustic velocity fluctuations.

The acoustic shift mode

The *mean field modification* theory really describes the existence of an *acoustic shift mode* analogous to the *shift mode* [101] described in Section 4.2.1. It operates through exactly the same mechanism whereby the acoustic forcing, rather than the hydrodynamic unsteadiness, feeds back to the mean flow through the modification of the flow effective viscosity. Just like the *shift mode* the *acoustic shift mode* can be calculated by subtracting the unforced mean field from the acoustically forced mean field. Figure 5.19 shows the acoustic shift mode between the unforced and 300 Hz forced simulations.

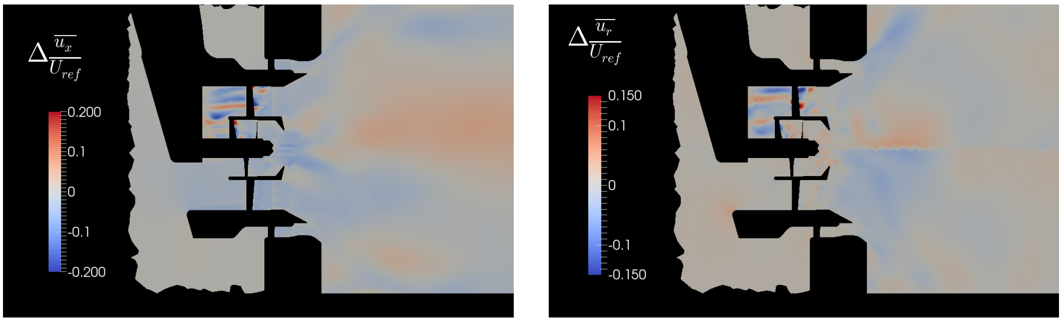

(a) Axial velocity.

(b) Radial velocity.

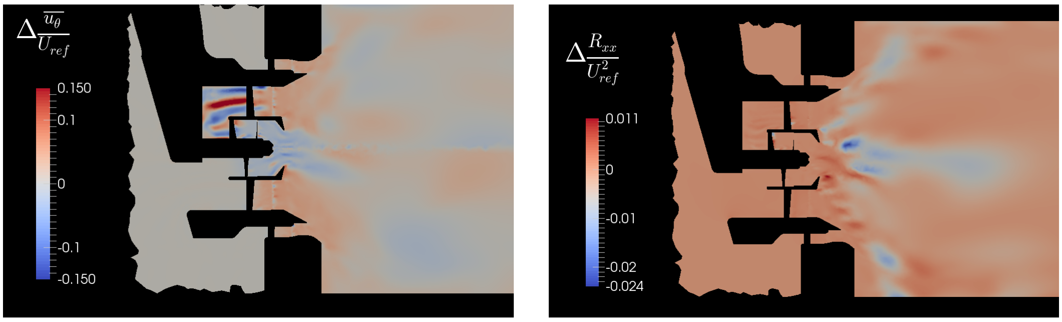

(c) Azimuthal velocity.

(d) R_{xx}

Figure 5.19: The acoustic shift mode

Along with reduced non-dimensional axial and azimuthal velocities, the largest difference in the velocity fields occurs behind the fuel feed arm in passage B, implying significant interaction between the acoustic waves and the recirculation zone behind the fuel feed arm. This results in the pilot jet being tilted upwards through the imposition of the acoustic forcing. The largest change to the fluctuating field is through the xx-component of Reynolds stress (Figure 5.19d). It shows that the majority of changes to the unsteady flow field occur on the outer edge of the pilot jet and inner edge of the mains jet. If the *mean field modification* mechanism was responsible for the change in helical mode stability then significant turbulent kinetic energy should be visible in the regions where the acoustic field causes direct variations in vorticity, the edges of the geometry. As the largest variation in the fluctuating flow field occurs further downstream, it makes more sense that the acoustic waves are modifying the growth rate of the instabilities.

5.6.2 Mean field modulation

The second of the two theories describing the observed change of the strength of the helical instabilities given by Terhaar et al. [152] is referred to as mean flow modulation. The mechanism is presumed active in cases where the mean flow field and therefore the mean flow stability remains relatively unchanged in the presence of forcing, however the strength of the observed helical instabilities is damped. Terhaar et al. [152] show that the some of the observed behaviour can be replicated using a parametric Van der Pol (VdP) oscillator [156]:

$$\ddot{q} - \mu(1 - q^2)\dot{q} - \gamma(t)\dot{q} + (1 - \beta(t))q = 0 \tag{5.48}$$

where the amount of non-linearity is governed by the choice of μ while $\beta(t)$ controls the eigenvalue of the system and $\gamma(t)$ alters the amount of system damping. If the damping and eigenvalue are both set such that:

$$\beta(t) = \gamma(t) = \hat{\zeta} sin(2\pi f_a^* t) \tag{5.49}$$

then the behaviour of the oscillator can be shown to replicate several effects of forcing on the behaviour of the forced combustion chamber flow. Numerical integration of Equation 5.48 with $\mu = 0.4$ showed that as the value of $\hat{\zeta}$ is increased the amplitude of the oscillator begins to decrease and begins to show several interaction frequencies ($f_h + f_a$,$2f_a - f_h$,$f_h - f_a$). This behaviour is most apparent when the forcing frequency is low, while when the forcing frequency is twice the natural frequency the system resonates to high amplitude. The natural frequency exhibits *lock-on* effects over the range of where resonance occurs. Part of the motivation for this theory comes from the observation that the temporal POD modes representative of the helical modes have a peak in their PSD spectrum at the forcing frequency, showing that they are being modulated. In contrast the Van der Pol oscilator did not reproduce the forcing frequency.

Phase averaged stability

Such a mechanism is also described by Terhaar [152] as being an oscillation of the eigenfrequency and growth rate of the modes due to the forcing. One possible interpretation of this mechanism is that while the mean flow remains constant the forcing causes the flow stability to change with respect to the phase averaged flow field. It is not easy to imagine how a phase averaged modified field stability contributes to the supression of helical modes because the evolution of the instability into a helical mode takes time. It is however easy to conclude that if during one part of the phase cycle the helical mode is stabilised, the mode is likely to be damped as this will lengthen the time it takes for an instability to develop. According to the mean flow modulation theory it would seem that the moments of increased flow stability contribute more then the alternate moments of increased instability. If the instability is enhanced by forcing then the opposite should be true. One test of this theory would be to calculate the stability of the flow at different phase averaged points in the acoustic cycle. If the instability is enhanced then it should be evident that for part of the acoustic cycle then the flow stability should be reduced and visa-versa.

5.6.3 Non-linear interaction and triggering

One very interesting possibility is the situation where a helical mode is unstable or marginally stable (convectively but not absolutely unstable). Acoustic forcing is applied but its not sufficiently strong to vary the mean field by itself. It does however interact with the helical mode and significantly amplifies its magnitude, or if the mode is marginally stable, provide enough energy to trip the mode from linear stability to non-linear growth. Such conditions may be rare in practice, however Kirthy et al. [60] have shown the such an interaction could be predicted using the method of multiple scales. This theory is also referred to as *Hemchandra's theory* in this thesis. The original analysis by Kirthy et al. was completed in Cartesian coordinates but when translated into polar coordinates for an axisymmetric burner results in:

$$q_1(x,r,t) = \omega_a \rho_u \sum_{m=0}^{M} R_m[q_{1m} : q_a]\tilde{q}_{1m}(x,r)[e^{-i\omega_a t} - e^{-i\omega_{gm} t}] \tag{5.50}$$

with:

$$R_m[q_{1m} : q_a] = \frac{\langle \tilde{q}_{1m}^{\dagger}, \tilde{q}_a \rangle}{\Omega_m \langle \tilde{q}_{1m}^{\dagger}, \mathbf{B}\tilde{q}_{1m} \rangle} \tag{5.51}$$

where $\mathbf{B}$ has diagonal elements $[1, \rho, 0]$, q_1 represents the evolution of the strength of helical instabilities in time. $\tilde{q}_{1m}$ is the mth global helical mode and $\tilde{q}_{1m}^{\dagger}$ is its adjoint. $\tilde{q}_a$ is the global acoustic mode. A full derivation of this equation is provided in Section A.5. Equation 5.50 shows that in the case of the global mode being unstable that the evolution of this mode will be dominated by the value of $\mathrm{Im}(\omega_{gm})$, its temporal growth rate. Kirthy et al. [60] suggest that this could lead to the global mode locking on to the acoustic forcing frequency or the flow dynamics being a combination of both the forced and natural flow oscillations.

In the case where the mode is stable, its temporal variation is driven by the oscillatory $e^{-i\omega_a t}$ term. Stability of a mode is often presented as a function of a control parameter (such as the Reynolds number) or the amplitude of the mode. In this sense a mode that is driven significantly hard enough can break the threshold at which the mode goes from being stable to unstable. This behaviour is referred to as *triggering* and this analysis provides a theoretical framework to a situation where a normally stable mode becomes active once forced.

In either case the strength of interaction between the modes is given by the R_m term expressed in Equation 5.51. The term is composed of three factors. The first is the frequency difference between the modes (Ω_m). As the mode frequencies diverge the interaction between them is reduced by the factor $1/\Omega$ showing that forcing will interact more strongly with modes that have their natural frequency close to the forcing frequency.

The inner product $\langle \tilde{q}_{1m}^{\dagger}, \mathbf{B}\tilde{q}_{1m} \rangle$ shows the dependency of the response on increasing *non-normality* [22]. The final part of the numerator in Equation 5.51 is the inner product of the adjoint global mode with the acoustic forcing mode. This term is most active in the case where the global adjoint mode has a region of strong overlap with the acoustic forcing mode. Chomaz [22] talks extensively of the perfect flow controller being one where the sensor is located at the maximum of a global mode and the actuator located at a maximum of the global adjoint. From the point of view of passive control of global instabilities, this means modifying the geometry in regions where the adjoint is maximum. In the frame of looking at how instabilities may be modified by the acoustic forcing, the acoustic forcing mode must be strong in the regions where the adjoint of the global mode are maximum. In the case of our injector the local stability analysis was not as fruitful as hoped and as such the global modes and adjoints were not calculated. However, from the local analysis it has been identified that the instabilities are driven by regions of high shear close to the splitter plate separating passages A and B. As this is therefore likely to be the *wave maker* region [58] and as the adjoint mode must overlap with the global mode here, the adjoint mode must also have a reasonable amplitude in this area. Such an observation leads to the conclusion that interaction between the acoustic and unstable modes requires that the acoustic mode be active in this same region. Furthermore from the point of view of passive control of these instabilities, modifying the splitter plate geometry seems to be a reasonable method of changing the prevalence of this mode.

If a large part of the flow field and heat release is governed by large scale instabilities then the strength of these flow structures will dominate the thermoacoustic response. Currently full scale unsteady reacting simulations are too expensive to be used routinely within industry and hence if the thermoacoustic response is to be considered early in the design process, alternative methods must be developed. The helical modes begin life as linear instabilities in the shear layers between the jets emanating from the fuel injector and can, at least in principle, be found using linear stability analysis that requires as its input the mean field only. The acoustic impedance has been found using fully non-linear URANS simulation however the linearity of the acoustic field opens the possibility of linearised methods for deriving the same result. If the theory of Kirthy et al. [60] is dominant, a designer may in the future be able to run a single steady simulation to derive the mean field, followed by a linearised acoustic solver that approximates the acoustic impedance and gives the magnitude of fluctuation inside the injector and combustor. Such a simulation would be solved in frequency space making it significantly faster than stepping through each frequency with the URANS methodology. The same mean field could then be analysed using a linear stability analysis to find the most dominant unsteady flow structures and global modes.

The Kirthy et al. [60] theory could then be used to check which unstable flow structures would be most sensitive to acoustic forcing by combining the results of the linear stability analysis and the acoustic field simulation. The designer may then opt for a design with a very strong helical mode that competes with the $m = 0$ acoustic mode, a very quiet design with no large scale instabilities, or a design which is relatively quiet which then becomes unstable with respect to many different helical modes when forcing is applied. Such a design would lead to a significant amount of turbulent production that may help to dampen the instability. Our current understanding of the relationship between helical modes and combustion instabilities is not sufficient for us to know what the desirable characteristics of an injector are which make full scale simulations and experiments important to our future understanding before simplified methods can be used effectively within industry.

5.6.4 The dynamic mode decomposition (DMD)

DMD is a method used to decompose a spatially and temporally varying set of data into a series of linearly evolving modes first introduced by Schmid [131].:

$$q(\mathbf{x}, t) = \sum_{j=0}^{N} \hat{q}_j(\mathbf{x}) e^{i\omega_j t} \tag{5.52}$$

where $\hat{q}$ is the DMD eigenvector and ω is the complex value of the mode frequency and growth

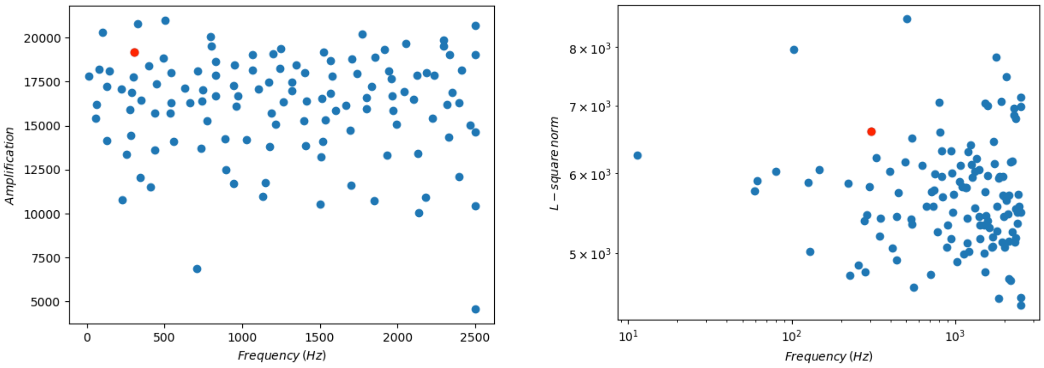

(a) The mode amplification rates.

(b) The mode amplitudes.

Figure 5.20: The DMD modes calculated from the 300 Hz forced calculation. The mode best representative of the acoustic effects is highlighted in red.

rate. This decomposition has a strong resemblance to the forcing function assumed by Kirthy et al. [60] and this makes DMD an excellent candidate for identifying the net effects of the acoustic forcing. The DMD is calculated here following the work of Motheau et al. [95] however $q = [\mathbf{u}, p]$ which doesn't include temperature as the flow is non-reacting. The first step in calculating the DMD using this method is to compute the POD. In order to prevent the POD being dominated by the variables with the largest values, the input data is first normalised by its temporal and spatial mean:

$$q^* = [\mathbf{u}^*, p^*] \tag{5.53}$$

$$\mathbf{u}^* = \frac{\mathbf{u} - \overline{\mathbf{u}}}{\langle |\overline{\mathbf{u}}| \rangle} \tag{5.54}$$

where the $\langle\rangle$ symbols designate a spatial average. The DMD was then calculated of the 300 Hz forced simulation using the method of Chen et al. [21] over 250 snapshots. DMD provides a set of complex frequencies for each mode and the DMD eigenvector. The complex frequencies contain the fluctuation frequency and amplification rate of each mode. These two things along with the norm of the eigenvector can be used to classify the mode. As the forcing was applied at 300 Hz, the DMD mode corresponding to that forcing should have a fluctuation frequency of around that value. Due to the temporal resolution and the number of snapshots used there was no mode at exactly 300 Hz, however a mode fluctuating at 304 Hz was identified as being representative of the effects of acoustic forcing. This mode has a strong axis of symmetry along the x-axis in agreement with the forcing methodology. The norm and amplification rate of each DMD mode is shown in Figure 5.20. If the modes were perfectly converged then the amplification of each of them would be negative as any modes with a positive amplification rate will diverge to infinity as time continues. The significantly positive values found in this DMD is most likely a result of insufficient snapshots although it is possible that some minor transients are causing this error. The result may be improved by using the POD as a filter and computing the DMD over a limited number of the higher energy containing modes.

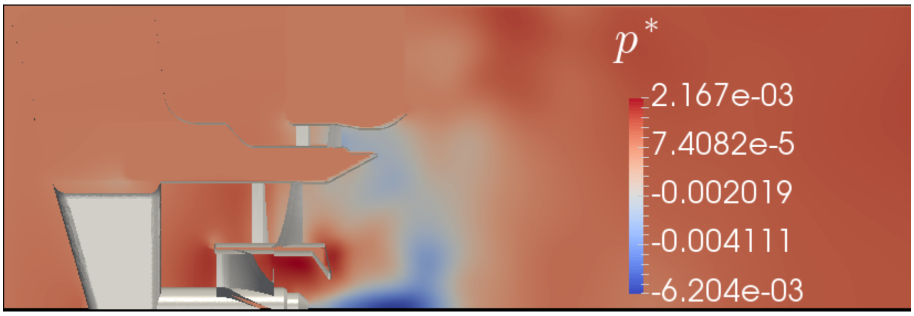

(a) Pressure.

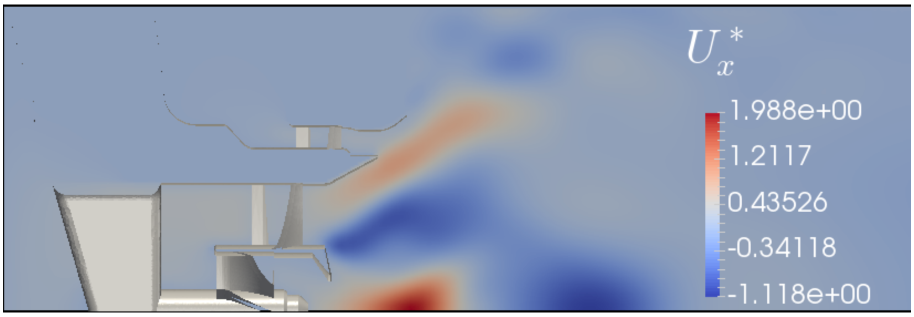

(b) Axial velocity.

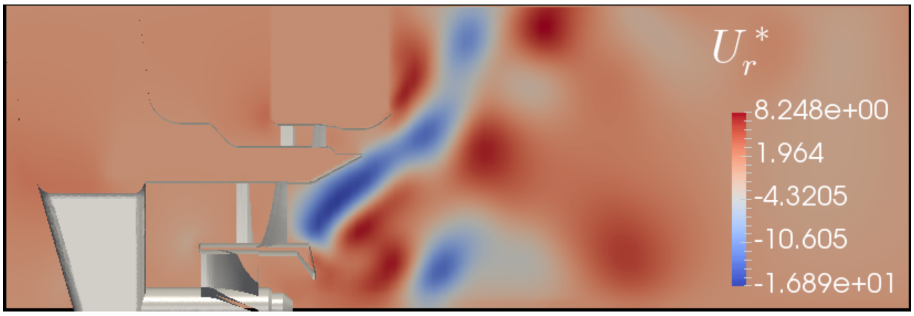

(c) Radial velocity.

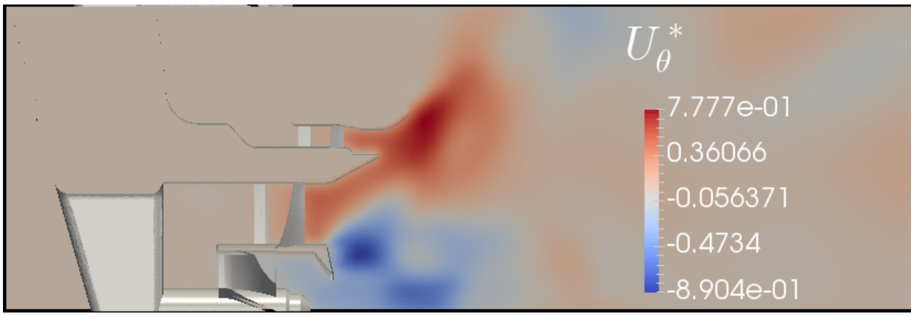

(d) Azimuthal velocity.

Figure 5.21: The DMD mode identified as being best representative of the effects of acoustic forcing averaged in the azimuthal direction.

Because the forcing and chosen mode is close to axisymmetric, the mode has been averaged in the azimuthal direction and can be seen in Figure 5.21. It shows the the acoustic forcing accelerates the velocity through the injector, reduces the cone angle and increases the magnitude of the swirl number coming from each passage. The pressure field is mostly constant over the field due to the long wavelength of the acoustic wave but shows however, two out of phase regions. The first is along the centreline where the vortex core, induced by the swirl is located. The strength of this vortex is fluctuating with the unsteady azimuthal fluctuations of velocity. Because the azimuthal fluctuations lag behind the axial fluctuations of velocity induced by the acoustic wave, at this frequency of excitation the vortex strength is maximum when the acoustic pressure is maximum resulting in the low pressure region in the centre of the domain. The other region of interest is located at the exit of passage B and C. This region is dominated by large scale helical structures which appear to have become emboldened by the acoustic forcing. These structures convect out from the splitter plate towards these passages. The high velocity flow contained within these structures caries with it a region of low pressure that due to the convective delay, can also be out of phase with the acoustic wave. The periodic enhancement of these unstable helical modes is described further in the next section.

5.6.5 Non-linear acoustic feedback

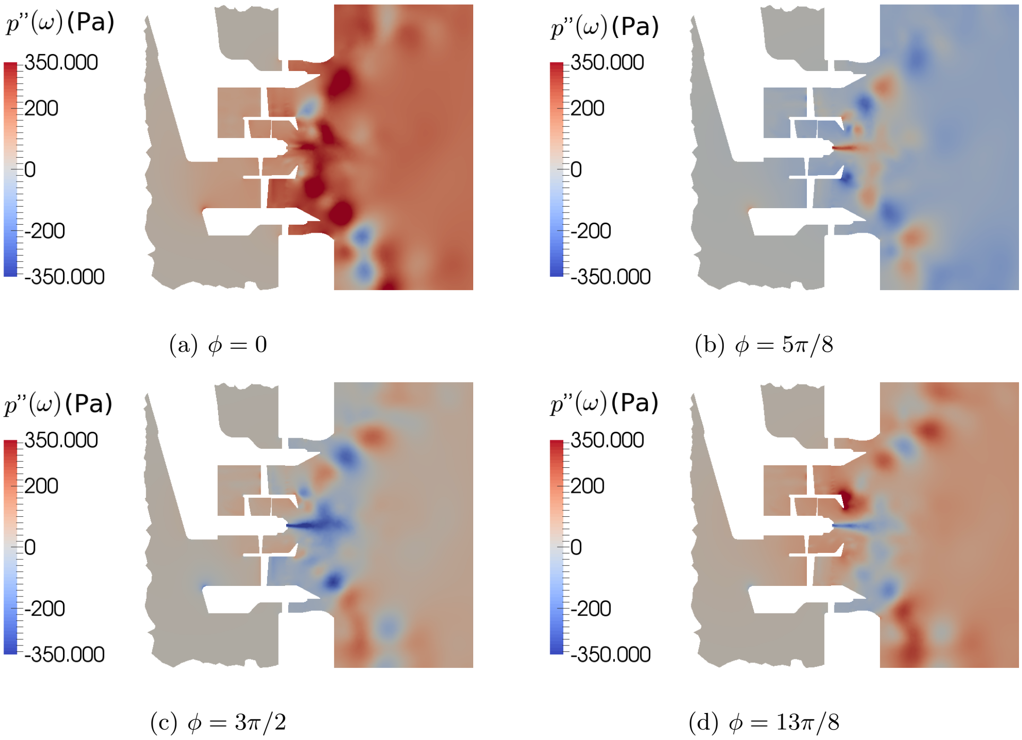

(a) $\phi = 0$

(b) $\phi = 5\pi/8$

(c) $\phi = 3\pi/2$

(d) $\phi = 13\pi/8$

Figure 5.22: Some of the phase averaged pressure fields for the compressible OpenFOAM simulation forced at 300 Hz highlighting regions of hydrodynamic feedback.

Performing a Reynolds decomposition on the Bernoulli equation [8] and assuming that the total pressure field and density are constant results in:

$$p' = -\frac{1}{2}\rho(2\overline{u}u' + u'^2) \tag{5.55}$$

This shows that there is a local modification of the pressure field due to a fluctuation in the velocity field. A localised fluctuation of pressure close to the exit of a flow passage as a result of a vortex convecting from the splitter plate in a radially outwards direction, might then cause an acceleration of fluid through the injector in a similar manner to the incident acoustic forcing.

Because the helical modes result in the shedding of vortices that convect downstream from the splitter plate, they rotate around the central axis at a similar rate to the fluid. It is therefore expected that these structures cause an appreciable azimuthal, radial and temporal variation in the mass flow rate through the injector. Furthermore, if the injector hydrodynamic stability is augmented according to either the *mean flow correction* or *non-linear interaction* then it can be expected that the strength of this natural radial, azimuthal and temporal variation will be increased.

Due to the convective nature of these vortices, the wavelength of these pressure fluctuations are short and the phase velocity small in comparison to the acoustic wave. This means that as the vortices pass each flow passage, the pressure fluctuations they produce can be either in or out of phase with the incoming acoustic wave. If they are in phase then the closest part of the injector will experience a higher than normal pressure drop while an out of phase response will results in a smaller than normal mass flow rate change. Again, if the strength of these structures have been modified by the acoustic wave then a pathway of non-linear acoustic feedback has been established.

Figure 5.22d shows a large low pressure region caused by a vortex passing close to the exit of passage C. This vortex causes the passage exit pressure to be lower than the pressure in the bulk of the duct. There are also several high pressure regions caused by spots of stagnating flow that appear to be raising the pressure of passage B. In Figure 5.22c the same effect appears to be lowering a section of passage B while in Figure 5.22a regions of stagnating flow are seen to raise the pressure in passages B and C.

As the flow is pulled through the injector by the fluctuating acoustic field, the flow is accelerated azimuthally. This azimuthal acceleration has the effect of strengthening the vortex at the centre of passage A with an associated drop in pressure. Similar to the mass flow rate through the injector, this process has an associated time lag that leads to this reduction of pressure due to the vortex acceleration to be out of phase with the incoming pressure field. This effect can be seen in figures 5.22c and 5.22d. Similarly, when the flow through the injector decelerates, the vortex is weakened, leading to a pressure increase which is out of phase with the acoustic wave as seen in Figure 5.22b.

5.6.6 Mean pressure field modification

One of the most surprising discoveries of this chapter is that the acoustic forcing appears to be able to modify the effective area of the individual passages. Table 5.3 shows the change in discharge coefficient between an unforced simulation and the same simulation forced at 300 Hz. Unfortunately the forcing has caused some pressure drift which makes it difficult to compare the pressure fields directly, however one possible explanation is through the modification of the mean pressure field due to the increased hydrodynamic fluctuation. This hypothesis is supported by the

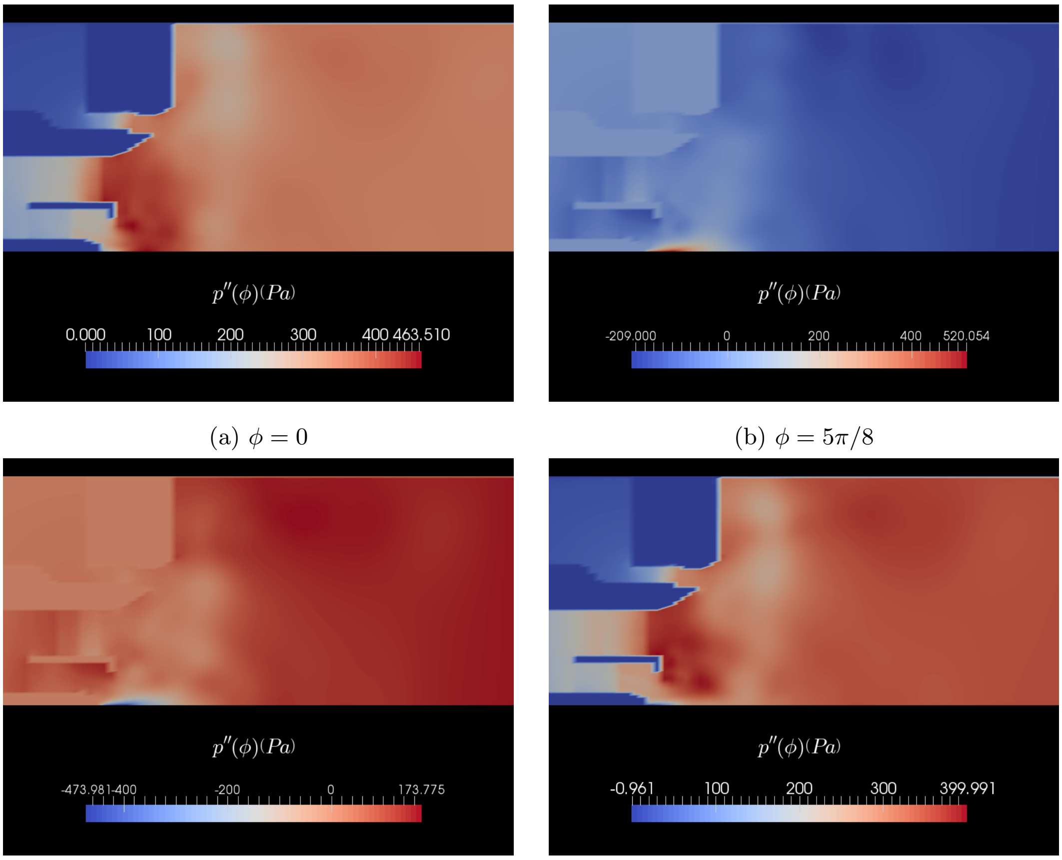

(a) $\phi = 0$

(b) $\phi = 5\pi/8$

(c) $\phi = 13\pi/8$

(d) $\phi = 15\pi/8$

Figure 5.23: The phase and azimuthally averaged pressure fields from the OpenFOAM simulation forced at 300 Hz.

Table 5.3: The percent change in discharge coefficient between the unforced and 300 Hz forced simulation for each flow passage and the total injector.

Passage	A (%)	B (%)	C (%)	total (%)
Cd change (%)	+9.0	-0.4	+18.3	- 0.6

observation that the discharge coefficients of each passage change by differing amounts. The large difference in how each of the flow passages responds may be due to the effect of phase between the acoustic and hydrodynamically induced fluctuations of pressure. Passage B appears to be located close to the splitter plate and the time it takes for the hydrodynamic instability to be modified and for the modification to convect to the passage tend to correspond to a pressure maximum. Passage C appears to located such that this convection time lines up with an acoustic pressure minimum, further accelerating the flow through this passage. Figure 5.23 shows the phase and azimuthally averaged pressure field at several points in the acoustic cycle. Figure 5.23a shows that the highest pressure is located just downstream of passage B and lowest just downstream of passage C. The highest pressure found in the domain when the acoustic wave is at a minimum can be found in the centre of passage A due to the lag of the central vortex as seen in Figure 5.23b and the lowest pressure when the acoustic pressure is high as seen in 5.23c. Finally Figure 5.23d shows that the maximum pressure can be found in passage B and minimums in passages A and C. This translates to a significant increase in the discharge coefficients of passages A and C and a reduction in the discharge coefficient of passage B. Some of these effects are due to vortices that are formed directly from the fluctuating acoustic pressure, however the strong increase in the presence of hydrodynamically unstable helical modes suggests that the interaction between the acoustic response and hydrodynamics is what drives the majority of this process. It is also worth taking note that the phase and azimuthally averaged pressure fields resemble the DMD mode identified as being representative of the effects of acoustic forcing.

5.7 Conclusion

The acoustic response of three different injector geometries have been analysed using a variety of experimental and numerical methods. The multi-microphone technique was used to derive the acoustic impedance experimentally of two fuel injectors with very different geometries and very different thermoacoustic responses. The acoustic impedance shows, in contrast, very little geometric sensitivity.

Using the method of Su et al. [149], the acoustic impedance of one of the injectors was obtained numerically using OpenFOAM and the multi-microphone technique. This result was then compared with the acoustic impedance as calculated by measuring the mass flow rate and pressure fluctuations at the exit plane of the injector directly. This showed the calculation of resistance to be consistent across both measurement techniques but with the incorrect slope of the reactance. This shows that the magnitude and phase of the velocity fluctuations are higher at the exit plane of the injector than those calculated by only considering the acoustic waves.

The pressure, mass flow rate, angular velocity and swirl number were calculated for each passage and the entire injector over the acoustic cycle. The magnitude of the pressure and mass flow rate appears to average across the passages while the fluctuations of angular velocity and swirl number sum. The fluctuations in swirl number are a function of both axial and azimuthal velocity fluctuations. The relative strength of the swirl number fluctuations is dependent on the phase shift between the axial and azimuthal fluctuations which in turn is a function of the frequency of excitation, the swirl vane response time and the distance between the swirl vanes and the exit of the passages. Given the mechanisms described in Candel et al. [16], it seems likely that the FTF may be tuned by varying the location of the swirl vanes inside the passages.

POD analysis has shown that the unstable helical structures identified in the unforced case are also apparent in the case of acoustic forcing. Furthermore, the strengths of each mode appears to be modified by the addition of forcing. Three different possible mechanisms for this phenomenon are discussed including two linear and one non-linear mechanism. The non-linear mechanism opens an interesting possibility for future design methods.

Two methods by which the modified hydrodynamics can feed back to the pressure field have also been identified: The first is from localised changes of velocity that convect past flow passages modifying the local pressure field causing azimuthal variations in the velocity field through the injector passages. The second is that even when these azimuthal variations are averaged in the azimuthal direction, the presence of these structures can alter the mean pressure field locally leading to a different mass flow rate through one passage relative to another.

Characteristic boundary conditions have been developed for the compressible version of PRECISE-UNS and some validation of the code is attempted with simple geometries and the SCARLET rig. The code was still very much in development for the majority of this project and as such, further validation is required. The difficulties involved with correctly resolving acoustic waves using URANS or LES are a primary motivation for the next chapter where incompressible simulations with appropriate upstream boundary conditions are used as a surrogate for compressible simulation.

Chapter 6

The reproduction of acoustic forcing in incompressible simulations

Up until this point, the upstream boundary condition has been of relatively little importance. In the plenum, injector and duct system, the injector generates the vast majority of interesting flow features. It is the interaction of acoustic waves with the fuel injector that generates the phase varying change to mass flow rate, angular velocity and swirl number while the complex shape that the swirler vanes and flow passages present to the flow that cause the production of turbulence. Due to the injector geometry being included in the simulations so far, the inflow boundary conditions specified at the inlet plenum boundary make very little difference to the observed flow inside the injector.

The direct modelling of the flow upstream and within the injector geometry has significant disadvantages. Firstly, the complexity of the injector shape requires that a very large number of cells must be used within the flow passages, making such simulations very costly. The flow upstream of the injector is essentially uninteresting, as there are no significant flow structures present that relate to the heat release of the flame. Unfortunately, when modelling a plenum to plenum simulation some of the plenum must be included to ensure that the upstream potential field of the fuel injector is captured correctly. In the case of a simulation aimed at resolving the acoustic impedance some of the upstream plenum must be resolved to ensure that the acoustic waves that pass through the injector are dissipated or released through the boundaries without reflection. The modelling of the fuel injector is required to provide the necessary acoustic impedance to the flow, resolve the large scale flow structures and to generate the correct amount of turbulence. If the acoustic impedance has been resolved sufficiently well using the compressible URANS method it would therefore be desirable to stop modelling the plenum and air swirler to significantly reduce the complexity and size of the computational grid. This requires however that a suitable way of representing the acoustic impedance and generation of turbulence is included within the upstream boundary condition.

The fluctuating velocity at the upstream boundary will be decomposed into two components,

$\mathbf{u}''(\mathbf{x}_{in}, \phi)$ represents the fluctuation associated with the acoustic forcing while $\mathbf{u}'''(\mathbf{x}_{in}, t)$ represents the random fluctuations associated with turbulence. These quantities can be related to the instantaneous velocity at the inlet through:

$$\mathbf{u}(\mathbf{x}_{in}, t) = \overline{\mathbf{u}}(\mathbf{x}_{in}) + \mathbf{u}''(\mathbf{x}_{in}, \phi) + \mathbf{u}'''(\mathbf{x}_{in}, t) \tag{6.1}$$

To avoid having to model the plenum or the air swirler, the upstream boundary can be switched from being a total pressure inlet to an inlet of fixed velocity. In the case of a non-acoustically forced simulation this upstream boundary condition must include suitable approximations for the velocity as well as any additional quantities required by the chosen turbulence modelling approach (k and ϵ for the $k - \epsilon$ model, k and ω for the $k - \omega$ model etc.). In the case of a URANS simulation most of the $\mathbf{u}'''(\mathbf{x}_{in}, t)$ term is approximated by the turbulence model, however there may be non-phase locked variations in inlet velocities due to the presence of natural hydrodynamic instabilities operating at characteristic frequencies different to the forcing frequency. The upstream TKE or other fluctuations of interest might also be a function of the phase angle of the incident acoustic wave and contain random, non-phase related fluctuations. Therefore equation 6.1 must be accompanied by (in the case of the $k - \epsilon$ model):

$$\mathbf{u}(\mathbf{x}_{in}, t) = \overline{\mathbf{u}}(\mathbf{x}_{in}) + \mathbf{u}''(\mathbf{x}_{in}, \phi) \tag{6.2}$$

$$k(\mathbf{x}_{in}, t) = \overline{k}(\mathbf{x}_{in}) + k''(\mathbf{x}_{in}, \phi) \tag{6.3}$$

$$\epsilon(\mathbf{x}_{in}, t) = \overline{\epsilon}(\mathbf{x}_{in}) + \epsilon''(\mathbf{x}_{in}, \phi) \tag{6.4}$$

The phase sensitive part of the inlet velocity (along with any additional turbulent quantities) needs to be represented by the inlet boundary condition in both the LES and URANS simulations. In the case of reacting simulations, it is likely that the URANS methodology will not provide sufficient accuracy and as such an LES methodology will be preferred. In the case of the LES simulation, the upstream boundary conditions should reproduce a significant portion of the turbulent flow emanating from upstream. The turbulent flow, is by its nature, random, however, turbulence can be classified using a number of statistical measures.

In addition to the correct statistical representations of turbulence as introduced in chapter 3, a good unsteady upstream boundary condition should reproduce a significant proportion of the coherent structures within the flow. Coherent structures are broadly classified as flow features of interest that have been extracted from the flow using some kind of filtering technique. The most commonly observed coherent structures are vortices, but other structures have been identified such as streaks and recirculation bubbles [5]. Coherent structures contain a significant amount of the kinetic flow energy and tend to dominate the evolution of the flow. One such example is the interaction between a vortex sheet from the preceding blade row interacting with the downstream blade boundary layer in a low pressure turbine [129]. It is for this reason a good inlet boundary will contain these significant flow structures to ensure that downstream interactions between these structures and other flow phenomenon in the domain are accurately captured.

This chapter focuses on methods to reproduce the acoustic response of the injector in an incompressible simulation. Compressible and incompressible URANS simulations using various techniques are compared for accuracy to the forced compressible simulations completed in the previous chapter. The methods tested are the *Strong recycling method* (Section 6.2.1), *mass flow forcing method* (Section 6.2.3) and the newly developed *PODFS method* (6.2.4). Methods for producing additional turbulent fluctuations for LES are also described in this chapter.

6.1 Computational Domain

In order to reduce the cost of simulations, it is desirable to reduce the size of the computational domain and switch to an incompressible methodology.

The first thing that can be changed from the mesh used for the computation of the acoustic impedance is that the outlet duct can be shortened. The length of the duct is important for two reasons, in the case of the acoustic impedance computations it is important that the acoustic waves arrive at the injector in an almost planar manner; in the case of incompressible simulations, this requirement is no longer relevant. However, the constant static pressure boundary condition at the exit forces the pressure across the duct to be homogeneous, and this is generally not the case. This distortion of the pressure field extends upstream and can cause the velocity field to be incorrectly calculated in the region of interest. In simulations using incompressible methods the downstream duct was shortened to $6.5D$ and because the acoustic waves do not need to be resolved at the exit, the duct cells can be geometrically expanded in size such that the cell closest to the exit has a resolution of $0.175D$. This is more than sufficient to capture the mean pressure field gradient.

In the case of an incompressible simulation aimed at simulating the effects of acoustic forcing on the downstream flow, the inlets must be located at the exit plane of the injector passages. This is because the acoustic impedance must be accurately represented at this plane. Using planes further upstream will result in a phase shift between the imposed impedance and the true impedance due to the speed of sound being finite in the compressible simulation against the infinite sound velocity in the incompressible simulation. A phase shift could be applied to correct this effect, but then the magnitude of the fluctuating flow at the outlet would need to be checked and modulated as the velocity fluctuation magnitude is not the same at the beginning and end of the flow passages [149].

Two different computational domains were tested. The first is for mass flow forcing simulations (Section 6.2.3). It is similar to the domain generated for the acoustic simulations however the large hemispherical plenum is replaced with a smaller cylindrical one (diameter $3.9D$) with the inlet defined over the circular end and the outlet duct is shortened. The second domain is for the recycled simulations (Sections 6.2.1 and 6.2.4) which impose the velocity at the exit of each passage. Both meshes were generated using the RRBP method (Section 2.5.1), with the same cell density settings as the LUBPAI mesh (Section 2.5.2). The three passages are given their own inlets. The first is aligned with the tip of the splitter plate and extends radially from the centre of the injector to the edge of passage A. The second is aligned with the exit off passage B. The third

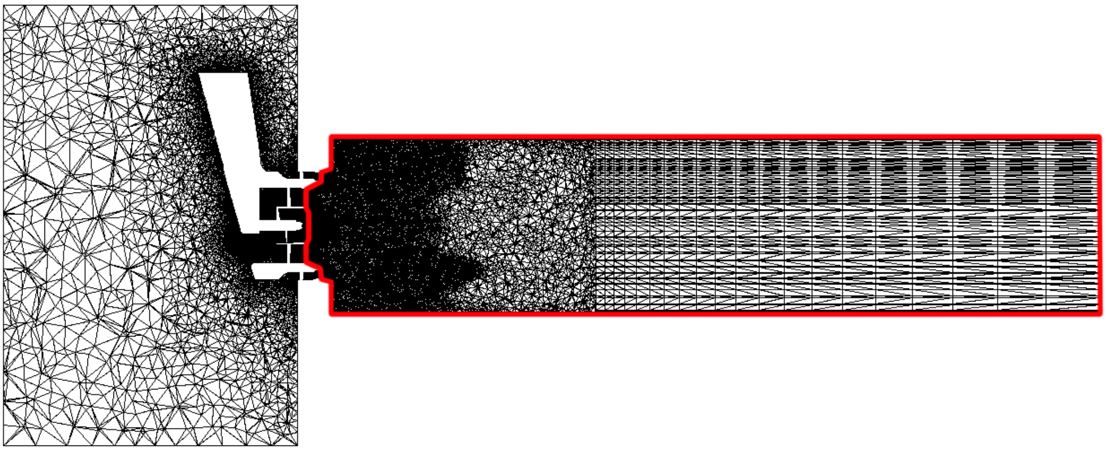

Figure 6.1: The mesh used for the mass flow forcing and the mesh used for the recycling methods shown with the red line for comparison.

inlet is located such that it is aligned with tip of the pre-filmer and covers the radii corresponding to passage C. As these three inlets are at separate x-locations they are modelled and applied separately. The domains are shown in Figure 6.1. The mass forced simulation domain resulted in a cell count of 9.0 million cells while the cut down mesh was only 2.6 million cells.

6.2 Modelling Approaches

6.2.1 The strong recycling method

The most intuitive method for the imposition of accurate time varying inlet conditions is to paste the velocity onto the inlet taken directly from another simulation. This involves the extraction of the relevant planes of velocity data from the unsteady, acoustically forced simulations before re-pasting those planes of data onto the inlet of the incompressible simulation [165]. The process is complicated by the fact that the two meshes are not the same and both unstructured. Therefore the information from the compressible simulation must be interpolated onto a plane with a certain resolution and then re-interpolated onto the mesh of the incompressible simulation. The resolution of the plane and interpolation methodology is an important factor however the plane resolution is fixed in this section and the interpolation in all cases is performed using a simple linear interpolation method.

This methodology, whilst simple, has the advantage that no modelling is required and apart from errors introduced during the interpolation, should reproduce the field accurately [165]. There are significant disadvantages with this method however, some that are replicated in some of the other methods presented here. The first problem is that this method does not provide an interface between simulations that use different methods of turbulence closure. For example, the compressible URANS simulations used in the determination of the acoustic impedance cannot be used to directly provide information to an LES simulation about the state of the turbulence at the boundary except through the estimation of some of the components of Reynolds stress tensor. This problem also exists even between URANS simulations as different turbulence models rely on the computation of different scalars. These scalars may be related through some relation (as is the case for ϵ and ω) however due to different equations being solved for these quantities a simple relationship between them cannot be assumed.

The next disadvantage involves the limitation of applying known quantities to the inlet of the simulation. In order to do this, the values of each quantity must be known at each timestep. This causes problems when the first simulation is completed using a different timestep to the first. Temporal interpolation can be used to ease this restriction however another related problem occurs when it is desirable to continue a simulation longer than the simulation used to generate the inflow. The data can of course be re-applied however this introduces a non-physical periodicity in time and a discontinuity in the velocity field. This problem can be minimised by using a blending zone where the end and start of the inlet data are overlapped and blended together [169].

The final problem with this method is that it lacks generality. Ideally an inlet boundary condition supplies all of the necessary information to generate the correct downstream result without explicitly introducing any invalid assumptions about the upstream flow including any errors or bias. The strong recirculation method does not provide any randomness to the simulation and as such is as accurate or inaccurate as the donor simulation.

The strong recirculation method has been used extensively for LES computations where the inlet turbulence is provided through the use of an auxiliary simulation of homogeneous isotropic turbulence (HIT) [70], coaxial pipe flow [111] and fully developed channel flow [166]. It has been stressed that the initial and final simulation should be computed with similar Reynolds numbers and spatial resolution [25]. All of these examples rely on auxiliary simulations that are periodic in the stream wise direction that introduces additional streamwise periodicity in the inlet data. In the work presented here the auxiliary simulation contains no such periodicity and as such contains none of the associated downsides of the method.

6.2.2 Weak recycling methods

Weak recycling methods [165] are a family of methods that include the recycling and rescaling method [167]. There is no donor simulation but instead an upstream section of the simulation becomes a turbulence adjustment zone where the effects of the imposed inlet velocity profile can be filtered out of the flow field. At the end of this adjustment zone, or further downstream, depending on the simulation, the velocity field is extracted at each timestep. The extracted velocity field is then reapplied at the inlet. Because any boundary layers present will have grown over this distance, the velocity field must be scaled to prevent the infinite growth of the boundary layer thickness. The main disadvantage with this method is that it induces non-physical periodicity into the flow but has been used successfully for LES and DNS simulations in a number of studies [165, 167]. In the case of representing the acoustically forced flow, there is an additional complication due to the different propagation speeds of azimuthal and axial fluctuations. Axial fluctuations travel downstream at the speed of sound, which being infinite in incompressible flow means that the necessary upstream extension of the domain would not cause problems with capturing the correct phase shift between pressure and velocity fluctuations at the exit of the injector. Unfortunately the azimuthal fluctuations convect at the local flow velocity and hence the phase shift between these and the axial fluctuations would be disrupted by the upstream extension of the domain. This could be corrected by calculating the shift and adding a suitable delay to the axial fluctuations however this complication is the primary reason this method was not tested. Furthermore the extension of the domain upstream, although less complicated than the flow swirler, somewhat negates the advantages of truncating the domain in the first place.

6.2.3 Mass flow forcing

Unlike the above techniques that can be applied generally to any inlet flow type, mass flow forcing is achieved simply by scaling a base flow velocity by an appropriate factor such that the mass flow through the domain increases or decreases appropriately. This method is often used to approximate the effects of acoustic forcing on a flame to derive the FTF or FDF [38].

This approximation is often made as a mass flow forced simulation may be undertaken assuming incompressible flow therefore making the numerical methodology significantly simpler. The disadvantage of this method is that it does not accurately capture the phase shifts between axial velocity fluctuations induced by incoming acoustic waves and azimuthal velocity fluctuations that convect downstream at the flow velocity, as demonstrated in this chapter.

The implementation made in this study modifies the already hard coded mass flow rate inlet condition within PRECISE-UNS. This condition, when applied, forces the inlet flow to be oriented in a direction defined by the user, however the inlet velocity is scaled such that the user defined mass flow rate is maintained across the inlet. In the mass flow rate forcing method and a mass flow fluctuation amplitude (A_m), scaled by the mean flow rate and fluctuation frequency (f) are also chosen. At each timestep the inlet mass flow is set by:

$$\dot{m} = \overline{\dot{m}}(1 + A_m sin(2\pi f t)) \tag{6.5}$$

6.2.4 POD Fourier series (PODFS) model

The PODFS model has been developed here to alleviate several weaknesses of the strong recycling method. Firstly, in order to significantly reduce the amount of data that must be stored, the POD is applied to the data collected from the donor simulation and used as a filter. This filtered data then consists of a set of spatial modes that are independent of time and a set of temporal modes that are independent of space. The spatial modes can be used as a convenient basis for a model that does not suffer from issues of non-physical periodicity, discontinuity or other time related problems however the temporal modes must be modelled somehow.

In the PODFS model, introduced here, the temporal modes are each modelled using a series of analytical functions of time, this removes any timestep dependence from the inlet condition. Additionally, if the analytical functions are chosen to be periodic, then no discontinuities will exist at the time corresponding to the end of the donor simulation data and the solution will remain bounded at all points in time during the simulation. As such, each temporal mode is modelled using a Fourier series. The disadvantage of this model is the imposed periodicity as compared to the LODS approach however the PODFS is considerably more robust.

The POD must be computed such that it includes all of the variables required at the inlet boundary. In the case of the $k-\omega$ turbulence model, that includes the extracted values of k and ω at the inlet. This presents a problem as the units and magnitudes of velocity and turbulent kinetic energy are vastly different. This can cause the POD to over-represent the variable with the highest magnitudes. As such, normalisation of the inlet variables must be achieved. The inlet variables will be denoted U_{in} where:

$$U_{in}(\mathbf{x}_{in}, t) = \{u(\mathbf{x}_{in}, t), v(\mathbf{x}_{in}, t), w(\mathbf{x}_{in}, t), k(\mathbf{x}_{in}, t), \omega(\mathbf{x}_{in}, t)\} \tag{6.6}$$

the normalisation leads to U_{in}^* defined as:

$$U_{in}^*(\mathbf{x}_{in}, t) = \frac{U_{in} - \overline{U}_{in}}{\langle U_{in} \rangle} \tag{6.7}$$

where the $\langle\rangle$ brackets denote a spatial average and the over line represents the temporal average. From here the snapshot POD method is employed by first generating the covariance matrix $\mathbf{C}$:

$$\mathbf{C}_{ij} = U_{in}^*(\mathbf{x}_{in}, t_i) U_{in}^{*T}(\mathbf{x}_{in}, t_i) \tag{6.8}$$

This equation is not generally correct but can be used in the case where each cell of the input mesh is of equal area. The eigenvalues of the covariance matrix are then computed such that:

$$\mathbf{CA} = \lambda \mathbf{A} \tag{6.9}$$

As the covariance matrix is real and symmetric then the eigenvectors ($\mathbf{A} = \{a_1, a_2 ... a_n\}$) are orthogonal and represent the temporal POD modes while the eigenvalues (λ) describe the relative magnitude of each POD mode. This allows the modes to be sorted as to how significantly each mode contributes to the original data. It can also be used to decide how many POD modes should be used to represent the data. The array $\mathbf{C}$ has dimensions $N_s \times N_s$ where N_s is the number of temporal snapshots. As such, there will be N_s POD modes. If the POD modes are ordered according to the size of their eigenvalues, the eigenvalue energy ($E_\lambda(N_P)$) is defined as:

$$E_\lambda(N_P) = \frac{\sum_{i=1}^{N_P} \lambda_i}{\sum_{i=1}^{N_S} \lambda_i} \tag{6.10}$$

with $N_P \leq N_S$ and is chosen such that E_λ is sufficiently close to 1. Once the number of modes has been chosen, the spatial POD modes are obtained by projecting the temporal modes onto the original data set:

$$\phi_i(\mathbf{x}_{in}) = \frac{1}{N_S \lambda_i} \sum_{k=1}^{N_S} a_i(t_k) U_{in}^*(\mathbf{x}_{in}, t_k) \tag{6.11}$$

The POD approximation of the normalised inlet variables is then given by:

$$\tilde{\mathbf{U}}^*_{in}(\mathbf{x}_{in}, t, E_\lambda) = \sum_{i=0}^{N_P} a_i(t)\phi_i(\mathbf{x}_{in}) \tag{6.12}$$

As already stated, the spatial modes are invariant in time and ready for use in the model however the temporal modes must be interpolated using an analytical function. In order to make the signal periodic and continuous, a Fourier series is chosen to reproduce each temporal POD mode. An ordinary Fourier representation would result in the following approximation of each temporal mode:

$$\tilde{a}_i(t) = \sum_{j=0}^{N_F-1} b_{ij} e^{2\pi\imath jt/\Delta t N_S} \tag{6.13}$$

with the coefficients defined by:

$$b_{ij} = \frac{1}{N_S} \sum_{k=0}^{N_S-1} a_i(t_k) e^{-2\pi\imath jt_k/\Delta t N_S} \tag{6.14}$$

N_F is the number of Fourier terms to be used in the approximation and as N_F approaches N_S the error between the Fourier representation and the original temporal POD modes will decrease, however not all Fourier terms contribute equally to the error. Equation 6.13 is therefore modified to enable an arbitrary selection of whatever Fourier terms reduce the error the fastest:

$$\tilde{a}_i(t) = \sum_{j=0}^{N_{F_i}-1} b_{ik(j)} e^{2\pi\imath k(j)t/\Delta t N_S} \tag{6.15}$$

In this equation the index j is replaced by the function $k(j)$ that allows for an arbitrary ranking of each term in the Fourier series for each temporal mode. The method of ranking each term in the Fourier series is borrowed by the snapshot POD method. If the Fourier coefficient is a representation of how much each Fourier series term contributed to the error between the original signal and the approximation, then the Fourier coefficients are sorted by their magnitude and a Fourier series energy (E_b) can be defined as:

$$E_b(N_F, i) = \frac{\sum_{i=1}^{N_F} |b_i|}{\sum_{i=1}^{N_S} |b_i|} \tag{6.16}$$

In order to reduce the chance of spurious oscillations the Fourier series energy was chosen to be lower than the eigenvalue energy. The Fourier series representation imposes periodicity and therefore continuity of the POD temporal mode. Spurious oscillations are driven by high frequency components of the Fourier series that appear due to the discontinuity of the POD temporal mode between the last and the first snapshots. Finally a scaling coefficient α is added that can be used to correct any changes in the mean flow due to the interpolation error. The approximation for the inlet variables ($\tilde{U}_{in}(\mathbf{x}_{in}, t, E_\lambda, E_b)$) is therefore given by:

$$\tilde{U}_{in}(\mathbf{x},t,E_\lambda,E_b) = \alpha\overline{U}_{in} + < U_{in} > \sum_{i=0}^{N_P} \phi_i(\mathbf{x}_{in}) \sum_{j=0}^{N_{F_i}-1} b_{ik(j)} e^{2\pi\imath k(j)t/\Delta t N_S} \tag{6.17}$$

The reconstructed inlet variable vector is written as a function of the eigenvalue and Fourier coefficient energies as a reminder that both of these quantities will affect the accuracy of this method. The two parameters that also affect the accuracy that are not considered in this equation are the resolution of the 2D inlet planar mesh and the order of interpolation between the original 3D mesh and the 2D inlet planar mesh that the POD is computed over.

6.2.5 A low order empirical model or dynamical system approach

The PODFS method is in some ways a simplification of the low order empirical model or dynamical system approach (LODS) developed by Perret et al. [109]. The method begins with the POD of a data set, either experimental or numerical, collected at the inlet plane. The evolution of the temporal modes in time is then assumed to be governed by a low order dynamical system.

The Galerkin projection of the Navier-Stokes equations for incompressible flow (as computed in the finite element method (FEM)) results in the following system of equations ([101]):

$$\dot{a}_i(t) = \sum_{j=1}^{N}\sum_{k=1}^{N} q_{ijk}a_j(t)a_k(t) + \sum_{j=1}^{N} l_{ij}a_j(t) + d_i \tag{6.18}$$

The coefficients q_{ijk}, l_{ij} and d_i can be calculated directly or approximated from empirical data using any number of methods. In FEM the number of equations is equal to the number of grid points, however as POD is capable of reproducing the dynamics of the flow using significantly fewer modes than grid points, only a small number of these equations need to be solved resulting in a fast low order dynamical model of the flow. There are several significant advantages to this method. Because it is based on a dynamical system, the model can be integrated infinitely in time and at any timestep provided the model remains stable. As the model is based on the Navier-Stokes equations and calibrated to reproduce the donor flow data, it should reproduce the dynamics accurately. One disadvantage is that it cannot take the incoming fluctuations of pressure into account however integration of the governing equations by parts allows for separation of the boundary conditions resulting in:

$$\dot{a}_i(t) = \sum_{j=1}^{N}\sum_{k=1}^{N} q_{ijk}a_j(t)a_k(t) + \sum_{j=1}^{N} l_{ij}a_j(t) + c_i p(\phi) + d_i \tag{6.19}$$

where $p(\phi)$ is the pressure at the injector face. This model could then be trained to respond correctly to the incoming acoustic wave. The original model was applied to (low speed) PIV data which resulted in a deficiency of high frequency, small scale structures in the flow. In order to remedy this, a SDFM (see section 6.3.2) was added to add additional turbulence to the flow. This situation is analogous to when an LES simulations is run with an inlet model built from URANS data implying that such a model may still require some additional synthetic turbulence to be added. The downside of such an approach is that it can be difficult to guarantee model stability [125].

6.3 Synthetic turbulence generators

In this chapter only URANS simulations are considered which limits the requirement that the inflow boundary condition be able to reproduce the turbulent flow fluctuations. However two methods are described here that could be used in the future to generate inflow turbulence for LES simulations. The inlet turbulence is seen to be an important factor in the correct prediction of the pilot jet penetration in chapters 7 and 8. In chapter 8 an additional PODFS model is built from the fluctuating velocities at the injector outlets as captured using the large incompressible LES simulation from chapter 3. This additional PODFS model is seen to provide a sufficient reproduction of the upstream turbulence resulting in the correct prediction of the downstream pilot jet. Synthetic turbulence generators have an advantage however as they are less computationally expensive as compared to auxiliary LES or DNS simulations.

6.3.1 Random noise generation

Random noise generation is the first and simplest of a series of *synthetic inflow turbulence generation* methods that aim to generate a velocity field reminiscent of fully developed turbulence without relying on recycling data from a previous or auxiliary computation. The method works by generating at each timestep and for each boundary cell, a random number that is then scaled by a given Reynolds stress tensor and added to mean velocity. Such a method produces turbulence that is statistically correct in terms of the Reynolds stresses but not correct in terms of the relative temporal and spatial scales. This results in a turbulent field without any large scale coherence which breaks down much faster than homogeneous isotropic turbulence should.

6.3.2 The synthetic digital filter method (SDFM)

The random noise generation method does not impose any spatial or temporal correlation on the imposed velocity fluctuations. As a consequence of this large scale coherent structures are not present at all or break down very quickly in the domain. The SDFM [61] imposes temporal and spatial coherence on the fluctuations to remedy this. An array of random data points r_m is convoluted with a digital filter with support extend N and filter coefficients b_i resulting in:

$$u'''_m = \sum_{i=-N}^{N} b_i r_{m+i} \tag{6.20}$$

Xie and Castro [168] assumed an exponential spatial correlation of $R_{uu}(j\Delta x) = exp[-\pi|j|/(2M)]$ with a length scale of $M\Delta x$ to find:

$$b_j = \tilde{b}_j / \sqrt{\sum_{i=-N}^{N} \tilde{b}_i^2} \text{ with } \tilde{b}_j \approx exp(-\frac{\pi|j|}{M}) \tag{6.21}$$

The temporal correlation is maintained by assuming an exponential decaying form:

$$u_m(t+\Delta t) = u_m(t)e^{-\frac{\pi\Delta t}{2T}} + u_{m\beta}(t)\left(1 - e^{-\frac{\pi\Delta t}{T}}\right)^{0.5} \tag{6.22}$$

$u_{m\beta}$ is calculated using exactly the same method as u_m but using a different input of random values. This method cannot be used to represent the phase averaged fluctuating velocity however it was hoped that it could be used to provide the fluctuating turbulent quantities for an LES simulation in the case where an inlet boundary condition has been built from a URANS simulation.

6.4 URANS of forced injector at $St = 0.3$

Three methods were tested for their ability to reproduce the effects of acoustic forcing on a cut down domain using the PRECISE-UNS incompressible solver. The donor simulation was the OpenFOAM simulation forced at $St = 0.3$. Because forcing by mass flow is unlikely to capture the fluctuations in swirl number necessary for a correct prediction of heat release, each of the methods are tested for their ability to reproduce the swirl number at the exit plane of the injector.

6.4.1 Strong Recycling Method

The strong recycling method was applied by extracting the velocity and turbulence variables at the planes of interest from the acoustically forced OpenFOAM calculation using a 200×200 structured grid resulting in a resolution of $0.008D$. The extraction was performed using a linear interpolation technique and the values were applied at each timestep to the inlets of the domain using the same methodology. The acoustically forced simulations are run with a timestep of $1 \times 10^{-3} t^*$ and the incompressible simulations run with a timestep of $2 \times 10^{-2} t^*$, therefore the extraction was applied at every 20th timestep. The simulations were run using a linear upwind convection scheme and a second order temporal scheme. The $k - \omega$ SST model was used for turbulence closure in all of the simulations.

In order to test the accuracy of the method, the mass flow rate and swirl number were measured across each inlet to test the quality of the reproduction of the acoustically forced values. The same variables were also measured at the injector exit plane to test how close the flow field is to the original data. The results per passage are shown in Figure 6.2. The mass flow rate and swirl number are accurately captured by the method for passages A and B, however they appear to be a under-predicted for passage C. These results could be improved with better interpolation techniques and resolution however some of the error is likely due to changes in density between the incompressible and compressible formulation. One possible option is to simply scale applied velocities to artificially increase the mass flow rate as is done in the case of the PODFS simulations.

The total injector mass flow rate appears to be closely aligned with the original simulations however the swirl number appears to be slightly lower than desired (see Figure 6.3). Some of this variation may be due to the increased timestep or the meshing strategy. Again, artificially increasing the velocity or azimuthal component at the exit of passage C may solve this problem. Despite the slightly lower swirl numbers achieved, the recycled simulation reproduces most of the small time scale features despite the much larger timestep and incompressible method. This shows that most of the fluctuation must be due to convective rather than acoustic phenomena.

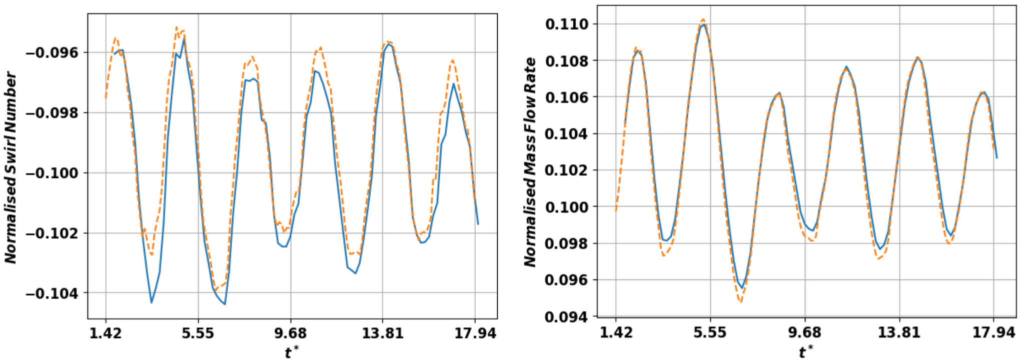

(a) Swirl number through passage A.

(b) Mass flow through passage A.

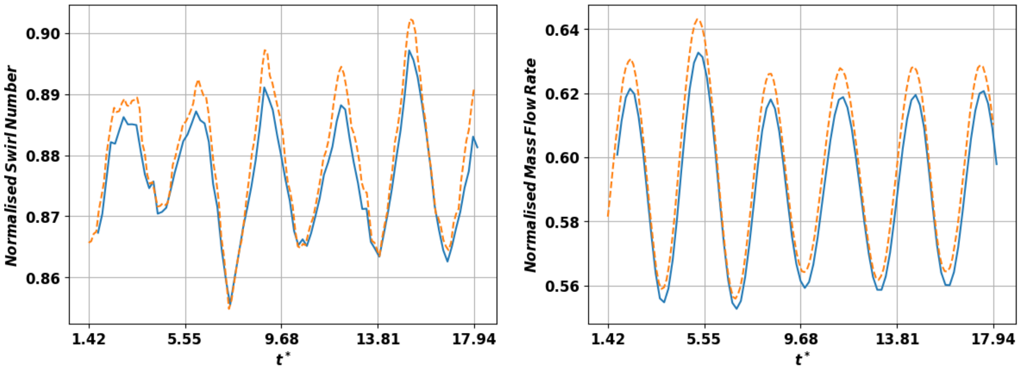

(c) Swirl number through passage B.

(d) Mass flow through passage B.

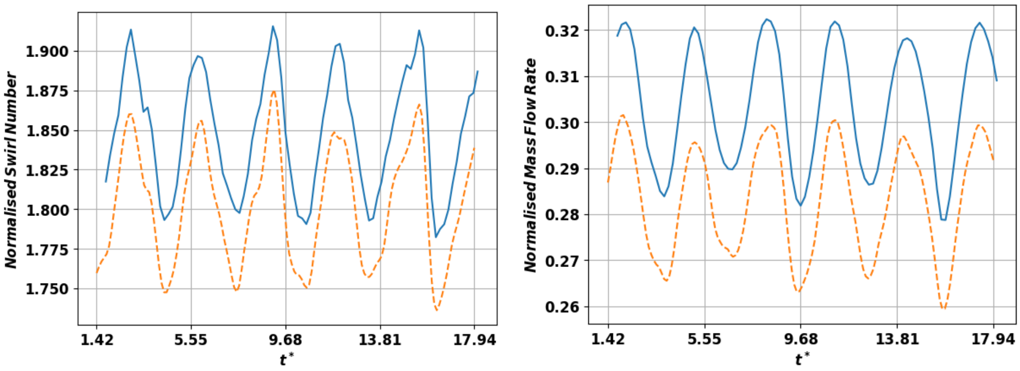

(e) Swirl number through passage C.

(f) Mass flow through passage C.

Figure 6.2: The mass flow rate and swirl number through the injector passages for the acoustically forced simulation (blue) and the recycled simulation (orange).

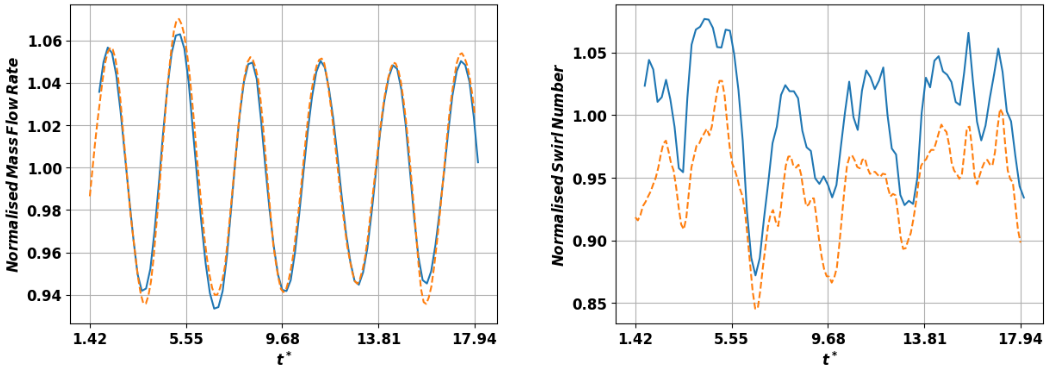

(a) The mass flow rate.

(b) The swirl number.

Figure 6.3: The mass flow rate and swirl number at the injector exit plane for the acoustically forced simulation (blue), the recycled simulation (orange).

6.4.2 Mass flow forcing

The mass flow forcing simulations were run in exactly the same manner as the recycled simulations except that the upstream boundary condition was set by defining a mass flow rate that fluctuated in a sinusoidal manner. The mean and fluctuating mass flow rate was set to be equal to that observed in the acoustically forced case. The results are shown in figures 6.4 and 6.5 along with the PODFS results.

Despite the total mass flow rate being accurately matched for the entire injector, the mass flow rate for each passage is significantly different in the case of the mass forced simulation. This reflects the findings of the previous chapter that the acoustic forcing can alter the relative mass flow rate of each passage due to the *mean pressure field modification* mechanism. The effects of the *non-linear acoustic feedback* mechanism are also visible in the acoustically forced case as the mass flow rate through the injector shows significant cycle-to-cycle variation. The incompressible simulation is not capable of reproducing these interactions between acoustics and hydrodynamics. Looking at the swirl number fluctuations it is obvious that the mass forcing technique does not capture the correct phase or magnitude response. The phase is incorrect due to the azimuthal velocity fluctuations being generated at the same time as the axial fluctuations at the flow swirlers with the azimuthal fluctuations convecting at the flow speed towards the passage exit and the axial fluctuations arriving at the passage exit instantaneously. In the compressible simulation with downstream forcing, azimuthal fluctuations lag behind the axial fluctuations at a time of the distance between swirler and passage exit divided by the speed of sound plus the same distance divided by the mean flow velocity:

$$\Delta t_{x>\theta} = \frac{d_s}{c} + \frac{d_s}{u_x} \tag{6.23}$$

The incorrect capturing of phase also results in an incorrect amplitude of swirl number fluctuations as the amplitude is a function of the phase lag. At the exit of the injector, this poor representation of swirl number at the inlets results in swirl number fluctuations of a lower magnitude and incorrect phase that do not contain the same high frequency fluctuations as in the acoustically forced case.

6.4.3 PODFS

Table 6.1: Mean flow scaling factor, fourier energy, number of POD modes, corresponding POD energy and total energy captured by the PODFS model normalised by total energy.

Inlet	α	E_b	N_P	E_λ	$E_b E_\lambda$
A	1.027	0.9	30	0.985	0.8865
B	1.007	0.8	30	0.963	0.7704
C	1.122	0.9	30	0.995	0.8955

The PODFS inlets were built by extracting the velocity fields at the planes of interest from the acoustically forced simulation in exactly the same manner as in the strong recycling method case. The POD of the resulting set of data was then computed using the snapshot POD method one plane at a time. Once completed, the number of modes to be used and the target energy of the Fourier representation are chosen and the Fourier series of the relevant temporal modes calculated. In the results shown here the number of modes was chosen to be 30 per inlet with the Fourier target energy and total resolved energy per passage shown in Table 6.1. The PODFS mean flow has then been scaled such that the mean flow produced by the model corresponds with the source data.

Because the model has been calibrated, the mass flow rate through each passage (Figure 6.4) and through the injector exit (Figure 6.5) are closely aligned with the acoustically forced simulation. The swirl number is also accurately reproduced for passages B and C with some small deviation for passage A; again, further scaling of the azimuthal velocity could remedy this deficit. The swirl number at the exit plane also shows good approximation of the original data, with the correct phase and with most of the high frequency fluctuations included. It is not as accurate as the recycled case but as it is only an approximation of this input data, it can only ever approach the quality of this result. Increasing the number of modes used or increasing the number of Fourier coefficients will increase the accuracy, however using a very high number of Fourier coefficients runs the risk of spurious high frequency oscillations in the reconstructed temporal mode. The reconstructed velocity at an instant for passages A and B is shown in Figure 6.6. It shows that the model faithfully reproduces most of the interesting flow features including the vane wakes and the influence of the downstream helical modes.

The mean axial velocity and two principle components of the resolved Reynolds stress are shown in Figures 6.7 and 6.8 for the donor simulation, mass forced simulation and the PODFS reconstruction. They show that the PODFS reconstruction captured the mean velocity profile and two components of velocity fluctuations accurately while the mass forced simulation shows significant differences, especially the fluctuations of azimuthal velocity. The cutdown mesh was made without prism layers at the walls and hence a small discrepancy can be seen in the mean velocity field along the duct walls. The velocity fluctuation magnitudes also seem to be a little stronger in the PODFS case, especially in the axial component, this could be due to the slightly different mesh or numerical method.

The failure of the mass forced simulation to capture either the mean or fluctuating fields

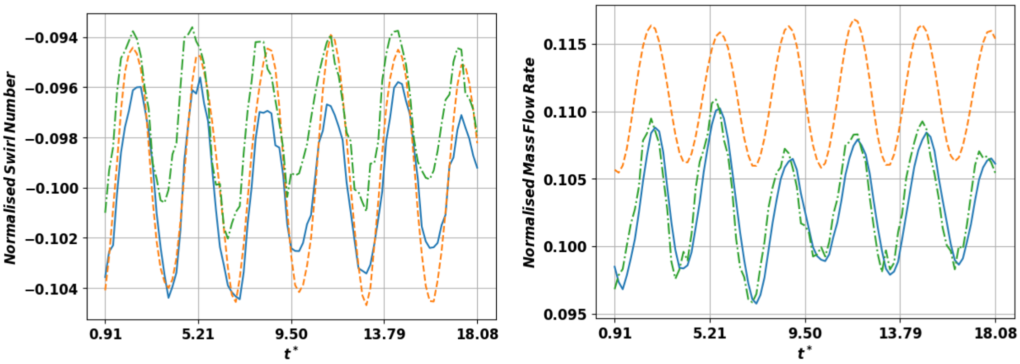

(a) Swirl number through passage A.

(b) Mass flow through passage A.

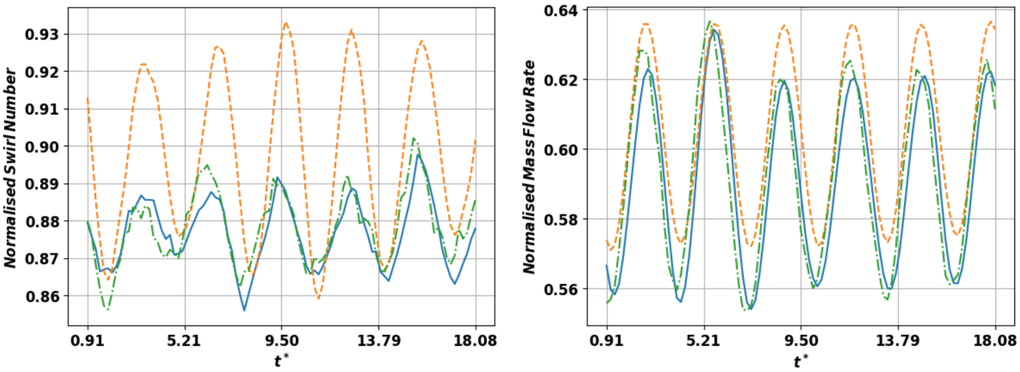

(c) Swirl number through passage B.

(d) Mass flow through passage B.

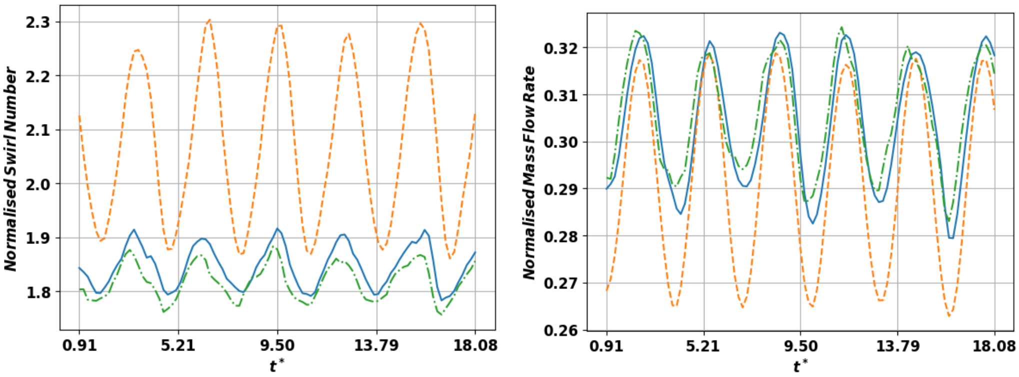

(e) Swirl number through passage C.

(f) Mass flow through passage C.

Figure 6.4: The mass flow rate and swirl number through the injector passages for the acoustically forced simulation (blue), the mass forced simulation (orange) and the PODFS simulation (green).

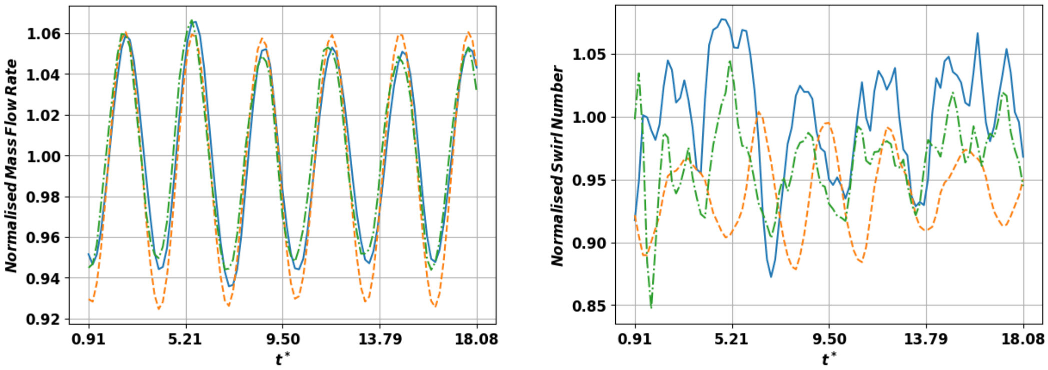

(a) The mass flow rate.

(b) The swirl number.

Figure 6.5: The mass flow rate and swirl number at the injector exit plane for the acoustically forced simulation (blue), the mass forced simulation (orange) and the PODFS simulation (green).

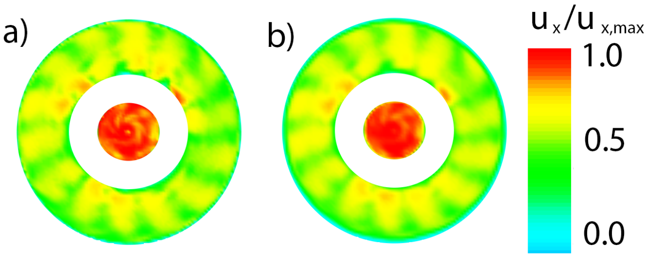

Figure 6.6: The raw axial velocity field extracted from the acoustically forced simulation for passages A and B at one instant (a) and the same for the PODFS approximation (b).

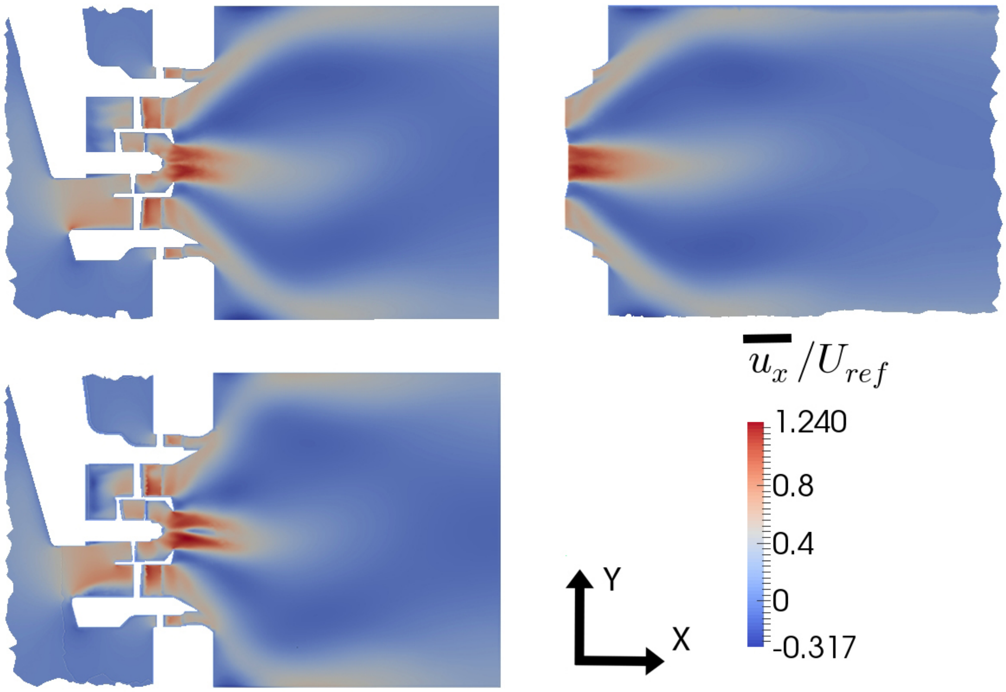

(a) Mean axial velocity.

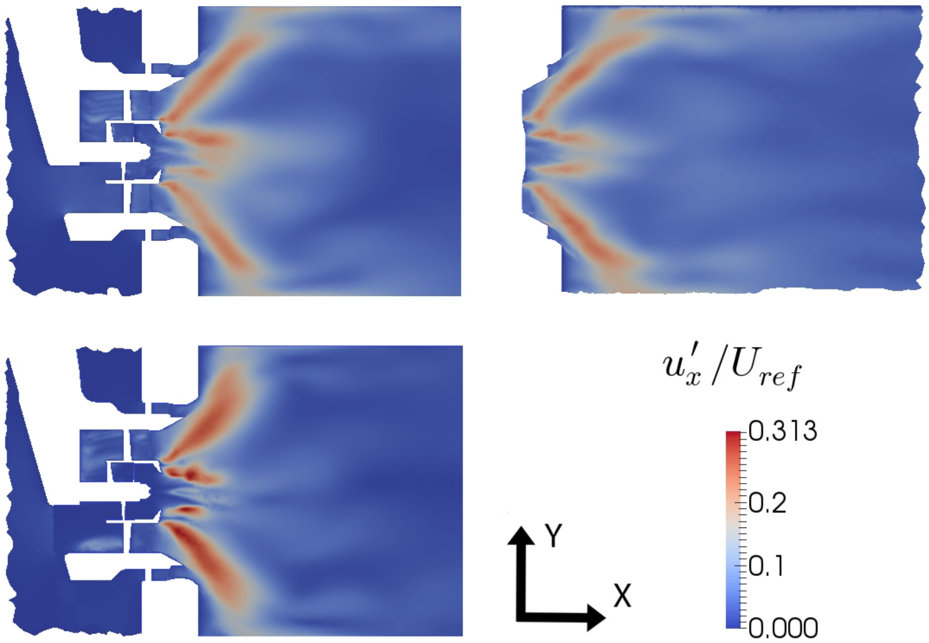

(b) RMS of axial velocity.

Figure 6.7: The mean and RMS axial velocity from the donor simulation (top left), mass forced simulation (bottom) and the PODFS simulation (right).

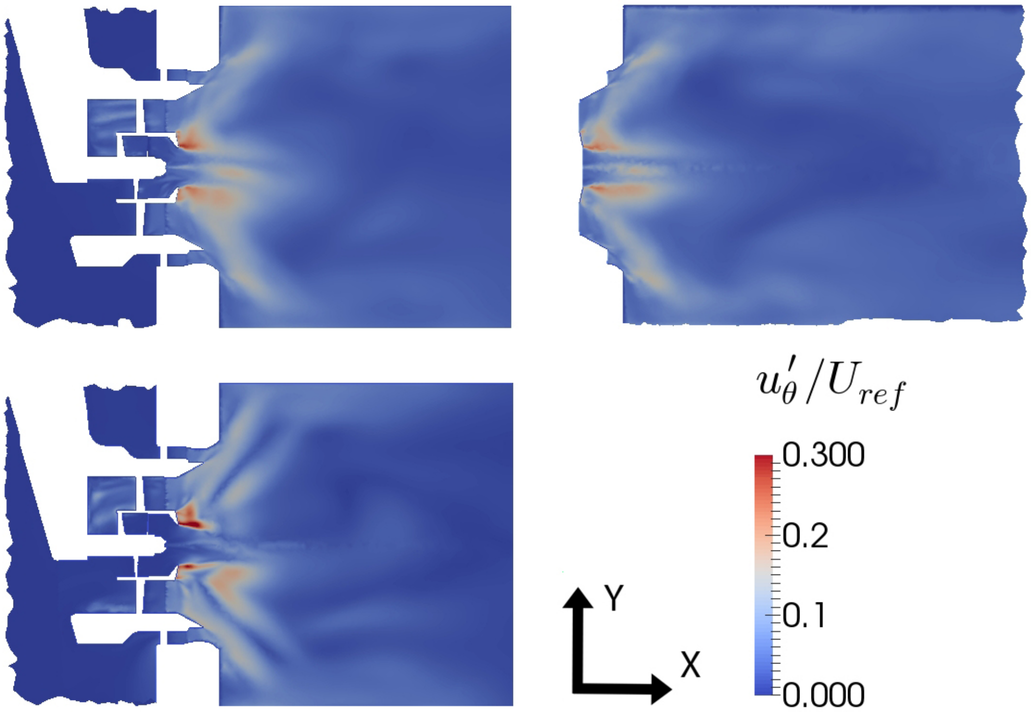

Figure 6.8: RMS of azimuthal velocity from the donor simulation (top left), the mass forced simulation (bottom) and the PODFS simulation (right).

correctly is due to a few separate mechanisms: firstly, the wavelength of the mass forced fluctuations is much smaller than in the case of the acoustic forcing which leads to the double peak in azimuthal velocity fluctuations in the axial direction seen in Figure 6.8. Secondly, because the speed of sound is infinite, the directionality of the acoustic wave is not captured leading to the incorrect relationship between the azimuthal and axial fluctuations of velocity in the swirler passages and at the exit of the injector. Finally, because the simulation is incompressible, neither the *non-linear acoustic feedback* mechanism nor the *mean pressure field mechanism* as described in section 5.6 can be active in this simulation.

In addition to the already stated benefits of the model, Table 6.2 shows the cell count, number of timesteps and CPU hours required for each simulation type. The mass flow forcing method benefits from a larger timestep, reducing the computational time by a factor of two. The strong recycling method benefits from a smaller timestep and a much smaller mesh however the time it takes for the loading of the plane data for each timestep makes the method more expensive than the mass forcing method. The PODFS method must calculate the Fourier series for each mode at each timestep which may include hundreds of coefficients but because the spatial data is restricted to 30 modes per inlet that are loaded at the beginning of the simulation, the PODFS method is 2.73 times faster than the recycled simulation and 6.3 times faster than the original acoustically forced simulation.

Table 6.2: Comparison of mesh sizes and running times of the 4 simulations.

Forcing	Cell count	Time steps	CPU hours
Acoustic	9.2 million	40000	3872
Mass flow	9.0 million	2000	1428
Recycled	2.6 million	2000	1678
PODFS	2.6 million	2000	615

6.5 Conclusion

Some of the conventional methods for imposing turbulent and acoustic fluctuations of velocity onto the inlet of a simulation have been discussed with the aim of reproducing these effects in an incompressible framework. A set of simulations have been run to test the effectiveness of two of these methods along with a third novel approach called the PODFS method in reproducing the effects of acoustic forcing. The two conventional methods are the strong recycling method where the results of an auxiliary simulation are applied directly to the inlet, and the mass forcing method which simply scales the inlet velocity sinusoidally in time. The strong recycling method has been shown to be quite accurate, however, comparatively expensive given the significantly simplified domain while the mass forced method fails to predict the mean and fluctuating velocity fields, relative magnitude of mass flow through each injector flow passage or the correct swirl number at either the passage exits or the injector exit plane.

The new PODFS method uses the POD to significantly reduce the size of the data set required to apply the strong recycling method. The temporal POD modes are approximated using a Fourier series that gives the model flexibility in running time and timestep that the strong recycling method does not have. Despite the code having to compute hundreds of Fourier coefficients at each timestep, because all of the reduced data set is loaded at the beginning of the simulation, the model performs much faster than every other method tested. By scaling the mean velocity using a small coefficient, the model has been shown to effectively reproduce both the mass flow rates and swirl numbers for each flow passage and the injector as a whole.

The generation of a turbulent inflow using any of the methods described has not been tested in this chapter. In the following chapter, the PODFS method as generated from the acoustically forced URANS simulation, is applied to the inlets of a LES simulation, the aim of which is to analyse the convection of the fuel spray. Due to time constraints, a turbulent inflow method is not applied or tested in addition to the PODFS. While this is a likely source of error, it was the opinion of some members of industry that the large scale structures captured by the PODFS method would be sufficient to generate the required turbulence in the combustion chamber [172]. This was based on the observation that a large proportion of the turbulence is generated downstream of the injector where the counter-rotating streams are brought into contact.

Chapter 8 contains a test of using the PODFS method to apply additional turbulence to each inlet taken from the large incompressible LES computed earlier in the thesis. The results show that inlet turbulence is important for stabilising the pilot jet and flame and that using the PODFS for this purpose is an effective strategy.

Chapter 7

Simulating the two-phase flow and spray dynamics

PRECISE-UNS uses the Eulerian-Lagrangian framework for predicting the two-phase flow field. This method resolves the main working fluid (air, fuel vapour and products) as a Eulerian field with the liquid fuel spray being represented by Lagrangian particles. The Eulerian field imparts forces and temperatures onto each Lagrangian particle and each Lagrangian particle imparts a force, mass flux and heat flux back onto the Eulerian field.

7.1 Atomisation

In order to produce a combustible mixture from the liquid fuel, the fuel must be atomised into a fine spray that will then evaporate and mix with the surrounding air. The model injector chosen for this study has two atomisers, one for the main flame and one for the pilot. The main is fed with an airblast atomiser while the pilot is fed with a pressure swirl atomiser.

Atomisation is difficult to model, however the critical condition for drop breakup is when the surface tension force of the fuel is equal to the aerodynamic drag [73]. By equating these two forces and using a representative drag coefficient of 2/3, an estimate for the largest droplet produced in the atomisation process is given as:

$$D_{\max} = \frac{12\sigma}{\rho_A U_R^2} \tag{7.1}$$

where σ is the surface tension of the fluid, ρ_A is the air density and U_R is the relative velocity between air and fuel. Such an expression shows that the droplet diameter is inversely proportional to the square of the velocity while all the other parameters are fixed by the type of fuel and the operating point of the engine.

7.1.1 (Prefilming) Airblast atomisers

Prefilming airblast atomisers (or just airblast atomisers) rely on high velocity air and turbulence to generate the required atomisation of the fuel. The fuel is injected onto a conical prefilming surface where it forms a thin film that is pushed downstream and away from the prefilming surface by the air. Two air streams are usually employed, one that is directed onto the prefilming lip and one that interacts with the first to generate high levels of turbulence that help to break the liquid sheet into small droplets. In this type of atomiser the velocity of the air passing through the injector has a strong effect on the atomisation diameter. The maximum achievable velocity is set by the pressure drop across the combustor liner and hence it is sensitive to the operating point of the engine, however effective atomisation is achievable across a broad range of fuel flow rates giving this type of atomiser a very good operating range. As turbulence also contributes to the atomisation [73], the promotion of additional turbulence through the use of counter-rotating jets, trips or vortex generators could be used to modify the particle diameter distribution. As the pressure drop across the combustor liner is an important factor in the atomisation process it can be assumed that the addition of the acoustic forcing will affect the atomisation process and hence change the distribution of fuel droplets produced.

7.1.2 Pressure-swirl atomisers

In a pressure-swirl atomiser the relative velocity U_R is provided by the high injection velocity of the fuel. The fuel is injected into a swirling chamber terminating in an orifice that opens into the combustion chamber. As the fuel is incompressible the injection velocity is proportional to $\Delta p^{1/2}$. Provided that the injection pressure is around 0.1 MPa higher than the ambient pressure, the swirling that is imposed on the fuel results in a conical sheet that terminates in a hollow-cone spray [73]. The advantage of such an atomiser is that the atomisation is somewhat independent of the air velocity. The disadvantage is that these types of injectors have a narrow range of operating flow rates. This is because large flow rates require extremely high pressures that are not achievable with conventional pumps. Airblast atomisers are also preferred because instead of redirecting mechanical work to high pressure fuel pumps, the fuel is broken up by the high velocity air that is required for flame stabilisation regardless of the atomisation method.

7.1.3 Secondary atomisation

The processes described above can be considered to be part of the *primary atomisation* process. Secondary atomisation occurs once the droplets have been formed but may interact with other spray droplets to coalesce into larger, or break up into smaller, droplets. As the atomisation is not modelled directly in this study, it is assumed that the effects of secondary atomisation are taken into account in the primary atomisation processes already discussed.

7.2 Modelling

The modelling of the fuel spray within the combustion chamber requires a method to represent atomisation, a method to represent convection and a method to represent evaporation. This study chooses not to focus of the physics of atomisation but instead to use semi-empirical relations. Convection and evaporation are achieved using a standard Lagrangian approach.

7.2.1 Spray injection

In order to avoid modelling the atomisation directly, the fuel droplets have been injected using a known size distribution at a velocity representative to those found in experimental studies [18]. In order to inject an experimentally representative spray into the combustion chamber, it is important that sufficient randomness of location, angle and spray velocity is captured by the method.

7.2.2 Fixed diameter distribution with randomised injector location

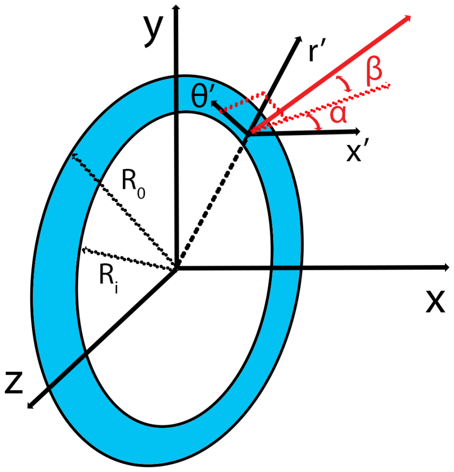

Figure 7.1: The injection zone for spray droplets

The spray is injected at a number (N_{SL}) of locations at each timestep at a random location on a ring as defined in Figure 7.1. The location of the centre of the ring and the mean diameter is defined at the beginning of each simulation along with the ring orientation and ring thickness. The radius and azimuthal location of each fuel injection point is selected using a random number generator and a droplet of fuel injected into the flow.

Each droplet of fuel is given an initial velocity magnitude ($|\mathbf{u}_d|$) as defined by the user with a radial injection angle (α) that is selected from a truncated Gaussian distribution with a selected mean, variance, minimum and maximum value. The values for each of these parameters was taken from values used in previous work by Su et al. [148] using OpenFOAM. The parameters used in this study are included in Table 7.1. PRECISE-UNS does not allow the radial location of injection to be totally random. Instead, a number of equispaced rings can be defined between the outer and inner ring diameter. The ring that is used for droplet injection is chosen at random for each parcel. OpenFOAM also allows for the droplet velocity to be defined in both the radial and azimuthal directions at the same time. PRECISE-UNS can add swirl to the droplets but at the cost of setting $\alpha = 0$. Given the large variation in α used in the OpenFOAM simulations it was decided that reproducing α was of higher importance than the swirl angle β. The diameter of each droplet is chosen from a predefined list of diameters with a probability of each diameter size given by the user.

Table 7.1: Injection parameters for droplets for various simulations.

Simulation	OpenFOAM	PRECISE-UNS
x_0	$-0.13D$	$-0.13D$
y_0	0	0
z_0	0	0
R_i	$0.40D$	$0.40D$
R_o	$0.41D$	$0.41D$
α_{min}	-15	-15
α_{max}	30	30
β	63	0
$\lvert\mathbf{u}\rvert$	$0.1U_{ref}$	$0.1U_{ref}$
Number of rings	N/A	10

Whilst this is the simplest method used for injecting the spray particles and only loosely models reality, it is useful for decoupling the effects of convection from the effects of atomisation. As the velocity is highly unsteady close to the atomiser and as the velocity has a large influence over the diameter of the spray droplets, the presence of large coherent structures and forcing will likely alter the atomisation process and hence the droplet diameter and velocity. The droplets, once formed, are then convected by the flow where they are also influenced by the unsteady flow field. The question is therefore which of the two effects influence the mixture fraction and flame response more: The affect of a changing diameter driven by the atomisation process, or the effect of large scale structures on the dynamics of the spray? By fixing the atomisation diameter, or diameters, the effect of the velocity field on the spray dynamics may be investigated.

7.2.3 Rosin-Rammler diameter distribution

In a real atomisation process there is a natural variation in the diameter of the produced droplets. This distribution of droplet sizes is often approximated by the Rosin-Rammler distribution [73]:

$$1 - Q = exp\left(-(D_s/D_{32})^n\right) \tag{7.2}$$

Q is the proportion of the total volume contained in drops with a diameter less than D_s, n is a constant that defines the spreading of the distribution and D_{32} is the Sauter Mean Diameter (SMD). n and D_{32} are usually found by fitting the Rosin-Rammler function to a set of experimental data. Figure 7.2 shows the Rosin-Rammler distribution using various n values against an experimentally derived distribution using PDA. It is noticeable that in the experimental case the distribution appears to be biased towards the smaller droplets and flatter overall. This experimental data was recorded using injector I1 at sea level temperatures and pressures. It is likely that at higher operating pressures, the increased peak velocities cause the distribution to move towards smaller particles. This is because given a higher operating pressure with the same percent combustor pressure drop, U_R in Equation 7.1 will increase.

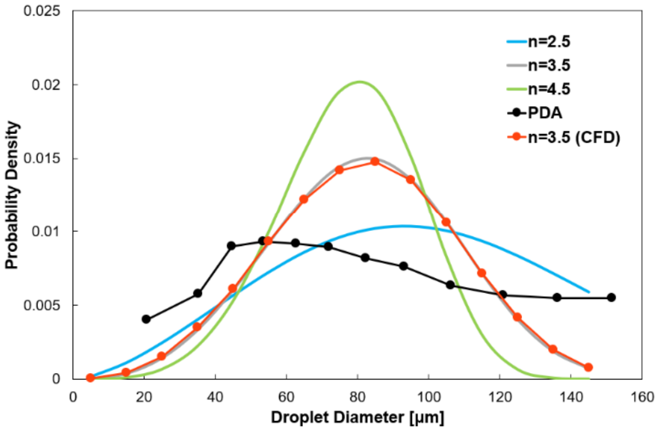

Figure 7.2: The histogram of droplet sizes produced from the LDA data of the fuelled chosen injector. Reproduced from Su et al. [148].

7.2.4 Phase locked Rosin-Rammler diameter distribution

Using the standard Rosin-Rammler distribution as defined above does not take into account any fluctuations in droplet statistics as a function of time. As the acoustic forcing generates a fluctuation in the mass flow rate through the injector, it is likely that the diameter of the fuel droplets will change over the acoustic cycle. It therefore makes sense to make the parameters of the Rosin-Rammler distribution a function of phase angle of the acoustic fluctuation:

$$1 - Q(\phi) = exp\left(-(D_s(\phi)/D_{32}(\phi)^{n(\phi)}\right) \tag{7.3}$$

again, $n(\phi)$ and $D_{32}(\phi)$ may be found by fitting the distribution to phase averaged data.

Su et al. [148] used correlation functions developed by Lefebvre [71] and Jasuja [57] to reproduce the fluctuations of droplet SMD caused by the acoustic excitation. The Jasuja correlation is governed by:

$$D_{32} = \left[A\frac{(\sigma\rho_L)^{0.5}}{\rho_A U_A} + B\left(\frac{\mu_L^2}{\sigma\rho_A}\right)^{0.425}\right]\left(1 + \frac{\dot{m}_L}{\dot{m}_A}\right)^{0.5} \tag{7.4}$$

where A and B are constants, σ, ρ_L and $\dot{m}_L$ are the surface tension, density and mass flow rate of the fuel and ρ_A, U_A and $\dot{m}_A$ are the density, velocity and mass flow rate of the air. Figure 7.3 shows the SMD calculated using this correlation with the same A and B values as calibrated by Su et al. [148]. In this case the air velocity fluctuations are extracted from the phase and circumferentially averaged velocity field of a forced simulation just behind the prefilmer lip while the air mass flow rate is taken as the combined mass flow through passages B and C. This simulation is of the SCARLET rig at engine representative conditions being acoustically forced at $St = 0.14$. The variation of SMD shows a mean of around $2.22 \times 10^{-4} D$ microns which is consistent with values used in industry [143] with a fluctuation of around 1.2%. Previous work at sea level conditions has shown much larger fluctuations due to acoustic forcing of around 16% [18, 148] but whether or not such high fluctuations are seen in engine representative conditions remains an open question. The

Lefebvre correlation [71] was also tested, however at these elevated pressures the SMD's produced appeared to be too high for engine representative conditions. There appears to be a deficiency in experimental data of spray distributions at elevated temperatures and pressures let alone under acoustically forced conditions. The SCARLET rig presents a good opportunity in the future to help with this problem.

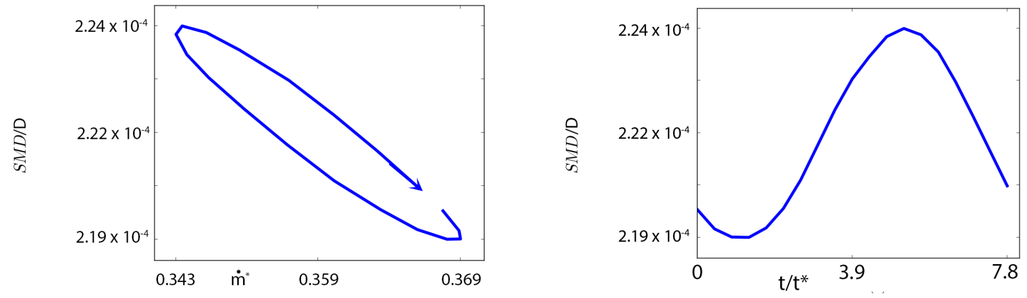

(a) SMD as a function of mass flow through passages B and C.

(b) SMD as a function of time.

Figure 7.3: Phase and azimuthally averaged SMD from a simulation of the SCARLET rig forced at $St = 0.14$ using the Jasuja correlation.

7.2.5 Velocity correlated diameter distribution

The use of the phase and circumferentially averaged velocity and mass flow rates results in quite a small variation of SMD. Given that the atomisation and fuel transport are governed by non-linear processes, it is possible that local flow structures and their associated large fluctuations in flow velocity may cause large changes in SMD that may amplify the effects of the acoustic forcing, especially when the forcing may amplify the magnitude of these helical structures. It would therefore be preferable to extract the flow velocity at each injection location and use that to determine the local SMD during the simulation. Figure 7.3a shows the same characteristic *phase portrait* seen by Su et al. [148]. However this cycle may collapse onto itself or reverse direction as the frequency of excitation is increased. This is because there is a time delay between the increase in air velocity and it affecting the atomisation of the droplet. The droplet must also convect to the point where the droplets are injected. These delays mean that the relationship between velocity, mass flow and droplet SMD should be expressed as a complex transfer function. Su et al. [148] made an attempt to model this transfer function using two different delay filters with mixed results.

Whilst this method is attractive and in the long term may significantly improve the prediction of the flame transfer function, it has been neglected in this study because of the difficulties in modelling the response and the lack of data available at elevated temperatures and pressures. Su et al. [148] found that simple delay filters only provided a weak representation of the data. The phase relationship between fluctuations of velocity, mass flow rate and droplet SMD as a function of frequency is likely to be dependent on local pressure which means that sea-level experimental data should be used with caution. Rather than assume that the Su et al. correlations are accurate for the real engine case, the effects of the acoustically modulated air flow through the injector on heat release and the effects of SMD fluctuations on heat release are tested separately in Chapter

8. Future work will aim to combine these two effects in one simulation which accounts for both mechanisms.

7.2.6 Spray transport and evaporation

The mass, momentum and temperature (or enthalpy) is calculated for each droplet at each timestep using standard Lagrangian particle tracking. Further details may be found in Appendix A.6 or in [4].

7.3 Spray Statistics

PRECISE-UNS contains as part of its Lagrangian solver the ability to calculate some statistics on the fuel spray as viewed in the Eulerian frame of reference. The quantities that are calculated include the SMD (equation 1.12), number density defined as:

$$N_\rho = \frac{\sum_i D_{d,i}}{D_{32}^3} \tag{7.5}$$

the spray velocity defined as:

$$\mathbf{u}_s = \frac{\sum_{d,i} \mathbf{u}_{d,i} D_{d,i}^3}{\sum_i D_{d,i}^3} \tag{7.6}$$

the spray mass defined as:

$$m_s = \rho_f \pi D_{32}^3 N_\rho / 6 \tag{7.7}$$

where ρ_f is the density of the fuel, and the spray momentum is defined by:

$$\mathbf{p}_s = m_s \mathbf{u}_s \tag{7.8}$$

These statistics have been augmented with an additional one as used in PDA measurements at Loughborough University, the spray flux, defined as:

$$f_s = \frac{m_s}{\delta t S_P} \tag{7.9}$$

where δt is the time between samples and S_P is the volume of the Pth cell. It is often desirable to divide the droplets into groups or bins. The bins may be based on the individual droplet diameter. Such an approach is interesting for seeing how individual droplets respond to the forces imposed upon them by the airflow or how long it takes before the droplet evaporates. Another approach is to divide the droplets into groups based on how much they contribute to the total mass of droplets passing through the interrogation volume. This approach is interesting as it is the mass of fuel, not the volume of each droplet, which will contribute to the heat release rate. Binning of droplet statistics by contribution to spray mass or diameter was implemented in PRECISE-UNS along with the spray flux calculation.

The binning by diameter function works by dividing the maximum specified droplet diameter by the number of bins and sorting each droplet accordingly when required. In the case of the droplet diameter being chosen by a Rosin-Rammler distribution, the largest droplet size is chosen as the mean SMD plus three times the standard deviation.

If the bins are chosen to be of equal mass, every cell containing droplets must sort each droplet present by diameter and the total droplet mass must be calculated. Then the droplets are put into the first bin until the mass present is greater or equal to the total droplet mass divided by the number of bins. This process is computationally expensive because of the sorting operation. Both binning operations, and spray statistics more generally, collected using CFD calculations can sometimes be difficult to draw conclusions from due to the relatively small number of droplets tracked and the small number of temporal snapshots when compared to experiments.

If a simulation (or experiment) is run at ambient conditions with no combustion then there is unlikely to be much evaporation of the fuel. Because there are regions of reverse flow, droplets may recirculate several times before exiting the region of interest. In the case of reacting flow, droplets will evaporate over time which means that any statistics collected will differ between hot and cold conditions. For this reason an additional feature in the code was developed to track the time since injection of each droplet and then to stop including the droplets in the statistics after a certain time, calculated to correspond with the droplet evaporation time in the reacting case. Because the tracking of droplets that are no longer being used for evaluating the spray statistics is a waste of computing resources, the option to stop tracking particles after a certain time was also included.

7.3.1 Droplet evaporation times

The droplet evaporation times are first estimated in Section 1.5.3. In order to check the relevance of these numbers, the time till evaporation for a range of droplets from a reacting LES simulations by Andrea Giusti at Cambridge University were acquired.

Figure 7.4 shows the distribution of droplet diameters injected in this simulation and the evaporation times of each droplet as compared to the estimated values. It shows that despite the droplet diameters chosen for the approximation being outside the range of values used in the simulation, that the estimated evaporation times are roughly consistent with the values obtained from Cambridge.

7.4 The effect of the flow field on spray transport

The purpose of this test case was to highlight the impact of the fluctuating flow field on the transport of the fuel spray. The acoustically forced fluctuating flow field is already known to affect the atomisation but it is not well understood as to the importance of the fluctuating flow field to the transport of the fuel spray. The flow field is extremely chaotic inside the combustion chamber which has the desired effect of distributing the fuel spray throughout the volume. As the phase averaged, acoustically forced, flow field is only a fraction of the total flow fluctuation energy (see Chapter 5) it may be the case that the acoustic forcing only causes a small change to the

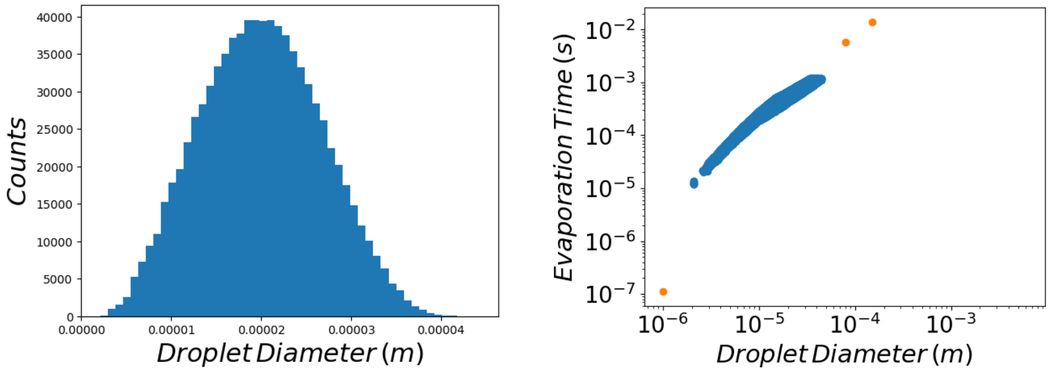

(a) The histogram of droplet diameters.

(b) The evaporation times for each droplet as a function of droplet diameter for the simulation (blue) and the three approximations made in Section 1.5.3.

Figure 7.4: Droplet diameters and evaporation times from the Cambridge data.

dispersion of the fuel droplets when compared to the effect of the forcing on atomisation or direct flow field-flame interaction.

In order to test the effect of forcing on spray transport, an LES simulation with the same truncated domain as used in the PODFS simulation in Chapter 6 was run with the injection of three different diameter fuel droplets. The droplet sizes were chosen from what has been visualised in PDA experiments at Loughborough University such that they represented the very smallest ($1.4 \times 10^{-5}D$), average ($7.0 \times 10^{-4}D$) and very largest droplets ($2.1 \times 10^{-3}D$) produced during atomisation. At each timestep 100 nozzles are chosen at random locations and a parcel of droplets injected for each injection size. The simulation fields were saved every 200 timesteps where the droplets were sorted into their three different sizes and three different mass bins. In order to track the droplets throughout their life, the evaporation mechanism was deactivated and each particle was only counted in the statistics for $14t^*$ to ensure that only relevant particles were included in the statistics. Particles older than this were also no longer tracked to improve the simulation efficiency.

The computational domain was exactly the same as the recycled and PODFS simulations run in the previous chapter except that the mesh was refined in the primary zone to an LES standard, this resulted in a tetra-dominated mesh of approximately 38 million cells.

The timestep was set to $1 \times 10^{-3}t^*$ using a second order backwards temporal integration scheme with the spatial differencing achieved using a second order linear upwind methodology for the velocity variables and pressure using a 2nd order symmetrical method. A Van Albada limiter was set for the gradients of pressure and pressure correction variables. Turbulence closure was achieved using the constant Smagorinsky LES model. The simulation was first run as a steady flow calculation for 4000 iterations before being run for 25,000 steps unsteady to get through any transients. The flow statistics were then calculated over the next 40,000 steps.

In order to simulate forcing the upstream boundaries were forced using the PODFS method at $St = 0.3$ with the same input data used in the previous chapter. No additional modelling was used to take account of the turbulence that is not resolved in the donor URANS simulation as it was hoped that sufficient turbulence would be generated at the point where the air streams come into contact.

7.5 Results and Discussion

7.5.1 Mean field

The mean velocity field was constructed from every 200th snapshot over the 40,000 timesteps run. The mean axial velocity is shown in Figure 7.5a and the mean azimuthal velocity is shown in Figure 7.5b. They show that the jets coming from passages B and C have been captured accurately by the method. The penetration of the jet emanating from passage A has been significantly diminished as compared to the previous simulations and the PIV data (see Chapter 3). Furthermore an additional recirculation zone has opened in the centre of the pilot jet. This flow configuration is referred to as being *flipped* and has been found in previous reacting and non-reacting experiments to be dependent on the upstream flow field [142, 162]. Given that both the PODFS forced URANS simulation, compressible URANS simulations, PIV experiments and the large unforced LES reproduced the mean flow consistently, it would seem likely that the error experienced in this simulation is due to insufficient turbulence in the pilot zone.

In Chapter 8 additional PODFS modes were computed from the large incompressible LES simulation to account for this missing inlet turbulence which was then shown to correct the pilot jet penetration. Regardless of whether or not the pilot jet flow is captured accurately, the flow in the region of the main atomiser is captured correctly, and so the spray dynamics in this region are likely to be representative of the non-flipped flow configuration. The simulation is also simply seeking to demonstrate the link between flow structures and the fuel spray and as such the accuracy of the flow itself is of secondary importance.

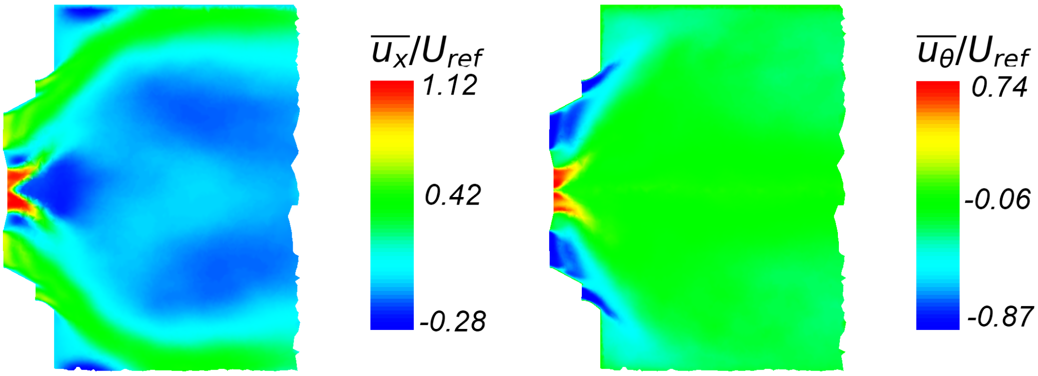

(a) Axial velocity. (b) Azimuthal velocity.

Figure 7.5: The mean velocity field of the LES spray simulation.

Figure 7.6 shows the spray parcels at a single snapshot in time coloured by diameter. It shows that the smallest droplets are dispersed throughout the field and the ORZ. The largest droplets follow the main airflow more closely before being broken up by large scale flow structures. The mid size particles show a behaviour that lies between the two. They do not remain in the main jet for as long and appear to be broken up by large scale structures, however the droplets are not dispersed throughout the flow.

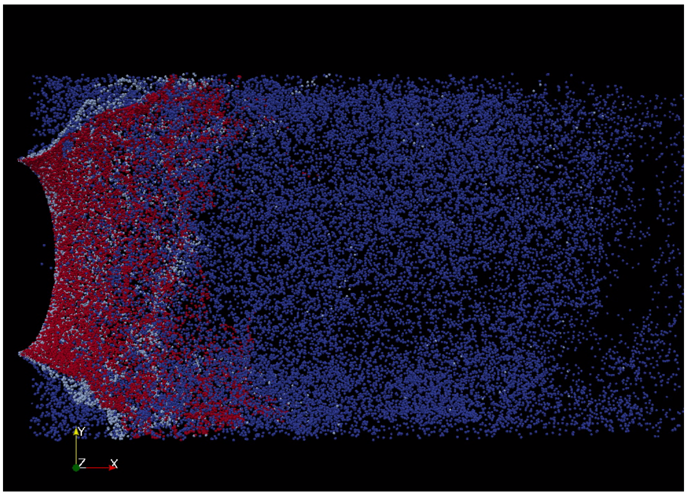

Figure 7.6: A snapshot of the fuel spray coloured by diameter. Large droplets are coloured red, mid size droplets are coloured white and the smallest droplets are coloured blue.

The time and azimuthally averaged normalised spray mixture fraction, or the mass fraction of liquid fuel in each cell, is shown for each size of droplet in Figure 7.7. This figure confirms that the smallest droplets are dispersed throughout the flow with the mean path taking them outwards in comparison to the other droplets. They also seem to accumulate in the ORZ. This is because as they are so light, they quickly accelerate up to the flow azimuthal velocity. Whilst the particles accelerate easily, their newly obtained azimuthal velocity acts as a centrifuge and pulls them outwards out of the main jet and into the ORZ.

The largest droplets also enter the main stream but take longer to accelerate, their high inertia prevents them from being accelerated azimuthally and centrifuged out of the main. Because they stay in the main flow for longer, they end up being flung out further downstream. The mid-size particles also remain in the main flow for longer than the smallest ones and are pushed radially outward faster than the largest. The larger droplets are not dispersed in the flow as easily as the smallest ones because only large flow structures have the energy and timescale large enough to significantly affect the motion of these droplets.

7.5.2 POD analysis

It was seen in chapters 4 and 5 that the unsteady flow was dominated by large scale helical structures. Whilst the small scale droplets are likely to be influenced by smaller scale turbulent fluctuations, it might be expected that larger droplets will be disproportionately affected by these large flow structures whose characteristic frequencies are closer aligned with the droplets themselves.

In order to test how droplets of different sizes respond to flow structures of different characteristic frequencies, the POD of velocity and different sized droplets were performed. As the flame heat release rate will depend most strongly on the mixture fraction of fuel to air, this was chosen

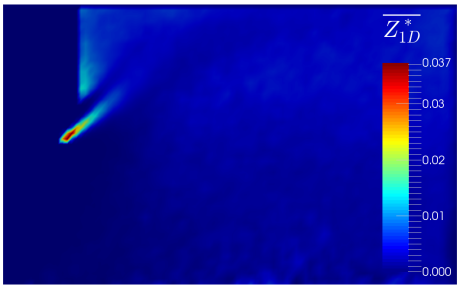

(a) Smallest droplets.

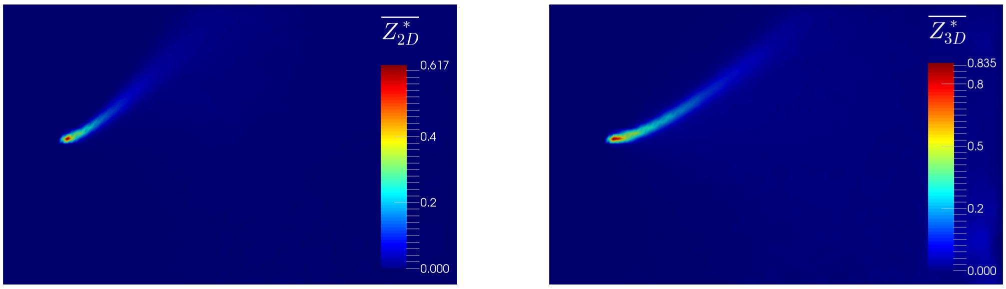

(b) Mid size droplets.

(c) Largest droplets.

Figure 7.7: The mean mixture fraction field for the three droplet sizes in the LES spray simulation.

as the most representative quantity for testing. The POD was performed using the snapshot POD method first defining the square of velocity to be the inner product used:

$$C_{i,j} = \frac{1}{N_s} \int_{\mathbf{x}} \mathbf{u}(\mathbf{x}, t_i) \cdot \mathbf{u}(\mathbf{x}, t_j) d\mathbf{x} \tag{7.10}$$

Following this, the POD was computed three more times with the inner product defined using the normalised velocity and mixture fraction for each of the three droplet sizes. The inner product is therefore defined for the smallest droplets as:

$$C_{i,j} = \frac{1}{N_s} \int_{\mathbf{x}} \mathbf{u}^*(\mathbf{x}, t_i) \cdot \mathbf{u}^*(\mathbf{x}, t_j) + D^*_{1D}(\mathbf{x}, t_i) D^*_{1D}(\mathbf{x}, t_j) d\mathbf{x} \tag{7.11}$$

The azimuthal mode numbers, as identified using the method outlined in Section A.3.2, of the 10 most energetic POD modes for each droplet size are shown in Table 7.2 with the corresponding peak oscillation frequencies shown in Table 4.3. They show that the POD modes are consistent across the three POD different droplet sizes. The POD analysis of velocity only is slightly different in that the 5th most energetic mode is replaced by the 7th most energetic in the case of the POD including the droplets. This shows that the droplets are affected by the $St = 0.3$ forced $m = 0$ mode in preference to the $m = 2$ helical mode. This may be due to the lower mode frequency or that the mode strength is higher in the area where the particles are in high concentration levels.

Table 7.2: Dominant azimuthal wavenumbers for the first 10 POD modes for the simulations according to four different inner products.

Inner product	$\mathbf{u}$	$[\mathbf{u}^*, Z^*_{1D}]$	$[\mathbf{u}^*, Z^*_{2D}]$	$[\mathbf{u}^*, Z^*_{3D}]$
mode 1	3	3	3	3
mode 2	3	3	3	3
mode 3	1	1	1	1
mode 4	1	1	1	1
mode 5	2	0	0	0
mode 6	2	2	2	2
mode 7	0	2	2	2
mode 8	0	0	0	0
mode 9	0	0	0	0
mode 10	0	0	0	0

Figure 7.8 shows the eigenvalues of the POD modes for each droplet diameter. It shows that the mode energies are broadly consistent for the first 10 modes or so. For higher order modes there is a clear distinction with the smallest droplets containing more energy than the larger droplets. The lowest energy is held by the mid-size droplets in this range. These higher order modes are representative of small scaled structures and turbulence so it is expected that the smallest droplets are affected the most, however it is unexpected that the largest droplets seem to respond to these fluctuations more than the mid-size droplets.

Table 7.3: Dominant Strouhal number for the first 10 POD modes for four different inner products.

Inner product	$\mathbf{u}$	$[\mathbf{u}^*, Z^*_{1D}]$	$[\mathbf{u}^*, Z^*_{2D}]$	$[\mathbf{u}^*, Z^*_{3D}]$
mode 1	1.3	1.3	1.3	1.3
mode 2	1.3	1.3	1.3	1.3
mode 3	0.45	0.45	0.45	0.45
mode 4	0.45	0.45	0.45	0.45
mode 5	0.90	0.30	0.30	0.30
mode 6	0.90	0.90	0.90	0.90
mode 7	0.30	0.90	0.90	0.90
mode 8	0.30	0.30	0.30	0.30
mode 9	0.60	0.30	0.60	0.60
mode 10	0.30	0.60	0	0

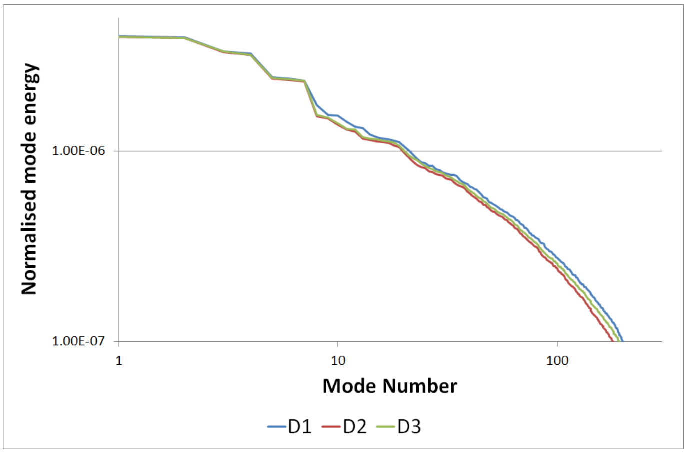

Figure 7.8: The POD mode energy for the POD of velocity and smallest droplets (blue), mid-sized droplets (red) and largest droplets (green).

In order to see the relationship between Stokes number and particle response, the modes were normalised by the total energy in each POD. The Stokes number of each mode and droplet combination can be calculated as [73]:

$$Stk(D,k) = \frac{\tau_d}{\tau_f} = \frac{\rho_d D^2 f_k}{18\mu} \tag{7.12}$$

where τ_d is the characteristic timescale of the droplet, τ_f is the characteristic timescale of the fluid, ρ_d is the droplet density, f_k is the characteristic frequency of the kth mode and μ is the air viscosity. Plotting the mode energies against Stokes number results in Figure 7.9. The correlation shows that in the majority of cases, the most energetic mode of each mode type belongs to the mid size particle POD modes, while the largest droplet POD modes are around 3% lower and the smallest particles are around 5% lower.

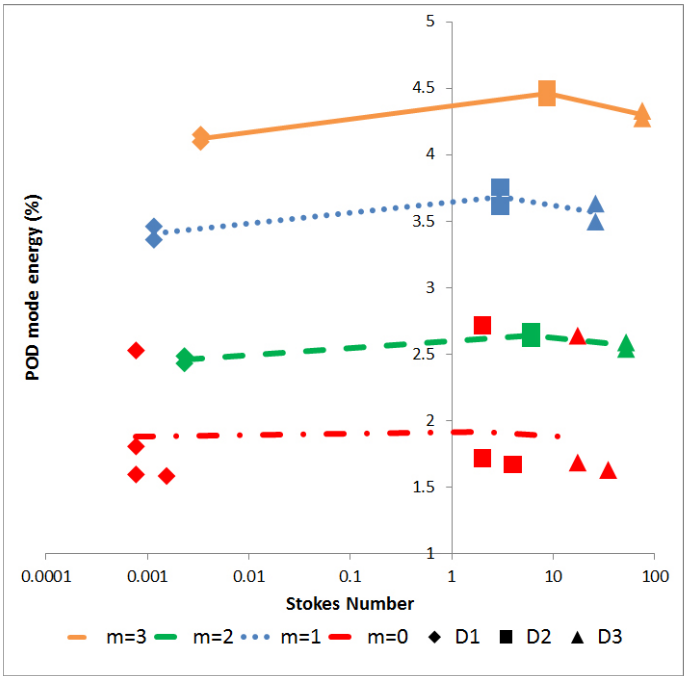

Figure 7.9: The normalised POD mode energies for the first 9 POD modes for the smallest droplets (diamonds), mid-sized droplets (squares) and largest droplets (triangles). The lines trace the average for each mode type.

These results are supported by previous work that has shown that droplets with a Stokes number between 1 and 10 are most likely to be dispersed by the flow [83]. If the droplets are atomised in the diameter range as used in the Cambridge data then it can be said that the atomisation is unlucky. Unlucky in the sense that the Stokes number of these droplets is in the range of the greatest susceptibility to the acoustic forcing frequency and large scale structures in the combustion chamber. Because atomisation diameter is inversely proportional to the local air velocity squared, it would seem more sensible to increase the size of the droplets rather than decrease them. The problem with this option is that larger droplets take longer to evaporate leading to localised regions of high mixture fraction which will increase soot formation [86]. It would therefore be preferable to decrease the size of atomised particles if possible. The problem is that the maximum velocity of the air available to atomise the fuel is set by the pressure drop across the combustion liner which is becoming smaller and smaller in an attempt to gain better thermal efficiencies of the engine. Another possible option is to atomise using a different method. The atomisation of the fuel could be assisted by using pressure swirl atomisers running at very high fuel pressures. This would have considerable downsides as it would increase the cost and weight of the fuel system while reducing the reliability and increasing the losses. Because pressure swirl atomisers are only functional over a small range of flow rates, it would be likely that additional stages of fuelling would be required further adding to the complexity of the fuel system. One possible advantage of such a system would be that the extra stages of fuelling would give additional methods of controlling instabilities as the relative flow rates of each stage could be used to alter the location and stability of the flame. Nevertheless given that a reduction of the droplet diameter of a factor of 50 would only result in a reduction of sensitivity of around 5%, it is likely that other methods of controlling the sensitivity of the combustion system will be employed in the future.

7.5.3 Phase and azimuthal average

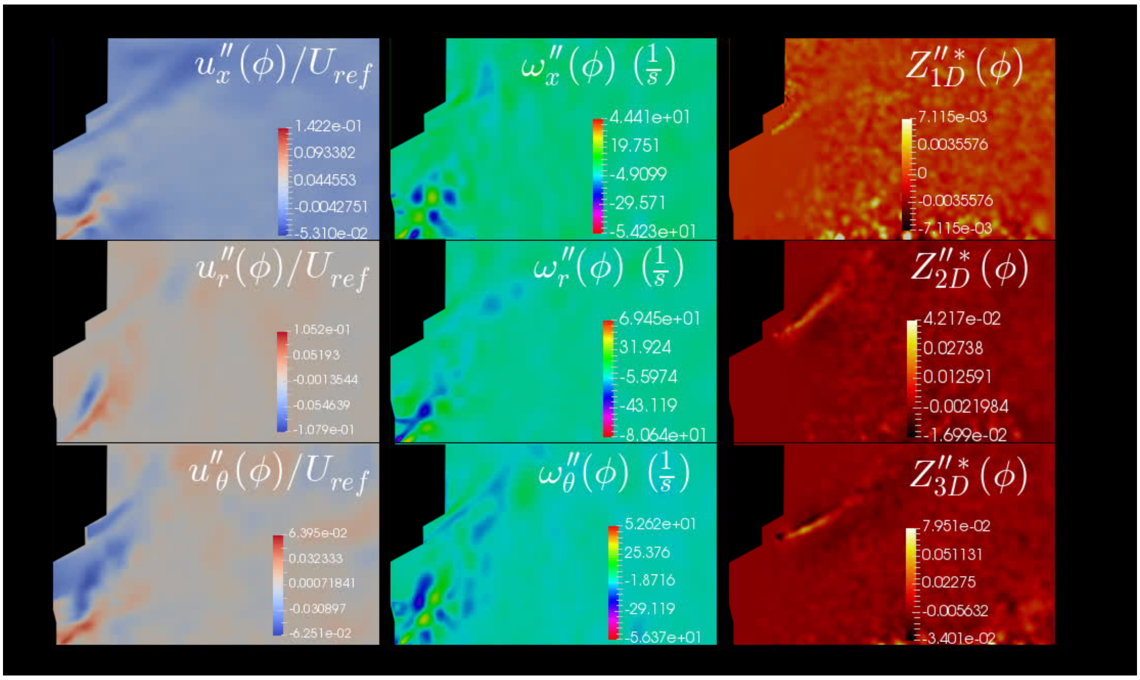

Figure 7.10: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 0$.

In order to see how the droplets are affected by the acoustic forcing, the simulation data was phase averaged at $St = 0.3$ by linearly interpolating the data between corresponding saved snapshots. Once phase averaged, the results were then azimuthally averaged to improve the resolution of the $m = 0$ acoustic induced fluctuations. Figures 7.10-7.19 show the phase averaged mixture fraction for the three different droplet sizes at various stages of the acoustic cycle along with the three components of velocity and vorticity to highlight some of the mechanisms by which the spray

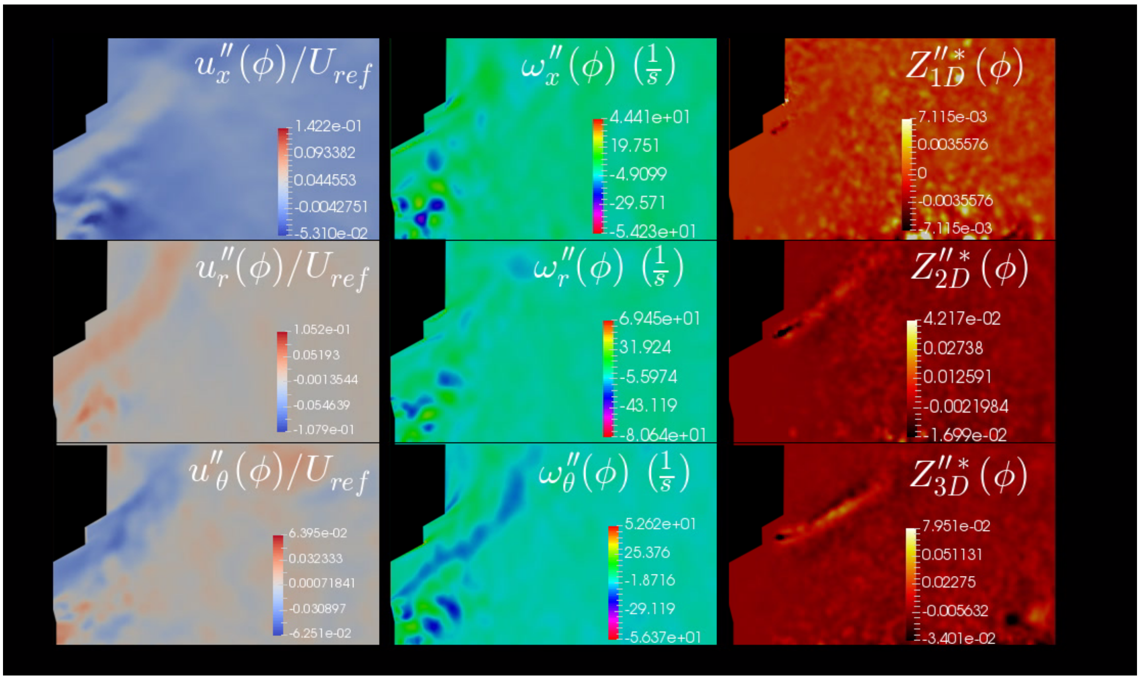

Figure 7.11: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 3\pi/8$.

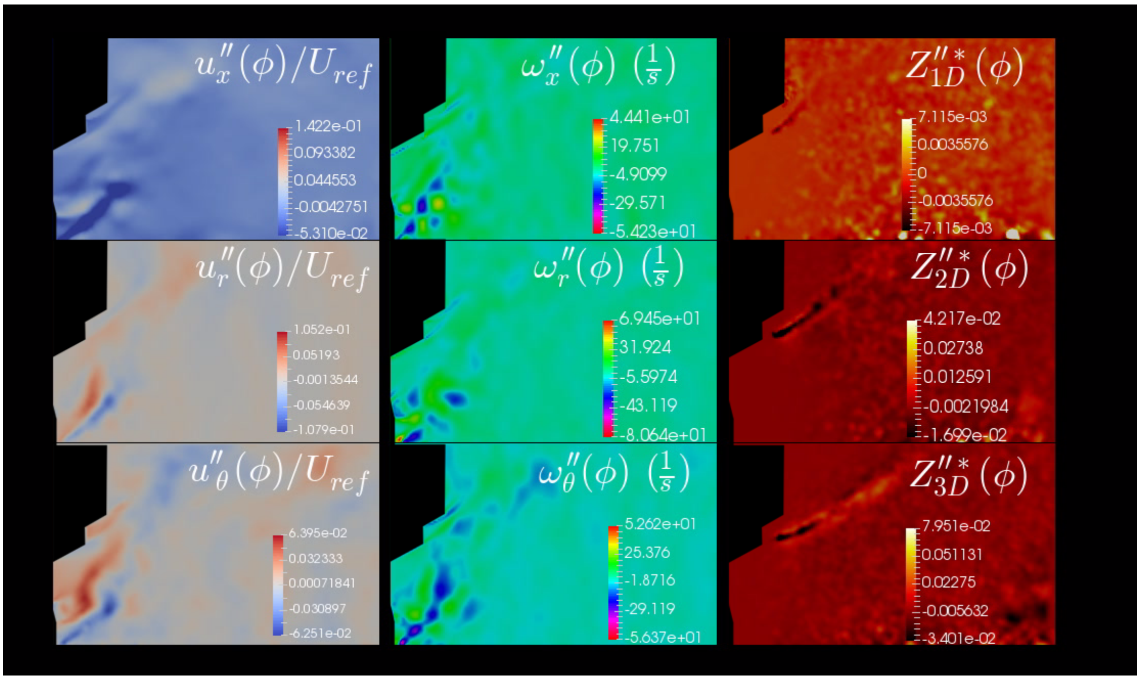

Figure 7.12: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 3\pi/4$.

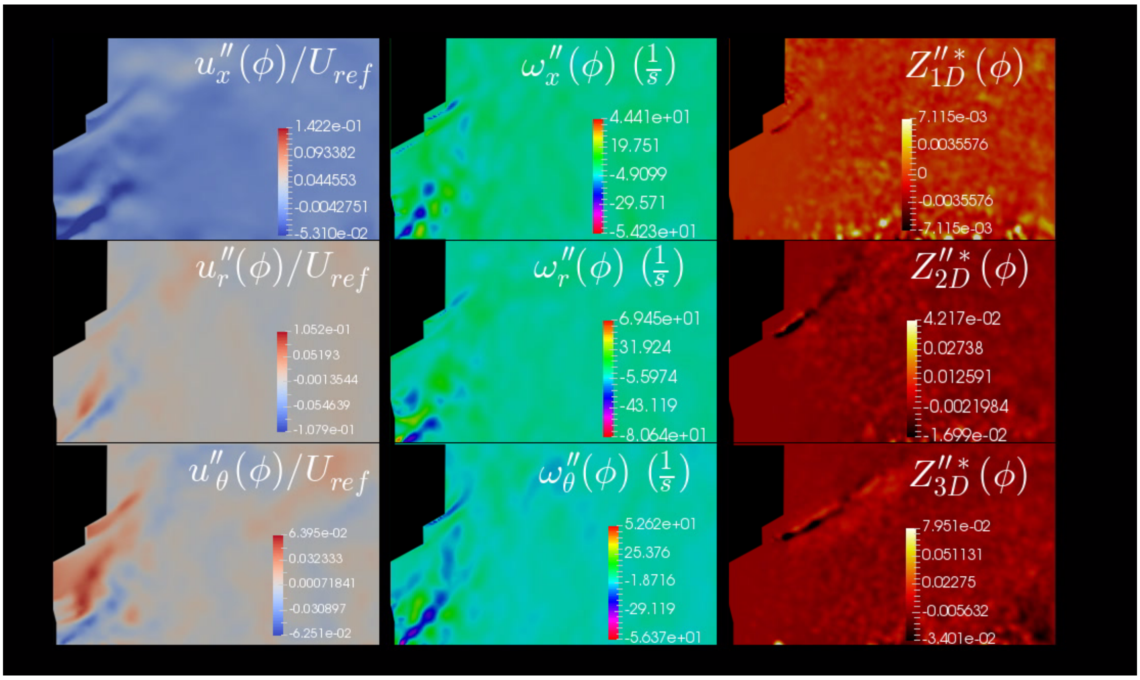

Figure 7.13: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 7\pi/8$.

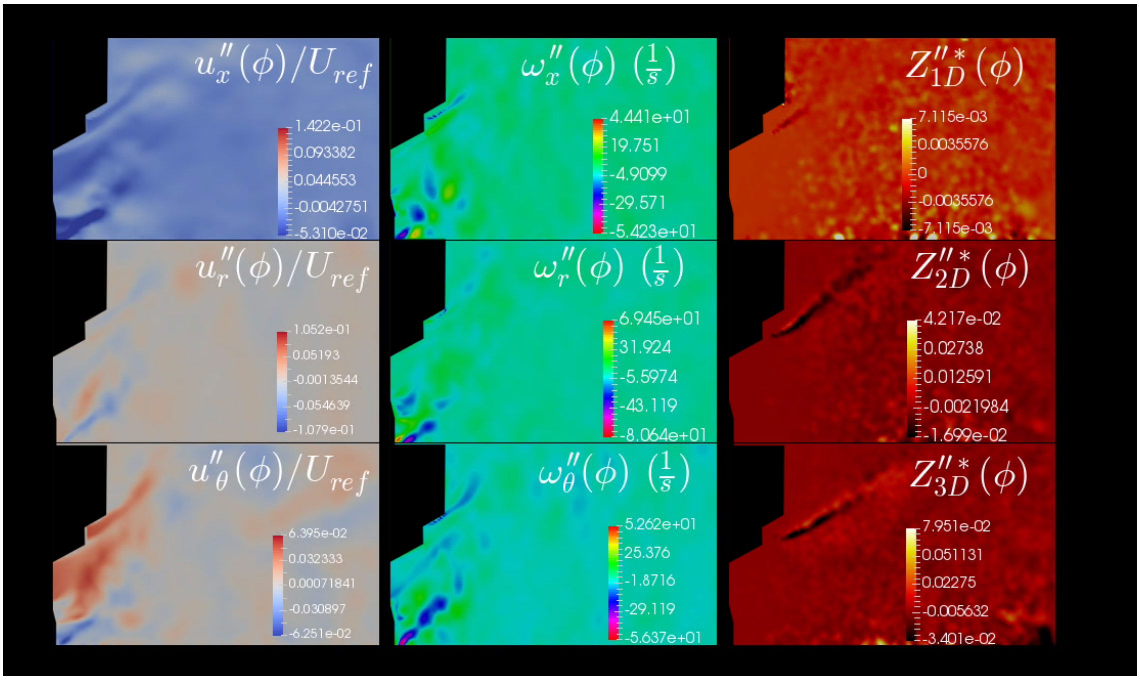

Figure 7.14: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = \pi$.

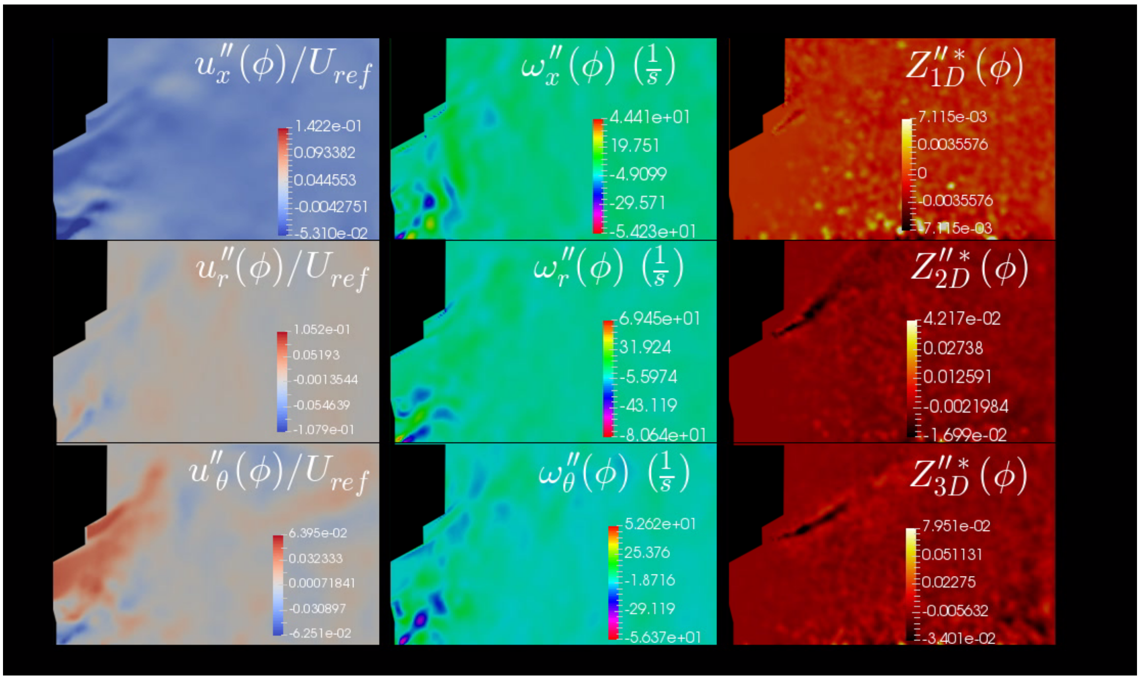

Figure 7.15: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 9\pi/8$.

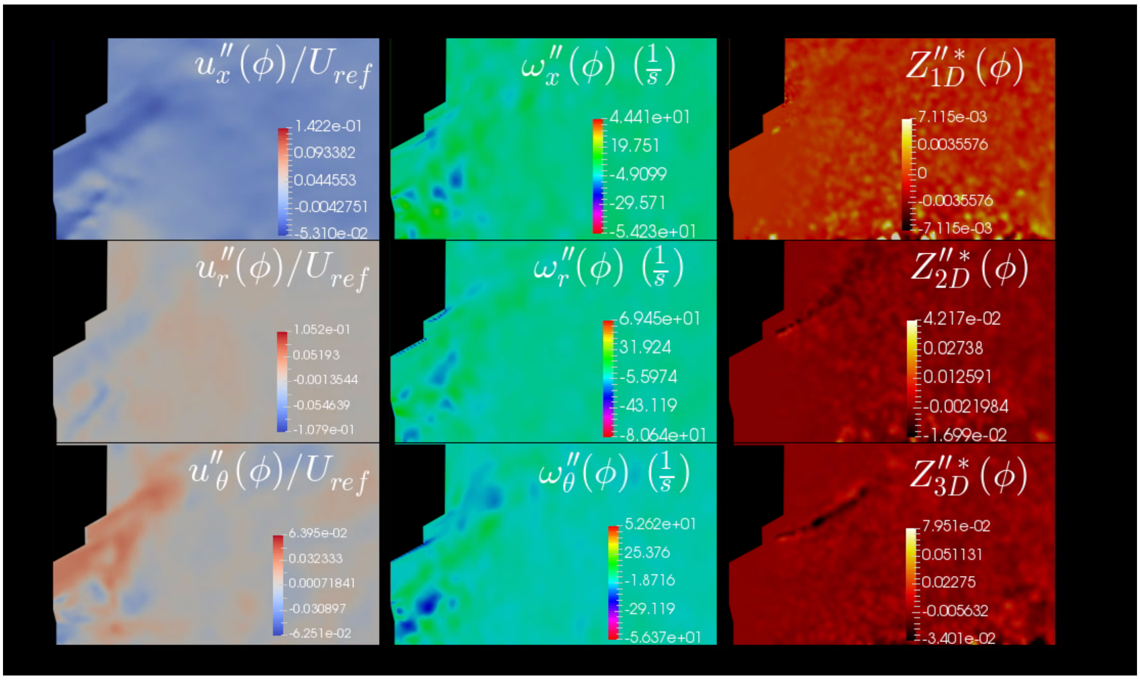

Figure 7.16: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 11\pi/8$.

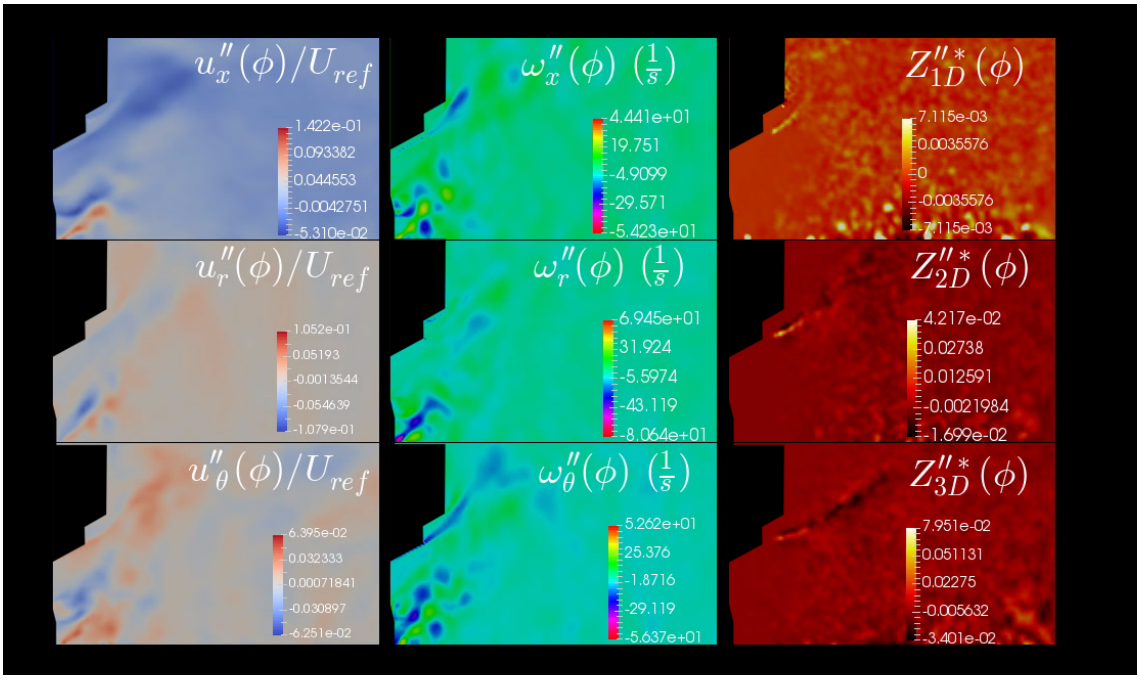

Figure 7.17: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 13\pi/8$.

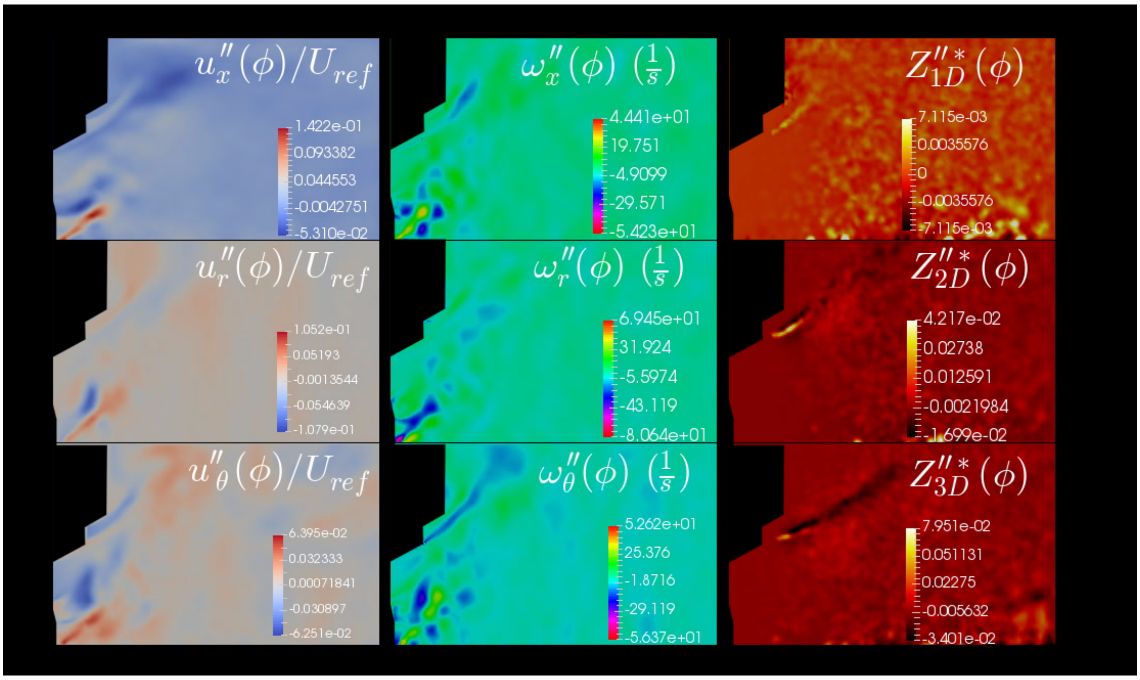

Figure 7.18: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 7\pi/4$.

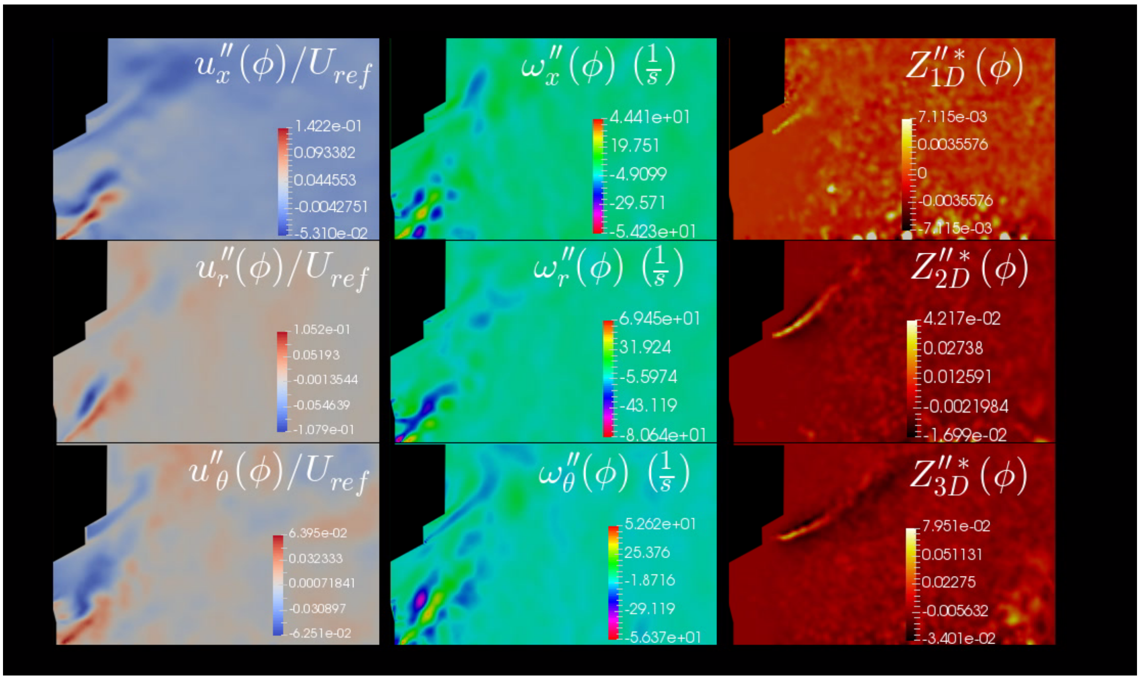

Figure 7.19: The phase averaged velocity, vorticity and mixture fraction of the three droplet groups at $\phi = 15\pi/8$.

mixture fraction is altered.

Figure 7.10 corresponds to the point in time where the velocity is maximum through the fuel injector and is designated the phase angle of $\phi = 0$. At this point in time the force experienced by the droplets pushing them downstream is close to a maximum, accelerating them away from the nozzle which by $\phi = 3\pi/8$ (Figure 7.11) leads to the emergence of a deficit in the spray mixture fraction close to the injection point. The minimum velocity through the injector appears close to $\phi = 9\pi/8$ (Figure 10). This velocity deficit tends to bunch the droplets up, increasing the mixture fraction close to the atomiser. Unlike the case when the droplets are spreading out this bunching up process takes different amounts of time depending on the droplet diameter. The smallest droplets react first leading to a growing mixture fraction at $\phi = 11\pi/8$ (Figure 7.16), the mid size droplet then show a growing mixture fraction at $\phi = 13\pi/8$ (Figure 7.17) and the largest droplets at $\phi = 7\pi/4$ (Figure 7.18). The maximum mixture fraction for all sizes in this region appears to be around $\phi = 15\pi/8$ (Figure 7.19).

This peak-to-peak variation in spray mixture fraction corresponds to around a 1% change for the smallest droplets, 4% change for the mid size droplets and 8% for the largest droplets. This is because the larger droplets contain so much more fuel than the smaller ones. The effect of the spray dispersion on fluctuations of mixture fraction are much larger for these droplets. It is therefore preferential to reduce the prevalence of these larger droplets as their sensitivity to large scale flow structures contributes much more to variations in mixture fraction. The cone angle of the main jet changes over the acoustic cycle which also causes the spray to move radially in and out, especially the largest droplets.

Apart from these primary spray dispersion processes it can also be seen occasionally that the spray sheet breaks up and disperses in a direction normal to the direction of the mean flow. This can be seen in Figure 7.12 where a bright, high mixture fraction halo forms around the otherwise dark, low mixture fraction located close to the atomiser. In figures 7.13 and 7.14 this halo can be seen to disperse outwards. The mechanism for such a dispersion can be identified by looking at the axial and azimuthal vorticity fields in Figure 7.16 and Figure 7.17 close to the droplet injection location. At these points in time there exists a strong velocity gradient between the velocities in passages B and C. This is because the mass flow and axial velocity variations are felt in each passage at almost the same time. Whilst the maximum flow rate is experienced at the same time in each passage, there exists a different distance between the exits of passage B and C and the droplet injection plane. This change in distance (primarily axial location) causes a phase shift between the maximum velocity of one passage to the next at the injection location because of the differing convection times of each passage which therefore generates a high quantity of vorticity.

These intermittent vortices may cause the spray to preferentially disperse at these points in the cycle leading to a shorter evaporation time or the fuel mixture fraction field to change through the cycle. This would therefore alter the location and intensity of the flame, leading to a variation in heat release rate. It would be relatively simple to design a new injector that eliminates this axial variation in the exit planes of passage B and C and test how the dispersion of the fuel spray or the heat release rate was altered, most likely damped. It is worth noting that, providing the passage length and location of the swirl vanes remains constant, the passage response will not change with this geometry modification. Because of this, the new geometry could be performed using the PODFS method based on the original flow simulation, significantly reducing the computational effort.

7.6 Conclusion

The effects of large scale flow structures and acoustic forcing on the convection of the fuel spray have been tested in this chapter using an LES simulation forced using the PODFS technique. The results show that mid size droplets are most responsive to large scale structures and acoustic forcing, consistent with previous work relating droplet response to the Stokes number of the droplet-flow feature combination. Droplets significantly larger or significantly smaller only show around 5% less sensitivity, therefore, changing the droplet size through alteration of the atomisation strategy seems to be relatively inefficient way of changing the flame response. Changing the atomiser location to help reduce the interaction of the spray with large scale structures may prove more useful. Phase and azimuthal averages of the spray mixture fraction for droplets of different sizes show two main mechanisms for alteration of the spray. The first is that the fluctuating axial velocity tends to bunch up the spray during the low velocity part of the cycle and spread the spray out during the high velocity part of the cycle. There is evidence to suggest that the largest droplets take more time to accelerate however their larger size causes them to have the most impact on the flow mixture fraction. The second mechanism is caused by a convective time delay between the exit of the individual flow passages and the point of injection. Passages B and C respond more or less in phase to the acoustic wave, however the different distance between each flow exit passage and atomiser means than the velocity fluctuations will be out of phase at the injection location. This causes a high velocity gradient that leads to the formation of intermittent vorticity close to the atomiser. This vorticity then causes the spray to break up much more quickly than at other parts of the cycle. It is suggested that moving the axial location of the exits of these two passages such that they are in line would reduce the prevalence of this mechanism.

Chapter 8

Combustion

This chapter is the culmination of all the work presented so far. The flame transfer function is derived using two different methodologies, a brute force compressible method, where the simulation aims to compute the acoustic and flow field simultaneously; and an incompressible method that takes acoustic information from a donor compressible simulation to represent the effects of forcing on a truncated domain. Because computation of the flow field allows for greater and simultaneous resolution of important flow variables with respect to experiments, the flow fields from these FTF computations have also been analysed to determine the driving mechanisms of the flame response.

As the fuel is injected and evaporates, the flow cools down and becomes more dense, while after combustion, the density field reduces significantly due to the greatly increased temperature. These density variations may cause large changes to the flow field which have not been considered up until this chapter.

Finally, a simulation has been run to investigate the effects of the periodic variation of fuel droplet SMD that may be caused by acoustic forcing. An attempt has been made to estimate the relative importance of this mechanism to the flame transfer function. The effects of different fuel spray injection velocities has also be considered separately.

The simulations have been completed for a case where the mean pressure drop was set to 3.8%, the air-to-fuel ratio (AFR) is 34 and the main is fed with 80% of the fuel.

8.1 Numerical methodology

8.1.1 Flamelet Generated Manifolds (FGM)

In addition to the equations governing mass (A.22) and momentum (A.19) combusting flows contain a great number of chemical species and a transport equation must be solved for each of them [157]:

$$\frac{\partial \rho Y_i}{\partial t} + \nabla \cdot (\rho \mathbf{u} Y_i) - \nabla \cdot \left(\frac{1}{Le_i} \frac{\lambda}{c_p} \nabla Y_i \right) = \omega_i^+ - \omega_i^- \tag{8.1}$$

where Y_i is the mass fraction of the ith species, $Le = \lambda / \rho D_{im} c_p$ is the Lewis number, the ratio of viscous to mass diffusion. D_{im} is the diffusion coefficient, λ is the thermal conductivity, c_p is the specific heat, ω^+ is the chemical production source term and ω^- is the chemical consumption source term. The enthalpy equation must now also consider the effect of the diffusion of these species. It is written here without the additional source terms due to pressure and turbulent heating [157]:

$$\frac{\partial \rho h}{\partial t} + \nabla \cdot (\rho \mathbf{u} h) + \nabla \cdot \left(\frac{\lambda}{c_p} \nabla h \right) = \nabla \cdot \left(\sum_{i=1}^{N} h_i \left(\frac{1}{Le_i} - 1 \right) \frac{\lambda}{c_p} \nabla Y_i \right) \text{ for } i = 1, ..., N \tag{8.2}$$

where N is the number of chemical species. The ideal gas law is re-written as:

$$p = \rho R T \sum_{i=0}^{N} Y_i / M_i \tag{8.3}$$

where R is the universal gas constant and M_i is the molar mass of the ith species. The specific enthalpy can be calculated from:

$$h = \sum_{i=1}^{N} Y_i h_i \tag{8.4}$$

with:

$$h_i(T) = h_i^0 + \int_{T_0}^{T} c_{p,i}(T') dT' \tag{8.5}$$

where h_i^0 is the enthalpy of formation of species i. The value of the chemical species source terms are calculated from a series of equations describing the possible reactions between the species present in the combustion chamber. The reactions can have characteristic time-scales that are orders of magnitude smaller than the flow. Computing the flow variables with a small enough time step to accurately capture all of these reactions would therefore be extremely expensive.

An additional problem occurs when considering the number of species and therefore reactions that must be used. A simple fuel such as methane has only a relatively small number of reaction mechanisms however kerosene may have thousands of species and tens of thousands of reactions. One option would be to simply remove many of these reactions and species or combine them into a set of pseudo-species that give the correct integral behaviour. Approaches such as this can usually be tuned to give the desired results at particular operating conditions or to focus on particular outcomes such as flame speed, flame temperature or mass fraction of a particular product but will fail to give the correct values at other conditions [12, 29, 31, 80].

Other techniques focus on the observation that as some of the reactions occur on comparatively small time scales, the balance between these species is essentially steady-state in the frame of the flow simulation. This allows for a steady state decoupling to occur between the flow and these species. In this case, the left had side of Equation 8.1 can be neglected leading to:

$$\omega_i^+ - \omega_i^- = 0 \text{ for } i = 1, ..., M \tag{8.6}$$

where M is the number of steady-state species. These equations can therefore be solved off-line as a function of the other species mass fractions. The solution of these equations can be stored in a look-up table of dimension $N - M$ as a manifold. This approximation works well except in the colder regions of the flow where these steady-state reactions may occur more slowly and where convection and diffusion of these species becomes important [157]. The FGM method [157] aims to improve the accuracy of these manifold methods through computation of the manifold using a one-dimensional flame where the diffusion and convection effects are modelled directly. Although the three-dimensional flame has a complicated topology, the reaction zone of the flame is extremely thin. When the flame is viewed locally, it appears to be very similar to a one-dimensional flame with a radius of curvature. The equation for the ith species in a one-dimensional flame is given as [157]:

$$\dot{m}\frac{\partial Y_i}{\partial n} = \frac{\partial}{\partial n}\left(\frac{1}{Le_i}\frac{\lambda}{c_p}\frac{\partial Y_i}{\partial n}\right) = \omega_i^+ - \omega_i^- + P_i \text{ for } i = 1, ..., N \tag{8.7}$$

where n is the direction normal to the flame and $\dot{m}$ is taken to be a constant mass flow rate. The non-one-dimensional effects on the flame are collected into the perturbation term $P_i(n,t)$. This term includes the effects of flame stretch, curvature effects and additional transport terms due to the flame-normal direction being inconsistent across different species. In the FGM method, at this stage of the analysis, these effects are assumed negligible and $P_i = 0$. In a pre-processing step this set of equations is then solved over a number of premixed flames that assume varying mixture fractions of fuel and air. This was done using the one-dimensional detailed chemistry solver CHEM1D [49] assuming that the fuel is dodecane [97]. The mixture fraction is defined as [121]:

$$Z = \frac{sY_f - Y_o + Y_o^0}{sY_f^0 + Y_o^0} \tag{8.8}$$

where s is the stoichiometric mass ratio, Y_f is the mass fraction of fuel, Y_o is the mass fraction of oxidiser, Y_f^0 is the fuel mass fraction in pure fuel and Y_o^0 is the oxidiser mass fraction in pure oxidiser.

As the reaction progresses, more and more reactants are consumed and replaced by combustion products, this process is described using a progress variable [121].

$$Y = Y_{con.} \tag{8.9}$$

where $Y_{con.}$ is the specific molar mass fraction of certain reaction products. The progress variable is equal to zero in an unburned mixture and equal to the mass fraction of certain reaction products in a fully burned one. All of the variables of interest such as the reaction rate, temperature, enthalpy and species mass fractions are linked via a table to the corresponding values of Z and Y. These two variables then become control variables that can be solved in the flow simulation and then every other variable of interest in the flame can be found in the corresponding table entry. The transport equations for the LES filtered control variables ($\tilde{Z}$ and $\tilde{Y}$) that must be solved along with the other flow variables in the simulation are [121]:

$$\frac{\partial \bar{\rho}\widetilde{Z}}{\partial t} + \frac{\partial \bar{\rho}\widetilde{u}_i\widetilde{Z}}{\partial x_i} = \frac{\partial}{\partial x_i}\left[\left(\frac{\mu + \mu_t}{\sigma_Z}\right)\frac{\partial \widetilde{Z}}{\partial x_i}\right] + \widetilde{\dot{\omega}_Z} \tag{8.10}$$

and:

$$\frac{\partial \bar{\rho}\widetilde{Y}}{\partial t} + \frac{\partial \bar{\rho}\widetilde{u}_i\widetilde{Y}}{\partial x_i} = \frac{\partial}{\partial x_i}\left[\left(\frac{\mu + \mu_t}{\sigma_Y}\right)\frac{\partial \widetilde{Y}}{\partial x_i}\right] + \widetilde{\dot{\omega}_Y} \tag{8.11}$$

where σ_Y and σ_Z are the Schmidt numbers for the progress variable and the mixture fraction respectively. The source term for the mixture fraction ($\widetilde{\dot{\omega}_Z}$) is equal to the amount of fuel spray that has been evaporated in the last timestep as calculated by the Lagrangian solver. The source term for the progress variable ($\widetilde{\dot{\omega}_Y}$) also comes from the FGM table however using the raw value calculated would be assuming that the flame is laminar. The highly turbulent flow increases the rate of mixing and combustion significantly in the combustion chamber. In order to take this into account the source term is modified by using the presumed probability density function (PPDF) method.

8.1.2 The presumed PDF method

The PPDF (presumed probability density function) method assumes that the relationship between the flame, control variables and their variances can be represented by simple probability density functions. This allows an extension of the FGM table to include the effects of turbulent fluctuations on the flame. The table is extended to include the variances of two control variables by computing:

$$\begin{gathered}\widetilde{\dot{\omega}_Y}(Y, Z, Y'^2, Z'^2) = \\ \int_0^1\int_0^1 P(Y, Z, Y'^2, Z'^2)\dot{\omega}_Y^{\dagger}(Y, Z)dYdZ\end{gathered} \tag{8.12}$$

where $\dot{\omega}_Y^{\dagger}(Y, Z)$ is the original reaction rate from the flamelet computation. This integration is performed before the computation of the flow and it is aided by the simplification of the PPDF to [56]:

$$P(Y, Z, Y'^2, Z'^2) = P_Y(Y, Y'^2)P_Z(Z, Z'^2) \tag{8.13}$$

where $P_Y(Y, Y'^2)$ is assumed to be a 3-delta PDF (at $Y = 0$, $Y = Y_{mean}$ and $Y = Y_{max}$) and $P_Z(Z, Z'^2)$ is assumed to be a β-PDF [4]. The transport equations for the variances of the mixture fraction and the scaled progress variable are [114]:

$$\begin{aligned}\frac{\partial \bar{\rho}\widetilde{Z'^2}}{\partial t} + \frac{\partial \bar{\rho}\widetilde{u}_i\widetilde{Z'^2}}{\partial x_i} = \frac{\partial}{\partial x_i}\left[\left(\frac{\mu + \mu_t}{\sigma_Z}\right)\frac{\partial \widetilde{Z'^2}}{\partial x_i}\right] + C_1\frac{\mu_t}{\sigma_Z}\left(\frac{\partial \widetilde{Z}}{\partial x_i}\right)^2 \\ -C_2\frac{\mu_t\widetilde{Z'^2}}{\sigma_Z\Delta^2}\end{aligned} \tag{8.14}$$

and,

$$\begin{aligned}\frac{\partial \bar{\rho}\widetilde{Y'^2}}{\partial t} + \frac{\partial \bar{\rho}\widetilde{u}_i\widetilde{Y'^2}}{\partial x_i} = \frac{\partial}{\partial x_i}\left[\left(\frac{\mu + \mu_t}{\sigma_Y}\right)\frac{\partial \widetilde{Y'^2}}{\partial x_i}\right] + C_3\frac{\mu_t}{\sigma_Y}\left(\frac{\partial \widetilde{Y}}{\partial x_i}\right)^2 \\ +C_4\widetilde{Y'^2\dot{\omega}_Y} - C_5\frac{\mu_t\widetilde{Y'^2}}{\sigma_Y\Delta^2}\end{aligned} \tag{8.15}$$

where C_1, C_2, C_3, C_4 and C_5 are model constants and the progress variable variance source term $\widetilde{Y'^2\dot{\omega}_Y}$ is derived in a similar way to the progress variable source term. Δ is the mesh spacing.

It is assumed that this PPDF approach takes account of effects related to the mean and oscillation of the strain rate. Furthermore the FGM method has been used to successfully model self-excited combustion dynamics in combustors which suggests that it provides a sufficient estimation of the fluctuating heat release rate [79].

8.1.3 Temperature

The combustion of the reactants is assumed to be adiabatic and therefore the total enthalpy of the mixture does not change. However, the enthalpy of the flow may differ locally from the enthalpy assumed during the computation of the one dimensional flames due to local changes in the pressure field. This local change in enthalpy can be taken into account by including the flow specific enthalpy as an additional control variable however this adds an extra dimension to the table and requires many more one dimensional flame calculations to be made at varying levels of flow enthalpy. Instead the approach of Vicquelin [158] is used: The temperature of the flow is corrected from the FGM value by making a linear interpolation based on the enthalpy of the flame and the enthalpy of the flow:

$$\tilde{T} = T_{FGM} + \frac{\tilde{h}(T) - h(T_{FGM})}{c_p(T_{FGM})} \tag{8.16}$$

where c_p, T_{FGM} and $h(T_{FGM})$ are read from the FGM table. Enthalpy must now also be solved as a flow variable, even in the case of incompressible flow.

8.1.4 Low resolution compressible LES

The low resolution mesh described in Section 5.5 was used both to generate upstream boundary conditions for incompressible calculations and also for some preliminary comparisons between the reacting and non-reacting flow fields and preliminary calculations of the flame transfer function. The non-reacting numerical methodology is detailed in Section 5.5, two sets of simulations were run: One set of simulations where the flow was acoustically forced at $St = 0.14$ and the other where the simulation was simultaneously forced at all of the frequencies tested in the experiment at once. Each set consisted of a simulation where the flow was forced from the downstream boundary and a simulation forced from the upstream boundary.

The reacting simulations were run mostly to test the methodology for resolving the flame transfer function. Fuel droplets were injected at 100 equidistant points around a single ring just behind the mains pre-filmer and 50 equidistant points behind the pilot swirler. This is the same strategy as used in Chapter 7 except that the radius and azimuthal location of injection remain unchanged for each droplet at each timestep, this was done to remain consistent with best practice used in Rolls-Royce. Each droplet was given a zero swirl angle velocity but the axial and radial components of velocity were allowed to fluctuate by ±60% and ±80% respectively. The mean absolute droplet velocity was set to $0.25U_{ref}$ with a small radial component. The SMD was chosen to be $2.2 \times 10^{-4}D$ consistent with the work of Su et al. [148]. Each droplet diameter was chosen from a Rosin-Rammler distribution assuming ten size groups. Evaporation was turned on according to the work of Lefebvre and Chin [72] as explained in Chapter 7. The spray boundary conditions remain a large source of error in the analysis due to insufficient experimental data available at engine representative temperatures and pressures.

In order to obtain the mean reacting flow field, a URANS simulation was run assuming pseudo-steady flow with a timestep of $2 \times 10^{-1} t^*$ using the $k - \omega$-SST model. The flow was initially run with combustion and spray turned off for around 4000 steps before the fuel spray was turned on. After a run of an additional 500 steps the flow was ignited by heating a small region of the flow on the outer edge of the pilot jet to 3000K. After a further 6000 steps the heat release rate had stabilised and the simulation could be continued as an LES. This large timestep requires the non-reflective boundary conditions to be turned off, but also helps to diffuse any acoustic waves. Each LES computation was then run for 100,000 timesteps with a timestep of $2 \times 10^{-3} t^*$ plus whatever time was needed to overcome any transients, this was in the region of 10,000 to 100,000 additional steps. Alternative start-up procedures were attempted such as assuming steady flow, however no start-up procedure was more effective in reducing the required number of steps to remove these transients. The large number of transient steps is linked to the physical length of the domain as it takes a very long physical time for waves to disperse to the boundaries. The simulations were completed using PRECISE-UNS with first order upwind spatial discretisation to maintain stability and second order temporal discretisation. The constant Smagorinsky model [140] was used for turbulence closure.

8.1.5 Medium resolution compressible LES

It was identified that the low resolution mesh was insufficient to resolve the flame. Furthermore, the small cooling hole features for the heatshield required extremely fine cells which resulted in very high CFL numbers in this region of the mesh. These high CFL number regions prevented the use of second order spatial differencing as the flow solution would become unstable in this case. To solve both of these problems the geometry was altered such that the cooling hole inlets were replaced by a new effusion boundary condition that acted to remove mass from this surface. The heatshield outlet was also modelled as an effusion boundary condition with the opposite mass flow rate. The mass flow rate was set to that found in the non-reacting simulations using the low resolution mesh. The mesh resolution was doubled in the region downstream of the fuel injector with respect to the low resolution mesh. The other meshing parameters remained the same resulting in a mesh of around 15 million cells. Improvements in the non-reflective boundary condition (see Section 5.3.3 removed the need for a large upstream sponge region and so the inlet was shortened.

Due to the removal of the small mesh features the simulation stability was greatly increased and as such a second order TVD scheme could be used for spatial differencing. All other simulation parameters remained the same. The mesh is visualised in Figure 8.1.

8.1.6 High resolution incompressible LES

The high resolution incompressible LES simulation was run to capitalise on the development of the PODFS method. If the injector outlet flow can be captured accurately by the method and

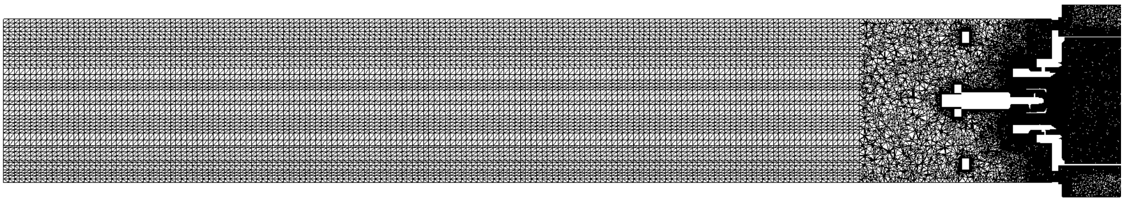

(a) Upstream section.

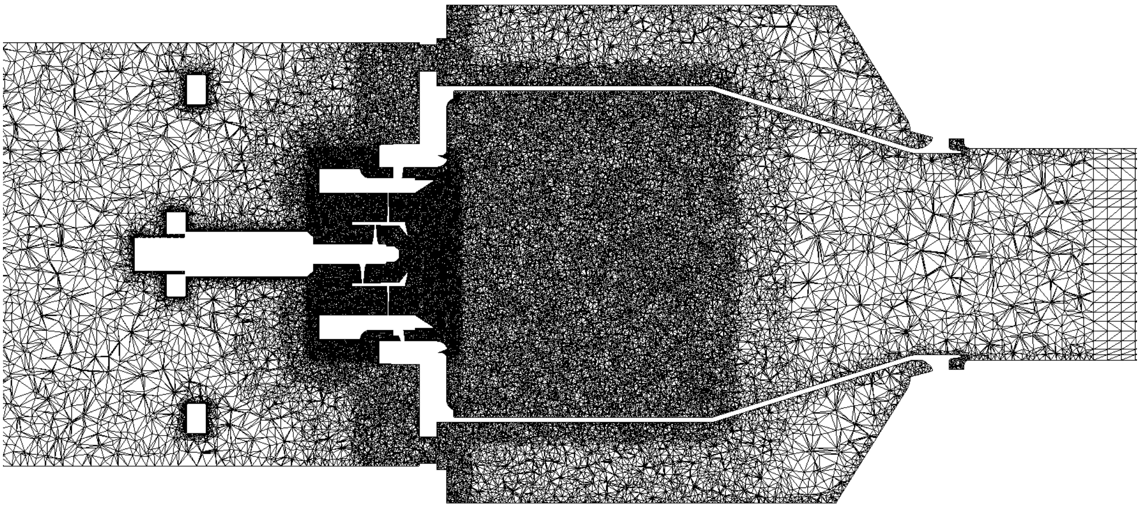

(b) Combustion chamber.

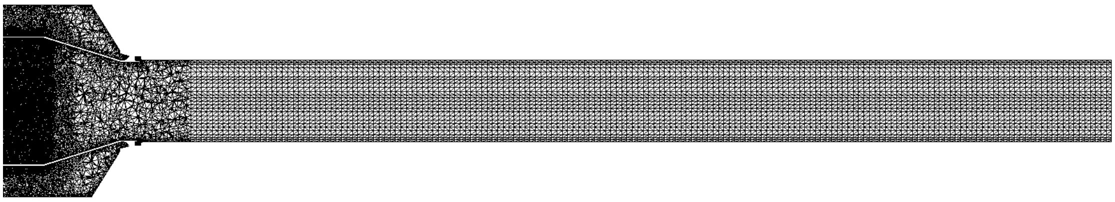

(c) Downstream section.

Figure 8.1: The medium density mesh used for LES FTF calculations.

the flame can be decoupled from the acoustic field, then an incompressible simulation of the flow downstream of the injector should recover the same result as the compressible simulation.

Because the injector passages are not included in the mesh, the number of mesh cells and the complexity the mesh drop dramatically. The reduction of cells required through removal of the injector geometry can therefore be used to boost the resolution of the flame. In this case the resolution was increased by a factor of two downstream of the injector with respect to the medium resolution mesh. The simplification of the geometry allows for a greater number of cells to be converted from tetragonal cells to hexagonal cells, further reducing the cell count and increasing the accuracy of the simulation. Whilst PRECISE-UNS is an unstructured solver it is worth noting that the geometry is now simplified to a level where a structured solver and mesh could be employed, further increasing the computational efficiency.

Changing from a compressible to an incompressible solver also removes the restrictions imposed by the need to resolve the acoustic waves at the boundaries and therefore the mesh size can be significantly increased far downstream of the combustor. Furthermore, the lack of non-reflective boundary conditions removes the restriction of CFL number at the boundaries, which means that provided the turbulent scales are resolved sufficiently well, the timestep may be increased.

A total of five simulations were run using this mesh: firstly, a baseline case was run where the inlet mass flow rates and spray SMD were held constant in time. Then the SMD of the fuel droplets was allowed to vary by 20% with respect to their mean value at a frequency of $St = 0.14$. Then the SMD was returned to the mean value and the inlet velocity was forced at the same frequency using the PODFS method built from data acquired during the non-reacting, downstream forced, compressible simulation. Finally, the inlet fluctuation frequency was increased to $St = 0.30$.

In each of these simulations the pilot jet split radially outwards and a pilot recirculation zone (PRZ) was formed in the centre of the domain similarly to the flow seen in Chapter 7. This ‘flipped’ flow structure has been seen before in reacting tests of a similar injector design by Williams et al. [162] and whether or not this structure forms has been identified as being sensitive to the upstream boundary condition. Furthermore, in non-reacting tests of similar geometries including experiments by Spencer [142], this flow structure has been induced by reducing the mass flow rate through passage A. Following the medium and low resolution compressible reacting simulations, it was apparent that this pilot flow structure is unlikely to be seen in the experiment or real engine conditions, however installation effects or violent changes in the combustor pressure or temperature could force the flow to this state. In order to recreate the pilot penetration seen in the compressible simulations, two strategies were pursued simultaneously. The first was to boost the quantity of upstream turbulence by adding additional modes to the PODFS. As the original PODFS was constructed from a URANS simulation, it is expected that the data does not contain sufficient high wavenumber structures to promote the production of turbulence in the pilot jet. This reduction in kinetic energy may therefore be responsible for the splitting of the pilot. In order to combat this, an additional PODFS was generated from data produced during the high resolution incompressible LES simulation of the Loughborough acoustic duct described in Section 3.2.8. This PODFS was constructed from the first 30 modes adding to the original 20 mode PODFS resulting in 50 modes per inlet. Because the geometry is different, the dimensions of the PODFS had to be modified and the values of the velocity rescaled.

The other strategy was to artificially boost the mass flow rate of the pilot jet by setting the initial mass flow rate to be 50% higher than normal. This resulted in a robust pilot jet that remained stable throughout the start-up procedure. Once the reaction rate had stabilised, the mean flow rate through the passages was then recalibrated to reflect what had been observed in the non-reacting compressible simulation by adjusting the value of the mean velocity scaling factor, α, appropriately. Subsequent tests showed that without the addition of the second, turbulent PODFS, the pilot flow structure would revert to the 'flipped' configuration despite the increased pilot mass flow. This shows that the additional PODFS model of the turbulent fluctuations can be used to correct the pilot flow penetration length.

8.2 The prediction of mean flow field

8.2.1 Comparison of reacting and non-reacting flow fields

The following section compares some of the mean field variables between the non-reacting compressible URANS simulations using the low resolution mesh and the reacting compressible LES using the medium resolution mesh. This specific comparison has been chosen because compressible URANS simulations have been shown to adequately represent the acoustic response of the injector without resorting to LES simulations that would likely require a finer mesh and potentially a smaller time step. It is therefore interesting to see what flow field differences occur between these URANS simulations and fully-reacting LES, which is chosen as the preferred numerical tool for combustion simulations. The mean velocity fields (Figure 8.2) are extremely similar upstream of the flame.

The main difference is the extension of the main and pilot jets. This is caused by the large increase in temperature across the flame that causes a corresponding drop in density that results in the flow accelerating downstream to maintain conservation of mass.

The fluctuations of velocity are altered in two fundamental ways (see Figure 8.3). The first is that the fluctuations downstream of the flame are greatly increased. The change in turbulence closure accounts for some but not all of this change, otherwise the same effects would be seen in the incompressible LES simulation in Section 3.2.8. The increased length of the pilot and main jet increases the region where the velocity gradients are sufficiently high such that turbulence is produced. The other noticeable change is that the high intensity fluctuations just downstream of the splitter plate between passages A and B have been markedly diminished. As discussed extensively in Chapters 4 and 5, this region is likely the wavemaker region where large helical structures begin as linearly evolving instabilities. The reduction of these large scale unstable structures implies that the flow has been partially stabilised with the introduction of combustion.

The flow stability is a function of the mean flow field [22], which has not changed in the wavemaker region. The only variable that can be responsible for this change is therefore the density. Figure 8.4 shows the density fields for both simulations where both fields are shown relative to the maximum density of the non-reacting simulation. In the case of the non-reacting simulation the density is almost constant except for the regions of highest velocity and the centre of the vortex coming from passage A due to the lower static pressure in these regions. In the case of reacting flow, the centre of the vortex has a density that is the highest in the flow because the

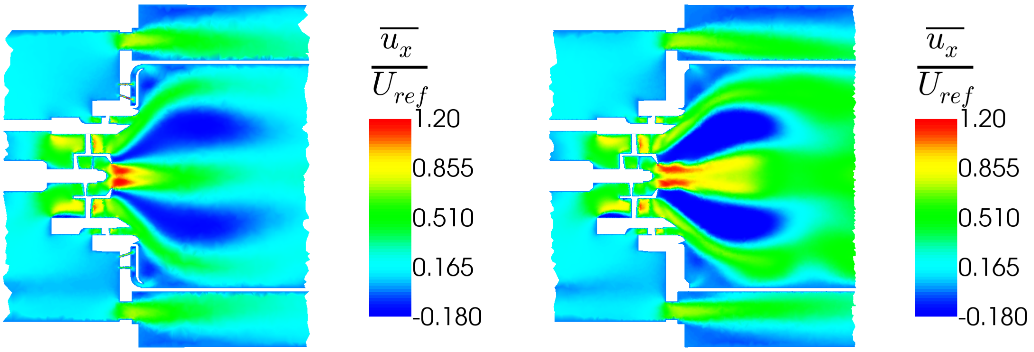

(a) Mean axial velocity for the non-reacting case.

(b) Mean axial velocity for the reacting case.

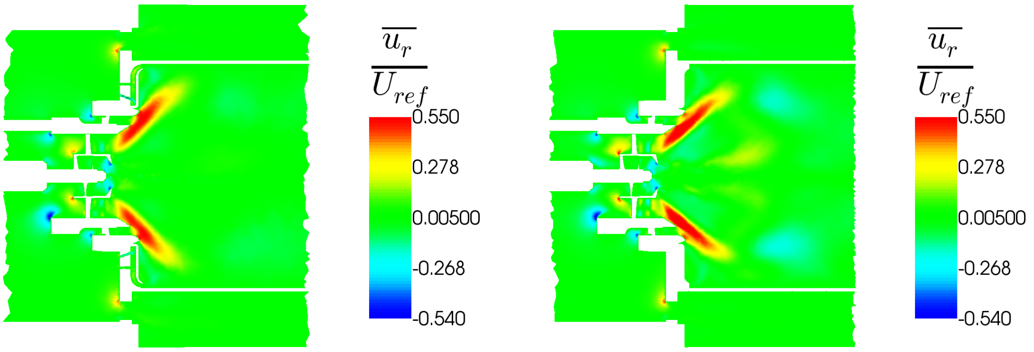

(c) Mean radial velocity for the non-reacting case.

(d) Mean radial velocity for the reacting case.

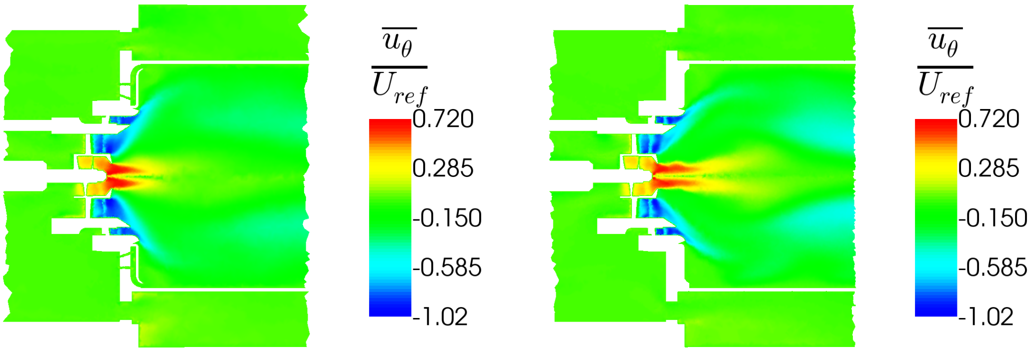

(e) Mean azimuthal velocity for the non-reacting case.

(f) Mean azimuthal velocity for the reacting case.

Figure 8.2: The effects of combustion on the mean flow field.

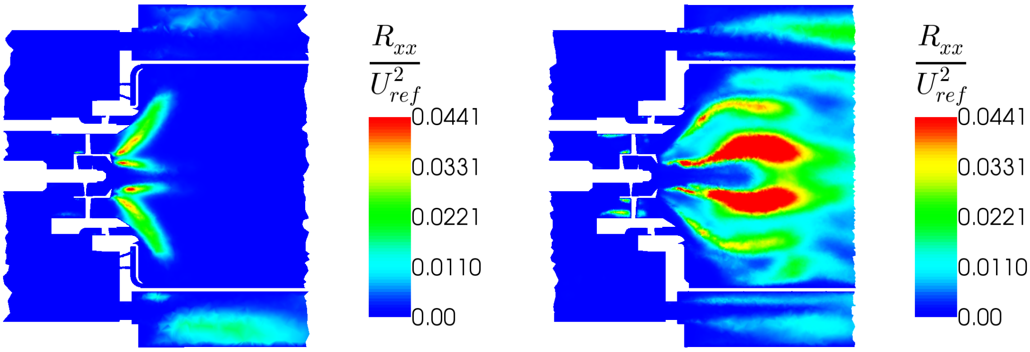

(a) R_{11} for the non-reacting case. (b) R_{11} for the reacting case.

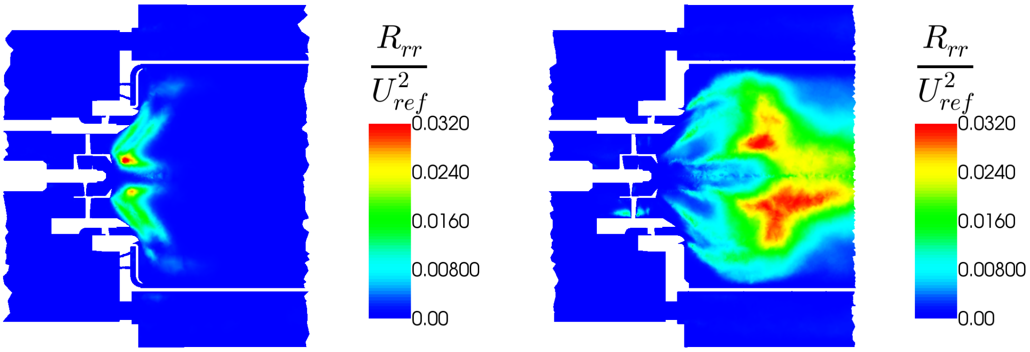

(c) R_{22} for the non-reacting case. (d) R_{22} for the reacting case.

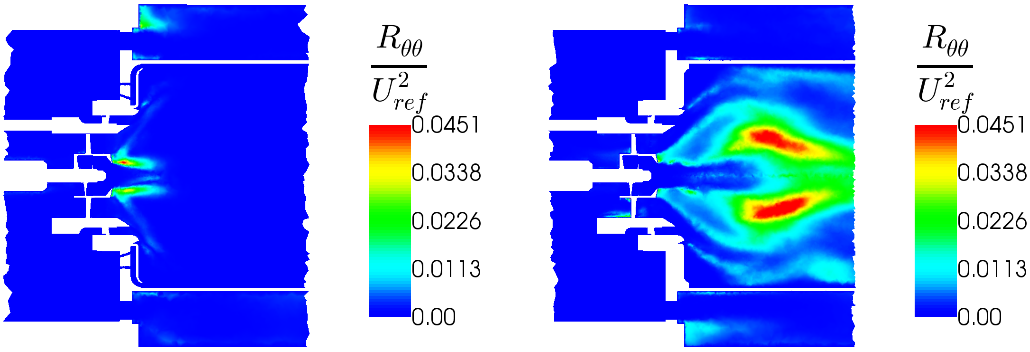

(e) R_{33} for the non-reacting case. (f) R_{33} for the reacting case.

Figure 8.3: The effects of combustion on the fluctuating velocity fields.

pilot atomiser is located in this region and the evaporation of the fuel causes the temperature to drop. In the region of the splitter plate, the hot combustion gasses are brought forward by the CRZ and therefore this region has a much lower density than in the case of the non-reacting flow.

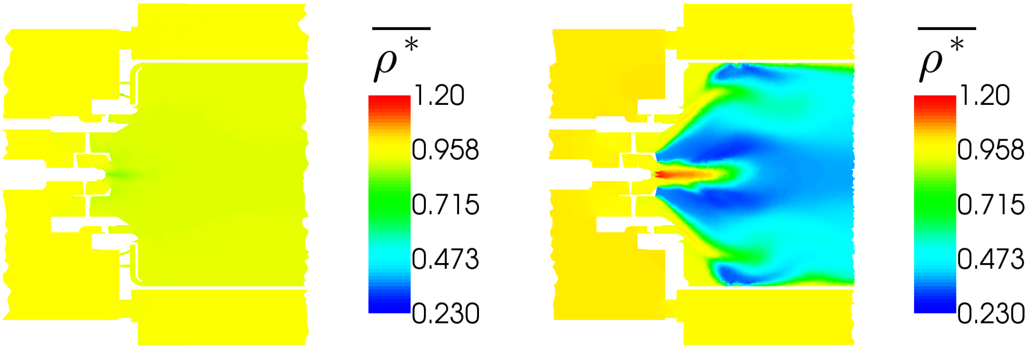

(a) The mean density field in the case of no combustion.

(b) The mean density field in the case with combustion.

Figure 8.4: The effects of combustion on the density field. In this case the density field in the non-reacting case has been normalised by the maximum density of the non-reacting case.

8.2.2 The effect of spray velocity

During the initial testing of the low resolution mesh, the velocity of the fuel droplets was altered to test the effects on the flame. Initially it was assumed that the newly atomised droplets would be travelling very close to the speed of the air passing the atomiser. Setting the droplet velocity to this level resulted in the fuel spray behaving as shown in Figure 8.5c. The main fuel droplets remain in the main jet and do not evaporate until they reach the end of the combustion chamber. This results in very few droplets being radially spun outwards into the ORZ as evidenced by the V-flame shape in Figure 8.5a. It would aslo likely result in a very low combustor efficiency as fuel would likely leave the combustion chamber before having time to react with the surrounding air. Experimental studies at ambient conditions suggest that the droplet initial velocity is much lower [148] and previous experience at Rolls-Royce Deutschland suggested a lower velocity of around $0.24U_{ref}$. This lower velocity gave a fuel spray as shown in Figure 8.5d which is much closer to what has been seen in experiments and what has been seen in Chapter 7. The problem is that there is little to justify this lower velocity except that for larger droplets, it is expected that they need sufficient time to accelerate. It is worth noting that only larger droplets will penetrate the flame zone and end up far downstream, as the smallest droplets will evaporate too quickly. These observed effects may be mesh sensitive and higher resolution meshes should be used to test the sensitivity of the flame to the droplet injection velocities.

These findings imply that in future it may be useful to choose the velocity of the fuel droplets based on a joint probability density function (PDF) between droplet velocity and droplet diameter. This would mean that small, lighter droplets would be injected at high velocity and larger droplets injected at lower velocities. The problem is that there is little experimental data at elevated temperatures and pressures available to build this PDF and numerical studies of droplet breakup are

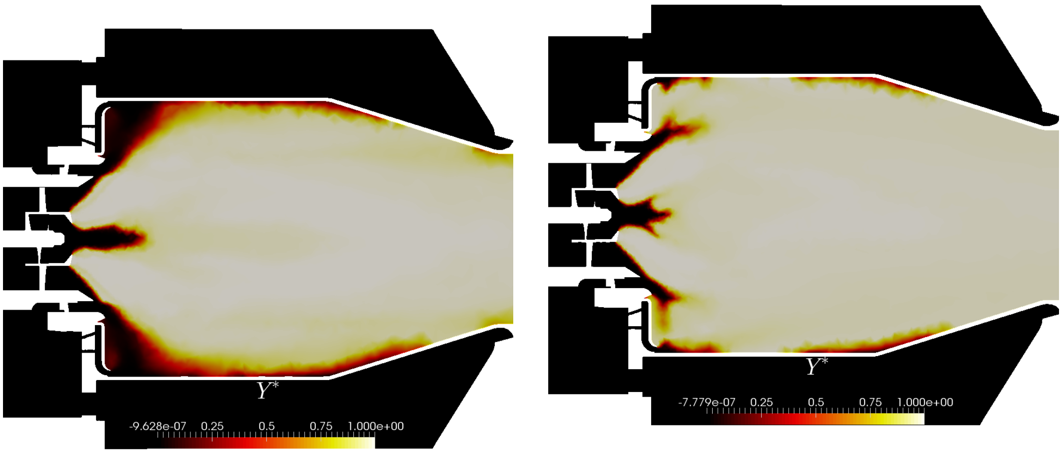

(a) Scaled progress variable in the case of high spray injection velocity.

(b) Scaled progress variable in the case of low spray injection velocity.

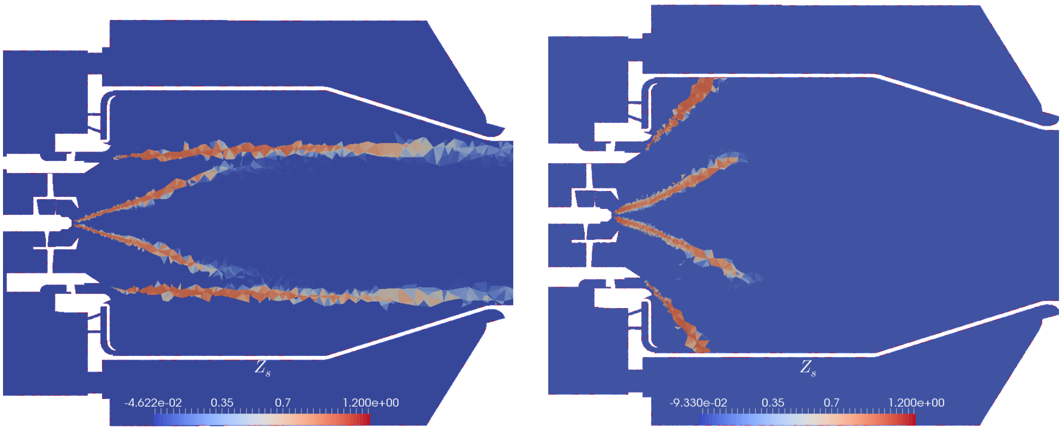

(c) Spray mixture fraction in the case of high spray injection velocity.

(d) Spray mixture fraction variable in the case of low spray injection velocity.

Figure 8.5: The effects of spray velocity on flame shape.

extremely expensive [32]. The SCARLET rig has optical access and therefore allows for measurements to be conducted of the spray and flame during operation. A PDA system could be used to measure the fuel spray under realistic and acoustically forced conditions therefore providing more accurate boundary conditions for future simulations.

8.2.3 Compressible and incompressible mean reacting flow fields

Figure 8.6a shows the azimuthally averaged mean axial velocity for the incompressible and compressible cases aerodynamically forced at $St = 0.14$. Firstly it can be seen that the pilot is penetrating in both cases showing that the addition of turbulence to the inlets of the truncated compressible domain has prevented the pilot from flipping outwards. The pilot jet axial velocity is clearly higher in the incompressible case and this may be due to density matching problems between the two simulations. The jet mass flow rate has been adjusted by altering the value of α in each PODFS model such that the mass flow rate is similar to the compressible case, however the compressibility of the fluid can also affect the mass flow rate through local density changes. It might therefore be interesting to investigate the effects of either matching the mass flow rate or the velocity. The central recirculation zone (CRZ) is also smaller in the incompressible case, however this may be due to the finer mesh which would likely generate more unsteady structures in the shear layers and hence a faster decay of the jets.

The azimuthal velocity is shown in Figure 8.6b and also shows that the pilot jet velocity is slightly higher in the incompressible case. In the recirculation zone downstream of the splitter plate, in the incompressible case, the azimuthal velocity appears to be closer to zero, while in the compressible case, the flow appears to be rotating at a speed closer to the main jet flows. The largest changes in the velocity field can be seen in the fluctuating velocity components. The xx component of the Reynolds stress is shown for both cases in Figure 8.7. The level of fluctuations coming from the passages for both cases is comparatively minor in relation to the fluctuations further downstream. Due to the finer mesh resolution of the incompressible method, the resolved fluctuation levels are higher.

The mean heat release rate fields are shown in Figure 8.8a and the fluctuations in Figure 8.8b. These figures show that the flame positions are quite consistent between the incompressible and compressible cases. The length of the pilot flame is shortened in the case of the incompressible flow most likely due to the increased mesh resolution. There appears to be an increased quantity of fuel hitting the wall in the compressible case. This may indicate that the evaporation rate of the droplets is mesh dependent. One possible mechanism is that when the mesh resolution is higher then the higher level of resolved velocity fluctuation will lead to a more complicated and longer integration path for each droplet. This in turn means that the particle takes longer to arrive at the wall and hence has a better chance of evaporating before reaching the wall.

The fluctuating heat release fields show that the pilot flame fluctuations are consistent with the changes to the mean. The main flame fluctuations are concentrated in the part of the flame which is detached from the wall in the incompressible case while they are also apparent at the wall in the case of the compressible flow. The fact that some fluctuations of heat release are close to the wall shows that the wall is a very efficient amplifier of fluctuations of the fuel spray. However these images do not necessarily reflect fluctuations at the forcing frequency and are therefore not necessarily responsible for the differences observed in the flame transfer functions of the two cases.

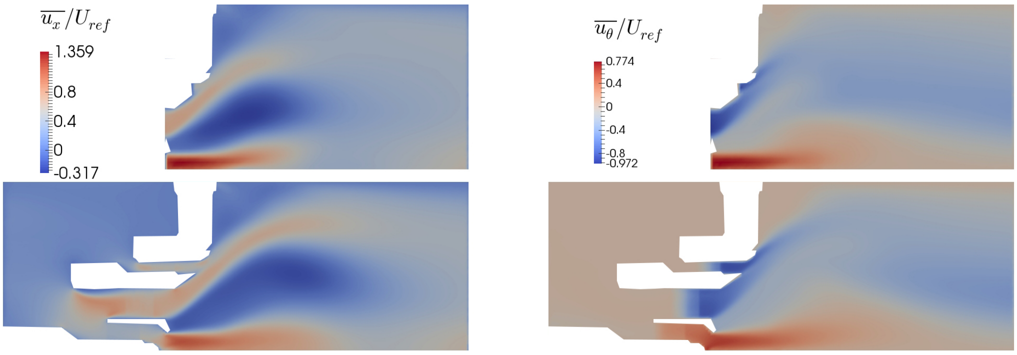

(a) Axial velocity. (b) Azimuthal velocity.

Figure 8.6: Mean and azimuthal average of axial and azimuthal velocity of the incompressible high density forced case (top) and the medium density compressible case (bottom).

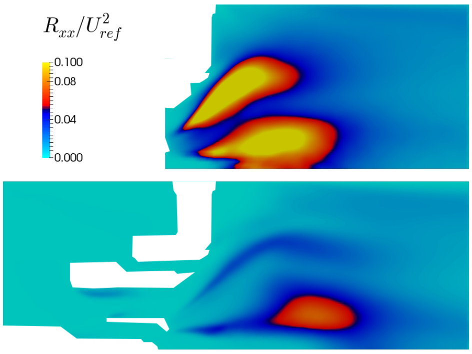

Figure 8.7: Azimuthally averaged R_{xx} of the incompressible high density forced case (top) and the medium density compressible case (bottom).

8.3 The prediction of the FTF

The flame transfer function has been calculated using two different methods. The first uses the fluctuating pressure signals from the microphones in the duct to calculate the amplification of sound waves by the flame. The second method measures the heat release directly from the flame and the mass flow rate directly from the flow passing through the exit of the injector.

8.3.1 The multi-microphone technique

Because the direct measurement of heat release is difficult in experiments, it has been measured indirectly in the SCARLET rig by analysing the change in the amplitude of the sound waves entering and exiting the combustion chamber. The acoustic field is linear and as such the response of the combustion chamber to acoustic excitation can be split between effects due to the damping of the sound by the impedance of the combustion chamber and the damping or excitation of the field by the flame. In order to derive the flame transfer function, four simulations or experiments

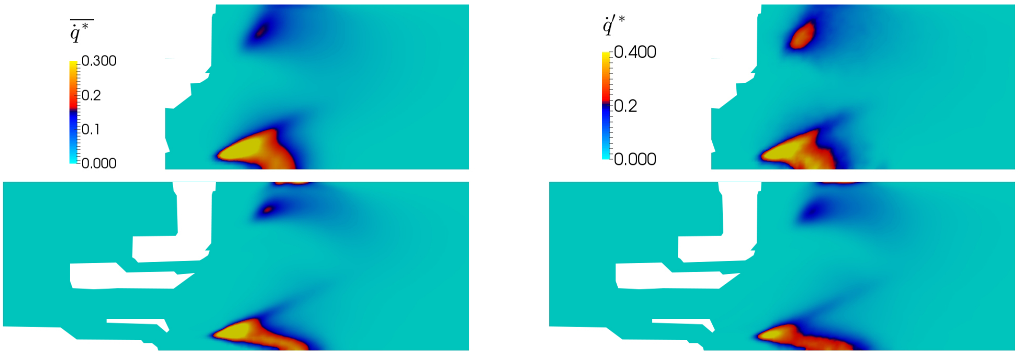

(a) Mean heat release rate.

(b) Heat release rate fluctuations.

Figure 8.8: Mean and RMS of azimuthally averaged heat release rate fluctuations of the incompressible high density forced case (top) and the medium density compressible case (bottom).

are required: The first two are conducted with non-reacting flow to obtain the acoustic transfer matrix of the combustion chamber in exactly the same way as in Section 5.5. The same two simulations are then run with the burner lit to obtain the flame and combustion chamber transfer matrix. Subtracting the initially derived acoustic transfer matrix of the combustion chamber from the flame and combustion chamber transfer matrix results in a transfer matrix for the flame itself. Linearisation of the Rankine-Hugoniot relations for a small acoustic perturbation results in the following equation linking the flame transfer matrix ($\mathbf{T}_{flame}$) to the flame transfer function ($\mathcal{F}$) [116]:

$$\mathbf{T}_{flame} = \begin{bmatrix} \frac{\rho_c c_c}{\rho_h c_h} & -\left(\frac{T_h}{T_c} - 1\right) M_h (1 + \mathcal{F}(\omega)) \\ -\gamma \left(\frac{T_h}{T_c} - 1\right) M_c & 1 + \left(\frac{T_h}{T_c} - 1\right) \mathcal{F}(\omega) \end{bmatrix} \tag{8.17}$$

where M is the Mach number and γ is the ratio of specific heats. The subscripts c and h refer to the values in the cold (upstream) and hot (downstream) ducts.

The first transfer function is presented in Figure 8.9 as a comparison between numerical and experimental results using the same post-processing methodology: the multi-microphone technique. The numerical results are shown using the low resolution and medium resolution meshes assuming compressible flow forced at $St = 0.14$. Because this method of obtaining the FTF is dependent on the resolution of acoustic waves in the upstream and downstream ducts, the incompressible simulations using the high resolution mesh cannot be used to obtain the FTF using this same methodology. Figure 8.9 shows that in the case of the low resolution mesh, the gain of the transfer function is greatly over-predicted and for the medium resolution mesh, slightly under-predicted. Given that the fuel spray fluctuations are not included in this prediction, it is possible that unaccounted fuel spray related mechanisms are responsible for this under-prediction. In the case of the low resolution mesh, it is likely that insufficient resolution of the flame front has resulted in the observed over-prediction. Looking at the phase response of the flame, it is interesting to note that both simulations capture the phase relationship reasonably accurately at this frequency. It would of course be interesting to obtain the full flame transfer function or additional frequencies to compare more fully the different techniques but as each frequency requires up to four simulations, this method is the most computationally expensive strategy for acquiring the flame transfer function and is likely impractical for determining the flame response during the design process.

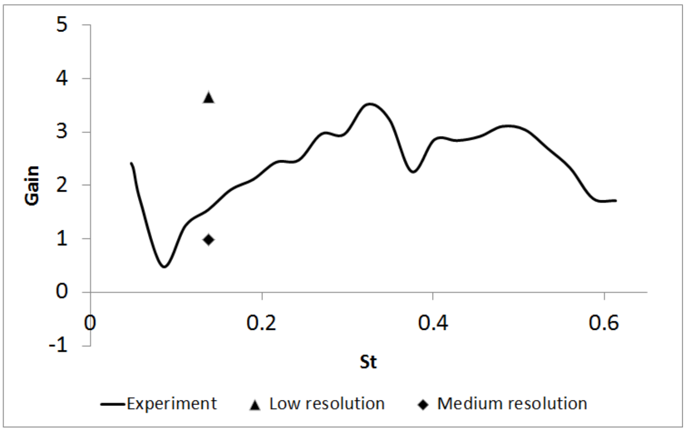

(a) Gain.

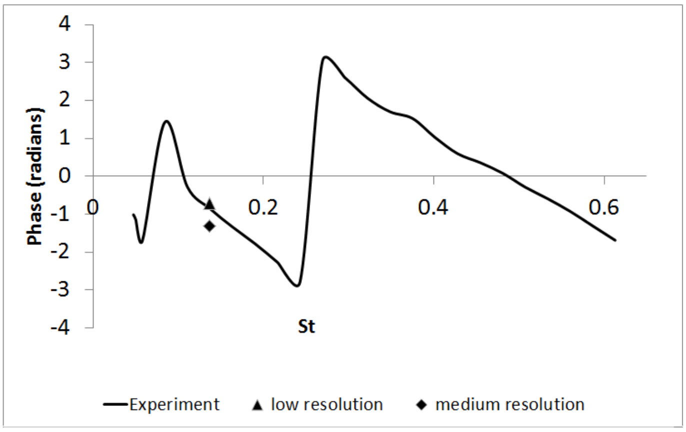

(b) Phase.

Figure 8.9: The FTF calculated by using the multi-microphone technique with the pressure signals from the microphones in the experiment and compared with pressure signals recorded from the cells lying along the centreline of the ducts in the numerical simulations.

8.3.2 Real-time extraction of heat release and mass flow rates

The compressible LES resolves the acoustic waves in the duct and by monitoring the pressure at the locations corresponding to the positions of the microphones in the experimental rig, the FTF may be derived in exactly the same fashion as in the experiments. This methodology requires a total of four simulations per frequency. The number of simulations can be reduced from four to one by extracting the heat release and mass flow rate from a reacting simulation directly.

Every 50 timesteps the flow solution was saved and a post-processing script run to calculate the mass flow rate. The mass flow rate was obtained by interpolating the flow solution at the exit plane of the injector onto the cell centres of a 200x200 point planar mesh. The mass flow was then found through each cell by multiplying the flow normal velocity by the local density and the area of the cell. These cell mass flow rates were then summed to achieve the total injector mass flow rate. The heat release rate was found at each timestep by multiplying the specific heat release rate as stored in the FGM table for each cell by the volume of that cell. The cell heat release rates were then summed across the domain to get the total heat release rate. Once the heat release rate and the mass flow rate was calculated the FTF was calculated using:

$$\mathcal{F} = \frac{S_{\dot{m}\dot{Q}}}{S_{\dot{m}\dot{m}}} \tag{8.18}$$

where $S_{\dot{m}\dot{Q}}$ is the cross-spectrum of mass flow and heat release rates and $S_{\dot{m}\dot{m}}$ is the power spectral density of the mass flow rate. The spectra were calculated using a fast Fourier transform with a Hanning window and Welche's method [159] where the frequency resolution was set to be $St = 0.02$.

The incompressible LES method does not resolve the acoustic waves and hence the flame transfer function can only be calculated from the direct measurement of the heat release and mass flow rates. The method has an advantage though in that because it is incompressible, the mass flow rate must be maintained from the inlets to the exit plane of the injector. For this reason the mass flow rate can be calculated by simply summing the mass of fluid that is entering the cells on the inlet boundaries. This can be done efficiently by the code itself and is computed at every timestep. The heat release rate is also calculated directly by summing the heat release rate from all of the cells in the combustion chamber at every timestep. The FTF was then calculated in a post-processing step using Equation 8.18.

The next flame transfer function shows in Figure 8.10 the experimental results as compared to the numerical results where the heat release rate and mass flow rate are measured directly. The low and medium resolution mesh results forced at $St = 0.14$ are shown together with a downstream forced calculation at $St = 0.30$. It can be seen that the downstream forced simulations give more consistent results to the microphone method derived results than the upstream forced results. This is likely due to the magnitude of the downstream forced response being significantly higher than the upstream response. This difference in response may be due to the difference in the time delay between the axial and azimuthal velocity fluctuations between the upstream and downstream forced cases. The higher gain in the case of downstream forcing suggests that combustion instabilities for this type of burner are likely to involve a feedback loop between the flame and the fuel injector without interacting with the engine components further upstream. If the upstream and downstream response of the flame is different, it may be sensible to complete future thermodynamic analysis using the full flame transfer matrix rather than the flame transfer function as this includes the full response of the flame from upstream and downstream waves.

Similarly to the microphone derived FTFs, the low resolution mesh results in a very high gain when forced from downstream whereas the medium resolution mesh results in an under-prediction. For the upstream forced simulations the gain is quite consistent across different mesh sizes. In the phase, the upstream forced simulations are consistent with each other and the downstream forced simulations also. The higher frequency simulation with the medium resolution mesh shows an under-prediction of gain and good prediction of phase. These results show that even poor resolution simulations are reasonably accurate at predicting the phase of the FTF and that the response of the flame is likely to be different for upstream and downstream forced cases.

The microphone derived results show a consistently higher prediction of gain and lower prediction of phase as compared to the directly measured FTFs. Such a persistent shift implies that there is a bias introduced either in the measuring of the heat release rate from the FGM table or due to the linear approximations made in the derivation of Equation 8.17.

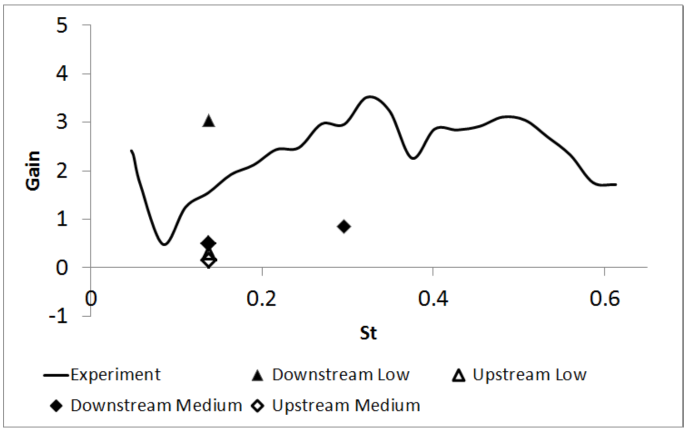

(a) Gain.

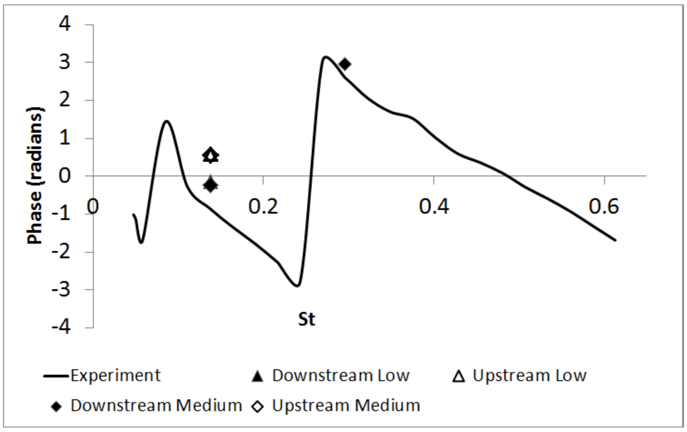

(b) Phase.

Figure 8.10: The FTF calculated by using the multi-microphone technique with the pressure signals from the microphones in the experiment and compared with the FTF as calculated from the directly calculated mass flow rate and heat release rate in the simulations.

8.3.3 Compressibility

Assuming that the mechanisms which cause the fluctuation of heat release are driven by convection of the flow from the fuel injector, then there is no reason why a suitably forced injector outflow cannot reproduce the effects of the acoustic forcing. Figure 8.11 shows the flame transfer function as calculated using the high resolution mesh forced by a PODFS built from acoustically forced non-reacting URANS simulations using the low density mesh combined with data from an incompressible high resolution LES simulation to take account of turbulence and maintain the pilot jet stability. The results show that the higher resolution mesh results in a slightly higher gain than the medium density mesh which shows that the mesh density could be improved for the compressible calculation. The phase is captured accurately by both methods in the case of the higher frequency but is better captured by the high density mesh for the lower frequency. This change in phase may be driven by the change in the location and size of the flame as shown in Section 8.2.3.

These results confirm that the PODFS method can effectively reproduce the convective effects of acoustically generated mass flow fluctuations through the fuel injector. The discrepancy of transfer function gain is most likely due to insufficiently modelled spray fluctuations such as those which effect the SMD and velocity of generated fuel droplets. It is also possible however that the choice of combustion model, grid resolution and pressure effects on combustion could also cause the observed deviation from experimental results. It would be useful to test the methodologies developed in this thesis on a premixed geometry as this would then remove the fuel spray fluctuations as a possible mechanism of flame response.

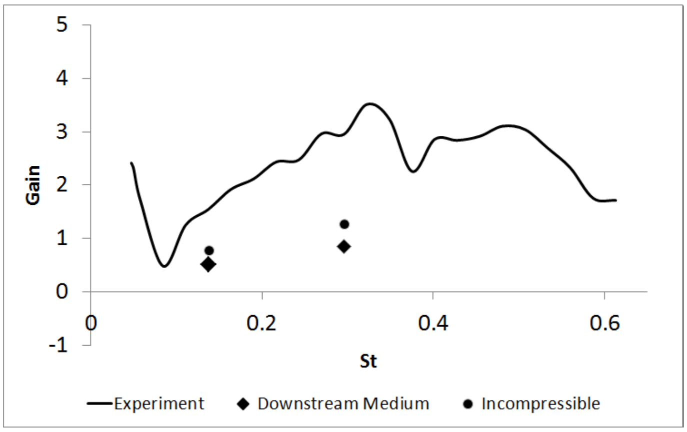

(a) Gain.

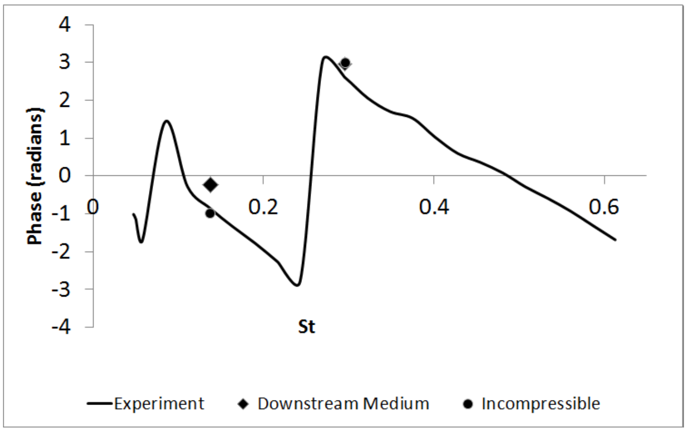

(b) Phase.

Figure 8.11: The FTF calculated by using the multi-microphone technique with the pressure signals from the microphones in the experiment and compared with the FTF as calculated from the directly calculated mass flow rate and heat release rate in the compressible and incompressible simulations.

8.3.4 Multiple frequency forcing

Finally, FTF results are presented from a simulation that was forced from downstream using multiple frequencies at once. This FTF is generated from direct measurements of the heat release and mass flow rates. The forcing was applied at the same frequencies as those chosen during experiments up to a maximum frequency corresponding to $St = 0.3$ to ensure that the mesh would accurately resolve the acoustic signal. This strategy is interesting because it is much cheaper than having to force the injector at many frequencies individually.

The flame transfer function is shown in figure 8.12 and shows that up to $St = 0.07$ the flame transfer function is reasonably well predicted using this multiple frequency forcing method. At higher frequencies however, the gain is over predicted and the phase badly captured. The problem with this method is that the fluctuations of the flame are not just caused by acoustically driven fluctuations of mass flow rate. The flame itself is unstable and may have many of its own naturally unsteady excitation frequencies [2]. These other instabilities result in fluctuations of heat release rate in addition to the fluctuations generated by the acoustic signal. Some of these instabilities are driven by convective processes both upstream from the flame to the injector via the recirculation zones and downstream via the main or pilot jets. Because the convection velocity upstream is different to the speed of sound, the characteristic frequency and delay of these fluctuations are different to thermoacoustic fluctuations and will not share the same phase relationship. Furthermore, the mass flow rate can fluctuate due to changes in density, driven by changes in local mixture fraction or temperature. If these naturally unstable fluctuations are stronger than the thermoacoustic type forced fluctuations then the gain will be over-predicted and the phase incorrect. It is important therefore that the magnitude of the mass flow rate fluctuations that are due to the acoustic forcing are of significantly higher magnitude than the natural unsteadiness of the flow. This ultimately requires the amplitude of acoustic forcing to be increased.

The increase of the acoustic fluctuation magnitude is limited however, to a value that is low enough to ensure a linear response from the injector. If the magnitude of the fluctuation exceeds the magnitude of the mean pressure drop then at the low frequency limit the flow through the injector will reverse, severely altering the shape and location of the flame. In the case of forcing the flow at a single frequency, the magnitude of the forcing does not need to be so high to ensure that the registered fluctuations of heat release rate are thermoacoustic in nature. However, when the flame is forced at many frequencies at once, the forcing magnitude must be much higher as the acoustic energy is split between all of the excitation frequencies. It may therefore be possible to force the flame at two or three frequencies at once and reproduce the results as seen while forcing at the same frequencies separately.

This particular injector was observed to have a high magnitude self excitation at $St = 0.2$ in experimental tests. This peak is also observed in the FTF calculated using the multiple frequency method. The unsteadiness of the injector is a function of its design. An injector with many naturally unstable fluctuations of heat release may be good or bad from a thermoacoustic point of view, however a quieter injector will be easier to test as lower magnitude forcing will be required to obtain the FTF for this design, which means that more frequencies can be tested simultaneously.

8.3.5 Azimuthal wavenumber of velocity component of POD modes

The POD and DMD modes are discussed and compared for the case where the flow was forced at $St = 0.14$. The POD has been computed for the whole field over a limited number of snapshots (224) for the high resolution mesh due to limitations of the available computer memory and the

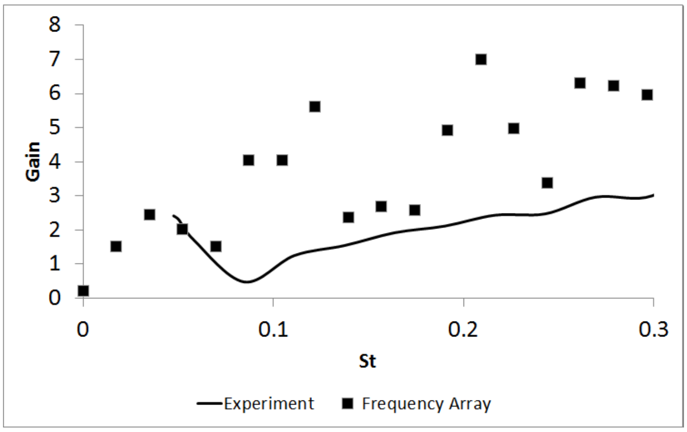

(a) Gain.

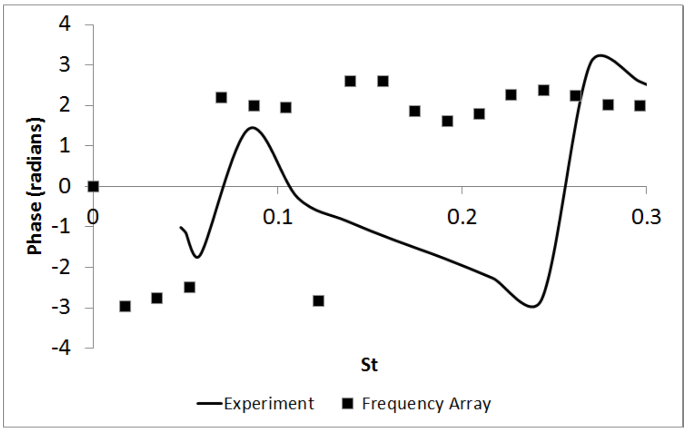

(b) Phase.

Figure 8.12: The FTF calculated by using the multi-microphone technique with the pressure signals from the microphones in the experiment and compared with the FTF as calculated from the directly calculated mass flow rate and heat release rate in the compressible simulations where forcing was applied at multiple frequencies at once from downstream.

DMD was computed over the low resolution azimuthally averaged plane using many more snapshots (400) for the medium resolution compressible case.

Because of the limited amount of computer memory available it was not possible to compute the POD of the full 3D field using all of the available snapshots. Three different PODs were considered in this section of various simulations forced at $St = 0.14$: Firstly, and in order to compare the results from Chapters 4 and 5, the POD was computed for the non-reacting low resolution case using the velocity variable only. Changes in the mode structures as compared to previous cases are therefore driven by the small changes to the injector geometry, mesh, forcing or operating point. The second POD was computed of the velocity field in the case of reacting flow using the medium density mesh. In this case, changes to the mode shapes are driven mostly by the changes in the density field. The third POD was computed of the reacting flow using the high density mesh but considered more variables than the velocity in the computation of the inner product. In this case, the velocity, spray mixture fraction, spray momentum, progress variable, mixture fraction, heat release, temperature, density and pressure were all used in the POD. Variations between the modes computed in this POD and those computed using only velocity are due to some structures being favoured by the reacting flow variables.

The azimuthal wavenumbers of the top 10 modes for each POD are shown in Table 8.1 along with the results from Chapter 4 (first column). In comparison with the results from Chapter 4, the POD of velocity for the non-reacting case in the SCARLET rig shows that the $m = 1$ mode has been suppressed with a subsequent promotion of the $m = 2$ modes. Low energy $m = 3$ modes also appear.

The introduction of fuel and reaction shows the promotion of a $m = 0$ mode which is expected as the FTF gain is greater than one at this frequency. The first of the $m = 3$ modes is also promoted at the expense of $m = 2$ modes while $m = 1$ modes reappear. The changing stability of these modes is likely due to changes in the mean density field.

In the case where the inner product is extended to include the reacting variables, the $m = 3$ and $m = 0$ modes are pushed further down the table and replaced by $m = 2$ modes. The results show that the addition of the fuel spray and heat release results in a preference for the $m = 1$ helical structures, and the flame itself shows a preference for $m = 2$ modes as the inclusion of reacting variables and spray promotes these modes in the table. As already mentioned, these fluctuations of heat release cannot affect longitudinal stability, however they may contribute to transversal (spinning) thermoacoustic instabilities.

8.3.6 DMD

As the medium density mesh results required less storage space than the high density results, more flow snapshots were saved. The increased number of snapshots made this data set ideal for DMD analysis. The field which was forced at $St = 0.14$ was azimuthally averaged to remove all non $m = 0$ structures and analysed using DMD to identify the main flow features responsible for the observed heat release fluctuations.

The DMD analysis shows that there is a single mode which contains the largest energy content fluctuating at the forcing frequency. The heat release rate and velocity components of this mode are shown in Figure 8.13. The DMD pilot flame shows a classical heat release fluctuation mechanism

Table 8.1: The azimuthal wavenumbers for the velocity components of the first 10 POD modes for the reacting and non-reacting SCARLET (SC) cases, together with the results from the compressible URANS computation from Chapter 4 using the Loughborough University Acoustic Duct rig (LUAD) and the high resolution cut-down SCARLET rig mesh (CSC). PODIP = POD inner product.

Geometry	LUAD	SC	SC	CSC
Compressible	Yes	Yes	Yes	No
Reacting	No	No	Yes	Yes
Turbulence	URANS	URANS	LES	LES
Mesh	LUBPAI	Low	Medium	High
PODIP	$\mathbf{u}$	$\mathbf{u}$	$\mathbf{u}$	all
Mode 1	1	2	0	2
Mode 2	1	2	3	2
Mode 3	2	0	1	1
Mode 4	2	0	1	1
Mode 5	0	0	2	0
Mode 6	0	3	2	0
Mode 7	0	0	2	0
Mode 8	0	3	1	3
Mode 9	0	3	2	3
Mode 10	0	3	1	3

as seen in laminar flames. As the axial velocity increases, the flame extends and when the axial velocity begins to reduce the momentum of the fluid at the tip of the flame causes it to detach and convect downstream. The main flame shows a different mechanism: The fluctuating swirl number causes the main jet cone angle to change which causes the flame to move radially inwards and outwards. The change in swirl number also affects the production of turbulence which then alters the flame speed which pulls the flame forwards or pushes it backwards. It can also be seen the rearmost tip of the flame near the wall also breaks away from rest of the flame, additionally changing the heat release rate.

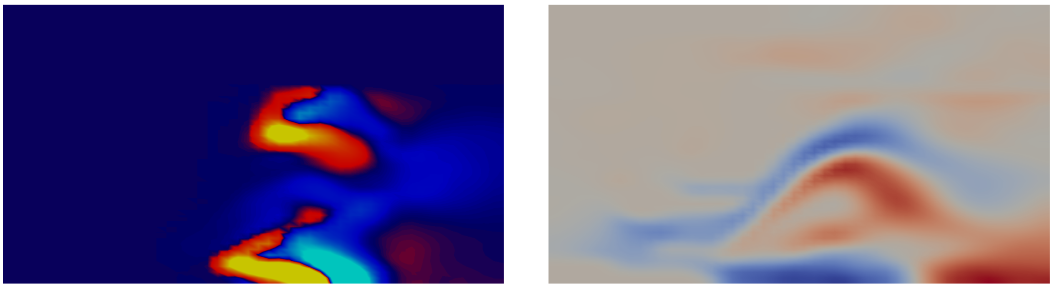

(a) Heat release rate. (b) Axial velocity.

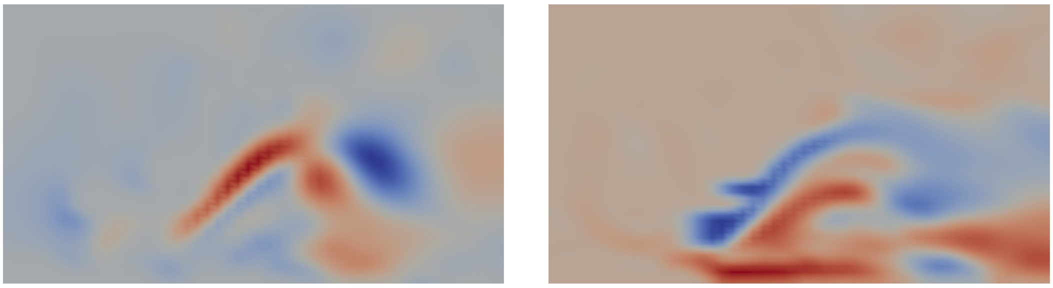

(c) Radial velocity. (d) Azimuthal velocity.

Figure 8.13: The most energetic DMD mode from the azimuthally averaged data of the medium resolution compressible simulation forced at $St = 0.14$. The mode has an eigenfrequency equal to the fluctuation frequency.

8.3.7 POD of the incompressible high density mesh simulation

This section further analyses the POD modes generated from the simulation of the forced flow at $St = 0.14$ using the high density mesh and incompressible method. Although the POD ranks the modes in terms of their contribution to the sum of the mean squared values of all the components, the ranking method does not pick out those modes that necessarily contribute the most to the fluctuations of heat release rate. In order to re-rank the modes in terms of their contribution to heat release rate, the mean square value of each component averaged across the field was calculated for each POD mode. The results are shown in Figure 8.14a. Despite the normalisation procedure before the POD is computed, certain components tend to contribute more to a mode than others. In order to take this into account, the mode contributions are then normalised by the mean values calculated across all of the modes which results in Figure 8.14b. The top 10 modes as ranked by their contribution to heat release are listed in Table 8.2 along with four additional modes that contribute greatly to the overall fluctuating flow energy.

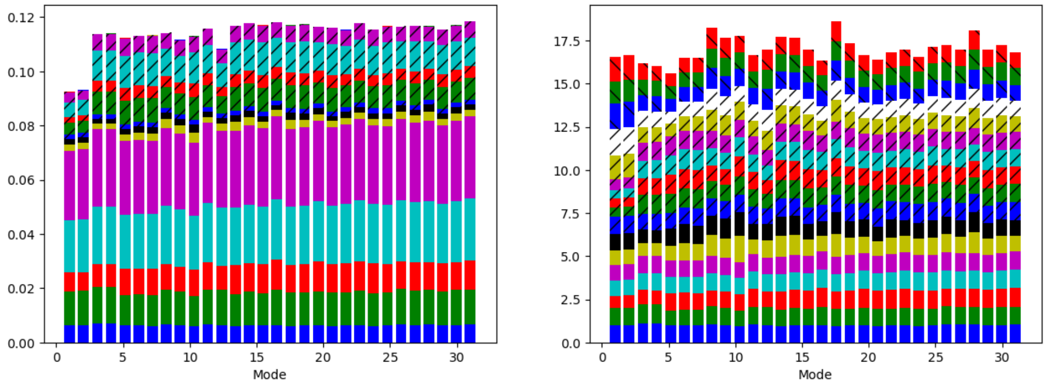

(a) Before normalisation of the mean contribution across the modes.

(b) After normalisation by the mean contribution across the modes.

Figure 8.14: The contribution of each component to each of the POD modes. Blue: axial velocity, green: radial velocity, red: azimuthal velocity, cyan: velocity magnitude, purple: velocity, yellow: SMD, black: spray mixture fraction, blue //: heat release rate, green // : temperature, red // : mixture fraction, cyan // : progress variable, purple // : density, yellow // : axial spray momentum, white // : radial spray momentum, blue \\: azimuthal spray momentum, green \\: spray momentum magnitude, red \\: spray momentum.

The first column in this table gives the rank of the mode according to how much it represents the fluctuating heat release rate. A smaller number corresponds to modes that correlate highly with fluctuations of heat release rate. In contrast, the second column is the mode ranking by how much the mode contributes to the fluctuating flow energy across all variables (velocity, density etc.). While the third column gives the proportion of total flow energy represented by that individual mode.

The table shows that the modes where the heat release contributes strongly are those with quite high mode numbers and therefore low energy. This suggests that the flame has a preference for small, low energy, high wavenumber flow structures. The table shows the six most active variables for each mode. These are the six variables which contribute the greatest to the normalised mode energy shown in Figure 8.14b. These modes can be loosely broken into two groups: Modes 30, 23, 24, 20 and 26, which show a low correlation between heat release and spray related variables and the remaining modes which do. A further observation can be made that most of these modes depend strongly on azimuthal velocity fluctuations. This can be explained in the same way as in Section 8.4. When the azimuthal velocity changes, the centrifugal force on the droplets changes and this drives the droplets outwards at different rates in time. This results in the areas of high mixture fraction and low progress variable varying in time, causing the heat release to change. Azimuthal velocity fluctuations are also responsible for changes in the turbulent flame speed through their contribution to swirl number [16]. Fluctuations of swirl number generate different amounts of turbulent kinetic energy which alters the effective turbulent flame speed, changing both the location of the flame and the reaction rate while fluctuations of axial velocity change the amount of large scale vorticity in the flame region, altering the flame area.

Table 8.2: The top 10 POD modes ranked by the normalised contribution of heat release rate fluctuations with the top 6 contributing variables.

Mode rank	Mode Number	Energy (%)	1st	2nd	3rd	4th	5th	6th
1	30	0.63	u_θ	$\dot{q}$	$\lvert\mathbf{u}\rvert$	$\mathbf{u}$	u_r	u_x
2	23	0.70	u_θ	$\dot{q}$	Z	ρ	$\lvert\mathbf{u}\rvert$	$\mathbf{u}$
3	24	0.69	Z	$\dot{q}$	Y	u_θ	T	ρ
4	10	0.96	D_{32}	Z_s	Z	Y	$\dot{q}$	$P_{s,x}$
5	20	0.74	Z	Y	T	ρ	$\dot{q}$	u_θ
6	28	0.65	Z	u_θ	$\dot{q}$	Y	$\lvert\mathbf{u}\rvert$	$\lvert\mathbf{P_s}\rvert$
7	26	0.66	u_θ	$\dot{q}$	Z	$\lvert\mathbf{u}\rvert$	$\mathbf{u}$	Y
8	25	0.67	D_{32}	u_r	$\lvert\mathbf{u}\rvert$	$\mathbf{u}$	$\lvert\mathbf{P}_s\rvert$	$\dot{q}$
9	27	0.66	D_{32}	$P_{s,x}$	Z_s	$\mathbf{P}_s$	$P_{s,\theta}$	$\lvert\mathbf{P}_s\rvert$
10	8	0.98	$P_{s,x}$	$\mathbf{P}_s$	$P_{s,\theta}$	$P_{s,r}$	D_{32}	$\lvert\mathbf{P}_s\rvert$
13	13	0.85	D_{32}	Z_s	Z	Y	T	u_θ
15	9	0.81	$P_{s,x}$	Z_s	$\mathbf{P}_s$	$P_{s,\theta}$	$P_{s,r}$	D_{32}
17	1	2.65	$P_{s,r}$	$P_{s,\theta}$	$\mathbf{P}_s$	$P_{s,x}$	$\lvert\mathbf{P}_s\rvert$	u_r
21	2	2.51	$P_{s,r}$	$P_{s,\theta}$	$\mathbf{P}_s$	$P_{s,x}$	$\lvert\mathbf{P}_s\rvert$	u_r

Modes number one and two are of particular interest because of their high contribution to the fluctuating flow energy and were identified as having an azimuthal wavenumber of $m = 2$. Figure 8.15 shows contours of axial velocity, axial fuel spray momentum, mixture fraction and heat release for the most energetic mode. The images are to scale which shows that the strongest velocity fluctuations in this mode are near the centreline of the combustor while the subsequent effects on the fuel spray, mixture fraction and heat release are most active further downstream and radially outwards. The strong $m = 2$ helical mode can be seen in the velocity field, which then causes the fuel spray to accelerate periodically leading to regions of high and low spray momentum and spray mixture fraction. The regions which contain a higher spray mixture fraction then produce regions of higher mixture fraction which consequently results in regions that produce a higher heat release rate. Because these fluctuations of heat release are axisymmetric they cannot directly contribute to longitudinal thermoacoustic instabilities. They can however contribute to transverse instabilities as the rotation of these structures could become coupled with an acoustic wave that travels around the annulus of the combustor [164].

Figure 8.16 shows the PSD of the temporal component for the 1st, 8th, 9th and 10th most energetic modes. The most energetic mode is a $m = 2$ mode and while it contains a lot of the fluctuating flow energy it is not a mode that contains a comparatively large heat release component. Modes 8, 9 and especially 10 contain significantly less flow energy but are heavily weighted towards the heat release rate and hence the flow structures seen in these modes are relatively efficient in generating large fluctuations of heat release. The PSD of the most energetic mode shows a single large peak at $St \approx 1$ which shows that this mode is independent from the forcing. The other three modes show a broad peak around the forcing frequency but also show a distinct secondary peak at the same frequency as the $m = 2$ mode which suggests that these modes contain a proportion

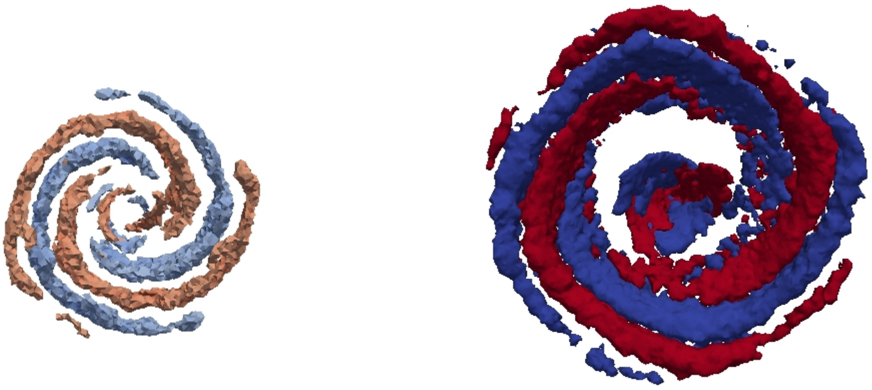

(a) The axial velocity component for the 2nd most energetic POD mode.

(b) The axial momentum component of the fuel spray momentum for the 2nd most energetic POD mode.

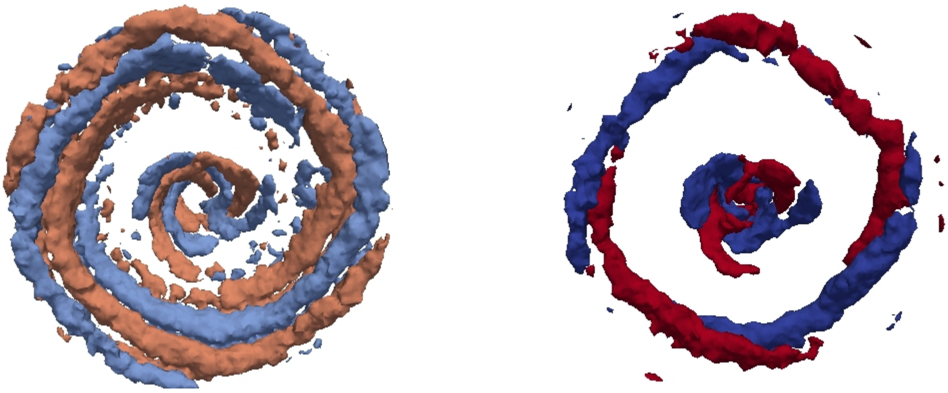

(c) The mixture fraction component for the 2nd most energetic POD mode.

(d) The heat release component for the 2nd most energetic POD mode.

Figure 8.15: The axial velocity, axial spray momentum, mixture fraction and heat release components of the 2nd most energetic POD mode. Blue are areas of low magnitude and red are areas of high magnitude.

of this $m = 2$ mode.

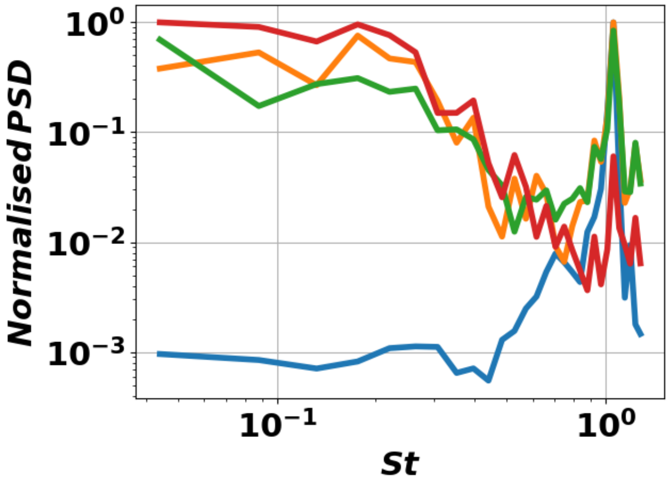

Figure 8.16: The normalised PSD of the temporal component of the most energetic POD mode (blue), 8th most energetic POD mode (orange), 9th most energetic POD mode (green) and 10th most energetic POD mode (red) for the incompressible high resolution reacting case forced at $St = 0.14$.

Four of the largest contributors to mode number 10 are shown in Figure 8.17. The mode azimuthal wave number was calculated to be $m = 3$ however the mode is visualised here by taking the azimuthal transform that isolates the $m = 0$ component. Figure 8.17a shows the radial velocity component of the mode which shows an alternating opening and closing of the cone angle of the jets similar to the DMD mode shown in Section 8.3.6. The axial momentum of the fuel spray and the spray mixture fraction show that the fuel spray is flicking between a state of being concentrated near the wall or dispersed downstream further towards the centreline. This flicking backwards and forwards of the spray can be seen to directly affect the heat release rate by changing the relative mixture fraction between the region next to the wall and the further downstream position. The dispersion of the spray downstream also seems to lead to some fuel being fed to the pilot flame while at the reverse phase there is a high magnitude flame at the centreline which may indicate that some droplets take a longer route around the outside of the combustion chamber and back to the pilot flame.

This rather simple mechanism which is likely responsible for a large part of the forced fluctuations is clearly augmented by the helical $m = 2$ mode due to the strong peak in the PSD at the natural mode oscillation frequency (Figure 8.16). This is the same case for the other two modes identified as being efficient at producing fluctuations of heat release. Due to the additional observation that smaller, lower energy structures are most highly correlated with heat release, it seems that although the helical modes are not directly adding to longitudinal heat release fluctuations, their interaction with other flow structures can augment the amount of heat release produced by $m = 0$ shaped modes. One possible explanation for this is that the acoustic forcing periodically increases the strength of helical modes which then generates structures that are of small enough scale that they greatly affect the heat release rate and are also able to affect the longitudinal heat

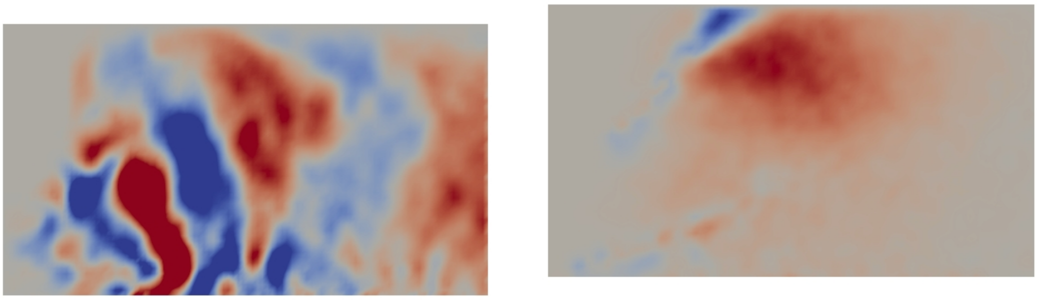

(a) The radial velocity component.

(b) The axial momentum component of the fuel spray momentum.

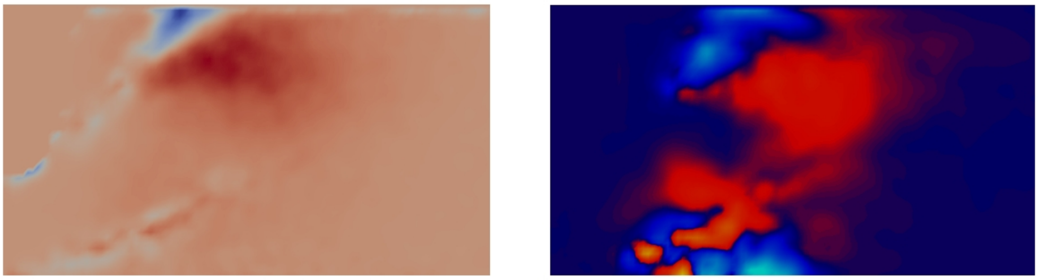

(c) The mixture fraction component.

(d) The heat release component.

Figure 8.17: The radial velocity, axial spray momentum, mixture fraction and heat release components of the 10th most energetic POD mode. Blue are areas of low magnitude and red are areas of high magnitude.

release fluctuations.

An interesting test would be to apply chevrons to one or more of the passage outlets. Chevrons are small deviations in the axial or radial direction of the otherwise circular shape of the jet outlet designed to promote mixing. Because the DMD and POD structures identified as being most important to the fluctuating heat release cause an opening and closing of the main jet, chevrons placed around the outlet of the injector may then induce periodic stream-wise vorticity that breaks up this large coherent structure and hence might reduce the mechanism behind the associated heat release fluctuations. Unfortunately the periodic increase of stream-wise vorticity will also periodically alter the turbulent flame speed in the main flame and hence amplify the effects of the flapping of the jet. Which of these effects are more important could only be decided after testing. Chevrons could also be added to the splitter plate between passages A and B as this might help to disrupt the formation of helical modes and hence reduce the prevalence of small scale structures that can most effectively alter the heat release rate.

8.4 The effect of fluctuations of SMD

The numerically predicted FTF's shown earlier in this chapter are reasonably accurate in phase but significantly underestimate the gain. Some of this discrepancy could be due to the numerical method however such a large discrepancy suggests that some of the fundamental mechanisms of heat release fluctuations are not being modelled. One possibility is that the acoustic wave could be altering the SMD of the atomised droplets as seen in [148].

Studies at Loughborough university suggest that the SMD may fluctuate by up to 16% in ambient conditions however Equation 7.1 suggests that the gradient of SMD with respect to velocity is given by:

$$\frac{dD_s}{dU_R} = \frac{-24\sigma}{\rho_A U_R^3} \tag{8.19}$$

where U_R is the relative velocity between the fuel and the air, σ is the fuel surface tension, and ρ_A is the air density. This equation shows that as the air velocity increases the gradient of the droplet diameters decreases with respect to one over the third power of velocity. In an engine, increases in the overall pressure ratio of the engine results in a larger reference velocity for a given pressure drop across the combustor liner. This results in a significant increase in the air velocity in the combustion chamber and a subsequent reduction in the sensitivity of the fuel spray to fluctuations in velocity caused by acoustic waves. This explains the approximately 1.2% in SMD fluctuation predicted using the correlations tested by Su et al. [149] and shown in Section 7.2.4. Measuring the fluctuations directly in the SCARLET rig would confirm whether or not spray velocity or SMD fluctuations are significant or not at elevated temperatures and pressures.

If it can be ascertained that SMD fluctuations are significant at engine operating conditions, it would be useful to know whether such fluctuations have a large effect on the heat release rate of the flame. Because the mass flow rate of fuel is fixed, the heat release fluctuations must be caused by periodic alteration of the time it takes for a droplet of fuel to convect, evaporate, mix and combust.

In order to test the response of the flame to fluctuations in SMD, the high resolution mesh was used to run an LES simulation where the SMD of the main fuel spray varied by 20%. An initial baseline computation was performed to obtain the mean flow and to act as a comparison to the results from the forced computation. In these two simulations only steady profile inlets were used as the upstream boundary condition resulting in a the pilot jet remaining in the ‘flipped’ configuration.

Comparison of the mean flow fields (temporally and azimuthally averaged) between the baseline case and the fluctuating SMD case (figures 8.18a and 8.18b) show identical results. Plots of mean droplet SMD (Figure 8.19a), mean spray mixture fraction (Figure 8.20a), mixture fraction (Figure 8.21a), and heat release rate (Figure 8.22a) show similarly identical results between the simulations with only a slight thickening of the inside edge of the main flame.

The only difference that can be seen are in the plots of the RMS values of the spray and combustion variables. The plot of the RMS of the spray SMD (Figure 8.19b) shows an expected noticeable increase in intensity in the main regions but without any major increase in the area where the spray is active. There is also a small increase in the intensity of fluctuations in the pilot spray which is surprising given that the pilot spray SMD is fixed. The fuel spray mixture fraction fluctuations (Figure 8.20b) show a similar pattern with the intensity of fluctuations in the main fuel spray most affected, however in this case, the footprint of the main fuel spray has been enlarged slightly, as has the pilot spray too. This variation translates to a slightly higher magnitude of the mixture fraction at the wall but no discernible difference to the fluctuations of mixture fraction in

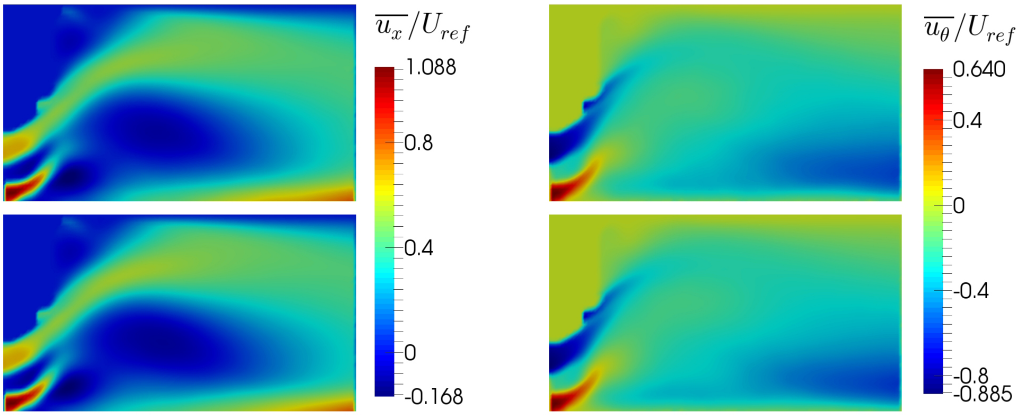

(a) The mean axial velocity.

(b) The mean azimuthal velocity.

Figure 8.18: The mean and azimuthally averaged axial and azimuthal velocities for the cut-down geometry with a steady flow inlet and spray boundary conditions (top) and fluctuating SMD (bottom).

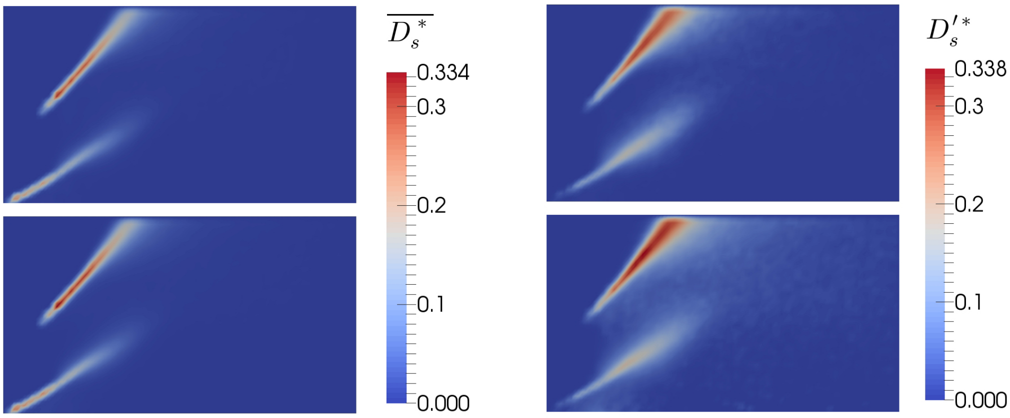

(a) The mean droplet SMD.

(b) The RMS of the droplet SMD.

Figure 8.19: The azimuthally averaged mean and RMS of the fuel spray SMD for the cut-down geometry with a steady flow inlet and spray boundary conditions (top) and fluctuating SMD (bottom).

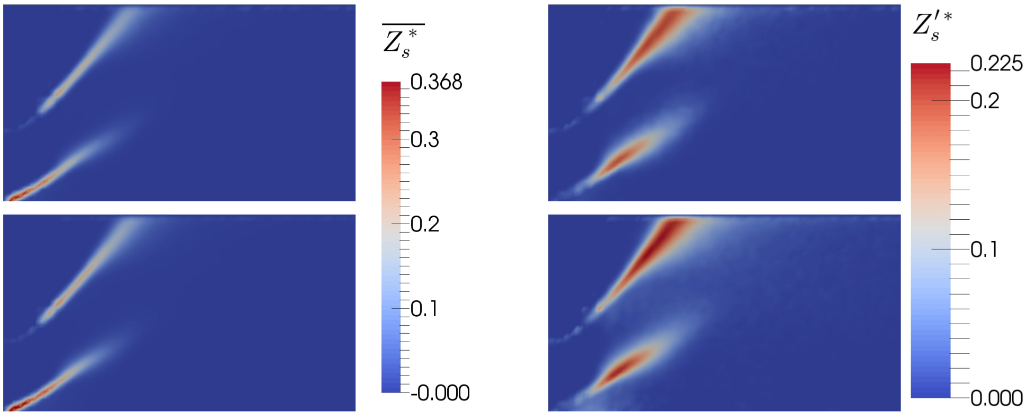

(a) The mean spray mixture fraction.

(b) The RMS of the spray mixture fraction.

Figure 8.20: The azimuthally averaged mean and RMS of the spray mixture fraction for the cut-down geometry with a steady flow inlet and spray boundary conditions (top) and fluctuating SMD (bottom).

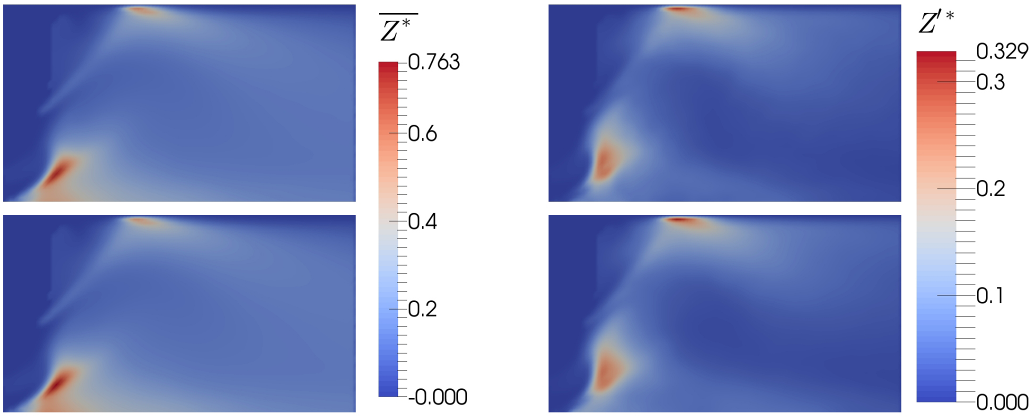

(a) The mean mixture fraction.

(b) The RMS of the mixture fraction.

Figure 8.21: The azimuthally averaged mean and RMS of the mixture fraction for the cut-down geometry with a steady flow inlet and spray boundary conditions (top) and fluctuating SMD (bottom).

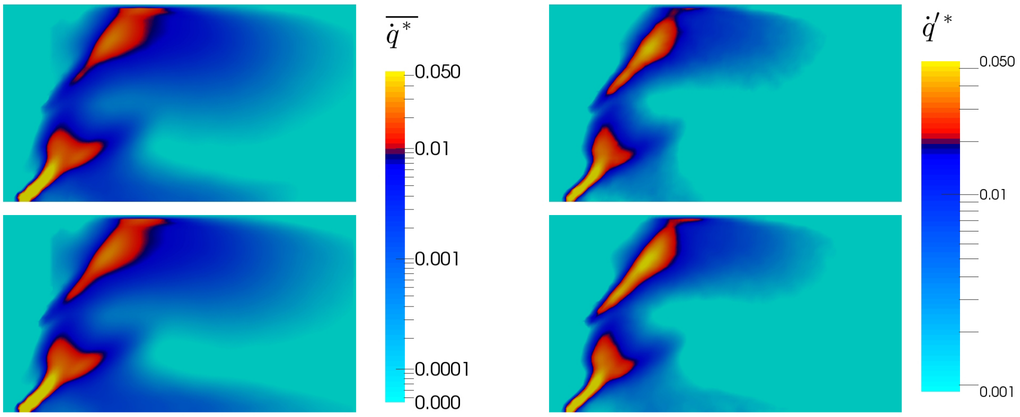

(a) The mean heat release rate.

(b) The RMS of the heat release rate.

Figure 8.22: The azimuthally averaged mean and RMS of the heat release rate for the cut-down geometry with a steady flow inlet and spray boundary conditions (top) and fluctuating SMD (bottom).

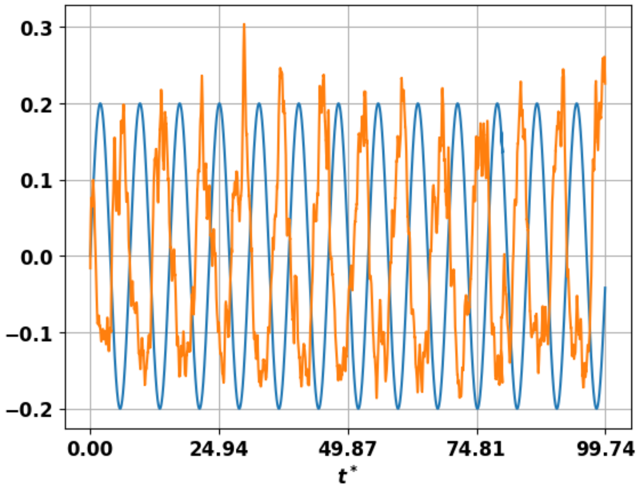

Figure 8.23: The normalised fluctuations of SMD (blue) and heat release (orange).

the pilot zone (Figure 8.21b). This translates to a noticeable increase in the area over which the main flame heat release fluctuates (Figure 8.22b). This relatively small change in the magnitude of the heat release fluctuations results in a not insignificant 20% fluctuation in the flame heat release. The heat release rate normalised by the mean heat release rate is shown together with the SMD fluctuation normalised by the mean SMD in Figure 8.23. This plot shows that there is significant cycle-to-cycle variation in the magnitude of the response with some underlying low frequency feedback mechanism. It appears that although the net effect of a higher SMD is negative with a higher SMD being correlated to a lower heat release rate, there is a brief spike of heat release when the SMD initially begins to rise.

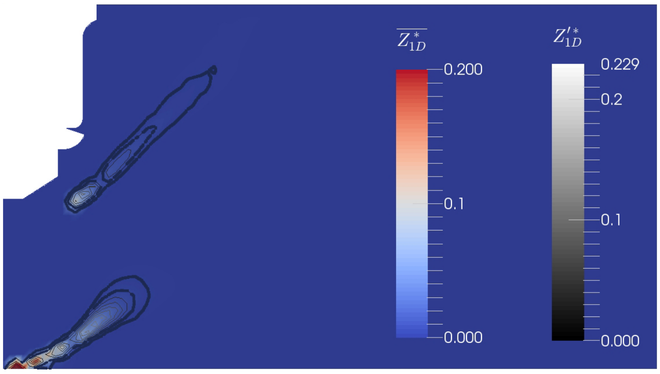

(a) The smallest 20% of droplets.

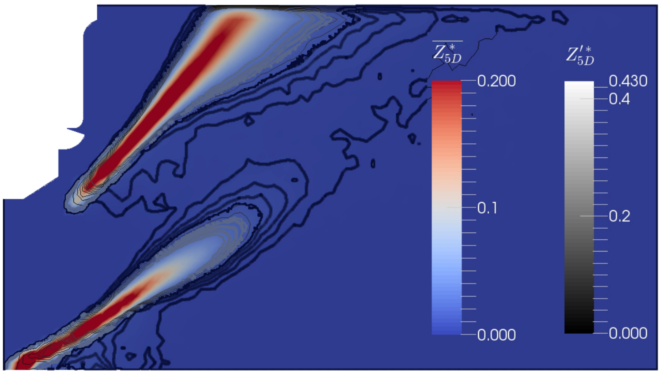

(b) The largest 20% of droplets.

Figure 8.24: The mean and RMS of the spray mixture fraction of different sized droplets in the case where the SMD was varied in time. The contour levels are the same for both figures.

In Figure 8.24 the spray mixture fraction and its RMS is plotted for the largest 20% and smallest 20% of droplets. In comparison to figures 8.20a and 8.20b it can be seen that the larger droplets drive the response of the spray mixture fraction. It can be seen that in the mean, there is not a significant difference in the angle that the droplets leave the atomiser for both droplet sizes however in the RMS field it can be seen that the larger droplets penetrate much further downstream. Consistent with what was seen in Chapter 7, the largest droplets remain in the main jet for longer and are therefore convected further downstream. It is noticeable from the RMS of the largest droplets that these droplets are disproportionately distributed downstream of the mean line compared to upstream. From this it can be ascertained that when the SMD reaches its maximum value the droplets move further downstream than they move upstream when the SMD is a minimum.

8.4.1 Phase average

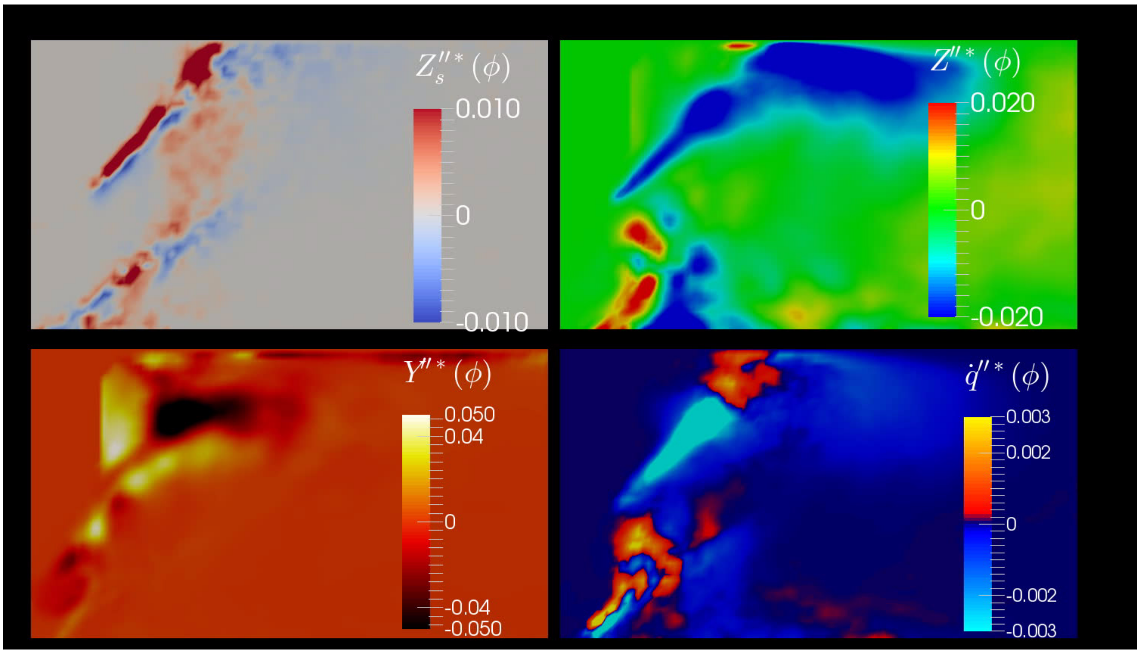

Figure 8.25: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 0.21\pi$.

In order to see how the flame responds to the forcing of the SMD a phase average was calculated at the forcing frequency. Figure 8.25 shows the phase and azimuthally averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 0.21\pi$. At this point of the cycle, the SMD in the main spray is low and the spray moves upstream slightly resulting in a small build-up of fuel mixture fraction where the droplets are forcibly evaporated at the hot wall. This results in a flame hotspot forming at the wall. At $\phi = 0.32\pi$ (Figure 8.26) the flame hotspot increases in size and as the SMD increases the fuel spray begins to pivot downstream. This pivot culminates in a large glut of fuel droplets at the wall further downstream by $\phi = 0.74\pi$ (Figure 8.27) and a corresponding increase of mixture fraction in the region of the wall and now extending downstream. The flame expands both downstream and upstream at this point to take advantage of the higher mixture fraction region. Some of the larger droplets that are not evaporated by the flame are convected downstream as can be seen in Figure 8.28 at $\phi = 0.94\pi$. The now smaller diameter droplets evaporate much faster creating a second area of mixture fraction concentration near the atomiser. The flame now extends far upstream, boosting the heat release rate which reaches a maximum near $\phi = 1.37\pi$ (Figure 8.29). There is now a deficit of droplets hitting the wall and the mixture fraction in this region decreases leading to a lower heat release rate, which splits the flame into a downstream and upstream branch. As the SMD of the droplets begins to increase the flame is forced back from the atomiser as it takes longer for these droplets to evaporate, as can

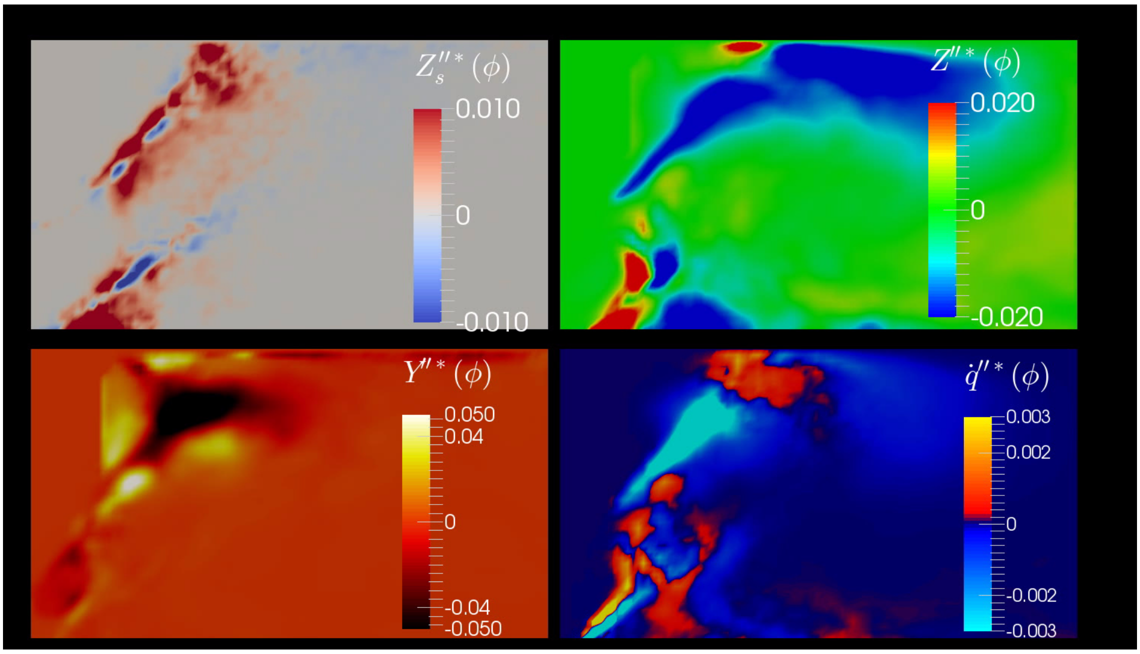

Figure 8.26: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 0.32\pi$.

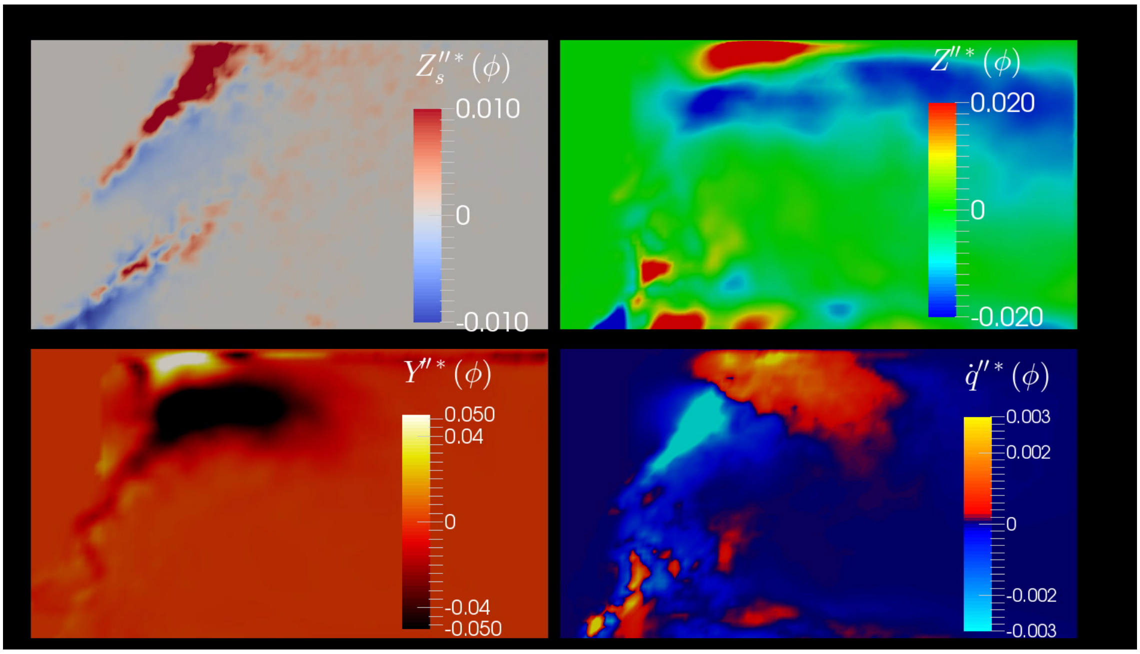

Figure 8.27: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 0.74\pi$.

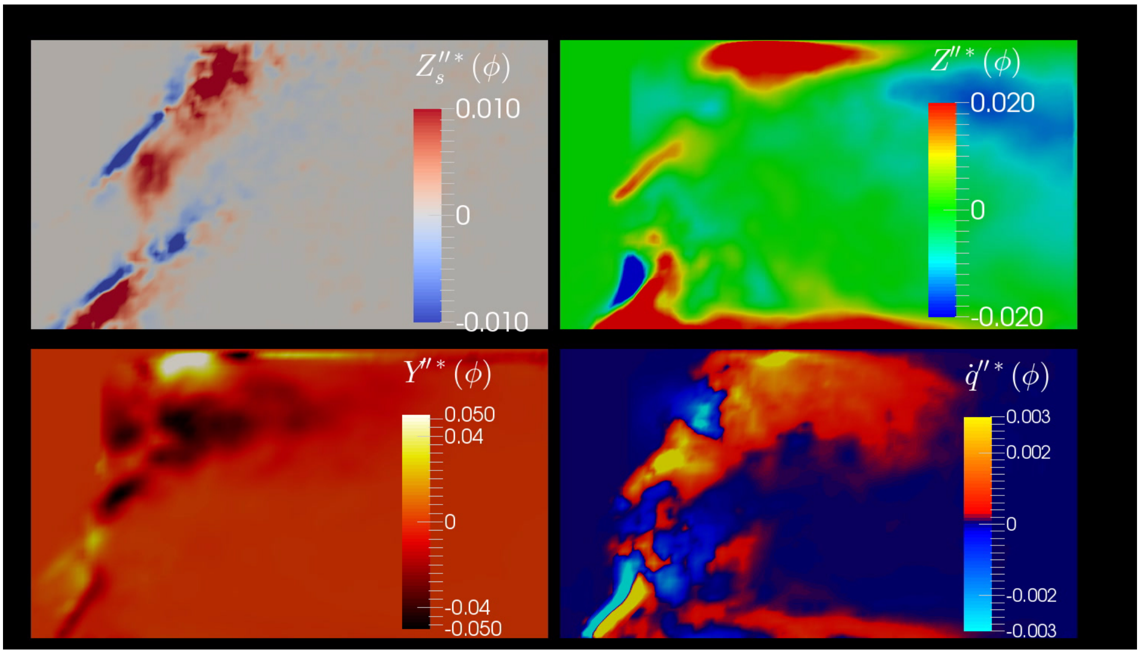

Figure 8.28: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 0.94\pi$.

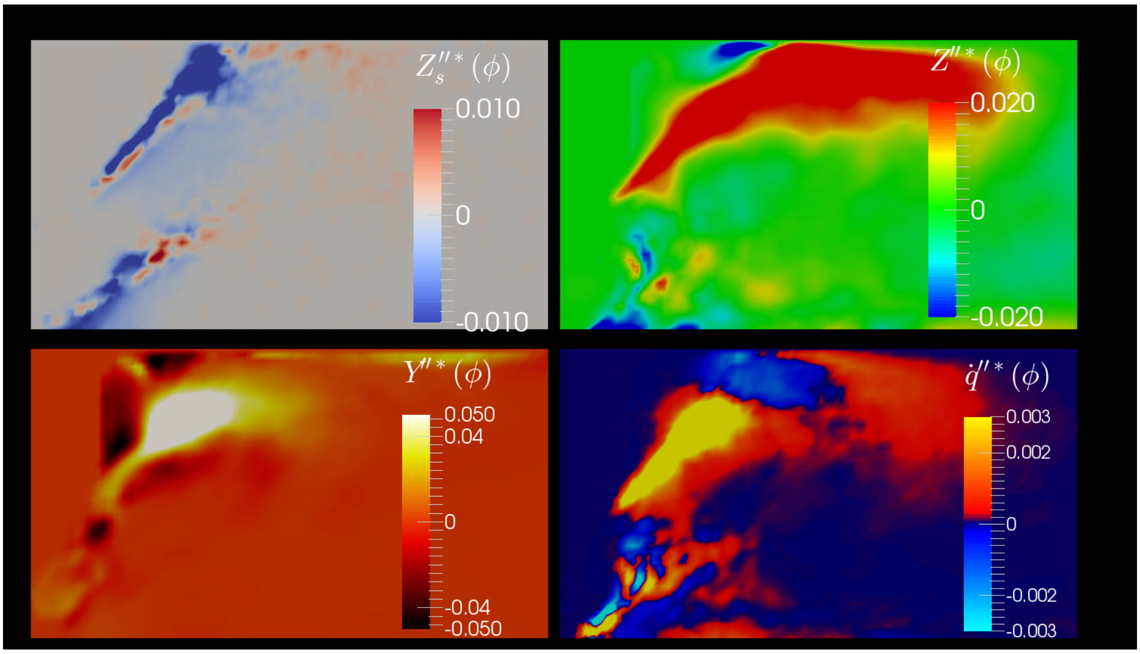

Figure 8.29: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 1.37\pi$.

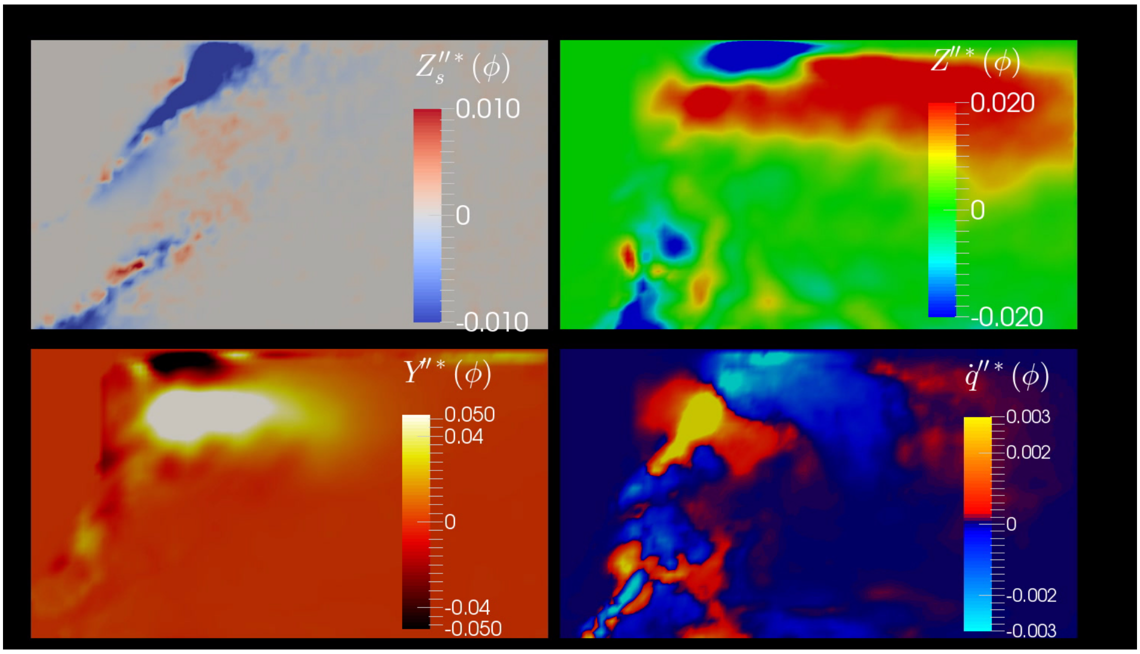

Figure 8.30: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 1.68\pi$.

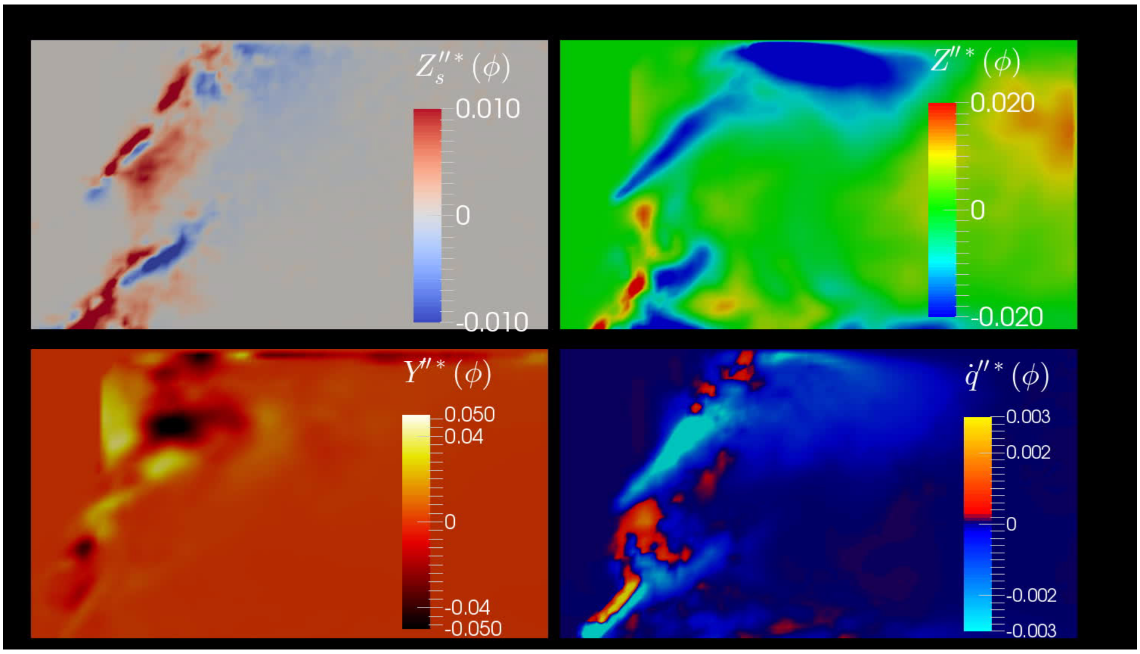

Figure 8.31: The phase averaged spray mixture fraction, mixture fraction, progress variable and heat release rate at $\phi = 0\pi$.

be seen in Figure 8.30 ($\phi = 1.68\pi$). Finally the flame is at a minimum as the SMD approaches a maximum at $\phi = 2\pi = 0\pi$ in Figure 8.31.

8.4.2 POD

Whilst the phase averaged results show a clear mechanism for the fluctuations of heat release in the main flame, it is plain to see that there is significant interaction between the pilot and main flames and that the pilot flame may either contribute to or detract from the fluctuations of the global heat release rate. In Figure 8.26 it can be seen that there is large glut of the fuel spray droplets in the centre of the pilot zone and an associated fluctuation in heat release here in this region. This is in contrast to Figure 8.30 where the reverse is true.

There is also a question of intermittent stability of the main fuel spray. At several points in the phase (see for instance Figure 8.28), there appears to be a breakup of the main spray that makes it through the flame front and causes a glut of spray mixture fraction downstream of the main flame. At first it was suspected that this breakup was driven by velocity fluctuations such as those described in Section 7.5.3 however as there are no imposed upstream velocity fluctuations, the velocity must be modulated by the fluctuating heat release itself. Analysis of the phase averaged velocity fields however, showed that while there are significant fluctuations of velocity in the pilot region, the main area is mainly free of phase locked changes in the velocity field. This means that these secondary spray structures must be altered by the unsteady evaporation rate of the particles.

In order to investigate more fully the different mechanisms responsible for the fluctuations of heat release in the pilot and main a POD of the unsteady flow field with the POD inner product defined by the fuel spray SMD, fuel spray mixture fraction, mixture fraction, progress variable, heat release rate, the temperature, the density and the velocity.

Figure 8.32 shows the area averaged and normalised spray mixture fraction, heat release rate, temperature and mixture fraction components of a variety of different POD mode combinations over two forcing cycles. Figure 8.32a shows the combination of the first 30 modes which contain 28% of the fluctuating flow energy, there are clearly two peaks in each of the four components plotted showing that there is a significant influence of the forcing frequency on the behaviour of these modes. Figure 8.32b shows the area averaged and normalised heat release for the main and pilot flames by dividing the combustion chamber by radius into a pilot and main region. It shows that the pilot and main flames are in phase although the pilot flame shows a much higher level of unsteadiness. It can also be seen that the main heat release is very similar to the total heat release showing that the main flame is driving the response as expected. In figures 8.32c and 8.32d the same results are presented but for the first two POD modes only. These figures show that the heat release is driven primarily by the action of these two modes. The pilot and main flames are again, in phase and the pilot flame has lost the high frequency unsteadiness showing that these fluctuations are driven by different flow mechanisms than the ones that cause the unsteadiness at the forcing frequency. Finally, figures 8.32e and 8.32f show the same results for the 7th and 8th most energetic modes. These modes are highlighted because the pilot and main heat release is out of phase showing that there is a mechanism of heat release fluctuation that involves an exchange of mixture between the main and pilot. These oscillations are not occurring at the forcing frequency and the main is dominating the pilot however modes that alternate the heat release rate between

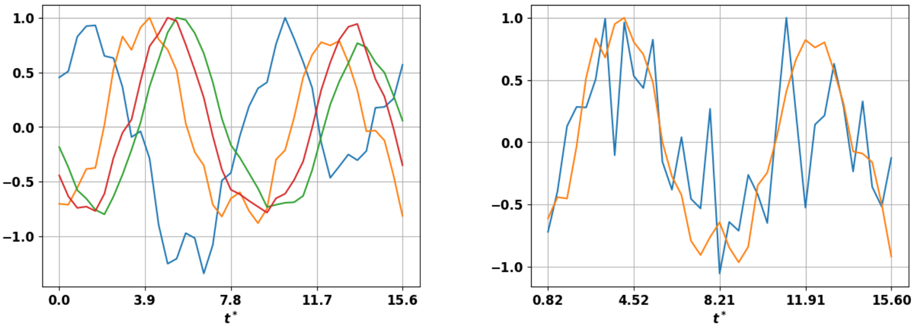

(a) The normalised spray mixture fraction (blue), mixture fraction (red), heat release (orange) and temperature (green) of the first 30 modes.

(b) The normalised heat release rate from the pilot flame (blue) and the main flame (orange) for the first 30 most energetic modes.

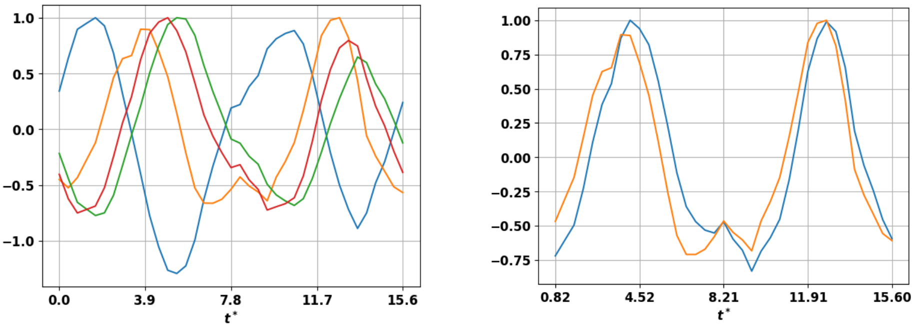

(c) The normalised spray mixture fraction (blue), mixture fraction (red), heat release (orange) and temperature (green) of the first two modes.

(d) The normalised heat release rate from the pilot flame (blue) and the main flame (orange) for the first two most energetic modes.

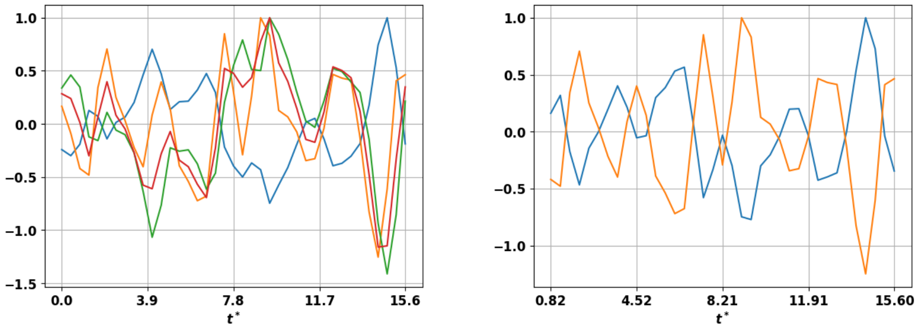

(e) The normalised spray mixture fraction (blue), mixture fraction (red), heat release (orange) and temperature (green) of the first 7th and 8th most energetic modes.

(f) The normalised heat release rate from the pilot flame (blue) and the main flame (orange) for the 7th and 8th most energetic modes.

Figure 8.32: The area averaged contributions of each collection of POD modes to spray mixture fraction, heat release rate, temperature and mixture fraction as well as the relative contributions made to the heat release rate fluctuations by the pilot and main flames.

main and pilot may be useful for dampening the overall system response.

The amplitudes of each component have been normalised to highlight the phase relationships between the different components. Removing velocity fluctuations from the possible driving mechanisms means that the fluctuations of heat release must be primarily driven by changes in mixture fraction. In order for mixture fraction to increase, there must be sufficient temperature and spray mixture fraction available to provide fuel to the reaction zone and evaporate it. It therefore follows that the expected phase relationship should be that a fluctuation in temperature and spray mixture fraction should result in a fluctuation of mixture fraction which should then precede the fluctuation of heat release which then creates a further fluctuation in temperature. What can be seen from Figure 8.32 is that this phase relationship depends on which POD modes are used to reconstruct the flow field and that in none of the mode combinations shown does the expected phase relationship hold exactly.

Firstly, Figure 8.32c shows that the spray mixture fraction leads the heat release rate which then leads the mixture fraction fluctuations. This delay implies that there is convection, evaporation and mixing between where the fuel spray is and where combustion occurs. As expected, the mixture fraction fluctuations are out of phase with the spray mixture fraction as the rate of fuel flow into the combustions chamber is constant. Because the heat release rate lies between the spray mixture fraction and mixture fraction fields, this implies that there is little perceivable delay between the appearance of the fuel spray and the heat release. In other words, when there is fuel present, it very quickly evaporates and burns without taking time to mix. This shows that the location of the fuel spray is a good way to change the location and rate of heat release, however this strategy would likely result in poor emissions performance as the flame is burning in localised regions of high fuel concentration. The in-phase behaviour of the mixture fraction and temperature shows that the temperature field inside the combustor is driven by the convection of hot reactant products, as well as fresh reaction. Regions where the mixture fraction is low and the spray mixture fraction is high results in high levels of evaporation and combustion. The fluctuations of heat release rate are therefore driven by areas of the combustor that are usually fuel lean, which, through whatever mechanisms, suddenly become fuel rich. This observation has an interesting consequence: If the combustion chamber is well mixed and the level of premixed combustion is high, then the concentration of fuel is therefore fairly constant throughout. In this case, small fluctuations of spray mixture fraction would not result in significant changes in the heat release rate because the combustion chamber mixture fraction would be close to equal and hence the evaporation rate of the fuel would be quite low. In this case, the largest contributions to fluctuations in heat release rate would be due to velocity fluctuations and the flame transfer function would converge towards that seen in premixed flames.

In the case of the out-of-phase 7th and 8th modes, the temperature field is driven mostly by mixture fraction fluctuations. This may be due to the choice of combustion modelling strategy as the temperature of the field is directly linked to the mixture fraction field through the FGM table. In this case it can be seen that the flame flicks between being premixed (high correlation with mixture fraction) and non-premixed (high correlation with spray mixture fraction).

The first and second modes' heat release rate, spray mixture fraction and azimuthal velocity components are shown in Figure 8.33 and Figure 8.34 respectively. These two modes contribute to

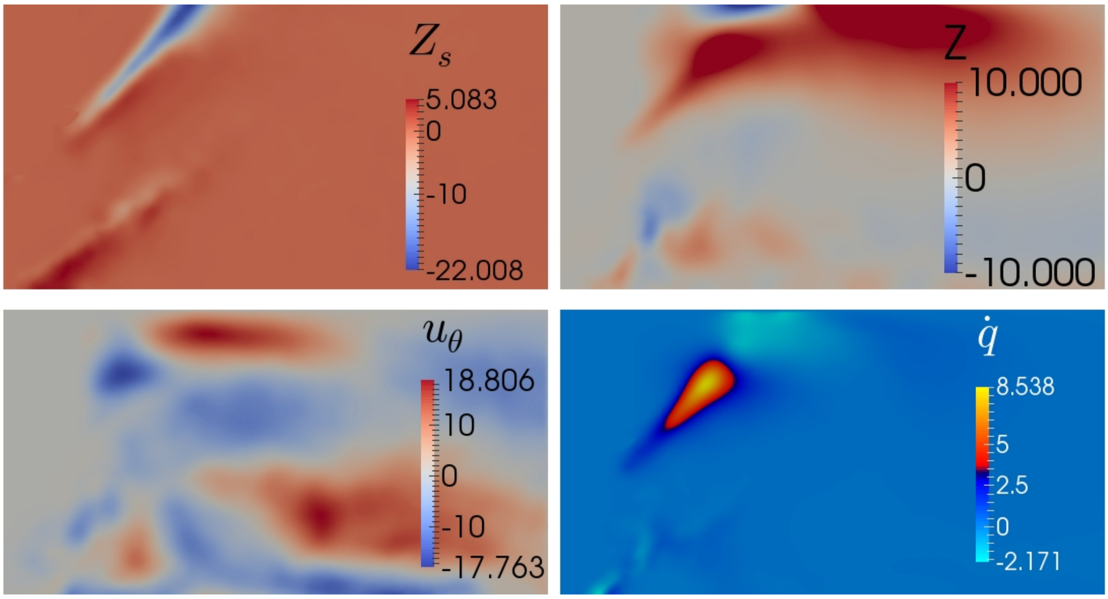

Figure 8.33: The azimuthally averaged spray mixture fraction, mixture fraction, heat release rate and temperature components for the most energetic POD mode.

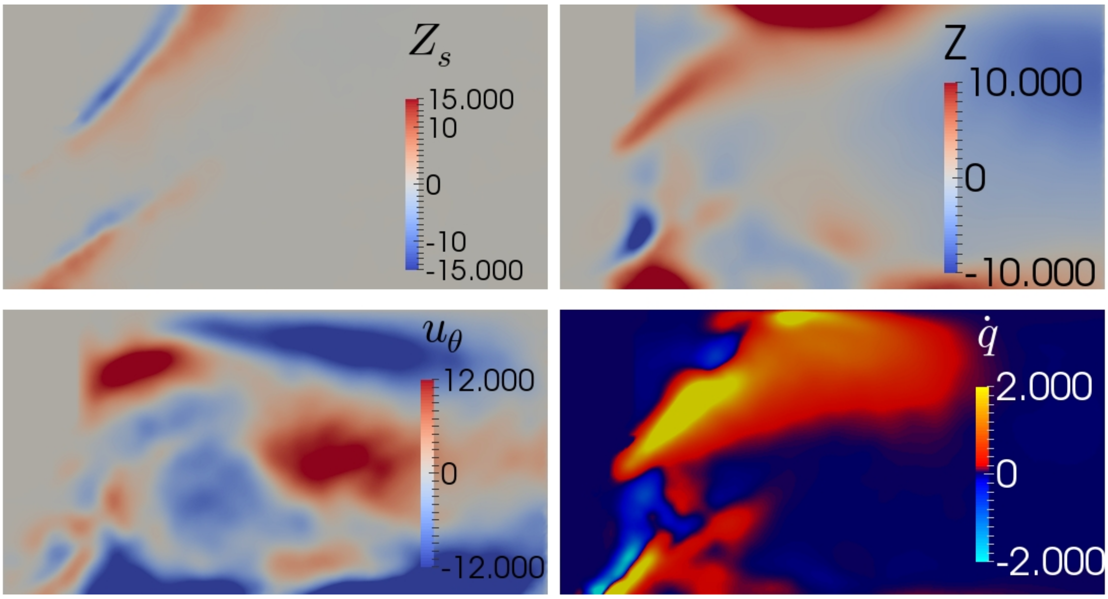

Figure 8.34: The phase and azimuthally averaged spray mixture fraction, mixture fraction, heat release rate and temperature components for the second most energetic POD mode.

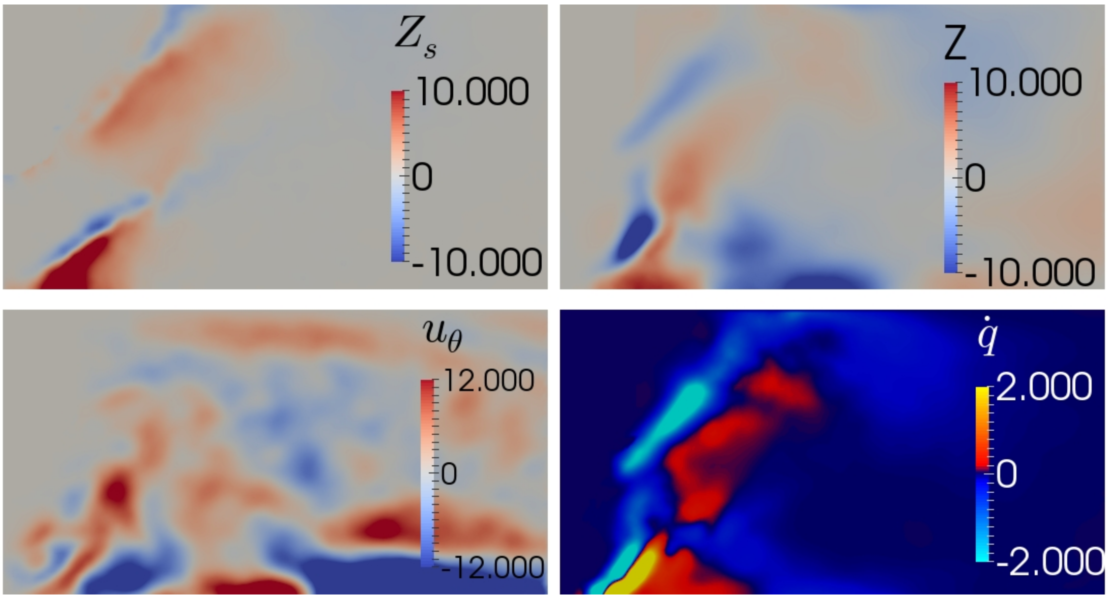

Figure 8.35: The phase and azimuthally averaged spray mixture fraction, mixture fraction, heat release rate and temperature components for the seventh most energetic POD mode.

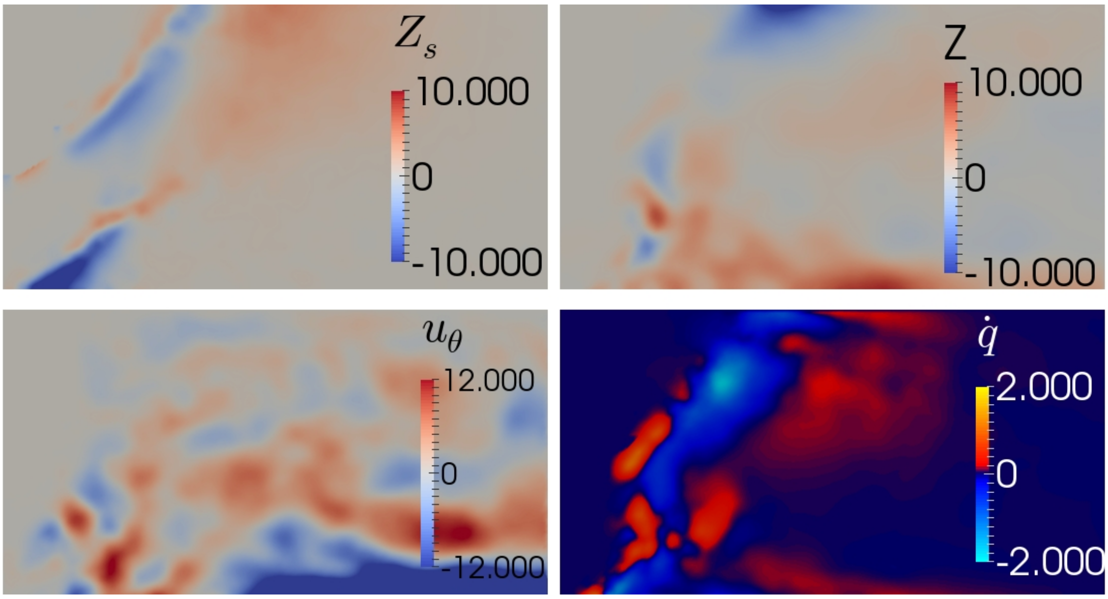

Figure 8.36: The phase and azimuthally averaged spray mixture fraction, mixture fraction, heat release rate and temperature components for the eighth most energetic POD mode.

around 4% of the fluctuations seen and as already mentioned, contribute to most of the in phase fluctuations of heat release. The first mode shows a strong negative peak of spray mass fraction near the wall where it hits the wall and starts to evaporate and burn. The positive peak corresponds to an earlier point in the cycle where the main flame is being augmented mid-way between the injection point and the wall by smaller droplets that have evaporated before the wall. The second most energetic mode shows the point where the heat release is maximum. Fresh mixture that has been created at the wall causes a heat fluctuation at the wall while at the same time the droplets at the atomiser edge are at their minimum diameter leading to fresh mixture being produced at the atomiser edge also. The flame grows in this case, both upstream and downstream.

Whilst the fluctuations in the main fuel spray are driven by the changing SMD, the pilot fuel spray is also highly unstable. The mechanism behind this pilot fuel response is driven by the changing azimuthal velocity. The fluctuating main heat release causes a corresponding fluctuation in density. As the flow heats up, the density reduces and the flow accelerates. As the flow accelerates the strength of the vortex at the centre of the domain grows stronger, increasing the strength of the recirculation zone. Droplets of fuel that may have escaped downstream are therefore drawn into the centre of the domain leading to a flame being forced along the centreline. At the same time, the increase in azimuthal velocity acts on the pilot fuel spray. As the spray accelerates, the centrifugal force on the droplets increases and the droplets are centrifuged outwards.

This relatively simple mechanism shows how the pilot flame spray may be altered by the modulation of the main SMD but says nothing about how the intermittent fluctuations of fuel spray mass fraction appear downstream of the main flame. Figures 8.35 and 8.36 show the 7th and 8th most energetic modes. These modes show a strong feedback mechanism from the main to the pilot with excess mixture fraction and spray from the main being pulled into the pilot flame by the recirculating flow. There is evidence to suggest that these pockets of high spray mixture fraction that end up bypassing the main flame are driven by small scale flow structures at the interface between passages A and B. This can be seen in the mixture fraction and azimuthal velocity components of the two modes. These flow fluctuations may be naturally unstable $m = 0$ modes or helical modes that are augmented by the forced heat release or flow fields. Furthermore it can be seen in Figure 8.36 that the out-of-phase response seen in Figure 8.32f is driven by an alternating fluctuation of mixture fraction between main and pilot. Some of the main fuel spray that escapes the flame ends up being drawn into the pilot zone where it evaporates and burns.

These interactions between the pilot and main are essential as the pilot must be able to relight the main in the case of moving from a low to high power condition. The recirculation zone is designed such that hot reaction products from the pilot move forward to the splitter plate and then outwards to the inside of the main jet before being pulled back down towards the pilot. This cycle has a characteristic time and characteristic times exist between a particularly hot fluctuation in the pilot then helping to ignite a unusually rich pocket in the main flame. These characteristic times are a function of the strength of the recirculation zone and the flame positions. At certain flame positions these characteristic times will help to re-enforce instabilities, while others may result in the pilot and main flames pulsating out of phase, reducing the sensitivity of the system at certain frequencies. The location of the flames in the combustion chamber is largely a function of the swirl number and investigations on the sensitivity of flame location and stability with respect to swirl number may provide an additional method for altering the response of the flame.

It is clear that the wall has an effect in producing local fluctuations of mixture fraction. Large droplets that hit the hot wall are assumed to evaporate instantly resulting in a periodic hot spot that moves up and down the combustion chamber wall. Moving the main spray atomiser inboard would allow more time for the spray to evaporate before it hits the wall, leading to smaller localised fluctuations of heat release.

Despite what was said earlier about the relative insensitivity of the fuel spray to mass and velocity fluctuations through the injector at higher operating pressures, it is clear from the results of this section that even small changes in the atomisation can have relatively large effects on the heat release of the flame. If it is considered that the total flame response is the sum of the aerodynamic response of the flame and any fuel spray related processes then:

$$\mathcal{F} = \mathcal{F}_{aero} + \mathcal{F}_{spray} \tag{8.20}$$

and the spray transfer function can be decomposed into a part relating the mass flow rate to the spray and a part that relates the spray to the heat release:

$$\mathcal{F}_{spray} = \mathcal{F}_{\dot{m}S}\mathcal{F}_{S\dot{q}} \tag{8.21}$$

From the results of this section the gain of the transfer function relating the spray to the heat release is in the vicinity of being equal to one which means that the transfer function relating the spray to the heat release is equal in magnitude to the response of the spray to the fluctuations of mass flow through the injector:

$$|\mathcal{F}_{spray}| = |\mathcal{F}_{\dot{m}S}| \tag{8.22}$$

Taking the example from Section 7.2.4, the 1.2% fluctuation in SMD corresponds to a fluctuation of less than one micron, however the fluctuation of mass flow rate through the injector is only 0.5% which means that the gain between the fluctuations of mass and the fluctuation of SMD is around 2.4. This shows that the fluctuations of mass flow rate induced by the incoming acoustic wave are an extremely efficient method of changing the SMD and therefore the heat release.

8.5 Lessons from pre-mixed combustion

The work of Candel et al. [16] and others has highlighted the two main mechanisms for heat release fluctuations in swirl stabilised premixed combustion. The first is due to the fluctuations of axial velocity that generate additional vorticity at the exit of the injector, which convect to the flame and alter the flame area, generating a fluctuation in heat release. The other mechanism is driven by a fluctuation of swirl number, this fluctuation alters the turbulent flame speed leading to a different flame angle. These two mechanisms can constructively of destructively interfere leading to a flame describing function with a gain that contains a characteristic ‘V’ shape. Whilst the lean burn injector studied has a much more complicated structure and the flame transfer function is also driven by other mixture fraction related processes, the flame transfer function and premixed flame describing function exhibit several similarities:

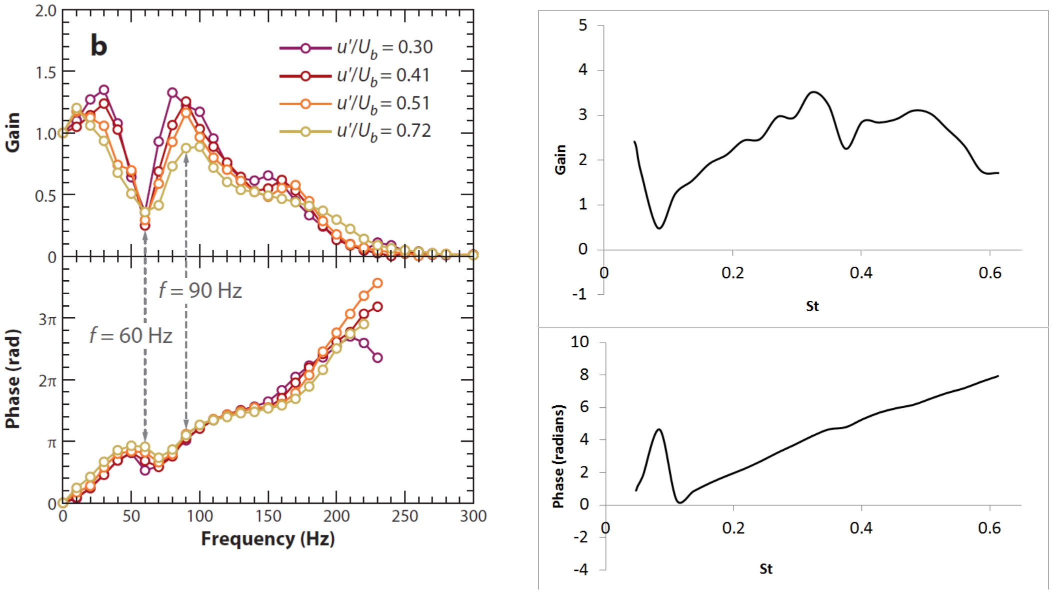

(a) The FDF for a premixed swirl stabilised flame (reproduced from [16]). U_b is the mean flow velocity measured by the HWA probe.

(b) The FTF for the lean burn fuel injector as measured experimentally on the SCARLET rig.

Figure 8.37: The FDF of a premixed, swirl stabilised flame and the FTF of the lean burn injector.

Figure 8.37 shows the FDF of the swirl stabilised lab scale combustor next to the FTF of the lean burn combustor. The gains of both the FDF and FTF show characteristic troughs at 60 Hz and $St = 0.08$ respectively, followed by peaks at 90 Hz and $St = 0.32$. A second trough can be seen 140 Hz and $St = 0.39$ which are close to the sum of the peak and trough frequencies for both flames (150 Hz and $St = 0.4$). The phase of the FDF shows a peak at 50 Hz and a trough at 70 Hz. The same features can be seen in the FTF with a peak at $St = 0.08$ and a trough at $St = 0.12$. The location of the peaks and troughs in the gain of the FDF have been linked to the constructive and destructive interference of the effects of swirl number and axial velocity fluctuations. The fact that several smaller peaks and troughs are apparent in the FTF is likely caused by the effects of multiple passages. These similarities point to similar mechanisms being present within each flame type.

Furthermore, the strength of the swirl number fluctuations and their phase relationship to the axial fluctuations have been shown to be caused by the location and response of the swirlers inside the injector passages for both the premixed flame and the lean burn injector. This means that there is a relatively simple analytical function linking the location of the swirl vanes to the phase and magnitude of the swirl number fluctuations and therefore to the location of the peaks and troughs in the flame transfer function. It should therefore be possible to calibrate this simple model to the results coming from experiment or numerical simulation. This would then allow a designer to test how moving the location or changing the response of the swirl vanes would lead to a change in the location of the peaks and troughs of the transfer function extremely quickly. The relationship between heat release, swirl number and axial fluctuations is discussed in more detail in Chapter 5.

8.6 Conclusion

This chapter has shown that while calculating the FTF for a practical lean-burn gas turbine injector still remains technically challenging, there exist a number of techniques that provide consistent results and that may be used to target particular elements of the problem. The predicted FTFs show that the gain calculation is weakly dependent on the quality of the mesh but may also be successfully computed using an incompressible method with PODFS inlets which significantly reduces the cost of doing high resolution simulations. The calculation of the phase of the FTF however is remarkably consistent with even very poor quality simulations being able to arrive at reasonable result. The main challenge for predicting the gain correctly appears to be the unknowns involving the fuel spray. It can be seen from additional simulations that fluctuations in the spray velocity or droplet size can significantly alter the shape of the flame. It is suggested that fuel spray measurements using the SCARLET or other rigs could be used to build a joint PDF between droplet size and velocity that can be used in future simulations at elevated temperatures and pressures. The relative location of the combustor walls has also been seen to be important variable and adjusting the atomiser location away from the walls may allow more of the fuel to evaporate before wall interaction can occur and hence reduce the sensitivity of the flame to SMD variations.

The stability and penetration of the pilot jet appears to be sensitive to both the mass flow rate and the turbulence level in the pilot jet. In the simulations that used the cut-down geometry and the incompressible method. Only those simulations that used the PODFS method to generate upstream turbulence predicted the correct penetration of the pilot jet.

There is also evidence to suggest that in the current design the flame location and strength is driven by the location of the fuel spray which is bad for both flame stability and emissions performance. Improving the atomisation strategy would lead to a more premixed flame response by the burner which would improve stability and emissions performance. This would also make prediction easier as fluctuations of the fuel spray would have less effect on the heat release rate of the flame.

It is also clear that azimuthal velocity fluctuations are an important driver of heat release rate fluctuations both through direct interaction with the flame and also due to interactions with the fuel spray. This makes it especially important that these velocity fluctuations are modelled and accurately captured. This highlights the importance of using either fully compressible simulations or the use of the PODFS method for generating the flow field upstream of the flame.

Finally it is suggested that breaking up the symmetry of the injector with chevrons could help to break up the large scale flapping of the main jet which appears to drive the fluctuations of heat release caused by acoustic excitation. Using POD analysis, small scale structures appear to be most efficient at altering the heat release rate which means that the addition of chevrons may inadvertently increase the sensitivity of the flame to acoustic forcing. Alternatively, chevrons may be added to the splitter plate between passages A and B in an attempt to breakup the formation of helical modes which are shown to augment the fluctuations of heat release produced by axisymmetrical structures and may also directly contribute to spinning modes in annular combustion chambers.

Chapter 9

Conclusion and future work

9.1 Conclusion

This thesis has achieved a number of important outcomes. In Chapter 3, various different methods for computation of the non-reacting single phase flow were compared. The results showed that unsteady URANS calculations gave broadly consistent results across different codes, turbulence models and between incompressible and compressible flow simulations which also agreed reasonably well with PIV measurements. The $k-\omega$ SST model results in lower levels of turbulent viscosity and represents velocity fluctuations better than the $k-\epsilon$ RNG model. Steady RANS was unable to capture the mean flow correctly even though it remains a popular method in industry for comparing injector designs. Large scale LES was also tested which shows broad agreement with the URANS simulations but with an increased cost.

The main advantage of the LES method is that it captures not just the large scale flow structures but also a reasonable proportion of the fluctuating turbulent field. How important this is to the injector design process is questionable as the turbulent kinetic energy is also approximated by the use of turbulence models. Analysis of the spatial field correlations for an LES simulation indicate that even meshes that are considered by conventional methods to be sufficiently well refined, may cause a non-physical build up of TKE at the mesh resolution which may then affect the prediction of heat release in the combustion chamber. This would indicate that meshes should be refined further, in some cases making routine LES simulation too expensive for industrial applications. LES is considered to be the best method for reacting flow simulations in industry and academia although meshes are often not refined to the level tested in this study. This should prompt a re-think about how mesh resolution is defined as being sufficient in industrial settings.

Chapter 4 focuses on the reasons why steady simulation does not capture the mean field accurately and sets up a further discussion on the generation of large scale structures in the combustion chamber such as the PVC. Steady simulation cannot predict the formation of large scale structures that promote mixing and modify the mean field in a way that cannot be captured using conventional turbulence models. Large scale structures have been observed to affect the heat release rate in combustion chambers although in the case of non-axisymmetric unstable modes, they cannot contribute directly to longitudinal thermoacoustic instabilities in combustion chambers. POD is described as a method for extracting these large scale structures from the unsteady flow field in an efficient manner. It can be seen from analysing the temporal evolution of these structures that they act in competition with each other, with modes of one azimuthal wavenumber exchanging energy with modes of another azimuthal wave number. The appearance of these large scale unsteady structures is driven by the mean flow velocity and density gradients and linear stability analysis is introduced as a method for predicting these structures from the mean field. It is seen however that the prediction of these structures is not straightforward.

The mechanisms driving heat release in premixed flames are driven by the interaction between acoustic waves and the fuel injector. Chapter 5 describes in great detail the methods used for simulating and analysing this interaction. The acoustic impedance and acoustic transfer matrix of the fuel injectors are computed using compressible URANS simulations and compared to experiment. The flow fields are analysed in the case of one forced simulation to show the relative fluctuations of pressure, axial velocity, azimuthal velocity and swirl number through the injector passages. The axial velocity and pressure fluctuations which are associated with the acoustic impedance were shown to average across the injector passages while the azimuthal velocity and swirl number were seen to sum. This is important due to the importance of swirl number fluctuations to the transient response of premixed and swirl stabilised flames. The magnitude of the swirl number response was also seen to be dependent on the relative distance between the passages exits and the swirl vanes. The flow fields of two simulations were also analysed using POD to show that the acoustic forcing can modify the strength of the large scale structures described in Chapter 4. This interaction may be important given the importance of these structures to the heat release rate. Several interaction mechanisms between acoustics and hydrodynamics are described in the literature and an analysis of the applicability of each is also included in this chapter.

Because the interaction of acoustic waves with the injector is localised to the region around the injector body and because it is considered that the acoustic waves have negligible direct interaction with the flame, an incompressible methodology can be considered for forcing the flame. The simplest method is to simply force the upstream mass flow rate at the desired frequency. In Chapter 6, comparisons of the flow field between this method and the case where the flow field is acoustically forced using a compressible method shows that the downstream flow field evolution is markedly different. Simulations that ‘cut’ velocity information from a compressible simulation and then ‘paste’ that information onto the inlets of a cut-down incompressible domain show that the correct evolution of mass flow rate and swirl number can be captured using an incompressible methodology. The cut-down geometry includes the domain downstream of each flow passage and as such, the three flow passage outlets, must be updated at each timestep. This constant loading and application of the data is expensive and limits the efficiency that is gained from using a much smaller computational domain. Use of the PODFS method improves this computational efficiency greatly with only a small penalty in accuracy.

Whilst the fuel injector is designed to enhance the mixing of fuel and air, the use of a liquid fuel is likely to lead to significant mixture fraction related flame fluctuations unlike a premixed flame. Chapter 6 describes how the atomisation of the fuel may be altered due to the acoustic forcing while simulations are run to show how large scale flow structures and the atomised fuel spray may interact inside the combustion chamber. The results show that the response of fuel droplets of a size typically found inside combustion chambers is larger than droplets an order of magnitude larger or smaller in size. However, the difference in response is probably not sufficient to justify the large cost associated with changing the atomisation strategy. It is suggested that it may instead be sensible to somehow protect the atomiser from these large scale structures. It was also shown that the different axial locations of the exits of the passages on either side of the main atomiser generates intermittent vorticity that then causes a breakup of the spray. It is hypothesised that such a break up will lead to additional fluctuations of mixture fraction and heat release in the combustion chamber.

Finally Chapter 8 deals with how the other topics studied in each previous chapter can affect the heat release rate. Firstly, the effects that the addition of combustion has on the mean flow field are discussed. The addition of heat release causes a sudden decrease in density and an increase in flow velocity leading to longer jets as compared to the single phase case. Additionally the density field is noticeably different as compared to the cold case in the region of the splitter plate, previously identified as the location where the large scale structures are formed. This results in less fluctuations of velocity upstream of the reacting zone and suggests that capturing the mean density field will be important in any future attempts at predicting the appearance of large scale structures. The affects that spray atomisation velocity can have on the flame are also shown. The PODFS method is also shown to be capable of capturing the fluctuating flow produced through the interaction between the fuel injector and the acoustic field. In URANS simulations, this information is sufficient to reproduce the flow in a truncated domain however for LES additional turbulence must be introduced at the inlets to ensure the correct penetration length of the pilot jet. This was achieved in Chapter 8 by constructing additional PODFS modes from the non-reacting LES simulation described in Chapter 3.

The FTF for the chosen injector was found at two frequencies using various different methods with the incompressible PODFS method found to give an FTF that is in agreement with the compressible simulation results using a significantly less expensive simulation method. The results suggest that direct computation of the FTF by measuring the heat release and mass flow rates directly from the simulation causes a slight under-prediction of the FTF gain as compared to the measurement using the multi-microphone method, as used in experiments. Upstream forcing appears to give a much weaker response to downstream forcing with a noticeably different phase response. Additional error seen between experiments and simulations are hypothesised as being caused by unaccounted fluctuations of the atomisation of the fuel spray. This hypothesis was then tested by running a simulation where the spray SMD was varied by 20% while the inlet flow was assumed not to fluctuate. This simulation resulted in heat release fluctuations of around 20% showing that the effects of acoustic forcing on atomisation are an extremely strong mechanism for altering the heat release rate of the flame. These results show that through improving the fuel mixing, the injector response gain is likely to reduce. This is also consistent with other combustion system design aims such as reducing the formation of soot.

9.2 Injector design recommendations

9.2.1 Lessons from pre-mixed combustion

As highlighted in section 8.5 the characteristic 'V' shape observed in the gain of the FTF is most likely due to swirl number and axial velocity fluctuations as described in Candel et al. [16]. The depth and location of this 'V' is therefore dependent on the response of the swirl vanes to axial velocity fluctuations and the location of the swirl vanes inside the injector passages. Analysis of flame transfer functions and velocity fields from either simulations or experiments could be used to build simple models that should be able to predict how changing the response or location of a set of swirl vanes will alter the FTF of the fuel injector.

9.2.2 Cleaning up the geometry

In Chapter 5 the *acoustic shift mode* is calculated for the lean burn injector forced at 300 Hz. This mode describes the alteration of the mean flow field due to the acoustic forcing. The largest changes in the mean field occurred in the region within passage B between the fuel feed arm and swirl vanes. The phase averaged velocity field shows an interaction between the acoustic field and the separated flow behind the fuel arm in this region. This interaction at some frequencies may in fact be beneficial to the acoustic response whilst at other frequencies it may cause problems. The choice of fuel arm profile is probably made on the basis of cost and that choice may in time be proven to be justified, however this obviously bad aerodynamic design makes the analysis of instabilities very difficult. It would be better in the future to focus on aerodynamically clean designs for the analysis of the mechanisms leading to a high heat release response. This is achieved somewhat through the use of the TOS injector that does not have a fuel feed arm.

The TOS injector does however contain the same swirl vanes as the fuel injector with the fuel arm fitted. These vanes are manufactured from sections of flat metal plate that are prone to developing separations. These separations have been shown to be beneficial to the generation of downstream turbulence, leading to better mixing and lower soot emissions [48]. Once again the design of the swirl vanes is likely driven by cost, and may be justified given the effects on soot production but from the point of view of understanding the mechanisms behind heat release fluctuations, the geometry makes the analysis difficult.

9.2.3 Reduction of intermittent vorticity

The intermittent vorticity that is identified as being responsible for the breakup of the spray in Chapter 7 is caused by the different axial locations of the exits of passages B and C which sit either side of the fuel atomiser. In order to reduce the intensity of this vorticity the exit of these two passages should be aligned. Practically this requires that the injector face be flatter instead of concave which will have ramifications for the shape of the mains and pilot jet which would have to be assessed separately. Furthermore it would also require a similar axial location of the swirl vanes for these passages which would restrict one of the other parameters that could be adjusted to alter the thermoacoustic response of the injector. The aligning of the injector passages exit plane does not necessarily require the injector passages to be of equal length as they could be extended in the opposite direction.

9.2.4 Wall-spray interaction

The strong response seen in Chapter 8 of the flame to SMD fluctuations is driven by the interaction of the fuel spray and the wall. This makes things difficult as the wall-spray interactions are tricky to model correctly. One possible solution to the is problem is to remove the interaction by moving the prefilmer radially inwards. Practically this means that passage C will get larger and passage B smaller, affecting the flow field, however the main motivation for placing the atomiser edge outboard appears to be to keep the mains and pilot flames separated. Due to the fuel being centrifugally spun outwards by the swirling flow, this would still be achieved with an atomiser much closer to the centreline and would give the spray more time to evaporate before it hits the wall or approaches the flame front. The observation that the spray interacts with the wall is based on simulations with assumed fuel spray properties. High velocity, large droplets would likely avoid the wall while smaller droplets will be sucked into the ORZ. Better measurements of the fuel spray at elevated temperatures and pressures would enable more accurate simulations to better test the spray wall interactions seen in this study. Other strategies to improve evaporation and mixing would also likely improve the stability of the flame.

9.2.5 A less sensitive design?

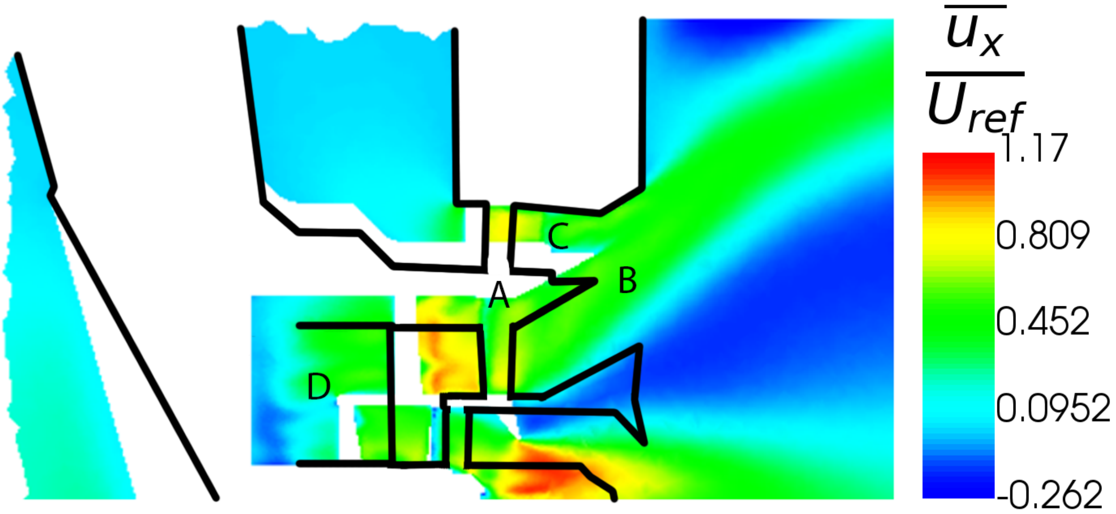

Figure 9.1: A new injector with the suggested design modifications drawn in black over the top of the current injector design.

A new injector design incorporating all of the suggested design features in shown in figure 9.1. The new design includes: A: Swirl vanes located such that the desired swirl number fluctuations are achieved, with the outer two flow passage swirl vanes aligned. B: The atomiser is now sheltered from helical instabilities coming from the splitter plate between passages A and B. C: Passage C has been enlarged, pushing the atomiser closer to the centreline and D: The fuel feed arm has been elongated and smoothed to prevent separations that interact with the swirl vanes and the acoustic field.

9.3 Future work

Ultimately, the calculation of the FTF for arbitrary fuel injector geometries is likely to remain too expensive for practical use in the foreseeable future. A rough estimate of the cost per FTF (based on 24 frequencies) is in the range of $0.5 - 3.0$ million Euros depending on the method used. This cost may be reduced by forcing the simulations at multiple frequencies simultaneously however the applicability of such methods is likely dependent on the injector design. Despite this cost simulations that are aimed at identifying particular instability mechanisms of specific designs will prove to be useful in the years ahead. Future work should therefore focus on ways that the design can be evaluated quickly to provide a *ballpark* estimate of the injector response which can then be tested and improved using more expensive simulations and experimental methods.

Firstly, the acoustic field is linear which opens the possibility of using linearised methods to capture the fluctuating acoustic field inside the combustion chamber. Such methods can be solved in Fourier space which significantly improves the calculation time as compared to URANS simulations. The problem is that the acoustic response of the injector involves a complex non-linear interaction between the acoustic waves, large scale flow structures and the flow turbulence which typically requires a transient and non-linear methodology. However it is not impossible, through the use of analytical models and data assimilation techniques, that the transient response can be extrapolated from the mean flow and linearised acoustic fields. The interaction between the growth of large scale structures and the acoustic field may, for instance, be predictable using Hemchandra's method (see Section 5.6.3). Luckily, non-reacting URANS simulations aimed at resolving the injector acoustic response do not require a highly refined mesh and may be forced at multiple frequencies simultaneously, making them relatively inexpensive.

Unfortunately, the use of linearised methods is predicated on the correct mean field prediction which currently requires a transient solution method. The reason RANS simulations are unable to correctly capture the mean field is because large scale structures that promote mixing are not accounted for in the turbulence modelling. Turbulence models are built and calibrated against certain test cases that are chosen to be universal, making them safe to use in a broad range of areas. Recently however, a new class of specialised turbulence models have been developed using data assimilation to better predict specific flows. Turbulence models could therefore be developed that better predict the reacting mean flow field in swirl stabilised combustors which would reduce the need to run expensive URANS or LES simulations.

The use of Hemchandra's method also requires the global adjoint modes which are theoretically obtainable from the mean flow field but were not found in this thesis due to the time and expertise required to calculate them in this complex industrial geometry. Better and more dependable methods of calculating the flow stability will lead to a better understanding of how particular designs will promote and dampen large scale structures that have been shown to interact with the fuel spray and affect heat release without requiring transient calculations.

Whilst it would be useful to further validate the methods tested here, the largest unknown in calculating the FTF accurately is the response of the fuel spray atomisation under acoustically forced conditions at elevated pressures and temperatures. It would therefore be extremely helpful if the fuel spray droplet diameters and velocities were known at the injection point as a function of acoustic phase angle. Future experiments may be able to provide this information.

In the future it is hoped however that the fuel injector design will tend further towards a premixed approach where the fuel spray has less influence on the transient response of the flame and the production of emissions are reduced. In this case the FTF will be very similar to a premixed swirl stabilised response where the use of simple analytical models combined with data assimilation from experiments should give a good prediction of how the FTF will change as a function of simple control variables such as the location and response of swirl vanes, mean swirl number and mass split between passages.

References

[1] Abo-Serie, E., Carrote, J., and Spencer, A. (2007). The impact of acoustic excitation on jet mixing. Technical Report TT07R05, Department of Aeronautical and Automotive Engineering Loughborough University, Loughborough UK.

[2] Al-Abdeli, Y. M., Masri, A. R., Marquez, G. R., and Starner, S. H. (2006). Time-varying behaviour of turbulent swirling nonpremixed flames. *Combustion and Flame*, 146(1-2):200–214.

[3] Alne, K. S. (2007). Reduction of NOx emissions from the gas turbines for Skarv Idun. Master's thesis, Institutt for energi-og prosessteknikk.

[4] Anand, M., Eggels, R., Staufer, M., Zedda, M., and Zhu, J. (2013). An advanced unstructured-grid finite-volume design system for gas turbine combustion analysis. In *ASME 2013 Gas turbine India conference*, pages V001T03A003–V001T03A003. American Society of Mechanical Engineers.

[5] Aubry, N., Holmes, P., Lumley, J. L., and Stone, E. (1988). The dynamics of coherent structures in the wall region of a turbulent boundary layer. *Journal of Fluid Mechanics*, 192:115–173.

[6] Aurifeille, E. (2014). Precise-UNS best practice for the isothermal CFD modelling of lean burn injectors. *Rolls-Royce Technical Report*, DNS198825.

[7] Barker, A., Carrotte, J., and Denman, P. (2005). Analysis of hot-wire anemometry data in an acoustically excited turbulent flow field. *Experiments in Fluids*, 39(6):1061–1070.

[8] Batchelor, G. (1967). *An introduction to fluid dynamics.* Cambridge University press.

[9] Bechmann, A. and Sørensen, N. N. (2010). Hybrid RANS/LES method for wind flow over complex terrain. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 13(1):36–50.

[10] Beér, J. M. and Chigier, N. A. (1972). Combustion Aerodynamics. *New York.*

[11] Bellucci, V., Paschereit, C. O., Flohr, P., and Magni, F. (2001). On the use of Helmholtz resonators for damping acoustic pulsations in industrial gas turbines. In *ASME Turbo Expo 2001*, pages V002T02A006–V002T02A006. American Society of Mechanical Engineers.

[12] Bennett, B. A. V. and Smooke, M. D. (1998). Local rectangular refinement with application to axisymmetric laminar flames. *Combustion Theory and Modelling*, 2(3):221–258.

[13] Bergmann, M. and Cordier, L. (2008). Optimal control of the cylinder wake in the laminar regime by trust-region methods and POD reduced-order models. *Journal of Computational Physics*, 227(16):7813–7840.

[14] Boxx, I., Stöhr, M., Carter, C., and Meier, W. (2010). Temporally resolved planar measurements of transient phenomena in a partially pre-mixed swirl flame in a gas turbine model combustor. *Combustion and Flame*, 157(8):1510–1525.

[15] Campa, G. and Camporeale, S. M. (2010). Application of transfer matrix method in acoustics. In *European COMSOL Conference*.

[16] Candel, S., Durox, D., Schuller, T., Bourgouin, J.-F., and Moeck, J. P. (2014). Dynamics of swirling flames. *Annual Review of Fluid Mechanics*, 46:147–173.

[17] Caraeni, M., Devaki, R., Aroni, M., Oswald, M., and Caraeni, D. (2009). Efficient acoustic modal analysis for industrial CFD. In *47th AIAA Aerospace Sciences Meeting*.

[18] Carrotte, J. and Barker, A. (2012). Two phase flow and acoustic excitation. *Loughborough University Technology Centre Report*, TT12R12.

[19] Cassell, M. A., Regunath, G. S., and Carrotte, J. F. (2017). Capturing fuel system unsteady flow behaviour using 1D/3D coupled methods. *Loughborough University Technology Centre Report*, TTUTC17R16.

[20] Chen, H. and Patel, V. (1988). Near-wall turbulence models for complex flows including separation. *AIAA journal*, 26(6):641–648.

[21] Chen, K. K., Tu, J. H., and Rowley, C. W. (2012). Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses. *Journal of Nonlinear Science*, 22(6):887–915.

[22] Chomaz, J.-M. (2005). Global instabilities in spatially developing flows: non-normality and nonlinearity. *Annual Review of Fluid Mechanics*, 37:357–392.

[23] Chomaz, J.-M., Huerre, P., and Redekopp, L. G. (1991). A frequency selection criterion in spatially developing flows. *Studies in applied mathematics*, 84(2):119–144.

[24] Chong, L. T. W., Komarek, T., Kaess, R., Foller, S., and Polifke, W. (2010). Identification of flame transfer functions from LES of a premixed swirl burner. In *ASME Turbo expo 2010*, pages 623–635. American Society of Mechanical Engineers.

[25] Chung, Y. M. and Sung, H. J. (1997). Comparative study of inflow conditions for spatially evolving simulation. *AIAA Journal*, 35(2):269–274.

[26] Crocco, L. and Cheng, S.-I. (1956). Theory of combustion instability in liquid propellant rocket motors. Technical report, Princeton University NJ.

[27] Cummings, A. (1984). Acoustic nonlinearities and power losses at orifices. *AIAA Journal*, 22(6):786–792.

[28] Cyclone Fluid Dynamics (2019). Dolfyn, open-source cfd. https://www.dolfyn.net/.

[29] de Goey, L. and de Lange, H. (1994). Flame cooling by a burner wall. *International Journal of Heat and Mass Transfer*, 37(4):635–646.

[30] De la Cruz Garcia, M., Mastorakos, E., and Dowling, A. (2009). Investigations on the self-excited oscillations in a kerosene spray flame. *Combustion and Flame*, 156(2):374–384.

[31] de Lange, H. and de Goey, L. (1993). Two-dimensional methane/air flame. *Combustion Science and Technology*, 92(4-6):423–427.

[32] Desjardins, O., Pitsch, H., et al. (2010). Detailed numerical investigation of turbulent atomization of liquid jets. *Atomization and Sprays*, 20(4):311.

[33] Dianat, M., McGuirk, J. J., Fokeer, S., and Spencer, A. (2014). LES of unsteady vortex aerodynamics in complex geometry gas-turbine fuel injectors. *ETMM10: 10th International ERCOFTAC Symposium on Engineering Turbulence Modelling and Measurements.*

[34] Dickson, N. and Bureau, I. A. T. (2014). Local air quality and ICAO engine emissions standards. *Action Plan pn Emmission, Reduction, Kenya.*

[35] Dowling, A. P. and Stow, S. R. (2003). Acoustic analysis of gas turbine combustors. *Journal of Propulsion and Power*, 19(5):751–764.

[36] Dunham, D., Spencer, A., McGuirk, J. J., and Dianat, M. (2008). Comparison of URANS and LES CFD methodologies for air swirl fuel injectors. In *ASME Turbo expo 2008: Power for land, sea and air*, pages 187–196. American Society of Mechanical Engineers.

[37] Eggels, R. (2018). LES simulation of combustor including NGV's. *Internal Rolls-Royce Deutschland technical presentation.*

[38] Febrer Alles, G. (2012). *A hybrid approach for inclusion of acoustic wave effects in incompressible LES of reacting flows.* PhD thesis, Loughborough University.

[39] Ferziger, J. H. and Peric, M. (2012). *Computational methods for fluid dynamics.* Springer Science & Business Media.

[40] Gallimore, S., Vickers, R. M., and Boyns, M. B. (1997). The design modifications of the EGT Tornado industrial gas turbine to incorporate a dry low emissions combustion system. In *ASME 1997 International Gas Turbine and Aeroengine Congress and Exhibition*, pages V003T11A002–V003T11A002. American Society of Mechanical Engineers.

[41] Garnier, E., Adams, N., and Sagaut, P. (2009). *Large Eddy Simulation for compressible flows.* Springer Science & Business Media.

[42] George, W. K. (2013). *Lectures in Turbulence for the 21st Century.* Turbulence Research Lab: www.turbulence-online.com.

[43] Giauque, A., Selle, L., Gicquel, L., Poinsot, T., Buechner, H., Kaufmann, P., and Krebs, W. (2005). System identification of a large-scale swirled partially premixed combustor using LES and measurements. *Journal of turbulence*, (6):N21.

[44] Grinstein, F. and Guirguis, R. (1992). Effective viscosity in the simulation of spatially evolving shear flows with monotonic FCT models. *Journal of Computational Physics*, 101(1):165–175.

[45] Gunasekaran, B. and McGuirk, J. (2011). Mildly compressible pressure-based CFD methodology for acoustic propagation and absorption prediction. In *ASME 2011 Turbo Expo*, pages 355–365. American Society of Mechanical Engineers.

[46] Hamming, R. W. (1998). *Digital filters.* Courier Corporation.

[47] Harten, A. (1983). High resolution schemes for hyperbolic conservation laws. *Journal of Computational Physics*, 49(3):357–393.

[48] Hazenberg, T. and Eggels, R. (2018). Improvements in rich burn fuel injectors. *Rolls-Royce Deutschland technical report.*

[49] Hermanns, R. (2001). Chem1D, a one-dimensional laminar flame code. *Report, Eindhoven University of Technology.*

[50] Hirsch, H. (1990). Numerical computation of internal and external flows. *Computational Methods for Inviscid and Viscous Flows*, 2:536–556.

[51] Ishizuka, S. (2002). Flame propagation along a vortex axis. *Progress in Energy and Combustion Science*, 28(6):477–542.

[52] Israel, D. (2006). URANS and VLES: using conventional RANS models for time dependent flows. In *36th AIAA Fluid Dynamics Conference and Exhibit*, page 3908.

[53] Issa, R. I. (1986). Solution of the implicitly discretised fluid flow equations by operator-splitting. *Journal of Computational Physics*, 62(1):40–65.

[54] Iudiciani, P. and Duwig, C. (2011). Large eddy simulation of the sensitivity of vortex breakdown and flame stabilisation to axial forcing. *Flow, Turbulence and Combustion*, 86(3-4):639–666.

[55] Jakobsen, P. (2013). Introduction to the method of multiple scales. *arXiv preprint arXiv:1312.3651.*

[56] Janicka, J. and Kollmann, W. (1978). Ein rechenmodell für reagierende turbulente scherströmungen im chemischen nichtgleichgewicht. *Wärme-und Stoffübertragung*, 11(3):157–174.

[57] Jasuja, A. (1979). Atomization of crude and residual fuel oils. *Journal of Engineering for Gas Turbines and Power*, 101(2):250–258.

[58] Juniper, M. P. (2012). Absolute and convective instability in gas turbine fuel injectors. In *ASME Turbo Expo 2012: Turbine Technical Conference and Exposition*, pages 189–198. American Society of Mechanical Engineers.

[59] Kinsler, L. E., Frey, A. R., Coppens, A. B., and Sanders, J. V. (2000). *Fundamentals of acoustics.* Wiley.

[60] Kirthy, S. K., Hemchandra, S., Hong, S., Shanbhogue, S., and Ghoniem, A. F. (2016). Role of shear layer instability in driving pressure oscillations in a backward facing step combustor. In *ASME Turbo Expo 2016*, pages V04BT04A010–V04BT04A010. American Society of Mechanical Engineers.

[61] Klein, M., Sadiki, A., and Janicka, J. (2003). A digital filter based generation of inflow data for spatially developing direct numerical or large eddy simulations. *Journal of Computational Physics*, 186(2):652–665.

[62] Kolmogorov, A. N. (1941). The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. In *Dokl. Akad. Nauk SSSR*, volume 30, pages 299–303.

[63] Kosović, B. (1997). Subgrid-scale modelling for the large-eddy simulation of high-reynolds-number boundary layers. *Journal of Fluid Mechanics*, 336:151–182.

[64] Krediet, H., Beck, C., Krebs, W., Schimek, S., Paschereit, C., and Kok, J. (2012). Identification of the flame describing function of a premixed swirl flame from LES. *Combustion Science and Technology*, 184(7-8):888–900.

[65] Kuenne, G., Ketelheun, A., and Janicka, J. (2011). LES modeling of premixed combustion using a thickened flame approach coupled with FGM tabulated chemistry. *Combustion and Flame*, 158(9):1750–1767.

[66] Lacarelle, A., Faustmann, T., Greenblatt, D., Paschereit, C., Lehmann, O., Luchtenburg, D., and Noack, B. (2009). Spatiotemporal characterization of a conical swirler flow field under strong forcing. *Journal of Engineering for Gas Turbines and Power*, 131(3):031504.

[67] Launder, B. E. and Spalding, D. B. (1983). The numerical computation of turbulent flows. In *Numerical Prediction of Flow, Heat Transfer, Turbulence and Combustion*, pages 96–116. Elsevier.

[68] Laurence, D. (2002). Turbulence modelling. In *Proc. second QNET-CFD Workshop, Luzern, Switzerland,*.

[69] Laval, J.-P. (2014). Private Communication, Laboratoire de Mecanique de Lille. We discussed his PhD defense where Pierre Sagaut made the comment that his PhD work to improve the performance of LES models was possibly better directed at improving the the cell count on the basis of the computational cost.

[70] Lee, S., Lele, S. K., and Moin, P. (1993). Direct numerical simulation of isotropic turbulence interacting with a weak shock wave. *Journal of Fluid Mechanics*, 251:533–562.

[71] Lefebvre, A. (1988). *Atomization and Sprays*, volume 1040. CRC press.

[72] Lefebvre, A. and Chin, J. (1983). Steady-state evaporation characteristics of hydrocarbon fuel drops. *AIAA Journal*, 21(10):1437–1443.

[73] Lefebvre, A. H. (2010). *Gas turbine combustion.* CRC Press.

[74] Lieuwen, T. C. and Yang, V. (2005). *Combustion instabilities in gas turbine engines: operational experience, fundamental mechanisms, and modeling.* American Institute of Aeronautics and Astronautics.

[75] Lighthill, M. J. (1952). On sound generated aerodynamically I. General theory. *Proc. R. Soc. Lond. A*, 211(1107):564–587.

[76] Ljung, L. (1999). System identification. *Wiley Encyclopedia of Electrical and Electronics Engineering*, pages 1–19.

[77] Lumley, J. L. (1967). The structure of inhomogeneous turbulent flows. *Atmospheric Turbulence and Radio Wave propagation*, pages 166–178.

[78] Luong, T., Howe, M., and McGowan, R. (2005). On the Rayleigh conductivity of a bias-flow aperture. *Journal of Fluids and Structures*, 21(8):769–778.

[79] Ma, F., Proscia, W., Ivanov, V., and Montanari, F. (2015). Large eddy simulation of self-excited combustion dynamics in a bluff-body combustor. In *51st AIAA/SAE/ASEE Joint Propulsion Conference*, page 3968.

[80] Mallens, R., de Lange, H., van de Ven, C., and de Goey, L. (1995). Modeling of confined and unconfined laminar premixed flames on slit and tube burners. *Combustion Science and Technology*, 107(4-6):387–401.

[81] Manceau, R. (2007). Modlisation statistique de la turbulence. http://remimanceau.gforge.inria.fr/Publis/PDF/Manceau2007.pdf.

[82] Manoharan, K., Hansford, S., O‘Connor, J., and Hemchandra, S. (2015). Instability mechanism in a swirl flow combustor: precession of vortex core and influence of density gradient. *Proceedings of ASME Turbo Expo, GT2015–42985.*

[83] Menon, S. (2005). Acoustic-vortex-flame interactions in gas turbines. *Combustion Instabilities in Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, and Modeling*, page 277.

[84] Menter, F. and Egorov, Y. (2005). A scale adaptive simulation model using two-equation models. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, page 1095.

[85] Menter, F., Kuntz, M., and Langtry, R. (2003). Ten years of industrial experience with the SST turbulence model. *Turbulence, heat and mass transfer*, 4(1):625–632.

[86] Meyer, T. R., Roy, S., Gogineni, S. P., Belovich, V. M., Corporan, E., and Gord, J. R. (2004). OH PLIF and soot volume fraction imaging in the reaction zone of a liquid-fueled model gas-turbine combustor. In *ASME Turbo Expo 2004: Power for Land, Sea, and Air*, pages 827–836. American Society of Mechanical Engineers.

[87] Meyers, J., Sagaut, P., and Geurts, B. J. (2006). Optimal model parameters for multi-objective large-eddy simulations. *Physics of Fluids*, 18(9):095103.

[88] Mittal, V. and Pitsch, H. (2013). A flamelet model for premixed combustion under variable pressure conditions. *Proceedings of the Combustion Institute*, 34(2):2995–3003.

[89] Moeck, J. P. (2010). Analysis, modeling, and control of thermoacoustic instabilities. *Doctoral thesis, Technische Universität Berlin, Berlin, Germany.*

[90] Moeck, J. P., Bourgouin, J.-F., Durox, D., Schuller, T., and Candel, S. (2012). Nonlinear interaction between a precessing vortex core and acoustic oscillations in a turbulent swirling flame. *Combustion and Flame*, 159(8):2650–2668.

[91] Moeck, J. P., Bourgouin, J.-F., Durox, D., Schuller, T., and Candel, S. (2013). Tomographic reconstruction of heat release rate perturbations induced by helical modes in turbulent swirl flames. *Experiments in Fluids*, 54(4):1498.

[92] Moeng, C.-H. and Wyngaard, J. C. (1988). Spectral analysis of large-eddy simulations of the convective boundary layer. *Journal of the Atmospheric Sciences*, 45(23):3573–3587.

[93] Mongia, H., Held, T., Hsiao, G., and Pandalai, R. (2005). Incorporation of combustion instability issues into design process: GE aeroderivative and aero engines experience. *Progress in Astronautics and Aeronautics*, 210:43.

[94] Monkewitz, P. A., Huerre, P., and Chomaz, J.-M. (1993). Global linear stability analysis of weakly non-parallel shear flows. *Journal of Fluid Mechanics*, 251:1–20.

[95] Motheau, E., Nicoud, F., and Poinsot, T. (2014). Mixed acoustic–entropy combustion instabilities in gas turbines. *Journal of Fluid Mechanics*, 749:542–576.

[96] Nayfeh, A. H. (2008). *Perturbation methods*. John Wiley & Sons.

[97] Nehse, M., Warnat, J., and Chevalier, C. (1996). Kinetic modeling of the oxidation of large aliphatic hydrocarbons. In *Symposium (International) on Combustion*, volume 26, pages 773–780. Elsevier.

[98] Nicoud, F., Benoit, L., Sensiau, C., and Poinsot, T. (2007). Acoustic modes in combustors with complex impedances and multidimensional active flames. *AIAA Journal*, 45(2):426–441.

[99] Nicoud, F. and Ducros, F. (1999). Subgrid-scale stress modelling based on the square of the velocity gradient tensor. *Flow, Turbulence and Combustion*, 62(3):183–200.

[100] Nicoud, F. and Poinsot, T. (2001). 3.2 boundary conditions for compressible unsteady flows. *Absorbing Boundaries and Layers, Domain Decomposition Methods: Applications to Large Scale Computers*, page 78.

[101] Noack, B. R., Afanasiev, K., Morzyński, M., Tadmor, G., and Thiele, F. (2003). A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *Journal of Fluid Mechanics*, 497:335–363.

[102] Noiray, N., Durox, D., Schuller, T., and Candel, S. (2008). A unified framework for non-linear combustion instability analysis based on the flame describing function. *Journal of Fluid Mechanics*, 615:139–167.

[103] Oberleithner, K., Sieber, M., Nayeri, C., Paschereit, C., Petz, C., Hege, H.-C., Noack, B., and Wygnanski, I. (2011). Three-dimensional coherent structures in a swirling jet undergoing vortex breakdown: stability analysis and empirical mode construction. *Journal of Fluid Mechanics*, 679:383–414.

[104] OpenCFD (2018). OpenFOAM. https://www.openfoam.com/.

[105] Palies, P., Durox, D., Schuller, T., and Candel, S. (2010). The combined dynamics of swirler and turbulent premixed swirling flames. *Combustion and Flame*, 157(9):1698–1717.

[106] Palies, P., Schuller, T., Durox, D., and Candel, S. (2011). Modeling of premixed swirling flames transfer functions. *Proceedings of the Combustion Institute*, 33(2):2967–2974.

[107] Pankiewitz, C. and Sattelmayer, T. (2002). Time domain simulation of combustion instabilities in annular combustors. In *ASME Turbo expo 2002*, pages 309–320. American Society of Mechanical Engineers.

[108] Paschereit, C. O., Gutmark, E., and Weisenstein, W. (2000). Excitation of thermoacoustic instabilities by interaction of acoustics and unstable swirling flow. *AIAA Journal*, 38(6):1025–1034.

[109] Perret, L., Delville, J., Manceau, R., and Bonnet, J.-P. (2008). Turbulent inflow conditions for large-eddy simulation based on low-order empirical model. *Physics of Fluids*, 20(7):075107.

[110] Pier, B., Huerre, P., and Chomaz, J.-M. (2001). Bifurcation to fully nonlinear synchronized structures in slowly varying media. *Physica D: Nonlinear Phenomena*, 148(1):49–96.

[111] Pierce, C. D. and Moin, P. (2004). Progress-variable approach for large-eddy simulation of non-premixed turbulent combustion. *Journal of fluid Mechanics*, 504:73–97.

[112] Piomelli, U. and Balaras, E. (2002). Wall-layer models for large-eddy simulations. *Annual Review of Fluid Mechanics*, 34(1):349–374.

[113] Poinsot, T. (2015). video of lecture. Lectures were recorded at the Princeton-CEFRC combustion summer school and are available at https://www.youtube.com/playlist?list=PLbInEHTmP9VYVRFGwl7tamDSYCvwnuT9h.

[114] Poinsot, T. and Veynante, D. (2005). *Theoretical and numerical combustion*. RT Edwards, Inc.

[115] Poinsot, T. J. and Lele, S. (1992). Boundary conditions for direct simulations of compressible viscous flows. *Journal of computational physics*, 101(1):104–129.

[116] Polifke, W. (2015). Six lectures on thermodynamic combustion instability. *21st CISM-IUTAM International Summer School on measurement, analysis and passive control of thermoacoustic oscillations.*

[117] Pope, S. B. (2001). *Turbulent flows*. IOP Publishing.

[118] Qiu, H.-H. and Sommerfeld, M. (1992). A reliable method for determining the measurement volume size and particle mass fluxes using phase-doppler anemometry. *Experiments in Fluids*, 13(6):393–404.

[119] Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., and Kompenhans, J. (2018). *Particle image velocimetry: a practical guide.* Springer.

[120] Rayleigh, L. (1878). The explanation of certain acoustical phenomena. *Roy. Inst. Proc.*, 8:536–542.

[121] Raynaud, F., Eggels, R. L., Staufer, M., Sadiki, A., and Janicka, J. (2015). Towards unsteady simulation of combustor-turbine interaction using an integrated approach. In *ASME Turbo Expo 2015: Turbine Technical Conference and Exposition*, pages V02BT39A004–V02BT39A004. American Society of Mechanical Engineers.

[122] Reynolds, W. and Hussain, A. (1972). The mechanics of an organized wave in turbulent shear flow. Part 3. Theoretical models and comparisons with experiments. *Journal of Fluid Mechanics*, 54(2):263–288.

[123] Richards, G. A., Straub, D. L., and Robey, E. H. (2003). Passive control of combustion dynamics in stationary gas turbines. *Journal of Propulsion and Power*, 19(5):795–810.

[124] Roux, S., Lartigue, G., Poinsot, T., Meier, U., and Bérat, C. (2005). Studies of mean and unsteady flow in a swirled combustor using experiments, acoustic analysis, and large eddy simulations. *Combustion and Flame*, 141(1):40–54.

[125] Rowley, C. W., Colonius, T., and Murray, R. M. (2004). Model reduction for compressible flows using POD and Galerkin projection. *Physica D: Nonlinear Phenomena*, 189(1-2):115–129.

[126] Rupp, J. (2013). *Acoustic Absorption and the Unsteady Flow Associated with Circular Apertures in a Gas Turbine Environment.* PhD thesis, Loughborough Unviersity, Loughborough.

[127] Rupp, J., Carrotte, J., and Macquisten, M. (2012). The use of perforated damping liners in aero gas turbine combustion systems. *Journal of engineering for gas turbines and power*, 134(7):071502.

[128] Sagaut, P. (2013). *Multiscale and multiresolution approaches in turbulence: LES, DES and hybrid RANS/LES methods: applications and guidelines.* World Scientific.

[129] Sarkar, S. and Voke, P. R. (2006). Large-eddy simulation of unsteady surface pressure over a low-pressure turbine blade due to interactions of passing wakes and inflexional boundary layer. *Journal of Turbomachinery*, 128(2):221–231.

[130] Schildmacher, K.-U., Koch, R., and Bauer, H.-J. (2006). Experimental characterization of premixed flame instabilities of a model gas turbine burner. *Flow, Turbulence and Combustion*, 76(2):177–197.

[131] Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28.

[132] Schneider, C., Dreizler, A., and Janicka, J. (2005). Fluid dynamical analysis of atmospheric reacting and isothermal swirling flows. *Flow, Turbulence and Combustion*, 74(1):103–127.

[133] Schroeder, W. J., Lorensen, B., and Martin, K. (2004). *The visualization toolkit: an object-oriented approach to 3D graphics*. Kitware.

[134] Seinfeld, J. H. and Pandis, S. N. (1998). From air pollution to climate change. *Atmospheric Chemistry and Physics*, page 1326.

[135] Selle, L., Lartigue, G., Poinsot, T., Koch, R., Schildmacher, K.-U., Krebs, W., Prade, B., Kaufmann, P., and Veynante, D. (2004). Compressible large eddy simulation of turbulent combustion in complex geometry on unstructured meshes. *Combustion and Flame*, 137(4):489–505.

[136] Seybert, A. and Ross, D. (1977). Experimental determination of acoustic properties using a two-microphone random-excitation technique. *Journal of the Acoustical Society of America*, 61(5):1362–1370.

[137] Shehata, M. S. and ElKotb, M. M. (2014). Combustion characteristics for turbulent prevaporized premixed flame using commercial light diesel and kerosene fuels. *Journal of Combustion*, 2014.

[138] Sieber, M., Paschereit, C. O., and Oberleithner, K. (2016). Spectral proper orthogonal decomposition. *Journal of Fluid Mechanics*, 792:798–828.

[139] Sirovich, L. (1987). Turbulence and the dynamics of coherent structures. I. Coherent structures. *Quarterly of Applied Mathematics*, 45(3):561–571.

[140] Smagorinsky, J. (1963). General circulation experiments with the primitive equations: I. the basic experiment. *Monthly Weather Review*, 91(3):99–164.

[141] Spalart, P. R. (2009). Detached-eddy simulation. *Annual Review of Fluid Mechanics*, 41:181–202.

[142] Spencer, A. (2004). Isothermal pilot main interaction of lean burn fuel injectors. *Loughborough University Technology Centre Report.*

[143] Staufer, M. and Eggels, R. (2018). Private Communication, Rolls-Royce Deutschland. We discussed the use of RANS in the injector design process numerous times during my time at RRD.

[144] Steinberg, A. M., Boxx, I., Stöhr, M., Carter, C. D., and Meier, W. (2010). Flow–flame interactions causing acoustically coupled heat release fluctuations in a thermo-acoustically unstable gas turbine model combustor. *Combustion and Flame*, 157(12):2250–2266.

[145] Stöhr, M., Boxx, I., Carter, C. D., and Meier, W. (2012). Experimental study of vortex-flame interaction in a gas turbine model combustor. *Combustion and Flame*, 159(8):2636–2649.

[146] Stöhr, M., Oberleithner, K., Sieber, M., Yin, Z., and Meier, W. (2018). Experimental study of transient mechanisms of bistable flame shape transitions in a swirl combustor. *Journal of Engineering for Gas Turbines and Power*, 140(1):011503.

[147] Su, J. (2016). *Unsteady Aerodynamic Response Characteristics of Gas Turbine Fuel Injectors*. PhD thesis, Loughborough University, Loughborough.

[148] Su, J., Barker, A., Garmory, A., and Carrotte, J. (2018). Spray response to acoustic forcing of a multi-passage lean-burn aero-engine fuel injector. In *ASME 2018 Turbine Technical Conference and Exposition*, pages GT2018–75554. American Society of Mechanical Engineers.

[149] Su, J., Garmory, A., and Carrotte, J. (2015a). The aerodynamic response of fuel injector passages to incident acoustic waves. In *ASME 2015 Turbine Technical Conference and Exposition*, pages GT2015–43248. American Society of Mechanical Engineers.

[150] Su, J., Rupp, J., Garmory, A., and Carrotte, J. F. (2015b). Measurements and computational fluid dynamics predictions of the acoustic impedance of orifices. *Journal of Sound and Vibration*, 352:174–191.

[151] Tammisola, O. and Juniper, M. (2016). Coherent structures in a swirl injector at Re= 4800 by nonlinear simulations and linear global modes. *Journal of Fluid Mechanics*, 792:620–657.

[152] Terhaar, S., Ćosić, B., Paschereit, C. O., and Oberleithner, K. (2016). Suppression and excitation of the precessing vortex core by acoustic velocity fluctuations: An experimental and analytical study. *Combustion and Flame*, 172:234–251.

[153] Treleaven, N. C. W., Su, J., Garmory, A., and Page, G. J. (2017). The response to incident acoustic waves of the flow field produced by a multi-passage lean-burn aero-engine fuel injector. pages GT2017–64527.

[154] Turton, S. E., Tuckerman, L. S., and Barkley, D. (2015). Prediction of frequencies in thermosolutal convection from mean flows. *Physical Review E*, 91(4):043009.

[155] Van Albada, G., Van Leer, B., and Roberts, W. (1997). A comparative study of computational methods in cosmic gas dynamics. In *Upwind and High-Resolution Schemes*, pages 95–103. Springer.

[156] Van der Pol, B. (1926). LXXXVIII. On relaxation-oscillations. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):978–992.

[157] van Oijen, J. and de Goey, L. (2000). Modelling of premixed laminar flames using flamelet-generated manifolds. *Combustion Science and Technology*, 161(1):113–137.

[158] Vicquelin, R., Fiorina, B., Payet, S., Darabiha, N., and Gicquel, O. (2011). Coupling tabulated chemistry with compressible CFD solvers. *Proceedings of the Combustion Institute*, 33(1):1481–1488.

[159] Welch, P. (1967). The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics*, 15(2):70–73.

[160] Wilcox, D. C. (2008). Formulation of the k-ω turbulence model revisited. *AIAA journal*, 46(11):2823–2838.

[161] Williams, F. A. (2018). *Combustion theory.* CRC Press.

[162] Williams, M. A., Carrotte, J. F., Moran, A. J., and Walker, A. D. (2018). Impact of upstream boundary conditions on fuel injector performance in a low TRL reacting flow experimental facility.

[163] Wolf, P., Balakrishnan, R., Staffelbach, G., Gicquel, L. Y., and Poinsot, T. (2012). Using LES to study reacting flows and instabilities in annular combustion chambers. *Flow, Turbulence and Combustion*, 88(1):191–206.

[164] Worth, N. A. and Dawson, J. R. (2013). Self-excited circumferential instabilities in a model annular gas turbine combustor: Global flame dynamics. *Proceedings of the Combustion Institute*, 34(2):3127–3134.

[165] Wu, X. (2017). Inflow turbulence generation methods. *Annual Review of Fluid Mechanics*, 49:23–49.

[166] Wu, X., Schlüter, J., Moin, P., Pitsch, H., Iaccarino, G., and Ham, F. (2006). Computational study on the internal layer in a diffuser. *Journal of Fluid Mechanics*, 550:391–412.

[167] Xiao, F., Dianat, M., and McGuirk, J. J. (2014). LES of turbulent liquid jet primary breakup in turbulent coaxial air flow. *International Journal of Multiphase Flow*, 60:103–118.

[168] Xie, Z.-T. and Castro, I. P. (2008). Efficient generation of inflow conditions for large eddy simulation of street-scale flows. *Flow, Turbulence and Combustion*, 81(3):449–470.

[169] Xiong, Z., Nagarajan, S., and Lele, S. K. (2004). Simple method for generating inflow turbulence. *AIAA Journal*, 42(10):2164–2166.

[170] Yakhot, V., Orszag, S., Thangam, S., Gatski, T., and Speziale, C. (1992). Development of turbulence models for shear flows by a double expansion technique. *Physics of Fluids A: Fluid Dynamics*, 4(7):1510–1520.

[171] Yang, D. and Morgans, A. S. (2016). A semi-analytical model for the acoustic impedance of finite length circular holes with mean flow. *Journal of Sound and Vibration*, 384:294–311.

[172] Zedda, M. (2017). Private Communication, Rolls-Royce plc and Loughborough University. We discussed the practical use of LES in industry many times, where the effects of upstream turbulence production are often ignored.

Appendix A

Supplementary derivations

This appendix contains supplementary derivations that are interesting but do not contain sufficient novelty or do not contribute greatly to the central thesis.

A.1 Single phase flow methods and algorithms

A.1.1 Steady incompressible flow governing equations

The governing equations are the incompressible Navier-Stokes equations and they are discretised in space by breaking the domain into a number of small volumes through the use of an unstructured mesh and integrating the governing equations over each individual cell which for the momentum equation becomes (for steady flow) [39]:

$$\int_{S_P} \rho \mathbf{u}(\mathbf{u} \cdot \mathbf{dS}) = \int_{S_P} \overline{\overline{\mathbf{T}}} \cdot \mathbf{dS} + \int_{V_P} \rho \mathbf{f} dV \tag{A.1}$$

where V_P is the volume and S_P are the surfaces of the Pth cell. $\overline{\overline{\mathbf{T}}}$ is the stress tensor containing the influence of pressure and viscosity and is defined as:

$$T_{ij} = \mu\Big(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\Big) - p\delta_{ij} \tag{A.2}$$

In the case of PRECISE-UNS, second order accuracy in space is provided by using second order estimates of the flux and gradients at the boundaries of each cell. In the case of the pressure, a symmetric stencil is used, however for velocity, a symmetric stencil may lead to the generation of instabilities and as such, in some cases a linear upwind stencil is used. Upwinding computes the flux by only considering the points that are upstream, as defined by the local velocity. In certain cases even this approximation can become unstable in which case 1st order unwinding is used instead, or a blending factor is used that blends the two methods together to achieve an order of accuracy and stability between the two. TVD schemes [47] use a blending function unique to each cell allowing for the order and stability of the scheme to vary across the domain.

A.1.2 SIMPLE algorithm

In order for the Navier-Stokes equations to be solved over the domain an iteration scheme must be chosen to drive the flow from the initial condition to the solution. The SIMPLE (Semi-Implicit Method for Pressure-Linked Equations) [39] implicit method is used as it allows for rapid convergence of steady solutions. Rather than compute the solution of both pressure and velocity simultaneously the pressure at the new iteration step is allowed to be equal to the pressure at the previous iteration step and the momentum equation is inverted to find the approximate future velocity $u_{i,P}^{m*}$. This solution does not contain the updated pressure field and it will usually not satisfy the continuity equation, as such, an additional equation needs to be solved to project this velocity onto its equivalent divergence free field. This is done by considering the difference between the divergent and divergence free velocities as well as the change in pressure between the old and new iteration:

$$u_i^m = u_i^{m*} + u' \tag{A.3}$$

$$p^m = p^{m-1} + p' \tag{A.4}$$

These small correctional quantities can be linked together using Equation (A.1) yielding:

$$u'_{i,P} = \tilde{u}'_{i,P} - \frac{1}{A_P^{u_i}} \Big(\frac{\delta p'}{\delta x_i}\Big)_P \tag{A.5}$$

with

$$\tilde{u}'_{i,P} = -\frac{\sum_l A_l^{u_i} u'_{i,l}}{A_P^{u_i}} \tag{A.6}$$

where $A_P^{u_i}$ is the matrix coefficient for cell P and $A_l^{u_i}$ are the coefficients for the neighbouring cell l which together define the relationship between the velocity in cell P and the neighbouring cells. Application of Equation (A.5) into the continuity equation results in the following equation for a pressure field that will lead to enforced continuity in velocity:

$$\frac{\delta}{\delta x_i}\Big[\frac{\rho}{A_P^{u_i}}\Big(\frac{\delta p'}{\delta x_i}\Big)\Big]_P = \Big[\frac{\delta(\rho u_i^{m*})}{\delta x_i}\Big]_P + \Big[\frac{\delta(\rho \tilde{u}'_i)}{\delta x_i}\Big]_P \tag{A.7}$$

The SIMPLE algorithm assumes that the last term in this equation is negligible. Through the solution of this equation a new velocity can be derived that satisfies the continuity equation through Equation (A.5). This then allows for computation of the velocity at the next iteration step. This process is repeated until the difference between the velocities and pressures calculated at various iterations falls below some predetermined convergence criteria. Between each of these interactions however, under relaxation is used to ensure stability of the solution.

A.1.3 PISO Algorithm

Whilst the SIMPLE algorithm allows for sufficiently rapid convergence of solution, neglecting the last term of Equation (A.7) causes the algorithm to require several inner iterations with significant under-relaxation before convergence is achieved. Whilst the last term is similarly neglected, the PISO (Pressure Implicit with Splitting of Operator) algorithm performs additional pressure correction steps that seek to find the pressure p'' defined by:

$$\frac{\delta}{\delta x_i}\Big[\frac{\rho}{A_P^{u_i}}\Big(\frac{\delta p''}{\delta x_i}\Big)\Big]_P = \Big[\frac{\delta(\rho\tilde{u}_i')}{\delta x_i}\Big]_P \tag{A.8}$$

with $\tilde{u}_i'$ found using Equation (A.5) after application of Equation (A.7) from the initial pressure correction step. Finally the second velocity correction can be found from:

$$u_{i,P}'' = \tilde{u}_{i,P}' - \frac{1}{A_P^{u_i}}\Big(\frac{\delta p''}{\delta x_i}\Big)_P \tag{A.9}$$

and the updated pressure and velocity from:

$$\mathbf{u}_i^m = \mathbf{u}_i^{m*} + \mathbf{u}' + \mathbf{u}'' \tag{A.10}$$

$$p^m = p^{m-1} + p' + p'' \tag{A.11}$$

This pressure correction loop can be repeated several times per iteration. Strictly speaking PISO refers to the case where the flow is steady and many iterations and pressure corrections are performed to obtain the solution or the flow can be unsteady but in this case only one interation of the algorithm (but with multiple pressure corrections) is performed per time step.

It has been noted in [39] that the advantage of this method is that the additional pressure correction means that the pressure equation does not need to be under relaxed and equivalently in the simulations run using OpenFOAM it has been found by Su [147] that the velocity and pressure need not be under-relaxed between timesteps either. The allows the simulations to be performed with only one iteration per timestep. This significantly reduces the amount of time required to perform simulations using OpenFOAM versus PRECISE-UNS despite a slower per iteration computation time due to the use of C++ instead of FORTRAN.

A.1.4 PIMPLE algorithm

Referring to the OpenFOAM nomenclature, PIMPLE is a hybrid PISO and SIMPLE algorithm. The PISO method is only applicable in the case of unsteady flows when the number of iterations per time step is 1. If the number of these iterations are increased then the method is transformed into the PIMPLE method as the additional iteration steps are similar to the SIMPLE algorithm.

The OpenFOAM acronym for the above general algorithm is *PimpleFoam* or *rhoPimpleFoam* for compressible flows and can be used for steady and unsteady flows while in the case where the number of inner iterations are restricted to one with no under-relaxation the method is refered to as *rhoPisoFoam* and can only be used for unsteady flows. It was has been found by Su [147] that the *rhoPisoFoam* solver is more accurate for acoustically forced simulations than the *rhoPimpleFoam* solver due to the lack of under-relaxation of either velocity or pressure. As such the *rhoPimpleFoam* solver was only used to find the unsteady swirling bias flow of the injector while the acoustically forced simulations were computed using *rhoPisoFoam*.

PRECISE-UNS uses a method equivalent to the PIMPLE method and may be forced to conduct extra pressure correction steps or as many iterations as required by the user.

A.1.5 Compressibility

In order to solve the steady flow equations in the case of compressible flows Equation (A.1) must now be considered in addition to an equation describing the evolution of the flow total energy (E_h) or an alternative form. OpenFOAM uses as default the equation for total energy [104]:

$$\int_{S_p} (\rho E_h + p)\mathbf{u} \cdot \mathbf{dS} = \int_{S_p} (\lambda \nabla T + \overline{\overline{\mathbf{T}}}\mathbf{u}) \cdot \mathbf{dS} \tag{A.12}$$

However the source term due to viscous heating is neglected due to the assumption that the Mach number is low. Garnier et al. [41] have identified that the total energy equation must be used if the equations are to remain conservative. This is important in the case of high Mach number flows where discontinuities exist. As the Mach number in the combustion chamber is significantly less than one, the total energy equation may be replaced with the equation for specific enthalpy h which is useful from the point of view of the integration of pre-existing combustion models in PRECISE-UNS [39]:

$$\int_{S_p} \rho h \mathbf{u} \cdot \mathbf{dS} = \int_{S_p} \lambda \nabla T \cdot \mathbf{dS} + \int_{V_P} [\mathbf{u} \cdot \nabla p + \overline{\overline{\mathbf{T}}} : \nabla \mathbf{u}] dV \tag{A.13}$$

where λ is the thermal conductivity and $\overline{\overline{\mathbf{T}}}$ is altered to include the viscous losses due to compression of the fluid:

$$T_{ij} = \mu\Big(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\Big) - \Big(\frac{2\mu}{3}\nabla \cdot \mathbf{u} + p\Big)\delta_{ij} \tag{A.14}$$

This equation is solved to resolve the specific enthalpy in each cell and therefore the temperature. The momentum equation remains the same and is solved in a similar method to the incompressible case however the continuity equation, while used in the incompressible case to enforce the divergence free velocity field, must now be used to calculate the density in each cell. If the momentum equation is solved with the previous iteration's pressure field, the momentum will in most cases not satisfy the continuity equation leading to:

$$\sum \dot{m}_e^* = \sum \rho_e^* \mathbf{u}_e^* \cdot \mathbf{S}_e = Q_m^* \tag{A.15}$$

where the subscript e refers to the eth cell face, $*$ refers to the pre-corrected value and at convergence $Q_m^* = 0$. The mass flux correction can be found to be [39]:

$$\dot{m}'_e = (\rho^{m-1} S v u'_n)_e + (u_n^{m*} S \rho')_e + (\rho' u'_n S)_e \tag{A.16}$$

where the last term is neglected as it is second order. The first term can be approximated in the same way as Equation (A.5) to link $\mathbf{u}'$ to p' while the second term must be approximated by the equation of state to close the set of equations. The mass flux correction now becomes [39]:

$$m'_e = -(\rho^{m-1} S V_P)_e \Big(\frac{1}{A_p^u}\Big)_e \Big(\frac{\delta p'}{\delta n}\Big)_e + \Big(\frac{C_p m^*}{\rho^{m-1}}\Big)_e p'_e \tag{A.17}$$

Substitution of this equation back into the continuity equation will yield the following pressure correction equation [39]:

$$A_P p'_p + \sum A_l p'_l = -Q_m^* \tag{A.18}$$

Hence the SIMPLE algorithm extended to compressible flows involves the solution of Equations (A.1), (A.18), (A.13) and the equation of state to get temperature and density iteratively until the values of p' drop below a predefined level.

Su [147] has modified the PISO algorithm in OpenFOAM to be similar to the algorithm proposed by Issa [53] where the energy equation is solved within the PISO pressure correction loop. The same approach was attempted within PRECISE-UNS however initial results did not show any benefit when compared to the SIMPLE approach where the energy equation (specific enthalpy) is solved only once per iteration.

A.1.6 Unsteadiness

As the flow is now unsteady, the equation that must be solved for each cell becomes [39]:

$$\frac{d}{dt}\int_{V_i} \rho \mathbf{u} dV + \int_{S_i} \rho \mathbf{u}(\mathbf{u} \cdot \mathbf{dS}) = \int_{S_i} \overline{\overline{\mathbf{T}}} \cdot \mathbf{dS} + \int_{V_i} \rho \mathbf{f} dV \tag{A.19}$$

The SIMPLE algorithm is applied again to the solution of this equation in exactly the same manner as in section A.1.2 with the addition of the unsteady term to Equation (A.1). However when convergence is reached or the number of inner iterations m exceeds a predetermined number u^{n+1} and p^{n+1} are updated with the values of u^m and p^m respectively where n is the timestep number. At this point the inner iteration scheme is reinitialised and the process begins anew. The new term includes the gradient of the velocity with respect to time. This term is dealt with either by using the first order accurate Euler finite difference method [39]:

$$\frac{d\mathbf{u}}{dt} = \frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\delta t} \tag{A.20}$$

or the second order accurate quadratic finite difference method:

$$\frac{d\mathbf{u}}{dt} = \frac{\mathbf{u}^{n+1} - 4\mathbf{u}^n + 3\mathbf{u}^{n-1}}{\delta t} \tag{A.21}$$

the second order quadratic method was preferred and employed in all simulations except where additional stability was required in which case the Euler method can be used to damp out some of the high frequency components of the pressure oscillations leading to smoother results and better convergence.

The compressible SIMPLE algorithm is used in exactly the same way except that the unsteady momentum equation is solved together with the mass and enthalpy equations [39]:

$$\frac{d}{dt}\int_{V_P}\rho dV + \int_{S_P}\rho\mathbf{u}\cdot\mathbf{dS} = 0 \tag{A.22}$$

$$\frac{d}{dt}\int_{V_P}\rho h dV + \int_{S_p}\rho h\mathbf{u}\cdot\mathbf{dS} = \int_{S_p}k\nabla T\cdot\mathbf{dS} + \int_{V_P}\left[\frac{\partial p}{\partial t} + \mathbf{u}\cdot\nabla p + \overline{\overline{\mathbf{T}}}:\nabla\mathbf{u}\right]dV \tag{A.23}$$

with the unsteady terms dealt with in exactly the same methodology as in the unsteady momentum equation. The compressible SIMPLE or PISO algorithm is applied within each sub-step until convergence is acquired whereupon the timestep and unsteady term are updated. The total energy equation used by OpenFOAM, with the source term removed, becomes [104]:

$$\frac{d}{dt}\int_{V_P}\rho E dV + \int_{S_p}(\rho E + p)\mathbf{u}\cdot\mathbf{dS} = \int_{S_p}(k\nabla T + \overline{\overline{\mathbf{T}}}\mathbf{u})\cdot\mathbf{dS} \tag{A.24}$$

A.1.7 Turbulence closure

Turbulence closure is required due to the appearance of the Reynolds shear stress term in the RANS equations and the residual stress tensor in the filtered LES equations. Although the mathematical justification is different, all turbulence closure methods seek to inject the correct amount of additional viscosity into each cell such that it has the same net effect as the turbulence that is not resolved in the simulation. As such, the numerically relevant RANS, URANS or LES equations can be developed by simply replacing μ with μ_e where:

$$\mu_e(\mathbf{x},t) = \mu_t(\mathbf{x},t) + \mu \tag{A.25}$$

where μ_t is the turbulent or SGS (Sub-Grid Scale) viscosity and is obtained by using one of the following methods of turbulence closure.

RANS: $k-\epsilon$ RNG

There are several $k-\epsilon$ turbulence models in use which define, through the use of dimensional analysis, the turbulent viscosity as being a function of the turbulent kinetic energy (TKE or k) and the dissipation ϵ:

$$\nu_t = C_\mu\frac{k^2}{\epsilon} \tag{A.26}$$

The equation for k can be derived from the Navier-Stokes equations exactly however several terms require additional closure. The equation first developed by Launder and Spalding [67] is:

$$\frac{\partial k}{\partial t} + \overline{u}_k \frac{\partial k}{\partial x_k} = 2\nu_t S_{ij} S_{ij} - \epsilon + \frac{\partial}{\partial x_k}\left[\left(\nu + \frac{\nu_t}{\sigma_k}\right)\frac{\partial k}{\partial x_k}\right] \tag{A.27}$$

where S_{ij} is the i, j component of mean rate-of-strain tensor. The terms are identified as being representative of unsteadiness, convection by the mean flow, production of TKE, dissipation of TKE and the diffusion of the TKE by the viscosity and turbulence. Similarly the equation for dissipation can also be derived exactly from the Navier-Stokes equations, however the equation contains additional unknowns that must be closed [81]. The dissipation equation is therefore replaced by a psuedo-dissipation equation where each term is equivalent to the same physical processes as in the TKE equation: time variation, convection, production, dissipation and diffusion:

$$\frac{\partial \epsilon}{\partial t} + \overline{u}_k \frac{\partial \epsilon}{\partial x_k} = 2C_{\epsilon 1}\nu_t S_{ij} S_{ij}\frac{\epsilon}{k} - C_{\epsilon 2}\frac{\epsilon}{k} + \frac{\partial}{\partial x_k}\left[\left(\nu + \frac{\nu_t}{\sigma_\epsilon}\right)\frac{\partial \epsilon}{\partial x_k}\right] \tag{A.28}$$

Yakhot [170] developed the Re-Normalisation Group version of the $k - \epsilon$ model to improve the performance of RANS simulations in the log-layer of the boundary layer and to correct the stagnation point anomaly [81]. The main change is the addition of an extra term $(-R)$ in the dissipation equation:

$$R = \frac{C_\mu \eta^3 (1 - \eta/\eta_0)}{1 + \beta\eta^3}\frac{\epsilon^2}{k} \tag{A.29}$$

This is equivalent to a modification of $C_{\epsilon 1} \rightarrow C^*_{\epsilon 1}$:

$$C^*_{\epsilon 1} = C_{\epsilon 1} - \frac{C_\mu \eta^3 (1 - \eta/\eta_0)}{1 + \beta\eta^3} \tag{A.30}$$

with $\eta = \sqrt{2} S_{ij} S_{ij} k/\epsilon$. The coefficients values are also modified from those used in the standard $k - \epsilon$ model: $C_\mu = 0.085$, $C_{\epsilon 1} = 1.42$, $C_{\epsilon 2} = 1.68$, $\sigma_k = 0.72$, $\sigma_\epsilon = 0.72$, $\eta_0 = 0.72$, $\beta = 0.72$.

This model is popular but the disadvantage of this model was shown by Laurence [68] to be that the improvement of the predictive capability of the model is due to the increase in dissipation the change in coefficients causes rather than correcting the over production in the stagnation regions of the flow.

RANS: $k - \omega$ SST

The $k - \omega$ model was developed as an alternative to the $k - \epsilon$ model by Wilcox [160]. It replaces the calculation of ϵ with the characteristic frequency of the large scale eddies ($\omega = \epsilon/k$). The equation for ω is formulated in exactly the same way as for k and ϵ:

$$\frac{\partial \omega}{\partial t} + \overline{u}_k \frac{\partial \omega}{\partial x_k} = 2\alpha\nu_t S_{ij} S_{ij}\frac{\omega}{k} - \beta\omega^2 + \frac{\partial}{\partial x_k}\left[\left(\nu + \frac{\nu_t}{\sigma_\omega}\right)\frac{\partial \omega}{\partial x_k}\right] \tag{A.31}$$

It has been observed that the model behaves better in the near-wall region than the k-ϵ model, especially in the cases where the flow separates [81, 117, 160]. The disadvantage of this model is that it is sensitive to the inlet boundary conditions, especially in the case of external flows

[81, 117, 160]. Menter [85] developed the $k-\omega$ SST model as a hybrid model that switches between a $k-\epsilon$ method far from the wall and a $k-\omega$ method near the wall. The equation for ω becomes:

$$\begin{aligned}\frac{\partial \omega}{\partial t}+\overline{u}_k\frac{\partial \omega}{\partial x_k}=\alpha S_{ij}S_{ij}-\beta\omega^2+\frac{\partial}{\partial x_k}\left[\left(\nu+\frac{\nu_t}{\sigma_\omega}\right)\frac{\partial \omega}{\partial x_k}\right]\\ +2(1-F_1)\rho\sigma_{w2}\frac{1}{\omega}\frac{\partial k}{\partial x_k}\frac{\partial \omega}{\partial x_k}\end{aligned}\qquad\text{(A.32)}$$

where F_1 is a blending function responsible for the switching between the two models, the details of which can be found in [85].

LES: Smagorinsky

In the case of RANS closure, the assumption is that all scales of turbulence must be modelled. Because large turbulent scales are significantly larger than the mesh size and turbulence time scales larger than the timestep, successful turbulence models utilise transport equations that allow turbulence to be generated in one place and dissipated elsewhere. In the case of Large Eddy Simulation (LES), the Navier-Stokes equations are filtered in space to remove only the smallest scales of motion. These small scales are responsible for dissipation of turbulent kinetic energy and assuming a suitable filter size is chosen, are size limited to being smaller than the mesh spacing. When using a simple SGS model such as the constant Smagorinski model, such an assumption limits any modelling strategy to being entirely local. The model should only consider the local effects of unresolved turbulence scales on the velocity field. This removes the need for transport equations making LES sub-grid models in general simpler that RANS turbulence models. The caveat is in the assumption that a sufficient percentage of the large scales of motion are captured accurately requiring a sufficiently fine mesh and time step. The Smagorinsky [140] model calculates the SGS viscosity using:

$$\nu_t=(C_s\Delta_c)^2\sqrt{2S_{ij}S_{ij}}\qquad\text{(A.33)}$$

with $C_s\approx 0.18$. Kolmogorov formulated his analysis for the case of homogeneous isotropic turbulence at infinite Reynolds number which is a long way from the flow inside a combustion chamber [117]. Such limitations make the Smagorinsky constant to be variable depending on the flow being considered. In a complex flow, the Smagorinsky “constant” should be assumed to be variable. Unfortunately tuning the Smagorinsky constant *a posteriori* is not a sensible strategy.

The size of the filter (Δ_c) is also an important choice. For unstructured meshes the most sensible choice is $\Delta_c=\Omega_P^{(1/3)}$ where Ω_P is the volume of the Pth cell.

There have been several attempts at improving the performance of the Smagorinsky model by turning the Smagorinsky constant into a dynamic quantity. These *dynamic* Smagorinsky models have many theoretical advantages, one of which being that they can result in a negative C_s which is representative of energy flowing from small scales to large scales. This process is referred to as *backscattering* and has been observed to be a significant source of error in LES simulations especially in wall-bounded flows. Practically, the constant Smagorinsky model has an advantage in that Equation (A.33) is always positive, and will therefore, always increase dissipation of kinetic energy in regions of high gradient such as those found due to numerical instabilities. Whilst the dynamic models may be more accurate they are often much more susceptible to instabilities and must be used carefully [81, 128].

There are modifications of the constant Smagorinsky model that improve the performance of the model in certain flows. For instance it has been shown that the model performs poorly for wall bounded flows unless the mesh resolution approaches that required for a Direct Numerical Simulation (DNS) [128]. The model tends to over-estimate the amount of turbulent viscosity at the wall. One such modification is the WALE (Wall-Adapting Local Eddy-viscosity) model [99] that replaces the square of the stress tensor ($S_{ij}S_{ij}$) with a function involving the *trace-less* stress tensor $S^d_{ij}S^d_{ij}$.

The question of the most appropriate sub-grid scale model for combustor flows is an interesting one but unfortunately it is beyond the scope of this thesis. It is true that the constant Smagorinsky model may not be the best model to use in this type of flow, however its broad use in industrial applications and guaranteed stability make it an attractive choice. In general LES suffers from addition problems, one of which being that the SGS viscosity generated by the sub-grid model can sometime be smaller than the effective viscosity induced by the numerical scheme (See appendix B). Such observations have lead to the use of Implicit LES, where no sub-grid scale model is used and the dissipation of the numerical scheme is the only energy sink. In such situations the grid resolution will play a large role in the amount of energy dissipated by the numerics. It has been observed that this is equivalent to there being a mesh dependence on the calculation of the Smagorinsky constant [87]. All of these observations make LES a careful balancing act between different sources of error, and given this, there is a temptation to conclude that the sub-grid scale model is not the main driver of error. Indeed one of the great recent developers of LES, Pierre Sagaut was quoted as saying that new sub-grid scale models are only useful if they reduce error faster and at lower computational cost than simply increasing the grid resolution [69]. In this sense the optimal LES strategy should be to increase the efficiency and accuracy of the CFD code such that the simulation approaches a DNS most rapidly where the sub-grid scale model does nothing at all.

A.2 Hydrodynamic damping

Classical linear stability theory is applied to a steady solution of the Navier-Stokes equations that is devoid of fluctuating components that are not of the assumed mode form. This implies that the base flow is laminar upstream of the initial perturbation. In most practical situations including swirl stabilised fuel injectors the base flow is not a steady solution of the Navier-Stokes equations as there is no such thing as a steady solution to this problem. Instead the base flow is a time-average of an unsteady simulation or a RANS average of a steady simulation. This means that the base flow and the perturbations assumed by the analysis are superimposed on the turbulent flow whether or not the turbulent flow is computed or simply modelled by the turbulence modelling strategy. The analysis is therefore only valid if the presence of the turbulence does not invalidate the assumption that perturbations grow linearly. This requires that the interaction between linear instabilities and non-linear turbulence can be safely neglected. Tammisola et al. [151], following the analysis of Turton et al. [154], present the following argument (completed in Cartesian coordinates for simplicity):

Firstly the triple decomposition of the flow is achieved in a similar way to that presented by Reynolds & Hussain [122]:

$$\mathbf{u} = \overline{\mathbf{u}} + \tilde{\mathbf{u}} + \mathbf{u}' \tag{A.34}$$

The over line represents the time average here and the tilde represents all of the large scale periodic motion the flow. $\mathbf{u}'$ is therefore all of the incoherent turbulent fluctuations. Substituting this decomposition into the non-dimensionalised Navier-Stokes equations and averaging in time results in:

$$\overline{\mathbf{U}}.\nabla\overline{\mathbf{U}} = -\nabla\overline{p} + \nabla.\left(Re^{-1}\overline{\mathbf{S}} + \overline{\tilde{\mathbf{u}}\tilde{\mathbf{u}}} + \overline{\mathbf{u}'\mathbf{u}'}\right) \tag{A.35}$$

The periodic large scale motions will satisfy the phase averaged Navier-Stokes equations, which when the mean flow equation is subtracted gives:

$$\frac{\partial\tilde{\mathbf{u}}}{\partial t} + \overline{\mathbf{U}}.\nabla\tilde{\mathbf{u}} + \tilde{\mathbf{u}}.\nabla\overline{\mathbf{U}} = -\nabla\tilde{p} + \nabla.\left(Re^{-1}\tilde{\mathbf{s}} + \widetilde{\tilde{\mathbf{u}}\tilde{\mathbf{u}}} + \widetilde{\mathbf{u}'\mathbf{u}'}\right) \tag{A.36}$$

It is now assumed that the coherent motions can be decomposed into into a set l of discrete limit cycles and their n harmonics that may be expressed as a double Fourier series:

$$\tilde{\mathbf{u}} \approx \sum_{l>0}\sum_{n\neq 0}\mathbf{u}_{l,n}exp(in\omega_l t) \tag{A.37}$$

In the first instance Tammisola et al. consider the case where there is no correlation between harmonics of different n. Equation (A.37) is substituted into Equation (A.35) with the above assumption resulting in:

$$\overline{\mathbf{U}}.\nabla\overline{\mathbf{U}} = -\nabla\overline{p} + \nabla.\left(Re^{-1}\overline{\mathbf{S}} + \overline{\sum_{l>0}\sum_{n>0}\mathbf{u}_{l,n}\mathbf{u}_{l,-n}} + \overline{\mathbf{u}'\mathbf{u}'}\right) \tag{A.38}$$

This equation demonstrates that the influence of each fundamental periodic mode on the mean flow is through its own interaction with its conjugate. Substitution of Equation (A.37) into Equation (A.36) results in the following relation for the evolution of each fundamental mode:

$$\begin{aligned} -i\omega\mathbf{u}_{l,1} + \overline{\mathbf{U}}.\nabla\mathbf{u}_{l,1} + \mathbf{u}_{l,1}.\nabla\overline{\mathbf{U}} = \\ -\nabla p_{l,1} \\ +\nabla.\left(Re^{-1}\mathbf{s}_{l,1} + \sum_{n\neq 0,1}\mathbf{u}_{l,n}\mathbf{u}_{l,1-n} + (\widetilde{\mathbf{u}'\mathbf{u}'})_l\right) \end{aligned} \tag{A.39}$$

This equation shows that the limit cycle mode l can be influenced by four things: The mean flow, its own shear stress tensor, the interaction between each conjugate harmonic and each higher harmonic and the incoherent turbulent motions that oscillate at the frequency ω_l. The linearised Navier-Stokes operator is defined as:

$$\mathcal{L}(\mathbf{u}) = \overline{\mathbf{U}}.\nabla\mathbf{u} + \mathbf{u}.\nabla\overline{\mathbf{U}} + \nabla p - \nabla.(Re^{-1}\mathbf{u}) \tag{A.40}$$

Equation (A.39) can be re-written as:

$$\mathcal{L}(\mathbf{u}_{l,1}) = i\omega_l\mathbf{u}_{l,1} + \mathcal{N}_1 + \widetilde{\mathbf{u'u'}} \tag{A.41}$$

Condensed into the non-linear term $\mathcal{N}_1$ are all of the interactions between the relevant fundamental periodic modes and their conjugates. In order for the above equation to be recast as a linear eigenvalue problem, one of two things must be true. The first is that:

$$\mathcal{N}_1 + \widetilde{\mathbf{u'u'}} = 0 \tag{A.42}$$

As first observed by Turton et al. [154] this leads to the fundamental modes being an exact eigenmode of the linearised Navier-Stokes equations:

$$\mathcal{L}(\mathbf{u}_{l,1}) = i\omega\mathbf{u}_{l,1} \tag{A.43}$$

In this case the eigenvalue ω will have a zero imaginary component and hence growth rate. It will also have a frequency equal to the fundamental limit cycle frequency. Instinctively this sounds reasonable as the instability will first grow linearly until there is sufficient dissipation of kinetic energy to prevent the instability from growing further and a limit cycle is reached. In most flows this is achieved through the transition to turbulence and hence the production of $\overline{\mathbf{u'u'}}$. This energy dissipation achieved through turbulence in this case balances the amplification of the fundamental modes through the non-linear interaction.
The second possibility is that:

$$\mathcal{N}_1 + \widetilde{\mathbf{u'u'}} = \mathcal{A}\mathbf{u}_{l,1} \tag{A.44}$$

where $\mathcal{A}$ is a linear operator. The fundamental mode can then be shown to be an exact eigenmode of the modified linearised Navier-Stokes operator:

$$(\mathcal{L} - A)\{\mathbf{u}_{l,1}\} = i\omega\mathbf{u}_{l,1} \tag{A.45}$$

In this configuration the combined effect of the turbulence and the non-linear periodic interactions is to add addition linear terms to the Navier-Stokes equations such as additional diffusion, however in order for these modified equations to have the same global modes as the original linearised Navier-Stokes equations the growth rate of the largest amplitude limit cycle mode must similarly be zero.

In order to deal with the non-linear periodic interactions, Turton [154] made use of the observation that if the amplitude of the fundamental mode is ϵ then the amplitude of the harmonics will decay as $\mathbf{u}_{l,n} \propto O(\epsilon^n)$. In this case the largest term in the expression for $\mathcal{N}_1$ reduces to size $O(\epsilon^3)$ and is therefore close to negligible.

The turbulence is then dealt with by assuming a suitable linear operator. Tammisola et al. [151] have chosen to model the turbulent fluctuations using an eddy viscosity hypothesis:

$$\mu_t(\mathbf{s}_{l,1}) = \nabla \cdot \widetilde{\mathbf{u'u'}} \tag{A.46}$$

Implicit in this assumptions is the idea that the turbulent fluctuations act only to damp the coherent periodic motions and that the addition of the eddy viscosity linear operator A does not have any major effect on the frequency or spatial orientation of the unstable modes. Tammisola et al. [151] admit that these assumptions can only be tested *a posteriori* from a fully non-linear simulation however if the a linear global stability mode has an amplitude close to zero and shares a similar spatial structure and frequency as the leading POD (proper orthogonal decomposition) mode then this serves as a good check as to whether the above assumptions are valid.

A.3 Proper orthogonal decomposition

The proper orthogonal decomposition (POD) is a method for breaking up a set of data into two sets of orthogonal modes. It was first described by Lumley [77] but is closely related to the Karhunen-Loeve decomposition, singular value decomposition (SVD) and the principal component analysis (PCA). It can be thought of as an extension of matrix diagnolisation to non-square matrices.

The matrix diagonalisation is reduction of a matrix $\mathbf{A}$ into three:

$$\mathbf{A} = \mathbf{D^{-1} \Lambda D} \tag{A.47}$$

The matrix $\mathbf{\Lambda}$ is a diagonal matrix composing the eigenvalues of the original matrix while $\mathbf{D}$ contains as its columns the eigenvectors. The eigenvectors are orthogonal and collectively span the N-dimensional Hilbert space created by $\mathbf{A}$. The data is therefore reorientated into a basis where each eigenvector can be thought of as being an independent dimension. The strength of each eigenvector is given by its eigenvalue. Given a vector in three dimensional space, its magnitude would be its eigenvalue while its orientation would be given by its eigenvector.

Such a diagonalisation is useful for many reasons, it allows for easy determination of the matrix $\mathbf{A}$ raised to the n^{th} power:

$$\mathbf{A}^n = \mathbf{D^{-1} \Lambda^n D} \tag{A.48}$$

It also allows for data compression: As the importance of each eigenvector is given by its eigenvalue, those eigenvectors with a small eigenvalue could be neglected with little consequence to a reconstructed matrix $\mathbf{A}^*$. If the original matrix was composed of a set of vectors, the relative magnitude of each eigenvalue will describe how diverse the data is. If all of the vectors are pointing in a similar direction, the diagonalisation will collapse into one eigenvector with a very large eigenvalue compared to the others. This measure of similarity is what allows the POD to be used for structure identification.

The SVD is the extension to non-square matrices. One searches for a decomposition of the matrix $\mathbf{B}$ such that:

$$\mathbf{B} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^* \tag{A.49}$$

$\mathbf{\Sigma}$ remains a diagonal matrix full of eigenvalues but the columns of $\mathbf{U}$ are the left-singular vectors that represent the columns of the original matrix and the columns of $\mathbf{V}$ are the right hand set of vectors that represent the rows of the original matrix. This decomposition becomes attractive as it separates the data into two independent sets. The POD seeks to approximate a set of data into two sets of eigenvectors or modes:

$$\mathbf{u}(\mathbf{x}, \mathbf{t}) \approx \sum_{i=0}^{N_s} a_i(t)\phi_i(\mathbf{x}) \tag{A.50}$$

The first set $a_i(t)$ are the temporal modes and are independent of space while $\phi_i(\mathbf{x})$ are the spatial modes. The spatial POD modes are a measure of statistical correlation and identify regions of the flow that behave in a similar manner with respect to time. These regions are referred to as structures and the temporal modes show how each of these structures fluctuate in time. The mode eigenvalues then show by how much each structure contributes to the to the flow field. A mode with a high eigenvalue denotes a structure that dominates the flow field.

Similarly to the SVD, the POD spatial modes are orthonormal and therefore:

$$\int_{\mathbf{x}} \phi_i \phi_j d\mathbf{x} = \delta_{ij} \tag{A.51}$$

and the temporal modes are orthogonal with their integral in time equal to their corresponding eigenvalue:

$$\int_T a_i a_j dt = \delta_{ij}\lambda_i \tag{A.52}$$

A.3.1 Methodology

The SVD may be directly computed directly using a number of algorithms however in the case of flow field data it is very inefficient as the number of spatial degrees of freedom may be several orders of magnitude higher than the number of timesteps. As such, a more efficient method referred to as the snapshot POD method [139] has been used in this body of work. The snapshot POD method takes advantage of some of the properties of the SVD. Multiplication of the SVD by $\mathbf{B}^*$ yields:

$$\mathbf{B}^*\mathbf{B} = \mathbf{B}^*\mathbf{U}\mathbf{\Sigma}\mathbf{V}^* \tag{A.53}$$

$$\mathbf{B}^*\mathbf{B} = \mathbf{V}\mathbf{\Sigma}^*\mathbf{U}^*\mathbf{U}\mathbf{\Sigma}\mathbf{V}^* \tag{A.54}$$

$$\mathbf{B}^*\mathbf{B} = \mathbf{V}\mathbf{\Sigma}^*\mathbf{\Sigma}\mathbf{V}^* \tag{A.55}$$

or if $\mathbf{B}^*\mathbf{B} = \mathbf{C}$ and $\mathbf{C}$ is both hermitian and definite positive then:

$$\mathbf{C} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^* = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^* \tag{A.56}$$

The eigenvalue decomposition of $\mathbf{C}$ can therefore be used to find the SVD of $\mathbf{B}$. It is important to note that the matrix $\mathbf{C}$ is an $N \times N$ matrix and it is therefore easier to compute these eigenvalues if $N < M$ where M is the other dimension of $\mathbf{B}$. The computation of the POD begins therefore with the construction of the correlation matrix $\mathbf{C}$:

$$C_{i,j} = \frac{1}{N_s} \int_{\mathbf{x}} \mathbf{u}(\mathbf{x}, t_i) \cdot \mathbf{u}(\mathbf{x}, t_j) d\mathbf{x} \tag{A.57}$$

where N_s is the number of snapshots.

It is noted here that the choice of this *scalar product* is completely arbitrary however the above relation is used most frequently because it is representative of the kinetic energy of the flow. In some chapters this scalar product is altered in order to better identify structures involving additional variables.

The *numpy.eigs* python package is then used to find the eigenvalues and eigenvectors of $\mathbf{C}$, the POD eigenvalues and temporal modes:

$$\mathbf{C}a_i = \lambda_i a_i \tag{A.58}$$

It can then be shown the spatial POD modes can be found using:

$$\phi_i(\mathbf{x}) = \frac{1}{T\lambda_i} \int_T \mathbf{u}(\mathbf{x}, t) a_i(t) dt \tag{A.59}$$

A.3.2 Post-Processing

The Power Spectral Density

The power spectral density (PSD) is an attempt to measure the energy content of a signal as a function of frequency. As the temporal POD modes are a function of time they can be analysed through construction of thier PSD. The PSD is a useful tool for determining if there are structures in the flow that oscillate at a particular frequency, such as the frequency of forcing or a natural hydrodynamic instability. The PSD is calculated in its most basic form as:

$$S_{aa}(f) = \mathcal{F}(a(t))\mathcal{F}^*(a(t)) \tag{A.60}$$

where $\mathcal{F}(a(t))$ refers to the Fourier transform of $a(t)$. The original time series of each temporal mode is a discrete digital signal and as such the Fourier transform is computed via a Fast-Fourier Transform (FFT). The above method is correct however the error using it can be of a similar magnitude to the signal. There are two options for increasing the signal to noise ratio. The first is to use a windowing function that reduces the effect of the finite time windowing of the signal. These functions are good at reducing the noise, especially at high frequencies however they also induce spectral blurring of the PSD which makes peaks more difficult to find. The preferred option used here is called Welche's method and reduces the error by cutting the sample into blocks, computing each blocks' PSD and then averaging across each block:

$$\hat{S}_{aa} = \frac{1}{M}\sum_{i=1}^{M} S_{aa}^{(i)}(f) \tag{A.61}$$

where M is the number of blocks. The restriction of how many blocks can be used is related to the sample length and the frequencies of interest. The minimum block length is related to the minimum frequency by:

$$f_{min} = \frac{1}{T} \tag{A.62}$$

The blocks may be overlapped to improve the number of samples however the benefits of Welche's method begin to degrade once the overlap is extended to beyond 50%.

Azimuthal wave number identification

The dominant azimuthal wavenumber is calculated by introducing a Fourier series representation of each spatial POD mode:

$$\phi_i(x,r,\theta) = \sum_{m=0}^{N_a-1} \psi_m(x,r)e^{im\theta} \tag{A.63}$$

where N_a is the number of azimuthal wave numbers considered, and ψ_m are the azimuthal Fourier coefficients that contain information about how much each wavenumber contributes to each POD mode. They can be calculated with:

$$\psi_m(x,r) = \frac{1}{N_a}\sum_{m=0}^{N_a-1} \phi_i(x,r,\theta)e^{-im\theta} \tag{A.64}$$

Integration across the domain then gives the total contribution of each wavenumber across the region of interest:

$$E_m = \int_{x_0}^{x_1}\int_{0}^{R_1} |\psi_m(x,r)|^2 r dr dx \tag{A.65}$$

The value of E_m gives the contribution of each azimuthal wave number to each POD mode. Before Equation (A.64) can be applied to the field, each POD mode must be interpolated onto a structured cylindrical mesh in polar coordinates along with the physical limits of the integration. The cylindrical mesh generation, interpolation and averaging are described in further detail in Section 4.5.1 however in the case of the calculations of azimuthal wavenumber, the dimensions and mesh resolution were chosen differently. The cylindrical mesh was chosen to extend from just behind the point at which passage A and passage B meet to the location where the high velocity air coming from passage C meets the duct wall ($x_1 = 0.64D$) as this region is where the flow structure fluctuations are strongest. The outer radius ($R_1 = 0.64D$) was chosen such that it included the majority of the duct but excluded the effect of the walls. The resolution of the mesh was changed between 6 and 50 in each principle direction (x, r, θ). It was found that the sensitivity of ψ_m to mesh resolution was negligible beyond 30 points and as such that was the resolution chosen for the analysis.

Velocity Reconstruction

The POD describes the flow field as a set of orthogonal modes and the flow may be recomputed from the POD modes as in Equation (A.50). Additionally the contribution to the entire flow field of a particular POD mode is given by:

$$\mathbf{u}_i(\mathbf{x}, t) = a_i(t)\phi_i(\mathbf{x}) \tag{A.66}$$

POD modes often come in pairs. These pairs often have extremely similar eigenvalues and the POD temporal modes will appear similar with a discernible 90 degree phase shift between them. When the temporal modes are plotted against each other they will form a circle (or spiral). Such mode pairs are representative of convective processes and have been observed in a wide range of flows such as wakes and mixing layers [13]. An example of two mode pairs plotted in time and against each other is visualised in figure 4.4c.

In more complex flows, there may exist more than two modes that exhibit similar properties such as peak oscillation frequency or azimuthal wavenumber. These modes are often describing the evolution of one or more coherent structures in the flow that contain a significant proportion of the flow kinetic energy. At various point is this thesis the motion of a coherent structure will be highlighted using a linear combination of POD modes justified on the similarity of the temporal or spatial components of the POD modes involved.

A.4 The acoustic impedance of an orifice

As an example, consider a simple orifice in a wall, subjected to an incoming pressure wave. A fuel injector is not an orifice, however whilst it is significantly more complicated, it is fundamentally a hole in a wall separating a pressure differential. As such, one expects similar mechanisms of interaction between the injector and the sound wave. As the pressure wave approaches, the pressure differential across the orifice will change as the wave approaches and then either passes through or is reflected by the wall. As the pressure differential decreases, the fluid residing inside the orifice will accelerate in the negative direction and then as the wave passes, it will accelerate in the positive direction. At the limit of infinite wavelength or zero frequency limit the flow through the orifice will act in a quasi steady manner with any induced flow rate through it being equivalent to that induced when a pressure differential is applied across the orifice and left for an infinite amount of time.

As the frequency of the incoming wave increases the pressure differential is applied for progressively smaller amounts of time. The system can be modelled as a solid mass being subjected to an oscillating force. If the pressure is applied over a uniform area and the fluid inside the orifice is considered incompressible then:

$$\delta p = \rho L \frac{du}{dt} \tag{A.67}$$

where L is the length of the orifice, u is the acoustic particle velocity and ρ is the fluid density. If δp is a simple sinusoidal signal ($\delta p = cos(\omega t)$) then this can be represented as the real part of the complex pressure signal:

$$\delta p = \mathrm{Re}(\tilde{p} e^{j\omega t}) \tag{A.68}$$

leading to the differential equation:

$$\rho L \frac{du}{dt} = \tilde{p} e^{j\omega t} \tag{A.69}$$

by assuming a solution of the form $u = Ae^{j\omega t}$, the particular solution can be found to be:

$$u = -\frac{\tilde{p} j}{\omega \rho L} e^{j\omega t} \tag{A.70}$$

The specific acoustic impedance is formally defined as:

$$\mathcal{Z}(\omega) = R + jX = \frac{\hat{p}}{\hat{u}} \tag{A.71}$$

and for the simple example given can be shown to be equal to:

$$\mathcal{Z}(\omega) = L\omega j \tag{A.72}$$

The real part of this transfer function (R) is referred to as the resistance and describes the in-phase response of the system. The imaginary part (X) is referred to as the reactance and describes the out of phase response of the mass of air to to the pressure fluctuation. Together they describe both the system gain and the phase shift or time lag between the peak incoming pressure wave and the peak disturbance to mass flow rate. As mentioned in the introduction, the time lag is especially important for the prediction of acoustic instabilities. The resistance, reactance, gain and phase shift can be linked through:

$$\phi_Z = tan(X/R) \tag{A.73}$$

and:

$$Z = |Z|e^{j\phi_Z} \tag{A.74}$$

If the gain ($|Z|$) is calculated incorrectly then the sensitivity of the system will be affected, if the phase shift is significantly different then the time it takes for a pressure wave to disturb the heat release rate will be calculated incorrectly and hence frequencies identified as those being most responsible for instability will be incorrect.

In the above analysis it is assumed that the fluid is inviscid and that there is no bias flow. The result therefore shows a viable value of reactance but no resistance. in order to find a theoretical expression for the resistance the effect of viscosity must be included in the analysis.

Kinsler et. al. [59] approximate the effect of the laminar boundary layer on a velocity fluctuation travelling through a pipe as:

$$R_w = \frac{m}{a}\sqrt{2\nu\omega} + \frac{m}{a}\sqrt{2\nu\omega}j \tag{A.75}$$

where m is the mass of fluid in the orifice and a is its radius. Another neglected effect relates to the radiation of sound from either end of the orifice. Again Kinsler et. al. [59] provides an estimate of:

$$2R_r = \frac{\rho\omega^2 S^2}{2\pi c} \tag{A.76}$$

where S is the surface area of the orifice and c is the speed of sound. It has also be noted that the vibration of the fluid within the orifice has an effect on the fluid on either side of the orifice, thus causing an apparent lengthening of the orifice which, for an orifice without a flange leads to [59]:

$$L' = L + 1.2a \tag{A.77}$$

Substituting these three relations back into Equation (A.72) forms a new dynamical system representative of an orifice with viscous and radiation effects included:

$$L'\frac{du_p}{dt} + \frac{2R_r + R_w}{S}u_p = \frac{\hat{p}}{\rho}e^{j\omega t} \tag{A.78}$$

This equation leads to an impedance equation:

$$\mathcal{Z}(\omega) = \frac{1}{\pi}\Big[\frac{\omega^2}{c} + \frac{L'}{a^3}\sqrt{2\nu\omega} + \frac{L'}{a^2}\big[\omega + \frac{1}{a}\sqrt{2\nu\omega}\big]j\Big] \tag{A.79}$$

where the first term is the effect of radiation, the second term represents the resistance due to viscosity, the third term represents the phase lag due to the mass inside the orifice and the 4th term is representative of the phase lag induced by viscosity. Up until this point the bias flow, induced by the mean pressure drop across the orifice, has been neglected. In order to take this effect into account it is best to consider the modified Cummings [27] equation developed by Luong et. al. [78]. this equation is developed by considering the Bernoulli equation across the orifice:

$$p_1 = p_{01} = p_{02} = p_2 + \frac{1}{2}\rho U_{ideal}^2 \tag{A.80}$$

with an associated ideal mass flow rate:

$$\dot{m}_{ideal} = \rho S U_{ideal} \tag{A.81}$$

Neglecting viscous or compressibility effects, the pressure drop across the orifice induces an ideal velocity at exit equalling U_{ideal} with a corresponding mass flow rate $\dot{m}_{ideal}$. However once viscosity is included the stagnation pressure downstream of the orifice will be lower than upstream leading to:

$$p_1 = p_{01} > p_{02} = p_2 + \frac{1}{2}\rho U^2 \tag{A.82}$$

with

$$\dot{m} = \rho S U \tag{A.83}$$

The ratio between $\dot{m}$ and $\dot{m}_{ideal}$ can be expressed as:

$$C_d = \frac{\dot{m}_{ideal}}{\dot{m}_{ideal}} \tag{A.84}$$

where C_d is the discharge coefficient and if rearranged and put in terms of the pressure difference:

$$p_1 - p_2 = \frac{\rho U^2}{2C_d^2} \tag{A.85}$$

If it is assumed that the acoustic velocity varies with pressure in a quasi steady way then in the absence of bias flow:

$$\tilde{p} = \frac{\rho u_p^2}{2C_d^2} \tag{A.86}$$

while in the case of bias flow:

$$p_1 - p_2 + \tilde{p} = \frac{\rho (U + u_p)^2}{2C_d^2} \tag{A.87}$$

Expanding and then subtracting the mean flow equation from the above equation leads to:

$$\tilde{p} = \frac{\rho (2U + u_p)}{2C_d^2} u_p \tag{A.88}$$

The quasi-steady assumption assumes that there will be no phase lag but this can be recovered by adding in the mass lag term to find:

$$L'\frac{du_p}{dt} + (U+\tilde{u})\frac{\tilde{u}}{C_d^2} = \frac{\hat{p}}{\rho}e^{j\omega t} \tag{A.89}$$

If it is assumed that $U >> |\tilde{u}|$ then the non linear part of this equation can be safely neglected yielding an impedance:

$$\mathcal{Z}(\omega) = \frac{1}{S}\left(\frac{U}{C_d^2} + \omega L' j\right) \tag{A.90}$$

An alternative approach was developed by Bellucci et al. [11] that casts the fundamental relationship between acoustic pressure and velocity as:

$$[j\omega L' + \zeta_{rc} + \zeta_{vis}]\tilde{u} = \frac{\tilde{p}}{\rho} \tag{A.91}$$

where ζ_{ac} is the contribution of the discharge coefficient and ζ_{visc} is the contribution of the boundary layers in the orifice. The Impedance in the case is given by:

$$\mathcal{Z}(\omega) = \frac{1}{\pi}\Big[\frac{U}{C_d r^2} + \frac{L'}{r^3}\sqrt{2\nu\omega} + \frac{L'}{r^2}\big[\omega + \frac{1}{r}\sqrt{2\nu\omega}\big]j\Big] \tag{A.92}$$

Comparing this equation to Equation (A.79) that neglects bias flow, it can be seen that the main effect of the neglected non-linear term might be the loss of acoustic energy due to the radiation of sound. The final equation is also aided by noticing that the effect of bias flow will almost certainly lead to the flow becoming turbulent, at least in the boundary layer or across the orifice face, therefore the viscosity ν might be replaced by ν_e the effective viscosity, a combination of the fluid viscosity and the apparent turbulent viscosity that approximates the net effect of turbulence on the flow:

$$\mathcal{Z}(\omega) = \frac{1}{\pi}\Big[\frac{U}{C_d r^2} + \frac{\omega^2}{c} + \frac{L'}{r^3}\sqrt{2\nu_e\omega} + \frac{L'}{r^2}\big[\omega + \frac{1}{r}\sqrt{2\nu_e\omega}\big]j\Big] \tag{A.93}$$

In this equation each of the terms represent, the bias flow effect on resistance, radiation effect on resistance, viscosity effect on resistance, mass effect on reactance and the viscosity effect on reactance.

A.5 Non-linear interaction and triggering between acoustic and hydrodynamic modes

The original analysis by Kirthy et al. [60] was completed in Cartesian coordinates for the case of a backwards facing step but is repeated here symbolically in polar coordinates for an axisymmetric burner. The Navier-Stokes equations in the low Mach number limit are written in the following form:

$$\mathbf{B}\frac{\partial q}{\partial t_1} + \epsilon\mathbf{B}\frac{\partial q}{\partial t_2} + \mathbf{N}(q,q) = \frac{1}{Re}\mathbf{L_v}(q) + \epsilon^2 F(\mathbf{x},t) \tag{A.94}$$

where q is now including fluctuations of density such that $q = [\rho, \mathbf{u}, p]$, F is the forcing of density, velocity and pressure, $\mathbf{B}$ has diagonal elements $[1, \rho, 0]$, $\mathbf{N}$ is the non-linear convective operator and $\mathbf{L}_v$ is the linear viscous operator. Using the method of multiple scales [96], the time dependence of q has been split into a fast timescale $t_1 = t$ and a slow timescale $t_2 = \epsilon t$. ϵ is a small number that enables the dynamics to be split into two different time scales and allows the effects of forcing to be separated from the natural unsteadiness of the flow. q is now expanded into an asymptotic series in ϵ:

$$q(x, r, t_1, t_2) = \overline{q}(x, r) + \epsilon q_1(x, r, t_1, t_2) + \epsilon^2 q_2(x, r, t_1, t_2) + ... \tag{A.95}$$

q_1 is assumed to be composed of a linear combination of global hydrodynamic instability modes q_{1m}:

$$q_{1m}(\mathbf{x}, t) = \tilde{q}_{1m}(x, r)e^{i(m\theta - \omega_{gm}t_1)} \tag{A.96}$$

Substitution of Equation (A.95) into Equation (A.94) and collecting the time averaged mean terms results in the Navier-Stokes equations for the time averaged mean flow. Collecting terms of order ϵ and ignoring the non-linear terms results in:

$$-\omega_{gm}\mathbf{B}\tilde{q}_{1m} + \mathbf{L}(\tilde{q}_{1m} : \overline{q}) = 0 \tag{A.97}$$

This is the same equation as Equation (4.14) in Chapter 4. The linear combinations of global hydrodynamic modes are weighted such that:

$$q_1(x, y, t_1, t_2) = \sum_{m=0}^{M} A_m(t_2)\tilde{q}_{1m}(x, y)e^{-i\omega_{gm}t_1} \tag{A.98}$$

The weightings $A_m(t_2)$ are the slowly varying amplitude of each global mode due to the presence of acoustic forcing. Finally, collecting the terms of order ϵ^2 and neglecting the non-linear terms results in:

$$\mathbf{B}\frac{\partial q_2}{\partial t_1} + \mathbf{L}(q_2 : \overline{q}) = -\frac{\partial q_1}{\partial t_2} + F(\mathbf{x}, t) \tag{A.99}$$

The form of the forcing term is assumed to be:

$$F(x, y, t_1, t_2) = -i\omega_a \rho_u \tilde{q}_a(x, y)e^{-i\omega_a t_1} \tag{A.100}$$

and using Equation (A.98) for an expression of q_1 results in:

$$-\omega_{gm}\mathbf{B}\tilde{q}_{1m} + \mathbf{L}(\tilde{q}_{1m} : \overline{q}) = -\sum_{m=0}^{M}\left[e^{i(\omega_a - \omega_{gm})t_1}\frac{dA_m}{dt_2}\mathbf{B}\tilde{q}_{1m}\right] - i\omega_a\rho_u\tilde{q}_a \tag{A.101}$$

Provided that the frequency difference between the global mode and the forcing is small one can write that:

$$\omega_a = \omega_{gm} + \epsilon\Omega_m \tag{A.102}$$

resulting in:

$$-\omega_{gm}\mathbf{B}\tilde{q}_{1m} + \mathbf{L}(\tilde{q}_{1m} : \overline{q}) = -\sum_{m=0}^{M}\left[e^{i\Omega_m t_2}\frac{dA_m}{dt_2}\mathbf{B}\tilde{q}_{1m}\right] - i\omega_a\rho_u\tilde{q}_a \tag{A.103}$$

whilst Equation (A.101) is a function of both t_1 and t_2, Equation (A.103) is a function of t_2 only. According to the *Fredholm Alternative Theorem* [55], A linear system:

$$Mx = b_0 \tag{A.104}$$

has a solution only if:

$$\langle \mathbf{u}, b_0 \rangle = 0 \tag{A.105}$$

for all vectors $\mathbf{u}$ such that:

$$M^\dagger \mathbf{u} = 0 \tag{A.106}$$

where $M^\dagger$ is the adjoint operator of the operator M and the angled brackets denote a suitable inner product. Translating this to our system of equations, Equation (A.103) only has a solution if:

$$\langle q^\dagger_{1m}, \sum_{m=0}^{M}\left[e^{i\Omega_{gm} t_2}\frac{dA_m}{dt_2}\mathbf{B}\tilde{q}_{1m}\right] + i\omega_a\rho_u\tilde{q}_a\rangle = 0 \tag{A.107}$$

for all $q^\dagger_{1m}$ such that they are solution to the adjoint stability equation:

$$-\omega^\dagger_{gm}\mathbf{B}^\dagger\tilde{q}^\dagger_{1m} + \mathbf{L}^\dagger(\tilde{q}^\dagger_{1m} : \overline{q}) = 0 \tag{A.108}$$

where the adjoint of the operators are defined such that they satisfy the following equation:

$$\langle -\omega^\dagger_{gm}\mathbf{B}^\dagger\tilde{q}^\dagger_{1m} + \mathbf{L}^\dagger(\tilde{q}^\dagger_{1m} : \overline{q}), \tilde{q}_{1m}\rangle = \langle \tilde{q}^\dagger_{1m}, -\omega_{gm}\mathbf{B}\tilde{q}_{1m} + \mathbf{L}(\tilde{q}_{1m} : \overline{q})\rangle \tag{A.109}$$

and the inner product is defined as:

$$\langle q_1, q_2 \rangle = \int_{\mathbf{X}} q_1^H \mathbf{P} q_2 d\mathbf{x} \tag{A.110}$$

where $\mathbf{P}$ is a diagonal matrix with elements $[1, \rho, 1]$ and H represents the conjugate transpose. Luckily, the adjoint modes satisfy the bi-orthogonality condition:

$$\langle \tilde{q}_{1m}^{\dagger}, \mathbf{B}\tilde{q}_{1n} \rangle = C\delta_{nm} \tag{A.111}$$

where C is a constant and δ_{nm} is the Kronecker delta symbol. This relation removes the sum in Equation (A.107) resulting (after integration in time) in an equation for A_m:

$$A_m(t_2) = \frac{\omega_a \rho_u \langle \tilde{q}_{1m}^{\dagger}, \tilde{q}_a \rangle}{\Omega_m \langle \tilde{q}_{1m}^{\dagger}, \mathbf{B}\tilde{q}_{1m} \rangle}[e^{-i\Omega_m t_2} - 1] \tag{A.112}$$

Finally, substitution of this into Equation (A.98) results in:

$$q_1(x, r, t) = \omega_a \rho_u \sum_{m=0}^{M} R_m[q_{1m} : q_a]\tilde{q}_{1m}(x, r)[e^{-i\omega_a t} - e^{-i\omega_{gm} t}] \tag{A.113}$$

with:

$$R_m[q_{1m} : q_a] = \frac{\langle \tilde{q}_{1m}^{\dagger}, \tilde{q}_a \rangle}{\Omega_m \langle \tilde{q}_{1m}^{\dagger}, \mathbf{B}\tilde{q}_{1m} \rangle} \tag{A.114}$$

A.6 Spray transport and evaporation

The mass, momentum and temperature (or enthalpy) must be calculated for each droplet at each timestep. The mass of each particle is controlled by the evaporation rate that is assumed to follow the response of Chin and Lefebvre [72]:

$$\frac{dm_d}{dt} = 2\pi D_d \left(\frac{\lambda}{c_p}\right)_g \ln(1 + B_M) \tag{A.115}$$

where dm_d is the mass of the droplet, λ is the mean thermal conductivity of the gas, c_P is the specific heat, the subscript d refers to the droplet, the subscript g refers to the gas and B_M is the mass transfer number defined as:

$$B_M = \frac{Y_{F_s}}{1 - Y_{F_s}} \tag{A.116}$$

The mass fraction of fuel vapour at the droplet surface (Y_{F_s}) can be related to the fuel vapour pressure (p_{F_s}) the air pressure (p), and the ratio of the molecular weights of the air and fuel (M_A/M_F):

$$Y_{F_s} = \left(1 + \left(\frac{p}{p_{F_s}} - 1\right)\frac{M_A}{M_F}\right)^{-1} \tag{A.117}$$

The vapour pressure of the fuel at the drop surface can be estimated from the Clausius-Clapeyron equation:

$$p_{F_s} = exp\left(a - \frac{b}{T_s - 43}\right) \tag{A.118}$$

where T_s is the surface temperature and a and b are constants that are a function of the fuel used. This model assumes that the fuel droplets are spherical and that the effect of convective cooling, caused by a velocity difference between the fuel and air, is negligible. The fuel droplets are assumed to have a constant temperature throughout, equal to the droplet surface temperature, the change of which is calculated as:

$$\frac{dT_d}{dt} = \frac{Q - Q_c}{C_{p,f} m_d} \tag{A.119}$$

Q is the heat transferred from the air to the droplet of fuel through conduction while Q_c is the heat absorbed through evaporation of the fuel:

$$Q = 2\pi D\kappa(T - T_d)ln(1 + B_M)/B_M \tag{A.120}$$

$$Q_c = \frac{dm}{dt}L \tag{A.121}$$

The diameter of each droplet is approximated from the droplet mass using:

$$D_d = [6m_d/(\pi\rho_d)]^{1/3} \tag{A.122}$$

Any heat transferred from the flow to the droplet must also be subtracted from the flow. In order to achieve this an additional source term must be added to the energy equation to be solved in each cell. as such Equation (A.23) becomes:

$$\frac{d}{dt}\int_{V_P} \rho h dV + \int_{S_p} \rho h \mathbf{u} \cdot \mathbf{dS} = \int_{S_p} k\nabla T \cdot \mathbf{dS} + \int_{V_P} [\mathbf{u} \cdot \nabla p + \overline{\overline{\mathbf{T}}} : \nabla \mathbf{u}] dV + Q_P \tag{A.123}$$

where

$$Q_P = \sum_{i=1}^{N_d} Q_{c,i} - Q_i \tag{A.124}$$

where N_d is the number of particles within the Pth cell. Similarly the loss of droplet mass must result in a gain of fluid mass. Therefore Equation (A.22) must be altered with a source term $\dot{m}_P$:

$$\frac{d}{dt}\int_{V_P} \rho dV + \int_{S_P} \rho \mathbf{u} \cdot \mathbf{dS} = \dot{m}_P \tag{A.125}$$

with:

$$\dot{m} = \sum_{i=1}^{N_P} \left(\frac{dm_d}{dt} \right)_i \tag{A.126}$$

The change in droplet momentum is calculated as:

$$\frac{d(m_d \mathbf{u}_d)}{dt} = m_d \frac{d\mathbf{u}_d}{dt} + \mathbf{u}_d \frac{dm_d}{dt} \tag{A.127}$$

where the second term represents the momentum transfer of fuel being evaporated and is calculated using Equation (A.115). The first term is due to the body forces on the droplet of the fuel and can be calculated as:

$$m_d \frac{d\mathbf{u}_d}{dt} = \mathbf{F}_d + \mathbf{F}_v + \mathbf{F}_b \tag{A.128}$$

where $\mathbf{F}_d$ is the drag force, $\mathbf{F}_v$ is the virtual mass force and $\mathbf{F}_b$ are the body forces. The body forces are assumed to be negligible. The drag force is calculated from:

$$\mathbf{F}_d = C_d(Re_d) \frac{1}{2} \rho_a |\mathbf{u} - \mathbf{u}_d| (\mathbf{u} - \mathbf{u}_d) S_d \tag{A.129}$$

where S_d is the frontal area of the fuel droplet and Re_d is the Reynolds number based on the droplet diameter and velocity difference between the droplet and the air:

$$Re_d = \frac{\rho_a |\mathbf{u} - \mathbf{u}_d| D_d}{\mu_a} \tag{A.130}$$

The drag coefficient (C_d) is a function of the droplet Reynolds number and can be approximated by:

$$C_d = \begin{cases} 24,000 & Re_d \leq 0.001 \\ 24/Re_d & 0.001 \leq Re_d \leq 0.2 \\ 24(1 + 0.15 Re_d^{0.687})/Re_d & 0.2 \leq Re_d \leq 1000 \\ 0.44 & 1000 \leq Re_d \leq 200,000 \\ 0.10 & Re_d \geq 200,000 \end{cases} \tag{A.131}$$

The virtual force $\mathbf{F}_v$ is due to the acceleration of the additional mass associated to the fluid surrounding the droplet that must be accelerated as the particle accelerates:

$$\mathbf{F}_v = -C_{vm} \rho_a \Omega_d \frac{d(\mathbf{u}_d - \mathbf{u})}{dt} \tag{A.132}$$

C_{vm} is the virtual mass coefficient and is assumed to be equal to 0.5.

The temperature of the droplet, and therefore its surface can be related to the droplet enthalpy as:

$$h_d = c_{p,d} T_d \tag{A.133}$$

and the change in enthalpy per droplet is calculated as:

$$\frac{d(m_d h_d)}{dt} = m_d \frac{dh_d}{dt} + h_d \frac{dm_d}{dt} \tag{A.134}$$

The first term on the right describes the heat removed from the flow due to the heating up of each droplet and the second term relates to the addition of heat to the flow due to evaporation induced mass transfer from droplet to the gas phase. Using Equations (A.115), (A.127) and (A.134) the mass, momentum and enthalpy of each droplet can be updated at each timestep. Finally the location of each droplet is updated by:

$$\frac{d\mathbf{x}_d}{dt} = \mathbf{u}_d \tag{A.135}$$

In order to reduce the computational cost of the simulation it is often considered that the spray and the flow are decoupled meaning that while the flow velocity has an effect on the spray, the spray velocity has no effect on the flow. This approach requires no modification of Equation (A.19), however if it is assumed that the force imparted to the spray results in an equal but opposite force on the flow then Equation (A.19) becomes:

$$\frac{d}{dt}\int_{V_p} \rho \mathbf{u} dV + \int_{S_i} \rho \mathbf{u}(\mathbf{u}\cdot \mathbf{dS}) = \int_{S_i} \overline{\overline{\mathbf{T}}} \cdot \mathbf{dS} + \int_{V_i} \rho \mathbf{f} dV + \mathbf{F}_P \tag{A.136}$$

where

$$\mathbf{F}_P = \sum_{i=1}^{N_P} \left(\frac{d(m_d \mathbf{u_d}}{dt} \right)_i \tag{A.137}$$

Finally, as equations A.115, A.127, A.135 and A.134 must be solved for each droplet tracked, a large fuel mass flow rate may result in a prohibitively large number of droplets and as such, a large number of equations to be solved. One solution is to fix the number of droplets to be injected and simply consider that each droplet represents the mass, velocity, location and energy of a parcel of droplets. The number of droplets within each parcel is defined by the number density $N_d(m_d, \mathbf{u}_d, \mathbf{x}_d, h_d)$.

Appendix B

Assessing the Dynamic Response of the Code PRECISE-UNS

In order for the acoustic impedance of the injector to be correctly obtained the code used to calculate the interaction between acoustic waves and velocity fluctuations must be able to advect the incoming acoustic waves at the correct propagation speed, maintain the amplitude of the wave and maintain the correct frequency. The ability of a code to maintain the correct amplitude of wave as it propagates through a medium is referred to as the dissipation of the numerical scheme. In the case of a finite difference scheme the dissipation of the method can be calculated by considering the size of the neglected even terms in the taylor expansion, however due to the more complicated numerics involved with the finite volume method, it is significantly easier to test the dissipation of the scheme through the use of carefully designed numerical experiments. If the pressure at one location in space is $p_m(k\Delta x)$ and the pressure at the next grid point is $p_{m+1}(k\Delta x)$ where k is the wavenumber of the incoming pressure wave and Δx is the grid spacing, then a gain function can be defined as the transfer function that links these two measurements of pressure:

$$G(k\Delta x) = \frac{\hat{p}_{j+1}(k\Delta x)}{\hat{p}_j(k\Delta x)} \tag{B.1}$$

Assuming that this transfer function is linear allows it to be calculated by comparing the wave power at two points and taking the reciprocal power of the number of grid points.

Waves may not just be damped, but also dispersed to other frequencies. This can be more of a problem than the dissipation of waves as without the assumption that the waves maintain their frequency and hence wavelength there is no way to track individual waves once they are in an environment with a number of other waves of other frequencies. Both dissipation and dispersion ultimately limit how large the grid spacing can be before waves of the desired frequency cannot be propagated throughout the mesh. For the regions of the mesh where the velocity scales are of interest, the mesh size will be limited by the need to resolve them rather than the pressure fluctuations. However velocity scales may only be of interest in very small parts of the mesh, such as the regions where there is significant dissipation of fluid momentum such as boundary layers and wakes or in the case of two phase flows, where the convection, evaporation and mixing of the fuel needs to be accurately resolved. As such, it is ideal to design a mesh where the grid spacing within the primary zone and flow passages of the injector is decided by the size of the velocity scales of interest while in the upstream plenum and downstream duct, the size of pressure waves dominate the choice of grid spacing. By doing this, the mesh is only as resolved as it needs to be, reducing computational cost.

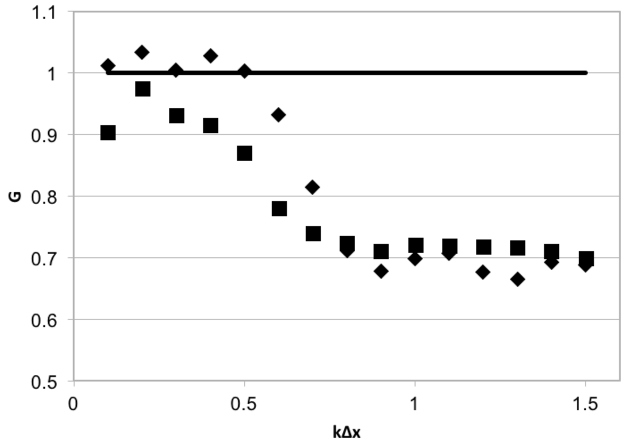

Figure B.1: Wave gain amplitude per grid point as a function of $k\Delta x$. Line: Ideal, Diamonds: Laminar with slip walls. Squares: Turbulent with no-slip walls.

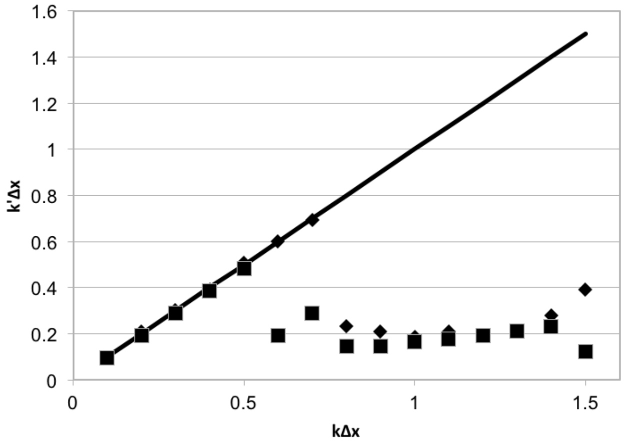

Figure B.2: Wave output normalised wavenumber after moving grid point $k'\Delta x$ as a function of input normalised wavenumber $k\Delta x$. Line: Ideal, Diamonds: Laminar with slip walls. Squares: Turbulent with no-slip walls.

In order to choose cells of the correct size where pressure fluctuations dominate the unsteady flow field, two experiments were performed to test the response of the code PRECISE-UNS, both were conducted using the same low resolution structured duct mesh as described in Section 3.5.3 The first experiment assumes that there is no bias flow and hence no mean pressure gradient along the duct and as the reynolds number is zero, the flow can assumed to be laminar. A pressure wave with a magnitude of 63 Pa was then injected from the outlet of the duct and allowed to propagate upstream. Each simulation was allowed to run for 12 cycles with the pressure being measured at the injection plane and a point 10 grid cells upstream. As the dispersion properties of the code may modify the wavelengths of the pressure waves, the experiment had to repeated for every desired frequency to be tested to ensure that there was only one wave propagating in the duct at once. The quantities of interest when describing the wave at any point are its wavenumber multiplied by the grid spacing ($k\Delta x$) and its pressure amplitude ($|p|$). A $k\Delta x$ value of .1 corresponds to roughly 63 grid spaces per wavelength while a $k\Delta x$ value of 1 corresponds to just over 6 grid spaces per wavelength. The highest possible $k\Delta x$ value is equal to π as this corresponds to the Nyquist frequency of the grid however as only very high order, low dispersion schemes can capture theses high values of $k\Delta x$, the mesh was only subjected to a maximum $k\Delta x$ of 1.5.

Results from this first experiment show that PRECISE-UNS operating with 2nd order linear upwind spacial differencing for velocity and 2nd order symmetric differencing for pressure with a second order temporal scheme is non dispersive for waves with a $k\Delta x$ smaller than around 0.7 and non dissipative for waves with a $k\Delta x$ smaller than around 0.5. This corresponds to 9 grid points and 12.5 grid points per wave respectively. Somewhat worryingly, the dissipation was found to be negative at some wavelengths meaning that the pressure waves were growing slightly as the wave moved upstream which could lead to numerical instabilities however this gain is only slight and the waves can be seen not to continue growing in magnitude as they move further upstream. In order to ensure that the effects of dispersion and dissipation were independent of timestep and hence CFL number, the experiment was repeated for some of the frequencies with a timestep half of what was used originally.

As described in Section A.4 the presence of a turbulent boundary layer will cause additional dissipation and dispersion of an acoustic wave. In order to test this effect, the above numerical experiment was repeated except that a bias total to static pressure ratio of equal to that described in Section 5.3.1 was applied across the duct with the k-epsilon turbulence model selected and wall functions enabled. The extremely poor spatial resolution in the wall normal direction is likely to exacerbate the dispersion and dissipation effects of the turbulent boundary layer hence indicating that the dissipation and dispersion effects present in the acoustic impedance tests is likely to lie somewhere between these two cases described in this section. In the case of bias flow and turbulence model engaged, the limiting dispersion and dissipation $k\Delta x$ was found to be 0.5 or roughly 12.5 grid points per wavelength.

Appendix C

Assesing the quality of characteristic boundary conditions in PRECISE-UNS

Using the same mesh as in Section 5.3.1 a test was carried out to determine the effectiveness of the charateristic boundary conditions, with particular emphasis on the inlet boundary conditions and the values of the relaxation coefficients. In this case the inlet total pressure was set to be 108000 Pa and the downstream static pressure to 101325 Pa. The simulation was run for 2000 iteration assuming steady flow with the $k - \epsilon$-RNG turbulence model. This resulted in a uniform flow through the duct corresponding to a mass flow rate of approximately 20.5 kg/s. Once this steady state solution was found the simualtion was then run for a total of 5000 iterations with a non-reflective pressure outlet boundary condition with a 1000 Pa 50 Hz acoustic wave being injected from the same boundary. The upstream boundary was given a constant mass flow condition while the relaxation factor was altered.

C.1 $K_v = 0$

This makes the upstream truly non-reflective but results in no fixing of the upstream velocity or pressure. Figures C.1a and C.1b show the pressure and axial velocity at several locations along the duct as a function of time. It can be seen that the pressure and the axial velocity in the duct remains constant until the wave travelling upstream reaches the monitoring point. The magnitude of the acoustic wave can be seen to remain constant in time and in space showing that the wave sucessfully enters and exits the domain without reflection.

C.2 $K_v = 0.1$

Here the relaxation coefficient is set relatively high to force the inlet to the mass flow rate set by the user. It can be seen from the pressure and velocity traces (Figures C.2a and C.2b) that the pressure and velocity are both pulled down to the desired quantity at the inlet sending a shock

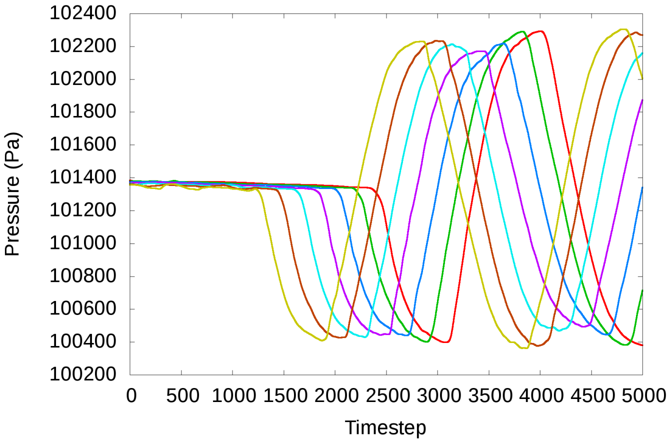

(a) The pressure at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.0$.

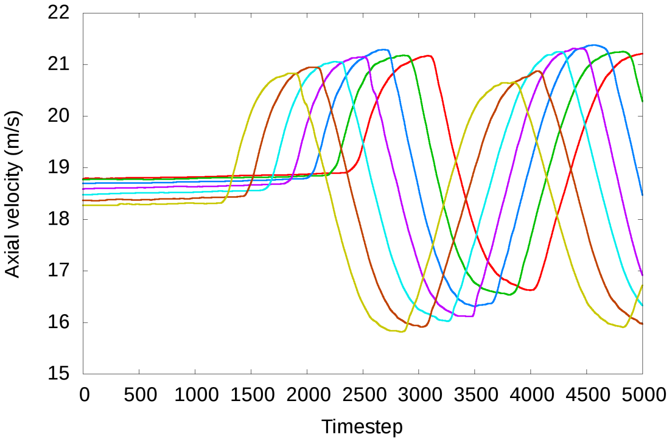

(b) The axial velocity at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.0$.

Figure C.1: velocity and pressure probes for the case where $K_v = 0.0$.

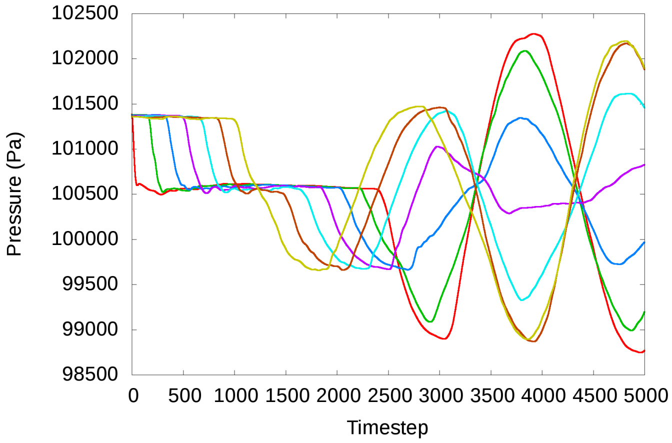

(a) The pressure at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.1$.

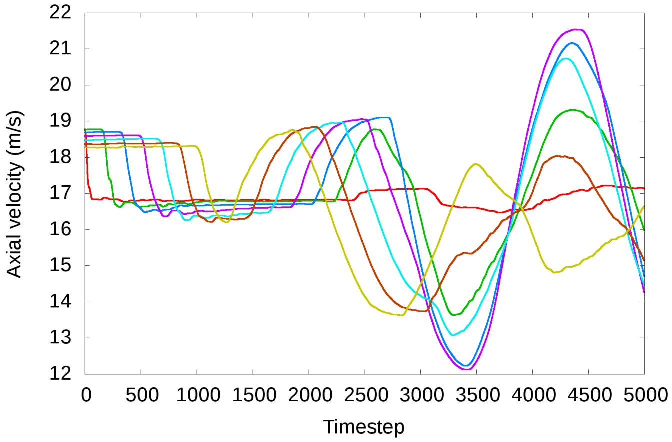

(b) The axial velocity at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.1$.

Figure C.2: velocity and pressure probes for the case where $K_v = 0.1$.

wave from the inlet towards the exit. The velocity fluctuations at the inlet are seen to be minimum while pressure fluctuations are maximum showing that the inlet has become a psuedo velocity node.

C.3 $K_v = 0.001$

In this case the relaxation factor was reduced by two orders of magnitude such that it was more non-reflective. Figures C.3a and C.3b show that the inlet pressure and velocity are slowly forced to the desired values while the inlet velocity is still allowed to oscillate as the acoustic wave passes out of the inlet. The magnitude of the acoustic wave is seen to be somewhat consistent across the length of the duct and in time.

C.4 $K_v = 0.001$ without initial bias flow

In this case the simulation was run without the initial flow generated from the steady solution as a test to see how stable the boundary condition is. The velocity at the inlet jump quickly from zero and then asymptotically approaches the desired inlet velocity (Figure C.4b). Figure C.4a shows that there is a slight spike in pressure but the inlet remains relatively non-reflective.

C.5 $K_p = 0.00001$ without initial bias flow

This is another case where the initial velocity was set to zero except this time the inlet boundary was set to approach a desired inlet pressure rather than velocity. Initially the pressure spikes at the inlet before falling quickly to the desired inlet value (Figure C.5a). This figure also shows that the relaxation constant K_p is set too high as the pressure remains constant at the inlet leading to the wave been reflected at the inlet. Figure C.5b shows that the velocity initially overshoots before slowly returning towards the mean value. When the downstream acoustic wave approaches the inlet is causes the inlet velocity to fluctuate, a standing wave then forms in the duct.

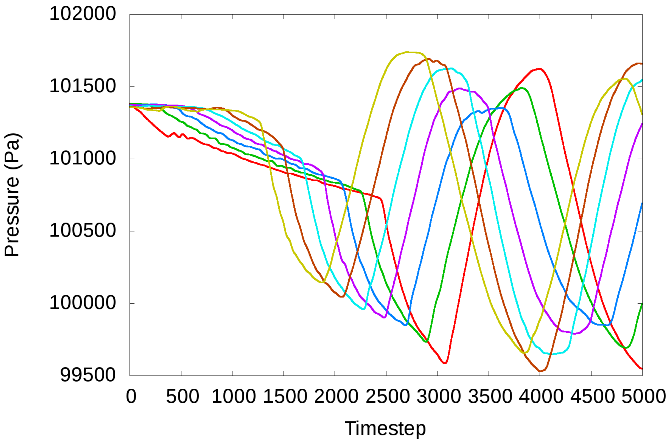

(a) The pressure at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.001$.

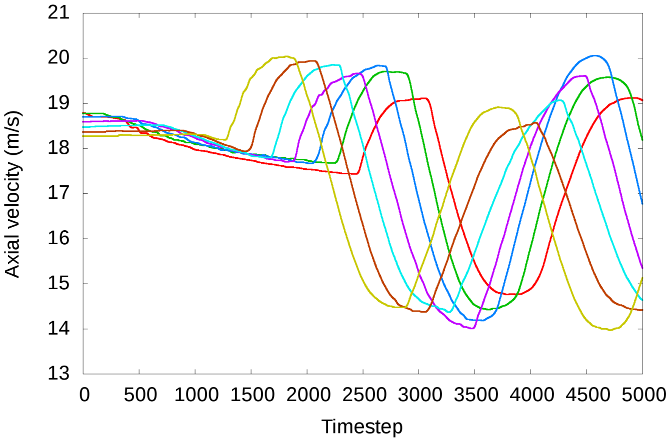

(b) The axial velocity at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.001$.

Figure C.3: velocity and pressure probes for the case where $K_v = 0.001$.

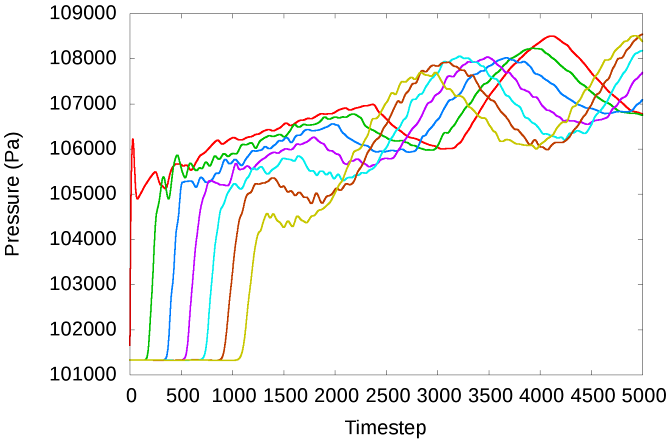

(a) The pressure at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.001$.

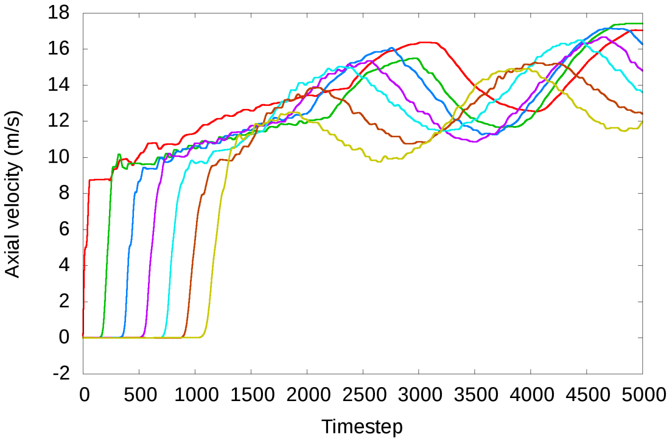

(b) The axial velocity at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow) for the case where $K_v = 0.001$.

Figure C.4: velocity and pressure probes for the case where $K_v = 0.001$ without bias flow.

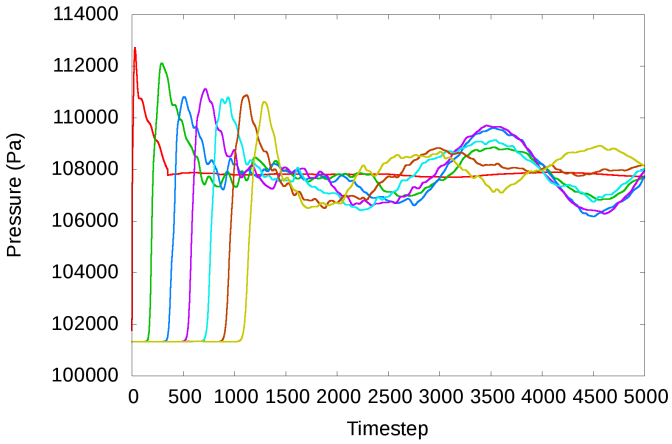

(a) The pressure at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow).

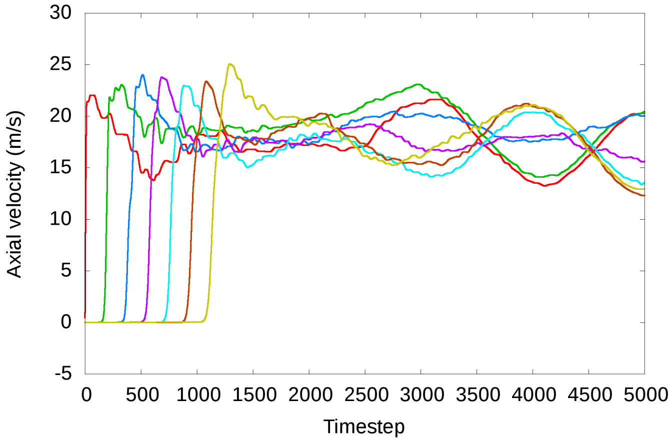

(b) The axial velocity at x=0.0 m (red), x=0.5 m (green), x=1.0 m (blue),1.5 m (purple), 2.0 m (cyan), 2.5 m (brown) and 3.0 m (yellow).

Figure C.5: velocity and pressure probes for the case where $K_p = 0.00001$ without initial bias flow.

Appendix D

SCARLET rig mass split

The SCARLET rig is an experimental facility located at DLR Cologne but owned and run by Rolls-Royce. It is composed of a single sector lean-burn combustion chamber appended at each by a long duct, acoustic sirens and acoustic dampers. The purpose of this rig is to determine the Flame Transfer Function of lean-burn fuel injector designs. The rig is operated such that the upstream pressure and mass flow rate are maintained to be constants and acoustic waves are introduced from downstream or upstream as required, the dampers are designed to prevent acoustic waves from reflecting from the upstream or downstream boundaries. As the rig is a hot combusting facility the burner must be thermally protected through the use of cooling flows that are introduced both beside the injector and through the flame tube liner. Whilst the total mass flow is known the mass flow split between the injector and liner is not known. Luckily the experimental facility contains pressure probes that are capable of measuring the pressure drop across the injector, liner and the meter plate placed between the upstream duct and outer annulus.

D.1 Geometry and Probe locations

The flow first enters the rig through a metered inlet past a large upstream damper and through an orifice plate to the upstream duct section. Two acoustic sirens are connected to two sides of the upstream duct that allow acoustic waves to be generated at very high amplitudes. The acoustic waves then travel down the duct towards the injector and past five microphones that are mounted in the wall of the upstream duct. Once the air passes through the combustion chamber there is a further duct with another five wall mounted microphones. At the end of this duct sits another two acoustic sirens, an orifice plate and downstream acoustic damper. It is noted in [147] that the measurement of pressure at the wall in simulations often contains high frequency noise due to the presence of large scale flow structures and a much cleaner pressure signal can be obtained by placing the probes in the centre of the duct instead. Unlike a microphone a pressure probe placed in the computational domain can also be used to measure mean pressure and as such the probe at position ‘p5’ (see Figure D.1) is used to record the upstream ‘P30’ pressure value. Two additional pressure probes were placed into the computational domain as shown by Figure D.2. ‘I22’ is located just inside the flame-tube and is used to measure the pressure drop across the injector and liner while I22 in located in the outer annulus and is used to measure the pressure

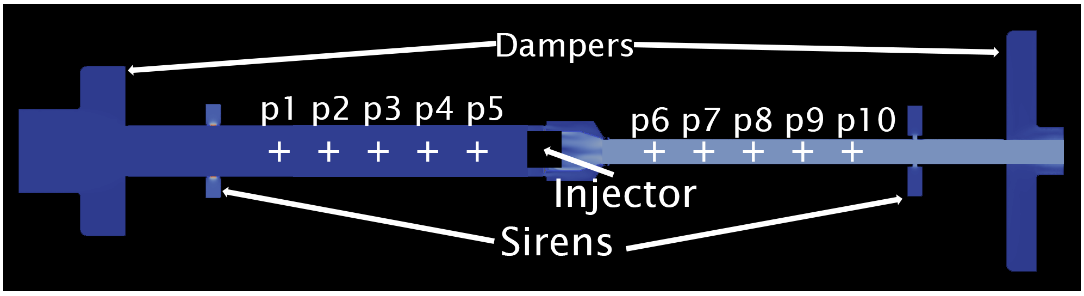

Figure D.1: The SCARLET rig showing the positions of the injector, dampers, sirens and pressure probes used to analyse the acoustic waves. Flow is from left to right.

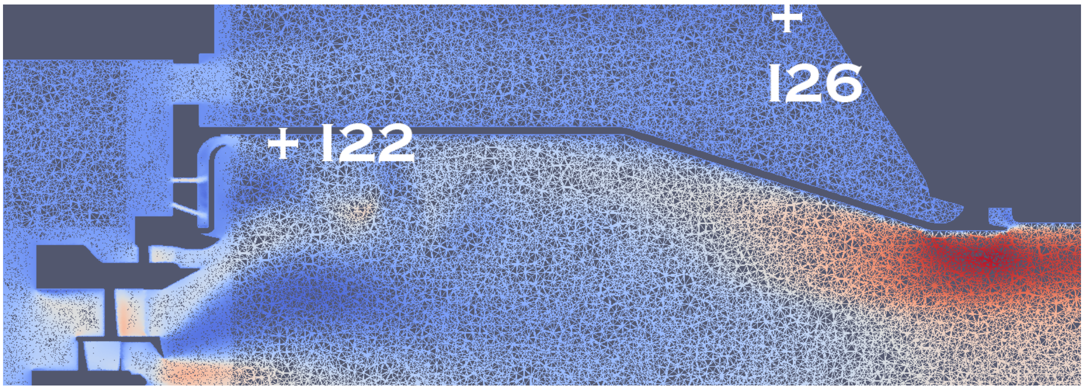

Figure D.2: The SCARLET rig combustion chamber showing the locations of the two probes used to measure the pressure inside the combustion chamber (I22) and the outer annulus (I26).

drop across the meter panel and across the liner.

D.2 Test simulation

It took some time for the correct geometry to be obtained for the SCARLET simulation cases and as such a similar but slightly different geometry was used first to test the computational methodology and find the mass split between the injector and cooling flow. A mesh suitable for acoustic simulations was generated of this preliminary geometry using the RRD method (section 2.5.3) with the upstream and downstream ducts neglected. Once a satisfactory mesh was obtained the mesh was extruded upstream and downstream using a regular cell spacing of $0.072D$, this resulted in a mesh of approximately 25 million cells. The simulation was run for 6000 timesteps of $2 \times 10^{-1} t^{*}$ in order to obtain a psuedo-steady state flow. The pressure was measured at the ‘p5’, ‘I22’ and ‘I28’ locations at each timestep (see Figure D.3). The total mass flow of the rig was known and the mass split between the cooling flow and injector flow is set by an effusion boundary condition that subtracts mass from the outer side of the combustor liner and injects it into the combustion chamber in order to simulate the effect of the cooling holes. This cooling flow was set as 20% of the total mass flow

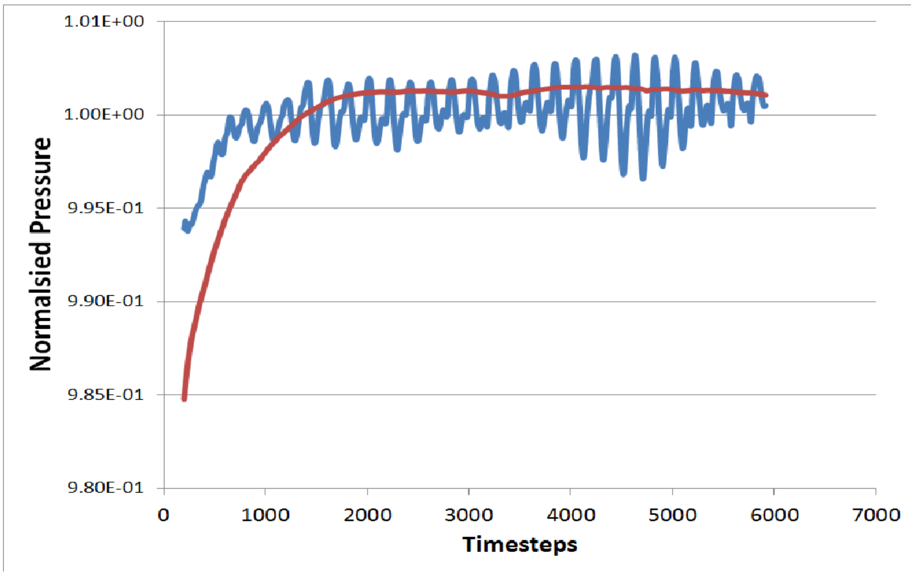

Figure D.3: The pressure inside the combustion chamber (I22) (blue) and the outer annulus (I26) (red) normalised by thier mean values throughout the simulation.

D.3 Calculation of effective area

The effective area is a measure of how much of a pressure drop is required to generate a certain mass flow though a component. It is related to the discharge coefficient:

$$A_{eff}(Re, Ma) = C_D(Re, Ma)A \tag{D.1}$$

Due to this dependence, if the cross sectional area of the device (A) is a constant, then the effective area is a function of Reynolds number and Mach number only. To simply analysis we will assume that the Reynolds number is sufficiently high to ensure independence of Reynolds number and it will be assumed that the Mach number is sufficiently low that the flow may be considered to be incompressible. The pressure drop across the meter panel can therefore be related to the mass flow by:

$$\dot{m}_{MP} = \rho A_{eff,MP}\sqrt{2\Delta P_{MP}/\rho} \tag{D.2}$$

and hence, this equation can be re-arranged to find the effective area of the meter plate:

$$A_{eff,MP} = \frac{\dot{m}_{MP}}{\sqrt{2\rho\Delta P_{MP}}} \tag{D.3}$$

As the mass flow is given by the cooling flow split and total mass flow, pressures obtained during a test simulation can be used to determine the effective are of the meter panel. The same procedure can be used to determine the effective area of the injector.

D.4 Determination of mass flow split

From the trial simulation the pressure drop across the meter panel was recorded as being 0.27% of the upstream pressure and across the liner as being 6.78% of the pressure in the outer annulus. Experimental results taken at the same conditions show that the meter panel pressure drop should be 1.41% and the liner pressure drop 1.5%. This indicates that there is not enough mass flowing through the cooling system and the mass split should be increased.

Taking the effective area of the meter panel as being constant and by assuming that the pressure drop across the meter panel will be the same as in the experimental case a new estimate of the mass flow through the cooling system can be generated using Equation (D.2). This then lead to a new estimate of the cooling flow mass split of 35%.