Focused Laser Differential Interferometry

Thesis by
Joel Michael Lawson

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2021
Defended 8th April, 2021

Joel Michael Lawson
ORCID: 0000-0002-3042-0909

ACKNOWLEDGEMENTS

Thanks must first go to my advisor, Prof. Joanna Austin, for all her guidance and encouragement during my time in GALCIT, and of course for accepting me into her research group! Beyond her academic advising, Joanna is unfailingly kind and pleasant to work for. She genuinely cares about the well-being of her students, which was made particularly clear during all the disruption due to the coronavirus pandemic: her first concern was always the mental and financial health of each one of us. Next is Prof. Joe Shepherd: he stepped in to help our group in a time of need, and then he never quite stepped out! Despite having a research group of his own, plus heavy responsibilities in his administrative role, he not only managed to find the time to meet regularly with each of us, but also took a strong interest in all of our projects. A never-ending fount of new ideas (some of which I attempted!), with a deep knowledge of the literature, and loads of practical experience: I have been very fortunate to have him as another advisor. Finally, to complete the hypersonic professorial triad: Prof. Hans Hornung. A legendary figure in the community, and for good reason. His insight and physical intuition are amazing, I have learned a lot from him over the course of our chats. He is always willing to help, and despite his huge knowledge never condescending in the face of some very silly questions from me.

I would like to thank all the other faculty I have interacted with or been taught by. In particular: Prof. Beverley McKeon, especially for entrusting me as the mentor for the Base11/Aerospace Mentorship Program. That was a daunting task, but I thoroughly enjoyed it and all the students I got to teach and work with. Thanks also to Prof. Guillaume Blanquart for five generous years of tasty group lunches, good advice on how to negotiate grad student life, incisive and helpful observations from a different point-of-view, and of course his inability to walk past any student he knows without making a humorous remark. Lastly, thanks to Prof. Tim Colonius, for being a member of my candidacy and thesis committees.

To the past and present members of the Caltech Hypersonics Group, and our “sister” group, the Explosion Dynamics Lab: it has been a pleasure working (or slacking off!) with you. The members of the 2015-16 GALCIT/MCE cohort, too numerous to name, just as all those late nights doing problem sets were too numerous to count! I am glad to have formed an enduring group of friends largely from among these students.

My time at Caltech has been greatly facilitated by many administrative staff: Christine Ramirez, Denise Ruiz, Liza Bradulina, Barbara Slater, and Jamie Meighen-Sei. Our group's engineer, Bahram Valiferdowsi, has been helpful for all manner of practical issues, as well as irrepressibly good-humored. Laura Flower Kim and Daniel Yoder in ISP for all your visa-wrangling and friendliness.

Beyond the boundaries of the Caltech campus, I am very grateful to The Chapel at Pasadena [formerly Prism Church] and all the friends we have made there. Arriving in the US with no friends or family, Angela and I got plugged into a great community within our first few weeks and have been there ever since. The intensity of the first year plus having no car at that time meant that church and community group were some of the few chances to get out of the "Caltech bubble".

Further away, across the width of the Pacific Ocean, I am thankful to many people back home in New Zealand. Thank you to the many friends who have stayed in touch. Luckily Los Angeles is a major transport hub and as a result I have had the chance to show many of these friends around our adopted home city. I have missed out on many milestones in my friend group, but whenever we go home, it's like going straight back to old times.

To my family, especially Mum and Dad: thank you for all your love and support throughout my life. I believe that me finishing my Ph.D. can ultimately be traced back to you instilling in me a love of reading, and cultivating my curiosity for scientific and historical knowledge. The theoretical part of grad school is only part of the equation, however; lab work and engineering design requires a quite distinct set of skills. I am greatly indebted to all the time spent in Grandad's workshop where some of his meticulousness and creativity rubbed off on me!

Lastly and most importantly, thank you to my wife Angela. I know how much you have sacrificed to follow me over here in my pursuit of a Ph.D. We arrived in the US just 4 months after getting married, and now our 6$^{\text{th}}$ anniversary has recently gone by. As this chapter draws to a close, I hope you will find much joy in whatever comes next for us. I have missed you sorely over the last several months that we have been apart, and I look forward to your return soon.

This work was partially supported by the Office of Naval Research award N00014-16-1-2503 with Dr. Eric Marineau and the Air Force Research Laboratory award STTR FA8651-17-C-0071 with Dr. Daniel Reasor.

This thesis is dedicated to:
my grandpa, James Ronald “Jim” Lawson (1935–2002),
and my grandfather-in-law, Jia Chen 贾琛 (1933–2017).
Both were teachers, then trained other teachers.
Lovers of education, it would please them to see me obtain my Ph.D.

On the 9th of April, 2021 (the day after my defense),
my Pasadena church pastor, Dr. Chuck Ryor,
tragically died in a motorcycle accident.
Having done a Ph.D. himself, he often encouraged me.
Despite cheerfully admitting he wouldn’t understand any of it,
he took the time to attend my defense—the last time I ever spoke with him.
Hence, I want to additionally dedicate this thesis to Chuck.

Of making many books there is no end,
and much study wearies the body.

ECCLESIASTES 12:12

ABSTRACT

The focused laser differential interferometer (FLDI) is a non-imaging optical diagnostic that is sensitive to density disturbances. A distinguishing feature is reduced sensitivity away from the focal plane of its beams. The spatial resolution is sub-mm, and the temporal resolution is restricted only by photodetector bandwidth, typically > 10 MHz. These traits make FLDI particularly suited to measurements in hypervelocity ground-testing facilities, where the low densities, short time-scales, and harsh environments preclude the use of intrusive diagnostics. Line of sight integration issues associated with other optical techniques are therefore minimized, a distinct advantage for measurements in impulse facilities, where the core flow of interest is often surrounded by highly-turbulent shear layers.

The systematic design principles for single and double FLDI systems are discussed, based on ray transfer matrix analysis combined with Gaussian optics. A detailed guide is presented for the practicalities of aligning, calibrating, and operating an FLDI.

A modular numerical implementation of Schmidt and Shepherd's FLDI ray-tracing model is developed, capable of accepting arbitrary flow-fields defined via analytical expressions, simulation coupling, or experimental datasets. This numerical implementation is used to perform the first comprehensive experimental validation of the model, using known static and dynamic phase objects. Quantitatively-accurate predictions of the response of real FLDI systems are obtained. Importantly, the spatial sensitivity of the instrument is found to be dependent on disturbance wavelength, with scaling matching that predicted analytically from the model. Propagating shock waves are used as another highly-dynamic test phase object, and it is shown that FLDI maintains its theoretical performance at sub-μs time-scales.

The validated ray-tracing model is used to develop analytical expressions for the response of FLDI to propagating plane waves, extending on the results of Schmidt and Shepherd, and Settles and Fulghum. For the first time, the inverse problem is solved for this class of flow-field, allowing the density fluctuation spectrum to be recovered quantitatively from FLDI phase shift data. This approach is validated using synthetic flow-fields with the numerical ray-tracing scheme, and is also compared with the approximate approach introduced by Parziale et al.

FLDI is used to make freestream density fluctuation measurements on two facilities:

a conventional blowdown tunnel, and an expansion tube. On the conventional tunnel, a comparison is made between pitot-probe and FLDI measurements after converting both to freestream pressure fluctuation spectra. A modification of Stainback and Wagner's theory, incorporating recent numerical results from Chaudhry et al., is used to interpret the pitot data, while the new inversion algorithm is applied to the FLDI data. Close agreement is found between the two sets of spectra, showing that accurate quantitative data can be obtained with FLDI, and used to extend spectra beyond the pitot bandwidth.

On the expansion tube, the theory of Paull and Stalker for freestream noise originating in the driver gas is investigated. Their proposed relationship between freestream density fluctuations and the primary interface sound speed ratio is not observed. Spectral banding is also absent, however this is expected due to the relatively low secondary expansion strengths. The envelope of accessible conditions is somewhat restricted due to the low mean freestream densities that lead to signal-to-noise issues.

Significant performance improvements can still be made to FLDI, in terms of its noise and bandwidth limitations, and to the spatial localization of its sensitive region; suggestions are given for possible approaches. With the ray-tracing model now validated, it can be used to optimize FLDI, or even to suggest derivative instruments based on similar principles.

PUBLISHED CONTENT AND CONTRIBUTIONS

Lawson, J. M. and Austin, J. M. (2021). "Focused laser differential interferometer response to shock waves". In: *Measurement Science and Technology* 32.5, p. 055203. DOI: 10.1088/1361-6501/abdbd3.

Lawson, J. M., Neet, M. C., Grossman, I. J., and Austin, J. M. (2020). "Static and dynamic characterization of a focused laser differential interferometer". In: *Experiments in Fluids* 61.187. DOI: 10.1007/s00348-020-03013-6.

Lawson, J. M. and Austin, J. M. (2020). "Expansion tube freestream disturbance measurements using a focused laser differential interferometer". In: *AIAA Scitech 2020 Forum*. Orlando, FL. DOI: 10.2514/6.2020-1064.

Lawson, J. M., Neet, M. C., Grossman, I. J., and Austin, J. M. (2019). "Characterization of a focused laser differential interferometer". In: *AIAA Scitech 2019 Forum*. San Diego, CA. DOI: 10.2514/6.2019-2296.

Lawson, J. M. and Austin, J. M. (2018). "Design of conventional and detonation-driven hypervelocity expansion tubes". In: *AIAA Aviation 2018 Forum*. Atlanta, GA. DOI: 10.2514/6.2018-3566.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

INTRODUCTION

1.1 Motivation

1.1.1 Fluctuating Flows

In the supersonic and hypersonic flow communities, fluctuating quantities are at the core of multiple active areas of research with practical implications for vehicle design. Laminar boundary layers (LBLs) are receptive to fluctuations in the freestream flow, leading to the growth of instabilities and ultimately to transition to turbulence. Turbulent boundary layers (TBLs) cause greatly-increased thermal loading to the vehicle surface, so prediction of the timing and extent of transition during a flight trajectory is very important. Shock-wave boundary-layer interaction (SWBLI) is a phenomenon found particularly near corners and protrusions of supersonic vehicle geometries, where a complex and often unsteady interplay leads to boundary layer separation and reattachment, additional shock waves, shear layers and so on. SWBLI is often associated with regions of peak heat load, again making prediction important for vehicle survivability—this prediction in turn relying on knowledge of the external freestream disturbance levels. The study of fluid-structure interactions (FSI) is concerned with the coupling between fluctuating flows and compliant surfaces (such as thin metal panels comprising an aircraft skin). Knowledge of the frequency modes and amplitudes in the flow are required in order to predict which structural modes are likely to be excited in these panels.

1.1.2 The Need for Ground Testing

Hypersonic flows have arguably the most complex physics of any regime in aerospace: non-equilibrium real-gas chemistry, radiative heating, turbulent boundary and shear layers, and complex shock interactions that all must be measured or modeled. Including even a subset of these phenomena can render numerical simulation infeasible for practical engineering design, even if they are understood well enough to model in the first place. However for lower-enthalpy supersonic, and indeed even subsonic flows, the aforementioned range of scales for fluctuating flows can still make computation very expensive, even if the underlying physics are comparatively simpler to model. These challenges make experiments necessary.

High-speed free-flight experiments are very expensive, and the possible types of on-board instrumentation limited. Such experiments often involve years of planning from large-scale university and military collaborations, e.g. HIFiRE and BOLT (Juliano et al., 2015; Wheaton et al., 2021). The precursor to flight testing is ground testing, where key flow parameters are replicated in the laboratory, allowing a wider range of diagnostics to be employed. Ground testing can be done either by moving a solid model through quiescent gas at hypersonic speeds (as in a ballistic range), or, more commonly, by inducing the gas to flow over a stationary model. There are many methods by which such flows can be generated, although they tend to have in common finite (and sometimes very short) test times and low freestream densities. This is due to these flows typically having specific enthalpies h_0 in excess of $1\,\mathrm{MJ\,kg^{-1}}$: the power needed to continuously sustain such a flow scales like $\rho u A h_0$ (where ρ is density, u is velocity, and A is cross-sectional area), and can easily exceed 1 GW. Hence, continuous wind tunnels are usually limited to lower velocities.

1.1.3 Facility Types

A distinction can be made between hypersonic (Mach number $M \gtrsim 5$) and hypervelocity (velocity $u \gtrsim 3\,\mathrm{km\,s^{-1}}$). In the former case, expanding a gas causes a reduction in temperature and therefore sound speed: high M can be achieved without especially large u. The high-enthalpy effects that distinguish true hypervelocity flows from low-enthalpy supersonics come from the dominance of the kinetic energy contribution, $\frac{1}{2}u^2$.

Non-continuous high-speed ground-testing facilities can be categorized into two general types: blowdown and impulse. Blowdown facilities have a high-pressure source reservoir connected via a Laval nozzle to a low-pressure sink, either atmosphere or a vacuum tank. Once a valve is opened, the reservoir gas expands and accelerates to supersonic speeds through the nozzle, "blowing down" to the sink. With large enough source and sink volumes, steady flow can be maintained for seconds or even many minutes. These facilities are usually low-enthalpy, even with heat addition in the reservoir. Impulse facilities use gasdynamic processes to accelerate the test gas to high-enthalpy "true" hypervelocity states, although usually at the expense of test times on the order of µs to ms. In this thesis, two particular ground-testing facilities will be considered: a conventional blowdown tunnel, and a type of impulse facility known as an expansion tube. More detail on these will be provided in Chapter 2.

1.1.4 Diagnostics

Intrusive sensors necessarily alter the flow merely by their presence, and have finite inertia (e.g. the mass of a diaphragm or the heat capacity of a thermocouple). Furthermore, some sensors exhibit resonant behavior that limits their high-frequency response; however, at hypervelocity, the frequencies of interest (such as the second Mack mode for BL instabilities) can exceed 1 MHz (Bitter, 2015). Generally, all three of these undesirable traits (intrusiveness, inertia, and resonance) can be minimized by reducing the size of the sensor, however, this makes them fragile—and impulsively-started hypersonic flows are a very harsh environment.

In contrast, optical techniques do not suffer from inertia or resonance, only being limited by the response time of the detector or camera. Most techniques deposit negligible amounts of energy into the flow, hence being considered non-intrusive. The main drawbacks of optical techniques are line-of-sight integration and difficulties in quantitative interpretation. When a ray of light traverses the test section, it is altered (e.g. by refractive index or species concentration changes) all along its path, making it difficult to determine the spatial distribution of the quantity of interest without exploiting symmetry or other assumptions about flow geometry. Many optical techniques are also difficult to calibrate, leaving their results qualitative, such as schlieren.

Focused laser differential interferometry (FLDI) is an instrument that attempts to address the key drawback of optical techniques, i.e. the line-of-sight integration that leads to delocalization of the response. In its most fundamental form, FLDI comprises two coherent, orthogonally-polarized beams of light whose principal rays run parallel, separated by a small distance Δx. The two beams are focused down from initially large diameters to give two small foci. Beyond this focal plane, the beams expand again, then are overlapped and allowed to interfere with each other. There are many ways to achieve this beam geometry; the basic optical configuration used in this work is shown in Fig. 1.1. FLDI is a non-imaging interferometer: its output is a single voltage–time signal, proportional to the interference intensity integrated over the recombined beams.

FLDI is the main topic of this thesis; both the behavior of the instrument itself and its application to hypersonic ground facilities will be addressed in detail. A more detailed introduction to FLDI and its developmental history will be given next in Section 1.2. Following this, literature reviews are presented concerning freestream noise studies in conventional and impulse facilities (Sections 1.3 and 1.4).

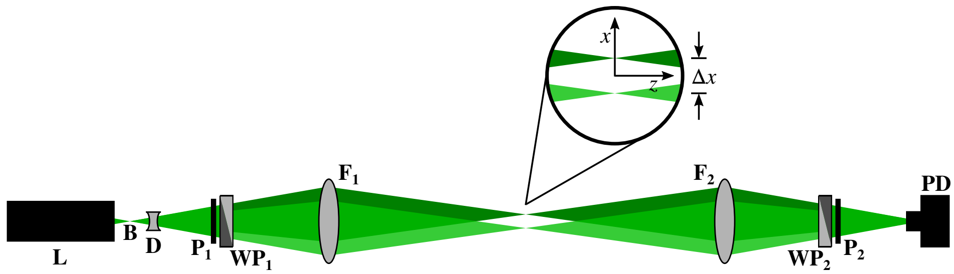

Figure 1.1: FLDI schematic. Component annotations in boldface: **L** = laser, **B** = beam waist, **D** = diverging lens, $\mathbf{P_1}$ = quarter-wave plate, $\mathbf{WP_i}$ = Wollaston prisms, $\mathbf{F_i}$ = focusing lenses, $\mathbf{P_2}$ = linear polarizer, **PD** = photodetector.

1.2 Focused Laser Differential Interferometry

1.2.1 Introduction

Like all optical interferometers, FLDI responds to optical path differences (OPD) between the two beams; for FLDI, the OPD are due to variations in the refractive index field n. In gases (and liquids), n is linearly related to the density ρ by the Gladstone-Dale relation (Gladstone and Dale, 1863; Merzkirch, 1987):

$$n = K\rho + 1 \tag{1.1}$$

where K is a coefficient that depends on the wavelength of the light. Eq. (1.1) was derived empirically; a more fundamental approach from first principles led to the Lorenz-Lorentz relation, after its two independent but confusingly-named discoverers (Gardiner et al., 1981; Kragh, 1991):

$$R_L = \frac{n^2 - 1}{n^2 + 2} \cdot \frac{1}{\rho} \tag{1.2}$$

Again, the constant R_L depends on wavelength; Gardiner et al. contains extensive tabulations of R_L at common laser wavelengths, and is used as the reference for this thesis. For gases, where ρ is small and n is close to unity, Eq. (1.2) approaches Eq. (1.1), with $K \approx 3R_L/2$. For either law, mixtures of gases can be treated by taking linear combinations of R_L or K, weighted by the mole fraction of each component. These relationships mean that FLDI indirectly measures density fluctuations, and does so with high spatial resolution due to the small magnitude of Δx, typically $O(100\,\mu\text{m})$. The instrument also has high temporal resolution because it is only limited by the bandwidth of modern photodetectors, usually $O(10\,\text{MHz})$. The output voltage $V(t)$ corresponds to an overall optical phase shift $\Delta\Phi(t)$, measured in radians. Fundamentally, the instrument response is related to a finite-difference approximation to the local density gradient, i.e. $\Delta\Phi(t) \sim \Delta\rho/\Delta x$. As $\Delta x \to 0$, the response approaches the true derivative, but the magnitude of the signal is also diminished.

FLDI can be considered a variant of laser differential interferometry (LDI), a simpler configuration where the beams remain a constant diameter instead of focusing. It is this focusing that leads to the key distinguishing feature of FLDI: it is less sensitive to refractive index disturbances further away from the focal plane. This alleviates the issues caused by line-of-sight integration: contributions to the overall signal largely come from a “sensitive region” in the vicinity of the foci. In the ideal limit, this sensitive region would be made so small that the FLDI would give point-like measurements. This feature is particularly useful in hypersonic ground testing, where the uniform core flow is often surrounded by highly turbulent free shear layers, or boundary layers on the test section windows. The intensity of these turbulent outer flows may overwhelm the signals of interest in the core; if FLDI can attenuate these sufficiently, then the core can be probed optically. FLDI can be extended to a “dual” or “double” FLDI (DFLDI) where two foci pairs are generated, usually aligned in the streamwise direction. This allows cross-correlation of the two interference signals to compute a time-of-flight and thus a velocity, in addition to the density information from the individual channels.

1.2.2 History

FLDI was first conceived by Smeets and George (1973) at the French-German Research Institute of Saint-Louis, as one of several optical designs of laser differential interferometers. However, it appears that the relative infancy of laser technology, along with limitations in photodetection and acquisition instruments, prevented its widespread uptake at that time. In the mid-1990s there was interest in novel optical methods for sensitive disturbance measurements in high-speed flows (Collicott et al., 1996). Multiple reports from Smeets and George were translated into English by A. G. Goetz in 1995-6 at the behest of Schneider*, with assistance from both Smeets and George themselves who visited the USA at different times during this period. The non-focusing LDI was pursued rather than the FLDI by these groups (Salyer et al., 2000).

*Personal communication with S. P. Schneider.

Parziale et al. (2012) demonstrated a simple, robust implementation of FLDI for hypersonic flows (also with a basis in the work of Smeets and George), after considering a variety of other optical techniques with the potential for quantitative measurements[†]. Since then, there have been a number of efforts to further develop and analyze the technique. Parziale et al. (2013b, 2014, 2015) and Parziale (2013) further used FLDI to make both free-stream turbulence and cone boundary layer instability measurements in T5.

Fulghum (2014) and Settles and Fulghum (2016) obtained free-stream turbulence power spectra in the Penn State Supersonic Wind Tunnel and the AEDC Hypervelocity Wind Tunnel 9, and made comparisons with hot-wire anemometry. Ceruzzi and Cadou (2017, 2019) obtained spectra in turbulent jets, then utilized FLDI to study turbulent wall boundary layers in a Mach 2.6 tunnel (Ceruzzi et al., 2020). Benitez et al. (2021) and Bathel et al. (2020) also adapted FLDI for instability measurements along cones in the Mach 6 Quiet Tunnel at Purdue University and the 20-Inch Mach 6 Air Tunnel at NASA Langley, respectively. Houpt and Leonov (2018, 2019) used cylindrical lenses to produce a "planar" FLDI variant in order to make measurements closer to solid surfaces without beam clipping.

Very recently (at the time of writing) Ceruzzi et al. (2021a) presented freestream density and velocity measurements made in Tunnel 9 at $M = 18$. In a separate work (Ceruzzi et al., 2021b), a "sensitivity function" for FLDI was developed and validated, based on the transfer function approach of Schmidt and Shepherd.

1.2.3 Interpretation of FLDI Response

Many of the above-cited works deal with the raw FLDI data, i.e. the voltage signal from the photodetector, sometimes converted to optical phase shift. For many applications this is sufficient: for example, when cross-correlating a pair of signals from a DFLDI. Parziale et al. introduced a simple method for estimating the average density difference $\Delta\rho$ between the two beams:

$$\frac{\Delta\rho}{\overline{\rho}} = \frac{\lambda_L}{2\pi K\zeta\overline{\rho}} \sin^{-1}\left(\frac{V}{V_0} - 1\right) \tag{1.3}$$

where $\overline{\rho}$ is the local average density, λ_L is the laser wavelength, K is the Gladstone-Dale constant, V is the FLDI output voltage, V_0 is a reference voltage taken when the interferometer is set to the middle of a fringe (see Section 3.4), and ζ is some

[†]Personal communication with N. J. Parziale.

empirically-determined “sensitive length”. Note some notation has been changed from the original publication in order to avoid confusion within this thesis. They also introduced a coefficient that compensates for changes in FLDI response with disturbance wavelength. The limitations of this modeling approach will be addressed in Section 5.7. Recently, Hameed and Parziale (2021) used a vibrating cylindrical lens as a phase object to experimentally test this model.

At around the same time as each other, Schmidt and Shepherd (2015) and Settles and Fulghum (2016) presented more complex analytical models of FLDI. The results of the latter paper are also given with more detail in Fulghum (2014). These two groups took quite different approaches to the optical modeling: Schmidt and Shepherd used geometrical optics with the paraxial approximation to derive a ray-tracing equation; Settles and Fulghum instead modeled the propagation of polarized electric fields using Jones vectors, alongside a simplified finite difference model. Nevertheless, they both arrived at the same important result: a transfer function for the response of FLDI to a one-dimensional sinusoidal disturbance. It is given by:

$$H(k) = \frac{2}{k\Delta x} \sin\left(\frac{k\Delta x}{2}\right) \exp\left(-\frac{w^2 k^2}{8}\right) \tag{1.4}$$

The meaning of this transfer function will be revisited in greater detail at several points later in this thesis. For now, the key observation comes from the exponential term, the argument of which contains the product of the local beam width w and the disturbance wavenumber k. This can equally be expressed as the ratio of two lengths: the beam width over the disturbance wavelength, and this is the reason for the spatial filtering effect that makes FLDI distinct. Far from the beam foci, where the beams are wide, their sensitivity to wavelengths much smaller than the local w is greatly diminished. This leads to a small sensitive region centered about the focal plane—although Eq. (1.4) also demonstrates that the sensitive length is wavenumber-dependent. Both groups also developed related transfer functions for other classes of disturbance geometry. Again, these will be reviewed later, and compared with new results.

Another important result of these analyses is that FLDI response is only dependent on three optical parameters: the foci separation Δx, the Gaussian beam radius of the foci w_0, and the laser wavelength λ_L, where $\lambda_L << w_0 << \Delta x$. Typically λ_L falls in the narrow band of visible wavelengths (400–700 nm), and has only a scaling effect on the response. Δx gives the minimum spatial resolution, for fluctuation wavelengths shorter than this there is a strong roll-off in FLDI response. The most important parameter is w_0: a smaller w_0 means the beam width grows more rapidly away from the focus. As just discussed, this focusing effect is what gives FLDI its spatial selectivity: hence reducing w_0 shortens the sensitive region. Note that throughout this work, the radius w_0 will be used interchangeably with the corresponding diameter d_0.

1.2.4 Comparison to Pitot Data

A recent study that has parallels with the work performed herein is that of Birch et al. (2020). FLDI was used along with a pitot probe to make measurements in the freestream of a $M = 6$ piston-driven Ludwieg tube. Geometrical limitations due to their wide test section meant the focusing angle of their FLDI was relatively small; initial testing showed significant contributions to the signal from the outer turbulent shear layers. Consequently, aerodynamic beam shrouds were constructed to shield the outer portions of the FLDI beams.

Conversion of the FLDI signal to density was performed by using elements of theory from both Schmidt and Shepherd, and Settles and Fulghum. The pitot data were interpreted using the classical approach of Stainback and Wagner (1972) with the assumption made of purely acoustic disturbances with the same orientation as per Laufer (1961) (these theories are addressed in Section 1.3). These various pieces of theory have some inconsistent assumptions undergirding them: one transfer function used for the FLDI assumes plane waves propagating parallel to the flow, while another transfer function as well as their pitot conversion both assume acoustic waves with some angle of incidence to the flow direction. As will be shown in this thesis, these inconsistencies will lead to quantitative errors.

The main conclusion of interest to this work comes from their comparison of fluctuating static pressure p' as recovered from FLDI vs. pitot. They present RMS values of p' evolving over the 200 ms run-time of the facility. Until 75 ms there is close agreement between the two methods, after which p' as calculated from the pitot data sharply increases, while the FLDI data does not. This was interpreted as being due to entropy-mode disturbances being present in the flow; these influence the pitot response differently to FLDI. In turn, the arrival of these entropic disturbances midway through test time was thought to be due to transition processes occurring in the piston barrel, upstream of the nozzle.

1.3 Conventional Facility Noise

1.3.1 Overview

Much work has been done on characterizing the freestream noise of conventional supersonic wind tunnels, both through theoretical developments and experiments with intrusive diagnostics. Broadly speaking, the theoretical foundation was provided in the 1950s by Kovásznay (1953) and Morkovin (1957, 1959), who also demonstrated the use and interpretation of hot-wire anemometry (HWA) in supersonic flows. In the 1960s through to the early 1970s, these techniques were applied in detail to a variety of facilities; one of the prominent works of this period is by Laufer (1961). Stainback and Wagner (1972) used a pitot pressure probe in addition to HWA; their interpretation method for the pitot data remains in widespread use. The main interest in understanding the tunnel noise environment was in order to study boundary layer transition. Pate (1980) gives an extensive review of the work done to that point both on characterizing the freestream disturbance environment, and the effect thereof on boundary layers.

It appears there was something of a lull in research into tunnel noise characterization after this time, although work on the receptivity, stability, and transition of boundary layers certainly continued in earnest (these topics are however outside the scope of this thesis). The understanding gained in the 1960s about the nature and origin of tunnel noise inspired the pursuit of quiet supersonic and hypersonic tunnels, the developmental history of which is given by Schneider (2001, 2008).

More recently, the 2010s saw many new freestream characterization studies. These included "traditional" HWA-pitot surveys, often supplemented by more advanced probe diagnostic designs. Renewed attention was given to interpreting the response of intrusive diagnostics, with consideration given to high-frequency and probe geometry effects. These modern studies have the benefit of a variety of computational fluid dynamics (CFD) tools that are used to simulate the probe response to realistic disturbance field and thereby deduce detailed transfer functions.

1.3.2 Classical Results

The seminal work is Kovásznay (1953), where first-order perturbation theory was applied to the Navier-Stokes equations, yielding three disturbance fields often referred to in subsequent literature as "Kovásznay modes": the entropy, vorticity, and acoustic modes. To first order, these modes are independent, i.e. nonlinear interactions between modes are neglected. He also offered a second-order extension of the theory to allow for weak interaction; however, it was argued that the low intensities typical of supersonic wind tunnel fluctuations make the first-order theory sufficient. Further simplification via consideration of characteristic length and time scales demonstrated that the vorticity and entropy modes convect as "frozen patterns" (i.e. Taylor's hypothesis), while the acoustic mode is unattenuated and obeys the pure wave equation.

HWA does not respond directly to these three modes; instead its response is a function of fluctuations in mass flow $\dot{m}'$ and stagnation temperature T_0'. These in turn depend on fluctuations of density, velocity, and static temperature (ρ', u', and T', respectively), while pressure p' does not have a direct effect. By operating the HWA at multiple conditions and repeating experiments, one can recover (in a mean square sense) $\dot{m}'$ and T_0', as well as the correlation coefficient between these quantities, R_{mT}. This technique is often implemented graphically as a "fluctuation diagram" a.k.a. "mode diagram"; via these, Kovásznay showed that the three modes can be separated from the data for some special cases, although for the general case of all three modes coexisting with arbitrary correlation, they cannot be separated without further assumptions or information.

Morkovin (1957, 1959) expanded upon the nature and origin of the Kovásznay modes, and made arguments based on both theory and experiment as to the expected relative importance of each mode in conventional supersonic wind tunnels. The modes each behave as follows:

1. Entropy mode: variation of entropy s, ρ, and T at constant p.
2. Vorticity mode: variation of the solenoidal component of the velocity field ω.
3. Acoustic mode: variation of p, ρ, and T at constant s and constant irrotational component of the velocity field.

In incompressible flows, the vorticity mode is often simply referred to as "turbulence". Per Taylor's hypothesis, the entropy and vorticity modes convect along streamlines and so must originate from upstream of the nozzle (in the settling chamber or reservoir of the tunnel). Only the acoustic mode can cross streamlines and so can also originate from the walls of the nozzle and test section. Morkovin further sub-categorized these wall-sourced acoustic fluctuations:

a. Radiation from initial turbulent bursts within transitioning boundary layers.

b. Radiation from fully-developed turbulent boundary layers.

c. Diffraction and scattering from solid geometry (e.g. roughness, nozzle contours).

d. Radiation from wall vibrations.

Theoretical considerations suggest that in conventional tunnels, the acoustic mode should dominate over the other two. For the vorticity mode, all components of the velocity flucutation field decrease strongly with the speed ratio across the nozzle u_2/u_1. Because this ratio is very high especially in supersonic facilities, so long as the pre-nozzle settling chamber turbulence is kept reasonably low (e.g. via screens) then the vorticity fluctuations in the test section will be very low. The entropy mode originates from temperature spottiness upstream of the nozzle, which can be reduced via the same design considerations that help with settling chamber turbulence reduction. Measurements made in four different tunnels found reservoir temperature spottiness to be "inconsequential".

Regarding the various categories of acoustic noise, Morkovin postulated that Category (d) sound is negligible due to high acoustic impedance between the wall and flow. Categories (a) & (b) are the "aerodynamic generation of sound" as described in detail by Lighthill (1952, 1954); Morkovin however believed that Category (c) sound would be dominant in many facilities. This category is also referred to as "shivering" or unsteady Mach waves, and can be reduced by polishing the walls, smoothing joints, and carefully designing the nozzle contours.

Laufer (1961) performed HWA studies in the JPL 18×20 in. tunnel with flexible nozzles walls allowing for $M = 1.6$–5.0 . The mode diagrams showed pure acoustic noise fields, a conclusion supported both by additional measurements in the supply section, and by physical arguments similarly to Morkovin. A pure acoustic field obeys the isentropic relation:

$$\frac{p'}{\overline{p}} = \gamma \frac{\rho'}{\overline{\rho}} = \frac{\gamma}{\gamma - 1} \frac{T'}{\overline{T}} \tag{1.5}$$

Here, $\square'$ and $\overline{\square}$ represent fluctuating and time-averaged quantities, respectively. The data further showed that the preferred orientation of the sound was not the Mach angle, i.e. the majority of the measured p' was not due to fluctuating Mach waves, in opposition to Morkovin's belief. This meant the acoustic source could not be stationary in the laboratory reference frame; it was instead a moving eddy within the boundary layer, i.e. Category (b) noise above. Using an estimate of the wavelength range of the acoustic field, Laufer stated that the measurements were sufficiently distant from their source at the wall so that the assumption of plane waves could be made. The velocity of the sound source u_s and the inclination of the wave is given by:

$$\frac{u_s}{u} = 1 + \frac{1}{M \cos\theta} \tag{1.6}$$

where u and M are the freestream velocity and Mach number, and θ is the angle between the wave normal and the freestream flow. It was found that $0.4 < u_s/u < 0.5$, implying $90^\circ < \theta < 180^\circ$, i.e. a backwards-facing wave. Note that for $u_s = 0$ the Mach angle $\mu = \theta - 90^\circ$ is recovered. Laufer stated that Eq. (1.6) gives a lower bound on u_s, since it assumes a single source but in reality Mach wave [Category (c)] noise may also be present, albeit to a lesser extent. Besides the aforementioned acoustic impedance argument, wall vibration [Category (d)] noise was further ruled out because the observed frequencies were much too high. Several supplementary experiments were performed that gave additional support to the acoustic field being dominated by TBL [Category (b)] noise, including: shielding the HWA from acoustic radiation from one of the four tunnel walls, which gave a corresponding decrease in signal, and operating the tunnel at a LBL condition, where the overall fluctuations greatly decreased, with the small residual signal having a mode diagram indicative of entropy fluctuations. Similar results were obtained in a smaller 12×12 in. tunnel.

In a follow-up work, Laufer (1964) continued his investigation into the nature of this radiated acoustic field. Ballistic free-flight schlieren studies suggested the field comprised backward-facing wavelets of limited spatial extent rather than infinitely-long plane waves. The variation in the orientation of these wavefronts decreased as M increases, i.e. the radiated field became more coherent. One important conclusion for square cross-section tunnels was that each wall radiated equally and in an uncorrelated fashion; furthermore, reflection of the acoustic waves from the other walls was found to be negligible. Data were given for the far-field pressure spectrum and a scaling for the radiated intensity with respect to the wall shear stress was provided.

Of the studies that followed Laufer, the most important for this thesis is the report by Donaldson and Wallace (1971), because it was conducted in VKF Tunnel D, one of the two facilities studied in this work (see Section 2.6 for more information on this facility, and Chapter 6 for results). For now, it suffices to say that Tunnel D is of a fundamentally similar design to the tunnel used by Laufer, with a square cross-section and flexible nozzle that allows for variation in M. Donaldson and Wallace performed HWA studies at a range of unit Reynolds numbers Re_m and $M = 4$, with assumptions and data analysis techniques following those of the earlier works. As with Laufer, they found that the mode diagrams indicated pure acoustic fields at all conditions, and that the dominant orientation of the waves was not equal to that of stationary Mach waves. They found a higher value of u_s than Laufer, reporting it in terms of the wave-normal angle: $\theta = 122$–$128°$, i.e. backwards-facing waves. Facility-specific calculations were done to show the validity of the assumption of negligible vorticity and entropy mode fluctuations in the Tunnel D test section.

Wagner (1971) performed a HWA study in the Langley 60-in. High Reynolds Number Hypersonic Helium Tunnel, which unlike other works cited here used helium instead of air, and operated at a much higher Mach number of 17.5. Measurements were taken in both the freestream and in the shock layer of a slender cone (half-angle of 2.87°). The mode diagrams showed that the fluctuations were due to moving sound sources in the turbulent boundary layers. Furthermore, the shock layer did not show evidence of redistribution of energy from the acoustic mode into the vorticity or entropy modes, although this was attributed to the shock being weak—redistribution would be expected from a stronger shock.

The last work for us to review from the early period is that of Stainback and Wagner (1972), who compared pitot-probe measurements against HWA data for freestream fluctuations in a $M = 5$ air flow from an axisymmetric nozzle. The theory they developed for interpreting unsteady pitot data remains in widespread use. Again, the HWA mode diagrams showed a noise field of a purely acoustic nature, generated by non-stationary sources with reasonably similar u_s to Laufer. When the facility was operated at a LBL condition there was almost no signal, as expected.

A pitot probe in supersonic flow measures the total pressure behind the bow shock that stands off from the probe, denoted $p_{t,2}$. This data needed to be converted to freestream static pressure p_1 (as obtained from the HWA analysis); an outline of Stainback and Wagner's derivation is as follows. For $M > 2.5$ the following is valid to within 0.2%:

$$p_{t,2} = G\rho u^2 \tag{1.7}$$

for some $G = G(\gamma)$ where γ is the specific heat ratio. Assuming Eq. (1.7) holds in an instantaneous sense, it can be used to derive a second-order fluctuation equation (Harvey et al., 1969), which holds generally:

$$\left(\frac{\tilde{p}_{t,2}}{\overline{p}_{t,2}}\right)^2 = 4\left(\frac{\tilde{u}}{\overline{u}}\right)^2 + 4\left(\frac{\tilde{u}\tilde{\rho}}{\overline{u}\overline{\rho}}\right) R_{\rho u} + \left(\frac{\tilde{\rho}}{\overline{\rho}}\right)^2 \tag{1.8}$$

where $R_{\rho u} \equiv \overline{\rho' u'}/\tilde{\rho}\tilde{u}$, and $\overline{\square}$, $\tilde{\square}$, $\square'$ denote mean, root mean square (RMS), and instantaneous values, respectively. The isentropic relations of Eq. (1.5) can be modified slightly to account for inclined plane waves; substituting these into Eq. (1.8) gives a relationship specific to purely acoustic fields:

$$\left(\frac{\tilde{p}_{t,2}}{\overline{p}_{t,2}}\right)^2 = \left(\frac{\tilde{p}}{\gamma\overline{p}}\right)^2 \left[1 - 4\frac{n_x}{M} + 4\left(\frac{n_x}{M}\right)^2\right] \tag{1.9}$$

where $n_x \equiv \cos\theta$, with the wave orientation θ defined by Laufer as per Eq. (1.6). Eq. (1.9) did not produce good agreement with the HWA data, with the pitot data giving about double the values of $\tilde{p}_{t,2}/\overline{p}_{t,2}$. Amending Eq. (1.9) by a multiplicative factor of 2 improved the agreement; the rationale being that Eq. (1.9) was derived under quasi-steady assumptions, but under unsteady conditions the reflection of the compression wave from the probe face must be accounted for. The unsteady pitot equation is then simply:

$$\left(\frac{\tilde{p}_{t,2}}{\overline{p}_{t,2}}\right)^2 = 2\left(\frac{\tilde{p}}{\gamma\overline{p}}\right)^2 \left[1 - 4\frac{n_x}{M} + 4\left(\frac{n_x}{M}\right)^2\right] \tag{1.10}$$

The factor of exactly 2 in Eq. (1.10) relies on several simplifying assumptions: no wave diffraction around the probe geometry, thin boundary layers on the probe (relative to the acoustic wavelengths), and simple reflection of waves that are parallel to the probe face. In particular, the last assumption ignores that Eq. (1.10) allows for arbitrary wave orientation n_x. Stainback and Wagner performed additional comparisons in another facility (helium, $M \approx 20$) which produced better agreement than in the $M = 5$ air cases, although it was not understood why this was. The authors warned that Eq. (1.10) cannot be relied upon for useful quantitative data, and that a more general theory and further calibration experiments are required. In particular they noted that the wavefront angles would be important. Despite these reservations, later works by other authors tend to use Eq. (1.10) as a quantitative result without critique. Some of these studies will be addressed below.

1.3.3 Modern Developments

Masutti et al. (2012) performed a characterization of the freestream disturbances in the VKI-H3 tunnel[‡] ($M = 6$, air). They used a dual HWA and a fast-response pitot probe, although not simultaneously—it was assumed that the statistics did not change between runs. A small-perturbation approximation was used to develop a system of equations:

$$\frac{m'}{\overline{m}} = \frac{\rho'}{\overline{\rho}} + \frac{u'}{\overline{u}} \tag{1.11a}$$

$$\frac{p'}{\overline{p}} = \frac{\rho'}{\overline{\rho}} + \frac{T'}{\overline{T}} \tag{1.11b}$$

$$\frac{T_0'}{\overline{T}_0} = \alpha\frac{T'}{\overline{T}} + \beta\frac{u'}{\overline{u}} \tag{1.11c}$$

where $\alpha, \beta = f(M, \gamma)$, m is the mass flux, and T and T_0 are the static and stagnation temperatures, respectively. $m'/\overline{m}$ and $T_0'/\overline{T}_0$ are obtained from the dual HWA, while $p'/\overline{p}$ comes from the pitot. In this way, Eqs. (1.11a) to (1.11c) can be solved simultaneously to yield all fluctuating flow quantities. In order to compute $p'/\overline{p}$ from the raw pitot pressure data, the authors used Stainback and Wagner's approach, i.e. Eq. (1.10). To determine the disturbance orientation n_x, Laufer's Eq. (1.6) was used, and it was assumed that $u_s/u = 0.6$ as taken from Stainback and Wagner, i.e. Masutti et al. assumed without evidence that the acoustic waves are at the same angle as found in a rather different facility, operated at a different value of M. Their results showed higher fluctuations in total temperature than could be explained purely by the acoustic mode. This was attributed to VKI-H3 lacking an upstream thermal equalizer system, therefore leading to entropy spottiness in the supply. Their RMS pitot data showed an inverse linear relationship with the unit Reynolds number Re_m, attributed to stabilization of the TBL.

[‡]VKI = von Kármán Institute, in Belgium. Not to be confused with VKF = von Kármán Gas Dynamics Facility, in Tennessee!

In addition to trends in the raw data, once all the fluctuations were computed per Eqs. (1.11a) to (1.11c), the three Kovásznay modes could be recovered:

$$\Theta' = \frac{T'}{\overline{T}_0} - \frac{\gamma - 1}{\gamma}\frac{\overline{T}}{\overline{T}_0}\frac{p'}{\overline{p}} \tag{1.12a}$$

$$\omega' = \frac{u'}{\overline{u}} + \frac{1}{\gamma M^2}\frac{p'}{\overline{p}} \tag{1.12b}$$

$$\sigma' = \left(1 - \frac{1}{M^2}\right)\frac{p'}{\gamma\overline{p}} \tag{1.12c}$$

where Θ, ω, and σ are the entropy, vorticity, and acoustic modes, respectively. All three modes were found to be present and uncorrelated with each other as Re_m varies. They concluded that the acoustic mode is dominant in VKI-H3 over the whole Re_m range, however their own provided data do not appear to back up this conclusion. $\sigma'_{RMS} \approx 0.71\%$ while Θ'_{RMS} and ω'_{RMS} are close to this at 0.6%; furthermore for several values of Re_m, Θ'_{RMS} is either equal to or exceeds σ'_{RMS} (within their wide error bounds of up to ±40%). While σ'_{RMS} is on average the largest contribution, it certainly cannot be stated to be dominant, as even the largest ratio between any two modes is no more than a factor of two and usually much less than this. Another issue with their analysis is the reliance on Stainback and Wagner's pitot conversion equation, despite the warnings of the original authors that such a simplified approach cannot be relied on quantitatively.

Gromyko et al. (2013) used pitot probes, heat-flux gauges, and HWA to study freestream disturbances in the Transit-M facility. This is an electrically-heated Ludwieg tube with a fast-acting valve, with test times of 110–200 ms at $M = 6$, differentiating this study from those previously discussed, which were more conventional blowdown facilities with much longer test times. Various possible sources of vortical and entropic disturbances in the pre-chambers were discussed; it was expected that the high-frequency components of these would be damped out due to the distance scales involved, and that the high-frequency portion of the overall spectrum should largely be due to the acoustic mode. Using wall boundary layer thickness as the characteristic length scale, it was estimated that the lowest-frequency components of the acoustic mode should be 20–40 kHz. Their measurements had a bandwidth of 350 kHz based on noise floor considerations.

To aid with interpretation of the pitot data, the authors performed flow simulations using commercial software (ANSYS) with single-frequency acoustic waves superimposed on the mean flow. These results showed that the pitot response depends on the frequency and orientation of the wave, due to wave interference effects that lead to non-uniform pressure distributions across the sensor face.

Their mean p' data were comparable to a variety of other facilities (both conventional blowdown and short-duration tunnels) but lower than in T4 (a reflected shock tunnel). However, T_0' was considerably higher than in conventional facilities. This was attributed to short-duration facilities having less mixing time after heating. Substantial azimuthal asymmetry was also observed in p'. Subsequent tests implied this was likely due to asymmetry in the laminar-turbulent transition line on the nozzle wall, which in turn were attributed to fabrication inaccuracies in the transonic portion of the nozzle, or possibly complex effects within the pre-chamber. An increase in fluctuations was observed far from the centerline, due to getting close to the developed TBL, as also observed by Rufer and Berridge (2012). Decreasing Re_m caused an increase in low-frequency fluctuations, stated to be due to a thicker nozzle boundary layer (a larger characteristic dimension means a lower characteristic frequency).

As part of a larger investigation, Mai and Bowersox (2014) conducted freestream pitot fluctuation studies in the Texas A&M ACE tunnel, a conventional blowdown facility, here operated at $M \approx 5.9$. The authors assumed that the noise was acoustic in nature. There was a low-noise regime observed at low Re_m, with a sharp rise to a high-noise regime as Re_m increased, attributed to transition of the nozzle boundary layer. Within the high-noise regime, the noise decreased gradually with Re_m as seen in other works; likely due to stabilization effects.

The pitot surveys were performed at two streamwise locations: the nozzle exit plane, and then 95 mm further downstream. In the laminar regime, the downstream location was noisier than the nozzle plane for a given Re_m. This was thought to be due to a cumulative effect of Mach waves from imperfections in the nozzle and test section (i.e. Morkovin Category (c) acoustic noise). Conversely, in the turbulent regime both locations had the same noise level, which the authors took to imply that the TBL radiation (Morkovin Category (b) acoustic noise) had "saturated the freestream noise levels". The integral time scales of the noise were also plotted against Re_m. Around the inter-regime range (corresponding to turbulent transition) there was a notable spike in scale. This was due to intermittent large-scale structures in the boundary layer as it transitioned. At full turbulence, the energy cascade is established and the size of structures reduces and stabilizes. This was also investigated via spectrograms.

Schilden et al. (2016) studied the freestream of the Hypersonic Ludwieg Tube Braunschweig (HLB) using two methods: the "classical" HWA-pitot, and a cone probe they developed. Similarly to Transit-M in Gromyko et al. above, HLB is a heated Ludwieg tube employing a fast-acting valve, giving $M = 5.9$ for about 80 ms. Their cone probe had a half-angle of 30°, with flush-mounted pressure transducers along the surface. Direct numerical simulation (DNS) was used to deduce the transfer functions for this probe. The freestream was modeled using acoustic and entropy Kovásznay modes, with the vorticity mode neglected. For the analysis of the HWA-pitot data, they followed the approach of Masutti et al. discussed above. They elaborated on several of the assumptions inherent in this approach, such as the use of Stainback and Wagner's unsteady pitot equation, the choice of the slow acoustic mode, and the need to assume a value of the sound source velocity. They also noted that Masutti et al. performed their decomposition in an RMS sense only, not spectrally as done here.

The HWA-pitot data showed that all three modes all decreased with increasing Re_m. In an RMS sense, the acoustic mode was an order of magnitude larger than the other two modes, with the entropy mode about double the vorticity mode. Issues were encountered with the decomposition of the cone probe data being overly sensitive to small errors either in the data or the DNS-derived transfer functions; this necessitated the introduction of additional assumptions, wherein good agreement was found with the HWA-pitot results. The main set of results were taken at a location 100 mm from the centerline; on the centerline itself there was a significant increase in the vorticity mode, thought to be from the valve, while close to the walls the acoustic mode became stronger.

Wagner et al. (2018) presented a study on three facilities: two Ludwieg tubes (one being HLB from Schilden et al.) and a shock tunnel (HEG). A slender wedge-shaped probe was used, because HWA is not suitable for high-enthalpy facilities such as HEG. The probe was instrumented with pressure, temperature, and heat-flux gauges in a similar manner to the cone probe of Schilden et al. The authors stressed the need to represent fluctuations spectrally, since RMS values tend to emphasize low-frequency contributions whereas boundary-layer receptivity can be at higher frequencies. Again, DNS was used to simulate the wedge probe, with results supporting the conclusion that slow acoustic waves (rather than fast) are the dominant disturbance in such facilities.

Conventional pitot measurements were also taken and compared with the wedge probe data by expressing the latter as RMS values integrated up to 50 kHz (the full wedge probe bandwidth was 500 kHz). The approach of Stainback and Wagner was used to convert the pitot data, including the same assumption of $u_s/u = 0.6$. Good agreement was found at $M = 6$ (both Ludwieg tubes) and $M = 7.4$ (HEG), but at $M = 3$ (HLB) the wedge probe gave double the RMS of the pitot, for reasons unknown.

Duan et al. (2019) is an overview paper summarizing a recent collaborative effort regarding hypersonic freestream disturbances. This includes the aforementioned work of Schilden et al. Other contributing works from this collaboration include the DNS studies of acoustic radiation from supersonic and hypersonic boundary layers (Duan et al., 2014, 2016), and the studies of pitot probe transfer functions (Chaudhry and Candler, 2017; Chaudhry et al., 2019). These are to be discussed next.

Duan et al. (2014) used DNS to simulate a fully-developed TBL over a flat wall at freestream $M = 2.5$. The height of the domain above the wall was much larger than the TBL thickness, i.e. the freestream far-field was simulated. Non-reflecting boundary conditions were used on the relevant portions of the domain, i.e. there was no attempt to include reflections of acoustic noise from opposing walls of the test section, although the authors mentioned that this is a possible complication in reality. Kovásznay modal analysis was applied to the freestream from these DNS results, and it was found that the magnitudes of all fluctuating quantities were small relative to their mean values, validating the small perturbation approach taken in previous works. Additionally, the acoustic mode was found to be "overwhelmingly dominant" over the negligible entropy and vorticity modes, when the only noise source was a TBL. Detailed spectra were provided for the pressure fluctuation field. The instantaneous pressure field was described as "plane-wave-like", albeit with limited spatial coherence. Within a streamwise–wall-normal plane, a preferred orientation could be computed; for the conditions of this study the angle between the wavefront normal and the freestream velocity was found to be $\theta_n \approx 132°$, i.e. a slow acoustic mode corresponding to $u_s/u \approx 0.4$. Visual inspection of pseudo-schlieren images agreed with this value of θ_n. The waves were not truly planar in the spanwise direction, as seen both from images and correlation studies. This finite spanwise extent was attributed to the acoustic sources within the TBL also being of finite size. The TBL itself was also investigated, and it was found that these acoustic sources were largely located in the buffer layer.

Duan et al. (2016) extended the study to $M = 5.86$, with results being compared to the previous $M = 2.5$ case. Here it was again found that the acoustic mode was dominant in the freestream. The normalized RMS magnitude of the pressure fluctuations increased from 0.4% to 2%, and θ_n decreased to $\approx 120°$; however, qualitatively the wavepackets remained the same. Within the TBL, the frozen-eddy assumption still held and the sources were again mostly in the buffer layer.

Chaudhry and Candler (2017) pointed out that many studies using Kovásznay modal decomposition rely on Stainback and Wagner's unsteady pitot formulation (such as Masutti et al.; Schilden et al.). DNS was used to compute the transfer function for pitot probes with the goal of ultimately replacing the Stainback and Wagner method; this paper was a first step in that it only considered flow-parallel disturbances. Three types of disturbance were considered separately: fast and slow acoustic waves $\sigma\pm$, and entropy waves Θ. These disturbances were linear combinations of sinusoids:

$$q'(t) = \sum_{k=1}^{N} q'_k \cos(\alpha_k x - \omega_k t + \phi_k) \tag{1.13}$$

where q' is a general fluctuating quantity. For acoustic waves, $q' = p'$, and the other quantities were given by:

$$\begin{bmatrix} \rho'(t) \\ u'(t) \\ T'(t) \end{bmatrix} = \begin{bmatrix} \frac{1}{\bar{c}^2} \\ \pm\frac{1}{\bar{\rho}\bar{c}} \\ \frac{\bar{T}(\gamma-1)}{\bar{\rho}\bar{c}^2} \end{bmatrix} p'(t) \tag{1.14}$$

and:

$$\alpha_k = \frac{\omega_k}{\bar{u} \pm \bar{c}} \tag{1.15}$$

with + and − corresponding to fast and slow modes, respectively. ω is an angular frequency, while in their notation c represents sound speed.

The entropy disturbances convect along streamlines, and were considered to be locally planar at the length scale of the pitot probe. For these, $q' = T'$ and the dependent quantities were:

$$\begin{bmatrix} \rho'(t) \\ u'(t) \\ T'(t) \end{bmatrix} = \begin{bmatrix} -\frac{\overline{\rho}}{\overline{T}} \\ 0 \\ 0 \end{bmatrix} T'(t) \tag{1.16}$$

with:

$$\alpha_k = \frac{\omega_k}{\overline{u}} \tag{1.17}$$

Frequency-wise transfer functions were defined as the ratio of the pitot transducer response to the freestream disturbance, in terms of the normalized power spectral density (PSD):

$$\chi(f) \equiv \frac{PSD\{p_{t,k}/\bar{p}_t\}}{PSD\{q_{1,k}/\bar{q}_1\}} \tag{1.18}$$

where subscripts t and 1 denote transducer face average and freestream quantities, respectively.

The authors derived a linearized Rayleigh pitot equation that held in the low-frequency limit, to be used for validation of the DNS results. This was:

$$\frac{p'_0}{\overline{p}_0} = \frac{p'_1}{\overline{p}_1} + \left(2\frac{u'_1}{\overline{u}_1} - \frac{T'_1}{\overline{T}_1}\right)\left(1 - \frac{1}{1 - \gamma + 2\gamma\overline{M}_1^2}\right) \tag{1.19}$$

where subscript 0 denotes stagnation point conditions—i.e. Eq. (1.19) used stagnation-point pressure, not transducer surface-average pressure; the difference was small, especially for low frequencies. The equation is valid for arbitrary fluctuations. Eq. (1.14) or Eq. (1.16) were substituted to get specific forms for the transfer functions, $\chi^*_{\sigma\pm}$ and χ^*_{Θ}, where * denotes a low-frequency limit analytical transfer function. These were found to be functions of $\overline{M}_1$ and γ only. The simulations of each disturbance case agreed well with their corresponding χ^* at low frequency (defined to be components between 0–20 kHz).

The 1D theory due to Morkovin suggests that resonance should be controlled by the standoff frequency $f_s \equiv c_0/2\Delta$, where Δ is the shock standoff distance. This should occur at half-integer multiples of f_s, with the primary resonance at $0.5f_s$. It was found that the simulated transfer function plots collapsed if normalized as f/f_s and

χ/χ^*. The global minimum response was $\chi/\chi^* = 1$, as $f/f_s \to 0$, with a global maximum due to resonance of $\chi/\chi^* \approx 4.5$ at $(0.415 \pm 0.006)f_s$ due to various 2D effects. Subsequent minima were near integer multiples of f_s, and had values $1 < \chi/\chi^* < 2$. These results held for all three types of disturbance considered. A general functional form for this normalized relationship was not given.

Chaudhry et al. (2019) built on the work of Chaudhry and Candler by extending the simulations to disturbances angled with respect to the pitot probe, and also with experimental comparisons. Only acoustic disturbances were considered, with fluctuations defined by a 2D analog of Eq. (1.13):

$$p'(t) = \sum_{k=1}^{N} p'_k \cos(\alpha_k x + \beta_k y - \omega_k t + \phi_k) \tag{1.20}$$

with the wavenumbers given by:

$$\alpha_k = \frac{\omega_k \cos\theta}{\overline{u}\cos\theta + \overline{c}} \tag{1.21a}$$

$$\beta_k = \frac{\omega_k \sin\theta}{\overline{u}\cos\theta + \overline{c}} \tag{1.21b}$$

Simulations were performed over the range $\theta = 100$–$130°$ as well as at $180°$ (this being the flow-parallel slow acoustic mode from the previous study). These were repeated for various pitot probe geometries. The results of these simulations showed that the frequency and magnitude of the resonance peak strongly depended on the disturbance angle θ; the transfer function also depended on the transducer geometry, more so than in the flow-parallel cases of Chaudhry and Candler. The low-frequency limit for χ monotonically increased as θ decreased, and was independent of geometry. This implied the existence of an extended analytical solution for χ^* as was done for the flow-parallel case, although a derivation was not attempted.

The experimental results showed that the χ simulated at $\theta = 120°$ matched quite well with selected experimental data for larger pitot probe geometries, but less so for the smallest size trialed. Additionally, unexplained differences were seen between the two experimental campaigns that suggested a Reynolds number dependence that did not manifest in simulations.

1.4 Impulse Facility Noise

1.4.1 Overview

Notwithstanding the brief reference to the HEG shock tunnel, all the facilities discussed in Section 1.3 were either conventional long-duration blowdown tunnels, or Ludwieg tubes[§]. Much of the development efforts for quiet tunnels have been targeted towards these types of facilities (Schneider, 2008). Quiet tunnels rely to a large extent on having a smooth, highly-polished nozzle throat; the high-temperature, particulate-laden flows of high-enthalpy impulse tunnels are inimical to maintaining such a surface finish.

Furthermore, the classical view presented above—of a well-conditioned reservoir smoothly connected via a nozzle to the test section—does not necessarily apply to impulse facilities: shock tunnels do have an effective reservoir, but thermodynamic conditions therein are extreme. Expansion tubes may not even have a nozzle, and therefore no reservoir. The nozzle or tube exit may open abruptly into the test section, giving strong turbulent shear layers around the core flow, which will radiate acoustic noise in addition to the nozzle boundary layers. Many impulse facilities rely on rupturing diaphragms, which involve complex flow interactions. These phenomena and others show that such facilities are likely to have additional sources of freestream disturbance when compared to blowdown tunnels, and various types of impulse facility may also be quite distinct from each other: e.g. the impulsive mechanism of flow acceleration in a shock tunnel is very different to that of an expansion tube.

As already discussed, a fundamental difference between conventional and impulse facilities is the much higher freestream enthalpy attained in the latter. Fujii and Hornung (2001) showed that some gases, such as air and CO_2, have significantly enhanced acoustic absorption at elevated temperatures typical of impulse facilities, including at high-frequency bands of relevance to transition studies. This result implies that acoustic radiation from wall and nozzle boundary layers may be more attenuated at the facility centerline, compared to equivalent conditions for a conventional tunnel.

Studies of freestream noise in impulse facilities appear to be rare in comparison with blowdown tunnels. This section will address a selection of work pertaining

[§]These operate in a blowdown fashion, but do not require control mechanisms to maintain time-independent reservoir conditions. Instead, the reservoir state remains constant until the expansion wave returns from its round-trip reflection from the far end of the storage tube; this results in short-duration test times (Igra and Seiler, 2015).

to expansion tubes and shock tunnels. The details of how these types of facilities operate will not be provided here; expansion tubes will be discussed in more depth in Chapter 2, while shock tunnels will not be described as they are not studied further in this thesis—the interested reader is instead referred to Part II of Igra and Seiler (2015).

1.4.2 Expansion Tubes

Historically, the first expansion tube designs found limited use due to unacceptably high freestream disturbance levels during test time. Norfleet et al. (1966) constructed a standard constant-diameter expansion tube, with a 680 atm helium driver, and a modified expansion tube which had a nozzle between the driven and acceleration sections and a driver pressure of 1700 atm. For both facilities, although the mean flow properties were in line with theoretical expectations, pitot pressure measurements showed large disturbances, which were strongest on the centerline. These disturbances were described as being "of sufficient magnitude to preclude meaningful aerodynamic testing". A range of tests were performed to eliminate several possible causes of these disturbances; it was concluded that the secondary diaphragm was the most likely source.

Spurk (1965) used streak interferometry to obtain time-resolved density measurements during start-up and the steady test time. He determined that the secondary contact surface was broadened by mixing. The temporal duration of this interface region was similar to the shock-contact surface interval time. Spurk attributed the majority of the flow irregularities to secondary diaphragm effects, and suggested that the acceleration and heating of the small diaphragm fragments extracts momentum and energy from the test gas slug.

Paull and Stalker (1992) proposed a theory of expansion tube freestream noise, based on the concept of a lateral acoustic wave originating in the driver section. This theory will be elaborated upon in Section 2.3, following the prerequisite introduction to expansion tube operation provided in Sections 2.1 and 2.2. For now, it suffices to summarize its main claims: that the freestream noise should reduce as the primary sound speed ratio a_3/a_2 decreases, and that the unsteady expansions cause a focusing of the initially-broadband noise toward distinct frequency bands.

Erdos and Bakos (1994) presented an overview of potential sources of noise in expansion tubes and tunnels. These were: primary and secondary diaphragm rupture, tube and nozzle wall boundary layers (the latter for tunnels only), wall surface roughness, steps and gaps, and lastly, stress waves and vibrations.

The primary diaphragm petals impact and cause vibrations within the tube walls, and were also thought to induce turbulence in the driver gas that flows over them—fast-acting valves have been suggested as a replacement for the diaphragm. The authors emphasized that the secondary diaphragm rupture process was still uncertain at the time of writing, although it was thought that the rupture is capable of inducing additional turbulence beyond that which originates from the driver gas. Indeed, data from the HYPULSE expansion tunnel showed the presence of lateral acoustic modes that could not have come from the driver; these were conjectured to originate instead from the secondary diaphragm. Erdos and Bakos referred to the then-recent theory of Paull and Stalker for this part of the analysis. Wall-static and freestream-pitot measurements yielded spectra with distinct peaks, compatible with the frequency-focusing aspect of the theory, although the quantitative match was not close enough to be conclusive. As an aside: the secondary diaphragm rupture is particularly difficult to model accurately; finite burst times and mass effects can lead to complex wave interactions (Furukawa et al., 2005; Shinn and Miller III, 1978). Various proposals have been made to either actively rupture the diaphragm ahead of the shock, or even do away with it entirely (Parziale et al., 2013a).

The boundary layer along the acceleration tube wall was found to have a large effect on the flow quality in HYPULSE: LBLs gave good flow quality, while fully-developed TBLs had noisier flows that were nevertheless considered still usable for some types of study. But if the BL was transitional, unacceptably-large pitot pressure fluctuations were observed. The authors determined the character of the BL via heat transfer measurements at the walls, and found the requirement on the unit Reynolds number was $Re_m \lesssim 7 \times 10^5\ \mathrm{m}^{-1}$. However, even with LBL the measured values of wall-static and freestream-pitot fluctuations were large enough to require an additional noise source other than wall boundary layers. This additional noise was thought likely to be due to surface roughness, and possibly also steps and gaps at the flanged segments. In expansion tubes, it is very difficult to maintain smoothness levels typically specified for quiet operation, due to particulates from the primary diaphragm.

The initial design and commissioning of the Hypervelocity Expansion Tube (Dufrene et al., 2007) drew from Paull and Stalker's theory. Two conditions were presented: one at $a_3/a_2 = 0.74$ showing large relative fluctuations in pitot pressure and the other at $a_3/a_2 = 0.44$ with substantially lower fluctuations. This latter condition was taken to be acceptably low-fluctuation, and the design criterion of $a_3/a_2 < 0.55$ was instated.

1.4.3 Shock Tunnels

Schneider (2001) briefly addressed freestream noise in shock tunnels, and reported on the few experimental characterizations available at that time, with inconclusive results. Hornung (2000) showed that the abrupt area reduction from the end of the shock tube to the nozzle throat causes the reflected shock wave to converge axisymmetrically onto the centerline. This process generates a strong vortex ring in the reservoir (Fig. 1.2), which is thought to ultimately lead to vorticity-mode noise in the freestream.

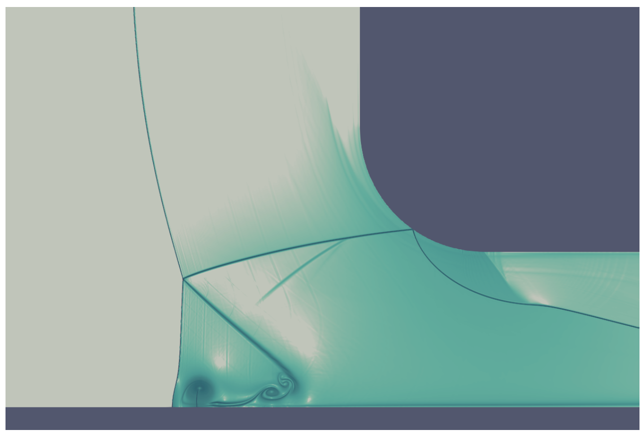

Figure 1.2: Vortex ring formation at shock tube–nozzle interface.

Grossir et al. (2013) performed a comprehensive survey with a T-shaped pitot rake on the VKI Longshot facility, which is a (non-reflected) shock tunnel with a piston in the driven section with an $M = 14$ nozzle. Pitot pressure fluctuations were presented at many locations, although due to resonance the data were lowpass-filtered at only 80 kHz. Away from the boundary layer, the RMS fluctuations in this bandwidth were in the range 4.5–7.5 %. When normalized by a Reynolds number based on the nozzle exit diameter, these data were slightly higher than, but comparable in trend, to data at the same M from AEDC Tunnel 9 (a conventional blowdown tunnel).

Parziale et al. (2014) used FLDI to make measurements of the density fluctuations in the T5 free-piston reflected shock tunnel at Caltech. They reported spectra over the band 5 kHz–20 MHz, and computed RMS values in various sub-bands of this range. These data showed that in the range of frequencies corresponding to slender-body hypervelocity boundary layer instabilities, the freestream noise was $\leq 0.5\%$. This experiment demonstrated both the need to appeal to optical techniques in order to surpass the frequency limitations of pitot probes, and that facilities like T5 may be quieter than assumed, particularly in bandwidths of interest to transition studies.

1.5 Project Scope & Outline

This work aims to:

1. Develop mature methodologies for the optical design and calibration of the instrument that expand on the successful pioneering works.
2. Extend our understanding of the capabilities of FLDI by developing and validating analytical and computational models for the instrument response.
3. Demonstrate multiple applications of the instrument to hypersonic ground testing facilities.

The thesis is structured as follows: Chapter 2 introduces the two hypersonic facilities on which FLDI was used to make measurements, and gives details of the fixtures and flow conditions used in each experimental campaign. Chapters 3 to 5 concern FLDI itself: Chapter 3 addresses the practicalities of designing and calibrating an FLDI system, as well as post-processing the raw signal. Chapter 4 contains validation studies for the ray-tracing model of the instrument, while Chapter 5 takes this validated model and uses it to derive analytical solutions for the recovery of quantitative density field data. In Chapters 6 and 7, FLDI is applied to the facilities introduced in Chapter 2, with the results being interpreted using the models of Chapters 4 and 5. Finally, Chapter 8 summarizes the thesis, and suggests directions for further research into FLDI.

Chapter 2

HYPERSONIC FACILITIES & EXPERIMENTAL PROCEDURES

This chapter gives details on the operating principles of two hypersonic facilities: Caltech's Hypervelocity Expansion Tube (HET), and the von Kármán Gas Dynamics Facility Wind Tunnel D. Experimental campaigns using FLDI were carried out on these two facilities; the design and goals of each campaign are discussed following the explanation of the respective facility. HET is addressed in Sections 2.1 to 2.5, with the corresponding results found in Chapter 7; Tunnel D is covered in Sections 2.6 and 2.7, with results in Chapter 6.

2.1 Hypervelocity Expansion Tube

Originally constructed and operated at the University of Illinois at Urbana-Champaign (Dufrene et al., 2007), the Hypervelocity Expansion Tube was subsequently relocated to Caltech, where it is one of the three main ground-testing facilities of the Caltech Hypersonics Group (CHG). Due to not having a nozzle, the freestream state is not restricted to a fixed Mach number. Expansion tubes are a type of impulse facility, where energy is transferred from the driver gas into the driven gas by a shock wave, followed by an unsteady expansion wave. This arrangement does not stagnate the test gas at any stage, unlike a reflected shock tunnel, and thereby avoids dissociation which could persist in the final test gas. However, test times are 160–500 μs, much shorter than the T5 facility.

HET comprises three sections of tube (referred to as the driver, driven/test, and acceleration/expansion sections), all with a constant internal diameter of 152 mm. A primary diaphragm of aluminum, with thickness selected from the range 0.8–1.6 mm, separates the driver and driven sections, while a much thinner (8 μm) secondary diaphragm of mylar separates the driven and acceleration sections. The flow exits the acceleration section as a free jet into the test section, which is square in cross-section with an internal dimension of ~ 230 mm. Optical access is possible from all four sides of the test section, although typically the bottom port is obscured by a mounting baseplate for test articles. The test section terminates with a dump tank, which together are fixed in place, while the tube is free to move axially along its mounting beam via linear bearings. The relative position of the tube exit and the windows can be varied between shots. An illustration of HET is given in Fig. 2.1.

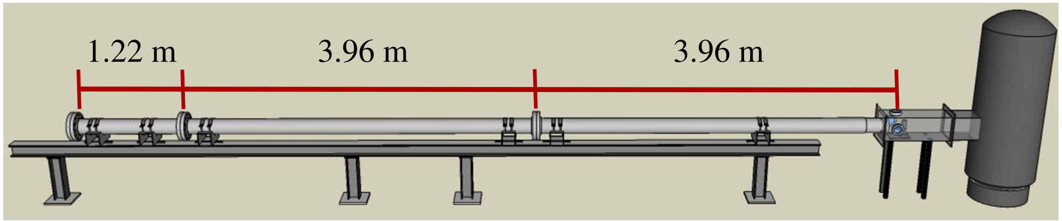

Figure 2.1: HET, with lengths of the three sections given.

The basic principle of operation of HET is as follows: the driven section is filled with the test gas of interest to $\mathcal{O}(1\,\mathrm{kPa})$, and the acceleration section is evacuated to $\mathcal{O}(100\,\mathrm{mTorr}) \approx \mathcal{O}(10\,\mathrm{Pa})$. The driver is then pressurized, usually with helium, until the primary diaphragm bursts at pressures given by Table 2.1. A knife-blade device in the primary flange aids with repeatability and causes the diaphragm to neatly rupture with four attached "petals", avoiding hazardous free-flying metal fragments (Sharma, 2010; Yanes, 2020). The rupture of the primary diaphragm due to passive driver pressurization results in a shock wave that provides the primary acceleration of the test gas in the driven section. When this primary shock reaches the secondary diaphragm, a secondary shock is transmitted into the acceleration gas, while the reflected wave, an unsteady expansion fan, is convected downstream in the supersonic flow. This unsteady expansion fan provides the secondary acceleration of the test gas to the final hypervelocity test condition. The steady test time begins with the arrival of the contact surface between the test and acceleration gases, and ends when either one of the expansion wave characteristics reaches the test section. These inviscid wave processes are summarized using an x–t diagram in Fig. 2.2.

2.2 HET Test Condition Simulation

Expansion tube solvers have previously been developed by other groups. The PITOT code is used with the X2 and X3 expansion tube facilities at The University of Queensland (James et al., 2018; James et al., 2013), alongside a 1D code, L1d3. Similarly, MacLean et al. (2010) developed the CHEETAh code for the LENS-XX expansion tube at CUBRC, which is used with their quasi-1D code, Jaguar. Each of these codes is tailored for in-house use with a particular facility, e.g. PITOT models the piston drivers and diverging nozzles of the X tubes, while CHEETAh includes high-density equation-of-state and rotational non-equilibrium effects that arise from the use of a heated hydrogen driver.

A inviscid, 1D, gas-dynamic calculator was developed with an accessible GUI interface. This code, known as LETS, is modular and extensible; for a more

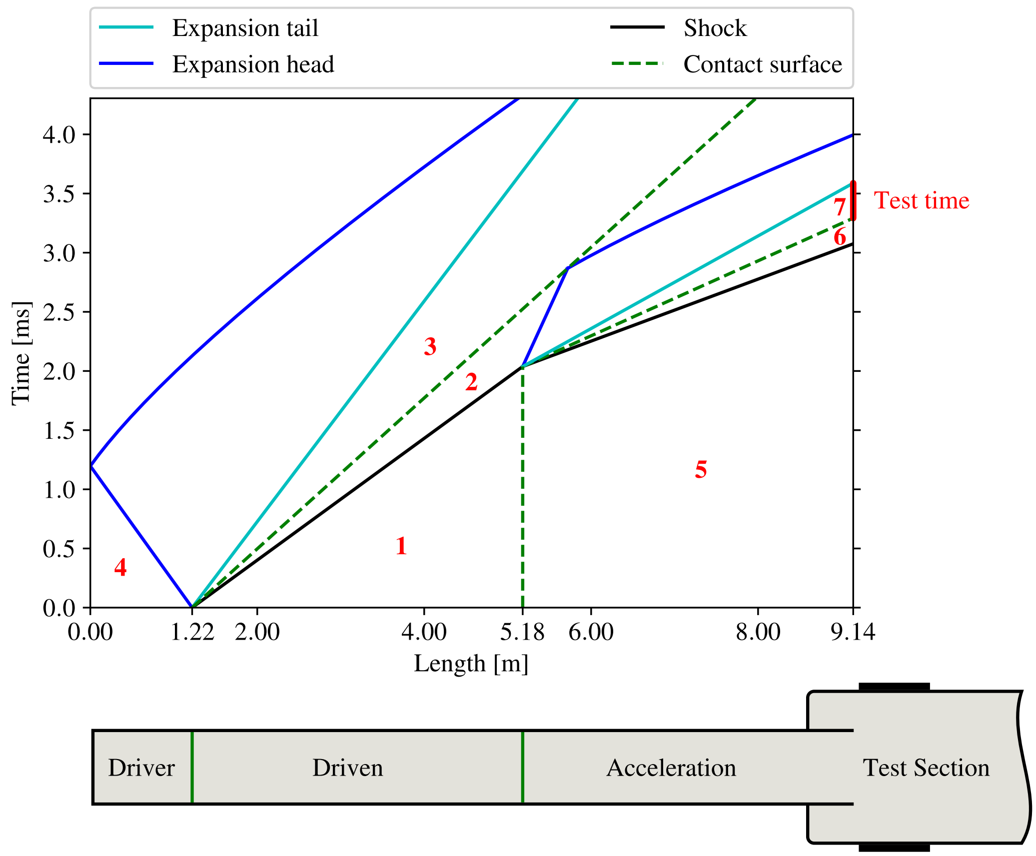

Figure 2.2: A representative x–t diagram for HET, with the conventional numbering system for the various gas states. The period of steady test time is indicated in red. Below the x-axis is a cartoon of HET showing the relative lengths of each section.

Table 2.1: Standard HET primary diaphragm thicknesses with corresponding burst pressures.

Thickness [mm (in.)]		Burst pressure [kPa]
0.8	(0.032)	1250 ± 40
1.0	(0.040)	1740 ± 70
1.3	(0.050)	2480 ± 140
1.6	(0.062)	3300 ± 90

complete description of its features and operation, see Lawson and Austin (2018). LETS performs perfect-gas and equilibrium calculations, the latter using Cantera (Goodwin et al., 2018) and the Shock & Detonation Toolbox (Browne et al., 2008). It features a full GUI where all input parameters are selected and where results are displayed immediately; at no point does the user need to interact with a command line. LETS is also able to run and visualize more complex multi-parameter sweeps by writing scripts that bypass the GUI and access the underlying functions directly. These visualizations can show the operating envelope of a given facility in terms of any combination of two input and two output variables, taking inspiration from Mollier diagrams. An example of this is shown in Fig. 2.3, where the input variables are the primary and secondary pressure ratios, with the performance envelope plotted in an output space of freestream Mach number versus total enthalpy.

LETS was validated via several comparisons with other solutions: a previous expansion tube code used within the group (Dufrene, 2006) and the example run conditions listed in MacLean et al. (2010)—which were in turn validated against CHEETAh. Additionally, a 1D Euler computation was performed* with equilibrium

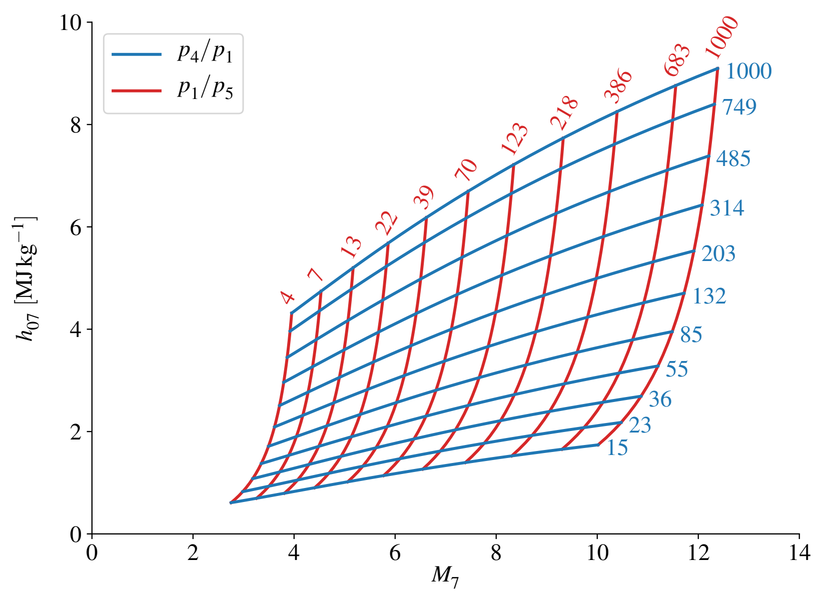

Figure 2.3: An example of the multi-parameter sweep capabilities of the LETS code. Plotted is the effect on freestream Mach number M_7 and total enthalpy h_{07} of changing the two pressure ratios of HET. Driver: He, driven: N_2, acceleration: He. Perfect gas assumed.

*By Prof. H. G. Hornung.

chemistry, using the Eilmer code from the University of Queensland (Jacobs and Gollan, 2016). The comparison between this result and the prediction of LETS is shown in Fig. 2.4. All wave trajectories are in quantitative agreement, including the non-linear reflections of the expansion heads. The only disagreement is in the primary contact surface: Eilmer predicts that it accelerates after its interaction with the secondary expansion head; LETS does not attempt to model the primary contact surface because it has no influence on the steady freestream state of interest. Predictions of relevant thermodynamic parameters were also validated, with p, T, ρ, and u all agreeing to within 0.5 %.

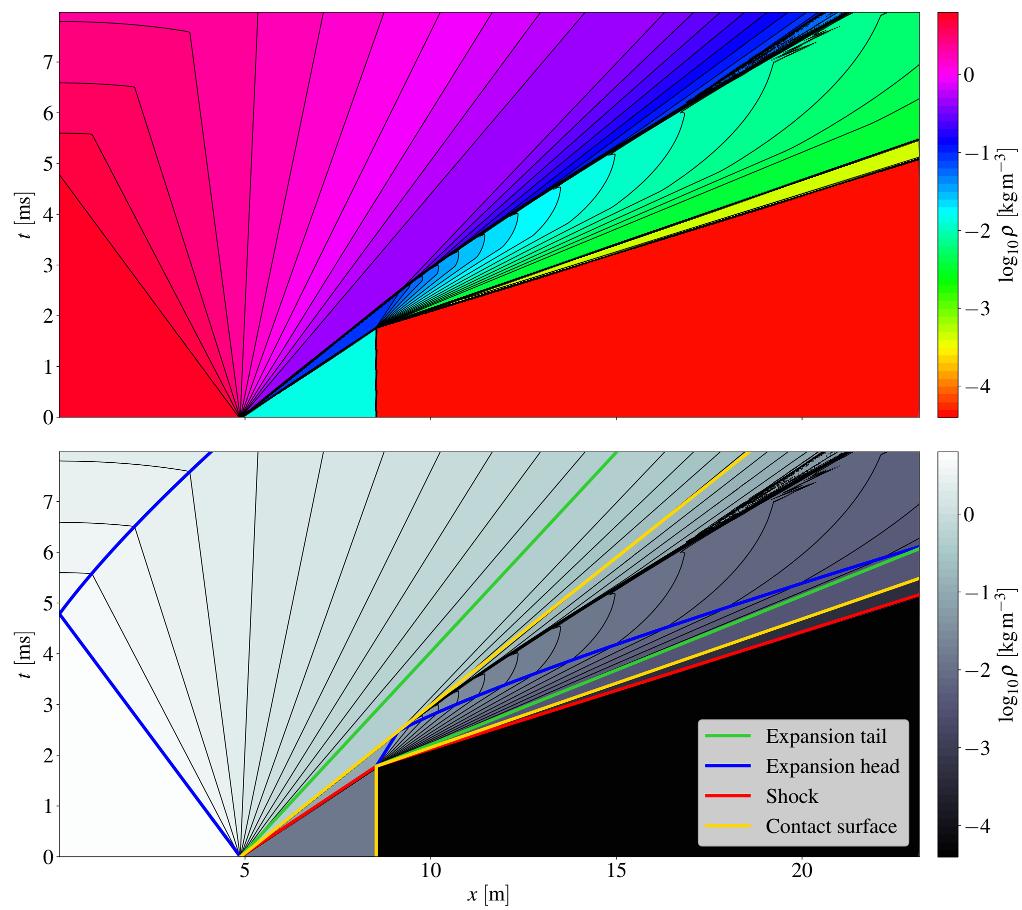

Figure 2.4: Comparison of predicted x–t diagrams for Eilmer and LETS. The upper plot reveals the wave trajectories using isocontours of density from the Eilmer data. The lower plot has the same contours in greyscale, with the semi-analytical wave trajectories from LETS overlaid in color.

For more accurate estimates of the time-averaged freestream state, viscous effects need to be taken into account. Mirels (1963, 1964) showed that in shock tubes, the growth of boundary layers along the tube walls leads to velocity changes for the shock and contact surface. These features trace out curved trajectories in the x–t plane, rather than straight lines as in the inviscid case; instead of diverging indefinitely, the shock–contact surface separation distance tends to a constant maximum value l_m. This in turn alters both the post-shock thermodynamic state and the duration of test time. Efforts are currently underway in the research group to account for these viscous effects using simulations in Eilmer.

2.3 Expansion Tube Lateral Acoustic Wave Theory

Expansion tubes have several possible sources of noise, as detailed by Erdos and Bakos (1994) and summarized in Section 1.4. Here, more detail is provided on one particular noise source, the theoretical implications of which were extensively derived by Paull and Stalker (1992).

They postulated that a dominant source of noise in expansion tubes is lateral acoustic waves in the driver section. These waves are transmitted across the primary contact surface into the test gas, where they persist during the freestream test time flow. The details of these waves are given in Chapter 7. This model indicates that freestream noise can be mitigated by reducing the sound speed ratio a_3/a_2 across the primary contact surface. The ratio of expanded driver to shocked driven gas, a_3/a_2, is a strong function of the initial driver pressure p_4, and high p_4 corresponds with both a reduction in a_3/a_2 and an increase in freestream total enthalpy h_{07} (Fig. 2.5). Hence, expansion tubes operated in a so-called "high-enthalpy" mode also profit from decreased a_3/a_2 ratios; an upper bound of $a_3/a_2 < 0.55$ was found to produce acceptably low noise levels during test time (Dufrene et al., 2007). However, a_3/a_2 is not a function of p_4 alone as it also depends on the specific heat ratios of both gases, γ_1 and γ_4; nor is it the only determining factor (in the framework of Paull and Stalker's theory) concerning the magnitude and spectrum of noise transmitted to the test gas.

Paull and Stalker further showed that the primary contact surface effectively acts as a low-frequency filter, i.e. significant reflection of noise back into the driver only occurs below a threshold frequency. This frequency is proportional to the initial sound speed in the driver gas, a_4, which also suggests heating the driver in order to raise a_4. After the acoustic waves transmit from the driver gas into the test gas,

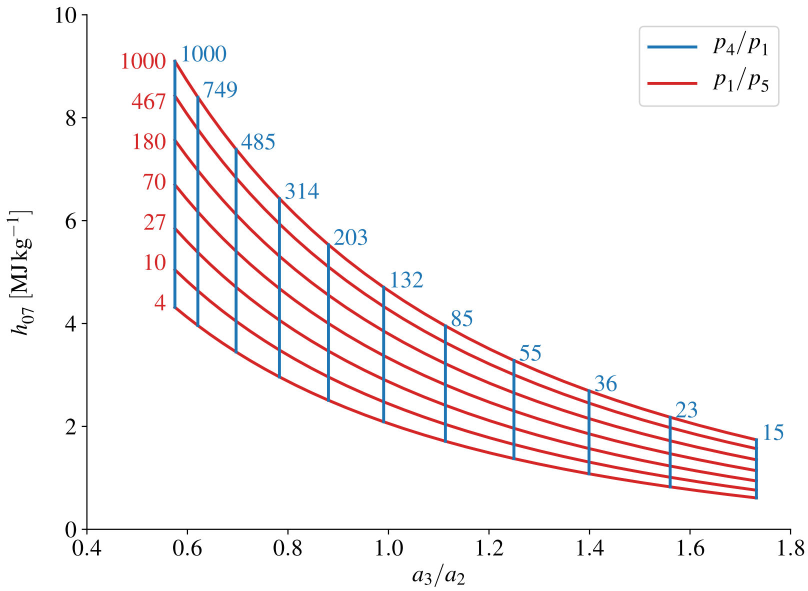

Figure 2.5: The relationship between freestream total enthalpy h_{07} and the primary contact surface sound speed ratio a_3/a_2, over an arbitrary operating range defined by the primary and secondary pressure ratios. Driver: He, driven: N_2, acceleration: He. Perfect gas assumed.

they are further processed as the test gas expands through the unsteady expansion resulting from the rupture of the secondary diaphragm. They demonstrated that this expansion causes a “frequency-focusing” effect on the disturbances, where the initial broadband frequencies of the acoustic wave converge towards “discrete narrow bands of frequencies”. This effect is a function of the strength of the secondary expansion wave (addressed further in Chapter 7).

Hence according to this theory, the magnitude and spectral content of the freestream noise depends on the compositions and initial pressure of all three sections of an expansion tube, and three quantities of relevance were proposed:

1. The primary contact surface sound speed ratio, a_3/a_2
2. The initial driver gas sound speed, a_4
3. The secondary unsteady expansion sound speed ratio, a_2/a_7

These factors were predicted to influence the following aspects of the freestream noise spectrum, respectively:

1. The overall magnitude of disturbances at all frequencies
2. The high-pass frequency threshold
3. The degree to which the spectra are "focused" into discrete frequency bands

In conclusion, Paull and Stalker developed a theory that thoroughly explored the implications of the propagation of acoustic waves initiated in the driver gas. Seven conditions were measured experimentally using pitot probes, including both shock-tube and expansion-tube shots. The rationale for performing shock-tube shots is that it isolates the first Riemann problem, i.e. it allows investigation of noise transmission across the primary contact surface only, without complicating effects from the secondary wave system. These preliminary experiments showed spectral features consistent with their theory, but conclusive proof from an extensive experimental campaign is not yet available in the literature. For example, there are no studies that test for systematic trends in freestream noise metrics when a_3/a_2 is varied.

2.4 HET Experimental Campaign Design

A study of the HET freestream noise was performed using FLDI. Being impulsive and high-enthalpy, HET presents a harsh environment for intrusive diagnostics, and its short test times mean the lowest relevant frequencies are $O(10\,\mathrm{kHz})$. This makes FLDI an attractive alternative for obtaining highly-resolved test time data.

2.4.1 Goals

The campaign firstly aimed to test the theory of Paull and Stalker, by looking for trends in the freestream noise as a_3/a_2, a_4, and a_2/a_7 were varied. Arguably, a_3/a_2 is the most important of these quantities for many applications, since it determines how much driver noise is transmitted into the test gas in the first place. The driver gas sound speed a_4 is still important for this too: Paull and Stalker pointed out that argon as a driver gas will give better a_3/a_2 than helium, yet should still yield a noisier test gas because of a reduction in the bandwidth over which transmission is attenuated, thus allowing higher-frequency components from the diaphragm rupture process to be transmitted. In HET, a_4 can be varied only by the choice of driver gas—although in practice, all of the standard characterized conditions employ helium only. a_2/a_7 is of lesser importance because it modifies the spectral shape of the noise after it has already been transmitted into the test gas. With the above points in mind, the main objective was exploring the effect of varying a_3/a_2.

Paull and Stalker only considered noise originating in the driver, and then only of a particular form (lateral acoustic waves). Other possible sources of noise should not be discounted, e.g. turbulent boundary layer radiation, secondary diaphragm/contact surface effects. Hence, trends in parameters related to these effects, e.g. the Reynolds number Re, will also be examined.

2.4.2 Challenges

Some previous supersonic and hypersonic applications of FLDI were in: Caltech's T5 reflected shock tunnel ($0.034\,\mathrm{kg\,m^{-3}} < \rho_\infty < 0.075\,\mathrm{kg\,m^{-3}}$) (Parziale et al., 2014), the Penn State Supersonic Wind Tunnel ($\rho_\infty = 0.6\,\mathrm{kg\,m^{-3}}$), and the AEDC Hypervelocity Tunnel 9 ($\rho_\infty = 0.0385\,\mathrm{kg\,m^{-3}}$) (Fulghum, 2014). However, two of the most commonly-used HET conditions have $\rho_\infty < 0.004\,\mathrm{kg\,m^{-3}}$, an order of magnitude lower even than the Tunnel 9 condition, which Fulghum described as "very low" density for FLDI.

These values only offer an indication of likely signal levels, since FLDI responds to fluctuations in density, rather than the absolute value of the mean density—a very quiet tunnel might give no signal, even if the freestream density was high. Furthermore, the gas composition is important: an identical ρ' in helium and air will yield very different n'. Nonetheless, with such low average densities, HET poses a challenge for recovering adequate signal-to-noise ratios (SNR), even with high relative fluctuations.

2.4.3 Experimental Campaign

As already mentioned, the six independent variables for designing an HET condition are the initial compositions and pressures of each of the three sections. Although the facility has the capability for arbitrary gas mixtures in the driven and acceleration sections, this campaign uses four gases: He, Ar, CO_2, and air to examine possible real-gas effects. He and Ar are monatomic and near-ideal in behavior while having an order of magnitude difference in atomic mass, air is largely diatomic, and CO_2 has more degrees of freedom while being of similar mass to Ar.

The driver pressure p_4 takes on discrete values, determined by the primary diaphragm thickness (Table 2.1). The exact burst pressure is recorded and used to adjust the test condition calculation from the nominal values. Because HET was designed primarily for use with a helium driver, the use of heavier driver gases is restricted to the thinner diaphragms, otherwise the tube recoil is so large that it risks the nozzle pulling right out of the test section.

The driven and acceleration pressures (p_1 and p_5, respectively) are continuously variable. The pressure differential across the secondary diaphragm ($\Delta p = p_1 - p_5$) is restricted to $\lesssim$ 10 kPa to prevent premature rupture. p_1 and p_5 are bounded from below by the performance of the vacuum pumps, about 0.5 kPa and 5 mTorr $\approx$ 0.67 Pa respectively. The upper bound on p_5 is more arbitrary, although generally it should be kept low ($\leq$ 100 mTorr) for hypervelocity performance; here, 1 Torr $\approx$ 133 Pa is used, as a trade-off between increased freestream density ρ_7 (and thereby FLDI sensitivity) and maintaining conditions relatively similar to the standard operating range. This gives an upper bound on p_1 of 10 kPa.

These bounds on composition and pressure can be used to compute performance envelopes using LETS. A previous example of this was shown in Fig. 2.3. For clarity, just the bounding contours of this type of plot can be displayed, to give a performance envelope in various parameter spaces. This is done in Fig. 2.6 to show a_3/a_2 vs. n_7. The available gases give a total of 64 combinations, however only cases with air as an acceleration gas are shown, because this is found these give the highest n_7—fortuitously, as air is the simplest acceleration gas to work with, reducing the turnaround time between shots since less purging is required.

It is not possible to design campaigns where only one parameter varies while the rest remain constant. For example, with fixed compositions, one can move along a contour of fixed p_5 while varying p_1: this will change not only a_3/a_2, but also M_7, Re_7, h_{07}, etc. Instead, we attempt to keep n_7 as high as possible (for good SNR) while covering a wide range of both a_3/a_2 and other parameters of interest. A large number of shots were performed so that various subsets of data would be dense enough to show trends for each parameter. To reduce dependence on outlier shots or unforeseen correlations between parameters, the same value of a_3/a_2 was achieved in multiple shots using very different input conditions.

Because p_5 has no effect on a_3/a_2 (which only relies on the primary Riemann solution) it was kept higher than in usual HET conditions (1000 mTorr vs. a more typical 50 mTorr), in order to improve n_7. Test conditions were computed using a range of combinations of driver and driven gases; helium was never used as a driven gas because as shown by Fig. 2.6, its very low refractivity results in inadequate test time signal.

A full list of shots performed for this campaign is given in Appendix E; the post-processing and results are found in Chapter 7.

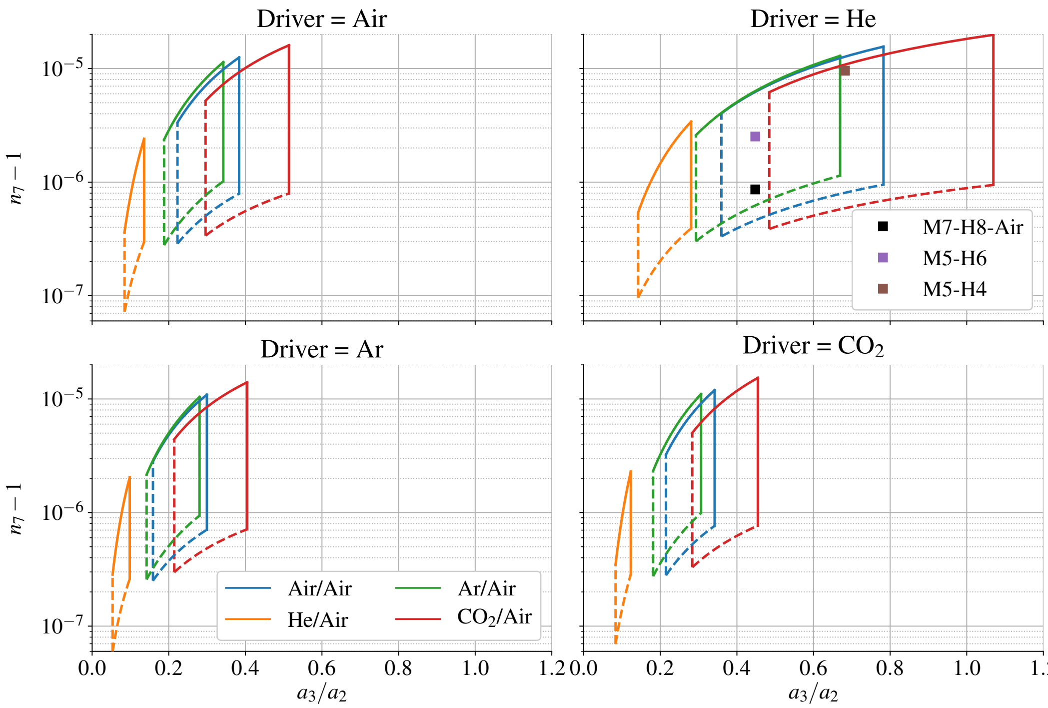

Figure 2.6: HET performance envelopes in an a_3/a_2–n_7 space, for fixed P_4 corresponding to the thickest primary diaphragm. Envelopes bounded by $0.5\,\mathrm{kPa} < p_1 < 10\,\mathrm{kPa}$ and $10\,\mathrm{mTorr} < p_5 < 1000\,\mathrm{mTorr}$. Each subplot is for a given driver gas composition, colors represent the driver/acceleration gas compositions. Dashed contours indicate the lower bounds of these ranges, solid contours the upper bounds. Some standard HET conditions are indicated for reference.

2.5 HET Optical Arms

During the HET flow-starting process, a complex outer flow-field is established due to an initial reflection of the transmitted shock from the window recesses, followed by further reflections in the cavity between the tube and test-section walls. These shock reflections occur outside the core flow, however spatial filtering is unable to fully remove their influence on the signal from the outer parts of the FLDI beams. This is because shock waves are both very thin, and have a large jump in density, giving very high local gradients in refractive index. Although the reflected shock signals are much-attenuated compared with that from the incident shock, they still yield a noisy signal that is indistinguishable from actual freestream noise. Furthermore, the length scales of the test section are such that the reflected shock interactions occur during the same period as test time when operating in expansion-tube mode. For further details on this shock reflection process and its influence on the FLDI response, see Section 4.4.

In order to mitigate the effects of these shock reflections on the FLDI signal, it was decided to construct a pair of optical arms that would shield the FLDI beams from the outer flow. Birch et al. (2020) also employed what they termed "flow shrouds" to shield the outer portions of FLDI beams, although their goal was to negate the influence of the nozzle boundary and shear layers. These layers were a concern because their facility geometry necessitated narrow FLDI beams with reduced spatial filtering, and their longer quasi-steady test times of 200 ms allowed measurement of lower-frequency components that are less attenuated in the outer beams.

Here, a "cookie cutter"-type geometry was used at the ends of hollow cylindrical optical arms to provide an undisturbed freestream core for the FLDI to measure (Figs. 2.7 and 2.8). A similar geometry was used by Parker et al. (2006) to shield beams for tunable laser diode absorption spectroscopy in the LENS I shock tunnel. In its normal configuration, HET has optical access through three glass windows. The optical arm assembly has a flange with the same dimensions as this original window that allows it to be clamped into position in the same way. The cookie cutters have sharp leading edges designed to keep the oblique shock attached over a relevant Mach number range. In order to keep the inner flow (between the parallel faces of the cookie cutters) as undisturbed as possible, the bolt heads are recessed and back-filled flush with the surface using low-viscosity RTV silicone, and the optical window face is also flush. For more information on the design rationale and specifications for these optical arms, see Appendix F.

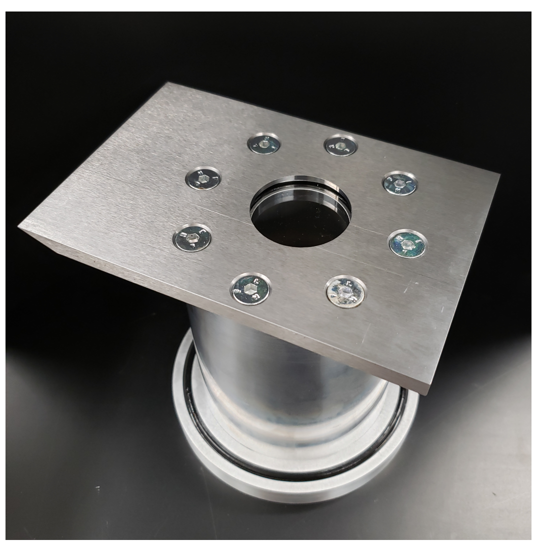

Figure 2.7: Photograph of optical arm assembly.

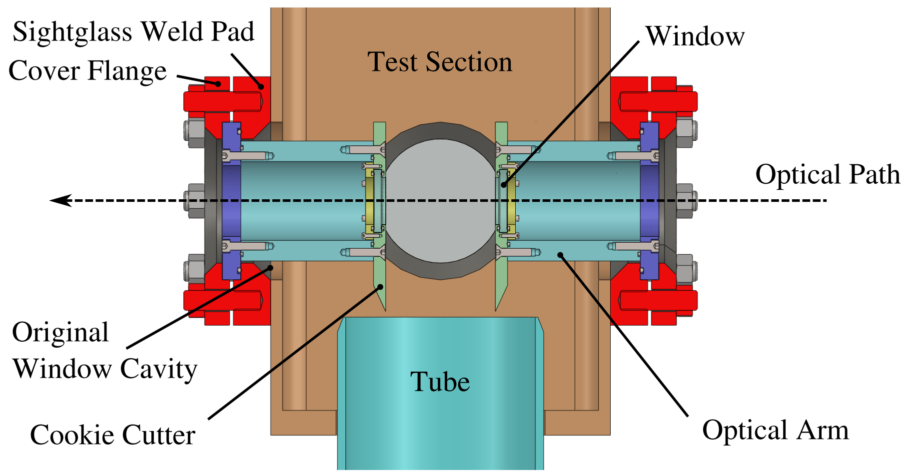

(a) Top-down cross-sectioned view.

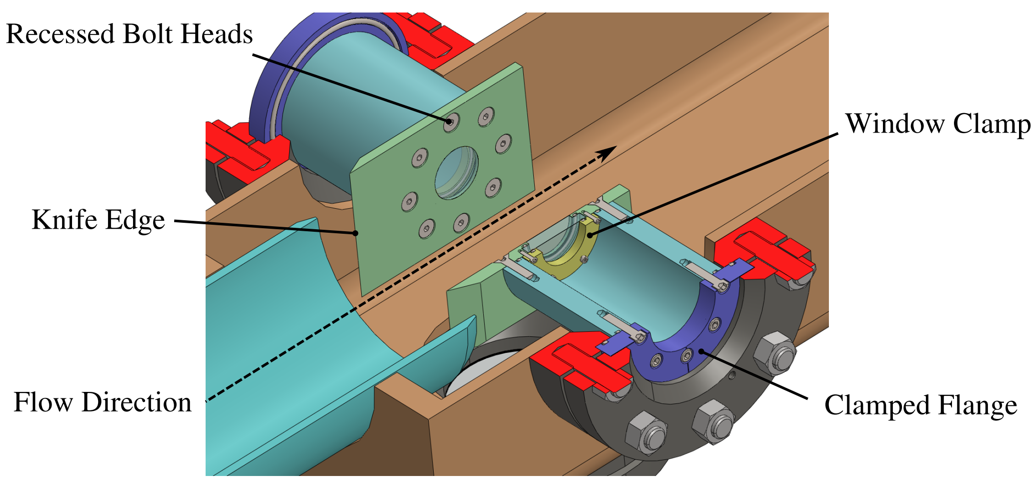

(b) Oblique view, with one optical arm intact and the other cross-sectioned.

Figure 2.8: CAD model of optical arms installed in HET test section.

2.6 VKF Tunnel D

The von Kármán Gas Dynamics Facility (VKF) Wind Tunnel D is a blowdown facility at the Arnold Engineering Development Complex (AEDC) in Tennessee. Originally in operation from the early 1950s through the late 1970s, it was reactivated and modernized from 2016–19 (Hofferth and Ogg, 2018, 2019). Tunnel D features a 12 in. (304.8 mm) square test section with flexible nozzle walls that allow air flows of Mach 1.5–5.0 with stagnation pressures 34–414 kPa. The high-pressure air supply is electrically heated at the inlet, although only enough to prevent liquefaction during expansion through the nozzle. Downstream of the test section is a large (22 m diameter) vacuum sphere that can be brought to pressures of 5–10 Torr; this allows for increased performance when compared to venting to atmospheric pressure. Tunnel D can sustain steady freestream flows for times on the order of minutes.

2.7 VKF Tunnel D Experimental Campaign Design

2.7.1 Goals

The goal of this campaign was to obtain simultaneous pitot and FLDI measurements of the Tunnel D freestream, over the full range of stagnation pressures (and therefore Reynolds number). The main objective of obtaining such data is to test hypotheses on how to convert each instrument's output into static pressure or density fluctuations. If consistent spectra are obtained between the two methods at moderate frequencies, then the faster response of FLDI will allow for extension of the spectra beyond the bandwidth of the pitot. Tunnel D is an ideal facility in which to test these concepts, because it is a conventional wind tunnel with relatively well-understood disturbance mechanisms. Furthermore, it is one of the few facilities where important aspects of the freestream noise have already been characterized experimentally via pitot-HWA (Donaldson and Wallace (1971), see Section 1.3 for further discussion). The long test times allow for extensive record lengths and correspondingly detailed spectra.

2.7.2 Experimental Design

The campaign comprised a total of six runs (see Table 2.2). In each run, the facility nozzle was set to a nominal Mach number of either 4.0 or 5.0, then a continuous blowdown lasting several minutes was executed. The reservoir pressure p_{res} was varied in a stepwise fashion during this blowdown, while maintaining a nominally constant reservoir temperature. At each pressure step, a 4 s burst of high-speed data was acquired, and so each of these datasets are designated as "Run x, Burst y". The bursts are numbered in ascending order of p_{res} (and hence Re_m, per Table 2.3).

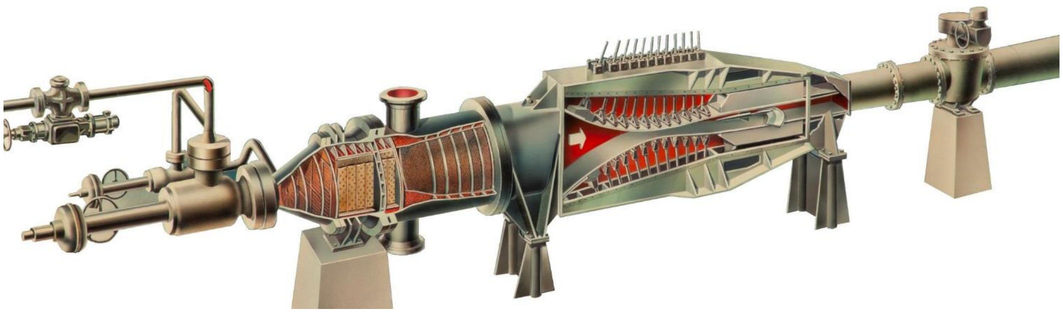

(a) Perspective-view cutaway illustration.

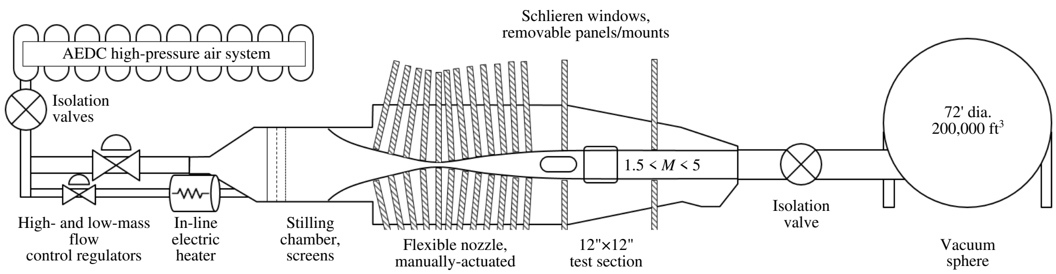

(b) Schematic.

Figure 2.9: Views of VKF Tunnel D. Source: AFRL/RQHX. Used/modified from original with permission from J. W. Hofferth.

The pitot probe rig comprised two pressure transducers [Kulite XCQ-062-10A], mounted at the horizontal midplane of the test section ±2 in. (50.8 mm) from the vertical midplane. These were oriented with the sensor face directly into the oncoming flow, i.e. operating as pitot pressure probes. The FLDI was constructed such that the optical axis also lay in the horizontal midplane, with its foci directly upstream from and coaxial with one of the pitot probes, i.e. the focal plane of the FLDI was offset 50.8 mm from the vertical midplane, towards the catch-side optics. A single isolated optical table spanned the test section, passing underneath the facility. A pair of rigid frames† were constructed to elevate each rail of the FLDI to be in line with the test section windows at either end of this optical table.

The high-speed data bursts were acquired at 3 MS/s with 16-bit depth [National Instruments PXIe-6368]. Both pitot probes were acquired without filtering; the

†Designed by B. Valiferdowsi

Table 2.2: VKF Tunnel D run summary

Run	M_∞	Notes
181	5.0	Shakedown run, data not used
182	5.0	
183	5.0	Conditioned FLDI channel data lost
184	5.0	Accelerometers added
185	4.0	
186	4.0	Original coaxial Kulite failed: swapped with other one

Table 2.3: Unit Reynolds numbers corresponding to the minimum and maximum Tunnel D reservoir pressures at each Mach number used in the campaign.

M_∞	Min. Re_m [$\times 10^6$ m^{-1}]	Max. Re_m [$\times 10^6$ m^{-1}]
4.0	1.4	17.9
5.0	1.0	10.4

FLDI signal was passed through a 50 Ω inline terminator then split into two channels: one raw and DC-coupled, the other AC-coupled, amplified (50× or 100×), and low-pass filtered at the Nyquist frequency of 1.5 MHz. A range of other quantities (e.g. reservoir condition) were recorded at a much lower rate of 10 S/s. The facility operating software uses these to automatically compute additional freestream quantities via gasdynamics. This acquisition was performed by a separate system to the high-speed sampling, although the two can be synchronized using an IRIG time reference. The result is that about 40 instances of freestream state parameters are available for every burst of high-speed data: sufficient to describe the time-averaged state, but not resolved enough to correlate any time-dependent features.

The pitot probe rig and data acquisition systems were designed and installed by the AEDC research engineers, who also operated the facility. The FLDI was constructed and operated together with M. C. Neet.

Chapter 3

FOCUSED LASER DIFFERENTIAL INTERFEROMETRY

This chapter is concerned with the design, operation, and data processing for FLDI systems. Section 3.1 gives the basic optical principles needed to select the components for a standard FLDI, and Section 3.2 extends these to cover double FLDI systems. Section 3.3 demonstrates how to assemble and align the components to obtain a working interferometer, and Section 3.4 gives procedures for calibration and post-processing that yield quantitative phase shift data.

The interpretation of the phase shift data in terms of flow quantities is left to Chapters 4 and 5; this chapter is chiefly concerned with the practicalities of the instrument itself.

3.1 Single FLDI Optical Design Procedure

The optical design of an FLDI is an optimization problem, with the objective of minimizing w_0, subject to dimensional constraints from the particular experimental layout and the available optical components. Here, a design procedure for the single FLDI configuration shown in Fig. 3.1 is given, although it is easily adapted for alternative optical arrangements.

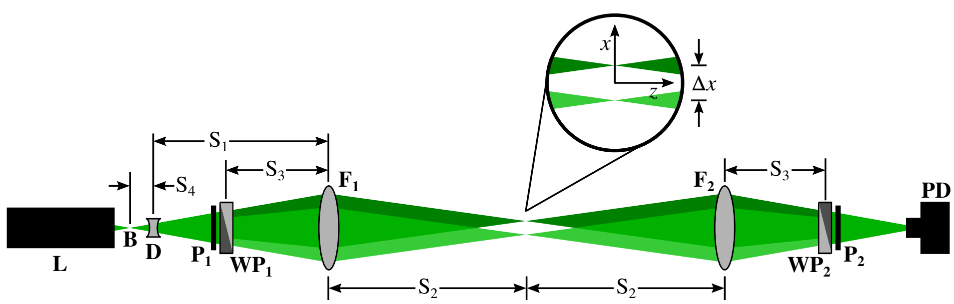

Figure 3.1: FLDI schematic repeated from Fig. 1.1, with dimensions S_i added for discussion in this chapter.

3.1.1 Ray Transfer Matrix Analysis

Ray transfer matrix analysis combined with Gaussian optics (Milonni and Eberly, 2010) can be used to model the transmitting side of the system. In this approach, a complex beam parameter $q(z_n)$ is modified by transfer matrices that represent the various optical components. z_n is the axial distance measured from the natural beam waist (labeled **B** in Fig. 3.1), so to use this method the location and Gaussian radius of this waist are required. These data are often provided by the laser manufacturer, but should be confirmed experimentally using a beam profiler. The fixed value q_n at the natural waist is related to q_f at the test-section focus by:

$$q_f = \frac{Aq_n + B}{Cq_n + D} \tag{3.1}$$

A, B, C, and D are the components of the overall transfer matrix $\mathbf{M}$, here composed from five elementary matrices representing three propagations and two thin lenses:

$$\begin{aligned} \mathbf{M} &= \mathbf{M_5 M_4 M_3 M_2 M_1} \\ &= \begin{bmatrix} 1 & S_2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/f_2 & 1 \end{bmatrix} \begin{bmatrix} 1 & S_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/f_1 & 1 \end{bmatrix} \begin{bmatrix} 1 & S_4 \\ 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} A & B \\ C & D \end{bmatrix} \end{aligned} \tag{3.2}$$

S_1, S_2, and S_4 are dimensions given in Fig. 3.1; f_1 and f_2 are the focal lengths of the diverging (**D**) and focusing (**F₁**) lenses, respectively. The beam parameter is given by:

$$q(z_n) = \left[\frac{1}{R(z_n)} + \frac{\mathrm{i}\lambda_L}{\pi w^2(z_n)} \right]^{-1} \tag{3.3}$$

The wavefront curvature R must be zero at a waist, and the test-section focus is by definition another beam waist. It can be shown that this leads to the requirement $\Re(q_f) = 0$, where $q_f \equiv q(z_n = S_1 + S_2 + S_4)$. The radius at the focus w_0 can then be extracted from $\Im(q_f)$*. Although the optimization is over a 5-dimensional space $\{S_1, S_2, S_4, f_1, f_2\}$, in practice lenses are only available with discrete focal lengths, giving instead a series of separate 3-dimensional optimizations for the lens positions only. Typical bounds on this optimization include: S_2 must be greater than the half-width of an enclosed test section, and local beam widths must not exceed some

*$\Re$ and $\Im$ denote the real and imaginary part, respectively.

percentage of the lens apertures. The natural beam waist can be quite close to the laser aperture, and if a periscope is used, this puts another bound on how short S_4 can be.

Note that the polarizer and Wollaston prism are not included in this analysis. For the small prism angles typically used for FLDI, the paraxial approximation holds, and the splitting of the beam by the prism does not significantly alter the size or axial location of the foci. However, when extending this method to a double FLDI system, the prisms do need to be included in the above derivation. This is addressed in the following Section 3.2. In order for the pair of beam centerlines to run parallel in the test section (i.e. a constant Δx), the Wollaston prism $\mathbf{WP_1}$ must be placed at the focal point of the lens $\mathbf{F_1}$, i.e. $S_3 = f_2$. Through simple trigonometry:

$$\Delta x = 2S_3 \tan\left(\frac{\theta}{2}\right) \approx S_3\theta \tag{3.4}$$

where θ is the full prism splitting angle, and the approximation holds for small θ. Because both the Wollaston prism and focusing lens have finite thicknesses in reality, it can be difficult to know where exactly to measure S_3 from, and some perturbation of the prism position while observing with a beam profiler may be required to ensure Δx is sufficiently constant. The true behavior of Wollaston prisms is significantly more complicated than implied by Eq. (3.4), particularly when the incident ray is not normal to the prism. Indeed, for a focusing or diverging beam (as encountered in FLDI), all rays except for the centerline are incident at an angle to the prism; in a double FLDI even the centerlines of the primary beams are not normal to the secondary prism (see Fig. 3.2). Nevertheless, for the small angles involved, it is found that the simplistic approach taken here is adequately accurate; comparisons of theoretical and actual Δx are given at the end of Section 3.2. If higher precision is desired, a useful treatment of the full behavior of Wollaston prisms is given by Soref and McMahon (1966).

3.1.2 Polarizer Selection

The simplest FLDI configuration is for $\mathbf{P_1}$ and $\mathbf{P_2}$ to both be linear polarizers (LP), oriented with their axes 45° to the principal axes of the Wollaston prisms. This causes a reduction in intensity by a factor of two at each polarizer. Another simple configuration uses a quarter-wave plate (QWP) for $\mathbf{P_1}$, exploiting the fact that most lasers are already strongly linearly polarized. After passing through $\mathbf{WP_1}$, the beams are still linearly, orthogonally polarized, albeit antiphase instead of in-phase as when

$\mathbf{P_1}$ is a LP. This makes no material difference, as the initial phase difference $\Delta\Phi_0$ gets selected during the calibration process anyway (see Section 3.4). Additionally, using a QWP avoids intensity loss. However, $\mathbf{P_2}$ still needs to be a LP in order to recombine and interfere the beams. Smeets and George (1973) provide many other viable prism/polarizer configurations.

The reason for the use herein of QWPs over LPs is the troublesome diffractive effects encountered when using a LP. The original FLDI system the author inherited from Parziale used two LPs of type Thorlabs LPVISB100, which functioned as expected. However, when constructing a duplicate system, that model of LP had been discontinued, with the recommended replacement being LPVISC100. It was found that this LP model caused the beams to diffract into a 2D pattern instead of remaining a single beam. Other replacement types of LP trialed also had this issue; through discussion with the manufacturer it was determined that the fabrication process had changed. It is suspected that differences in the lamination or nanoparticle structure of the LP led to this issue. Conversely, QWPs are constructed from simple birefringent materials and thus avoid the potential diffractive issues stemming from the LP microstructure. This work uses a Thorlabs WPQ10M-532. Although there must still be available LPs that behave correctly, this experience should serve to show why a beam profiler is an invaluable tool when constructing an FLDI, because beam issues such as this would be difficult to detect if aligning just by the unaided eye.

3.2 Double FLDI Optical Design Procedure

The ray-tracing matrix methods used in Section 3.1 for the design of the single FLDI can be extended to the double FLDI. Similarly, the design procedure has two stages: first, the lenses are positioned in order to minimize w_0; second, the Wollaston prisms are located after the lenses have already been fixed in place. The goal of the second step is still to ensure that all 4 beam centerlines run parallel in the test section. However, the computations are more involved compared to the single FLDI, because the primary prism $\mathbf{WP_A}$ should not simply be placed at the focal point of the primary lens $\mathbf{F_A}$. Instead, it needs to be positioned such that the centerlines of the two beams resulting from the primary split cross the optical axis exactly at the secondary prism $\mathbf{WP_B}$, as was independently discovered by Bathel et al. (2020). This is illustrated in Fig. 3.2. As a consequence of this requirement, the primary lens must be focusing rather than diverging, which in turn leads to needing larger distances in order to achieve the same w_0 because the source focal point now lies between the two lenses rather than (virtually) before the diverging lens. There are also practical limitations on how short the primary focal length can be, due to optical aberrations and the fact that the lens and prism mounts restrict how close the prism can be placed relative to the lens.

The pitch-side optics can be modeled as two sub-systems in series, each consisting of a Wollaston prism followed by a focusing lens as shown in Fig. 3.3. Consider only the centerline ray of each beam. A point on the ray can be described by a vector $\mathbf{r}$, composed of its orthogonal distance from the optical axis r and the local angle the ray makes to this axis, ϕ:

$$\mathbf{r} \equiv \begin{bmatrix} r \\ \phi \end{bmatrix} \tag{3.5}$$

The output rays $\mathbf{r_{out}}$ from each sub-system are linearly related to the input rays $\mathbf{r_{in}}$:

$$\mathbf{r_{out}} = \mathbf{M_i} \left(\mathbf{r_{in}} \pm \mathbf{r_{WP}} \right) \tag{3.6}$$

The additive term $\mathbf{r_{WP}}$ models the angular deviation of the Wollaston prisms:

$$\mathbf{r_{WP}} \equiv \begin{bmatrix} 0 \\ \theta/2 \end{bmatrix} \tag{3.7}$$

$\mathbf{M_i}$ models the combined effects of the displacements and the lens within each sub-system ($i \in \{A, B\}$):

$$\mathbf{M_i} \equiv \begin{bmatrix} 1 - \dfrac{d_{2i}}{f_i} & d_{1i} - \dfrac{d_{1i} d_{2i}}{f_i} + d_{2i} \\ -\dfrac{1}{f_i} & -\dfrac{d_{1i}}{f_i} + 1 \end{bmatrix} \tag{3.8}$$

Assuming that the single input ray is co-axial with the optical axis and thus normally-incident on the first prism, then the four output rays of the full system are given by:

$$\mathbf{r_{out}^{a,b}} = \mathbf{M_B} \left((-1)^a \mathbf{M_B} \mathbf{r_{WP_A}} + (-1)^b \mathbf{r_{WP_B}} \right) \ \forall \ a, b \in \{0, 1\} \tag{3.9}$$

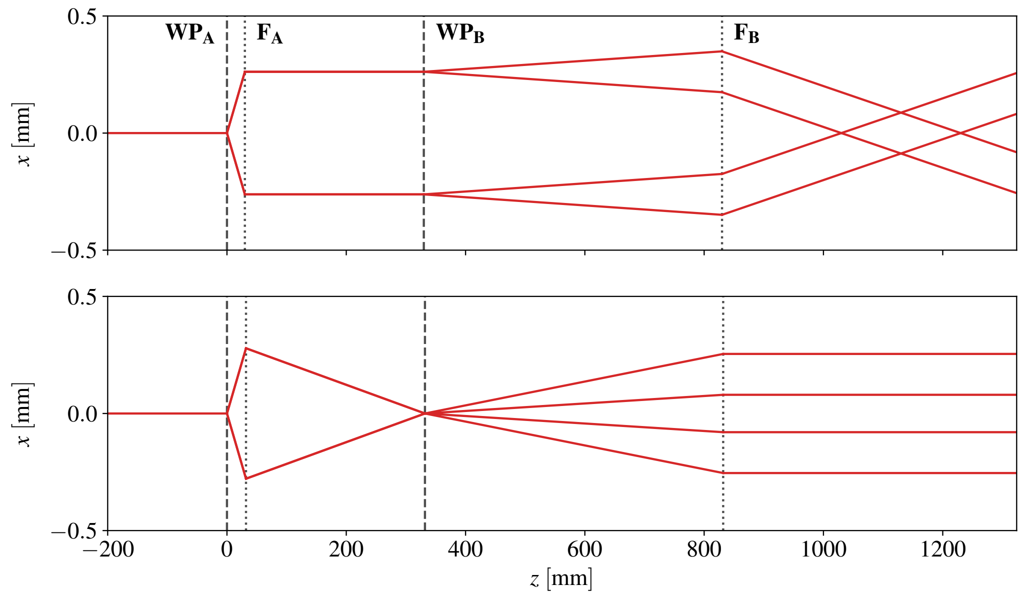

Figure 3.2: DFLDI pitch-side principal rays for incorrect (top) and correct (bottom) primary Wollaston prism (**WP_A**) positioning. In the correct configuration, all four rays are parallel in the test section. Note the transverse (x) scale is exaggerated.

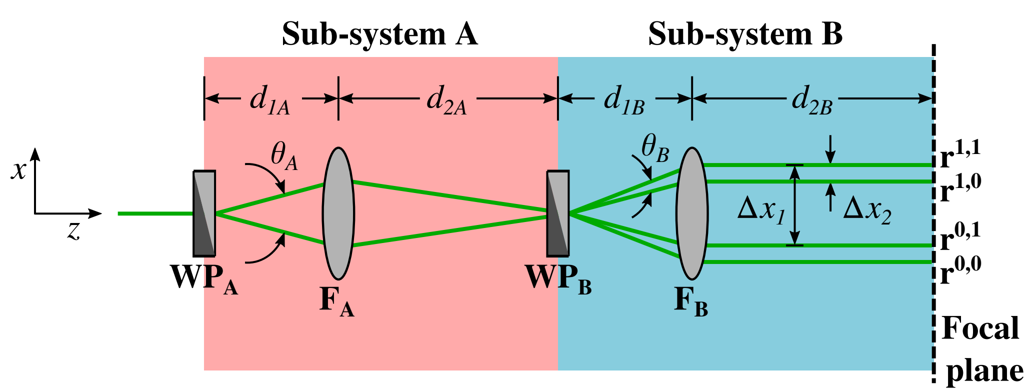

Figure 3.3: Schematic of DFLDI pitch-side optics conceptualized as two optical subsystems for application of ray-tracing matrix analysis.

The requirement that all rays must be parallel with each other and the optical axis in the test section gives $\phi^{a,b} = 0 \;\forall\; a,b \in \{0,1\}$. Due to symmetry, we can simply consider any one of the rays and set $\phi = 0$. Expanding out the relevant terms in Eq. (3.9) yields:

$$\frac{1}{2}\left\{-\frac{1}{f_B}\left(d_{1A} - \frac{d_{1A}d_{2A}}{f_A} + d_{2A}\right)\theta_A + \left(1 - \frac{d_{1B}}{f_B}\right)\left[\left(1 - \frac{d_{1A}}{f_A}\right)\theta_A + \theta_B\right]\right\} = 0 \quad (3.10)$$

Introduce the following non-dimensionalized variables:

$$F \equiv \frac{f_A}{f_B}\;,\; \Theta \equiv \frac{\theta_A}{\theta_B}\;,\; X \equiv \frac{d_{1A}}{f_A}\;,\; Y \equiv \frac{d_{1B}}{f_B}\;,\; K \equiv \frac{d_{1B} + d_{2A}}{f_B} \quad (3.11)$$

With these, Eq. (3.10) reduces to:

$$Y = \Theta\,(K - F - 1)\,X + [1 + \Theta\,(1 - K)] \quad (3.12)$$

It can be shown that the secondary prism can still be at the focal point of the secondary lens, which fixes $Y = 1$. Solutions with $Y \neq 1$ are possible, but this complicates calculations and should be avoided unless experimental constraints make it necessary. With this, the correct position of the primary prism can be determined:

$$d_{1A} = f_A X = f_A \cdot \frac{K - 1}{K - 1 - F} \quad (3.13)$$

Note that this is independent of the properties of either prism. Now that all optical component positions are known, the beam separations Δx_1 and Δx_2 can be computed:

$$\Delta x_1 = \left|r_{out}^{1,1} - r_{out}^{2,1}\right| = \left|r_{out}^{1,2} - r_{out}^{2,2}\right| \quad (3.14a)$$

$$\Delta x_2 = \left|r_{out}^{1,1} - r_{out}^{1,2}\right| = \left|r_{out}^{2,1} - r_{out}^{2,2}\right| \quad (3.14b)$$

Eq. (3.14a) evaluates to:

$$\begin{aligned}\Delta x_1 = \Bigg[&\left(d_{1A} - \frac{d_{1A}d_{2A}}{f_A} + d_{2A}\right)\left(\frac{d_{2B}}{f_B} - 1\right) \\ &-\left(d_{1B} - \frac{d_{1B}d_{2B}}{f_B} + d_{2B}\right)\left(\frac{d_{1A}}{f_A} - 1\right)\Bigg]\theta_A\end{aligned} \quad (3.15)$$

Note that Eq. (3.15) is more general, it holds even if $Y \neq 1$ and $|\phi| \neq 0$. A similar equation can be derived for Δx_2, but for $Y = 1$, the simpler trigonometric derivation as used for the single FLDI (Eq. (3.4)) still holds:

$$\Delta x_2 = 2 f_B \tan\left(\frac{\theta_B}{2}\right) \approx f_B \theta_B \tag{3.16}$$

The optimal relative positions d_{1A}, d_{1B}, d_{2A}, and d_{2B} of the pitch-side optical components are independent of θ_1 and θ_2. This makes the foci separations Δx_1 and Δx_2 (approximately) linearly dependent only on θ_1 and θ_2, respectively. This means, for example, that Δx_1 can be doubled simply by doubling θ_1, without re-positioning any of the components, nor affecting Δx_2, and vice versa.

The above model is still restrained by the same limitations as the single DFLDI design procedure: the paraxial assumption, and idealized Wollaston prism behavior. Nevertheless, it performs satisfactorily in regimes where these approximations are valid. For example, the first DFLDI constructed by the author had $\theta_A = 1°$, $\theta_B = 2'$, $f_A = 30\,\mathrm{mm}$, and $f_B = 300\,\mathrm{mm}$. The predicted values using this scheme were $\Delta x_1 = 334\,\mu\mathrm{m}$ and $\Delta x_2 = 175\,\mu\mathrm{m}$, compared with experimentally-measured values of $364 \pm 2\,\mu\mathrm{m}$ and $184 \pm 2\,\mu\mathrm{m}$, yielding errors of -8% and -5% respectively. A subsequent modification was made: the primary Wollaston prism was increased to $\theta_A = 5°$, with a predicted $\Delta x_1 = 1670\,\mu\mathrm{m}$ and Δx_2 unchanged. The measured values were $1739 \pm 2\,\mu\mathrm{m}$ and $183 \pm 2\,\mu\mathrm{m}$, i.e. the new Δx_1 was within 4% of the prediction, and Δx_2 remained unchanged to within measurement precision. These results are considered accurate enough for most design purposes, although it is still important to get precise values experimentally via beam profiler, especially when Δx_1 is being used for velocity calculations.

3.3 Optical Alignment

The practicalities of aligning all the optical components, and then performing a calibration to ensure the FLDI is behaving correctly, are lacking in the literature. The purpose of this section and the next is to provide details on how this can be done, once the component specifications have been chosen according to the procedures of Section 3.1 or 3.2.

It is helpful to mount each optical component on a translation stage to allow the optics to be traversed perpendicular to the beam, (i.e. in the x-direction). In particular, $\mathbf{WP_2}$ and $\mathbf{P_2}$ need to be mounted together on a single translation stage for the infinite-fringe configuration and calibration processes (discussed later). Furthermore, all Wollaston prisms and polarizers require rotational kinematics. In this particular setup, the system is mounted on two Newport X95-1 rails with associated rail carriers and precision linear translation stages. These rails were chosen due to the high stiffness afforded by their cross-sections, allowing the rails to be partially cantilevered when required by restrictions of the facility test section.

The following is a step-by-step guide to the optical mounting and alignment procedure, for a system where all components use standard optical posts and post-holders.

I **Rails**
The most desirable arrangement is a single, continuous, stiff optical rail spanning the full length of the FLDI system. This is usually not possible, due to available rails not being long enough, or obstruction by the facility test section. A normal system will have two rails, one pitch and one catch, whose axes should be made precisely parallel to one another. It is possible to "bend" the system using additional mirrors and rails if there is not enough space. It is also preferable to have both rails mounted to the same isolated optical table, because independent mounting can lead to relative motion between the rails, eliciting a spurious FLDI response.

II **Laser & Periscope**
The laser should be placed at the pitch end of the optical setup, with a periscope assembly to allow for precise directional control of the beam. First, the desired final beam height should be chosen. Align the beam parallel to the rails at this target height, and leave a target marker in place at the far end of the rails for use in subsequent steps.

III **Focusing Lenses**
Starting with the transmitting side of the system, place $\mathbf{F_1}$ such that the beam travels through its center. When this is achieved, the beam should strike the same target point that was set in Step 1 without the lens in place. Next, make the lens orthogonal to the optical path by aligning its back-reflection with the incident beam path. The lens should be initially positioned such that when it is twisted into alignment about its vertical axis, it tightens on the optical post thread (i.e. clockwise viewed from above). Repeat the process for $\mathbf{F_2}$ on the receiving side of the system. If the laser was well-aligned with the rails, the lenses can be initially located and aligned at arbitrary z-positions then

subsequently translated to their final positions without disturbing the beam direction.

IV **Diverging Lens**

D should be aligned in a similar manner to $\mathbf{F_1}$ and $\mathbf{F_2}$, ensuring that the beam is passing through the center of the lens and aligning the back-reflection. It is helpful to have **D** on its own stage as it needs to be temporarily removed from the system in subsequent steps. Construct a position reference on the rail to quickly re-locate it at the correct z-position.

V **Prisms & Polarizers**

For this step it is very helpful to have a beam profiler or suitable camera that can be used to observe the beam focus. Without this, it is very difficult to accurately measure the foci sizes and separation, or to notice any aberrations (such as those discussed regarding polarizer selection in Section 3.1). Remove **D** in order to have a narrow beam for alignment, then place a prism/polarizer pair on the transmitting side, such that $\mathbf{WP_1}$ is separated from $\mathbf{F_1}$ by a distance f_1. The exact offset between the prism and polarizer is not important so long as apertures are not exceeded. Align each component with the beam in the same way as the lenses.

Place **D** back into the system, and position the beam profiler near the system focal point. It may be necessary to temporarily introduce a neutral density filter or beam splitter to reduce the beam power to protect the profiler. If the transmitting-side optics are all in the correct positions, two foci will be observed. Adjust the beam profiler position to find the plane where the foci size is minimized; this should be equidistant between $\mathbf{F_1}$ and $\mathbf{F_2}$. The foci should have a separation Δx and Gaussian beam radius w_0 close to the nominal design values. With initial arbitrary rotational positions of $\mathbf{WP_1}$ and $\mathbf{P_1}$, the foci will not necessarily be of the same intensity, nor horizontally-oriented. Use the rotational mounts to control each of these aspects:

a **Foci Intensity**

This is controlled by polarizer rotation. Rotating through 360°, each foci will undergo a full cycle of intensity. At the extrema, one focus will not be visible, and the other will be at maximum intensity. Lock the polarizer rotation at the midpoint where both foci are of equal intensity. The beam profiler software used in this setup indicates the point of maximum intensity and the intensity centroid. At the desired configuration, the former should be jumping between the two foci, and the latter should be at the midpoint.

b **Foci Orientation**

This is controlled by prism rotation. Typically, orientation in the streamwise direction is desired.

Once this pair is aligned, the process should be repeated for $\mathbf{WP_2}$ and $\mathbf{P_2}$. Aligning both pairs in turn on the transmitting side of the system allows the use of the beam profiler to ensure that both pairs are behaving as expected; this makes the final recombination easier. As with the diverging lens, an extra rail carrier should be used for repeatable z-location after removal. Once both pairs have been aligned, they can be fixed into the system. If $\mathbf{P_1}$ and $\mathbf{P_2}$ are both linear polarizers, the pairs are interchangeable and it does not matter which is placed on the transmitting or receiving sides. If a quarter-wave plate is substituted for one of the linear polarizers, as in this setup, it must be used for $\mathbf{P_1}$ in order to achieve the infinite-fringe configuration.

VI **Infinite-Fringe Configuration**

At this point all optical components, excluding the photodetector, should be aligned and in place with their design separations. Now the FLDI system can be put into the infinite-fringe configuration. This is done by adjusting the x- and z-positions of the $\mathbf{WP_2}$–$\mathbf{P_2}$ pair. As previously noted, this pair must be mounted on a single translation stage. Place a screen down-beam of the pair at a location where the projected overlapping beams are large enough for easy viewing.

Fig. 3.4 shows the fringe patterns obtained by translating $\mathbf{WP_2}$–$\mathbf{P_2}$ along the beam path. When the pair are too close to $\mathbf{F_2}$, many straight, horizontal fringes are visible on the screen (Fig. 3.4, top-left). Translation across the beam path will cause the fringes to “scroll” across the screen, but the number of fringes will remain constant. As the pair is translated in z away from $\mathbf{F_2}$ the number of visible fringes decreases, (Fig. 3.4, upper row). Close to the infinite-fringe location the fringes will start to become curved and eventually only a single fringe will be visible (Fig. 3.5). Translating the pair across the beam path now will cause the single fringe to translate across the screen, however, regions of both light and dark will always be visible. At the correct z-location, this single fringe will expand to fill the entire projected beam, and x-translation will cause uniform brightening or darkening. If the pair is translated in z beyond the infinite-fringe location, the fringes straighten and the number of visible fringes increases again, (Fig. 3.4, lower row). The infinite-fringe location should be close to the nominal design position for the prism, i.e. at the focal length of $\mathbf{F_2}$.

Purely translating $\mathbf{WP_2}$–$\mathbf{P_2}$ along the beam path is not always enough to achieve the infinite-fringe configuration. Often the $\mathbf{WP_2}$ needs to be rotated from its alignment configuration, in very small increments. Fig. 3.6 shows the effects of this rotation. These fringe images were obtained by first putting the system into the infinite-fringe z-position, then rotating the prism up to $\pm 45°$ from alignment. The effect of this rotation on the fringes is similar to the effect of translation discussed above: more fringes become visible further from alignment.

VII **Photodetector**

The final step in aligning the FLDI system is to replace the viewing screen with the photodetector. It should be placed such that the sensor area is almost completely filled by the beam. This is usually slightly in front of the beam focal point to minimize the overall length of the system, but can also be behind the focal point. Because beam recombination and interference occur at $\mathbf{WP_2}$ and $\mathbf{P_2}$ respectively, mirrors can be used anywhere down-beam of the pair to redirect the light to the photodetector if space is limited.

3.4 Calibration & Post-Processing

Once the infinite-fringe configuration has been achieved as described in Section 3.3, the FLDI is functional. Introduction of a density disturbance, e.g. a jet sprayed from an air-duster near the foci, will yield a fluctuating voltage signal measured by the photodetector. For some qualitative work it may be possible to immediately employ the FLDI at this stage, but in most cases a calibration must be performed. This refers to establishing the relationship between the raw voltage signal V, and the FLDI phase response $\Delta\Phi$, as defined by Schmidt and Shepherd (2015). It will be shown in Section 4.2 that the function

$$V = A \sin(\Delta\Phi - \Delta\Phi_0) + D \tag{3.17}$$

produces quantitative agreement with the theoretical FLDI response. Note it is not important whether sin or cos is used, so long as a consistent convention is followed. The following calibration procedures establish the constant values of A, D, and $\Delta\Phi_0$. Translating the second prism-polarizer pair in the x-direction will trace out a sinusoid in V.

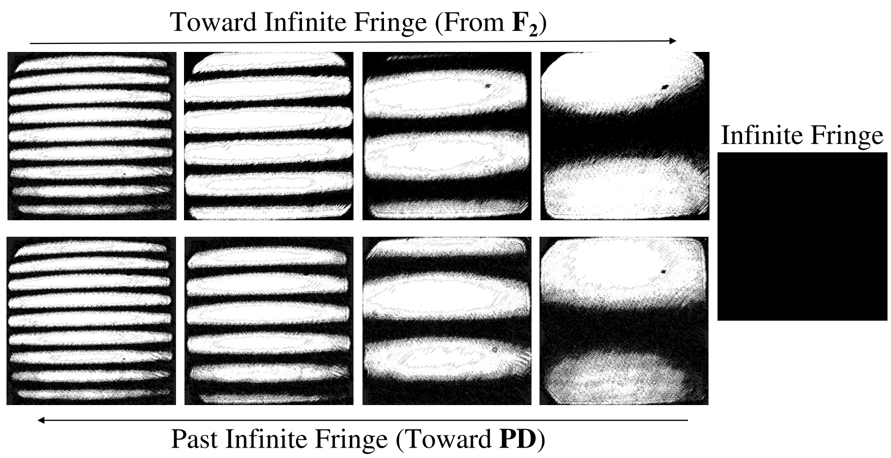

Figure 3.4: Fringe pattern obtained by translating receiving prism/polarizer pair along beam path (z-direction). Figs. 3.4 to 3.6 provided by M. C. Neet and used with permission.

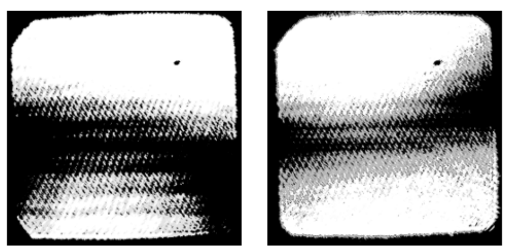

Figure 3.5: Curving of single fringe near the infinite fringe alignment location.

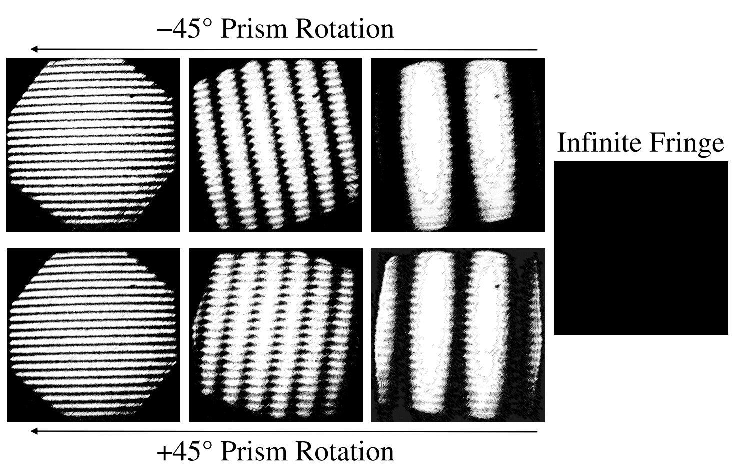

Figure 3.6: Effect of rotating Wollaston prism on fringes.

3.4.1 Manual Method

The simplest method is to perturb the system in this way until both the minimum and maximum values of V are found, and then compute $A = (V_{max} - V_{min})/2$ and $D = (V_{max} + V_{min})/2$. $\Delta\Phi_0$ is determined by the final value of V_0, i.e. during calibration there is no flow disturbance so $\Delta\Phi = 0$ and $V_0 = f(\Delta\Phi_0)$. Although this method does not rely on linearizing the sinusoid as done by other authors (Parziale et al., 2013b; Schmidt and Shepherd, 2015), V_0 is usually chosen to be near the middle of the range to give sensitivity to phase changes of either sign while avoiding the phase ambiguity that occurs when an extrema is crossed.

3.4.2 Automated Method

The method of calibration above simply moves the $\mathbf{WP_2}$–$\mathbf{P_2}$ translation stage manually. This relies on fitting a sinusoid based on only the voltage extreme, which is prone to error, e.g. due to voltage noise fluctuations. A more accurate and consistent method is to use a mechanized translation stage for these optics. If the stage is made to translate at a constant small velocity, then locally about the infinite fringe position this is equivalent to a constant rate of change of $\Delta\Phi_0$. This procedure traces out a smooth sinusoid on the oscilloscope, then curve-fitting tools can be used to fit a sine function to the whole calibration dataset, rather than relying on only a few extreme points. Additionally, the "home" position can be set so that after the calibration translation, the exact midpoint of the fringe can be returned to with high repeatability. This method was employed during the HET campaign using a Thorlabs MT1-Z8 translation stage and a custom driver script that allowed for one-touch calibration before every experiment.

3.4.3 Coupling

In applications where FLDI is being used to measure small-scale density fluctuations (e.g. a turbulent spectrum) it is often beneficial to acquire V in an AC-coupled mode in order to maximize bit depth. In this case, calibration should still be performed in DC-coupled mode, with V_0 being added back to V_{AC} in post-processing before inverting Eq. (3.17) to obtain $\Delta\Phi(t)$. Alternatively, the signal can be split into AC- and DC-coupled channels, with care being taken to keep terminating impedances matched to avoid reducing the signal bandwidth.

Chapter 4

RESULTS: FLDI MODELING & VALIDATION

This chapter addresses the modeling of the FLDI response as a function of the flow field being measured, and is largely comprised of material adapted and combined from two publications: Lawson et al. (2020) and Lawson and Austin (2021). The underlying analytical framework is introduced, along with a numerical implementation that allows for modeling of arbitrarily complex flow fields. The model is validated via three experiments that test both the static and dynamic response of FLDI. An important aspect of FLDI explored in all three of these experiments is the spatial extent of the sensitive region, and how this depends on the length scales and functional forms of the refractive index field.

The FLDI system characterized here is a refinement of the original instrument of Parziale et al. (2012), and was designed by I. J. Grossman. The light source is a 200 mW, 532 nm laser [Spectra-Physics EXLSR-532-200-CDRH] and the detector is a switchable-gain detector with 12 MHz bandwidth [Thorlabs PDA36A2]. All studies in this chapter were performed with a nominally-identical FLDI, although it was re-built for various experiments (e.g. tabletop work versus application to HET). The Gaussian diameters and separation of the foci were measured using a beam profiler [Thorlabs BP209-VIS].

The work in Section 4.2 is from Lawson et al. (2020), which gave $d_0 = 7 \pm 2\,\mu\text{m}$ and $\Delta x = 180 \pm 2\,\mu\text{m}$, while Sections 4.3 and 4.4 are from Lawson and Austin (2021) and uses $w_0 = 4 \pm 2\,\mu\text{m}$ and $\Delta x = 184 \pm 2\,\mu\text{m}$. The interchanging use of d_0 vs. $w_0 = d_0 \div 2$ is preserved here as per each parent work.

4.1 Modeling & Simulation

4.1.1 Analytical Model

Schmidt and Shepherd (2015) developed a ray-tracing theoretical model of FLDI response, along with a numerical scheme to allow simulations of the instrument. The two FLDI beams are modeled as parallel, focusing ray bundles, offset by a constant small distance Δx. A Cartesian co-ordinate system $\underline{x} = (x, y, z)$ is established with the origin in the focal plane, equidistant from the two foci, as shown in Fig. 4.1. Each ray in one beam is paired with a ray in the other beam, and is governed by the

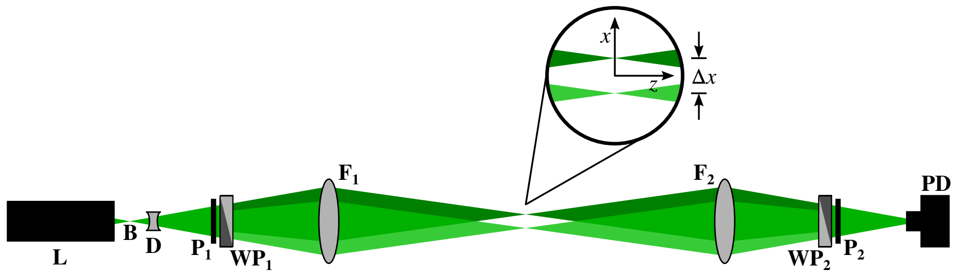

Figure 4.1: FLDI schematic repeated from Fig. 1.1 as a reminder of the co-ordinate system used throughout this chapter.

paraxial approximation. Conceptually, this ray pair originates from a single location at the source; rays are then separated a distance Δx by the pitch-side optics. The catch-side optics then recombine the ray pair at the detector, yielding an intensity contribution that varies based on the relative phase offset between the rays. $\Delta\Phi$ is the phase change integrated across the detector:

$$\Delta\Phi(t) = \frac{2\pi}{\lambda_L} \iint_D I_0 \sin\left[\frac{2\pi}{\lambda_L}\left(\int_S^D n(\underline{x}_1, t)\mathrm{d}s_1 - \int_S^D n(\underline{x}_2, t)\mathrm{d}s_2\right)\right]\mathrm{d}\xi\mathrm{d}\eta \tag{4.1}$$

where $n(\underline{x}_i, t)$ is the 3D refractive index field interrogated by the i^{th} beam, s_i are the parameterizations of the ray paths, and the source and detector planes, S and D, have the co-ordinate system (ξ, η). The assumption is made that ray deflection is negligible and thus these ray paths are constant and predetermined. The inner term in parentheses represents the optical path difference (OPD) between the rays in a pair. The contributions of all ray pairs are then integrated over the detector face. This model assumes a coherent light source of wavelength λ_L, with normalized intensity distribution $I_0(\xi, \eta)$. Invoking the small-angle approximation allows Eq. (4.1) to be linearized, yielding Eq. (4.2).

$$\Delta\Phi(t) = \frac{2\pi}{\lambda_L} \iint_D I_0 \left(\int_S^D n(\underline{x}_1, t)\mathrm{d}s_1 - \int_S^D n(\underline{x}_2, t)\mathrm{d}s_2\right)\mathrm{d}\xi\mathrm{d}\eta \tag{4.2}$$

To compute the temporal response $\Delta\Phi(t)$ for an unsteady flow field, a quasi-steady approximation must be made, where it is assumed the optical time-of-flight is short compared to flow timescales. An order-of-magnitude analysis based on the FLDI

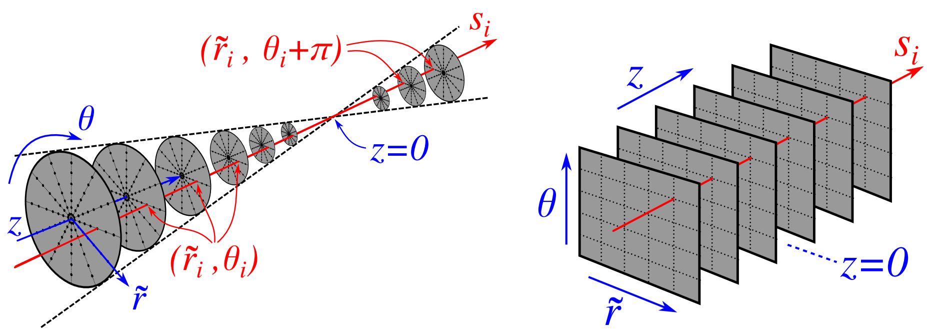

(a) Spatial domain.

(b) Computational domain.

Figure 4.2: Schematic of focused beam discretization using (r, θ, z) co-ordinates. A given ray s_i connects all points with the same $(\tilde{r}_i, \theta_i)$, where $\tilde{r} \equiv r/w(z)$. Upon passing through the focus at $z = 0$ all rays undergo image inversion such that $\theta_i \mapsto \theta_i + \pi$. In the computational domain, this inversion is accounted for during the integration step.

spatial resolution and dimensions of typical wind tunnel facilities shows that flow features would need velocities $u \sim \mathcal{O}\left(10^4\text{–}10^5\ \mathrm{m\,s^{-1}}\right)$ for this assumption to break down.

The default optical parameters used in calculations and simulations throughout this work are $\lambda_L = 532.6\,\mathrm{nm}$, $\Delta x = 180\,\mu\mathrm{m}$, and the Gaussian radius of the foci, $w_0 = 5\,\mu\mathrm{m}$. These are representative of the experimental parameters.

4.1.2 Numerical Discretization

The optics used to separate and re-combine the beams are not modeled. Instead, only the portion of the optical path between the two focusing lenses is simulated, with a source plane after the pitch-side lens and a detector plane before the catch-side lens. Note that Eq. (4.2) does not define a beam geometry, and is in fact descriptive of a much more general class of interferometers than just FLDI. Rather, the information about beam geometry is given by the function defining the ray paths s_i in Cartesian space $\underline{x}_i$. Here, the beams obey Gaussian optics, and their geometry can be entirely described *a priori* by Δx, λ_L, and w_0.

The beams are discretized into a series of circular slices along the z-axis (the direction of optical propagation). Each z-slice is itself discretized using a universal grid in polar co-ordinates $(\tilde{r}, \theta)$, where $\tilde{r} \equiv r/w(z)$ is the radius normalized by the local Gaussian beam radius. A ray is then simply the locus of all points with constant $(\tilde{r}, \theta)$, as illustrated in Fig. 4.2a. A refractive index field is sampled at each beam grid point, then numerical integrations are performed pairwise along the rays, and finally over the detector plane. This is equivalent to recasting Eq. (4.2) as:

$$\Delta\Phi = \frac{2\pi}{\lambda_L} \int_0^{2\pi} \int_0^{\infty} \tilde{I}_0 \left(\tilde{r}\right) \left(\int_{z_a}^{z_b} n_1 \mathrm{d}z - \int_{z_a}^{z_b} n_2 \mathrm{d}z \right) \tilde{r} \mathrm{d}\tilde{r} \mathrm{d}\theta \tag{4.3}$$

Note the shorthand $n_i = n(\underline{x}_i)$ is employed, and the time-dependence is left implicit. For a Gaussian beam, the normalized intensity profile is given by:

$$\tilde{I}_0 \left(\tilde{r}\right) = \frac{2}{\pi} \exp\left(-2\tilde{r}^2\right) \tag{4.4}$$

The Gaussian beam radius is:

$$w(z) = \sqrt{w_0^2 \left(1 + \left[\frac{\lambda_L z}{\pi w_0^2}\right]^2\right)} \tag{4.5}$$

The beams extend over the range $[z_a, z_b]$, and the co-ordinate transforms are:

$$x_1 = \tilde{r} w \cos\theta + \Delta x/2 \tag{4.6a}$$

$$x_2 = \tilde{r} w \cos\theta - \Delta x/2 \tag{4.6b}$$

$$y = \tilde{r} w \sin\theta \tag{4.6c}$$

The numerical integration in $\tilde{r}$ requires finite truncation of the upper limit, which is done here at $\tilde{r} = 2$. For further discussion on how this truncation is chosen, along with other details of the discretization scheme, such as the construction and validation of the grids, please refer to Schmidt and Shepherd (2015).

4.1.3 Implementation

The above discretization scheme was implemented in Python. The array structures provided by NumPy (Harris et al., 2020) allow for optimized array operations and much more readable code than looping over the grid points, at the cost of memory overhead. The beams are represented by 3D meshgrids R, Θ, and Z, each of shape (n_r, n_θ, n_z); along each axis, $\tilde{r}$, θ, and z are held constant, respectively. In order to interact with flow-field co-ordinate systems, which are most often defined as

Cartesian, one can either transform the flow-field to normalized polar co-ordinates and integrate using Eq. (4.3), or transform the beam grids to Cartesian co-ordinates and use Eq. (4.2). The former approach turns out to be very useful for deriving analytical expressions for $\Delta\Phi$, but this is outside the scope of this work. For the numerical implementation we use the latter approach. The beam grids R and Θ are transformed to Cartesian co-ordinate arrays $X(R, \Theta)$ and $Y(R, \Theta)$.

$\{X, Y, Z\}$ are used with the Pythagorean theorem to compute the ray trajectory array Σ, which stores the ray path length at every beam grid point, measured from the source plane. Because the two FLDI beams are identical in all ways except their x-offset, only one set of these arrays needs to be computed and stored. To account for this offset when computing the refractive index $n(\underline{x})$ on each grid, the two beams are passed to the assigning function as $\{X \pm \Delta x/2, Y, Z\}$. This results in different refractive index arrays for each beam, N_+ and N_-. To implement the inner line integrals of Eq. (4.2), Simpson's method is performed to integrate along the ray paths s_i using the integrand $N_\pm \cdot \Sigma$. This yields two arrays $J_\pm$ of shape (n_r, n_θ); $J_+ - J_-$ represents the OPD at each point in the detector plane. Simpson's method is invoked again, employing R and Θ to perform the outer double integral of Eq. (4.2), where the generic co-ordinate system (ξ, η) is chosen to be (r, θ). This returns the scalar FLDI response $\Delta\Phi$.

The functions that assign refractive index to the beam grids for each time instant are separate from the core routines, and only interface with the beam object via a universal wrapper function. This means the beams are constructed and integrated agnostic to how the flow geometry is defined. Several methods of defining $n(\underline{x})$ are available, including analytical functions (e.g. Section 4.3), HDF files (e.g. Section 4.4), and experimental datasets (e.g. Section 4.2). Generally, the flow is first defined or computed as a density field, which is converted to refractive index using the Lorenz-Lorentz equation (Eq. (1.2)) with constants taken from Gardiner et al. (1981). If an analytical function $n(\underline{x})$ is available, then this can just be invoked at each beam grid point. For interpolation of numerical or experimental datasets (which are usually Cartesian) onto the beam grids, $n(\underline{x})$ would need to be computed or measured within the Cartesian bounding box of the beams. The most conservative z-extent for this box ranges from the pitch-side focusing lens to the catch-side focusing lens, although this can be truncated to exclude portions of the beam where it is known that the $n(\underline{x})$ field is uniform and steady.

Besides the improvements offered by array operations, the current code includes some other changes from the original implementation of Schmidt and Shepherd. Firstly, the z-grid does not need to be linearly spaced: arbitrary distributions of grid-points are supported, allowing for mesh refinement close to the focal plane where $w(z)$ can be an order of magnitude smaller than the constant $\delta z = 100\,\mu\text{m}$ originally prescribed.

A subtler change is image inversion: the above definition of a ray as being the locus of constant $(\tilde{r}, \theta)$ is not physical, as it implies a ray originating in a given quadrant of the source will terminate in the same quadrant at the detector. In reality, the optical image must invert through the focus, so a ray terminates in the diagonally-opposite quadrant. More precisely, in this beam geometry, a ray through all $(\tilde{r}, \theta)$ before the focal plane must pass through all $(\tilde{r}, \theta + \pi)$ after the focal plane. For many classes of disturbance geometry, factors such as symmetry, localization of the disturbance to one side of the focal plane, or the very small beam widths near the focus mean that this inversion does not noticeably affect the instrument response. However, for one class of problem, the clipping of beams by solid objects, inversion is a key effect (see Appendix C). The most efficient way to implement optical inversion without disrupting the numerical framework is as follows: compute $N_{\pm} = f(X \pm \Delta x/2, Y, Z)$ as before, then for all slices of $N_{\pm}$ for which $z < 0$ (pitch side of focal plane) perform an array cycling operation along the θ-axis such that a row originally at $\theta = \theta_0$ is switched with the row at $\theta = \theta_0 + \pi$. No other arrays need to be altered. Now when the ray integrations are performed, the correct values of $n(\underline{x}_i)$ along the physically-accurate ray are used. The Gaussian optics assumption means that the beams have finite radius even at the focus, so this introduces a local discontinuity in connectivity in $\mathrm{d}s_i$, however with appropriate choice of δz this introduces negligible error. The only other restriction introduced by optical inversion is the requirement that n_θ needs to be even.

4.1.4 Grid Resolution

The original implementation of Schmidt and Shepherd defined the beam grid resolution using two parameters: the number of points in the $\hat{\theta}$-direction, n_θ, and the spacing of the z-slices, δz. The number and distribution of points in the $\hat{r}$-direction is a function of n_θ. They performed a convergence study using a 2D sinusoidal disturbance, with an analytically-determined solution as a convergence target, and found that $n_\theta = 300$ and $\delta z = 100\,\mu\text{m}$ were adequate for convergence.

In this work, the following grid resolution was chosen: $n_\theta = 256$ and $\delta z = 250\,\mu\text{m}$, and is used for all instrument simulations herein. As mentioned, this new implementation allows for non-uniform z-resolution. Thus, near the foci ($|z| < 20\,\text{mm}$), the grid is refined by a factor of 5, giving $\delta z = 50\,\mu\text{m}$. These parameters yield $(n_r, n_\theta, n_z) = (311, 256, 1441)$, giving 79,616 grid points per z-slice, and almost 115 million total grid points per beam.

A grid resolution study was performed using a propagating shock representative of those simulated in this work (Fig. 4.3). Unlike the sinusoidal disturbance used by Schmidt and Shepherd, there is no known analytical response function to compare with. The chosen parameters are deemed to give sufficiently precise results relative to the experimental uncertainties, using the root mean square (RMS) of the residual error as a metric (Fig. 4.4).

4.2 Characterization & Validation

In all previous experimental studies by other authors, the phase object (e.g. a gas jet) used to probe the FLDI response was either quantitatively uncharacterized, or indirectly characterized via simulation. The spatial positioning of the foci with respect to this phase object was generally not known to a precision comparable with the FLDI characteristic dimensions (typically the foci spacing is on the order of hundreds of µm, and the foci diameters themselves are tens of µm). Hence, this study aims to more fully and directly characterize the phase object by simultaneously making an independent measurement of the refractive index field using a Mach-Zehnder interferometer (MZI). The objectives of this characterization study are:

1. Experimentally obtain the response of an FLDI system to a known static refractive index field, and use these data to validate Schmidt and Shepherd's computational model. A steady laminar jet is used for this objective.

2. Experimentally obtain the response of an FLDI system to a dynamic refractive index field of known frequency, and use these data to validate Schmidt and Shepherd's predictions for wavenumber dependence based on their analytical transfer functions. An ultrasonic acoustic beam is used for this objective.

3. Use the validated model to explore the FLDI response as a function of the optical system parameters, such as the foci size and spacing. The purpose of this is to better understand how accurately these parameters need to be known, as well as informing the design of new FLDI systems for a particular application.

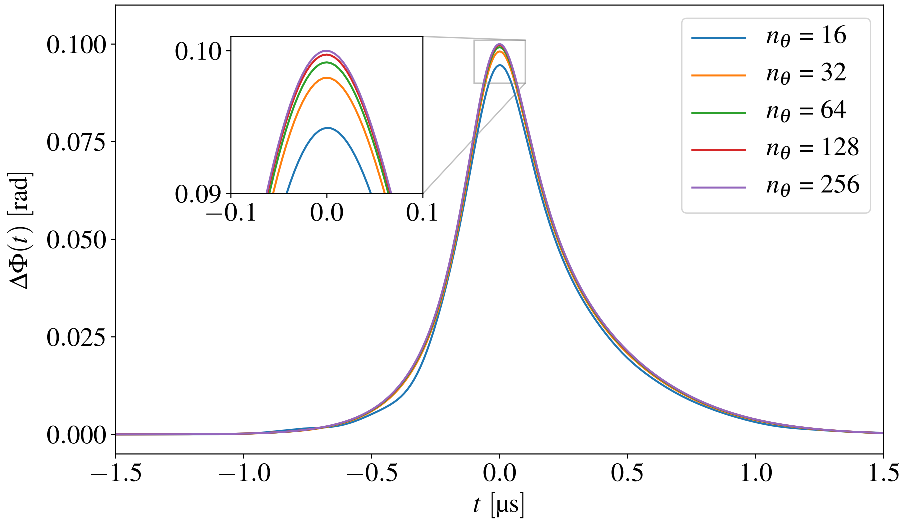

(a) Increasing n_θ with constant $\delta z = 250\,\mu\text{m}$.

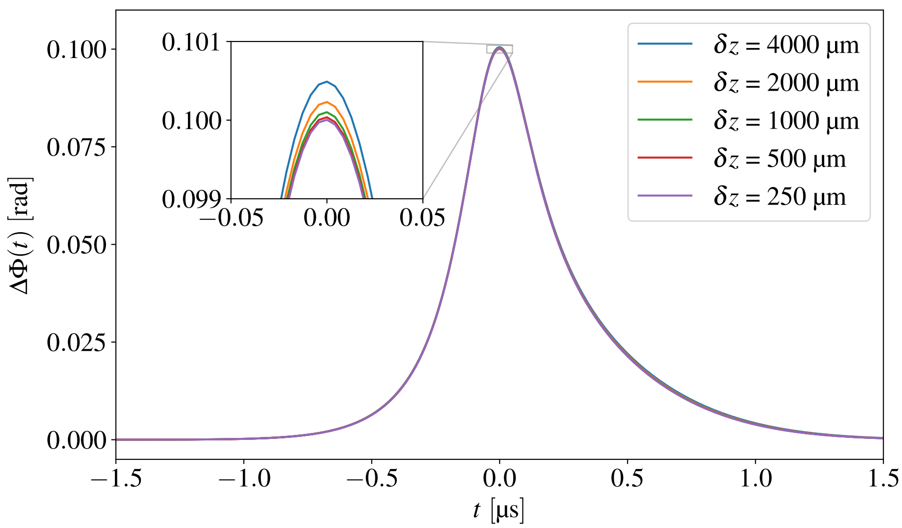

(b) Decreasing δz with constant $n_\theta = 256$.

Figure 4.3: Grid resolution studies for n_θ and δz. In all cases, δz is refined by a factor of 5 for $|z| < 20\,\text{mm}$.

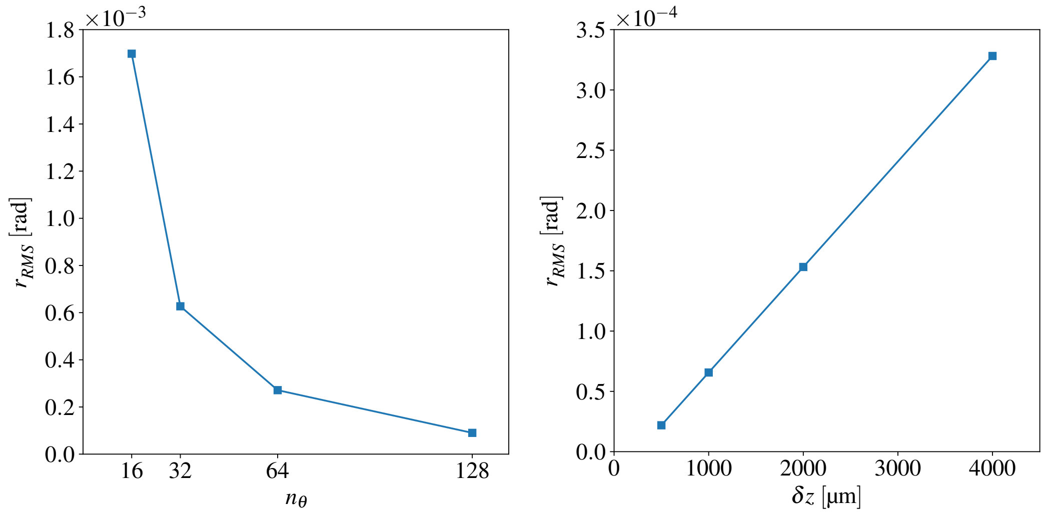

(a) n_θ (b) δz

Figure 4.4: RMS residual error trends for grid resolution, computed for cases in Fig. 4.3. Highest-resolution case used as reference.

4.2.1 Optical Design

The MZI uses the same type of laser as the FLDI system. The measurement beam passes through the measurement volume, intersecting orthogonally with the FLDI foci (Fig. 4.5), with a field of view 7.7 mm wide by 4.8 mm high. The reference beam is routed around the measurement volume and remains unperturbed. For this work, the MZI was set up with horizontal fringes with a density of 10.8 fringes/mm, yielding about 15 px/fringe. This fringe density is a compromise between the spatial resolution of the refractive index field, and the discretization artifacts that arise from too few pixels per fringe.

When the measurement beam is perturbed by a phase object the horizontal fringes are deflected. The response of a MZI has been extensively studied (Merzkirch, 1987), and the refractive index field can be calculated from fringe deflection using the post-processing tools from Coronel et al. (2018). In the Supplementary Material of that work, an error estimation is performed, using a synthetic refractive index field, corrupted with artificial noise. Even with 10% synthetic noise, the error within their thermal boundary layer (similar in geometry to our laminar jet) remained below 2% except very close to solid surfaces. Our refractive index gradients (and therefore fringe distortion) are considerably less severe than their simulated case, so we expect our uncertainties to also be below 2%, except possibly very close to the axis of rotation where the Abel inversion performs less well.

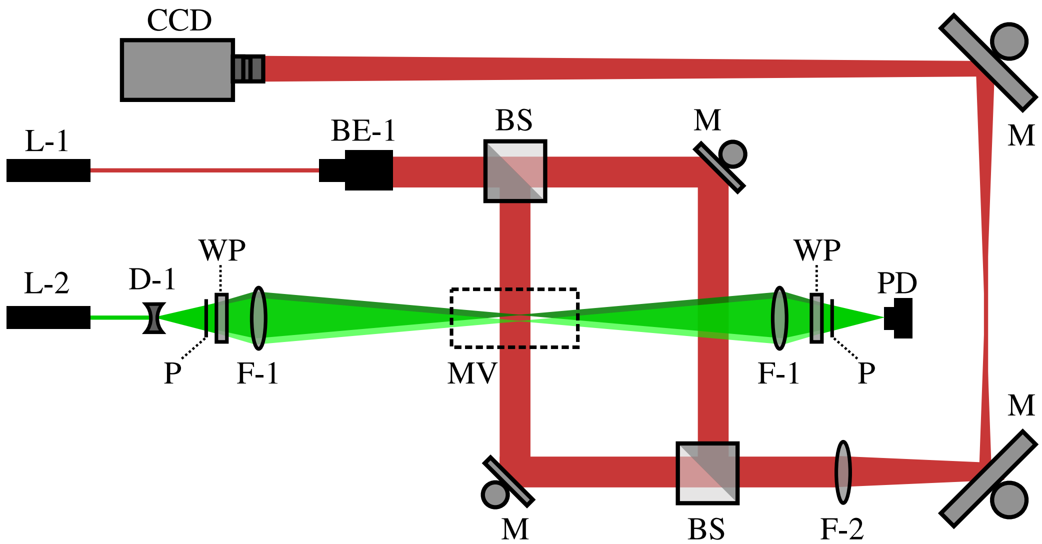

Figure 4.5: Schematic diagram of FLDI and MZI simultaneously observing the measurement volume (MV). FLDI beams shown in green. MZI beams shown in red for clarity. Annotations: L = laser, BE = beam expander, BS = beam-splitter, M = flat mirror, F = focusing lens, D = diverging lens, P = polarizer, WP = Wollaston prism, PD = photodetector, CCD = charge-coupled device (camera).

4.2.2 Laminar Jet

A jet was located in the measurement volume with its vertical axis perpendicular to the optical table, i.e. in the y-direction of the previously-defined FLDI co-ordinate system. A round jet was chosen as the phase object because the index of refraction of the axisymmetric flow field can be calculated via Abel inversion from the Mach-Zehnder interferograms. Additionally, this axisymmetry allows the MZI to observe the jet orthogonally to the FLDI beam direction (see Fig. 4.5). Helium was chosen as the jet gas because its refractive index is sufficiently different enough from air so that the MZI and FLDI results are of high quality. Other common gases have similar Δn with respect to air (e.g. CO_2), but He is lighter than air so the plume will not affect the measurement volume by returning via a "geyser" effect. Steady flow was achieved by controlling the He flow-rate using a two-stage pressure regulator and a rotameter. The jet was kept laminar via flow-straightening features in the plenum, followed by a contoured nozzle, designed and constructed* for this purpose, of outlet diameter 0.5 mm. The flow-rate was approximately 36 $\mathrm{cm^3\,s^{-1}}$, which corresponds

*By M. C. Neet, who was also involved in the construction of both interferometers, and helped conduct all experiments that used the laminar jet.

to a Reynolds number of 390 based on this diameter. As a check, the reference beam of the MZI was blocked, and a knife-edge placed near the focus located beyond focusing lens F-2 (Fig. 4.5). This turned the MZI into a schlieren system, which was used to directly observe the jet and validate that it was steady and laminar.

The laminar jet was mounted on motorized translation stages [Thorlabs MT1-Z8]. One stage was oriented to translate along the x-axis and the other stage orthogonally along the z-axis. These stages have a minimum bidirectional repeatability of 1.5 μm. The jet was translated 12 mm across the beams with 0.1 mm resolution and 37 mm along the beams with 1 mm resolution. Overlapping automated traverse patterns were executed in this plane, covering the entire domain multiple times so that several independent measurements were made at every (x, z) co-ordinate, with each location being visited between 5 to 17 times. FLDI data were acquired at each position for 1 second [Yokogawa DL850] at between 1–10 kS/s. This was then used to calculate an average voltage V at each location. V was then converted to the integrated phase difference, $\Delta\Phi$, using Eq. (3.17). The constants in this equation are determined using the calibration procedures described in Section 3.4.

Referring to Fig. 4.5, both interferometers are stationary with respect to the optical table. The test phase object (in this case, the laminar jet) moves. This means that the fringe image of the jet moves in the MZI field-of-view (FoV). The magnification of this FoV was chosen such that the jet diameter was at least several fringe widths. The jet traverses beyond the FoV in the z-direction, which means the jet was not visible to MZI at every test location. However, because the jet flow is nominally steady, a sufficient number of interferograms were obtained while the jet was visible, and these were used to build up a single composite image of the jet refractive index field. The symmetries of the jet and FLDI response were used to determine the location of the FLDI origin relative to the jet axis. A similar approach was taken by Schmidt and Shepherd (2015) in their experimental jet validation.

These two datasets, from the FLDI and MZI measurements are used to investigate the model of Eq. (4.2) for FLDI, using the new implementation of Schmidt and Shepherd's algorithm detailed in Section 4.1. The experimentally-measured values of Δx and d_0 are used to define the beam geometry. The jet origin is then fixed at a constant negative value of y_{jet}, i.e. below the FLDI foci, as per the experiment. The axisymmetric rotation of the MZI 2D Δn field about the vertical axis through the jet origin generates a cylindrical region in the domain, where jet data is available. Outside of this region, it is assumed that $\Delta n = 0$, i.e. that the far-field refractive index is attained by the edge of the MZI FoV. For every discretization point of the beams that lies within this cylindrical measurement volume, rectangular bivariate spline interpolation is used to calculate the refractive index value at that point. The integrations of Eq. (4.2) are then performed to yield a simulated FLDI $\Delta\Phi$ value from the experimental MZI data. $\Delta\Phi$ is computed in this way at every (x, z) jet location used in the experimental campaign.

4.2.3 Free Ultrasonic Acoustic Beam

A commercial electrostatic ultrasonic transducer [Senscomp Series 600], resonant frequency 50 kHz, was used to generate a dynamic phase object in the FLDI measurement volume. This particular transducer was chosen because it is designed for operation in air, and it is a broadband device which can be operated at frequencies quite far from its resonance frequency while still emitting significant acoustic power. The transducer was oriented with its surface orthogonal to the x-axis, such that a propagating wavefront would reach one FLDI beam before the other, thereby generating a response. It was driven using a continuous sinusoidal signal with a single frequency component. The manufacturer's data show a moderately flat transmission response above resonance out to 100 kHz, though with more rapid drop-off below resonance; in this study frequencies between 30–100 kHz are used. Per the manufacturer specifications, at 50 kHz the beam pattern is multi-lobed, with the majority of the acoustic power within $\pm 15°$ of the axis. It is not known how this pattern changes with frequency, and unfortunately, the density perturbations induced by the transducer are much too weak to be visualized by the MZI.

The transducer face was located approximately 60 mm from the FLDI optical axis, with the acoustic axis close to $y = 0$. The transducer was mounted on an optical rail that allowed it to be traversed along the z-direction up to the full distance between the FLDI focusing lenses, with a precision of 1 mm. At each traverse location and frequency, the AC-coupled FLDI photodetector signal was recorded at 5 MS/s over a record length of 50 µs.

4.2.4 Laminar Jet Results

The main experimental campaign was carried out with the FLDI foci 5.37 mm above the jet orifice. The traversals in this x–z plane were the most thorough, with each (x, z) location being visited between 5 to 17 times. The averaged $\Delta\Phi$ surface obtained is shown in Fig. 4.6, from two different viewpoints. The form of the surface is as expected: in the far-field, where both beams traverse nominally uniform air, there is no phase shift. When one beam begins to interact with the helium jet, $\Delta\Phi$ increases because this beam is now integrating a significantly lower refractive index than air. A maximum is reached, then $\Delta\Phi$ decreases again, until reaching zero at the point where both beams are symmetrically arranged about the jet axis—at this position, the beams integrate the same Δn field. As the jet continues to move in the x-direction, the same $\Delta\Phi$ signal is obtained, except with opposite sign. In the z-direction, the decrease in signal amplitude is especially apparent in Fig. 4.6b. The foci centroid is within the traverse range of the experiment, since the zero-crossing between the positive and negative peaks (the x-coordinate of the centroid) and the global extrema of $\Delta\Phi$ (the z-coordinate of the centroid) are both visible. Note the noise spikes that are visible in Fig. 4.6b: because the signal changes very rapidly near $x = 0$, very small alignment errors can cause an individual $\Delta\Phi$ data point to be assigned to a neighboring bin during the averaging process, skewing the cumulative average.

In addition to the expected drop-off in response as the jet is moved away from $z = 0$, it is also noted that the extrema diverge with increasing x. This is shown in Fig. 4.7, where slices at select values of z are taken from the data in Fig. 4.6. Although the symmetry about $x = 0$ persists, the peak and trough move further from the centroid as z increases. The x-location of these extrema depends on the convolution of the local beam intensity and jet refractive index profiles, which both broaden as functions of position along the jet or beam axes. It is not simple to deduce from characteristic dimensions of the system exactly where this will be; note that even at $z = 0$ the extrema are not at $x = \pm\Delta x/2$ (i.e. co-located with the beam centerlines). However, the qualitative trends can be understood by comparing Fig. 4.12a with Fig. 4.12b. In the former, $d_0 = 10\,\mu\text{m}$ is small, so the beam focusing angle is large, causing the beam widths (and thereby intensity profiles) to rapidly broaden, and the extrema move outwards in x by a considerable amount; in the latter, d_0 is much larger at $100\,\mu\text{m}$, the focusing angle is smaller, and the extrema diverge at a slower rate.

The averaged refractive index field of the jet is shown in Fig. 4.8. Note that it is presented as Δn (with respect to the far-field value). This is done because to calculate n absolutely, a value for the ambient room air n_0 is required. However, this is unnecessary, because Eq. (4.2) integrates a cumulative difference in n between the two beams; hence adding a constant n_0 will not change the calculated $\Delta\Phi$. Hecht (2002) gives n for air and helium at $0\,^\circ\text{C}$, 1 atm, and 589.29 nm, which are

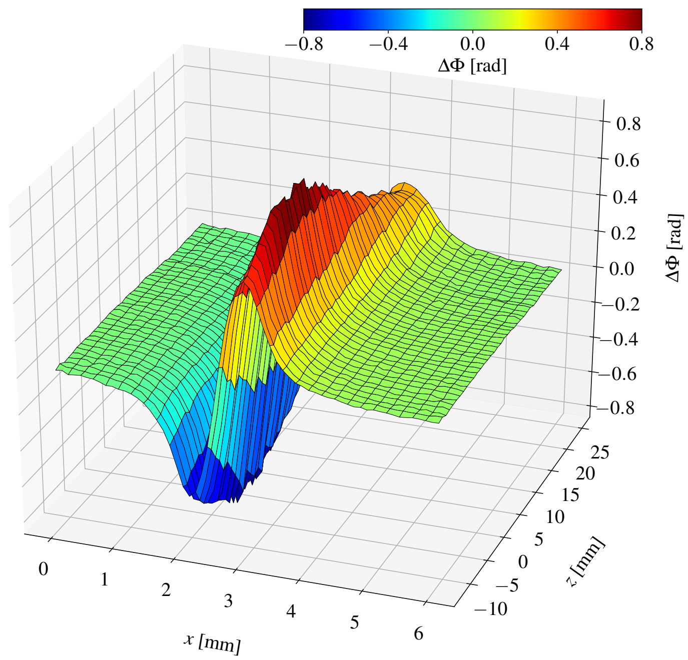

(a)

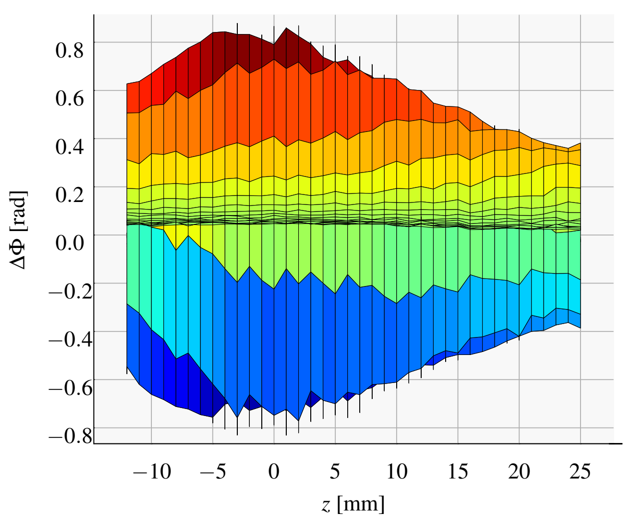

(b)

Figure 4.6: Averaged $\Delta\Phi$ profile obtained from traversing the jet in an x–z plane 5.37 mm below the FLDI foci. $\Delta x = 0.1$ mm, $\Delta z = 1$ mm.

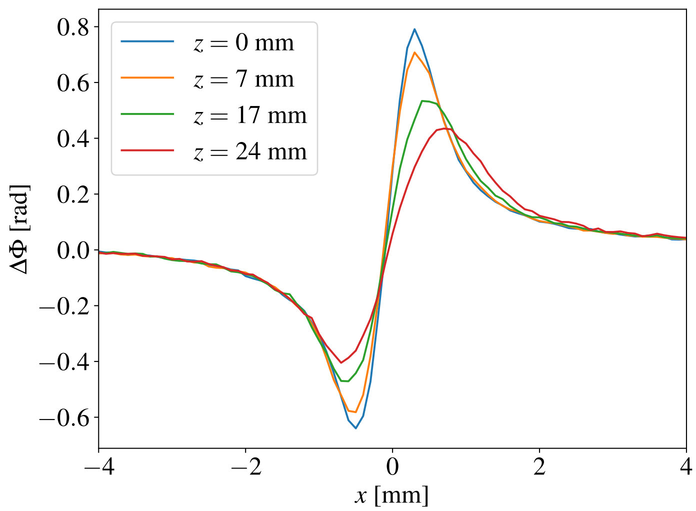

Figure 4.7: Experimental FLDI demonstration of variation in the x-location of the extrema of $\Delta\Phi$ as z increases.

comparable to the conditions in this work. These values yield $\Delta n = n_{He} - n_{air} = -25.7 \times 10^{-5}$. This is slightly higher than the maximum values obtained of around -19.0×10^{-5}, but this full difference would only be expected in regions of pure helium, i.e. near the jet centerline. The contour map and the overlaid radial cross-section show that although the data are still a little noisy even after averaging, both the spatial extent and magnitude of the Δn field are in line with expectations, given the gases used and the size of the jet orifice.

Shown in Fig. 4.9a is a simulation of a single x-traverse in the plane of best focus. The experimental location of $x = 0$ was found by choosing the midpoint of the x-locations of the peak and trough of the data; no further adjustments have been made. The $\Delta\Phi$ are the averaged raw values converted directly from the photodetector voltage measurements. We observe excellent qualitative and quantitative agreement, particularly on the left-hand side of the signal. The simulated FLDI peaks are, by definition, symmetrical about $x = 0$, but asymmetry is observed in the experimental peaks, with a difference in amplitude of about 0.1 rad between the positive and negative peaks. The asymmetry seen between the positive and negative peaks of the experimental data could be due to either asymmetry in the FLDI system (e.g. foci of slightly different diameters) or in the jet (e.g. the jet axis being tilted). To mitigate jet tilt due to drafts, an enclosure was constructed around the jet assembly for the second set of experiments. An additional set of FLDI and MZI measurements were made with the jet assembly raised up, such that the jet origin was very close to the

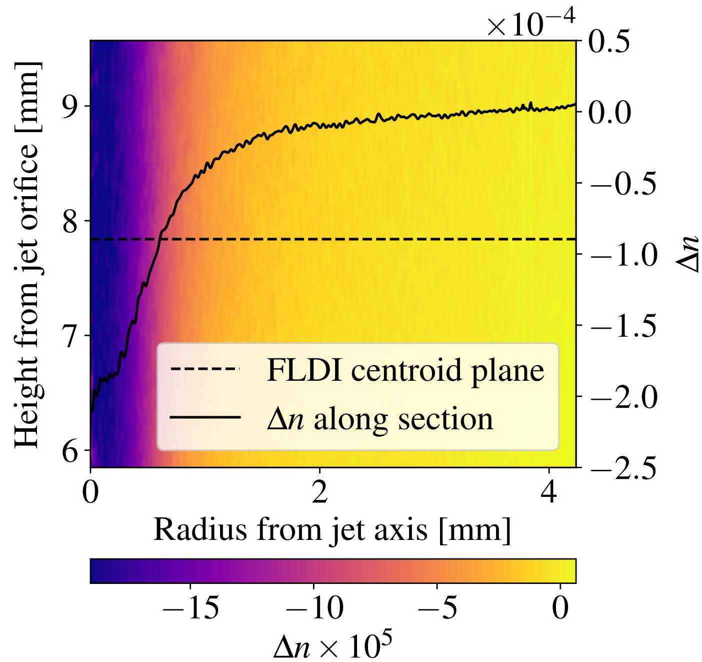

Figure 4.8: Experimental MZI data showing average Δn field of the axisymmetric He jet.

foci, only 0.15 mm below the FLDI plane. This was done in order to probe the core of the jet before it diverges or significantly mixes with the surrounding air. The corresponding comparison of experimental and simulated FLDI output is given in Fig. 4.9b. In this case, the simulation based on MZI measurements and the FLDI data agree within the uncertainty of the FLDI data. Qualitatively, the experiments at each of the two positions give the same response, the main difference being that in the second case, the peaks are narrower and taller. This is due to the jet gradually diverging with axial distance from the nozzle, so at the location of Fig. 4.9b (closer to the nozzle), the refractive index profile is narrower, with increased gradients. Note that in this case, the experimental signal is much more symmetric, i.e. the peaks are of similar amplitude. No changes were made to the FLDI setup between these experiments; the only difference was that the jet assembly was raised and repositioned, which may have subtly changed the orientation of the jet. This result points toward the asymmetry having originated from the jet rather than from the FLDI alignment.

The above results show that the computational model can quantitatively predict the response at beam focus to within experimental error. Next, the spatial sensitivity

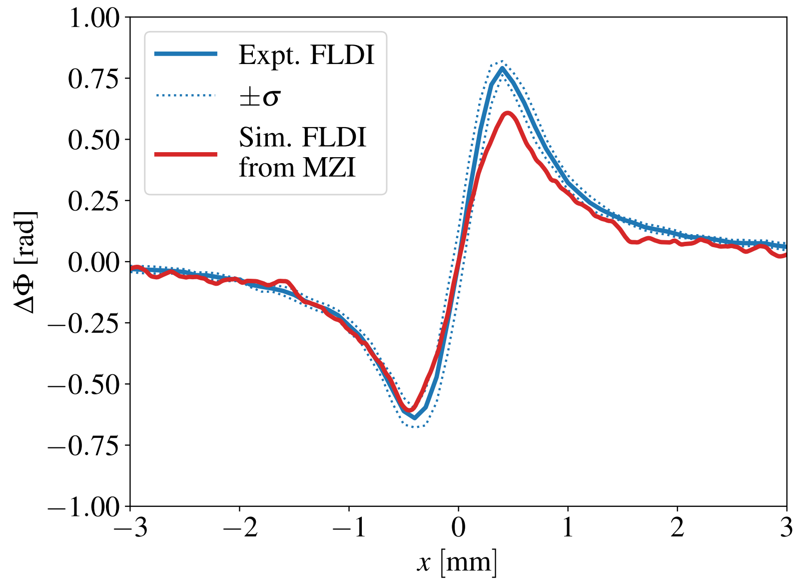

(a) $y_{jet} = -5.37$ mm (Main experimental campaign).

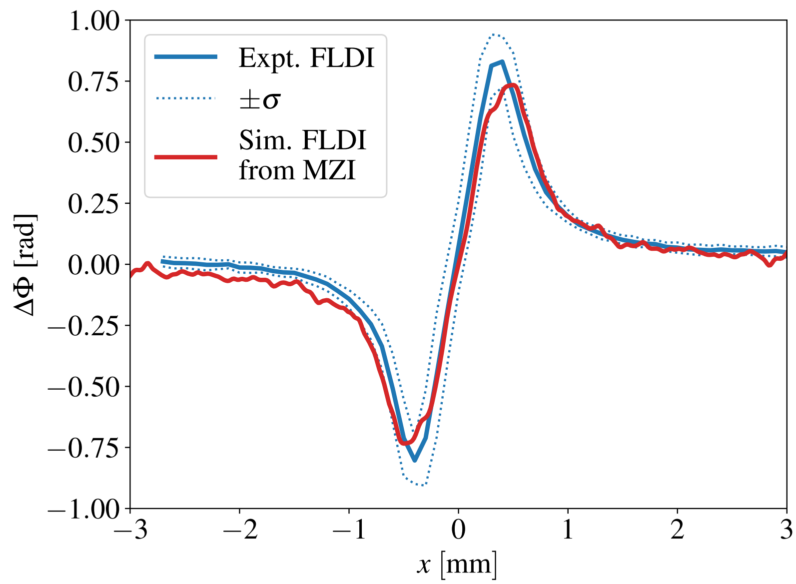

(b) $y_{jet} = -0.15$ mm (Auxiliary experimental campaign).

Figure 4.9: Comparison of experimental (blue) and simulated (red) FLDI responses to an x-traverse of the He jet at two positions below the foci, both at $z = 0$, i.e. in the plane of best focus. The experimental error bounds are quantified by the standard deviation σ of the average measurement at each location.

was examined, as quantified by the drop-off in $\Delta\Phi$ as z increases. At every z location, the maximum and minimum value of $\Delta\Phi$ was extracted. Comparison was then made with simulation in Fig. 4.10a. Various values of d_0 were simulated, still using the same experimental MZI input data for the jet. The rate of drop-off increases as d_0 decreases, and for this set of data, the experimental drop-off most closely matches the $d_0 = 10\,\mu\text{m}$ simulation, which is in reasonable agreement with the experimentally-measured focal size of $7 \pm 2\,\mu\text{m}$. Note that all the curves meet at $z = 0$, i.e. the magnitude of the simulated signal in the beam focal plane (e.g. Fig. 4.9) is not affected by d_0. The experimental decay of the positive peaks plotted in the upper half of Fig. 4.10a is offset from the simulated data by a constant value, due to the aforementioned asymmetry in the peaks present only in this dataset.

Following the two experimental campaigns using the translation stage setup described above (which had a total z range of 37 mm bounding the focal plane), a third setup was created, where a single mechanized stage (oriented in the x-direction) was mounted to a linear rail (oriented in the z-direction). This gave a z range of more than 200 mm in one direction—the jet assembly could be moved right up to the focusing lens (F-1 in Fig. 4.5). This was a less precise setup: the linear rail only offers precision of 1 mm. However, this was satisfactory for the purpose of studying whether the long-range behavior remained consistent with the simulated predictions. Also note that due to the different mounting arrangement, $y_{jet} = 8.90\,\text{mm}$ in this configuration. A new set of MZI interferograms were acquired for this height, to be used in the corresponding simulations. The asymmetry in amplitude between the peaks and troughs was again absent. It was also clear from both the FLDI and MZI data that the jet was less steady. For these data, we are substantially higher in the jet, at almost 18 nozzle diameters from the origin, where it is likely that there is substantial unsteadiness from entrainment. In Fig. 4.10b it is seen that the experimental drop-off agrees best with the simulation for $d_0 = 10\,\mu\text{m}$, which is the same as in Fig. 4.10a.

The effects of changing d_0 and Δx were explored in the simulations. Fig. 4.11 shows simulations in the focal plane. As before, d_0 has no effect on the gross shape of the signal in this plane, but the zoomed inset in Fig. 4.11a illustrates that the signal becomes less smooth as d_0 decreases. Recall from Fig. 4.8 that the input data from MZI has small-scale noise. As d_0 becomes smaller, there is less integration over the beam area of this noise. In contrast, Δx has a large effect on the signal: as Δx increases, the signal magnitude increases, and the peaks also broaden. This

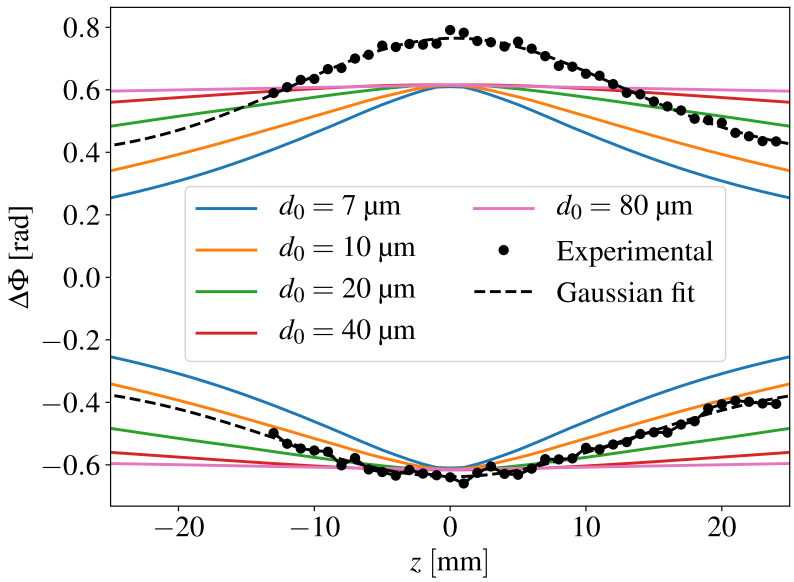

(a) Near-field.

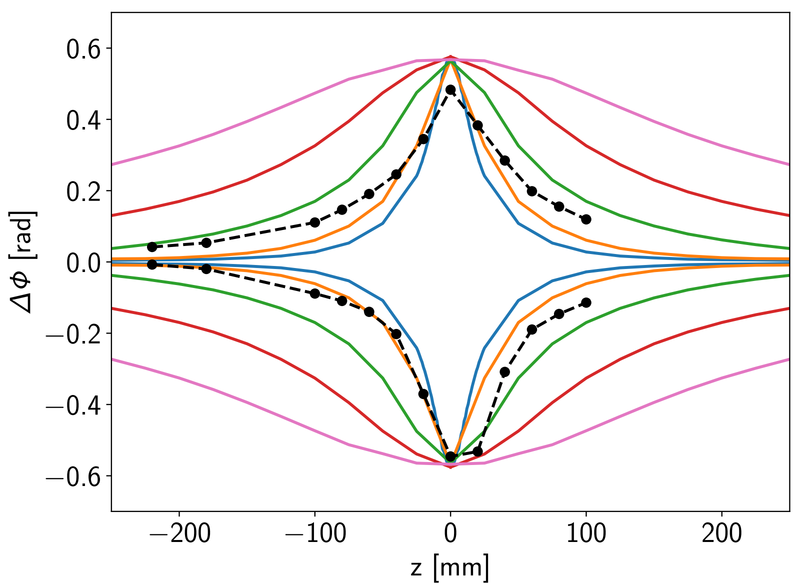

(b) Far-field.

Figure 4.10: Experimental and simulated sensitivity decrease as the jet moves further from the FLDI focal plane. Simulations repeated with a range of foci sizes (d_0).

also makes sense, as when the beams are further apart, they traverse regions of the refractive index field that have larger differences from each other, which is reflected in $\Delta\Phi$, being an integrated measure of these differences. Fig. 4.12 illustrates the strong effect d_0 has on spatial sensitivity: for $d_0 = 10\,\mu\text{m}$, the signal has roughly halved 25 mm from the focus, and has decayed almost to nothing by $z = 200\,\text{mm}$, whereas for $d_0 = 100\,\mu\text{m}$, the signal decrease is far more gradual, with significant response even 500 mm from the focal plane. This behavior is expected: under the assumption of Gaussian beam propagation in the model, a smaller d_0 requires the beams to converge and diverge more rapidly with z. This is alluded to in Settles and Fulghum (2016), where they discuss how the "elongation of the region of best sensitivity in z depends on the lens f/number", as this number determines how rapidly the beam converges to its best focus. Note that in the limit of constant beam diameter, the device is no longer "focused" and is now just a laser differential interferometer (LDI), with no spatial sensitivity.

4.2.5 Free Ultrasonic Acoustic Beam Results

Data from the ultrasound experiment were obtained over the full traverse range of approximately 1000 mm between the two field lenses. The distances in x and y of the acoustic source from the FLDI beam axis were kept constant, as were the driving voltage amplitude and bias. Because the dependence of acoustic beam power and shape on frequency is unknown, the FLDI response was normalized using the maximum response at each frequency. This was done by computing the Welch power spectrum for each dataset [Hann window, 50% overlap, $\Delta f = 305\,\text{Hz}$], then extracting the power spectral density (PSD) of the peak at the driving frequency.

These normalized data are shown in Fig. 4.13a. All frequencies show a Gaussian-like symmetrical decay away from the focal plane. Note that the lowest-frequency data are somewhat noisier, this is due to the transducer having a weaker response below its resonant frequency, and hence lower signal-to-noise ratios. The general trend is that higher frequencies are more spatially filtered than lower frequencies, with the 100 kHz signal decaying to the noise floor in approximately 25% the distance taken by the 30 kHz signal. This is the qualitative result predicted by previous works. In order to more quantitatively compare these data to the theoretical models, the analytical transfer functions of Schmidt and Shepherd are considered. These were derived from the same underlying theory as their numerical model, with analytical solutions obtained for particular disturbance geometries. An analytical treatment of FLDI that incorporates and extends upon these results of Schmidt and Shepherd is

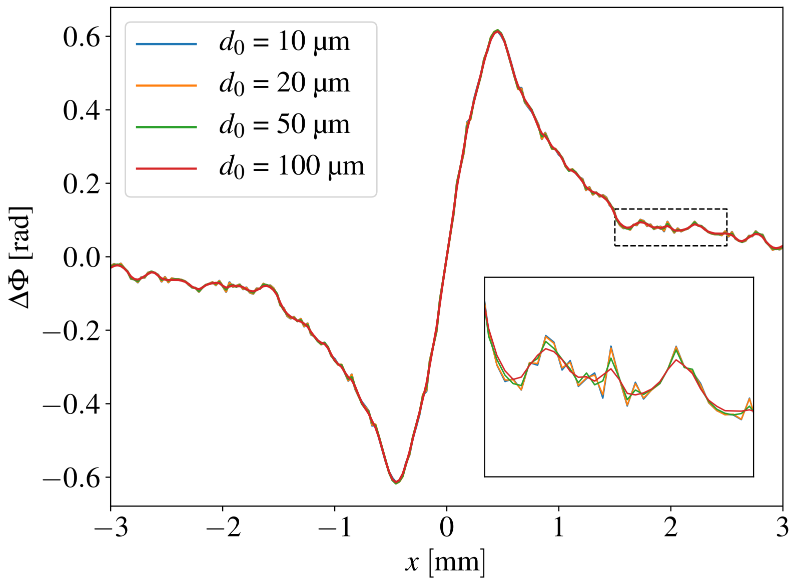

(a) Changing d_0 with $\Delta x = 180\,\mu\mathrm{m}$.

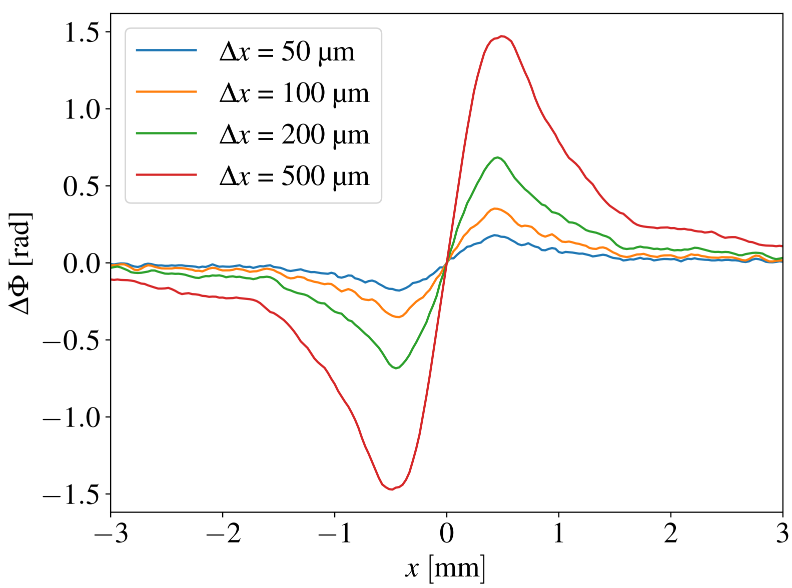

(b) Changing Δx with $d_0 = 80\,\mu\mathrm{m}$.

Figure 4.11: The effects of changing d_0 and Δx in the simulated FLDI, at $z = 0$.

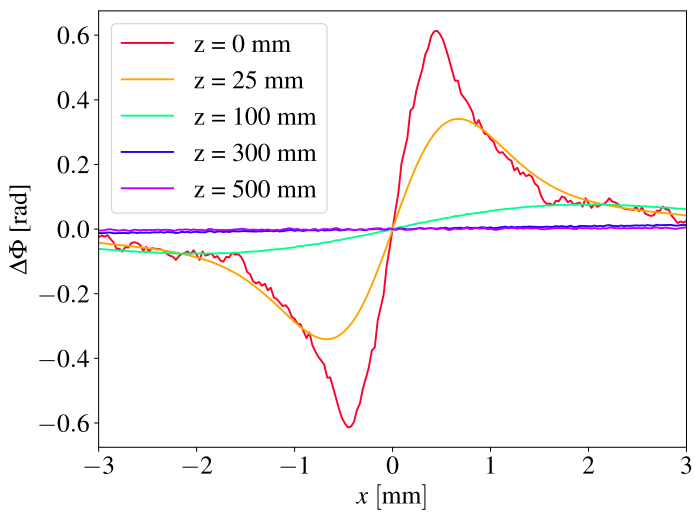

(a) $d_0 = 10\,\mu\text{m}$.

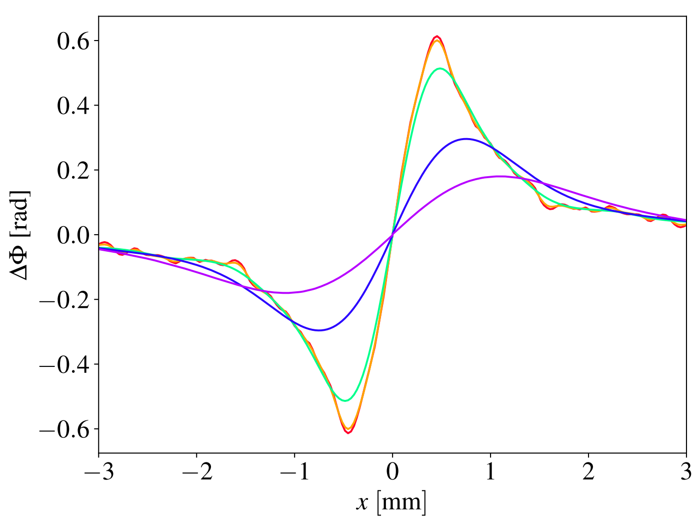

(b) $d_0 = 100\,\mu\text{m}$.

Figure 4.12: Effect on spatial sensitivity of changing d_0 in the simulated FLDI, for fixed $\Delta x = 180\,\mu\text{m}$.

presented in Chapter 5.

Schmidt and Shepherd normalized the response of the FLDI using the true derivative of the phase field, giving transfer functions of the form:

$$H \equiv \frac{\Delta\Phi/\Delta x}{\mathrm{d}\Phi/\mathrm{d}x} \tag{4.7}$$

Their transfer function for the fundamental case of a sinusoidal plane wave of infinitesimal thickness is:

$$\begin{aligned} H(k) &= H_s(k) \times H_w(k) \\ &= \frac{2}{k\Delta x}\sin\left(\frac{k\Delta x}{2}\right) \times \exp\left(-\frac{w^2k^2}{8}\right) \end{aligned} \tag{4.8}$$

where w is the Gaussian beam radius at the location of the disturbance, and k is the wavenumber of the single-frequency sinusoidal disturbance. $H(k)$ is the overall transfer function, which can be decomposed into the effects of the beam separation $H_s(k)$, and of the beam widths $H_w(k)$. It is this latter term that is responsible for the spatial filtering of FLDI, due to w increasing with z; this filtering is wavenumber-dependent.

The acoustic fields generated in this experiment are fundamentally sinusoidal, although with more complicated spatial distributions. However, the infinitesimally-thin plane wave is the fundamental model used as a starting point for derivations of the transfer functions of more complex disturbance geometries, and it is assumed that Eq. (4.8) will still predict the approximate scaling.

We compare the spatial filtering dependence on wave number predicted in Eq. (4.8) with dynamic experimental data. Over the range of axial positions z, at each ultrasound driving frequency f_0, a time-series $\Delta\Phi(t, z)$ was recorded, then transformed into a PSD $|\Delta\Phi(k, z)|^2$. As mentioned, the intensity of response of the FLDI to the driving frequency was quantified by isolating the value of the PSD at the k_0 corresponding to f_0, i.e. $|\Delta\Phi(k_0, z)|^2$. The value of the PSD at the focal plane was used for normalization, such that the quantity H' plotted in Fig. 4.13a is:

$$H'(k_0, z) = \frac{|\Delta\Phi(k_0, z)|^2}{|\Delta\Phi(k_0, 0)|^2} \tag{4.9}$$

As the true derivative of the phase field, $\mathrm{d}\Phi/\mathrm{d}x$, is independent of the z-position of the ultrasound transducer, Eq. (4.9) can be expressed as:

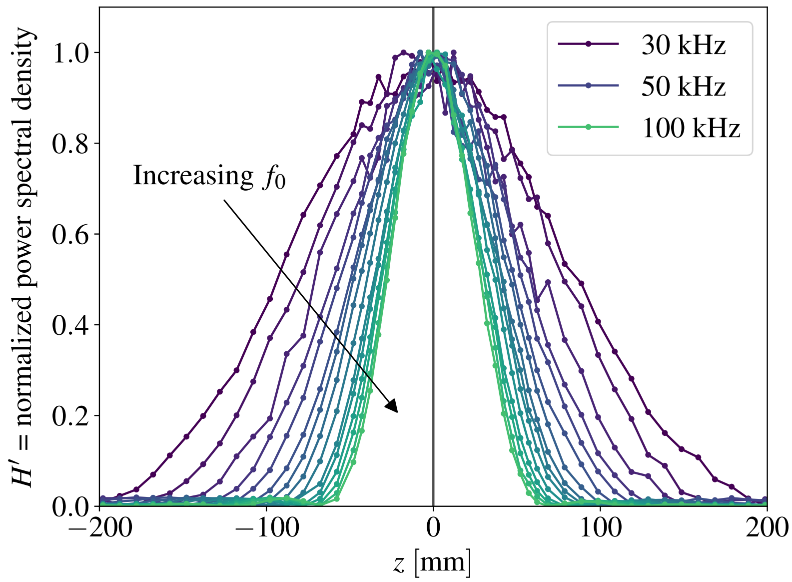

(a)

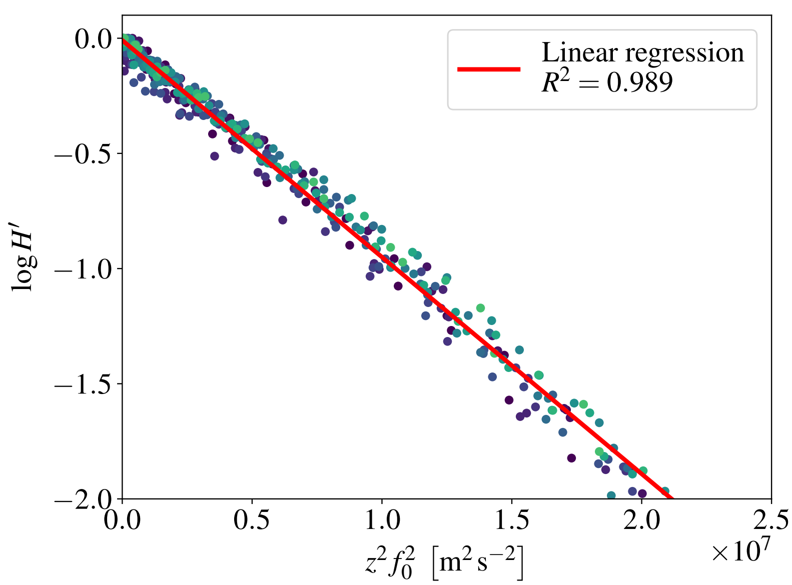

(b)

Figure 4.13: Results of the ultrasonic beam experiments. **(a)** shows the drop-off in FLDI response moving away from the location of the foci at $z = 0$. **(b)** shows the same data, plotting using the variables suggested by the analytical transfer function model. Frequencies f_0 range from 30–100 kHz in 5 kHz increments.

$$H'(k_0, z) = \frac{|\Delta\Phi(k_0, z)|^2/\Delta x^2}{(\mathrm{d}\Phi/\mathrm{d}x)^2} \cdot \frac{(\mathrm{d}\Phi/\mathrm{d}x)^2}{|\Delta\Phi(k_0, 0)|^2/\Delta x^2} \tag{4.10}$$

By comparison with Eq. (4.7), this is a ratio of two of Schmidt and Shepherd's transfer functions, i.e.

$$H'(k_0, z) = \frac{H^2(k_0, z)}{H^2(k_0, 0)} \tag{4.11}$$

Substituting Eq. (4.8) and simplifying, it is found that:

$$H'(k_0, z) = \exp\left(-\frac{\left[w^2(z) - w_0^2\right] k_0^2}{4}\right) \tag{4.12}$$

where $w_0 \equiv w(z = 0)$. Thus if the FLDI response to the ultrasonic field obeys Eq. (4.8), then the normalized experimental data should collapse under the appropriate transform implied by Eq. (4.12). The following approximations are made: except for very close to the foci ($z \to 0$), $w \propto z$ and $w >> w_0$. Additionally, $f_0 \propto k_0$. Hence, $H' \propto \exp(-z^2 f_0^2)$, or equivalently, $\log H' \propto -z^2 f_0^2$. This transformation is applied in Fig. 4.13b where $\log H'$ is plotted against $z^2 f_0^2$, for $\log H' \geq -2$. This limit is chosen because below this, the signal-to-noise ratio becomes small and the maximum response peak cannot be clearly distinguished in the spectra. The data for all frequencies collapse to a straight line with $R^2 = 0.989$, showing that this scaling fundamentally describes how the spatial sensitivity of FLDI depends on disturbance wavenumber, at least for sinusoidal disturbances.

Note that the same result is obtained regardless of whether H or H_w is used in Eq. (4.11), because the H_s component has no z-dependence. Thus the successful collapse of the data can only demonstrate agreement with the spatial filtering due to the varying beam width. The H_s transfer function predicts an aliasing effect when $k = \frac{2n\pi}{\Delta x}$ for integer n. Here, $f_0 = 30$–$100\,\mathrm{kHz}$. Assuming a sound speed of $330\,\mathrm{m\,s^{-1}}$, $k_0 \approx 570$–$1900\,\mathrm{m^{-1}}$, much smaller than the first critical wavenumber of $35\,000\,\mathrm{m^{-1}}$. Equivalently, the acoustic wavelengths 3.3–11 mm are too large compared to $\Delta x \approx 180\,\mu\mathrm{m}$ for aliasing to occur, and so $H_s \approx 1$ over the experimental range.

4.3 Response to Shock Waves

Subsequently to the static and dynamic validation cases discussed in Section 4.2, the FLDI code was used to compare against experimental measurements of hypersonic shock propagation in Caltech's Hypervelocity Expansion Tube (HET; refer to Section 2.1 for facility details). This represents a significant test of the model's predictive capabilities, given the highly dynamic nature of the flow. The previous dynamic validation case was an ultrasonic acoustic beam. However, due to limited knowledge of the magnitude of this acoustic field, only the correct frequency scaling could be validated; the response amplitude and shape were tested separately using the static case of a steady laminar jet. The use of propagating shocks represents a significant test of the predictive capabilities of the model, given the highly dynamic nature of the flow in which propagating shocks extend over much of the FLDI beam path length and have a highly discontinuous refractive index profile. The simplified analytical approach of Parziale et al. (2013b)—while a suitable approximation for the spatially-restricted disturbance fields such as conical boundary layers that they applied it to—does not work well for a flow geometry like a propagating shock because significant contributions are made to the signal even in the outer part of the FLDI beams. Hence, this work demonstrates the need to use the full ray-tracing model in order to accurately predict the instrument response for such flows.

4.3.1 Propagating Shock Waves

An extensive campaign of freestream FLDI measurements were made in HET, with $\Delta\Phi(t)$ acquired for a sample period spanning the entire start-up and test time. This dataset contains shock velocities in the range $1.8 < U_s < 3.6\ \mathrm{km\,s^{-1}}$, in air at ambient temperature and pressures between $145 < P_1 < 275$ mTorr. The majority of experiments are at the upper end of this range, corresponding to $9 < M_s < 10$.

The very fast response of FLDI, combined with sampling rates of 100 MS/s, allow for detailed resolution in the response to the propagating shock, as also observed by Benitez et al. (2020). The response of FLDI to an ideal planar shock can be simulated analytically using a moving discontinuity, with the pre-shock density on one side, and post-shock density on the other (Fig. 4.14). For the ambient density conditions encountered when HET is operated in expansion tube mode, the refractive index changes are small in an absolute sense, even for very strong shocks, and it is found that the linearized and non-linearized forms (Eqs. (4.1) and (4.2) respectively) give essentially identical responses. Although the shock is treated as infinitely thin, the FLDI signal will be a peak of finite width because the shock begins to interact with the outer portions of the conical beams before reaching $x = 0$. Δn across a shock can be large, so even regions of the beam far from the focal plane can contribute significantly to $\Delta\Phi$.

The experimental signals were asymmetric, unlike the symmetric signal predicted with the planar shock model. Additionally to this observed asymmetry, the peak magnitudes of the simulated signals were consistently greater than the corresponding experimental data by a factor of approximately 2.5. Accordingly, the shock was simulated with greater fidelity and the asymmetry and reduction in peak height were investigated.

4.3.2 Peak Asymmetry

First, the asymmetrical shape of the peak was considered. As mentioned in Section 2.1, HET has a free-jet configuration in which there is an abrupt increase in cross-sectional area where the constant-diameter tube opens into the test section. The shock therefore diffracts at the tube exit. If the diffraction creates a sustained interaction with the beam after passage of the planar shock, asymmetry of the signal will result. Inviscid perfect-gas simulations of this flow geometry were performed using the adaptive-mesh code AMROC (Deiterding, 2011), and coupled to the FLDI simulation code. It was found that if the FLDI beams were sufficiently downstream of the tube exit, the shock diffraction would indeed lead to an asymmetrical response of similar form to the experimental data. However, in the actual experiments, the FLDI was only 10 mm from the exit plane (7% of tube diameter), and since only the portions of the shock outside the Mach cone can deviate from planarity, there was negligible effect on the FLDI response (Fig. 4.15), and shock curvature from diffraction did not contribute to FLDI signal asymmetry in these experiments.

Viscous effects can cause shock curvature within the tube. Fig. 4.16 illustrates how a propagating curved shock interacts with the FLDI beam geometry, leading to an asymmetrical signal. Due to being slightly convex with respect to the propagation direction, the initial contact by the shock with the beams is delayed relative to a comparable planar shock. This decreases the characteristic time of the signal rise. Similarly, as the shock leaves the FLDI interrogation zone, a discontinuity remains for longer within the outer parts of the beam volume, lengthening the decay time after the signal peak.

Using the theory for shock curvature on a flat plate developed by Hartunian (1961) and extended to circular tubes by De Boer (1963), the shock shape (shown in

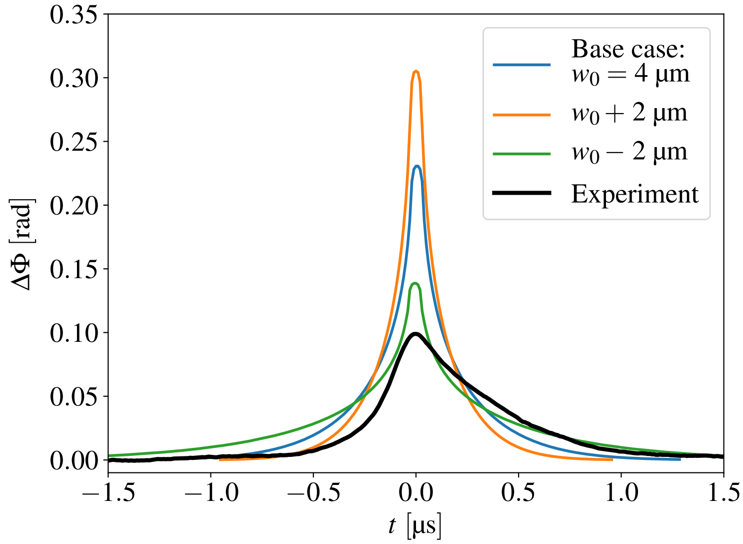

Figure 4.14: Simulated FLDI response to propagating planar shock. Base case optical and flow parameters match nominal experimental values. Additional simulations included for upper and lower uncertainty bounds on w_0.

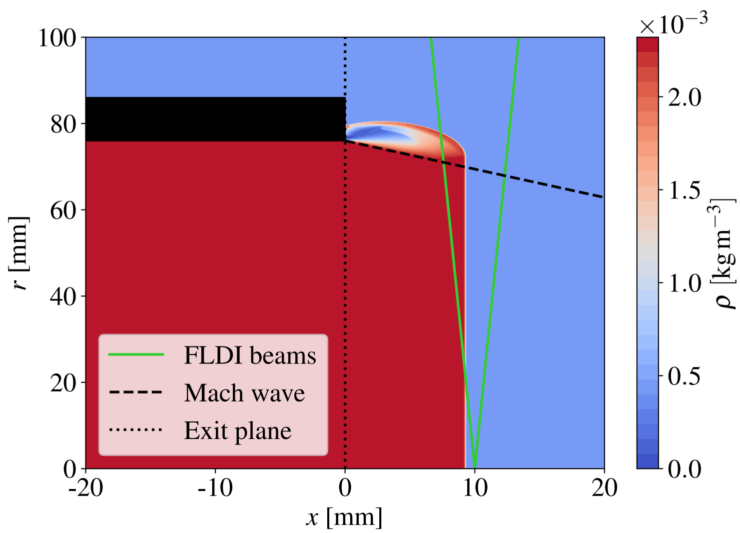

Figure 4.15: CFD simulation of shock diffraction from exit of HET (in black) at $M_s = 9.5$ into the $\rho_\infty = 4 \times 10^{-4}$ kg m^{-3} ambient conditions. FLDI beam extents defined by Gaussian ($1/e^2$) radius, with $w_0 = 5$ µm. The beam offset Δx is not visible at this scale. The Mach wave from tube corner shown to indicate extent of planar shock.

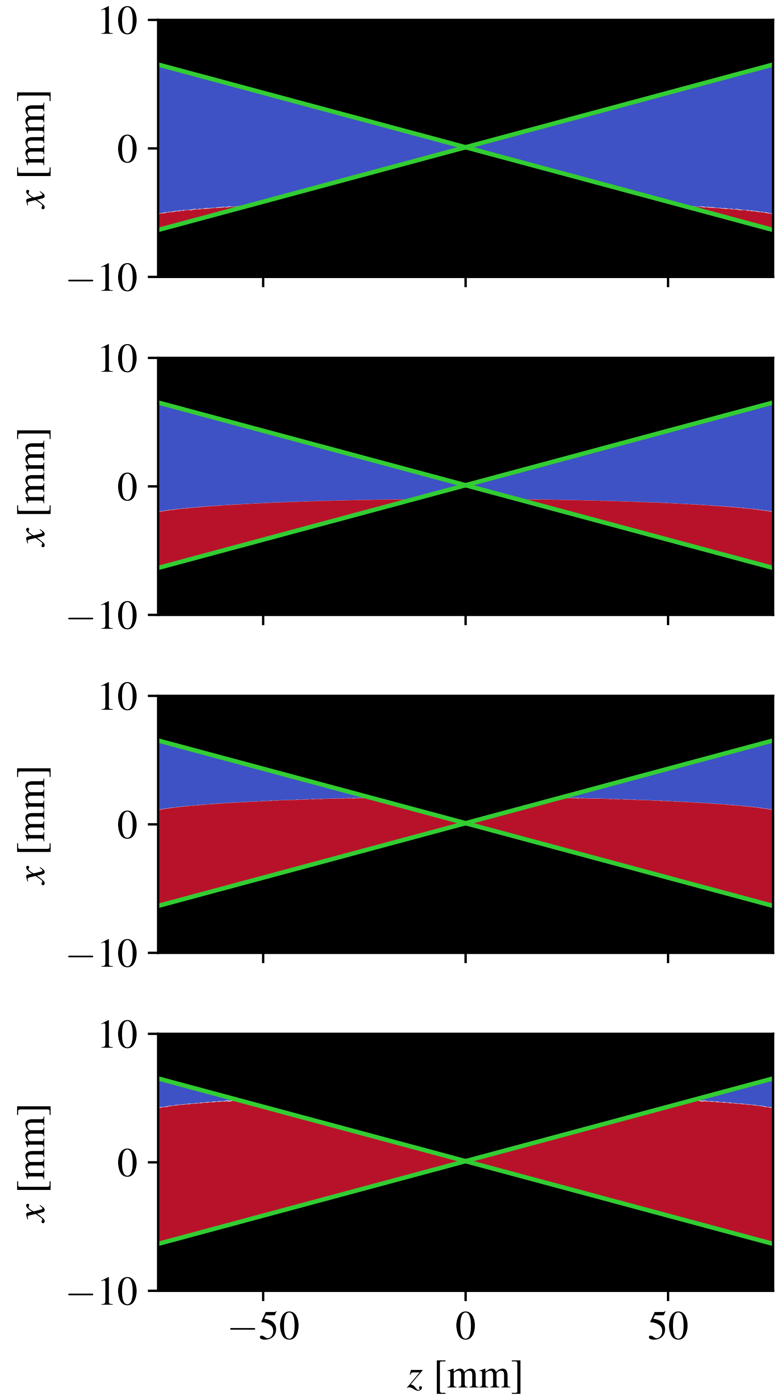

Figure 4.16: Analytical model of a curved shock ($\chi = 1\,\mathrm{mm}$) propagating across an FLDI beam ($w_0 = 4\,\mu\mathrm{m}$). Propagation is in the x-direction, from bottom to top of each figure, with low-density gas in blue and high-density in red. Time increasing from top to bottom frame. Beam extent (green) is here shown as $2w(z)$, the integration limits of the code.

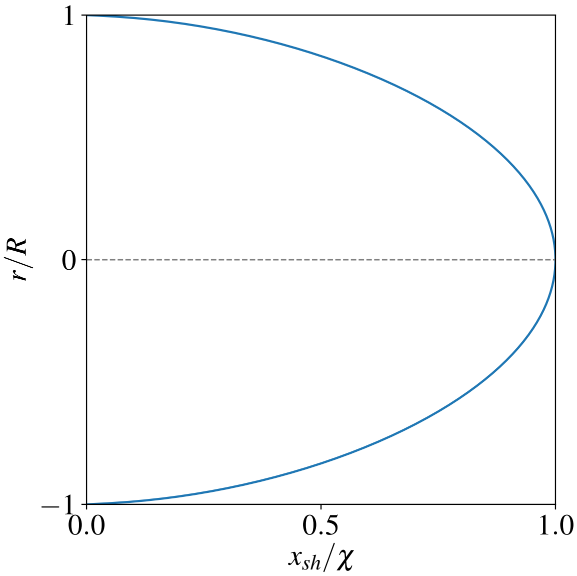

Figure 4.17: Normalized De Boer shock curvature profile for a laminar boundary layer in a circular tube.

Fig. 4.17) is given by:

$$x_{sh}(r) = \chi \left(1 - K \int_0^\infty \frac{I_0(kr) - 1}{k^{s+1} I_1(kR)} dk\right) \tag{4.13}$$

where r is the tube radial coordinate, R is the tube radius, I_n are modified Bessel functions of the first kind and n^{th} order, $s = \frac{1}{2}$ when the wall boundary layer is laminar in the region close behind the shock, and K is a scaling factor. χ represents the maximum axial displacement of the curved shock from planarity (i.e. at the centerline). Note that some symbols have been altered from the referenced work, and the result has been shifted to the lab-fixed frame rather than the shock-fixed frame of the original derivation.

Experimental data were used to estimate χ. Schmidt presents measurements in a 6 in. shock tube (the same radius as HET) that agree with De Boer's theory (Schmidt, 1976). These data are over a pre-shock pressure range of 10–110 mTorr in N_2, whereas our data span 145–275 mTorr in air; extrapolation predicts $\chi \approx$ 1.5–2.5 mm for these conditions. Propagating shocks with this profile and χ similar to the literature were simulated (Fig. 4.18). The density behind the curved shock is assumed to be uniform as per the planar shock base case. Even for curvatures exceeding the predicted range of the experimental conditions, the uniform post-shock density assumption is reasonable; for example at $\chi = 4$ mm, the most extreme oblique shock angle (at the wall) is 88°, yielding a local normal shock Mach number $M_{s,n} = 0.978M_s$. For $M_s > 9$, the post-shock density ρ_2 only differs from the centerline value by a maximum of 2%.

It was found that a curved shock with $\chi = 2$ mm consistently matched the experimental data closely on both the steep leading edge and broad trailing edge. This is in the expected range of χ. However, the central portion of the response around the peak was not matched; this discrepancy will be discussed below. Other experimental studies have found evidence for more complex shock shapes, e.g. "indentations" at the centerline, or non-axisymmetry (Kiefer and Manson, 1981; Liepmann and Bowman, 1964). These effects are outside the scope of this work, but could conceivably cause changes to the shock response shape.

Post-shock nonequilibrium may also need to be considered. We examine the effect of dissociation behind a planar shock on FLDI signal by calculating the nonequilibrium density profiles from a thermally- and chemically-frozen post-shock condition to equilibrium, using both single- and two-temperature models (Browne et al., 2008). Both models predict a very long relaxation region of $O(1\text{ m}) >> \Delta x$, giving very gradual refractive index slopes over distances $O(\Delta x)$. At the low pressure conditions of our experiments, the reactions are so gradual that on the timescales of the FLDI measurement, the post-shock state can be assumed to be quasi-steady. We note that analytical solutions have been developed in the case where gradients due to chemical reactions need to be considered at a curved shock wave, as reviewed by Hornung (2010).

4.3.3 Peak Magnitude

It can already be observed from Fig. 4.14 that the foci size w_0 has a strong effect on the peak magnitude. As w_0 decreases, the beam focusing angle increases, so the shock begins to interact with the beam earlier, and the entire interaction duration increases. Hence, the response broadens while decreasing in height. Very small absolute changes in w_0 cause large changes in the response magnitude, demonstrating the difficulties caused by the large relative uncertainty in the experimental measurement of w_0. A sensitivity study was also performed for the beam separation Δx, and the shock speed U_s. It was expected that increasing Δx should increase the peak width due to the longer time-of-flight between the beams. This is indeed the case, along with an increase in height, but these changes are slight, even for perturbations of

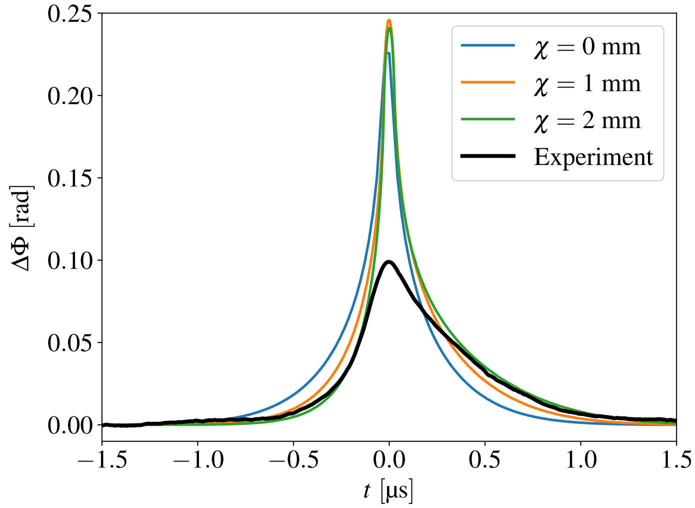

Figure 4.18: Simulated FLDI response to analytically-modeled propagating curved shocks with De Boer profile. $w_0 = 4\,\mu\text{m}$. χ is maximum extent of curvature; $\chi = 0$ is planar by definition, and identical to result in Fig. 4.14.

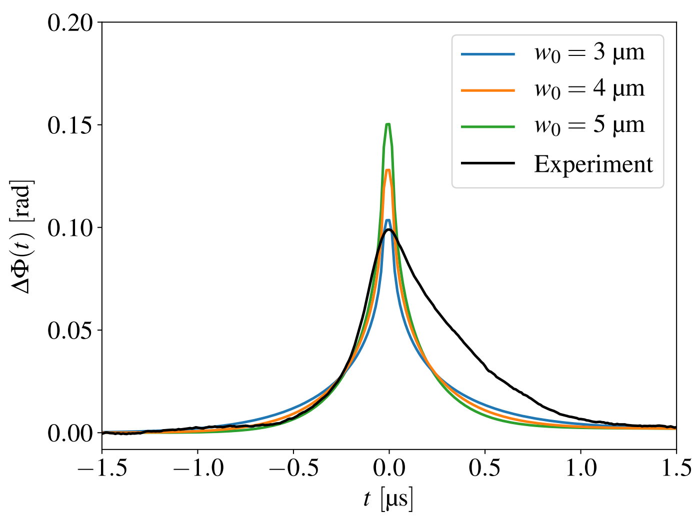

Figure 4.19: FLDI responses using post-shock $n(x)$ computed with two-temperature model.

$\pm 10\,\mu\text{m}$ which exceeds the experimental uncertainty of $\pm 2\,\mu\text{m}$. U_s also affects the time-of-flight and alters the post-shock state. Again it is found that even perturbations of $\pm 300\,\text{m s}^{-1}$ have marginal effect on the signal: since at $M_s > 9$ the density (and hence refractive index) ratios change little near the strong shock limit. From this study it is concluded that Δx and U_s have only secondary influence on the FLDI shock response; w_0 strongly controls the peak size.

All the simulated shock variants thus far have involved a discontinuous jump in refractive index, which gives a strong response since FLDI responds to gradients of n. The aforementioned range of P_1 for the present data correspond to mean free paths $190 < \Lambda < 350\,\mu\text{m}$. The shock structure has a sigmoid profile, modeled by Thompson et al. (1983) as:

$$f(x) = f_1 + \frac{f_2 - f_1}{2}\left[1 + \tanh\left(\frac{x}{\Delta_m/2}\right)\right] \tag{4.14}$$

where Δ_m is a measure of thickness based on the maximum slope. f can represent any of T, u, or ρ (and hence n), all of which have the same profile, merely displaced spatially. Thompson et al. show that an ideal-gas nitrogen model agrees well with experimental values of Δ_m, and presents two curves for ground-state and fully-vibrationally-excited N_2 ($\gamma = 7/5$ and $9/7$, respectively) as bounds to these experimental data. Here we assume that these N_2 bounds hold for air, thus over $9 < M_s < 10$, an approximate thickness range of $0.5 \lesssim \Delta_m \lesssim 1.5$ mm is predicted.

4.3.4 Combined Asymmetry and Magnitude Effects

A composite shock model was created, with curvature and thickness governed by Eq. (4.13) and Eq. (4.14), respectively. Thompson et al. show that Δ_m increases gradually and linearly with M_s above $M_s \approx 5$. It was previously demonstrated that curvature causes a maximum variation in $M_{s,n}$ of $\sim 2‰$. The scatter in the $\Delta_m = f(M_s)$ data of Thompson et al. is considerably larger than this, so we make the assumption that the shock thickness will not vary noticeably along the curved shock front. Thus the two effects can be treated independently. FLDI simulations were performed over many combinations in the uncertainty ranges of w_0, χ, and Δ_m. The latter two parameters show only minor variation over these ranges. Two of the best matches to experimental data are shown in Fig. 4.20. The height and leading-edge rise can be closely matched, and the asymmetry also shows good agreement.

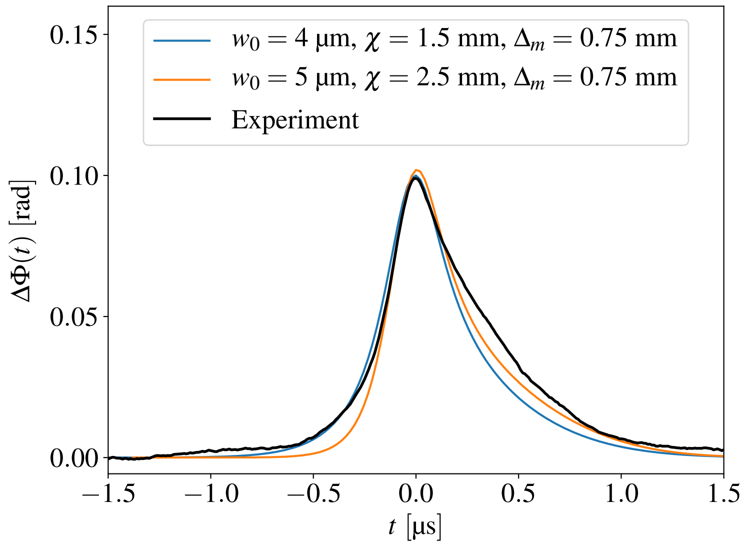

Figure 4.20: Selected FLDI responses to an analytical model incorporating shock curvature χ and thickness Δ_m within expected bounds. The post-shock refractive index corresponds to the chemically-frozen state.

4.4 Application To Complex Shock-Dominated Flow-Fields

The results of Section 4.3 demonstrate that a quasi-steady ray-tracing scheme is capable of accurate reproduction of the FLDI response to highly dynamic propagating shocks. This extends the library of analytical flow geometries for which FLDI response has been predicted.

Building on these results, the scheme was applied to a practical experimental case where FLDI was being used to make measurements in a flow-field dominated by complex inviscid shock interactions. Upon diffracting from the HET tube exit, the shock wave subsequently interacts with the complex geometry of the test section outside the core flow. For these experiments, the HET was operated as a shock tube with higher initial back pressure than is used in expansion tube mode, and repeatable features were observed in the FLDI signals following the primary shock that were not explained by one-dimensional shock-tube theory. A CFD simulation was performed using AMROC of the shock diffracting from the tube exit, similarly to that described in Section 4.3, but with the addition of a wall and window cavity matching the HET test section dimensions. The simulation was limited to 2D, although the true test section cross-section has 3D symmetry, being rectangular with a window on each face. The output of this simulation was coupled to the FLDI numerical code. The Cartesian CFD and polar-conical FLDI grids are independently-defined and not aligned with each other, requiring interpolation of the refractive index field from the CFD output onto the FLDI beams. Because AMROC uses adaptive-mesh refinement (AMR), the grid is not regular, instead comprising a complex arrangement of patches with different resolutions that change at each time-step. This necessitates the use of an unstructured interpolation algorithm, which is much slower and more expensive than the corresponding regular-grid interpolation. For this work, the `LinearNDInterpolator` function from the SciPy package is used (Virtanen et al., 2020). Additionally, in order to get adequate temporal resolution in the FLDI simulation, the full-field output of the CFD must be stored at a large number of time-steps. These computational time and storage requirements of both the CFD and interpolation process restrict us to 2D flow-fields.

The results were in agreement with experiment (Fig. 4.21). Referring to the numbered time instants in Fig. 4.22, the coupled simulation reproduces the initial transmitted shock (frame 1), which is followed by a dip then a gradual increase of the mean signal caused by the establishment of the steady expansion region (frames 2 and 3). The diffracting shock begins to reflect from the window cavity (frame 4) and a pair of reflected shocks generates the prominent "W"-shaped signal feature (frame 5). At later times (frame 6) a complex flow-field is established due to further reflections in the cavity between the tube and test-section walls.

Although this is a far more complicated flow-field than the simple propagating, near-planar shocks of the validation study, the FLDI code is able to quantitatively predict (both in t and $\Delta\Phi$) detailed features of the experimental signal. Ultimately the purpose of these instrument simulations was to show the amplitude and shape of FLDI response to dynamic, nonlinear wave interactions in the experiments were meaningful when interpreted using the ray-tracing approach. Fully 3D, Navier-Stokes simulations were outside the scope of this study, thus any discrepancies are likely due to the CFD simulation being performed with 2D geometry and without viscous effects. The smoother appearance of the simulated signal is due to a combination of temporal resolution limitations in the simulations, and high-frequency noise in the experimental signal acquisition.

The main reflected shocks cause oppositely-signed phase shifts with respect to the initial transmitted shock since they are traveling upstream, thereby interacting with the two FLDI beams in reverse order to downstream-propagating disturbances. Note that the reflected shock signals are far weaker ($\Delta\Phi \sim O(0.1\,\mathrm{rad})$) than that of the incident shock, which is strong enough to exceed $\Delta\Phi = \pi/2$, causing phase-

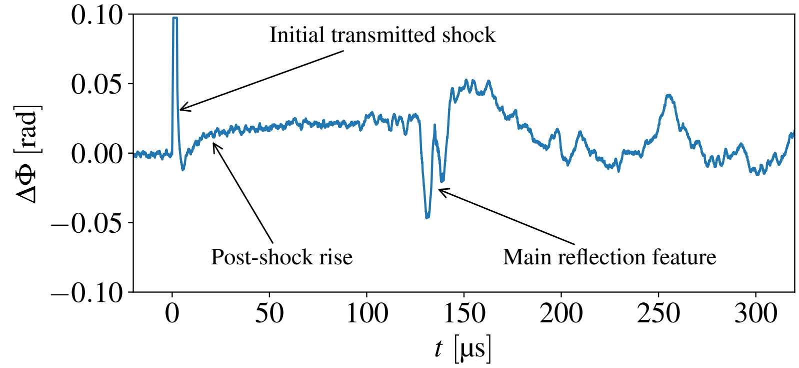

(a) Experimental FLDI signal.

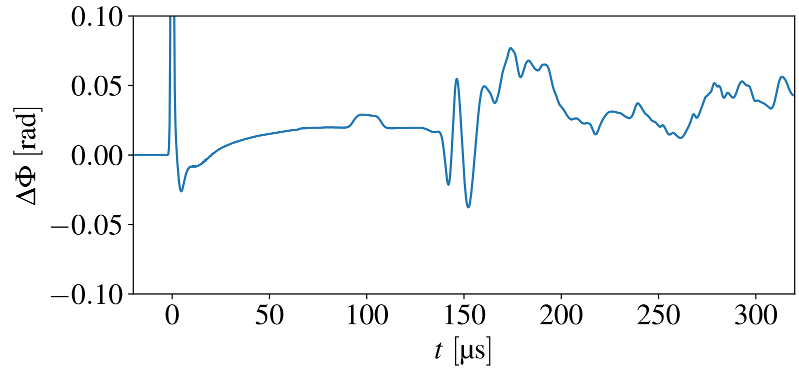

(b) Simulated FLDI signal.

Figure 4.21: Comparison between experimental data and a coupled AMROC/FLDI simulation of the shock reflection flow-field in HET.

wrapping near its peak (not visible at the scale of Fig. 4.21). Because these data are from shock-tube mode, densities are much greater than the usual expansion-tube operation of HET—hence the smaller transmitted shock magnitudes of Section 4.3. This means the reflected shock field in expansion-tube mode will be correspondingly weaker. However, these results show that due to the very steep refractive index gradients caused by propagating shocks, FLDI is not able to filter out contributions from shocks in the outer parts of the beams as effectively as for less severe gradients, such as turbulent fields.

In addition to the shock-related features, the smooth rise in $\Delta\Phi$ following the initial transmitted shock can also be attributed to a particular flow feature: the growth of the expansion region behind the curved diffracting shock. This is important to

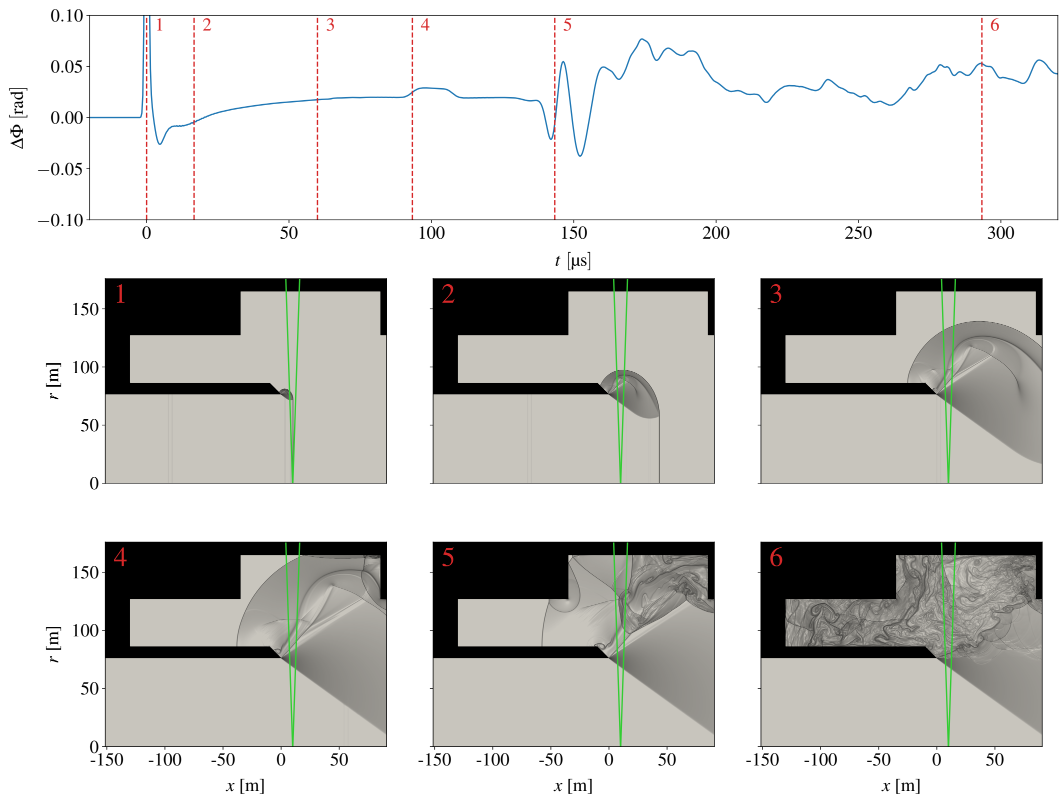

Figure 4.22: Instantaneous flow-field shown at several stages of development. The time instants for each frame are indicated by red numbers, FLDI beam extents are shown in green, and solid HET geometry is in black. The greyscale contours of the flow-field represent pseudo-Schlieren density plots, using a log scale to accentuate weaker flow features.

note, because once fully-established, the spatially-steady density gradient causes a constant offset from the baseline of $\Delta\Phi = 0$, which is relevant when doing spectral analysis of the test-time signal.

The success of this practical example gives credence to the concept of coupling the FLDI simulation code to more detailed CFD simulations, for example a viscous computation that includes the turbulent shear layers emanating from the nozzle trailing edge. Such shear layers are thought to be the main contributors of unwanted noise for line-of-sight-integrated optical techniques, and understanding the extent to which FLDI can reject these contributions is central to the technique's utility on these types of hypersonic facility.

The computational performance limitations discussed here may need to be addressed in order to move to larger-scale 3D couplings. If a regular-grid CFD scheme is used, the interpolation issues should be largely alleviated, although without AMR the CFD computational time itself may drastically increase in order to get sufficient spatial resolution, especially in shock-dominated flows. Re-implementing and optimizing the FLDI simulation scheme in a faster compiled language such as C++ may also yield improvements.

4.5 Conclusions

The goals of Section 4.2 were to: obtain precise spatial and temporal characterization data for an FLDI system, use these data to validate a model of FLDI, then use the validated model to explore the sensitivity of the instrument's response to various input parameters. FLDI spatial measurement of the refractive index field of a laminar jet phase object was independently characterized using a Mach-Zehnder interferometer, and found to be in quantitative agreement to within experimental uncertainty. Temporal characterization was obtained using a free ultrasonic acoustic beam, generated by an ultrasonic transducer, and traversed along the length of the beam pair.

Analysis shows that the simple ray-tracing scheme given by Eq. (4.2) along with the numerical discretization scheme detailed in Schmidt and Shepherd give accurate quantitative predictions for the static response of an FLDI system. A key result is that the Gaussian beam diameter d_0 strongly affects the rate of signal drop-off as the phase object moves away from the focal plane. A reduction in d_0 corresponds to a shortened sensitive length, in accordance with the predictions of previous works. Given this validation, the model can be utilized in the future to inform design decisions when constructing FLDI systems for a particular application. Further, as already mentioned, Eq. (4.2) is not specific to FLDI. Here, the beam geometry is defined to represent the parallel focused beam pair of FLDI. But by modifying this scheme, other optical arrangements can be explored.

The dynamic response of FLDI was obtained for a time-varying phase object generated by sinusoidal acoustic waves at pure frequencies in the ultrasound range. Firstly, it was found that FLDI is easily sensitive enough to detect these acoustic density fluctuations, which were too weak to be visualized at all by the MZI. Secondly, the drop-off in FLDI response was found to be strongly dependent on the wavenumber of these disturbances. This dependence agreed with the Gaussian scaling proposed by Schmidt and Shepherd, which was derived from the same theory underpinning the numerical simulation scheme used for the static response portion of this work.

Following this, quantitatively-accurate predictions of experimental FLDI results were obtained using simple analytical models of propagating shock waves coupled with ray-tracing calculations. Experiments were carried out using shock waves up to Mach 10 at low initial pressure conditions. A model that combines viscous shock curvature and shock thickness was in good agreement with the experimentally-measured phase change profile. Within the uncertainty bounds on the optical parameters, this chosen shock configuration is unlikely to be a unique solution, however the analysis illustrates that a physically-realistic shock model can produce accurate predictions of FLDI response. In order to get even closer agreement, it is likely that either full simulations of the shock and boundary layer structures, or more complex models would be required. One possibility for improved model agreement with observed signal asymmetry is the use of a more accurate shock thickness model, since at high M_s, the shock profile deviates from a symmetric sigmoid (Liepmann et al., 1962).

This result serves to demonstrate the capabilities of FLDI in resolving highly transient flow features with sub-microsecond timescales. The previous dynamic validation study in Section 4.2 showed that the spatial sensitivity is frequency-dependent, with the scaling predicted by Schmidt and Shepherd. Here, the validation is extended to demonstrate that the model can predict both the shape and magnitude of the FLDI response to dynamic flows; this was previously only validated for a static flow case.

This second validation study was performed using relatively simple analytical descriptions of shock waves. The capabilities of the FLDI simulation code were then further demonstrated by predicting the experimentally-measured FLDI response to the multiple reflections of a confined, diffracting shock. This complex flow-field cannot be described analytically, and so the FLDI code was coupled to a CFD simulation. The success of this coupling suggests employing CFD more widely to better understand and interpret the behavior of FLDI to a variety of spatially- and temporally-complex flows that are difficult to model using analytical transfer functions.

Overall, the results of this chapter show that if $n(\underline{x}, t)$ is known exactly, the FLDI response $\Delta\Phi(t)$ can be accurately computed using the ray-tracing method. If the general functional form is known to within a small number of scaling parameters (as in Section 4.3), these parameters can be adjusted until the simulated response matches that of the experiment. However in general, $n(\underline{x}, t)$ is not known *a priori*, which poses a problem when trying to interpret a flow-field via FLDI alone. Chapter 5 addresses this issue further.

Chapter 5

RESULTS: DENSITY FIELD DISTURBANCES FROM FLDI

Chapter 4 introduced the ray-tracing model of FLDI proposed by Schmidt and Shepherd (2015), described mathematically by Eq. (4.2). The validity of this model, and its associated numerical implementation, was demonstrated through three experiments. The key applicable result of this validation is that we now have the capability to compute the FLDI temporal response, in terms of phase shift $\Delta\Phi(t)$, to any arbitrary dynamic refractive index field. Expressed otherwise, the forward problem $\Delta\Phi(t) = f\left(n(\underline{x}, t)\right)$ is solvable for all $n(\underline{x}, t)$.

From a practical standpoint, the inverse problem $n(\underline{x}, t) = f^{-1}\left(\Delta\Phi(t)\right)$ needs to be solved, as a reconstruction of $n(\underline{x}, t)$ leads easily to the density field $\rho(\underline{x}, t)$. However, at each instant in time, the 3D scalar field input $n(\underline{x})$ yields a single scalar output $\Delta\Phi$: not being bijective, a large amount of information is lost in this process. In order to perform the inversion, some assumptions need to be made about the geometry and symmetry of the field. An example of this is seen in Section 4.3, where a certain shock model was assumed, and a limited number of parameters were tuned in order to match the computed response with experimental data.

The purpose of this chapter is to explore the analytical implications of Eq. (4.2) for an important class of disturbance field: sinusoidal plane waves. As discussed in Section 1.3, the freestream noise of conventional hypersonic ground test facilities is dominated by far-field acoustic radiation from wall boundary layers. This class of noise can be modeled using superimposed plane waves, resulting in a refractive index field described by:

$$n(\underline{x}, t) = \frac{1}{\sqrt{N}} \sum_{i=1}^{M} \sum_{j=1}^{N} A_j \cos\left(\underline{k}_{i,j} \cdot \underline{x} - 2\pi f_j t + \varphi_{i,j}\right) \tag{5.1}$$

This represents the superposition of waves propagating in M different directions, each composed of N frequency components; the meanings of the other terms will be introduced subsequently. This chapter will build towards this generalized result in stages.

We will derive relationships between the time- and frequency-domain responses of FLDI, and demonstrate that density fluctuation spectra can be ultimately be recovered from $\Delta\Phi(t)$ under some realistic assumptions. As alluded to in Section 1.2, previous works by Schmidt and Shepherd (2015) and Settles and Fulghum (2016) have derived results for the response of FLDI to certain classes of disturbance field, including some special cases of propagating sinusoidal waves. These results will first be reviewed in Section 5.1, before beginning the current approach from Section 5.2 onwards.

5.1 Results from Previous Works

5.1.1 Schmidt and Shepherd

Schmidt and Shepherd defined an FLDI transfer function H, already introduced as Eq. (4.7) and given here with more detailed notation:

$$H \equiv \frac{\max|\Delta\Phi|/\Delta x}{(\mathrm{d}\Phi/\mathrm{d}x)_{x=0}} \tag{5.2}$$

This compares the actual response to the ideal response—where the true derivative is measured at the instrument origin. Although not explicitly stated in their paper, it can be inferred from their provided intermediate derivation steps that Φ must be defined as follows:

$$\Phi(x) = \frac{2\pi}{\lambda_L} n(x) \tag{5.3}$$

They considered a sinusoidal refractive index wave propagating in the x-direction. Because the FLDI response is related to $\mathrm{d}n/\mathrm{d}x$, $\Delta\Phi$ is maximized when $n \sim \pm\sin(x)$ and is zero when $n \sim \pm\cos(x)$. Hence to obtain $\max|\Delta\Phi|$, the wave was fixed in space and chosen to be:

$$n' = A\sin(kx) \tag{5.4}$$

where $k \equiv 2\pi/\lambda_w$ is the wavenumber, with λ_w the disturbance wavelength. The prime indicates this is the fluctuating component of refractive index about some mean. Two geometry cases were considered. One variant was limited to an infinitesimal z-plane located at $z = z_a$:

$$n' = A\sin(kx)\delta(z - z_a) \tag{5.5}$$

where δ is the Dirac delta function. Note that they did not use this notation, instead just using $\delta(z)$, but they do mention the more general case where $z_a \neq 0$. The other spanned a finite z-extent:

$$n' = A\sin(kx)\left[U(z+L) - U(z-L)\right] \tag{5.6}$$

where U is the Heaviside step function, and the disturbance field has depth $2L$, symmetrical about the focal plane, i.e. $-L < z < +L$. For simplicity, from now on the primes will be dropped from n', since the differential response is unaffected by the mean refractive index.

The ideal FLDI comprises two point detectors with separation tending to zero. Hence, the overall transfer function was found by considering in turn each of the non-ideal effects of the real instrument: finite beam width, and finite beam separation. These effects were encapsulated by H_w and H_s respectively. To compute the former, they took $\Delta x \longrightarrow 0$, i.e. removing the effects of beam separation and only considering the finite beam widths. For the latter, the beams were treated as two point detectors, i.e. removing the effects of beam width and only considering their separation. Expressed mathematically:

$$H_w = \frac{\lim\limits_{\Delta x \to 0} \frac{\Delta\Phi}{\Delta x}}{\left(\frac{\mathrm{d}\Phi}{\mathrm{d}x}\right)_{x=0}} \tag{5.7a}$$

$$H_s = \frac{\frac{1}{\Delta x}\left[\Phi\left(x = +\frac{\Delta x}{2}\right) - \Phi\left(x = -\frac{\Delta x}{2}\right)\right]}{\left(\frac{\mathrm{d}\Phi}{\mathrm{d}x}\right)_{x=0}} \tag{5.7b}$$

To evaluate $\Delta\Phi$ for each geometry case, Eq. (5.5) or Eq. (5.6) were substituted into Eq. (4.2). For Φ, Eq. (5.4) was used for both cases, because point detectors have no y- or z-extent; also the Dirac delta is only well-defined in an integral sense.

Recall that the governing ray-tracing equation, Eq. (4.2), is defined using arbitrary in-plane co-ordinates (ξ, η). Schmidt and Shepherd showed that by using the properties of the Dirac delta function, and differentiation from first principles, an analytical solution could be obtained using Cartesian co-ordinates (x, y) for (ξ, η) if the infinitesimal wave Eq. (5.5) was used. The result is:

$$H_w(k) = \exp\left(-\frac{w^2 k^2}{8}\right) \tag{5.8}$$

where $w = w(z = z_a)$ is evaluated at the plane wave location. For $z_a = 0$, $w = w_0$. In Cartesian co-ordinates, substituting the finite-extent wave Eq. (5.6) directly into the integrals arranged as shown in Eq. (4.2) does not lead to a tractable solution. Instead, they integrated the infinitesimal solution Eq. (5.8) in z:

$$\begin{aligned} H_w(k) &= \frac{1}{2L}\int_{-L}^{L} \exp\left(-\frac{w^2(z)k^2}{8}\right) \mathrm{d}z \\ &= \frac{\pi w_0 \sqrt{2\pi}}{kL\lambda_L} \exp\left(-\frac{w_0^2 k^2}{8}\right) \mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \end{aligned} \tag{5.9}$$

Note that they included a normalization factor of $\frac{1}{2L}$ not indicated by any of the governing equations. This was required due to some dimensional discrepancies in their transfer function definition Eq. (5.2). $\Delta\Phi$ is an overall phase shift, measured in radians, and hence is dimensionless. However, per Eq. (5.3), Φ has dimensions of inverse length, and in fact represents a phase shift per unit length. This is due to modeling the ideal FLDI as a pair of point detectors, rather than line detectors. This means that without compensation from the $\frac{1}{2L}$ term, H_w is not dimensionless—which it should be.

However, inspection of Eq. (5.8) reveals that H_w for the infinitesimally-thin plane wave is dimensionless without needing a normalization factor. This is due to the fact that the Dirac delta function has dimensions that are the reciprocal of those of its argument, thus when using Eq. (5.8) in the ray-tracing equation, the result has dimensions of phase shift per unit length, i.e. it actually represents $\frac{\mathrm{d}(\Delta\Phi)}{\mathrm{d}z}$. This interpretation will be revisited later. With this, the numerator and denominator of H_w have equal dimensionality, and so it is dimensionless as required.

The evaluation of H_s using Eq. (5.7b) only involves $\Phi(x)$, and so yields the same result for both geometry cases:

$$H_s(k) = \frac{2}{k\Delta x} \sin\left(\frac{k\Delta x}{2}\right) \tag{5.10}$$

For $\frac{k\Delta x}{2} << 1$, $H_s \approx 1$ by the small-angle approximation. But when the disturbance wavelength gets shorter than the beam separation, then aliasing occurs whenever

$k\Delta x = 2m\pi$ for $m = \{1, 2, 3 \ldots\}$. The overall transfer function in wavenumber space is then simply given by $H = H_s H_w$.

5.1.2 Settles and Fulghum

While some final results are given in Settles and Fulghum (2016), the most relevant transfer functions for this discussion, along with more detailed derivations, can be found in Fulghum (2014). These derivations will be summarized here using the notation and terminology from the original work.

Fulghum began by considering a simple 1D model of FLDI responding to a sinusoidal disturbance—essentially a pair of point detectors. The response at each detector was given by:

$$F_{\pm} = \sin\left(2\pi f_0 \left[t \pm \frac{\Delta x}{2U_C}\right]\right) \tag{5.11}$$

where U_c and f_0 are the propagation velocity and frequency of the disturbance, respectively. The "difference signal" was then given by:

$$h_{f_0}(t) = \frac{F_+ - F_-}{\Delta x / U_c} \tag{5.12}$$

Eq. (5.12) was converted into the frequency domain:

$$\begin{aligned} H_{\Delta x}(k) &= \frac{U_c}{\pi \Delta x} \sin\left(\frac{\pi \Delta x f}{U_c}\right) \\ &= \frac{2}{\Delta x} \sin\left(\frac{k\Delta x}{2}\right) \end{aligned} \tag{5.13}$$

where the wavenumber definition $k \equiv \frac{2\pi f}{U_c}$ was used. It is unclear how the outer coefficient was derived—as the two alternate forms provided do not have consistent dimensions. Using the fact that a true derivative (which is what the ideal instrument measures) has the transfer function $H(k) = k$, Eq. (5.13) was divided by k to give a "normalized system transfer function":

$$H_{\text{sinc},\Delta x}(k) = \frac{2}{k\Delta x} \sin\left(\frac{k\Delta x}{2}\right) = \text{sinc}\left(\frac{k\Delta x}{2}\right) \tag{5.14}$$

Despite some issues with the derivation, Eq. (5.14) is identical to H_s from Eq. (5.10). Fulghum used a geometrical argument to extend Eq. (5.14) to waves propagating at an angle ϕ to the x-axis, although still confined to the x–y plane. The result was:

$$H_{\Delta x}(k, \phi) = \frac{2}{k\Delta x} \sin\left(\frac{k \cos\phi \cdot \Delta x}{2}\right) \tag{5.15}$$

We will show in Section 5.2 that Eq. (5.15) is not quite correct.

To account for the effects of finite beam width, Fulghum arrived at the same expression as Eq. (5.8) for a disturbance at a single z-location along the beam. This was then integrated in z to give the transfer function for a disturbance of finite z-extent (H_z in his notation). However, this integration was performed differently to the equivalent of Schmidt and Shepherd (Eq. (5.9)):

$$\begin{aligned} H_z &= \sqrt{\int_{-L}^{L} \left[\exp\left(-\frac{w^2(z)k^2}{8}\right)\right]^2 \mathrm{d}z} \\ &= \sqrt{2\pi^{3/2} \frac{w_0}{\lambda_L k} \operatorname{erf}\left(\frac{L\lambda_L k}{2\pi w_0}\right) \exp\left(-\frac{w_0^2 k^2}{4}\right)} \end{aligned} \tag{5.16}$$

Although Eq. (5.16) shares most of the same terms with Eq. (5.9), it is problematic because it is not dimensionless: the $\sqrt{\frac{w_0}{\lambda_L k}}$ factor means that H_z, has dimensions of $[\text{length}]^{1/2}$. Other transfer functions less relevant to this thesis were also considered, chiefly the response of FLDI to an axisymmetric turbulent jet (denoted H_σ). The lengthy expression will not be replicated here, but it shares the same issue of not being dimensionless.

5.2 Response to Single-Frequency Plane Waves

We extend the results of the previous works detailed in Section 5.1 by considering a generalized time-varying plane wave propagating in 3D space, of finite z-extent. Such a wave can be modeled by:

$$n(\underline{x}, t) = A \cos\left(\underline{k} \cdot \underline{x} - \omega t + \varphi\right) \tag{5.17}$$

Again, we combine this with Heaviside step functions as per Eq. (5.6) to bound the extent between $[z_a, z_b]$. The amplitude A, wavevector $\underline{k}$, angular frequency ω, and

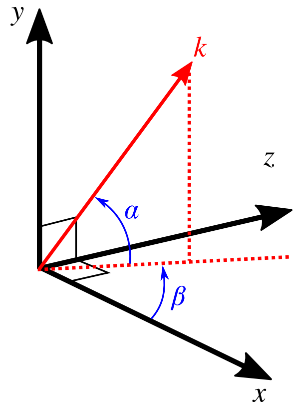

Figure 5.1: Co-ordinate system for inclined plane wave with wavevector $\underline{k}$.

phase φ are all constant. The wave has propagation velocity $\underline{a} = a\underline{\hat{a}}$, and $\underline{k} \parallel \underline{a}$. $\underline{k}$ is related to ω by the dispersion relation:

$$\underline{k} = \frac{\omega}{a}\underline{\hat{k}} \quad \Leftrightarrow \quad k = \frac{\omega}{a} \tag{5.18}$$

We establish the co-ordinate system shown in Fig. 5.1, where the wavevector direction can be described by the two angles (α, β). This is essentially a spherical co-ordinate system, although a little different to normal convention (e.g. it is a left-handed system, and the inclination is complementary to its usual definition). This system is chosen to be consistent with previous FLDI usage, i.e. x in the left-to-right flow direction, y upwards in the lab frame, and z in the beam propagation direction.

The unit vector projections are given by:

$$\underline{\hat{k}} \cdot \underline{\hat{x}} = \cos\alpha\cos\beta \quad \Leftrightarrow \quad k_x \equiv \underline{k} \cdot \underline{\hat{x}} = k\cos\alpha\cos\beta \tag{5.19a}$$

$$\underline{\hat{k}} \cdot \underline{\hat{y}} = \sin\alpha \quad \Leftrightarrow \quad k_y \equiv \underline{k} \cdot \underline{\hat{y}} = k\sin\alpha \tag{5.19b}$$

$$\underline{\hat{k}} \cdot \underline{\hat{z}} = \cos\alpha\sin\beta \quad \Leftrightarrow \quad k_z \equiv \underline{k} \cdot \underline{\hat{z}} = k\cos\alpha\sin\beta \tag{5.19c}$$

It was mentioned above that using Cartesian co-ordinates did not allow for tractable integration even for a simple wave geometry. However, we can instead use the polar co-ordinates first introduced as Eq. (4.3). Exchanging the order of integration and using the normalized intensity from Eq. (4.4), the governing ray-tracing equation can be equivalently expressed as:

$$\Delta\Phi = \frac{2}{\pi}\cdot\frac{2\pi}{\lambda_L}\int\limits_{z_a}^{z_b}\left(\int\limits_0^{\infty}\tilde{r}\exp\left(-2\tilde{r}^2\right)\left[\int\limits_0^{2\pi}n_1\mathrm{d}\theta\right]\mathrm{d}\tilde{r} - \int\limits_0^{\infty}\tilde{r}\exp\left(-2\tilde{r}^2\right)\left[\int\limits_0^{2\pi}n_2\mathrm{d}\theta\right]\mathrm{d}\tilde{r}\right)\mathrm{d}z \tag{5.20}$$

where $n_i = n_i(\tilde{r},\theta,z)$ as defined by Eq. (4.6). This form leads to a tractable integration using Eq. (5.17). The final result is:

$$\begin{aligned}\Delta\Phi(t) = {} & \frac{2\pi A}{\lambda_L}\cdot\frac{4\sqrt{2}\pi^{3/2}w_0}{\sqrt{k_x^2+k_y^2}\lambda_L}\cdot\sin\left(\frac{k_x\Delta x}{2}\right)\cdot\sin\left(\omega t-\varphi\right)\\ & \cdot\exp\left(-\frac{w_0^2}{8}\left[k_x^2+k_y^2+\frac{16\pi^2k_z^2}{\left(k_x^2+k_y^2\right)\lambda_L^2}\right]\right)\\ & \cdot\Re\left[\operatorname{erf}\left(\frac{\left(k_x^2+k_y^2\right)L\lambda_L^2+\mathrm{i}\cdot 4\pi^2k_zw_0^2}{2\sqrt{2}\pi\sqrt{k_x^2+k_y^2}\lambda_L w_0}\right)\right]\end{aligned} \tag{5.21}$$

where $\Re$ represents the real part, i is the imaginary unit, and erf is the error function. Please refer to Appendix A for a detailed derivation of Eq. (5.21). Note that Eq. (5.21) made the assumption that the wave is arranged symmetrically about the origin, i.e. $[z_a, z_b] \to [-L, +L]$. A similar result is obtained for the asymmetrical case—but it is even more unwieldy!

It is often found to be useful to abbreviate Eq. (5.21) by condensing all the wavenumber-dependent terms into a response function $h(\underline{k})$, designated as such due to its close relationship to the transfer functions $H(k)$ previously introduced. Only the original refractive index amplitude and the time-varying term are left exposed:

$$\Delta\Phi(t) = A\cdot h(\underline{k})\cdot\sin\left(\omega t-\varphi\right) \tag{5.22}$$

5.2.1 Special Cases & Comparison with Previous Works

In this thesis we prefer to work directly with $\Delta\Phi$ rather than H, because both experimental data and simulated results from the ray-tracing code are in this form. However, H is really just a normalized version of $\Delta\Phi$, and it is simple to interchange between the two. Recall that per Schmidt and Shepherd, the overall FLDI response to a wave propagating in the x-direction only is given by the product of Eqs. (5.9) and (5.10), yielding:

$$H = \frac{2}{k\Delta x}\sin\left(\frac{k\Delta x}{2}\right) \times \frac{\pi w_0\sqrt{2\pi}}{kL\lambda_L}\exp\left(-\frac{w_0^2k^2}{8}\right)\mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \tag{5.23}$$

To compare Eqs. (5.21) and (5.23), consider pure x propagation, i.e. $(\alpha,\beta) = (0°,0°) \Leftrightarrow (k_x,k_y,k_z) = (k,0,0)$. Noting that $\mathrm{erf}(x) \in \mathbb{R}\ \forall\ x \in \mathbb{R}$, and maximizing the response by using $\max\left(\sin\left(\omega t - \varphi\right)\right) = 1$, Eq. (5.21) reduces to:

$$\max|\Delta\Phi| = \frac{2\pi A}{\lambda_L}\cdot\frac{4\sqrt{2}\pi^{3/2}w_0}{k\lambda_L}\cdot\sin\left(\frac{k\Delta x}{2}\right)\cdot\exp\left(-\frac{w_0^2k^2}{8}\right)\cdot\mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \tag{5.24}$$

For conversion to H, consider the instantaneous spatial position of the wave that maximizes $(\mathrm{d}\Phi/\mathrm{d}x)_{x=0}$, which as before is $n(x) = A\sin(kx) \implies \Phi(x) = \frac{2\pi A}{\lambda_L}\sin(kx)$. Applying the definition of H, Eq. (5.2), Eq. (5.24) becomes identical to Eq. (5.23) if their $\frac{1}{2L}$ normalization is included. This shows that Schmidt and Shepherd's result is a special case recovered when the dimensionality of the wave propagation is reduced from 3D to 1D. Interestingly, the overall H is directly obtained without having to separately derive H_s and H_w as they did. This is because Eq. (5.20) incorporates both the Gaussian beam width behavior as the $\exp\left(-2\tilde{r}^2\right)$ term, and the finite beam separation included in $n_i = n(x \pm \Delta x/2, y, z)$.

We can methodically consider the remaining special cases of reduced dimensionality. The other two 1D cases are pure y- and z-propagation, where $(\alpha,\beta) = (\pm 90°, 0°) \Leftrightarrow (k_x,k_y,k_z) = (0,k,0)$, and $(\alpha,\beta) = (0°,\pm 90°) \Leftrightarrow (k_x,k_y,k_z) = (0,0,k)$, respectively. Both of these cases reduce Eq. (5.21) to $\Delta\Phi(t) = 0$. This is expected, because FLDI is only sensitive to refractive index gradients in the x-direction. Next, consider the 2D cases where $\underline{k}$ is confined to the xy, xz, or yz planes. Again, yz-propagation is found to give $\Delta\Phi(t) = 0$ because there are no x-gradients.

Disturbances propagating in the xy-plane have $\alpha \neq 0°$, $\beta = 0° \implies (k_x,k_y,k_z) = (k\cos\alpha, k\sin\alpha, 0)$:

$$\Delta\Phi(t) = \frac{2\pi A}{\lambda_L} \cdot \frac{4\sqrt{2}\pi^{3/2} w_0}{\sqrt{k_x^2 + k_y^2}\lambda_L} \cdot \sin\left(\frac{k_x \Delta x}{2}\right) \cdot \sin(\omega t - \varphi)$$
$$\cdot \exp\left(-\frac{w_0^2}{8}\left[k_x^2 + k_y^2\right]\right) \cdot \mathrm{erf}\left(\frac{\left(k_x^2 + k_y^2\right) L\lambda_L}{2\sqrt{2}\pi\sqrt{k_x^2 + k_y^2} w_0}\right) \tag{5.25}$$

Note that with $k_z = 0$, $k_x^2 + k_y^2 = k^2$, giving further simplification to:

$$\Delta\Phi(t) = \frac{2\pi A}{\lambda_L} \cdot \frac{4\sqrt{2}\pi^{3/2} w_0}{k\lambda_L} \cdot \sin\left(\frac{k\cos\alpha\Delta x}{2}\right) \cdot \sin(\omega t - \varphi)$$
$$\cdot \exp\left(-\frac{k^2 w_0^2}{8}\right) \cdot \mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \tag{5.26}$$

Conversion of Eq. (5.26) to H, using $n = A\sin(k_x x) = A\sin(k\cos\alpha \cdot x)$ gives:

$$H = \frac{2}{k\cos\alpha\Delta x}\sin\left(\frac{k\cos\alpha\Delta x}{2}\right) \times \frac{\pi w_0\sqrt{2\pi}}{kL\lambda_L}\exp\left(-\frac{w_0^2 k^2}{8}\right)\mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \tag{5.27}$$

There are two points to note in Eq. (5.27). First, only the H_s term is altered from Eq. (5.23). H_w remains the same because the spatial filtering effect is due to the relative sizes of the disturbance wavelength and local beam width; the beam cross-sections are circular and lie in the xy plane and so the spatial filtering effect is unchanged, so long as the disturbances are also confined to this plane. Secondly, the H_s term differs only in having k replaced by $k\cos\alpha$, i.e. the separation of the beams in x causes aliasing of the projected wavenumber k_x. This was predicted by Fulghum, although his H_s equivalent, Eq. (5.15), incorrectly only uses k_x inside the sin and misses it in the coefficient.

Lastly, xz-propagation has $\alpha = 0°, \beta \neq 0° \Rightarrow (k_x, k_y, k_z) = (k\cos\beta, 0, k\sin\beta)$. This case was not considered in any previous works.

$$\Delta\Phi(t) = \frac{2\pi A}{\lambda_L} \cdot \frac{4\sqrt{2}\pi^{3/2} w_0}{k_x\lambda_L} \cdot \sin\left(\frac{k_x \Delta x}{2}\right) \cdot \sin(\omega t - \varphi)$$
$$\cdot \exp\left(-\frac{w_0^2}{8}\left[k_x^2 + \frac{16\pi^2 k_z^2}{k_x^2\lambda_L^2}\right]\right) \cdot \Re\left[\mathrm{erf}\left(\frac{k_x^2 L\lambda_L^2 + \mathrm{i}\cdot 4\pi^2 k_z w_0^2}{2\sqrt{2}k_x\pi\lambda_L w_0}\right)\right] \tag{5.28}$$

Making the trigonometric substitutions and noting that $k_z/k_x = k\sin\beta/k\cos\beta = \tan\beta$ gives an alternate form in terms of β:

$$\begin{aligned}\Delta\Phi(t) = {} & \frac{2\pi A}{\lambda_L}\cdot\frac{4\sqrt{2}\pi^{3/2}w_0}{k\cos\beta\lambda_L}\cdot\sin\left(\frac{k\cos\beta\Delta x}{2}\right)\cdot\sin(\omega t-\varphi)\\ & \cdot\exp\left(-\frac{w_0^2}{8}\left[k^2\cos^2\beta+\frac{16\pi^2\tan^2\beta}{\lambda_L^2}\right]\right)\\ & \cdot\Re\left[\operatorname{erf}\left(\frac{kL\lambda_L\cos\beta}{2\sqrt{2}\pi w_0}+\mathrm{i}\frac{\sqrt{2}\pi w_0\tan\beta}{\lambda_L}\right)\right]\end{aligned} \tag{5.29}$$

5.2.2 Validation

To validate the general result Eq. (5.21) and its special cases, these analytical expressions are compared with the numerical $\Delta\Phi$ from the ray-tracing code. This code has itself been validated against experimental data (Chapter 4); furthermore it performs the numerical integrations using a different co-ordinate system, providing a useful comparison. As discussed in Section 4.1, the ray-tracing algorithm can couple with analytical input functions that define the refractive index at every point on the beam grids. Eq. (5.17) was implemented as such an input function, with the time-varying propagation performed in a quasi-steady manner. This is similar to how propagating shock waves were simulated in Section 4.3.

Validations were performed in both temporal and wavenumber spaces. For the temporal validations, arbitrary wave parameters were selected, and several wave cycles simulated. This was repeated for many different parameter sets, including all the special reduced-dimensionality geometries, as well as fully-general 3D propagation. In all cases there was excellent agreement between the analytical and numerical predictions; some of these cases are shown in Fig. 5.2.

For the wavenumber validations, max $|\Delta\Phi|$ was compared as a function of k. Recall that this maximum response is found analytically by setting $\sin(\omega t-\varphi) = 1$, so that:

$$\max|\Delta\Phi| = A\cdot h(\underline{k}) \tag{5.30}$$

Numerically, we can just set $\omega t = 90°$ and $\varphi = 0°$ so that the wave is frozen at the position of maximum response. For a given propagation direction defined by $\{\alpha,\beta\}$, this procedure is repeated for many k. The results are shown in Fig. 5.3, for the three reduced-dimensionality cases. Again, the agreement is very close, except at very high k where the disturbances begin to get small compared with the grid spacing of the numerical scheme, as was also observed by Schmidt and Shepherd. Note for reference that in Figs. 5.3b and 5.3c the blue curves are identical to the curve of Fig. 5.3a, as expected.

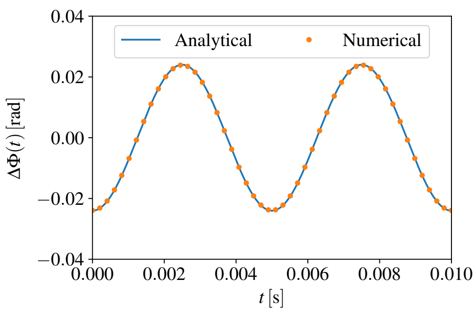

(a) $[x]$: $\alpha = 0°$, $\beta = 0°$, $\varphi = 90°$, $A = 1 \times 10^{-5}$, $L = 0.045\,\mathrm{m}$, $\lambda_w = 0.5\,\mathrm{m}$, $c = 100\,\mathrm{m\,s^{-1}}$

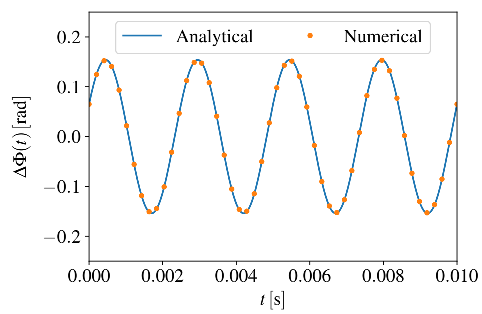

(b) $[xy]$: $\alpha = 30°$, $\beta = 0°$, $\varphi = -25°$, $A = 5 \times 10^{-5}$, $L = 0.04\,\mathrm{m}$, $\lambda_w = 0.3\,\mathrm{m}$, $c = 120\,\mathrm{m\,s^{-1}}$

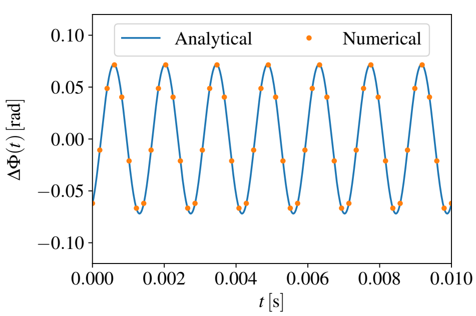

(c) $[xz]$: $\alpha = 0°$, $\beta = 15°$, $\varphi = 60°$, $A = 5 \times 10^{-6}$, $L = 0.07\,\mathrm{m}$, $\lambda_w = 0.1\,\mathrm{m}$, $c = 70\,\mathrm{m\,s^{-1}}$

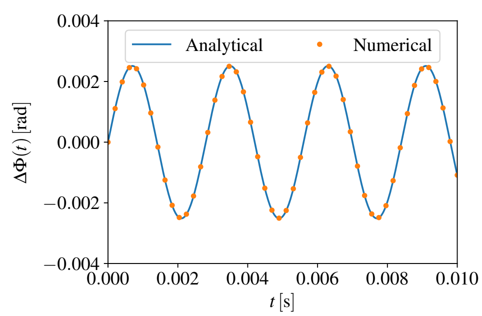

(d) $[xyz]$: $\alpha = 45°$, $\beta = 20°$, $\varphi = 0°$, $A = 1 \times 10^{-6}$, $L = 0.1\,\mathrm{m}$, $\lambda_w = 0.7\,\mathrm{m}$, $c = 250\,\mathrm{m\,s^{-1}}$

Figure 5.2: Temporal validation of single-frequency plane wave response. Selected cases shown for various propagation directions and wave parameters. λ_w and c are the wavelength and speed of the plane wave, respectively. The standard optical parameters are used, as given in Section 4.1.

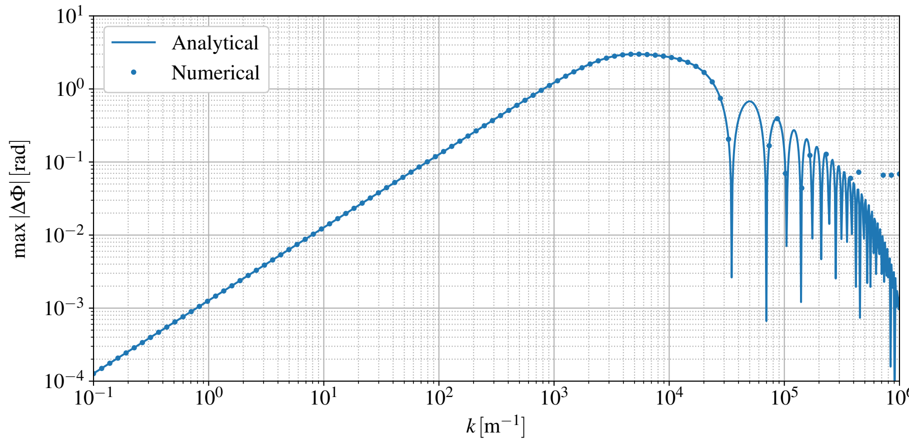

(a) $[x]$: $\alpha = 0°$, $\beta = 0°$.

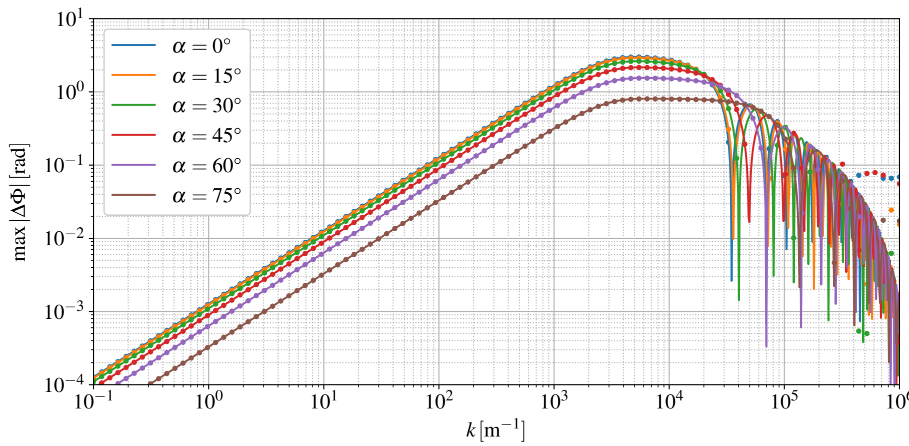

(b) $[xy]$: $\alpha \neq 0°$, $\beta = 0°$.

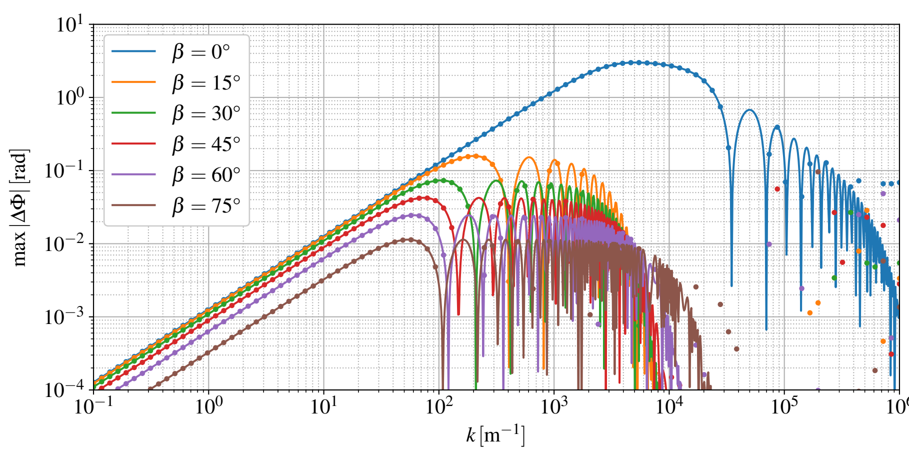

(c) $[xz]$: $\alpha = 0°$, $\beta \neq 0°$.

Figure 5.3: Wavenumber validation of single-frequency plane wave response. Reduced-dimensionality cases shown, for fixed $A = 10^{-5}$, $L = 0.06$ m, and standard optical parameters. Analytical solutions shown with lines, numerical simulations with markers.

With these results, Eq. (5.21) is considered validated. The implications and physical meaning behind the curves of Fig. 5.3 will be discussed shortly, but first there is an issue with the evaluation of Eq. (5.21) that needs to be discussed. $h(\underline{k})$ can be separated into four terms:

$$h(\underline{k}) = P(\underline{k})Q(\underline{k})R(\underline{k})S(\underline{k}) \tag{5.31}$$

where:

$$P(\underline{k}) = \frac{2\pi}{\lambda_L} \cdot \frac{4\sqrt{2}\pi^{3/2} w_0}{\sqrt{k_x^2 + k_y^2}\lambda_L} \tag{5.32a}$$

$$Q(\underline{k}) = \sin\left(\frac{k_x \Delta x}{2}\right) \tag{5.32b}$$

$$R(\underline{k}) = \exp\left(-\frac{w_0^2}{8}\left[k_x^2 + k_y^2 + \frac{16\pi^2 k_z^2}{\left(k_x^2 + k_y^2\right)\lambda_L^2}\right]\right) \tag{5.32c}$$

$$S(\underline{k}) = \Re\left[\operatorname{erf}\left(\frac{\left(k_x^2 + k_y^2\right) L\lambda_L^2 + \mathrm{i} \cdot 4\pi^2 k_z w_0^2}{2\sqrt{2}\pi\sqrt{k_x^2 + k_y^2}\lambda_L w_0}\right)\right] \tag{5.32d}$$

$P(\underline{k}) \propto k^{-1}$ and $-1 \leq Q(\underline{k}) \leq 1$. For $S(\underline{k})$, the argument of erf is real when $\beta = 0°$, and $-1 < \operatorname{erf}(x) < 1 \;\forall\; x \in \mathbb{R}$. However, an issue arises when $\beta \neq 0°$ (i.e. when there is a component of propagation along the optical axis). Complex-valued $\operatorname{erf}(z)$ is not bounded by unity as in the real case, and both $\Re[\operatorname{erf}(z)]$ and $\Im[\operatorname{erf}(z)]$ increase rapidly and without bound, in an oscillatory manner. For typical optical and wave parameters of interest to us, $S(\underline{k})$ grows so rapidly that arithmetic overflow occurs by about $\beta = 30°$, i.e. it exceeds the largest number that can be represented using standard 64-bit floating-point format (about 10^{308}). However, the corresponding output of the numerical simulation (which does not rely on evaluating the functions of Eq. (5.31)), is well-bounded, implying that the huge $S(\underline{k})$ term is counterbalanced by a similarly tiny term. This can only be $R(\underline{k})$, and indeed this term experiences arithmetic underflow as a 64-bit float. For $\alpha = 0°$, $R(\underline{k})$ can be alternatively expressed (see Eq. (5.29)):

$$R(\underline{k}) = \exp\left(-\frac{w_0^2}{8}\left[k^2\cos^2\beta + \frac{16\pi^2\tan^2\beta}{\lambda_L^2}\right]\right) \tag{5.33}$$

Because $\lambda_L \approx 10^{-7}$ m, the $R(\underline{k}) \sim \exp\left([\tan\beta/\lambda_L]^2\right)$ scaling is very strong, especially for $\beta > 45°$ where $\tan\beta > 1$. To avoid these evaluation issues, it would be ideal to find an analytical expression for the product $R(\underline{k})S(\underline{k})$ that is more well-behaved. This is difficult in the general case because R and S have different arguments, although it can be done in some limiting cases. Series and asymptotic approximations were also attempted but these could not be combined into bounded forms. A more brute-force method was used instead: the Wolfram language of Mathematica (Wolfram Research, Inc., 2019) allows for arbitrary-precision arithmetic, where numbers are not limited to a fixed bit width. This capability was incorporated into existing Python functions by using the Wolfram Client Library, which allows direct execution of Wolfram code from Python by connecting to a local kernel (e.g. if Mathematica is also installed). The analytical solutions shown in Fig. 5.3c were computed using this technique.

5.2.3 Implications

Comparison of Figs. 5.3a and 5.3b shows that waves inclined in the xy plane behave fundamentally the same as those restricted to propagation in x alone. Even for quite large inclinations α, the response curves are merely shifted down slightly (although they will eventually collapse to zero for all k when $\alpha \to 90°$). The high-k oscillatory roll-off is also similar, although the location of the zeroes due to the aliasing effect are dependent on α. These effects were interpreted physically by Settles and Fulghum: the instrument is simply responding to the component of the wave projected in the sensitive x direction.

On the other hand, Fig. 5.3c shows that waves inclined in the xz plane have quite different behavior. Even for mild inclinations β, the response curves shift drastically. Most noticeable is that the oscillatory roll-off begins more than two decades earlier in k. This will have important consequences for the interpretation of $\Delta\Phi$ data later. It is much more difficult to assign physical explanations to this oscillatory and roll-off behavior, as the differences in the xz case (as compared to the x or xy cases) come from the additional k_z-dependent terms in the arguments of $R(\underline{k})$ and $S(\underline{k})$ which do not have obvious physical interpretations—unlike, for example, the filtering effect of $\exp\left(-\frac{k^2w^2}{8}\right)$ or the aliasing effect of $\sin\left(\frac{k\Delta x}{2}\right)$. In fact, although superficially similar, the oscillatory region in Fig. 5.3c is not due to Δx aliasing at all: from Eqs. (5.19a) and (5.32a) aliasing should happen when $k_x\Delta x = k\cos\beta \cdot \Delta x = 2m\pi$ for integer m. This implies that the oscillatory region should shift to higher k as β increases (which is what happens with the analogous α behavior). Instead, these oscillations are not due to $Q(\underline{k})$ at all, but rather $S(\underline{k})$.

The source of the oscillations, along with other trends, can be seen by plotting each term in Eqs. (5.32a) to (5.32d) separately in Fig. 5.4. $P(\underline{k}) \sim k^{-1}$ for all k, while $Q(\underline{k}) \sim k$ for low k, with an oscillatory region beginning around $k = \pi/\Delta x$. $R(\underline{k}) \approx 1$ for low k, then rapidly rolls off. $\mathrm{erf}(x) \sim x$ for $x << 1$ so when $\beta = 0°$, $S(\underline{k}) \sim k$ for low k, then asymptotes to unity. To summarize for the $\beta = 0°$ case, the high-k roll-off is governed by R, while the low-k roll-off is governed by S. The behavior of P and Q is unchanged between Figs. 5.4a and 5.4b, but the effect on S of making $\beta \neq 0°$ is drastic. R actually retains much the same shape (flat then a roll-off at high k) but shifted far down in magnitude. The overall response function in Fig. 5.4b shows two distinct oscillatory regions: the first caused by S which also incorporates a roll-off, the second unchanged and due to Q with roll-off from R.

Both oscillating regions can cause issues in the inversion process (to be introduced subsequently). Non-oscillating envelopes can be defined:

$$Q_e(\underline{k}) = \begin{cases} \sin\left(\frac{k_x \Delta x}{2}\right) & k_x \Delta x \leq \pi \\ 1 & k_x \Delta x > \pi \end{cases} \tag{5.34}$$

$$S_e(\underline{k}) = \begin{cases} \Re\left[\mathrm{erf}(\Xi)\right] & k \leq k_{c1} \\ |\mathrm{erf}(\Xi)| & k > k_{c1} \end{cases} \tag{5.35}$$

In Eq. (5.35), Ξ represents the same argument as in Eq. (5.32d). The critical wavenumber k_{c1} where the piecewise function changes over has to be found numerically, as it is not easy to find analytically, unlike for Eq. (5.34). Examples of both these envelopes are illustrated in Fig. 5.5.

Putting aside for the moment the complexities of oscillation, we next address the low-k roll-off behavior. As mentioned, this is governed by the S term, which has the following simplified form when $\beta = 0°$:

$$S(\underline{k}) = \mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \tag{5.36}$$

For fixed optical parameters λ_L and w_0, the roll-off location depends on L (Fig. 5.6). Recall that this represents the extent of the disturbance field in the z-direction, and

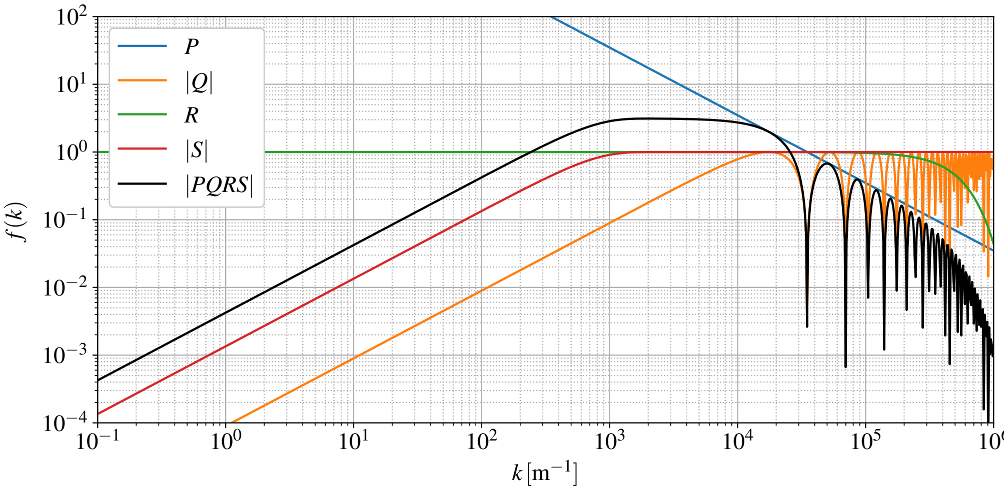

(a) $[x]$: $\alpha = 0°$, $\beta = 0°$.

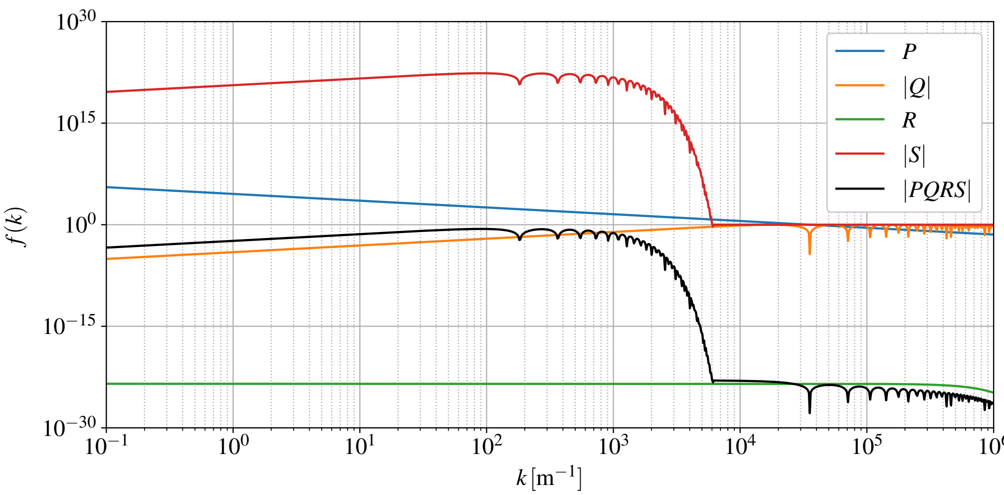

(b) $[xz]$: $\alpha = 0°$, $\beta = 10°$.

Figure 5.4: Contributing components for response function $h(\underline{k})$.

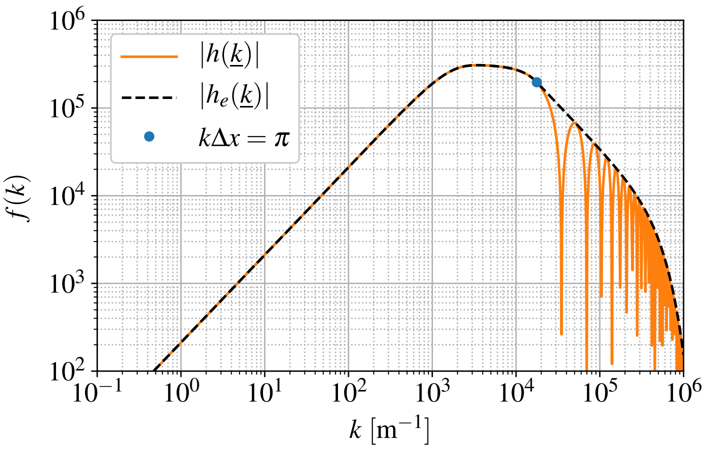

(a) $\beta = 0°$, envelope governed by $Q_e(\underline{k})$.

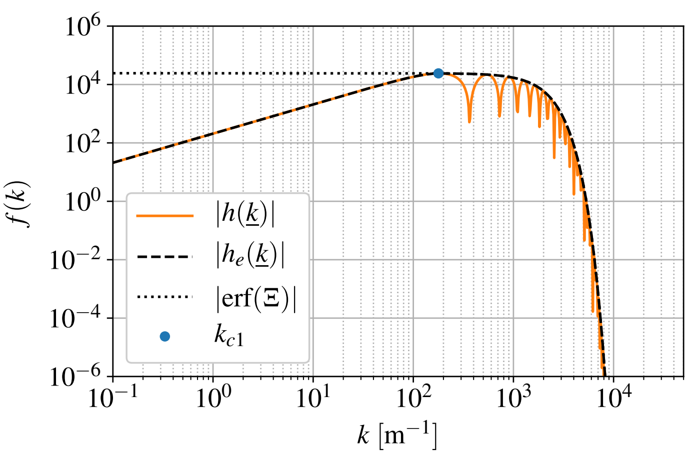

(b) $\beta \neq 0°$, envelope governed by $S_e(\underline{k})$.

Figure 5.5: Examples of envelope functions $h_e(\underline{k})$.

that it is assumed that the beams themselves are at least as long as L (they can be any value $> L$ because there are no contributions to the signal beyond this disturbance field). In the limiting case $L \to \infty$, $S \to 1 \; \forall \; k$, in other words, the erf function only manifests due to finite integration limits (note it is also absent in the other limit of an infinitesimally thin wave, e.g. Eq. (5.8)).

The physical explanation for this is as follows: a disturbance gets spatially averaged out when its wavelength λ_w becomes small compared to the local beam width $w(z)$. The sensitive length (denoted here as ζ) of the FLDI is the distance about the focal plane where a disturbance contributes significantly to the signal (e.g. one might

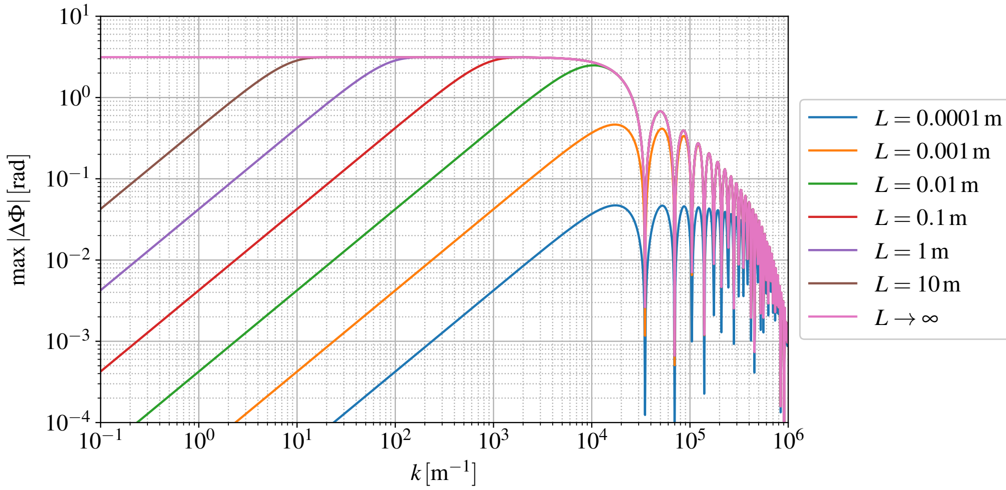

Figure 5.6: Influence of L on the low-k roll-off location.

arbitrarily define ζ as the length within which 99% of the total signal is generated—see more in Section 5.7). It follows, as pointed out by Schmidt and Shepherd, that ζ increases with λ_w: one must look further out from the focal plane to find a sufficiently large $w(z)$ to cause filtering. Hence, while $L < \zeta$ there is still "capacity" for $\Delta\Phi$ to increase as L increases. Once L exceeds ζ, there is negligible further increase in $\Delta\Phi$ due to spatial filtering in regions of the beam beyond $z = \zeta$. As L increases, ζ is exceeded for larger and larger λ_w (i.e. smaller k) and so the roll-off point shifts to lower k; hence in the limit $L \to \infty$ the maximum possible $\Delta\Phi$ is achieved at all k. The reason for this maximum $\Delta\Phi$ being constant for all k (i.e. why the pre-roll-off curve is flat) is that although larger-λ_w disturbances influence more of the beam, their signal contribution per unit length is smaller, because FLDI responds to refractive index gradients—and longer disturbances have shallower gradients. Referring back to Fig. 5.4a, this balance is encapsulated in the product PQ, which is independent of k.

Related to this topic is the discussion of how Schmidt and Shepherd's infinitesimally-thin wave solution fits in with the results given so far. Recall from Section 5.1 that due to dimensional considerations, their result (converted back from H) is really a phase shift per unit length:

$$\frac{\mathrm{d}(\Delta\Phi)}{\mathrm{d}z} = \frac{2\pi A}{\lambda_L} \cdot 2\sin\left(\frac{k\Delta x}{2}\right)\exp\left(-\frac{w^2k^2}{8}\right) \tag{5.37}$$

Again, $w = w(z_a)$ is evaluated at the z-location of the thin wave. For finite-extent waves of small thickness Δz centered about $z = z_a$:

$$\Delta\Phi \approx \frac{\mathrm{d}(\Delta\Phi)}{\mathrm{d}z}\Delta z \tag{5.38}$$

If the wave is symmetrical about the focal plane, then $\Delta z = 2L$ and $w = w_0$. We can then compare the approximate response of Eq. (5.38) with the exact solution. This is shown in Fig. 5.7, and as expected for small L there is a close match. For disturbance fields that extend significantly beyond the vicinity of the focal plane (where the infinitesimal solution is evaluated) there are larger discrepancies.

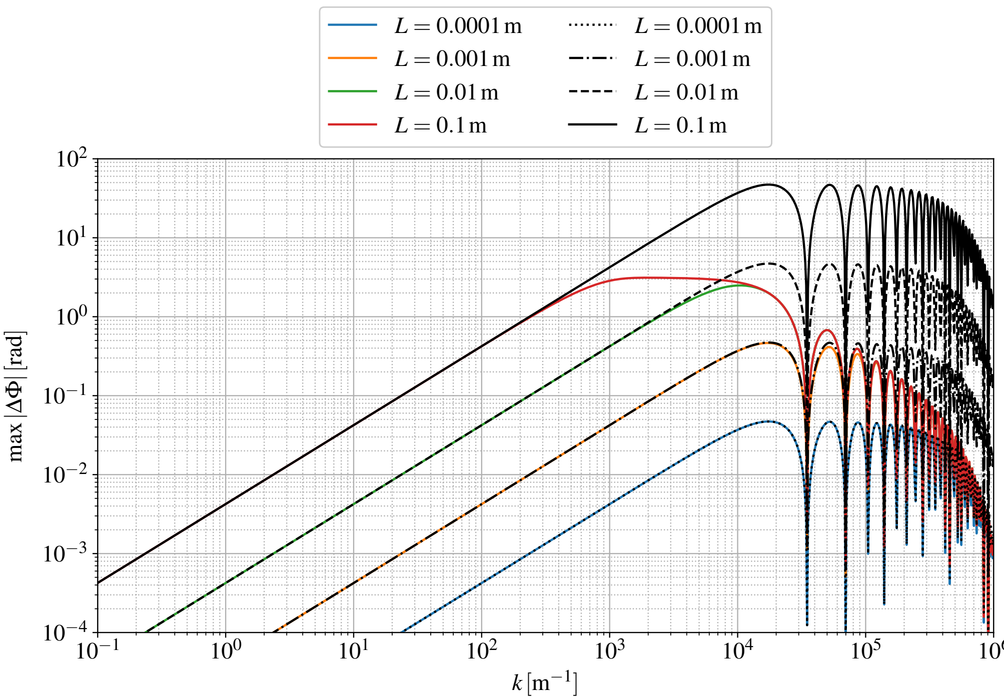

Figure 5.7: Comparison of exact response (solid colored lines) and approximate response using local infinitesimal solution (black patterned lines).

5.3 Response to Multiple-Frequency Plane Waves

Section 5.2 validated an analytical function for the response of FLDI to a propagating plane wave, of arbitrary orientation and a single frequency. Here this result is extended to waves with multiple frequency components. This development is done in two stages: at first all frequency components will still be confined to the same propagation direction (i.e. all sharing the same unit wavevector $\underline{\hat{k}}$), followed by relaxing this condition and allowing for superimposed waves each propagating arbitrarily.

The key to these developments is that Eq. (4.2) is linear, suggesting that composite signals can be built up through linear combinations of single-frequency results.

5.3.1 Multiple-Frequency, Single-Direction

Consider a refractive index wave built from N discrete frequency components:

$$n(\underline{x}, t) = \frac{1}{\sqrt{N}} \sum_{j=1}^{N} A_j \cos\left(\underline{k}_j \cdot \underline{x} - 2\pi f_j t + \varphi_j\right) \tag{5.39}$$

The frequency components f_j are selected within some bandwidth $[f_a, f_b]$, and these determine k_j via the dispersion relation Eq. (5.18). $\hat{\underline{k}}$ is fixed, and determined from the given $\{\alpha, \beta\}$. The amplitudes are sampled from some continuous function, $A_j = A(f_j)$, requiring a $\frac{1}{\sqrt{N}}$ scaling to keep a constant RMS as N increases for fixed bandwidth. Finally, the phases φ_j are drawn randomly from the uniform distribution $[0, 2\pi)$.

Linearity yields:

$$\Delta\Phi(t) = \frac{1}{\sqrt{N}} \sum_{j=1}^{N} A_j \cdot h\left(\underline{k}_j\right) \cdot \sin\left(2\pi f_j t - \varphi_j\right) \tag{5.40}$$

Similarly to the single-frequency results in Fig. 5.2, Eq. (5.40) was validated against the numerical scheme for many sets of parameters. The computational load is exacerbated for these multiple-frequency wave simulations. The refractive index function $n(\underline{x})$ must be evaluated at every beam grid point, at every timestep. The number of operations scales with N for Eq. (5.39). Furthermore, the total simulation period and timestep size have to be chosen in order to adequately resolve the lowest and highest frequency components, respectively. This leads to a large required number of timesteps. To speed things up, most validation cases were carried out with $n_\theta = 64$ instead of the usual 256, causing a resolution reduction of approximately 16×. However for all cases trialed there was still very close agreement in the analytical and numerical $\Delta\Phi(t)$ responses (for an example, see Fig. 5.9), and Eq. (5.40) is considered validated.

5.3.2 Multiple-Frequency, Multiple-Direction

Multiple instances of Eq. (5.39) can in turn be superimposed with different $\hat{\underline{k}}$:

$$n(\underline{x},t) = \frac{1}{\sqrt{N}} \sum_{i=1}^{M} \sum_{j=1}^{N} A_j \cos\left(\underline{k}_{i,j} \cdot \underline{x} - 2\pi f_j t + \varphi_{i,j}\right) \tag{5.41}$$

where $\underline{k}_{i,j} = k_j \hat{\underline{k}}_i$. Note that Eq. (5.41) is not the most general form: A_j and f_j are shared between each wave. It is done this way because it will be the most relevant for our later applications, where facility centerline noise is the sum of acoustic radiation from the four walls of a square test section. There, it is assumed that the noise from each wall is drawn from the same spectrum, i.e. $A(f)$ and the bandwidth are the same. However, the phases should remain uncorrelated. We do not normalize by M as done with N. This is because the summation over N discrete frequency components is meant to approximate an integral over a continuous spectrum of frequencies—and thus should not change RMS for different frequency sampling densities, whereas the summation over wavevector orientations is used to represent the superposition of genuinely discrete fields (e.g. the aforementioned radiating walls where $M = 4$).

Eq. (5.41) was implemented in the numerical scheme (an example of this is illustrated in Fig. 5.8), with the simulated result compared to the analytical expression. Again, very close agreement was observed, validating that the principle of superposition continues to hold for all combinations of frequency and propagation direction. We now have the ability to analytically calculate $\Delta\Phi(t)$ from an arbitrarily complicated $n(\underline{x},t)$ so long as it is composed of planar sinusoidal waves. However (as discussed in the introduction to this chapter) in real situations $\Delta\Phi(t)$ is known, and we wish to recover $n(\underline{x},t)$, at least partially. This is the topic of the next section.

5.4 Spectral Inversion

5.4.1 Continuous Spectral Representation

To facilitate a spectral treatment of the FLDI response, the single-component plane wave equation Eq. (5.17) should be changed to its complex representation:

$$n(\underline{x},t) = A \exp\left[\mathrm{i}\left(\underline{k} \cdot \underline{x} - \omega t + \varphi\right)\right] \tag{5.42}$$

where now $n(\underline{x},t) \in \mathbb{C}$, though it is still assumed that $\{A, \omega, t, \varphi\} \in \mathbb{R}$ and $\{\underline{k}, \underline{x}\} \in \mathbb{R}^3$. Applying the linearized ray-tracing procedure to Eq. (5.42) using the same transformations and integrations as done in Section 5.2, a very similar solution is obtained:

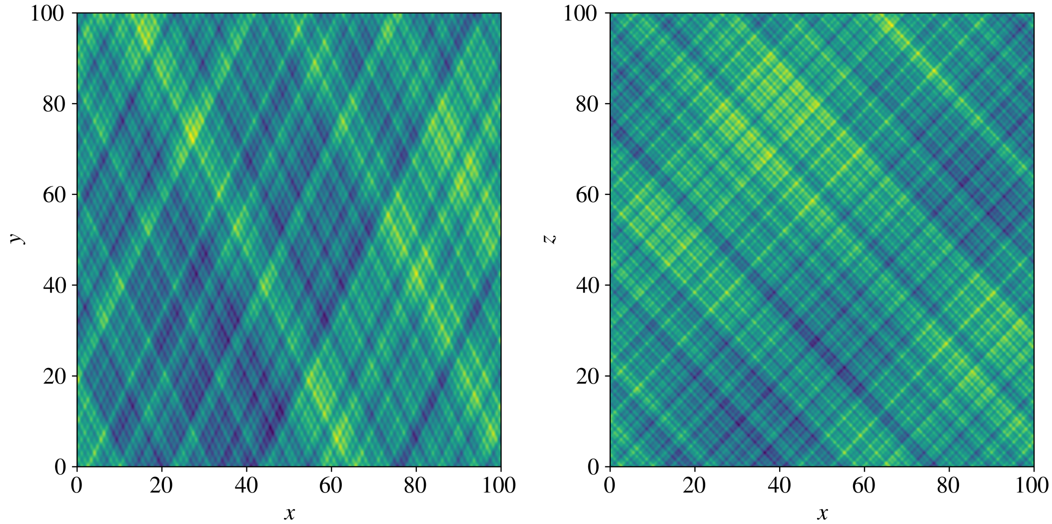

Figure 5.8: Superposition of four plane waves with different propagation directions given by $\alpha_j = \{-20°, 20°, 0°, 0°\}$, $\beta_j = \{0°, 0°, -45°, 45°\}$. xy and xz cross-sections of a 3D box, shown at a single time instant. Colors represent refractive index fluctuations (magnitudes and dimensions arbitrary).

$$\Delta\Phi(t) = A \cdot h(\underline{k}) \cdot \mathrm{i} \exp\left[\mathrm{i}\left(-\omega t + \varphi\right)\right] \tag{5.43}$$

The wavenumber-dependent response function $h(\underline{k})$ is the same as before; the only difference is in the temporal term. This term can be expanded:

$$\begin{aligned}\mathrm{i} \exp\left[\mathrm{i}\left(-\omega t + \varphi\right)\right] &= \mathrm{i}\left[\cos\left(-\omega t + \varphi\right) + \mathrm{i} \sin\left(-\omega t + \varphi\right)\right] \\ &= \sin\left(\omega t - \varphi\right) + \mathrm{i} \cos\left(\omega t - \varphi\right) \end{aligned} \tag{5.44}$$

Note then that the real part of Eq. (5.43) is equal to the real derivation Eq. (5.22), so the two approaches are consistent.

Now consider a composite plane wave with a continuous frequency spectrum:

$$n(\underline{x}, t) = \int_{-\infty}^{\infty} A(\omega) \exp\left[\mathrm{i}\left(\underline{k} \cdot \underline{x} - \omega t + \varphi\right)\right] \mathrm{d}\omega \tag{5.45}$$

Note that as before, $\underline{k} = \underline{k}(\omega)$; now also $\varphi = \varphi(\omega)$. This notation is omitted for clarity. $A(\omega)$ is now the amplitude spectral density with dimensions of $[\omega]^{-1}$. Eq. (5.45) can be substituted into the ray-tracing integrals, yielding by linearity:

$$\Delta\Phi(t) = \int_{-\infty}^{\infty} A(\omega) \cdot h(\underline{k}) \cdot \mathrm{i}\mathrm{e}^{\mathrm{i}\varphi} \cdot \mathrm{e}^{-\mathrm{i}\omega t}\mathrm{d}\omega \tag{5.46}$$

The general Fourier transform pair is (Osgood, 2007):

$$\hat{g}(s) = \frac{1}{a}\int_{-\infty}^{\infty} g(t)\mathrm{e}^{+\mathrm{i}bst}\mathrm{d}t \tag{5.47a}$$

$$g(t) = \int_{-\infty}^{\infty} \hat{g}(s)\mathrm{e}^{-\mathrm{i}bst}\mathrm{d}s \tag{5.47b}$$

where possible sign and normalizations conventions include:

$$a = \sqrt{2\pi} \quad b = \pm 1 \tag{5.48a}$$

$$a = 1 \quad b = \pm 2\pi \tag{5.48b}$$

$$a = 1 \quad b = \pm 1 \tag{5.48c}$$

So far we have been using f and $\omega = 2\pi f$ interchangeably for the temporal frequency. Most signal-processing libraries deal with f; accordingly, we choose $b = \pm 2\pi$ (and thus $a = 1$) so that the generalized frequency s represents f rather than ω. To choose the sign convention, inspect Eq. (5.46) to see we want $\beta = +1$ to give the following Fourier pair:

$$\hat{g}(f) = \int_{-\infty}^{\infty} g(t)\mathrm{e}^{+\mathrm{i}2\pi f t}\mathrm{d}t \tag{5.49a}$$

$$g(t) = \int_{-\infty}^{\infty} \hat{g}(f)\mathrm{e}^{-\mathrm{i}2\pi f t}\mathrm{d}f \tag{5.49b}$$

Comparing Eqs. (5.46) and (5.49b), the Fourier transform of $\Delta\Phi(t)$ is:

$$\widehat{\Delta\Phi}(f) = A(f) \cdot h(\underline{k}) \cdot \mathrm{i}\mathrm{e}^{\mathrm{i}\varphi} \tag{5.50}$$

We can simplify the phase:

$$\mathrm{i}\mathrm{e}^{\mathrm{i}\varphi} = \mathrm{e}^{\mathrm{i}\pi/2}\mathrm{e}^{\mathrm{i}\varphi} = \mathrm{e}^{\mathrm{i}(\varphi+\pi/2)} = \mathrm{e}^{\mathrm{i}\varphi'} \tag{5.51}$$

Note that by definition $A(f), h(\underline{k}) \in \mathbb{R}$ and $|\mathrm{e}^{\mathrm{i}\varphi'}| = 1$, so $|A(f)\cdot h(\underline{k})|$ is the magnitude and $\varphi'(f)$ is the phase of the complex Fourier transform. If the wave orientation is known *a priori*, and if $\Delta\Phi(t)$ and its Fourier transform $\widehat{\Delta\Phi}(f)$ were continuous, then we could in principle solve for both $A(f)$ and $\varphi(f)$, and could in fact reconstruct the original wave. However, real experimentally-obtained $\Delta\Phi(t)$ data are discrete, and the corresponding discrete Fourier transform (DFT) will only be known at discrete values of f. With sufficient sampling rate, this still allows for estimated reconstructions of the original temporal signal. Usually we are most interested in the spectral magnitude if working under the assumption that the phases are random and uncorrelated.

If we can obtain an estimate of $|\widehat{\Delta\Phi}(f)|$, and if $h(\underline{k})$ is known, then:

$$A(f) = \left|\frac{\widehat{\Delta\Phi}(f)}{h(\underline{k})}\right| \tag{5.52}$$

Knowing $h(\underline{k})$ means knowing the relationship $\underline{k}(f)$, which requires α, β, and c to be given, i.e. the direction and speed of the wave propagation. $A(f)$ is the spectrum of the refractive index fluctuation amplitudes, which can be used to obtain the desired density fluctuation spectrum using the Gladstone-Dale constant. To find an estimate of $|\widehat{\Delta\Phi}(f)|$, we turn to power spectral methods.

5.4.2 Power Spectra

The power spectral density (PSD), is formally defined for a signal $g(t)$ as follows (Goodman, 2000):

$$PSD\left\{g(t)\right\} = \lim_{T\to\infty}\frac{1}{T}|\hat{g}_T(f)|^2 \tag{5.53}$$

where $g_T(t)$ is the windowed function:

$$g_T(t) = \begin{cases} g(t) & -\frac{T}{2} \le t < \frac{T}{2} \\ 0 & \text{otherwise} \end{cases} \tag{5.54}$$

The PSD has units of (generalized) power per unit frequency: $[PSD\{\Delta\Phi(t)\}] = \text{rad}^2\,\text{Hz}^{-1}$. For finite-length discrete signals, the true power spectral density can only be estimated. A widely-used algorithm for estimating the PSD is Welch's method (Welch, 1967), which has implementations in many programming languages; here the function `scipy.signal.welch` from the SciPy package for Python will be used (Virtanen et al., 2020). The benefit of using Welch's method over directly computing the DFT is that it reduces the variance in the estimate by overlapping and averaging segments of the signal, although this comes at the cost of frequency resolution. Modern implementations allow for the choice of many different window functions to mitigate truncation effects, each with its own trade-offs.

Welch's method takes a time-series evenly spaced by $\Delta t = 1/f_s$, where f_s is the sampling frequency. It returns a spectrum at discrete frequencies on the interval $[0, f_{Nyq}]$ evenly spaced by Δf, where $f_{Nyq} = f_s/2$ is the Nyquist frequency, and:

$$\Delta f = \frac{f_s}{n_{ps}} \tag{5.55}$$

where n_{ps} is the number of samples per segment. The input parameters are the number of segments into which the full signal is split n_s, the degree of overlap, and the choice of window function. For a fixed signal length, increasing n_s reduces variance, but also reduces n_{ps} and hence Δf; this is the aforementioned trade-off. The optimal degree of overlap depends on the window; the default is 50% overlap with a Hann window.

The SciPy implementation can return either the PSD (here with units $\text{rad}^2\,\text{Hz}^{-1}$) or the power spectrum (PS) which has units rad^2, which is related to the PSD by:

$$PS = PSD \times \Delta f \times C_w \tag{5.56}$$

where C_w is a constant that depends on the window choice (e.g. $C_w = 1$ for a boxcar window—equivalent to no window—and $C_w = 1.5$ for a Hann window). The PS in turn can give the linear spectrum (LS) (Heinzel et al., 2002):

$$LS = \sqrt{PS} \tag{5.57}$$

The LS is so named because it is a direct estimate of $|\hat{g}(f)|$, having the same dimensions as the input signal, here simply $[\Delta\Phi(t)] = \text{rad}$. The PSD, PS, and LS are nothing but three representations of the same spectral data, and will be used interchangeably in the following depending on which is the most relevant.

5.4.3 Validation

In order to validate the inversion approach suggested by Eq. (5.52), synthetic refractive index signals $n(\underline{x}, t)$ are generated per Eq. (5.39). These are first evaluated at the single fixed point $\underline{x} = 0$, i.e. the FLDI origin, to give $n(t)$ as seen by a spatially-fixed observer. The full wave $n(\underline{x}, t)$ is then used to compute the FLDI response $\Delta\Phi(t)$, either by numerical simulation or the previously-validated analytical solution Eq. (5.40). Welch's method is applied to both the "input" $n(t)$ and the "output" $\Delta\Phi(t)$ time-series, using the same parameters, to give input and output spectra, $PS\{n(t)\}$ and $PS\{\Delta\Phi(t)\}$, respectively.

Referring to Eq. (5.39), each frequency component has magnitude $A_j/\sqrt{N}$. Thus $LS\{n(t)\}$ should comprise peaks of these magnitudes $A_j/\sqrt{N}$ at each $f = f_j$, and be zero elsewhere (i.e. at $f \neq f_j$). However, due to spectral leakage, these peaks will broaden in the actual implementation with finite Δf. With adequate resolution, this implies that an approximation to the continuous amplitude spectrum is given by:

$$A(f) \approx LS\{n(t)\} \times \sqrt{N} = \sqrt{PS\{n(t)\} \times N} \tag{5.58}$$

With *a priori* knowledge of α, β, and c, the continuous function $h(\underline{k})$ is fully-defined and can be evaluated discretely at $\underline{k}_j \Leftrightarrow f_j$. With this, Eq. (5.52) shows how to convert between input and output spectra:

$$PS\{n(t)\} = \frac{PS\{\Delta\Phi(t)\}}{|h(\underline{k}_j)|^2} \tag{5.59}$$

So the validation comprises two steps: recovering an approximation of the continuous amplitude function $A(f)$ from the discrete input spectrum, using Eq. (5.58), and converting the output spectrum (the FLDI response in terms of phase shift) back to the input spectrum (the wave in terms of refractive index), using Eq. (5.59). For the synthetic spectra, the following amplitude spectrum will be used, loosely based off Prasad (2017):

$$A(f) = \frac{B}{\sqrt{\left(1 + \frac{f^2}{F^2}\right)^r}} \tag{5.60}$$

which gives a flat spectrum at low f with constant magnitude B, that rolls off beginning at $f = F$ with decay slope controlled by r. It is not claimed that Eq. (5.60) is an accurate model of actual acoustic noise; however by choosing a bandwidth containing F, the spectrum will contain a wide range of magnitudes with rapid decay at higher f, which should be a good test of the inversion processes.

Fig. 5.9 shows the first such validation study. Here, the input wave has $N = 10$ frequency components, logarithmically spaced between $[5 \times 10^3, 5 \times 10^5]$ Hz. The low-f amplitude $B = 10^{-6}$, rolling off at $F = 10^5$ Hz with $r = 3$. For now there is only one propagation direction ($M = 1$, $\alpha = \beta = 0°$). The total period of the input $n(t)$ is shown at upper left, with a small segment magnified to show the high-frequency detail at center left. The corresponding plots in the right-hand column show the output $\Delta\Phi(t)$, both simulated with the ray-tracing code, and calculated analytically using Eq. (5.40). In the spectral plot at lower left, the input $A(f)$ from Eq. (5.60) is shown in blue. The discrete reconstruction of $A(f)$ using Eq. (5.58) is in orange. The 10 discrete frequency components are clearly visible, and as previously discussed show evidence of spectral broadening, with a non-zero spectral noise floor between each peak. The peaks closely follow $A(f)$ as required. At lower right, the output spectra $PS\{\Delta\Phi(t)\}$ are shown, computed from both the simulated and calculated time-series, using the same Welch parameters as each other and as the input spectrum. It is interesting to note that despite very little discernable difference between the simulated and calculated $\Delta\Phi(t)$, the spectrum of the simulated data shows a much higher and messier noise floor, although all the peaks are correctly replicated. This is due to additional error introduced with the simulation discretization scheme. This gives another reason to prefer the use of calculated over simulated $\Delta\Phi(t)$ where possible. Also plotted here is the response function $h(\underline{k})$, mapped to its corresponding values of f, along with a non-oscillatory version $h_e(\underline{k})$, the significance of which will be returned to later. Finally, in the lower left plot, the green curve shows the results of the inversion process performed by combining Eqs. (5.58) and (5.59). This maps the output spectrum back to a very close match with the input spectrum, as desired.

The next validation study, in Fig. 5.10, is more complex. It has $M = 4$ superimposed propagation directions, each with the same $N = 1000$ frequency components. For

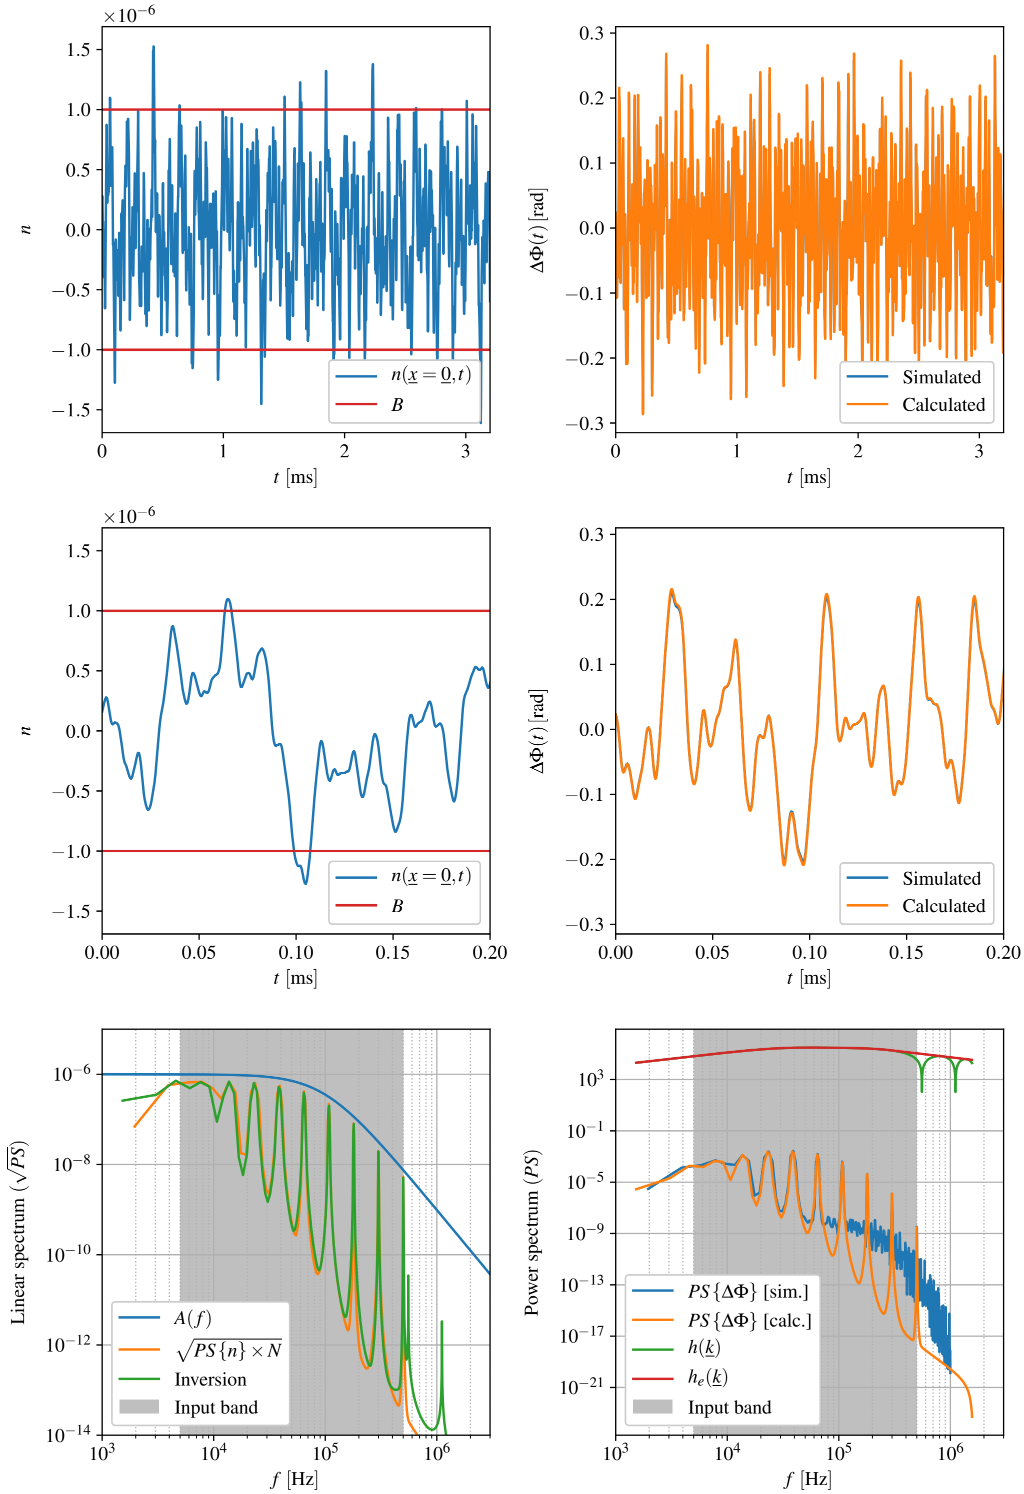

Figure 5.9: Spectral inversion validation study with $N = 10$, $\alpha = \beta = 0°$. Left column = input [$n(t)$], right column = output [$\Delta\Phi(t)$]. Middle row shows detailed zoom from full time-series in top row.

now, all directions $\hat{\underline{k}}_i$ are still confined to the xy plane, because of additional complications that arise when components of $\hat{\underline{k}}$ are in the z direction (this is the topic of Section 5.5). Recalling the assumption that contributing sources from all four walls share the same $A(f)$, then Eq. (5.58) is modified to become:

$$A(f) \approx LS\{n(t)\} \times \sqrt{N/M} = \sqrt{PS\{n(t)\} \times N/M} \tag{5.61}$$

The overall response function $h(\underline{k})$ in Eq. (5.59) is found using:

$$h(\underline{k}) = \sum_{i=1}^{M} h_i(\underline{k}) \tag{5.62}$$

Again, the output spectrum is mapped back very closely to the input spectrum using these transformations.

One small detail to note is that the inversion demonstrated in Fig. 5.9 uses the non-oscillatory envelope of the response function, $h_e(\underline{k})$, while that shown in Fig. 5.10 uses the original oscillatory $h(\underline{k})$. When the oscillatory region of $h(\underline{k})$ begins at high values of f beyond the bandwidth of the signal, as it does in these two examples, the differences between using $h(\underline{k})$ vs. $h_e(\underline{k})$ are small. $h(\underline{k})$ gives more accurate inversion, but the non-zero spectral noise floor gives rise to spurious spikes beyond

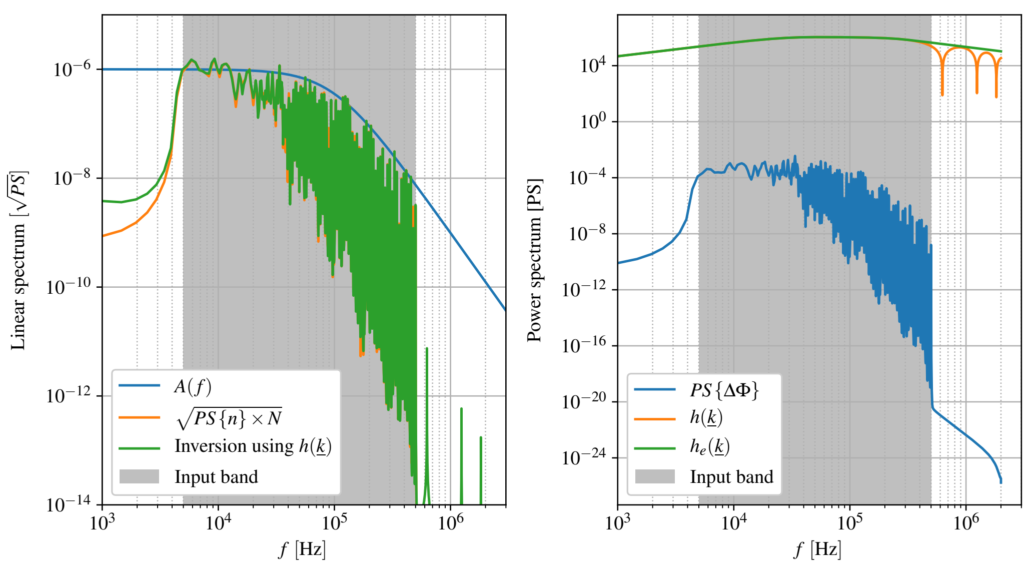

Figure 5.10: Spectral inversion validation study with $N = 1000$, $\alpha_j = \{0°, 15°, 30°, 45°\}$, $\beta_j = 0°$.

the bandwidth limit. $h_e(\underline{k})$ avoids such spikes, but can artificially damp actual signal components where $h_e(\underline{k})$ begins to diverge from $h(\underline{k})$; this is noticeable in the last frequency component peak of Fig. 5.9. These issues are not usually a problem for waves confined to the xy plane, because the oscillatory region governed by $Q(\underline{k})$ tends to be at very high f. However, when $\beta \neq 0$, the other oscillatory region due to $S(\underline{k})$ is at much lower f and cannot be neglected. This is the topic of the next section.

These oscillatory issues notwithstanding, the outcome of Sections 5.3 and 5.4 is that for plane wave disturbance fields, so long as the linearized ray-tracing equation is valid, we can both predict $\Delta\Phi(t)$ from known $n(\underline{x}, t)$, and recover the spectral magnitude of $n(\underline{x}, t)$ from a given $\Delta\Phi(t)$.

5.5 Propagation Along Optical Axis

Fig. 5.11 shows a similar inversion procedure to those in the previous section; this time, the propagating field has a component along the z axis ($\beta = 10°$). The parameters are such that the oscillatory roll-off in $h(\underline{k})$ occurs in the middle of the input bandwidth. Even in the temporal signal $\Delta\Phi(t)$ it is clear how strongly higher frequencies have been attenuated, despite the rather modest value of β. The spectral representation $PS\{\Delta\Phi\}$ has oscillations and the correct decay profile.

However, the discrete natures of both the signal and Welch's method mean the zeros do not always line up perfectly with their analytical locations, so inversion done via division by $h(\underline{k})$ causes spurious spikes due to slight misalignment. Conversely, inversion via $h_e(\underline{k})$ incorrectly preserves the oscillations when mapping back to the input spectrum. This misalignment of the function zeros causes the first issue for inversion when there is a component of propagation along the optical axis. A second issue is apparent in the output spectrum: even though $PS\{\Delta\Phi\}$ should continue to roll-off in magnitude, it eventually bottoms out on the spectral noise floor. This then leads to a spurious excursion in the inverted spectra, due to division by very small $h(\underline{k})$.

The first issue is mitigated somewhat in realistic applications, where we expect the incoming acoustic waves from each wall to have a range of angles about some preferred orientation. For example, the DNS studies of Duan et al. (2019) reported inclinations in the range $28° \pm 2°$. With all other parameters held constant, the locations of the zeros of $S(\underline{k})$ are dependent on β. This causes the mean response function to be smoothed, with less severe oscillations. This phenomenon was

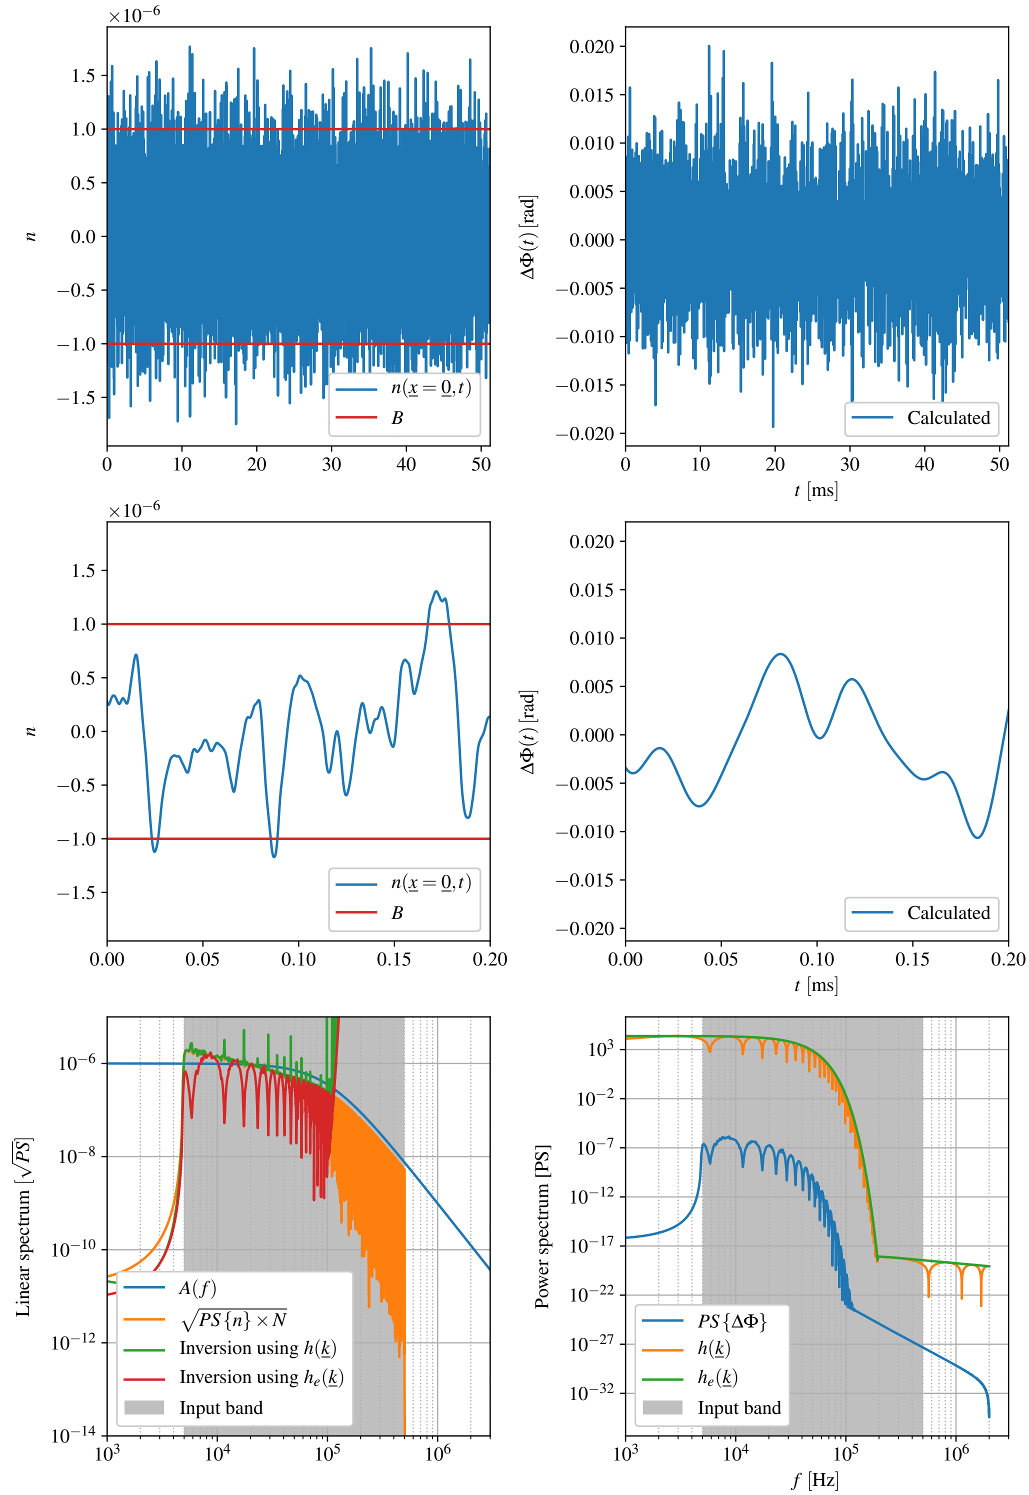

Figure 5.11: Spectral inversion with propagation along optical axis. $N = 1000$, $\alpha = 0°$, $\beta = 10°$. See Fig. 5.9 caption for subplot descriptions.

anticipated by Fulghum (2014) for xy propagation (i.e. the smoothing of oscillations in $Q(\underline{k})$). The effect is actually stronger in $S(\underline{k})$ because the zeros are not evenly spaced, getting closer together with increasing k. An example of this smoothing is shown in Fig. 5.12, for a disturbance containing a uniform distribution of angles in the range 35° ± 5°. While the oscillations are not entirely negated, they are damped substantially, and the non-oscillatory envelope $h_e(\underline{k})$ for the mid-point angle is a fairly good approximation for the overall behavior. This implies that the use of $h_e(\underline{k})$ for inversion should be quite accurate for real tunnel data. As mentioned, this concept doesn't work as well for oscillations caused by $Q(\underline{k})$; however these are often in a frequency band beyond the bandwidth of the experimental acquisition systems.

The second issue is ultimately caused by division of a small value by an even smaller one. Because $h(\underline{k})$ rolls off so steeply, any actual instrument response will be below the noise floor and thus negligible. Instead of inverting by dividing by $h(\underline{k})$, we can multiply by the reciprocal $h^{-1}(\underline{k})$, modified to include a cut-off, as well as being non-oscillating:

$$h_{ec}^{-1}(\underline{k}) = \begin{cases} \frac{1}{h_e(\underline{k})} & k \le k_{c2} \\ 0 & k > k_{c2} \end{cases} \tag{5.63}$$

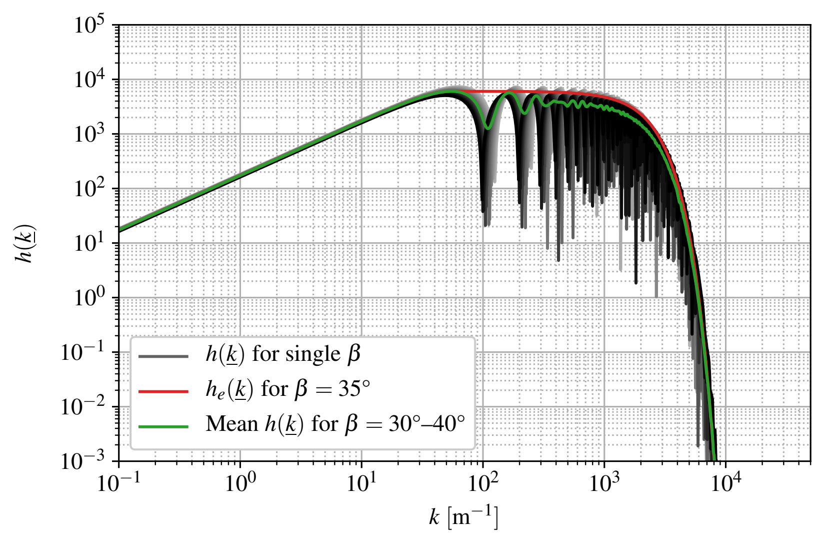

Figure 5.12: Effect of having a band of propagation directions on response functions. The shade of grey darkens for increasing β of each contributing component.

Again, the critical wavenumber k_{c2} needs to be determined empirically. For example, it could be chosen as the value where $h_e(\underline{k})$ falls to 1% of its maximum value.

5.6 Acoustic Waves in Convecting Flows

Thus far, the acoustic wave fields considered have all been plane waves propagating in a quiescent medium. For our applications in wind tunnels, we need to model a medium that is convecting at a temporally-steady, spatially-uniform velocity. Fortunately, this does not change the wave equation Eq. (5.17) and so all the derived results for $\Delta\Phi$ still hold. The only modification is in the dispersion relation that connects ω or f to k. Without convection, this was described by Eq. (5.18). With convection, it becomes:

$$k = \frac{\omega}{|\underline{u} \cdot \underline{\hat{k}} + a|} \tag{5.64}$$

where $\underline{u}$ is the convective velocity vector and a is the sound speed. By definition $\omega, k > 0$, hence the need for the absolute value of the denominator; signs are taken care of by $\underline{\hat{k}}$ as per Eq. (5.19). In most applications of FLDI to wind tunnels, $\underline{u} = u\underline{\hat{x}}$. Chaudhry et al. (2019) construct very similar acoustic-convective disturbance fields, albeit with propagation constrained to the xy plane; their dispersion relation Eq. (1.21) can be obtained as a special case of the fully-3D result given here.

Despite its simplicity, Eq. (5.64) can have strong effects on the interpretation and inversion of FLDI data. In high-Mach facilities $u >> a$ and so the denominator can change magnitude substantially. For the standard scenario of convective flow in the positive x-direction with Mach number M:

$$k = \frac{\omega}{|M \cos\alpha \cos\beta + 1|\, a} \tag{5.65}$$

For every $M > 1$ there is one wave orientation where Eq. (5.65) breaks down: if $M \cos\alpha \cos\beta = -1$ then the denominator goes to 0. Physically, this is a stationary Mach wave, i.e. no propagation occurs (and there would be no time-varying FLDI response).

$h(\underline{k})$ is unaffected by convection, but when the response function is plotted in f it will be shifted. An example is shown in Fig. 5.13. $\underline{\hat{k}}$ is compared for a wave propagating without any convection (i.e. $M = 0$) and for the same wave superimposed on a convective flow of $M = 5$. The so-called "fast" and "slow" acoustic modes are both

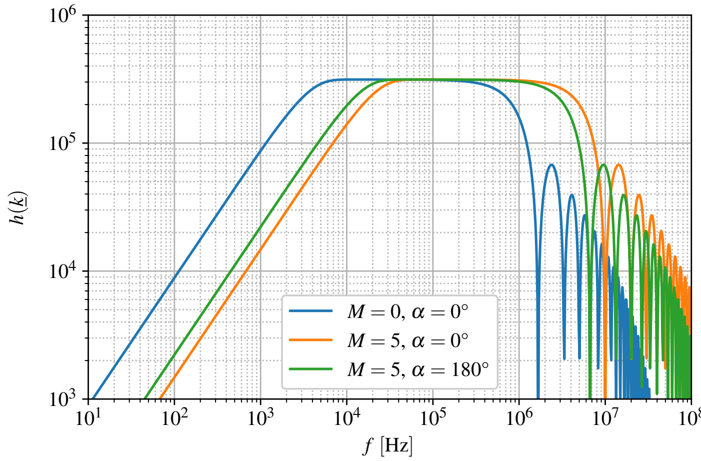

Figure 5.13: Frequency-shifting of response function due to convective mean flow. $\beta = 0°$, $a = 300\,\mathrm{m\,s^{-1}}$, $L = 1$ m, standard optical parameters.

shown, where the wave is facing downstream and upstream, respectively. In both cases, the wavefronts actually move downstream since $M > 1$, but different levels of shift are observed, since the effective wave speeds are $M \pm 1$.

For convection in the x direction with acoustic waves in the xy direction, the fast mode has $|\alpha| < 90°$, while the slow mode has $90° < |\alpha| \leq 180°$. More generally, the fast and slow modes are defined by $\underline{u} \cdot \hat{\underline{k}} > 0$ and $\underline{u} \cdot \hat{\underline{k}} < 0$, respectively.

5.7 Comparison with Method of Parziale et al.

5.7.1 Overview

The preceding lengthy exposition on the exact response functions invites comparison with the original method of Parziale et al. (2012). First introduced in this thesis as Eq. (1.3), it is repeated here for convenience:

$$\frac{\Delta\rho}{\overline{\rho}} = \frac{\lambda_L}{2\pi K \zeta \overline{\rho}} \sin^{-1}\left(\frac{V}{V_0} - 1\right) \tag{5.66}$$

with $\overline{\rho}$ as a local average density, λ_L the laser wavelength, K the Gladstone-Dale constant, V the FLDI output voltage, and V_0 the voltage at the middle of a fringe. ζ is a measure of the sensitive length (similar to that discussed qualitatively in Section 5.2). This method is attractive because it is simple to apply, either in temporal or frequency space, i.e. using $V(t)$ or $V(f)$. It was also published early in the modern development history of FLDI, and so has seen use by subsequent authors, e.g. Benitez et al. (2020) and Ceruzzi et al. (2020). Note that $\overline{\rho}$ appears on both sides and is only needed if one wishes to express the fluctuations in relative terms.

Parziale (2013) later introduced a correction coefficient:

$$c(\lambda) = \sin\left(\frac{\pi\Delta x}{\lambda}\right) \tag{5.67}$$

This was derived for sinusoidal disturbances of wavelength λ. The post-processing workflow is as follows:

1. Measure $V(t)$, with known $\{V_0, \lambda_L, \overline{\rho}, K\}$.
2. Compute $V(f)$ from $V(t)$ using Welch's method or similar.
3. Convert $V(f)$ to uncorrected $[\Delta\rho/\overline{\rho}](f)$ using Eq. (5.66).
4. Assume some disturbance propagation velocity with which to relate f to λ.
5. Convert $[\Delta\rho/\overline{\rho}](f)$ to $[\Delta\rho/\overline{\rho}](\lambda)$.
6. Compute $c(\lambda)$ using Eq. (5.67).
7. Correct the spectrum using $[\Delta\rho/\overline{\rho}](\lambda) \div c(\lambda)$.

This methodology was the one used in Parziale et al. (2014) to obtain density spectra of the T5 freestream.

5.7.2 Theoretical Relationship with Exact Solution

Although Eq. (5.66) does not appear to bear much resemblance to the ray-tracing Eq. (4.2), the former can be obtained as a simplified case of the latter. First recall that the overall post-processing procedure for FLDI consists of three conversions:

$$V \xrightarrow{A} \Delta\Phi \xrightarrow{B} n \xrightarrow{C} \rho \tag{5.68}$$

Conversions A and C are trivial, being given by Eq. (3.17) and Eq. (1.1) respectively. The complicated part, and the subject of this chapter, is conversion B. Eq. (5.66) represents the composite conversion ABC. To compare the methods we need to expose how Parziale et al. treat conversion B alone.

Assume an ideal interferometer with perfect contrast, such that $V_{min} = 0$, and with the reference point V_0 set to exactly the mid-fringe position. This reduces Eq. (3.17) to:

$$V = V_0 \left[1 + \sin\left(\Delta\Phi\right)\right] \tag{5.69}$$

Rearranging:

$$\Delta\Phi = \sin^{-1}\left(\frac{V}{V_0} - 1\right) \tag{5.70}$$

Using Eq. (5.70), Eq. (5.66) can be rearranged to:

$$\Delta\Phi = \frac{2\pi}{\lambda_L} K\zeta\Delta\rho \tag{5.71}$$

Consider Eq. (4.3) (the polar form of the ray-tracing equation) and substitute the Gladstone-Dale relation:

$$\Delta\Phi = \frac{2\pi}{\lambda_L} \int_0^{2\pi} \int_0^{\infty} \tilde{I}_0\left(\tilde{r}\right) \left(\int_{z_a}^{z_b} (K\rho_1 + 1)\mathrm{d}z - \int_{z_a}^{z_b} (K\rho_2 + 1)\mathrm{d}z \right) \tilde{r}\mathrm{d}\tilde{r}\mathrm{d}\theta \tag{5.72}$$

This simplifies to:

$$\Delta\Phi = \frac{2\pi}{\lambda_L} K \int_0^{2\pi} \int_0^{\infty} \tilde{I}_0\left(\tilde{r}\right) \left(\int_{z_a}^{z_b} \rho_1\mathrm{d}z - \int_{z_a}^{z_b} \rho_2\mathrm{d}z \right) \tilde{r}\mathrm{d}\tilde{r}\mathrm{d}\theta \tag{5.73}$$

Parziale et al. only consider sinusoidal plane waves propagating in the x-direction, i.e. $\rho \neq \rho(z)$. However, recall that the co-ordinate transformation to $(\tilde{r}, \theta, z)$ involves normalization by $w = w(z)$, so there is z-dependence in the integrations that originates from the focusing beam geometry. The approximation is made that the instrument is only sensitive within $|z| \leq \zeta/2$, even if the disturbance extent $[z_a, z_b]$ exceeds this. Furthermore, the focusing within the sensitive length ζ is neglected, so that:

$$\int_{z_a}^{z_b} \rho_i\mathrm{d}z = \rho_i \int_{z_a}^{z_b} \mathrm{d}z \approx \rho_i \int_{-\zeta/2}^{+\zeta/2} \mathrm{d}z = \rho_i\zeta \tag{5.74}$$

Then:

$$\Delta\Phi = \frac{2\pi}{\lambda_L} K\zeta \int_0^{2\pi} \int_0^{\infty} \tilde{I}_0 \left(\tilde{r}\right) \left[\rho_1 - \rho_2\right] \tilde{r} \mathrm{d}\tilde{r} \mathrm{d}\theta \tag{5.75}$$

Reverting to Cartesian co-ordinates we have $\rho = \rho(x)$ only, so:

$$\Delta\Phi \approx \frac{2\pi}{\lambda_L} K\zeta \int_0^{\infty} \int_0^{\infty} \tilde{I}_0 \left(x, y\right) \left[\rho_1 \left(x + \frac{\Delta x}{2}\right) - \rho_2 \left(x - \frac{\Delta x}{2}\right)\right] \mathrm{d}x\mathrm{d}y \tag{5.76}$$

The final approximation made is that this integral can be replaced with some $\Delta\rho$ representing a weighted average of the density differences between the two beams, thereby recovering Eq. (5.71). It can be viewed as a weighted average because the integral of the normalized Gaussian intensity distribution $\tilde{I}_0$ over all space is unity.

Note that the introduction of the Gladstone-Dale relation only alters Eq. (5.73) by a factor of K, so Eq. (5.71) can instead be expressed in terms of refractive index:

$$\Delta\Phi = \frac{2\pi}{\lambda_L} \zeta \Delta n \tag{5.77}$$

This isolates conversion B from Parziale et al.'s method and allows comparison with the exact result derived from Schmidt and Shepherd. To summarize, the simplifications made neglect the changing area of the focusing beams along with their Gaussian intensity profiles, which in the method of Schmidt and Shepherd are encapsulated in H_w. The only aspect of H_w that is preserved is the roll-off in response with increasing z, which is approximated by a sharp cut-off at $\pm\zeta/2$. For their optical parameters (which are similar to ours) Parziale et al. suggest $\zeta/2 \approx 10$ mm. This was originally based on a geometrical argument involving beam overlap and common-mode rejection, later deemed irrelevant to the operation of FLDI by Schmidt and Shepherd, although some simple benchtop experiments also gave similar sensitive lengths.

Although H_w does not enter into Parziale et al.'s method, $c(\lambda)$ can be seen as a predecessor to H_s, i.e. the effect of the finite separation Δx. Parziale et al. only address the roll-off effect of $c(\lambda)$ on low frequencies, not the high-frequency aliasing, likely because this aliasing only occurs at extremely high frequencies beyond the range of their experiments.

$\Delta\rho$ (or equivalently Δn) obtained by this method are treated by Parziale et al. as good approximations to the true fluctuations ρ' or n', since Δx is small.

5.7.3 Comparisons

With Eq. (5.77), direct comparison can be made with the exact solution, using the following procedure:

1. Define a known $n'(\underline{x}, t)$ and corresponding spectrum $n'(f)$.
2. Compute $\Delta\Phi(t)$ using the ray-tracing approach of Eq. (5.21).
3. Use $\Delta\Phi(t)$ or $\Delta\Phi(f)$ with Eq. (5.77) to compute Δn.
4. Compute $c(\lambda)$ or $c(f)$ and correct Δn.
5. Compare Δn with n'.

Beginning with a single-frequency wave, close agreement was found between the two methods, if ζ is chosen appropriately. This choice is very sensitive to the parameters of the wave and the FLDI. For many trialed frequencies, Parziale et al.'s suggestion of $\zeta \approx 20\,\mathrm{mm}$ resulted in Δn that were off by orders of magnitude.

For waves with x-propagation only, Eq. (5.24) shows that additional contributions to $\Delta\Phi$ diminish as:

$$\mathrm{erf}\left(\frac{kL\lambda_L}{2\sqrt{2}\pi w_0}\right) \to 1 \tag{5.78}$$

This can be used to determine ζ:

$$\zeta_\tau \equiv \frac{2\sqrt{2}w_0}{\lambda_L} \cdot \frac{u_c}{f} \cdot \mathrm{erf}^{-1}(\tau) \tag{5.79}$$

where τ is some threshold $\lesssim 1$, e.g. $\zeta_{0.99}$ is the length within which 99% of the response is generated. For $x \in \mathbb{R}$, $\mathrm{erf}(x)$ is one-to-one and so $\mathrm{erf}^{-1}(x)$ is defined, and is implemented in SciPy and other software packages. Note u_c is the convective velocity of the flow-field; for example, an acoustic wave superimposed on an underlying freestream flow would have $u_c = u_\infty + a$. This means that the sensitive length depends on this velocity when expressing ζ in terms of f.

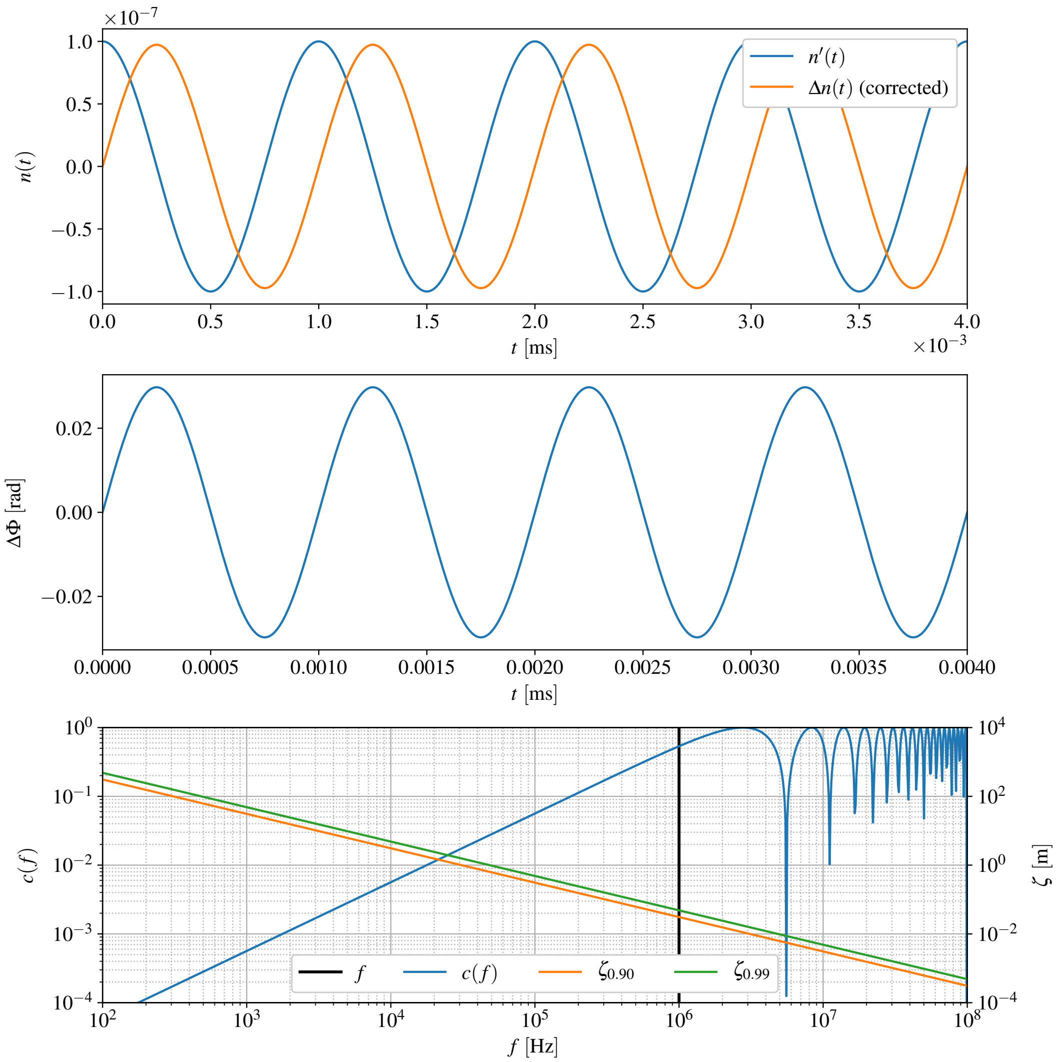

Figure 5.14: Comparison of exact n' with approximate Δn computed using Parziale et al. $u_c = 1\,\mathrm{km\,s^{-1}}$, $f = 1\,\mathrm{MHz}$, $A = 10^{-7}$, $\zeta_{0.99} \approx 48\,\mathrm{mm}$.

An example single-frequency case, using Eq. (5.79) to determine ζ, is shown in Fig. 5.14. This figure shows the original $n'(t)$ and the corresponding exact $\Delta\Phi(t)$, along with the approximate $\Delta n(t)$. The agreement in magnitude is quite good using $\zeta_{0.99}$. Note $\Delta n(t)$ is a quarter-cycle out-of-phase with $n'(t)$, because Parziale et al.'s method does not account for the phase shift that results from FLDI being a response to the density gradient; of course this is unimportant for the usual spectral representation of data.

The bottom plot of Fig. 5.14 shows the frequency dependence of $c(f)$ and ζ. Tests at other frequencies show that their $c(f)$ indeed provides the right magnitude correction if used in conjunction with the appropriate ζ. However, it can also be seen that ζ decreases by an order of magnitude for every decade of increase in f. This immediately implies that interpreting spectral data with a constant ζ will not be successful if the data span a bandwidth of any significance.

Fig. 5.15 extends the comparison to waves with multiple frequencies. The spectrum of the exact field $n'(t)$ is compared with three candidates: two use a fixed $\zeta = 20\,\mathrm{mm}$, with and without the correction factor $c(f)$, the third uses a variable $\zeta(f)$ from Eq. (5.79) as well as $c(f)$. It can be seen that this third method produces excellent agreement across the whole spectrum. Coincidentally, the value of ζ chosen here actually causes the original uncorrected method to also match quite closely, but the corrected method diverges strongly. However, other choices of ζ shift these spectra vertically.

The reason the uncorrected method (no $c(f)$, constant ζ) maintains the same shape as the true spectrum is that $c(f)$ and $\zeta(f)$ scale like f and f^{-1}, respectively, and so overall have no f-dependence. This means there is a single constant value ζ for a given flow-field geometry and velocity that will recover the true spectrum, but this ζ may not have an obvious relationship with any experimental length scales. In order to be sure that ζ is correct, Eq. (5.79) needs to be used, thereby drawing from the full ray-tracing theory. At this point, all the same parameters need to be known and assumed, such that one might as well just use Eq. (5.24) without much increase in complexity. There are also additional caveats to using Parziale et al.'s method: if $\zeta(f)$ exceeds the actual integration length $2L$ then inaccuracies will occur at those frequencies: essentially, high-frequency contributions will be over-counted.

The theory is strictly for plane sinusoidal waves propagating in the x-direction only, and if the disturbances either have some other component of motion or are not sinusoidal (e.g. a shock wave), the method will not give correct results, whereas the ray-tracing theory can give analytical results for any plane wave, and numerical results for arbitrary fields. For example, if the above comparison in Fig. 5.15 is repeated for a wave inclined in the xy plane at $\alpha = 30°$, there is no mechanism in Parziale et al.'s method to account for the inclination. Fig. 5.16 shows that even with the corrections to the method (variable ζ) the magnitude of the field will be underpredicted.

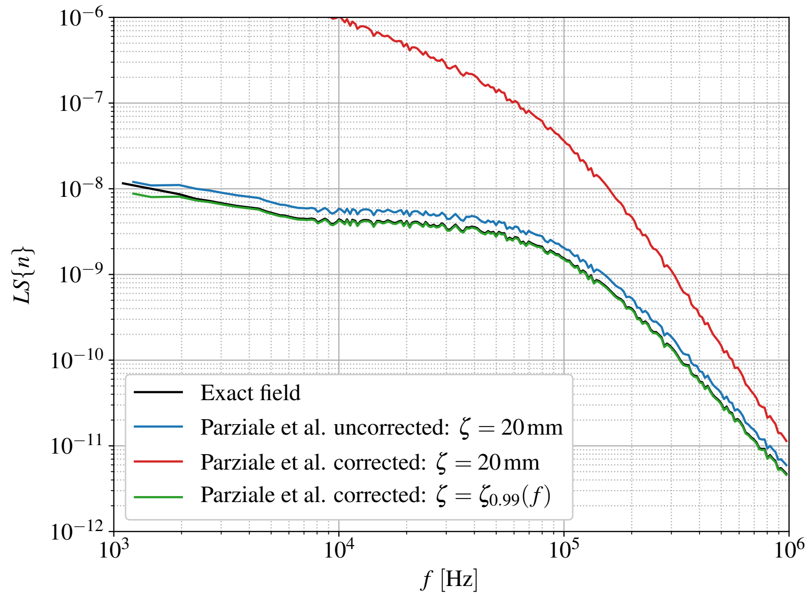

Figure 5.15: Comparison of exact refractive index spectrum with reconstructions using the method of Parziale et al. with various modifications. The flow-field is a plane wave propagating in the x-direction.

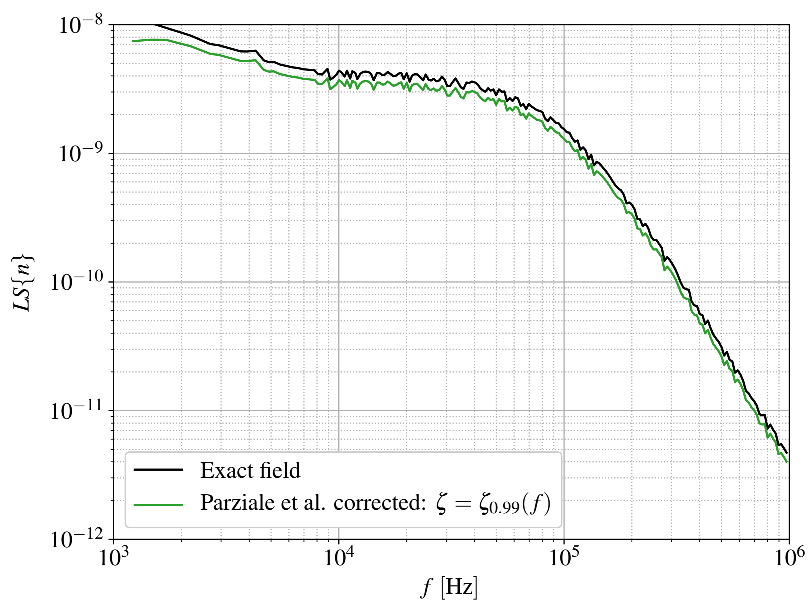

Figure 5.16: Comparisons for a plane wave propagating in the xy-direction with $\alpha = 30°$.

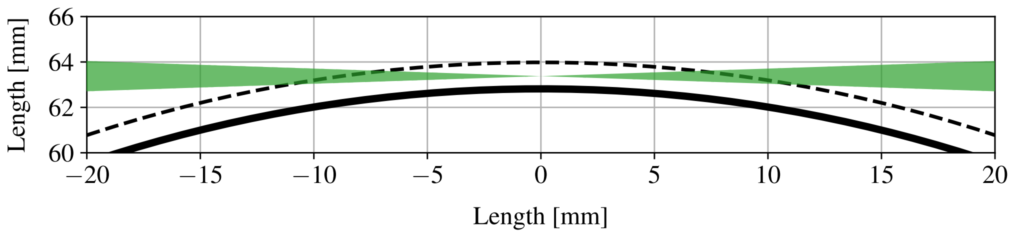

Figure 5.17: Cross-section of the geometrical interaction of FLDI beams [green] with the boundary layer [black dashed line] over a cone [black solid line]. Modified from Parziale et al. (2015) and used with permission from N. J. Parziale.

Besides its use for making freestream measurements in T5, FLDI was also used by Parziale et al. (2015) to observe instabilities in hypersonic boundary layers on a slender cone. The foci of the FLDI were aligned parallel with and very close to the surface of the cone. Geometrical calculations (Fig. 5.17) showed that a region of the beams roughly ± 10 mm about the focal plane were immersed in the boundary layer. A plot of $\zeta_{0.99}(f)$ can be used to determine whether observed instabilities of a given frequency lie within the boundary layer. For example, in Fig. 5.14, frequencies above $f \approx 5$ MHz have $\zeta_{0.99} < 10$ mm and so any signals in this bandwidth are guaranteed to originate within the boundary layer. Note that this is just an illustration of the method; $\zeta_{0.99}(f)$ in Fig. 5.14 was not computed for the same conditions as Parziale et al.'s experiment.

5.8 Conclusions

In the introduction to this chapter, the need to be able to perform the FLDI inverse problem $n(\underline{x}, t) = f^{-1}(\Delta\Phi(t))$ was described. It has been demonstrated here that this is possible in a spectral sense for all refractive index fields comprised of arbitrary superpositions of sinusoidal plane waves and uniform convecting flows. Furthermore, this inversion can be expressed as an explicit analytical function. The various terms of this function have been discussed as they give insight into the physical mechanisms behind FLDI operation. This particular class of field was not merely chosen because it yields an analytical solution. Multi-frequency sinusoidal plane waves, incident at some angle to a freestream flow, are a good model for acoustic radiation from wall boundary layers. From a practical standpoint, this theory can be used to recover density and pressure fluctuations from hypersonic wind tunnels—this will be demonstrated in Chapter 6.

The theory developed in this chapter was also compared with the previous developments of other authors. Specifically, the present results are shown to be a generalization of Schmidt and Shepherd's various transfer functions, and they were also used to make corrections to the functions of Settles and Fulghum (2016). Lastly, the original simple theory of Parziale et al., although appearing superficially unrelated, was shown to be recoverable by applying suitable approximations to the exact theory. The limitations of Parziale et al.'s model were discussed, as its use remains widespread in the community.

Chapter 6

RESULTS: CONVENTIONAL TUNNEL EXPERIMENTS

To reiterate from Section 2.6, the high-speed data obtained at Tunnel D comprised two pitot pressure signals from symmetrical locations near the facility centerline, and two signals (one AC-coupled/filtered, one DC-coupled/unfiltered) from the same FLDI, positioned upstream from one of the pitot stations, approximately on the stagnation streamline. The raw voltages can be converted into pitot stagnation pressures p_0 and FLDI phase shifts $\Delta\Phi$ via known calibration constants. The main goal for the analysis is to further convert these data into the same units of freestream pressure or density in order to directly compare them. This requires some prior knowledge or assumptions about the flow geometry, as well as transfer functions for each diagnostic.

6.1 Raw Data

An example raw dataset is shown in Fig. 6.1. All runs from Table 2.2 look qualitatively similar; additionally, there is very little difference between the filtered and unfiltered FLDI data, or between the two pitot locations.

In the middle band of frequencies, the spectra magnitudes are ordered by reservoir pressure, i.e. the absolute signals measured by either technique get larger in magnitude as p_{res} increases. This ordering breaks down somewhat for $f < 1\,\text{kHz}$, although this may be due to increasing numerical error in the PSD algorithm as the record length is approached. The pitot data are also displayed normalized by $\bar{p}_0$; this is usual in the literature as it allows comparison of different facilities in terms of Reynolds number. The magnitudes of $PSD\left\{p_0'/\bar{p}_0\right\}$ are compared in Fig. 6.3 with a compilation of results from various similar facilities given by Duan et al. (2019).

Each measurement method shows a characteristic peak, regardless of p_{res}, that is absent from the corresponding location in the spectra of the other method: for the FLDI, centered at 1.7 kHz; for the pitot, centered at 55.5 kHz. The FLDI peak is due to laser intensity noise, as discussed in Appendix D (compare to Fig. D.1). The raw pitot spectra for all the different run conditions are roughly parallel to each other over their useful range; the same would be expected from the FLDI data. However, it is seen that the laser intensity noise causes all the FLDI spectra to "bunch up": up to about 6 kHz all spectra are on top of each other, i.e. completely dominated by the noise floor, and even considerably above this the lower-pressure conditions (in blue and violet) are strongly affected by the noise and do not run parallel to the other spectra. The useful bandwidth for the FLDI therefore differs between conditions, but for simplicity 30 kHz is chosen as a value where all conditions appear to be sufficiently above the noise floor.

The pitot feature resembles a resonance peak, and while the natural frequency of the Kulite is 175 kHz, this can be shifted substantially lower when disturbances approach at an angle (this phenomenon was described by Chaudhry et al. (2019) and reviewed in Section 1.3). The presence of these peaks and the RF noise leads to the definition of useful bandwidths for these data: for the FLDI, this is taken to be $30\,\text{kHz} < f < 200\,\text{kHz}$; for the pitot, $100\,\text{Hz} < f < 40\,\text{kHz}$.

The experimental noise floors were characterized for the FLDI and pitot sensors by taking measurements in various no-flow configurations, using all the same acquisition settings. The results are summarized in Fig. 6.4. For the FLDI, it was found that the vast majority of the background noise is due to the laser. With the laser off, there is no difference observed when the photodetector is covered or powered off: background light, dark current noise, or AC power supply interference are negligible compared with the laser noise. With the laser on, the spectra remain the same regardless of whether the no-flow test section is at vacuum or ambient conditions: airborne dust or ambient air currents are also negligible. The noise floor with the laser on is orders of magnitude higher than the baseline noise, and the lowest-density flow condition (corresponding to the minimum reservoir pressure) is barely above this threshold. The FLDI signal is significantly above the noise floor for higher-density conditions across the full useful bandwidth.

Taken together, Figs. 6.1 and 6.4 show that the data from the lowest-pressure condition (Burst 1, in violet) is unreliable regardless of the diagnostic: for the pitot, its shape is anomalous, and does not follow the correct trend when plotted in relative terms; for the FLDI, as mentioned, it sits right on the noise floor. Therefore, subsequent plots will exclude this condition.

6.2 Conversion of Pitot Data

The state-of-the-art for interpreting supersonic pitot probe data is the work of Chaudhry and Candler (2017) and Chaudhry et al. (2019). This was discussed at length in Section 1.3, but to summarize, DNS studies were performed of both

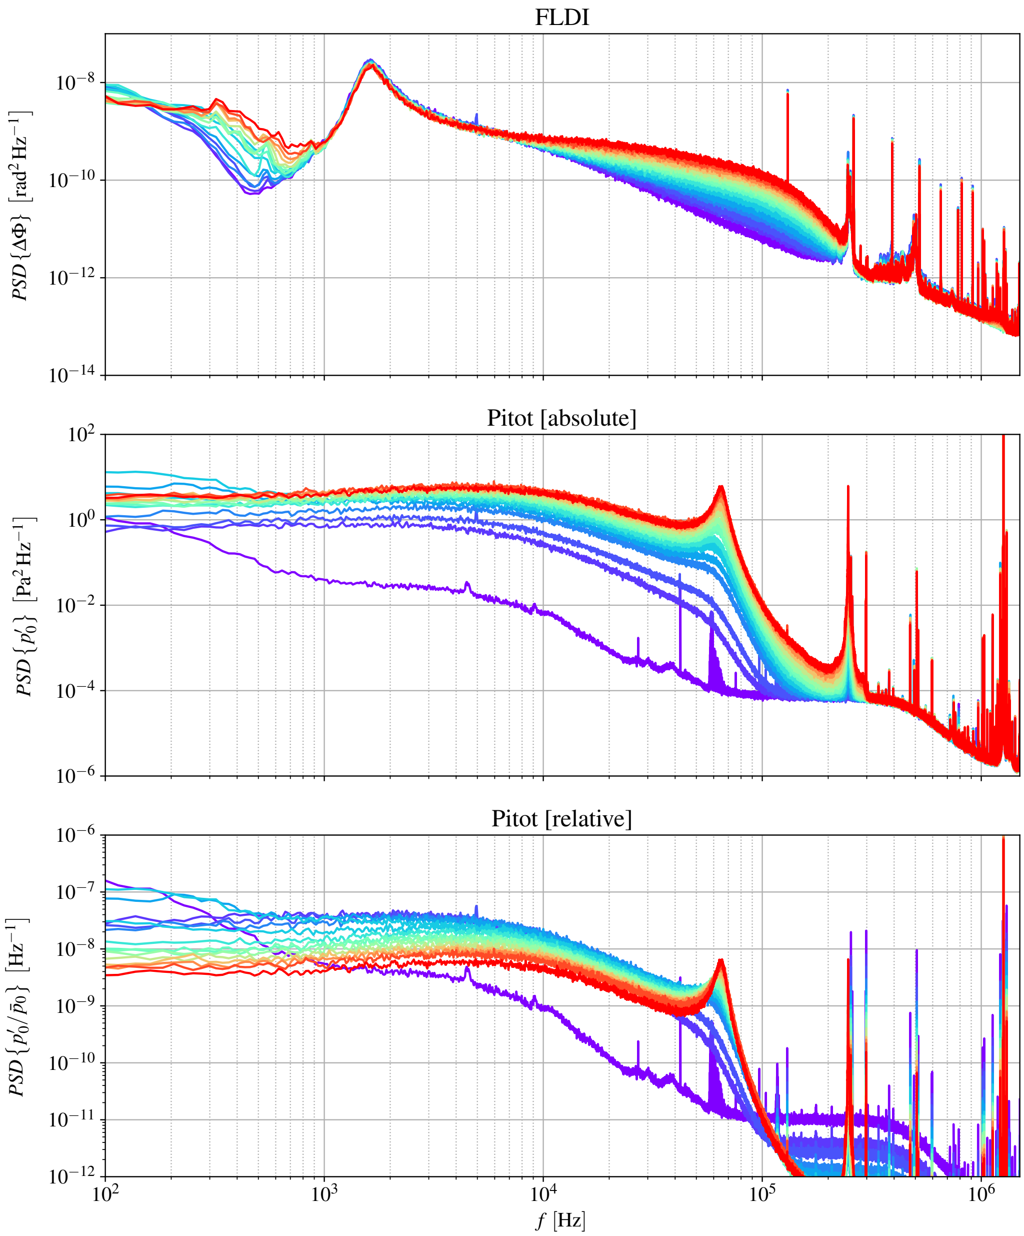

Figure 6.1: Example raw dataset from Tunnel D experiments. Shown is Run 184, with FLDI and pitot data (the latter displayed both in absolute and relative terms). Each color corresponds one of the 16 reservoir pressure setpoints, ranging from 34–414 kPa, see Fig. 6.2 for details.

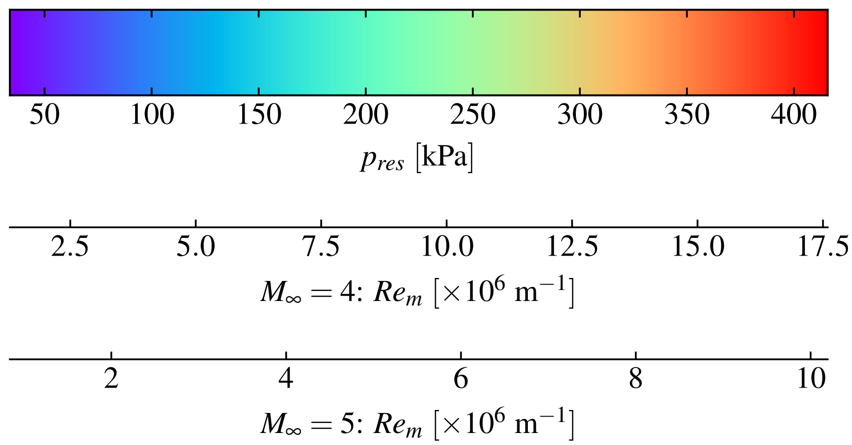

Figure 6.2: Color scale used throughout this chapter to represent p_{res} and corresponding Re_m.

flow-parallel and inclined acoustic plane waves incident on a pitot probe. It was found that the transfer functions were strongly dependent on the incident angle, frequency, and the probe geometry.

As yet, there are no analytical forms for these general transfer functions, but one was derived for the special case of flow-parallel waves in the low-frequency limit:

$$\chi^*_{\sigma\pm} = \left[\frac{M_\infty^2 \pm 2M_\infty \mp 1/M_\infty}{\gamma M_\infty^2 - (\gamma - 1)/2}\right]^2 \tag{6.1}$$

where M_∞ is the time-averaged freestream Mach number, $\sigma+$ and $\sigma-$ denote the fast and slow acoustic modes respectively, * the low-frequency limit, and the transfer function is defined as:

$$\chi(f) \equiv \frac{PSD\left\{p_0'/\bar{p}_0\right\}}{PSD\left\{p_\infty'/\bar{p}_\infty\right\}} \tag{6.2}$$

This result agreed with the DNS data at low frequencies (relative to the first resonant peak). However, there is no corresponding result for inclined waves—but it is known that the primary source of noise in Tunnel D are inclined acoustic waves Donaldson and Wallace (1971). The most widely-used approach remains to the present time the work of Stainback and Wagner (1972) (again, see Section 1.3), which despite having no dependence on frequency or geometry, does account for inclination. Repeated here with notational changes from Eq. (1.9):

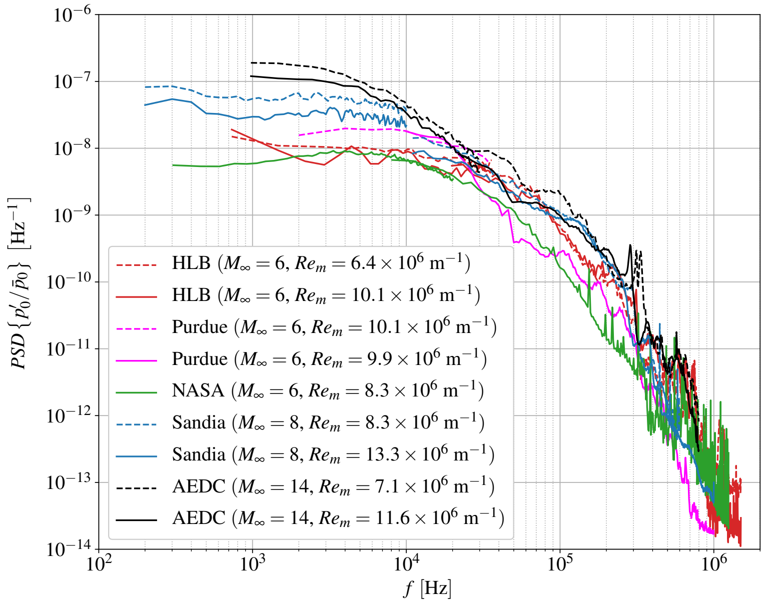

(a) Pitot pressure spectra for various facilities and conditions. Reproduced from Duan et al. (2019), with raw data and permission provided by L. Duan.

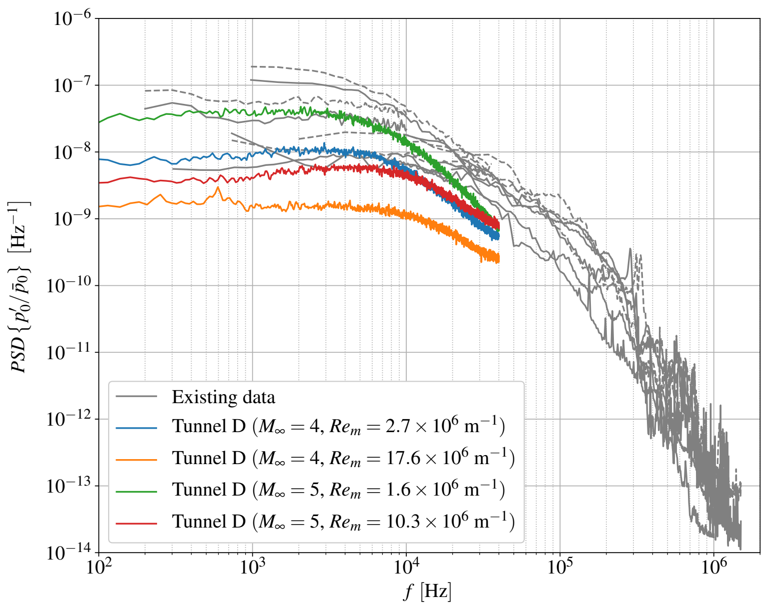

(b) Overlay of bounding Tunnel D conditions on existing spectra from Fig. 6.3a. Tunnel D data truncated before the resonance peak.

Figure 6.3: Comparison of Tunnel D pitot pressure spectra with literature data.

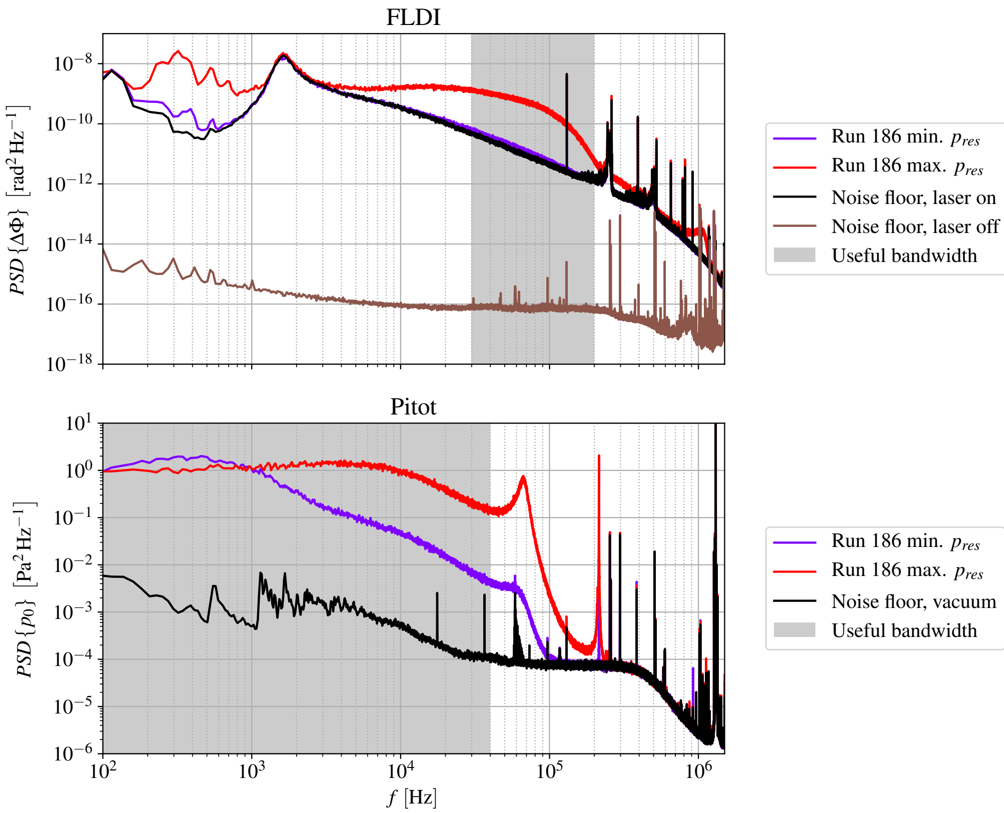

Figure 6.4: Background noise investigation for Tunnel D FLDI and pitot instruments. No-flow spectra are compared with the bounding cases of flow spectra; useful bandwidth for each instrument is also indicated.

$$\left(\frac{\tilde{p}_0}{p_0}\right)^2 = \left(\frac{\tilde{p}_\infty}{\gamma p_\infty}\right)^2 \left[1 - 4\frac{n_x}{M_\infty} + 4\left(\frac{n_x}{M_\infty}\right)^2\right] \tag{6.3}$$

Recall that the tilde represents RMS values rather than instantaneous fluctuations. However if Eq. (6.3) is assumed to hold in an instantaneous sense, then it can be rearranged to yield a transfer function χ in the sense of Chaudhry and Candler:

$$\chi^*_{sw\theta} = \frac{1}{\gamma^2}\left[1 - 4\frac{n_x}{M_\infty} + 4\left(\frac{n_x}{M_\infty}\right)^2\right] \tag{6.4}$$

Here, the subscript sw refers to "Stainback and Wagner" and the θ indicates that this holds for inclined waves, where $n_x \equiv \cos\theta$. Note that the additional factor of 2 that was added to account for unsteady shock reflection effects is not included, since here we wish to compare the quasi-steady low-frequency limit.

To directly compare with Eq. (6.2), we take the special case of flow-parallel disturbances (denoted with $\|$), where $\theta = 0°$ or $180°$, corresponding to $n_x = \pm 1$ for fast and slow mode respectively:

$$\chi^*_{sw\|} = \frac{1}{\gamma^2}\left[1 \mp \frac{4}{M_\infty} + \frac{4}{M_\infty^2}\right] \tag{6.5}$$

The comparison between Eqs. (6.2) and (6.5) is made in Fig. 6.5. Despite being derived by quite different methods, the two sets of transfer functions agree very closely, except below about $M_\infty = 2$; this is to be expected, since Stainback and Wagner make use of an approximation that holds for $M_\infty \gtrsim 2.5$ (Eq. (1.7)). However, there is a sign error: Stainback and Wagner's slow mode matches Chaudhry and Candler's fast mode, and vice versa. It is known from the DNS validation that the latter form is correct, so the sign error must lie with Stainback and Wagner.

Re-examining the derivation, the starting point is a general fluctuation equation, Eq. (1.8). The assumption is then made of pure acoustic waves, specifically slow-mode, which makes the correlation coefficient $R_{\rho u} = -1$ (it would be +1 for fast mode). The wave inclination is introduced via a modification to the isentropic relations:

$$\frac{\tilde{u}_\infty}{u_\infty} = \frac{n_x}{\gamma M_\infty} \cdot \frac{\tilde{p}_\infty}{p_\infty} \tag{6.6}$$

Consider the signs of each quantity in this equation. By definition the RMS quantities $\tilde{u}$ and $\tilde{p}$ are positive, as are the time-averaged pressure p, γ, and M_∞. Although for a slow-mode wave the instantaneous velocity fluctuations u' can be negative, the time-averaged (supersonic) velocity u must remain positive for the problem set-up to make physical sense. Thus $n_x > 0$ in order for Eq. (6.6) to be resolved. But this conflicts with the choice of slow-mode waves and $R_{\rho u} = -1$, since this requires $90° < \theta \leq 180°$ and so $n_x < 0$. Stainback and Wagner themselves use Laufer's method for obtaining n_x from the sound source velocity, re-expressed in their paper in the form:

$$n_x = \left(\frac{u_s - u_\infty}{u_\infty}\right)^{-1} M_\infty^{-1} \tag{6.7}$$

They use $u_s/u_\infty = 0.6$, which again requires $n_x < 0$. Hence the error may have arisen in Eq. (6.6) because of the use of RMS rather than instantaneous values, or even a “double negative” in choosing $R_{\rho u} = -1$ as well as using θ referenced from a fast-mode datum. In any case, a simple sign-change correction can be made to Eq. (6.4):

$$\chi^*_{sw\theta} = \frac{1}{\gamma^2}\left[1 + 4\frac{n_x}{M_\infty} + 4\left(\frac{n_x}{M_\infty}\right)^2\right] \tag{6.8}$$

This brings the flow-parallel cases $\chi^*_{sw\|}$ into agreement with $\chi^*_{\sigma\pm}$. It is now assumed that the angled cases $\chi^*_{sw\theta}$ will also hold in the low-frequency limit. The deviation from the validated flow-parallel result as θ changes is shown in Fig. 6.6.

As noted by Chaudhry and Candler, the fast- and slow-mode χ^* are substantially different in value even for quite high M_∞, as the two curves approach the high-M_∞ asymptotic limit slowly (illustrated in Fig. 6.5). Furthermore, at the conditions for Tunnel D, where $M_\infty = 4$ or 5 and $\theta \approx 125°$, the deviation in χ from flow-parallel is substantial.

The question now arises as to the data range over which these χ^* are valid. Chaudhry and Candler used bandwidths of $0 < f \lesssim 0.1 f_{peak}$ for their low-frequency validation studies, where f_{peak} is the resonant frequency found from the peak response. For example, Fig. 6.1 has $f_{peak} \approx 65$ kHz, allowing the use of χ^* only up to $f \approx 6.5$ kHz. The DNS results for the flow-parallel case show that at first, χ/χ^* deviates from unity relatively slowly as f increases: e.g. at $f \approx 0.25 f_{peak}$, $\chi/\chi^* \approx 1.2$, although it should be noted that these results were for the flow-parallel case; corresponding χ/χ^* normalizations were not provided since an analytical χ^* was not derived by Chaudhry et al. Conservatively, $f = 0.2 f_{peak}$ will be used as the cut-off. This represents a further restriction on the aforementioned useful bandwidth for the pitot, now $100\,\text{Hz} < f < 13\,\text{kHz}$. Note that since we divide by χ^* to recover $PSD\left\{p'_\infty/\bar{p}_\infty\right\}$, use of χ^* beyond this cut-off could produce a substantial over-estimation in magnitude, and the strong frequency dependence past this point will also result in incorrect spectral slopes.

Finally, because the measured pitot pressure p_0 consists of (assumed equal) contributions from all 4 walls, the static pressure spectrum of an inclined wave is recovered by division by $4^2 = 16$ since the spectral representation chosen is the PSD, which scales with the square of the signal.

This section has presented the methodology that will be used to convert $PSD\left\{p'_0/\bar{p}_0\right\}$ to $PSD\left\{p'_\infty/\bar{p}_\infty\right\}$. Some additional details of how to properly handle the different

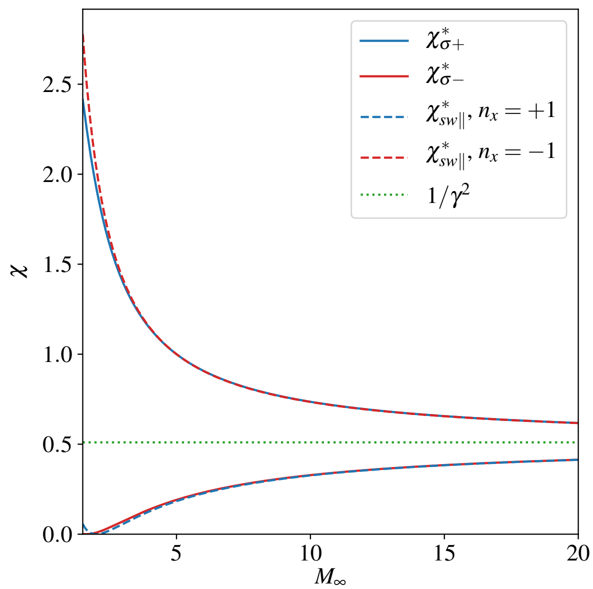

Figure 6.5: Comparison of pitot transfer functions for flow-parallel acoustic disturbances from Chaudhry and Candler [$\chi^*_{\sigma\pm}$] vs. Stainback and Wagner [$\chi^*_{sw\parallel}$]. High-M_∞ asymptotic limit shown in green.

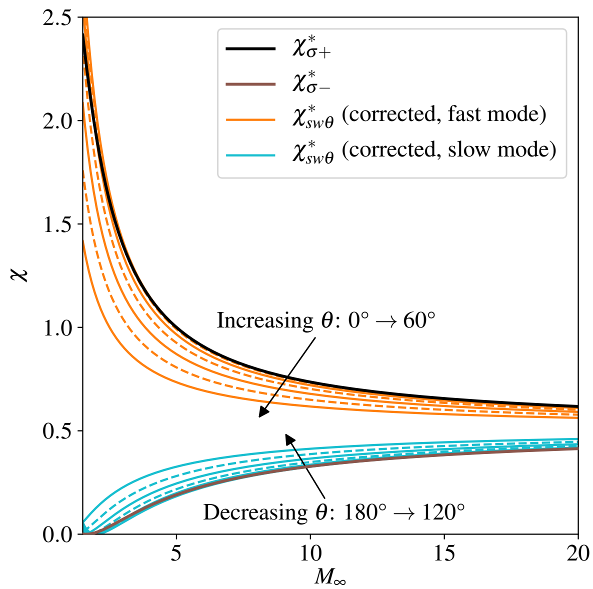

Figure 6.6: Effect of acoustic disturbance angle on pitot transfer functions for fast- and slow-mode waves. Computed using sign-corrected version of Stainback and Wagner's equation [$\chi^*_{sw\theta}$] and compared to flow-parallel limiting cases from Chaudhry and Candler [$\chi^*_{\sigma\pm}$]. $\Delta\theta = 10°$ between adjacent curves.

spectral representations are discussed in Appendix B. The results of applying this methodology to the data are in Section 6.4 along with the corresponding FLDI spectra, converted per Section 6.3.

6.3 Conversion of FLDI Data

Tunnel D has a square cross-section, and it is assumed that each wall radiates planar acoustic waves, uncorrelated but from the same spectrum and with the same preferred orientation θ. This assumption will be revisited later. The top and bottom walls have wavevectors constrained to the xy plane ($\alpha = \theta, \beta = 0$), while those from the left and right walls are in the xz plane ($\alpha = 0, \beta = \theta$).

To recover the spectrum of a single one of these four superimposed plane waves, we use:

$$PSD\left\{n'\right\} = \frac{PSD\left\{\Delta\Phi\right\}}{\left|\sum_{i=1}^{M} h_i(f)\right|^2} \tag{6.9}$$

Here, $M = 4$, $h_1 = h_2 = h_{xy}$, and $h_3 = h_4 = h_{xz}$:

$$PSD\left\{n'\right\} = \frac{PSD\left\{\Delta\Phi\right\}}{4\left|h_{xy}(f) + h_{xz}(f)\right|^2} \approx \frac{PSD\left\{\Delta\Phi(t)\right\}}{4\left[h_{xy,e}(f) + h_{xz,e}(f)\right]^2} \tag{6.10}$$

In the approximate equality, the oscillatory exact response functions are replaced with their non-oscillatory, always-positive envelope functions per the arguments made in Sections 5.4 and 5.5.

To evaluate Eq. (6.10), the following information is required:

1. The FLDI optical parameters Δx, w_0, λ_L. These are known.
2. The beam integration length L, here equal to the known test section width.
3. The acoustic wave inclination θ. From the measurements of Donaldson and Wallace, this is $\theta = 122$–$128°$ at $M_\infty = 4$.
4. The freestream Mach number and sound speed, M_∞ and a, in order to apply the convective shift Eq. (5.65). These are known; note that there are only two values of M_∞, but a differs at every p_{res} step during each run.

5. The frequency range f. This is generated discretely from the application of Welch's method to the data.

Although Donaldson and Wallace provided data over the full range of p_{res} at $M_\infty = 4$, corresponding data were not obtained for our other condition of $M_\infty = 5$. The assumption is made that the sound source velocity ratio u_s/u_∞ remains the same, allowing the relationship between θ and M_∞ to be determined using Eq. (1.6):

$$\frac{u_s}{u_\infty} = 1 + \frac{1}{M_{\infty,1}\cos\theta_1} = 1 + \frac{1}{M_{\infty,2}\cos\theta_2} \tag{6.11a}$$

$$\Rightarrow \theta_2 = \cos^{-1}\left(\frac{M_{\infty,1}}{M_{\infty,2}}\cos\theta_1\right) \tag{6.11b}$$

This yields $\theta = 115$–$120°$ for $M_\infty = 5$.

Consolidating all these parts, $h_{xy}(f)$ and $h_{xz}(f)$ are shown for the bandwidth of a representative Tunnel D condition in Fig. 6.7. It is important to remember when expressing h in terms of f (instead of $\underline{k}$) that the response function will be shifted differently for every combination of M_∞ and a. Defining an average enveloped response function $h_{av,e} \equiv \frac{1}{2}(h_{xy,e} + h_{xz,e})$, then Eq. (6.10) becomes:

$$PSD\{n'\} \approx \frac{PSD\{\Delta\Phi\}}{16h^2_{av,e}(f)} \tag{6.12}$$

$h_{av,e}$ is also shown in Fig. 6.7. Note that once $h_{xz,e}$ rolls off, the response is dominated by $h_{xy,e}$, so the numerical issues caused by $h_{xz,e}$ alone (Section 5.5) are not a concern. This is equivalent to stating that $h_{xy,e} >> h_{xz,e}$ over most of the k range, so that $h_{av,e} \approx \frac{1}{2}h_{xy,e}$. Returning to the assumption made earlier, although the square test section has four-fold rotational symmetry, the nozzle leading to it is two-dimensional and so the nature of the boundary layers on the vertical and horizontal walls of the test section may not be the same*. However, the dominance of $h_{xy,e}$ means that the FLDI is mainly responding to the waves radiating from the top and bottom walls only, and the interpretation will remain largely unchanged even in the extreme case of laminar boundary layers on the left and right walls (i.e. completely non-radiating). On the other hand, the pitot data would be affected by a breakdown in this assumption of four equally-radiating walls, as our model assumes an axisymmetric pitot transfer function.

*Personal communication from J. W. Hofferth

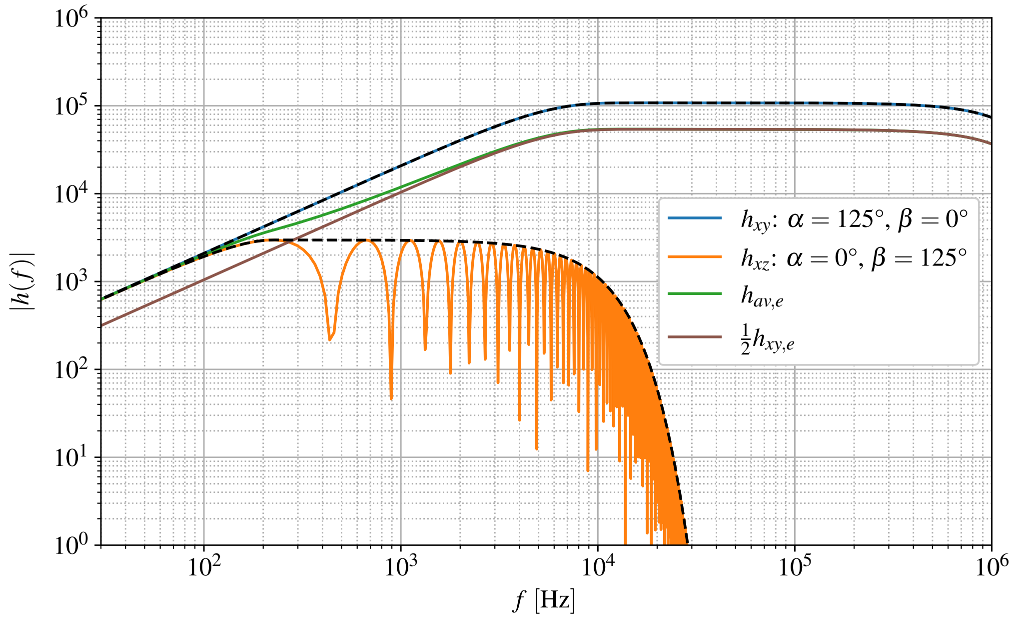

Figure 6.7: Example FLDI response functions $h(f)$ computed for Tunnel D conditions: Run 185, Burst 10 ($M_\infty = 4$, $p_{res} = 229\,\text{kPa}$, $a = 172\,\text{m}\,\text{s}^{-1}$). Envelopes $h_{xy,e}$ and $h_{xz,e}$ shown with dashed black lines.

With $PSD\{n'\}$ obtained via the above process, the Gladstone-Dale and isentropic relations can be used to convert to density and static pressure (in both absolute and relative terms):

$$PSD\{n'\} \underset{\times K^2}{\overset{\div K^2}{\rightleftharpoons}} PSD\{\rho'\} \underset{\times \bar{\rho}^2}{\overset{\div \bar{\rho}^2}{\rightleftharpoons}} PSD\{\rho'/\bar{\rho}\} \underset{\div \gamma^2}{\overset{\times \gamma^2}{\rightleftharpoons}} PSD\{p'/\bar{p}\} \underset{\div \bar{p}^2}{\overset{\times \bar{p}^2}{\rightleftharpoons}} PSD\{p'\} \tag{6.13}$$

Again, please refer to Appendix B for more discussion on these conversions.

6.4 Comparison of Pitot and FLDI Data

The conversion algorithms from Section 6.2 and Section 6.3 were applied to the full sets of pitot and FLDI data, respectively. Comparisons of selected individual conditions are shown in Fig. 6.8. In these plots, the dark shades show the bandwidths where each conversion is considered reliable, while the rest of the spectrum is shown in a lighter shade. Fig. 6.8a is an example of a very good match, where both the magnitude and shape of the spectra agree closely, even beyond the reliable bandwidths selected for each instrument. Fig. 6.8b shows the discrepancies that can arise due to the spectral distortion from the laser noise peak (for the FLDI) and the shock stand-off resonance (for the pitot). However, Fig. 6.8b also shows that although the reliable bandwidths do not overlap, it appears they should connect quite smoothly via extrapolation. This concept is used in the analysis of the full dataset that follows.

All spectra corresponding to each run are converted and overlapped in Fig. 6.9, which gives an example for each M_∞ (there was little shot-to-shot variation observed between runs at the same M_∞). Power-law fits are used to extrapolate the signal from each instrument across the gap in useful bandwidths. For most freestream conditions, these extrapolations match quite closely in terms of both slope and magnitude. The worst mismatches are found for the two lowest p_{res} conditions at $M_\infty = 5$ where, at the midpoint of the bandwidth gap, the FLDI extrapolations are larger in magnitude by about a factor of 3. These two conditions are the lowest-density of all remaining datasets (recall that the lowest p_{res} condition was discarded for both M_∞ as discussed in Section 6.1). Again it is possible that the mismatch may be due to simultaneous overestimation of the FLDI signal (being close to the noise floor) and underestimation of the pitot signal (based on the previously-discarded signal showing clearly anomalous pitot response due to low absolute density).

For each M_∞, sensitivity studies were performed for the chosen wave angle θ within the ranges given in Section 6.2, with little difference observed.

6.5 Conclusions

This chapter has presented methods for converting both pitot and FLDI raw data into static pressure fluctuations, by using a consistent model that assumes acoustic plane waves radiated from the wall boundary layers with a preferred orientation. The pitot conversion process involved a unification of the widely-used classical theory of Stainback and Wagner with the modern DNS work of Chaudhry and Candler; note that the *ad hoc* factor of 2 of Stainback and Wagner was not used in the low-frequency limit. Conversion of the FLDI data employed the novel method developed earlier in this thesis (Chapter 5).

This approach gave good agreement between spectra for most run conditions, although extrapolation was required due to the reliable bandwidths of each instrument not overlapping. This analysis provides a proof-of-concept for recovering quantitative spectral information from an optical technique. It is to be emphasized that the main outcome of this campaign is the development of the methodology for interpreting FLDI freestream measurements; the results are not yet proven enough to be used to quantify the facility performance.

It is clear that the useful bandwidth recoverable from FLDI is restricted at both the low and high frequency limits of the full sampled range. At the low end, the strong laser intensity noise overwhelms the signal response, and distorts the spectral shape even quite far from the noise maximum. At the high end, the true signal again appears to fall below the laser noise floor, and additional discrete electronic noise peaks become numerous. As discussed further in Appendix D.2, the laser used is actually considered “low noise” and advertised as suitable for interferometric applications. However, the low densities encountered in hypersonic wind tunnels still lead to poor signal-to-noise ratios.

The fundamental theory used to interpret the FLDI data as quantitative freestream static pressure fluctuations has proven reliable, using pitot as a reference, with close matches found in both magnitude and spectral shape. Further studies are recommended, perhaps coupled with simulation work, because as shown by Chaudhry and Candler even the interpretation of pitot data in hypersonic flows is not yet fully understood. Nevertheless, if the noise floor issues can be overcome with better-designed optics and acquisition systems, then FLDI has the potential to yield quantitative spectra well beyond the bandwidths reachable with pitot probes.

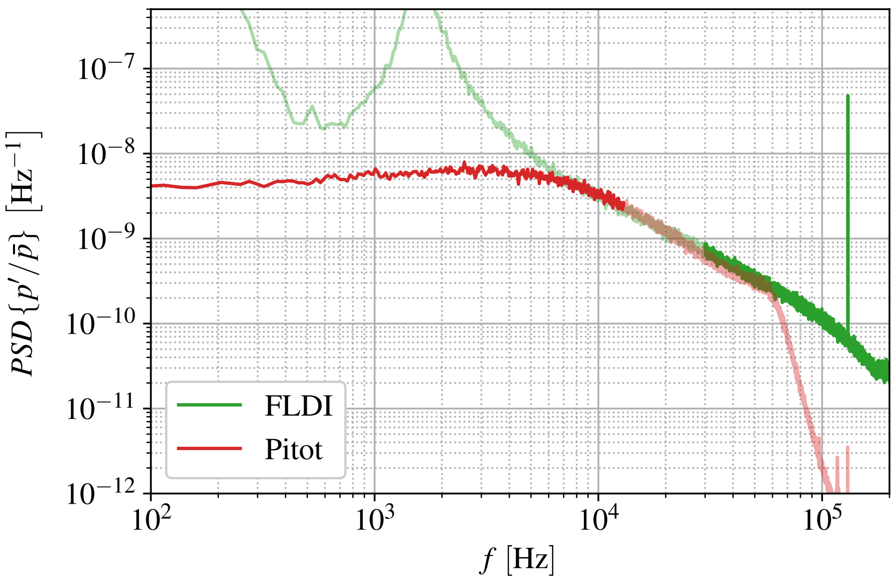

(a) Run 183, Burst 4: $M_\infty = 5$, $p_{res} = 102.7$ kPa

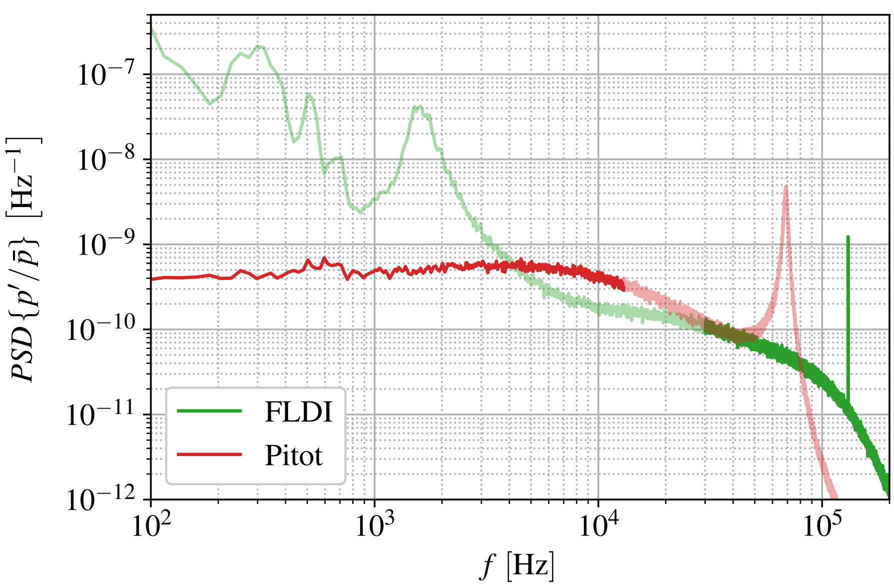

(b) Run 186, Burst 8: $M_\infty = 4$, $p_{res} = 206.8$ kPa

Figure 6.8: Comparison of converted FLDI and pitot spectra for two example cases with differing flow conditions.

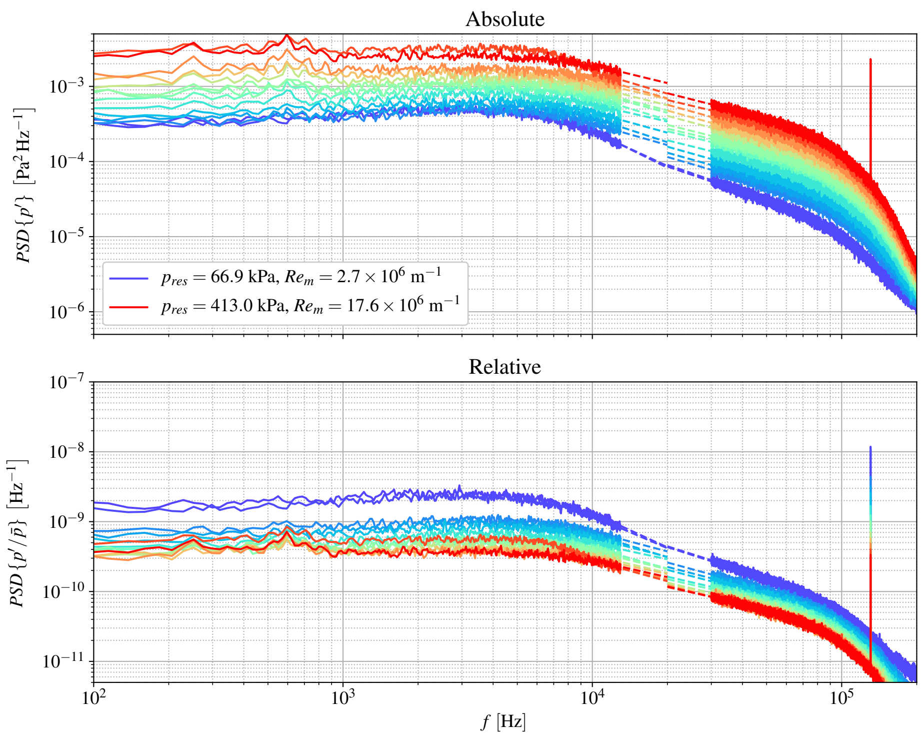

(a) Run 186: $M_\infty = 4$

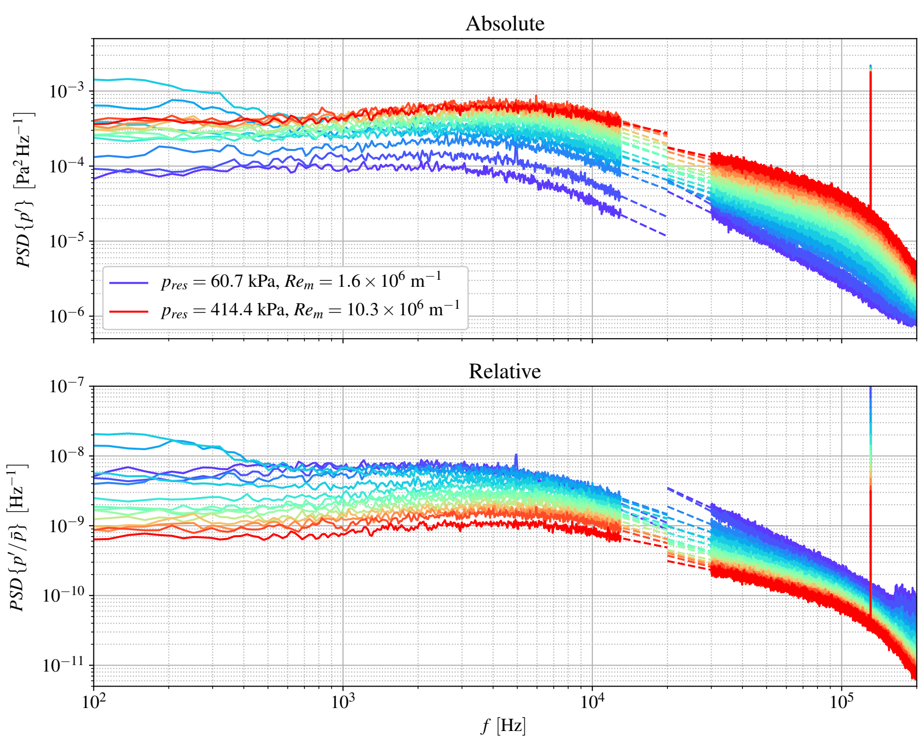

(b) Run 184: $M_\infty = 5$

Figure 6.9: Comparison of converted FLDI and pitot spectra for all conditions within a run. One example run given for each Mach number. Dashed lines indicate power-law extrapolations beyond the useful bandwidth of each instrument.

Chapter 7

RESULTS: HIGH-ENTHALPY EXPERIMENTS

This chapter presents the results of the experimental campaign conducted using FLDI with optical arms to measure the HET freestream. The campaign comprised 66 shots, tabulated in Table E.1; for information on the design of the test conditions, please see Section 2.4.

7.1 FLDI Signal Structure

A representative FLDI signal for an HET shot is shown in Fig. 7.1, along with the corresponding pitot trace. The same features are visible in each signal, albeit with a time lag since the pitot probe is mounted far enough downstream that its bow shock does not interact with any part of the optical arms. Specifically, the probe face is approximately 75 mm downstream of the FLDI beam centerline.

Referring to the generic x–t diagram in Fig. 2.2, the quiescent State 5 is processed to State 6 by a shock wave. In the pitot this manifests as a sudden increase in the time-averaged p_0; in the FLDI the shock appears as a sharp spike when viewed at this timescale, although it was shown previously that this shock signal has its own detailed internal structure (Section 4.3). The shocked State 6 is still relatively quiet, and is separated from State 7 (the test gas) by a contact surface. However, both measurements capture an indication of non-ideal interface effects. These interface effects generate a strong FLDI response which in some cases (such as here) approaches the magnitude of the shock signal. For Shot 1889, the interface is between two dissimilar gases, air and CO_2, which are also at quite different densities ($\rho_6 = 0.011\,\mathrm{kg\,m^{-3}}$ and $\rho_7 = 0.061\,\mathrm{kg\,m^{-3}}$, respectively). Instabilities and mixing at this interface can therefore lead to strong refractive index gradients.

The beginning of the quasi-steady test time Δt at State 7 is taken to be the end of this interface region. In the pitot, the test time appears as a relatively constant plateau. Test time is terminated by the arrival of the leading characteristic of the reflected expansion wave. Finally, the test gas is separated from the denser driver gas by the primary contact surface, the arrival of which is clearly visible in both the pitot and FLDI signals.

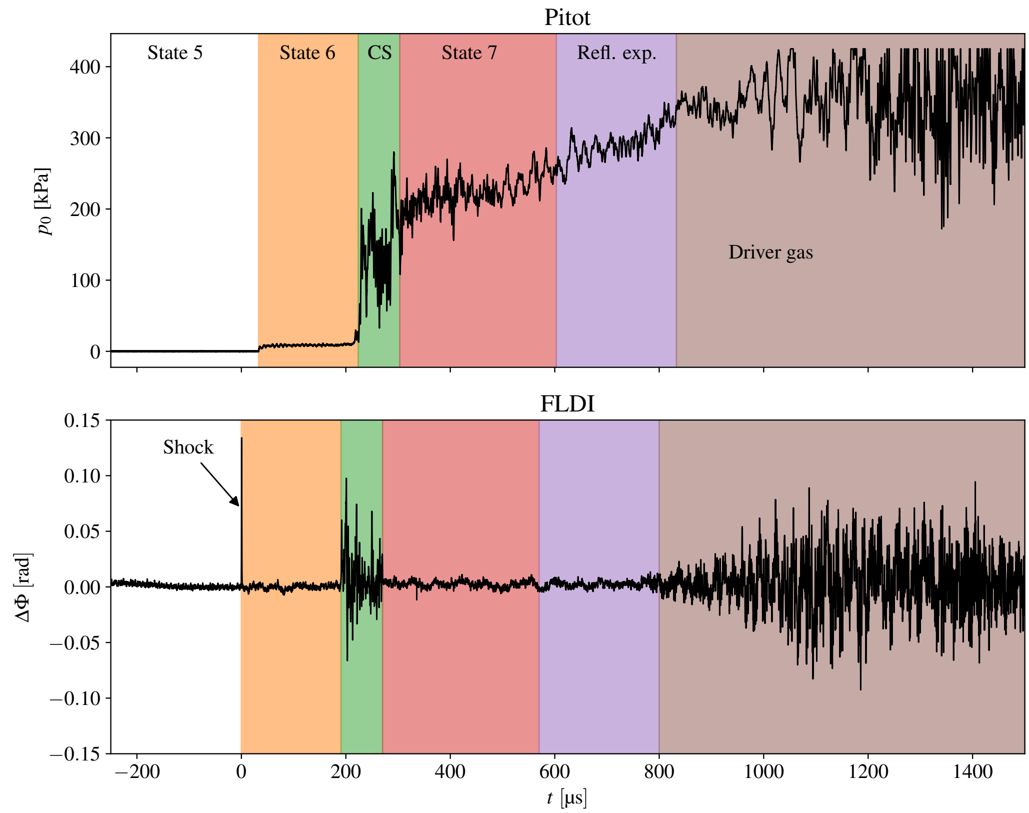

Figure 7.1: Pitot and FLDI signals for Shot 1889. The various flow states are indicated (CS = secondary contact surface, refl. exp. = secondary expansion wave, reflected from primary contact surface).

7.2 Post-Processing

The $\Delta\Phi$ signal from each shot is manually inspected to determine the beginning of test time. As discussed above, this is usually quite apparent due to the distinct nature of the signal during the broadened contact region. It is usually simple to determine the start of test time, but there is no clear indication of its termination, especially in the FLDI signal. Instead, some fraction of the ideal test time is used to extract the portion of the signal used for further analysis. In this work, 70% of the test time calculated with LETS was used, although the results showed little sensitivity to the exact fraction used. The extracted test time signal is used to compute a spectrum $PSD\,\{\Delta\Phi\}$ using Welch's method. The lower frequency limit scales like $1/\Delta t$, and so is roughly 10 kHz. This raw spectrum is subsequently converted and normalized (see Section 7.3).

Note that this entire campaign was performed using a double FLDI. This means that every dataset actually contains two nearly-identical $\Delta\Phi(t)$ signals, offset by

$\lesssim 1\,\mu$s. For this study, these signal pairs are used to give some quantification of measurement error, i.e. by applying the same post-processing to each and observing any differences.

7.3 Data Normalization

7.3.1 Assumptions & Modeling

To compare test-time data at different conditions, the FLDI signal must be normalized. Since the objective of this campaign is to detect trends in freestream noise, we convert $\Delta\Phi$ to $\rho'/\bar{\rho}$. This procedure was also performed to the Tunnel D data in Chapter 6. In the case of HET, additional assumptions must be made about the fluctuating field, compared to conventional tunnels where the nature of the noise sources are well-documented in the literature (see Section 1.3). The freestream noise in HET may be due to a combination of several sources: acoustic waves originating both from the driver and radiating from the wall boundary layers, vorticity from the primary diaphragm petals, knife-blades, and secondary diaphragm wire-cross, as well as other disturbances of uncertain nature from mixing at the secondary contact surface. The frequency content and relative magnitudes of each of these noise sources are unknown.

Currently, the only class of flow-field for which quantitative FLDI inversion can be performed is sinusoidal plane waves. The simplifying assumption is thus made here that the freestream noise can be modeled as comprising plane acoustic waves, propagating parallel to the freestream flow direction. The justification for this is based on several considerations. Firstly, Paull and Stalker considered the main source of noise in expansion tubes to be acoustic waves from the driver. These were modeled using a potential function ϕ:

$$\underline{u} = \underline{u}_0 + \nabla\phi \tag{7.1a}$$

$$p = p_0 - \rho_0 \frac{\partial \phi}{\partial t} \tag{7.1b}$$

$$\phi = J_0(\lambda r) \exp\left[i\omega \left(t \pm \frac{\beta x}{a} \right) \right] \tag{7.1c}$$

$$\beta = \sqrt{1 - \left(\frac{\lambda a}{\omega} \right)^2} \tag{7.1d}$$

Here, (x, r) are the axial and radial tube directions, a is the sound speed, λ is the radial wavelength, ω is the frequency, β is the dispersive term, and J_0 are zeroth-order

Bessel functions of the first kind. In their terminology, "longitudinal waves" are those without radial dependence ($J_0(0) = 1$), while all higher-order solutions with radial modes are "lateral waves"—although these lateral waves still have an axial component of vibration, as shown by Eq. (7.1c). It is shown that only lateral waves are transmitted into the test gas, and although their frequency content is shifted by the unsteady expansion (i.e. β is modified), at State 7 the fluctuating pressure (and hence density) field maintains the radial dependence $J_0(\lambda r)$. Experimental data from Paull and Stalker showed that the first-order mode was dominant. This mode only has a single radial node whose position is only a function of the constant tube radius. In the current experiments, the flow cutters extract only the central portion of this wave in the z-direction, and the beam width in the y-direction is small (the geometry is illustrated in Fig. 7.2). Hence, these waves from the driver, as seen by the FLDI, are assumed to be approximated by purely plane waves propagating in the x-direction.

Secondly, the acoustic radiation from the walls is assumed to have a preferred orientation, as seen widely in both experiments and simulations on conventional facilities (Section 1.3). The cumulative effect of the wall boundary layers is approximated in the vicinity of the centerline (where the FLDI is sensitive) as planar x-waves, the same geometry as from the driver. In this analysis, we examine the effect of acoustic waves, neglecting entropic and vortical disturbances. Because the goal is to simply compare between experiments performed on the same facility and measured using the same FLDI, rather than extract quantitative noise levels, this is considered an acceptable approach.

7.3.2 Inversion Process

Similarly to Eq. (6.12), the conversion is performed in a spectral sense:

$$PSD\{n'\} \approx \frac{PSD\{\Delta\Phi\}}{h_x^2(f)} \tag{7.2}$$

$PSD\{\rho'/\bar{\rho}\}$ is then recoverable from $PSD\{n'\}$ using Eq. (6.13). This requires the use of the time-averaged density and Gladstone-Dale constant of State 7, which is computed using LETS.

Note that $h_x(\underline{k})$ is the same for all conditions, because it only depends on the optical parameters and the integration length, which were unchanged throughout the campaign. However, $h_x(f)$ is not the same, as the conversion from wavenumber to frequency depends on M and a via the dispersion relation Eq. (5.65). These two values are also obtained using LETS.

This inversion process is applied to the extracted Δt segment, and also to a portion of the pre-shot quiescent signal. This is done to determine the noise floor in the same dimensions as the useful signal. Example spectra from both signals are shown in Fig. 7.3.

Note that the noise floor spectrum is more resolved than the test time. This is because a longer record length is available for the pre-noise spectrum, while the test time is limited by Δt. The spectra for all shots look qualitatively similar to this; broad peaks in the flow signal are attributable to poorly-resolved versions of the sharp peaks in the noise floor, believed to originate from the laser intensity noise. There is no evidence of frequency-focusing at other bands as postulated by Paull and Stalker.

For comparisons of fluctuation magnitude across many different conditions, it is useful to have a single-valued metric in addition to the full spectral representation. An RMS value can be computed for a bandwidth $[f_1, f_2]$:

$$RMS\left\{\rho'/\bar{\rho}\right\} = \sqrt{\int_{f_1}^{f_2} PSD\left\{\rho'/\bar{\rho}\right\} \mathrm{d}f} \tag{7.3}$$

For this work, $[f_1, f_2] = [10\,\mathrm{kHz}, 1\,\mathrm{MHz}]$ is selected since above this the signal falls below the noise floor.

A signal-to-noise ratio (SNR) is also defined:

$$SNR \equiv \frac{RMS_S\left\{\rho'/\bar{\rho}\right\}}{RMS_N\left\{\rho'/\bar{\rho}\right\}} \tag{7.4}$$

where RMS_S and RMS_N are computed from the test-time and noise floor data, respectively.

7.4 Results

The different test conditions are categorized by gas combination and primary diaphragm thickness. Again, these conditions are listed in Table E.1. For consistency across different plots, a universal color scheme is used throughout this section and given in Fig. 7.4.

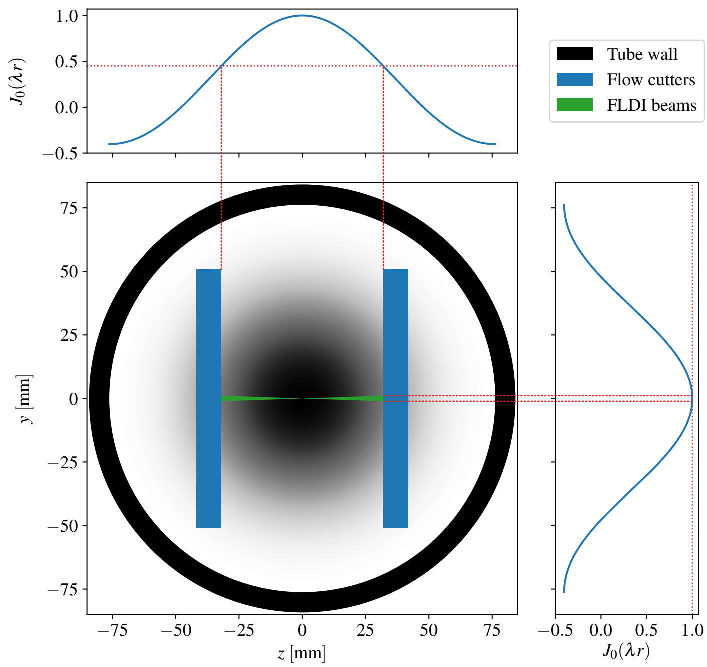

Figure 7.2: Geometry of the FLDI beams relative to the first radial acoustic mode of Paull and Stalker.

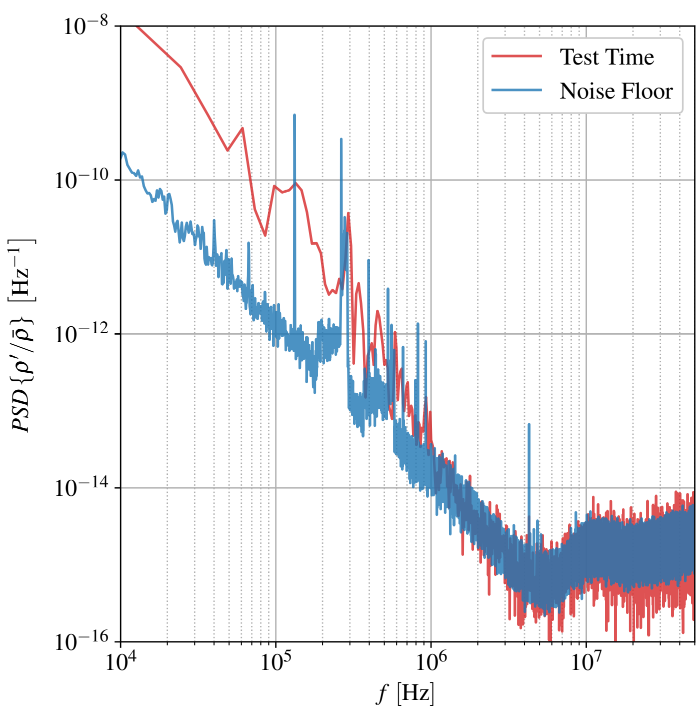

Figure 7.3: Test-time and pre-shot noise spectra for Shot 1887.

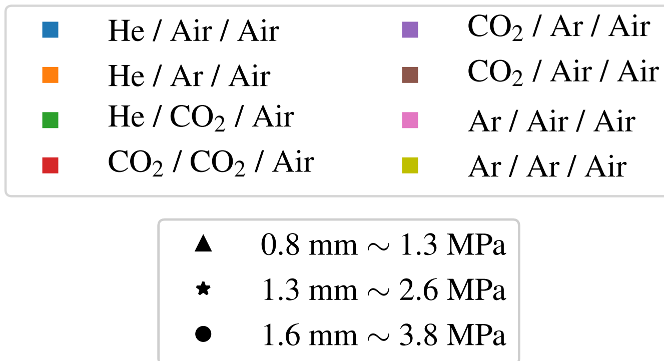

Figure 7.4: Universal legend for HET results. Colors represent different gas combinations; symbols represent primary diaphragm thicknesses and thus burst pressure.

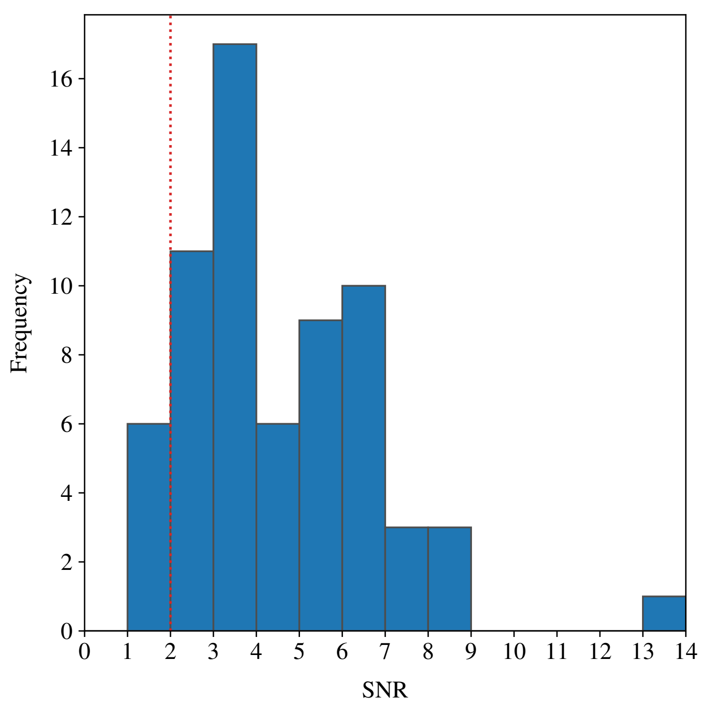

Figure 7.5: Signal-to-noise ratios for all datasets in HET campaign. Threshold $SNR = 2$ indicated with red dashed line.

7.4.1 RMS Trends

The signal-to-noise ratios computed using Eq. (7.4) for all shots are summarized in Fig. 7.5. The majority of shots have $3 < SNR < 9$. An arbitrary threshold of $SNR < 2$ was chosen to designate RMS data as being of suspect quality; the 6 conditions for which this is the case are omitted from the plots. The RMS fluctuation trends can be plotted as functions of many different variables, here only the ones most directly relevant to the assumed noise modes are shown. The transmission of the acoustic waves from the driver gas is hypothesized to be controlled by a_3/a_2, while the acoustic radiation from the wall boundary layer should depend primarily on $Re_{m,7}$. The RMS data for all conditions are shown as functions of these two parameters in Fig. 7.6.

Fig. 7.6a shows an overall decreasing trend in relative noise as a_3/a_2 increases, although there is a lot of scatter. Note that this is opposite to the trend predicted by Paull and Stalker, who state that the transmission should increase smoothly up to $a_3/a_2 = 1$, above which there should be a strong and discontinuous increase in the transmission of some frequency components. The helium driver data form one cluster, while below these, the heavy driver data form another cluster that appears to follow a similar decreasing trend. Again, this is in opposition to Paull and Stalker's theory, which states that the heavy driver conditions should be noisier for the same a_3/a_2 because the frequency cut-off is higher, i.e. more of the original driver noise bandwidth is transmitted through to the test gas.

Paull and Stalker's theory does not discuss the origins of the noise in the processed driver gas (State 3), rather it merely proposes a functional form. Their expansion tube used a free-piston driver, whereas HET simply fills the driver section from a high-pressure source until the diaphragm bursts. Possible noise sources for HET are thus the diaphragm rupture process itself, and possibly the turbulent jet mixing that results from the small-diameter inlet used to fill the driver. An obvious correlation to look for is between noise levels and primary diaphragm thickness. Fig. 7.7 extracts the He/Air/Air conditions only from Fig. 7.6a, as these are the most numerous set of conditions sharing the same gas compositions. These data show no grouping by diaphragm thickness.

Fig. 7.6b shows that the noise has little dependence on $Re_{m,7}$, with a slight downwards trend observable for both groupings (helium and heavy drivers). Because the relative noise $\rho'/\bar{\rho}$ is being considered, this decrease is in line with acoustic radiation results from conventional tunnels (see the relative spectra from the Tunnel D results,

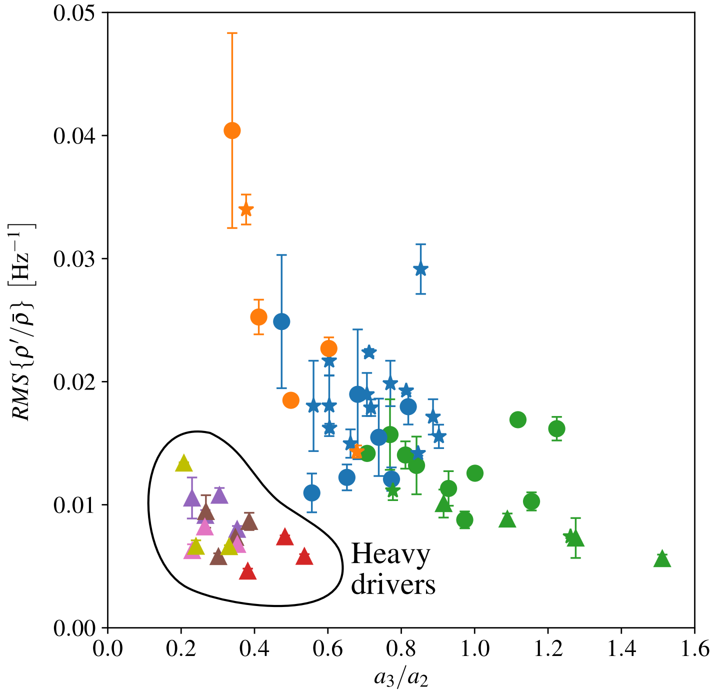

(a) $RMS\{\rho'/\bar{\rho}\}$ vs. a_3/a_2

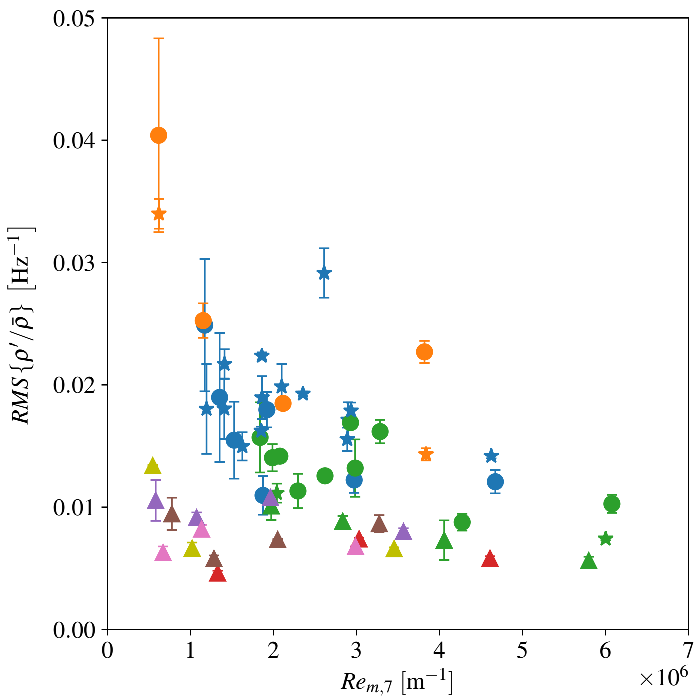

(b) $RMS\{\rho'/\bar{\rho}\}$ vs. $Re_{m,7}$

Figure 7.6: Estimated RMS freestream density fluctuations for all conditions in HET FLDI campaign. For color and symbol meanings, refer to Fig. 7.4.

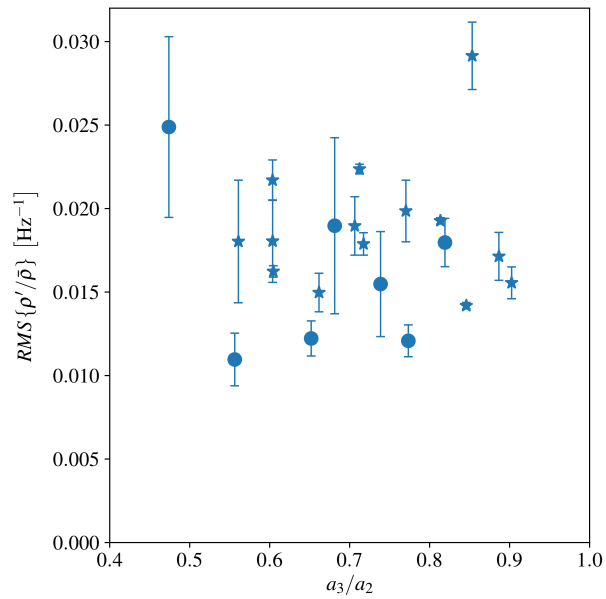

Figure 7.7: Estimated RMS freestream density fluctuations for He/Air/Air conditions only. Symbols represent different diaphragm thicknesses, per Fig. 7.4.

Fig. 6.9, and various results from the literature in Section 1.3). In their discussion of expansion tube noise, Erdos and Bakos (1994) used $Re_m \lesssim 7 \times 10^5\ \mathrm{m}^{-1}$ as the critical Reynolds number below which the boundary layer was found to usually be laminar per heat-flux measurements. According to this metric, the majority of these conditions are not laminar. Because $Re \sim \rho$, making FLDI measurements at substantially lower $Re_{m,7}$ would lead to severe SNR issues with the current laser noise floor.

As both discussed by other authors in Section 1.4, and observed in the present signal in Section 7.1, the secondary diaphragm rupture leads to complex flow processes resulting in a broadened interface region, likely with substantial mixing and instabilities. One such parameter that could be used to look for the influence of this on the freestream noise is the Atwood number across the interface, $A \equiv (\rho_7 - \rho_6)/(\rho_7 + \rho_6)$. However, it is expected that instabilities in the contact surface would generate vorticity- and perhaps also entropy-mode noise (the latter due to temperature spottiness since $T_6 \neq T_7$). Recall that the FLDI inversion process of Section 7.3 made the assumption that the freestream noise was dominated by acoustic-mode noise contributions; hence it would not be appropriate to look for trends in noise from the other modes when the underlying post-processing assumptions depend on neglecting these.

7.4.2 Spectral Trends

The spectral data over the useful bandwidth are presented for all shots in Fig. 7.8. The shots are categorized by gas composition using the same color scheme as previous plots. There are no clear peaks in the spectra that are not attributable to underlying spikes from the laser noise floor, particularly with the poor frequency resolution that is a consequence of the short test-time record length.

The reason we might expect to see prominent peaks was discussed in Section 2.3: as a radial-mode acoustic wave is processed by an unsteady expansion, its initially broadband frequency content is "focused" towards a single discrete frequency. This happens for each radial mode, and so the dominant peak should correspond to the first radial mode. Paull and Stalker's theory for this frequency-focusing behavior is shown in Fig. 7.9. The initial broadband wave contains a continuous spectrum of normalized frequencies $\beta_1 \in [0, 1]$. As the strength of the unsteady expansion increases (equivalent to a_7/a_2 decreasing), all frequency components converge towards a single post-expansion value β_2. However, full focusing only happens in the strong expansion limit, $a_7/a_2 \to 0$. Paull and Stalker provided some evidence for this frequency focusing, based on pitot data from two shots. For one of these conditions with $a_7/a_2 = 0.67$ they reported no focusing, while for the other at $a_7/a_2 = 0.46$ they claimed to observe focusing consistent with first-order radial-mode waves. However, this data was not analyzed spectrally, rather the authors simply measured the approximate period of the most prominent disturbance in the temporal signal, and compared this with the period corresponding to the computed β_2; this method cannot be considered conclusive. The "unfocused" signal appeared more stochastic with no obvious dominant period in the temporal representation. These two values of a_7/a_2 are indicated on Fig. 7.9a. The range of a_7/a_2 for all shots in this campaign is also shown. The conditions are in a region where both the theory and Paull and Stalker's experimental results imply frequency focusing should be slight, so it is not surprising that no strong peaks are observed in these data.

Note that the curves in Fig. 7.9a are specifically for test gases with $\gamma = 7/5$, which is not applicable for all the shots in the campaign. Fig. 7.9b shows the sets of focusing curves for the three values of γ used in the campaign: monatomic gases require stronger expansions for the same amount of focusing, while CO_2 requires slightly less, but the experimental a_7/a_2 range is still high enough that significant frequency focusing would not be expected. HET is capable of significantly stronger expansions, but the conditions used here had higher values of p_5 than usual, in order

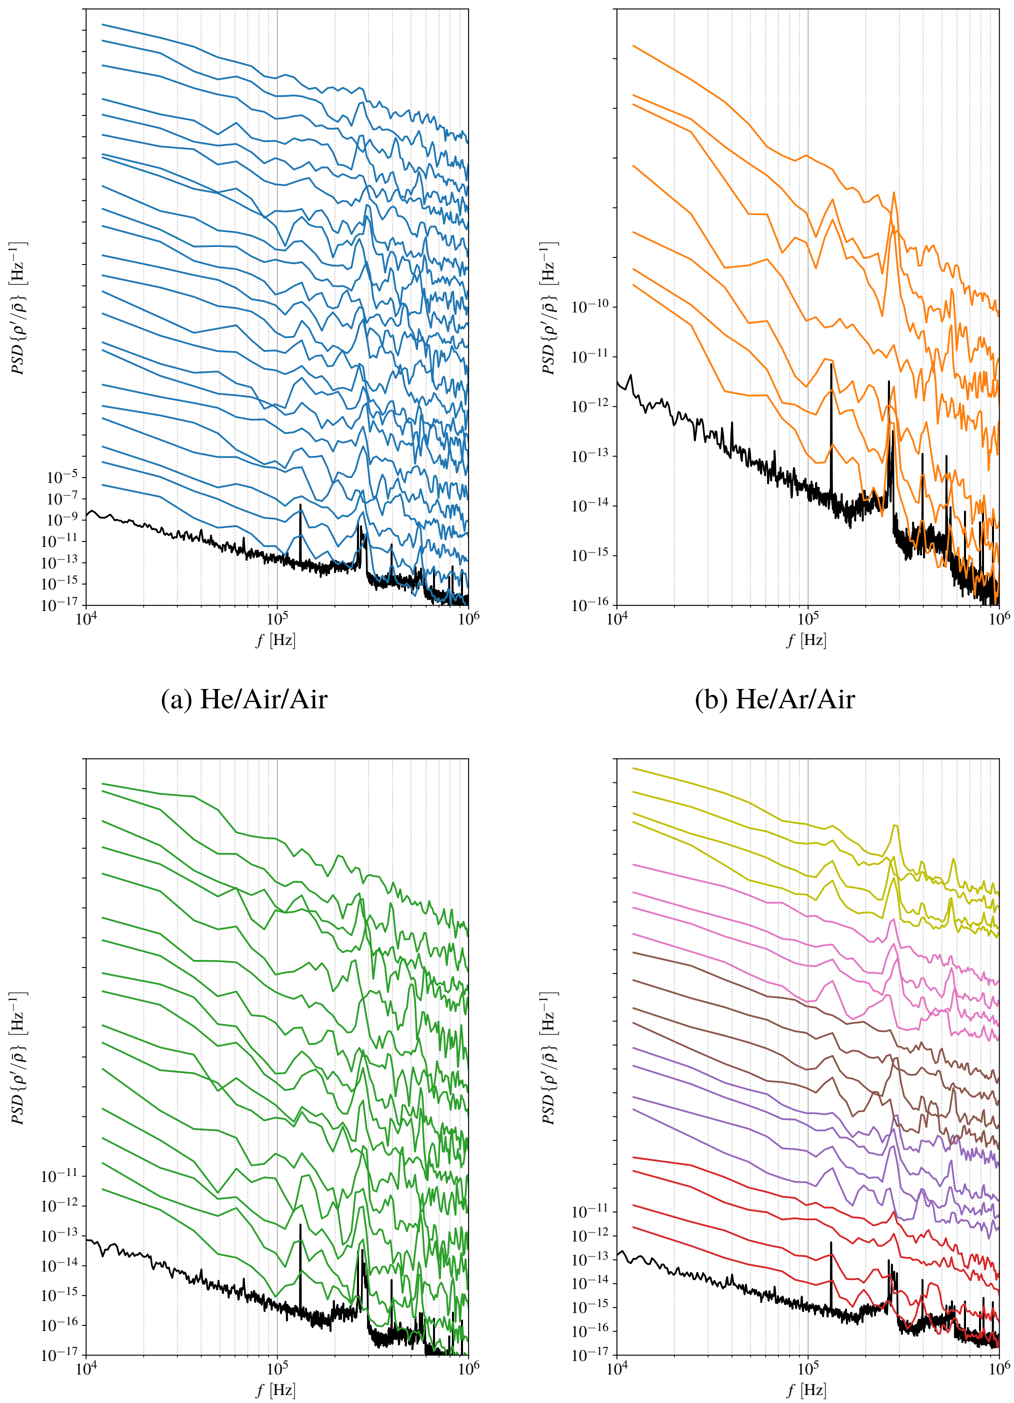

(a) He/Air/Air

(b) He/Ar/Air

(c) He/CO_2/Air

(d) Heavy drivers

Figure 7.8: Spectra for all conditions, grouped by gas combination. The noise floor (in black) and bottom-most spectrum are displayed at original magnitude, the other spectra are successively offset by factors of 10. Colors for each gas combination follow Fig. 7.4.

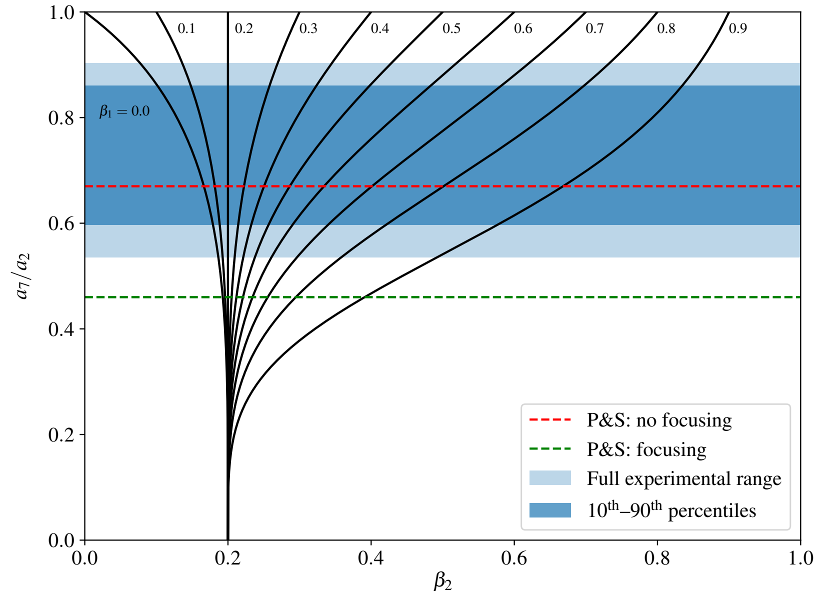

(a) Comparisons with experimental a_7/a_2 ranges from HET campaign, as well as Paull and Stalker's two cases. Focusing curves are for $\gamma = 7/5$ (diatomic gases).

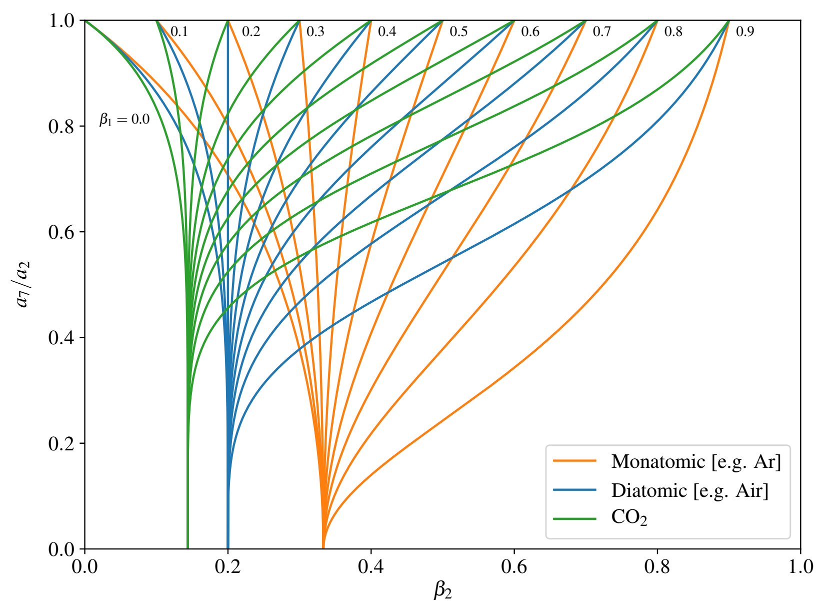

(b) Effect of γ on focusing curves.

Figure 7.9: Frequency focusing behavior across secondary expansion as hypothesized by Paull and Stalker.

to keep the test-time freestream density ρ_7 high enough to avoid SNR issues. Again, if the noise floor issues can be improved, FLDI may be able to be used to acquire useful data for lower a_7/a_2. The successful observation of frequency focusing at these conditions would be strong evidence towards Paull and Stalker's theories.

7.5 Conclusions

Compared to Tunnel D, HET poses a greater challenge to the effective application of FLDI for several reasons: lower densities, shorter test times, and less well-understood flow-field. Nevertheless, this chapter demonstrates that the use of FLDI on expansion tubes is still feasible, although performance improvements will be needed in order to draw stronger conclusions from the results.

The first result concerns the interpretation of the raw signal itself. As discussed at length in Section 4.4, the geometry of HET's test section is such that portions of the transmitted shock reflect back upstream and yield erroneous FLDI responses simultaneous with the facility test time. This necessitated the construction of optical arms to shield the outer portions of the FLDI beams, since the nearly-discontinuous refractive index change of a shock wave is not properly filtered out by FLDI. Similar phenomena are likely in other impulse facilities, which tend to share HET's geometrical features of an abrupt area change at the nozzle exit, and recessed window cavities. The key factor leading to the discovery of this behavior was the use of a double FLDI, which gives directional information and enabled us to determine that the anomalous feature was in fact moving upstream. The main unsteady flow features of HET are all identifiable in the FLDI signal, and correlate well with the pitot trace. This allows extraction of the temporal data for the steady test time (state 7), and subsequent calculation of the spectrum.

A methodology was presented for converting the raw spectra into dimensions of relative freestream density fluctuation, similarly to the procedure for Tunnel D in Chapter 6. However, the sources of noise in expansion tubes are less well-understood, whereas conventional blowdown facilities have a fairly comprehensive pool of experimental evidence as to the nature of their noise. This forces us to make more assumptions about the flow-field in order to interpret the FLDI signal for HET, limiting our ability to make conclusive statements about the quantitative noise magnitudes. The method is deemed sufficient for making comparisons between shots performed in the same facility and measured using the same instrument, as done here.

Due to the short record length and laser noise issues, the useful bandwidth of these spectra was found to be approximately 10 kHz–1 MHz. To put this bandwidth in perspective, Parziale et al. (2014) (studying BL instabilities on a 5° half-angle cone in the T5 shock tunnel) took 2–4 mm as the most-amplified wavelength range. Using 2–4 km s^{-1} as a representative range of acoustic-convective velocities attainable for standard HET conditions, and under the flow-parallel acoustic wave assumption used in this analysis, the most amplified frequencies would lie in the band 0.5–5 MHz. Parziale et al. computed RMS density fluctuations over several wider bands of wavelengths.The shots of the present HET campaign have a time-averaged acoustic-convective velocity ($u_7 + a_7$) of 3.0 km s^{-1}, for which the useful bandwidth of these data corresponds to wavelengths 3–300 mm. The closest range to this presented by Parziale et al. is 0.7–100 mm, for which they found $RMS\,\{\rho'/\bar{\rho}\} = 2.3\,\% \pm 0.7\,\%$. Referring to Fig. 7.6, most of the HET shots lie in the band $1\,\% \lesssim RMS\,\{\rho'/\bar{\rho}\} \lesssim 3\,\%$. This would make the HET freestream noise magnitude comparable to that of T5 over a roughly similar bandwidth, although caution should be taken with quantitative comparisons of this nature, given the assumptions inherent in the FLDI conversion process for both sets of experiments.

The overall trends observed in the RMS data were that the relative freestream density fluctuations decreased as the primary sound speed ratio a_3/a_2 increased, and also as the unit Reynolds number $Re_{m,7}$ increased. The primary clustering of data corresponded with the molecular weight of the driver gas. No clear trends were observed with respect to the primary diaphragm thickness, which may simply be due to the relatively small range of burst pressures on HET (1.3–3.3 MPa). In contrast, early expansion tubes often cited as being noisy had much higher burst pressures, or different driver mechanisms. Trimpi (1962) used an H_2/O_2/He combustion driver (burst pressures not given) while Spurk (1965) used a modified gun barrel for the driver, with an area reduction at the primary diaphragm and achieving pressures of 69 MPa in cold He. As discussed in Section 1.4, Norfleet et al. (1966) used He drivers at 69 and 172 MPa. Shinn and Miller III (1978) used a He driver with a steel double-diaphragm bursting at 33 MPa, and Paull and Stalker used a free-piston driver of He or Ar developing burst pressures of 34.5 MPa, also with an area reduction. These are all at least an order of magnitude higher than HET.

While the $Re_{m,7}$ trend (although weak) is consistent with literature results for acoustic radiation in conventional tunnels, the a_3/a_2 dependence is opposite to that postulated by Paull and Stalker. The reason for this is currently unknown, although it should be noted that the authors simply proposed a form of driver noise, then explored the logical consequences of such a form; a mechanism for the origin of such noise was not given. Paull and Stalker propose that most of the noise originates at the driver throat, a feature absent from HET. One possible factor that might contribute to the discrepancy with the predicted trend is real-gas acoustic absorption effects, as discussed by Fujii and Hornung (2001) (see also Section 1.4). Paull and Stalker did not account for absorption; their analysis was restricted to acoustic transmission effects across contact surfaces and unsteady expansions. In addition to the higher frequencies of interest, this high-temperature enhancement of absorption is another point of difference between conventional and impulse facilities. Spectrally, the frequency-focusing effect predicted by Paull and Stalker was also not observed; this is believed to be due to our conditions having insufficiently strong secondary expansions for disturbance focusing to have a measurable effect.

The noise environment in HET is likely to consist of the superposition of several sources, which could be any or all of acoustic, entropic, or vortical modes. It is recommended that further work be done, numerically and experimentally, to better determine the relative contributions of these. This will allow more accurate interpretation of the FLDI data. With regard to the FLDI itself, the same issues encountered with the laser intensity noise in Chapter 6 are exacerbated here. A significant reduction of this noise floor will be required for future application of FLDI to HET, especially if one wishes to investigate the more typical higher-enthalpy conditions that lie at lower densities, as well as extending the bandwidth into the MHz range most useful for transition studies.

Chapter 8

CONCLUSIONS & FUTURE WORK

This final chapter summarizes the contributions made by this thesis, which can be loosely grouped into results pertaining to the understanding of the FLDI instrument itself, and to its application on hypersonic ground-testing facilities, these being one of the key use-cases for FLDI. Following these summaries, recommendations are made for future avenues of research into FLDI.

8.1 Conclusions

8.1.1 Contributions to Understanding of FLDI Instrument

Chapter 3 presents a framework for using paraxial Gaussian optics to design a single FLDI with optimal optical parameters. This is then extended to allow for a double FLDI with the addition of another Wollaston prism. The latter half of the chapter gives a detailed set of practical instructions of the actual assembly, alignment, calibration, and post-processing of an FLDI. This fills a gap in the literature: while the standard FLDI is quite a simple instrument in principle, without prior experience there are some obstacles to getting one operational in practice. If one wishes to make quantitative measurements with FLDI, the instrument should be close as practical to the idealized configuration of the theory used to interpret the data. One emphasis of this chapter is the need to use a beam profiler during set-up, both to ensure that the beams are of correct polarization and equal intensity, and also to obtain accurate measurements of the beam separation Δx and the focal radius w_0, as these are crucial for quantitative applications.

Chapter 4 gives three experimental validations for the quasi-static ray-tracing model introduced by Schmidt and Shepherd (2015). First, a laminar helium jet was used as a steady-flow static phase object. This jet was simultaneously measured using FLDI and a Mach-Zehnder interferometer (MZI). The MZI response is already well-understood, and can be used to recover full-field refractive index fields in the case of axisymmetric phase objects such as this jet. The FLDI response predicted with the ray-tracing model using the known jet flow-field as an input gave a close quantitative match to the experimental data; this result comprised the static validation. Next, an ultrasonic transducer was used to validate the frequency dependence of the FLDI response, which was predicted analytically as a transfer function by Schmidt and Shepherd. The transformation implied by this transfer function was found to collapse all experimental data very well, thereby validating the ray-tracing model dynamically. A key consequence of this result is confirmation that the sensitive length of an FLDI depends on the wavelength of the disturbance being measured. Finally, a more highly-dynamic validation was conducted by measuring propagating shock waves with typical speeds of $\sim 3\,\mathrm{km\,s^{-1}}$. FLDI is capable of producing highly-resolved shock signals with durations in the sub-μs range. Using a shock model that incorporated viscous thickness and curvature, as well as chemical relaxation effects, the ray-tracing theory closely matched both the magnitude and shape of the experimental signal. This validated the continued use of the quasi-static assumption even for the short timescales associated with hypersonic flows.

The work of Chapter 4 necessitated a new numerical implementation of the ray-tracing model. This was made using Python, and features a modular design where the core optical computations are kept isolated from the definition of the refractive index field. Various modules allow this refractive index field to be defined via analytical functions, or by coupling with and interpolating CFD or experimental data in various formats. Examples of all of these use-cases are presented in this thesis. The code is also extensible, in that it is not restricted to use on FLDI, but by modifying the geometry of the rays, other classes of interferometer can also be simulated.

Chapter 5 addresses the central issue with FLDI: the recovery of quantitative information about the interrogated refractive index field $n(\underline{x}, t)$ from the FLDI phase-shift data $\Delta\Phi(t)$. The results of Chapter 4 provided thorough experimental validation that the analytical form of the forward problem $\Delta\Phi(t) = f\left(n(\underline{x}, t)\right)$ is solvable for all $n(\underline{x}, t)$. Building on this ground truth, an analytical form was derived for the forward problem, for a particular class of flow-field: sinusoidal plane waves, of arbitrary orientation and spectral content. This result is a generalization of some special cases previously derived by Schmidt and Shepherd, as well as Settles and Fulghum (2016). Following this, the corresponding inverse problem $n(\underline{x}, t) = f^{-1}\left(\Delta\Phi(t)\right)$ was solved in a spectral sense via the application of Fourier transforms. Specifically, if the spectrum of $\Delta\Phi(t)$ is computed, then the spectrum of $n(t)$ at the FLDI spatial origin can be recovered.

This inversion process was validated using synthetic flow-fields as inputs to the numerical ray-tracing code. Modifications to the process were introduced to deal with numerical issues encountered when the disturbance field has a component of propagation along the optical axis, and also to account for the superposition of an underlying convective flow. This latter point is crucial to bear in mind when interpreting FLDI results: although response functions may remain unchanged in a spatial (wavenumber) sense, they are highly dependent on the convective velocity in a temporal (frequency) sense.

Finally, this new and exact inversion process was compared with the approximate method introduced by Parziale et al. (2012). It was found that due to not accounting for the wavelength-dependence of the FLDI sensitive length, their method will not produce quantitatively-accurate results in the general case, and that their partial correction introduced later in Parziale et al. (2014) actually causes distortion of the spectral slope. Parziale et al.'s method can be modified using terms borrowed from the present theory, but it is still only narrowly applicable to disturbances propagating in the FLDI separation direction only, and does not offer any computational advantage over using the exact result. However, there are use-cases with geometrically-restricted flow-fields where a constant sensitive length can produce accurate results, and a guide to determining appropriate length scales was given.

8.1.2 Contributions to Understanding of FLDI Application

Chapters 6 and 7 discuss the application of FLDI to two different types of hypersonic ground-testing facility: VKF Tunnel D at AEDC, and HET at Caltech, respectively.

A significant portion of the novel contributions in Chapter 6 come from the synthesis of various strands of theory and experiment taken from the extensive literature review of Section 1.3. Firstly, the work of Kovásznay (1953), Laufer (1961), and Morkovin (1957, 1959) demonstrated that conventional blowdown supersonic wind tunnels have freestream noise environments dominated by acoustic radiation from turbulent boundary layers on the facility walls. In the far-field, this radiation can be modeled as plane waves with some preferential inclination relative to the freestream direction, and are slow-mode acoustic waves, i.e. their wavevector has an upstream-facing component. These results have been verified by modern numerical studies, such as those of Duan et al. (2014, 2016). Understanding the nature of the freestream fluctuations is crucial because the quantitative inversion of FLDI results at present still relies on some amount of prior knowledge of the field geometry. Fortuitously, a hot-wire anemometry study was performed on Tunnel D by Donaldson and Wallace (1971), providing accurate facility-specific information about the inclination angles.

Secondly, the present experiment on Tunnel D made simultaneous measurements of the freestream using FLDI and a pitot probe. Despite decades of widespread use, the response of pitot probes to high-frequency disturbances in supersonic flows is not fully understood, as shown by recent numerical and experimental studies (Chaudhry and Candler, 2017; Chaudhry et al., 2019; Duan et al., 2019). Use of a simple pitot transfer function due to (Stainback and Wagner, 1972) remains state-of-the-art despite it not accounting for the strong frequency dependence observed in the aforementioned numerical studies. To the author's knowledge, this thesis presents the first reconciliation of the two approaches by correcting a sign error in Stainback and Wagner's result, while emphasizing that the result is only valid at sufficiently low frequencies (relative to the resonance frequency caused by the pitot probe shock standoff distance).

These two developments, along with the results of Chapter 5, allowed for the FLDI and pitot data from the Tunnel D experiments to each be converted to pressure fluctuations. The FLDI spectra are affected by laser intensity noise, especially in the kHz band, and the pitot probe suffers from resonance peaking at 65 kHz. As a result, the useful bandwidths of the instruments do not overlap. However, a power-law extrapolation across this gap showed that the spectra match well, both in terms of magnitude and slope. Despite these limitations, the experiments on Tunnel D were considered a successful proof-of-concept for obtaining accurate quantitative data from FLDI on an actual facility, and the results show promise that noise spectra could potentially be extended far beyond the resonance- and inertia-limited bandwidth of pitot probes once the noise floor issues are resolved.

Chapter 7 demonstrates a more challenging application of FLDI to an expansion tube, HET. A pair of optical arms was designed and constructed to mitigate issues with reflected shock waves interacting with the outer parts of the FLDI beams. With these installed, a test campaign was conducted to measure the freestream noise during the quasi-steady test time of the facility. A similar approach to data analysis was taken as in Chapter 6, but because expansion tubes have more complex and less well-understood flow-fields and noise sources than conventional blowdown tunnels, more approximations had to be made. This highlights the current key limitation of FLDI: its reliance on adequate prior knowledge of the very flow-field it is trying to measure.

The assumption was made that the HET freestream noise is dominated by acoustic-mode fluctuations originating in the driver gas (as hypothesized by Paull and Stalker (1992)), and also radiated from the tube wall boundary layers similarly to conventional tunnels (as discussed by Erdos and Bakos (1994)). Under these assumptions, the RMS data showed the opposite trend with the primary sound speed ratio a_3/a_2 versus that predicted by Paull and Stalker. The dependence on the unit Reynolds number $Re_{m,7}$ was weak but in line with conventional tunnel results. The spectral data did not show any evidence of the frequency-focusing that should result from the secondary unsteady expansion; however the strength of this expansion as quantified by a_7/a_2 is not strong enough to expect focusing per Paull and Stalker, so this aspect of the theory cannot be confirmed or refuted by the current data.

As with the Tunnel D results, the useful bandwidth of the FLDI data remain restricted far below the photodetector limits due to the laser intensity noise floor. Compared with Tunnel D, the consequences of this are exacerbated due to the lower densities and shorter test times—leading to degraded signal-to-noise ratios and less-resolved spectra more easily contaminated by noise spikes. Nevertheless, it has been shown that FLDI can produce useful data even at such low densities and on such short time-scales—note that the high-dynamic validation case of Section 4.3 also originated from using FLDI on HET. Reduction of this noise floor may allow for significant improvements and could extend the range of the instrument down to the lower densities that correspond to some of the more well-characterized conditions of HET.

8.2 Recommendations for Future Work

8.2.1 FLDI Performance Improvements

Now that the governing equation for FLDI performance is understood and validated, it can be used to methodically improve the instrument. As already discussed, the response function only depends on three optical parameters: λ_L, w_0, and Δx. λ_L is generally restricted to the commercially-available visible laser wavelengths, unless one wishes to deal with the added complexities of working with non-visible lasers while Δx only has an effect when it begins to exceed the disturbance wavelengths. The greatest improvements therefore will come from reducing w_0, equivalent to increasing the focusing angle of the beams. This fact was pointed out by Settles and Fulghum, who expressed it in terms of increasing the optical f-number to shorten the sensitive region.

A larger f-number requires a stronger diverging lens and wider-aperture focusing lenses. It may be more practical to use focusing mirrors instead of lenses above a certain aperture; a Z-shaped FLDI could be designed similarly to schlieren. When using mirrors, care needs to be taken to preserve polarization states. At large f-numbers, the rays in the outer parts of the beam may be at a substantial angle to the centerline, leading to a breakdown of the paraxial assumption. Combined with aberrations from the off-axis layout, the ray-tracing approach as used in this thesis may become less quantitatively accurate if modifications are not made.

Both facility application campaigns in this thesis showed that for low-density hypersonic conditions, the signal-to-noise ratio becomes problematic. At relatively lower bandwidths, the laser intensity noise is the main issue. This noise needs to be reduced substantially, by at least an order of magnitude. While some possibilities have been explored for removing this in post-processing, it both more preferable and feasible to reduce it at the source. This will require the use of some active noise-control technique based on optical feedback. Intensity and frequency stabilization techniques have already been extensively developed (Nocera, 2004; Rollins et al., 2004), one application being interferometry for gravitational wave detection (although the frequencies of interest there are $O\,(1\text{–}100\,\text{Hz})$, much lower than in hypersonic flows).

8.2.2 Further Theoretical Developments

As has been re-iterated several times, a central limitation to extracting quantitative density information from FLDI is the need to know something about the flow-field geometry. If the geometry is known, then in theory any signal can be inverted; however, for arbitrary fields this can only be done numerically, i.e. by simulating the response to a range of fields to find the best fit to the experimental data. For large parameter spaces this approach may not be practical, and having an analytical form for the inversion function is much more efficient. This thesis presented such a function for only one class of flow-field: sinusoidal plane waves. Although this was done in quite a generalized manner, with wide-ranging applicability to acoustic radiation, there is a need to extend the library of analytical functions to other common geometries.

In particular, propagating wavepackets are of importance in boundary-layer stability studies, e.g. Parziale et al. (2015). Being able to quantify the magnitudes of these wavepackets from FLDI data would be a valuable future direction. Related to this is the need to perform inversions in temporal, rather than frequency, space since unlike the acoustic radiation, these are spatially-limited.

8.2.3 Extensions to the Base Instrument

This thesis largely focused on the "standard" single FLDI, although as mentioned in Section 1.2, various groups have already begun working with double FLDI and even instruments with higher numbers of focal pairs. The basic optical principles of DFLDI design were provided in Section 3.2, and the experiments performed on HET in Chapter 7 used a DFLDI. This dataset therefore still contains potentially useful information about the shock and contact surface velocity that could be recovered via cross-correlation.

Recall that the ray-tracing equation and its software implementation are applicable to two-beam interferometers in general, since its specific application to FLDI is only contained within a particular designation of the beam geometry. Hence these tools, that have already been validated quantitatively on FLDI, can also be useful in the design of variants. One such variant is a natural extension to the double, quad (and so on) FLDI systems: a dense 2D array of focal pairs, giving measurements in a sensitive region about a plane rather than just a single point.

The main factor that precludes FLDI from approaching a true point measurement is Δx, the beam separation, which is integral to the instrument function but also makes it sensitive to refractive index gradients in only that direction; as shown in Chapter 5 this greatly complicates the response function particularly when there is a component of motion along the optical axis. It may be advantageous to design instruments that have multiple foci pairs aligned in different directions, rather analogously to what is done with hot-wire anemometers. Consolidating the data carefully from each channel could yield a response function uniform in all directions. However, getting full 3-axis coverage might require sending beams through both orthogonal pairs of facing windows on a typical test section.

BIBLIOGRAPHY

Bathel, B. F., Weisberger, J. M., Herring, G. C., King, R. A., Jones, S. B., Kennedy, R. E., and Laurence, S. J. (2020). "Two-point, parallel-beam focused laser differential interferometry with a Nomarski prism". In: *Applied Optics* 59.2, pp. 244–252. DOI: 10.1364/AO.59.000244.

Benitez, E. K., Jewell, J. S., and Schneider, S. P. (2020). "Focused laser differential interferometry for hypersonic flow instability measurements with contoured tunnel windows". In: *AIAA Scitech 2020 Forum*. Orlando, FL. DOI: 10.2514/6.2020-1282.

– (2021). "Focused laser differential interferometry with contoured tunnel windows". In: *AIAA Journal* 59.2, pp. 419–429. DOI: 10.2514/1.J060081.

Birch, B., Buttsworth, D., and Zander, F. (2020). "Measurements of freestream density fluctuations in a hypersonic wind tunnel". In: *Experiments in Fluids* 61.158. DOI: 10.1007/s00348-020-02992-w.

Bitter, N. P. (2015). "Stability of hypervelocity boundary layers". Ph.D. thesis. California Institute of Technology.

Browne, S., Ziegler, J., and Shepherd, J. E. (2008). *Numerical solution methods for shock and detonation jump conditions*. GALCIT report FM2006.006. Graduate Aeronautical Laboratories, California Institute of Technology.

Ceruzzi, A. P. and Cadou, C. P. (2017). "Turbulent air jet investigation using focused laser differential interferometry". In: *53rd AIAA/SAE/ASEE Joint Propulsion Conference*. Atlanta, GA. DOI: 10.2514/6.2017-4834.

– (2019). "Simultaneous velocity and density gradient measurements using two-point focused laser differential interferometry". In: *AIAA Scitech 2019 Forum*. San Diego, CA. DOI: 10.2514/6.2019-2295.

Ceruzzi, A. P., Callis, B., Weber, D., and Cadou, C. P. (2020). "Application of focused laser differential interferometry (FLDI) in a supersonic boundary layer". In: *AIAA Scitech 2020 Forum*. Orlando, FL. DOI: 10.2514/6.2020-1973.

Ceruzzi, A. P., McManamen, B., and Cadou, C. P. (2021a). "Demonstration of two-point focused laser differential interferometry (2pFLDI) in a Mach 18 flow". In: *AIAA Scitech 2021 Forum*. Virtual event. DOI: 10.2514/6.2021-0983.

Ceruzzi, A. P., Neisess, C., McManamen, B., and Cadou, C. P. (2021b). "Investigation of focused laser differential interferometry (FLDI) sensitivity function". In: *AIAA Scitech 2021 Forum*. Virtual event. DOI: 10.2514/6.2021-1299.

Chaudhry, R. S. and Candler, G. V. (2017). "Computing measured spectra from hypersonic pitot probes with flow-parallel freestream disturbances". In: *AIAA Journal* 55.12, pp. 4155–4166. DOI: 10.2514/1.J055396.

Chaudhry, R. S., Candler, G. V., Gray, K. A., and Schneider, S. P. (2019). “Computations of measured pitot-probe spectra using angled freestream disturbances and comparisons to experiments”. In: *AIAA Scitech 2019 Forum*. AIAA Scitech 2019 Forum. San Diego, CA. DOI: `10.2514/6.2019-2149`.

Collicott, S. H., Schneider, S., and Messersmith, N. L. (1996). “Review of optical diagnostic methods for hypersonic low-noise facilities”. In: *AIAA 34th Aerospace Sciences Meeting and Exhibit*. Reno, NV. DOI: `10.2514/6.1996-851`.

Coronel, S. A., Melguizo-Gavilanes, J., Jones, S., and Shepherd, J. E. (2018). “Temperature field measurements of thermal boundary layer and wake of moving hot spheres using interferometry”. In: *Experimental Thermal and Fluid Science* 90, pp. 76–83. DOI: `10.1016/j.expthermflusci.2017.08.031`.

Coyne, D. (2011). *Design of the elastomeric sealed, high quality, viewports*. LIGO internal working note LIGO-T1100346-v2. LIGO Laboratory.

De Boer, P. C. T. (1963). “Curvature of shock fronts in shock tubes”. In: *Physics of Fluids* 6.7, pp. 962–971. DOI: `10.1063/1.1706852`.

Deiterding, R. (2011). “Block-structured adaptive mesh refinement - theory, implementation and application”. In: *ESAIM: Proceedings* 34, pp. 97–150. DOI: `10.1051/proc/201134002`.

Donaldson, J. C. and Wallace, J. P. (1971). *Flow fluctuation measurements at Mach number 4 in the test section of the 12-Inch Supersonic Tunnel (D)*. AEDC-TR-71-143 (AD723630). Arnold Air Force Station, TN: ARO, Inc.

Duan, L., Choudhari, M. M., and Wu, M. (2014). “Numerical study of acoustic radiation due to a supersonic turbulent boundary layer”. In: *Journal of Fluid Mechanics* 746, pp. 165–192. DOI: `10.1017/jfm.2014.116`.

Duan, L., Choudhari, M. M., and Zhang, C. (2016). “Pressure fluctuations induced by a hypersonic turbulent boundary layer”. In: *Journal of Fluid Mechanics* 804, pp. 578–607. DOI: `10.1017/jfm.2016.548`.

Duan, L. et al. (2019). “Characterization of freestream disturbances in conventional hypersonic wind tunnels”. In: *Journal of Spacecraft and Rockets* 56.2, pp. 357–368. ISSN: 0022-4650, 1533-6794. DOI: `10.2514/1.A34290`.

Dufrene, A. T., Sharma, M., and Austin, J. M. (2007). “Design and characterization of a hypervelocity expansion tube facility”. In: *Journal of Propulsion and Power* 23.6, pp. 1185–1193. DOI: `10.2514/1.30349`.

Dufrene, A. T. (2006). “Design and characterization of hypervelocity expansion tube”. M.S. thesis. University of Illinois at Urbana-Champaign.

Erdos, J. and Bakos, R. (1994). “Prospects for a quiet hypervelocity shock-expansion tunnel”. In: *25th Plasmadynamics and Lasers Conference*. Colorado Springs, CO: American Institute of Aeronautics and Astronautics. DOI: `10.2514/6.1994-2500`.

Fujii, K. and Hornung, H. G. (2001). *A procedure to estimate the absorption rate of sound propagating through high temperature gas*. GALCIT Report FM2001.004. Graduate Aeronautical Laboratories, California Institute of Technology.

Fulghum, M. R. (2014). "Turbulence measurements in high-speed wind tunnels using focusing laser differential interferometry". Ph.D. thesis. Pennsylvania State University.

Furukawa, T., Aochi, T., Saito, K., and Sasoh, A. (2005). "Realization of super-orbital entry flow using expansion tube". In: *ICIASF 2005: International Congress on Instrumentation in Aerospace Simulation Facilities*. IEEE, pp. 337–341.

Gardiner Jr, W. C., Hidaka, Y., and Tanzawa, T. (1981). "Refractivity of combustion gases". In: *Combustion and Flame* 40, pp. 213–219. DOI: 10.1016/0010-2180(81)90124-3.

Gladstone, J. H. and Dale, T. P. (1863). "XIV. Researches on the refraction, dispersion, and sensitiveness of liquids". In: *Philosophical Transactions of the Royal Society of London* 153, pp. 317–343. DOI: 10.1098/rstl.1863.0014.

Goodman, J. W. (2000). *Statistical Optics*. Wiley Classics Library Edition. New York: John Wiley & Sons, Inc. 550 pp. ISBN: 978-0-471-39916-7.

Goodwin, D. G., Speth, R. L., Moffat, H. K., and Weber, B. W. (2018). "Cantera: an object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes." In: *Zenodo*. DOI: 10.5281/zenodo.1174508.

Gromyko, Y. V., Polivanov, P. A., Sidorenko, A. A., Buntin, D. A., and Maslov, A. A. (2013). "An experimental study of the natural noise in the Transit-M hypersonic wind tunnel". In: *Thermophysics and Aeromechanics* 20.4, pp. 481–493. DOI: 10.1134/S0869864313040112.

Grossir, G., Paris, S., Bensassi, K., and Rambaud, P. (2013). "Experimental characterization of hypersonic nozzle boundary layers and free-stream noise levels". In: *51st AIAA Aerospace Sciences Meeting*. Grapevine, TX. DOI: 10.2514/6.2013-1130.

Hameed, A. and Parziale, N. J. (2021). "Focused laser differential interferometer response to a controlled phase object". In: *AIAA Scitech 2021 Forum*. Virtual event. DOI: 10.2514/6.2021-0602.

Harris, C. R. et al. (2020). "Array programming with NumPy". In: *Nature* 585, pp. 357–362. DOI: 10.1038/s41586-020-2649-2.

Hartunian, R. A. (1961). "Shock curvature due to boundary-layer effects in a shock tube". In: *Physics of Fluids* 4.9, pp. 1059–1063. DOI: 10.1063/1.1706450.

Harvey, W. D., Bushnell, D. M., and Beckwith, I. E. (1969). *Fluctuating properties of turbulent boundary layers for Mach numbers up to 9*. NASA Technical Note D-5496. Hampton, VA: Langley Research Center.

Hecht, E. (2002). *Optics*. 4th ed. Reading, MA: Addison-Wesley. ISBN: 978-0-8053-8566-5.

Heinzel, G., Rüdiger, A., and Schilling, R. (2002). *Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows*. Tech. rep. Hannover, Germany: Max-Planck-Institut für Gravitationsphysik.

Hofferth, J. W. and Ogg, D. R. (2018). "Reactivation of VKF Wind Tunnel D by AFRL at AEDC: overview and subsystem checkout results". In: *22nd AIAA International Space Planes and Hypersonics Systems and Technologies Conference*. Orlando, FL. DOI: 10.2514/6.2018-5263.

– (2019). "Reactivation of VKF Wind Tunnel D by AFRL at AEDC: supersonic performance and freestream characterizations". In: *AIAA Aviation 2019 Forum*. Dallas, TX. DOI: 10.2514/6.2019-2856.

Hornung, H. G. (2000). "Oblique shock reflection from an axis of symmetry". In: *Journal of Fluid Mechanics* 409, pp. 1–12. DOI: 10.1017/S0022112099007831.

– (2010). "Deriving features of reacting hypersonic flow from gradients at a curved shock". In: *AIAA Journal* 48.2, pp. 287–296. DOI: 10.2514/1.39993.

Houpt, A. W. and Leonov, S. B. (2018). "Focused laser differential interferometer for supersonic boundary layer measurements on flat plate geometries". In: *AIAA Aviation 2018 Forum*. Atlanta, GA. DOI: 10.2514/6.2018-3434.

– (2019). "Focused and cylindrical-focused laser differential interferometer characterization of SBR-50 at Mach 2". In: *AIAA Aviation 2019 Forum*. Dallas, TX. DOI: 10.2514/6.2019-3383.

Igra, O. and F. Seiler, eds. (2015). *Experimental methods of shock wave research*. New York, NY: Springer International Publishing Switzerland. DOI: 10.1007/978-3-319-23745-9.

Inverarity, G. W. (2003). "Numerically inverting a class of singular Fourier transforms: theory and application to mountain waves". In: *Proceedings of the Royal Society of London A* 459.2033, pp. 1153–1170. DOI: 10.1098/rspa.2002.1076.

Jacobs, P. and Gollan, R. (2016). "Implementation of a compressible-flow simulation code in the D programming language". In: *Applied Mechanics and Materials*. Vol. 846, pp. 54–60. DOI: 10.4028/www.scientific.net/AMM.846.54.

James, C. M., Gildfind, D. E., Lewis, S. W., Morgan, R. G., and Zander, F. (2018). "Implementation of a state-to-state analytical framework for the calculation of expansion tube flow properties". In: *Shock Waves* 28.2, pp. 349–377. DOI: 10.1007/s00193-017-0763-3.

James, C. M., Gildfind, D. E., Morgan, R. G., Jacobs, P. A., and Zander, F. (2013). "Designing and simulating high enthalpy expansion tube conditions". In: *APISAT 2013: Asia-Pacific International Symposium on Aerospace Technology*. Takamatsu, Japan: Japan Soc. for Aeronautical and Space Sciences.

Juliano, T. J., Adamczak, D., and Kimmel, R. L. (2015). "HIFiRE-5 flight test results". In: *Journal of Spacecraft and Rockets* 52.3, pp. 650–663. DOI: 10.2514/1.A33142.

Kiefer, J. H. and Manson, A. C. (1981). "Refractive index change and curvature in shock waves by angled beam refraction". In: *Review of Scientific Instruments* 52.9, pp. 1392–1396. DOI: 10.1063/1.1136779.

Kovásznay, L. S. G. (1953). "Turbulence in Supersonic Flow". In: *Journal of the Aeronautical Sciences* 20.10, pp. 657–674. DOI: 10.2514/8.2793.

Kragh, H. (1991). "Ludvig Lorenz and nineteenth century optical theory: the work of a great Danish scientist". In: *Applied Optics* 30.33, p. 4688. DOI: 10.1364/AO.30.004688.

Laufer, J. (1961). "Aerodynamic noise in supersonic wind tunnels". In: *Journal of the Aerospace Sciences* 28.9, pp. 685–692. DOI: 10.2514/8.9150.

– (1964). "Some statistical properties of the pressure field radiated by a turbulent boundary layer". In: *Physics of Fluids* 7.8, pp. 1191–1197. DOI: 10.1063/1.1711360.

Lawson, J. M., Neet, M. C., Grossman, I. J., and Austin, J. M. (2020). "Static and dynamic characterization of a focused laser differential interferometer". In: *Experiments in Fluids* 61.187. DOI: 10.1007/s00348-020-03013-6.

Lawson, J. M. and Austin, J. M. (2018). "Design of conventional and detonation-driven hypervelocity expansion tubes". In: *AIAA Aviation 2018 Forum*. Atlanta, GA. DOI: 10.2514/6.2018-3566.

– (2021). "Focused laser differential interferometer response to shock waves". In: *Measurement Science and Technology* 32.5, p. 055203. DOI: 10.1088/1361-6501/abdbd3.

Liepmann, H. W. and Bowman, R. M. (1964). "Shape of shock fronts in shock tubes". In: *Physics of Fluids* 7.12, pp. 2013–2015. DOI: 10.1063/1.1711113.

Liepmann, H. W., Narasimha, R., and Chahine, M. T. (1962). "Structure of a plane shock layer". In: *Physics of Fluids* 5.11, pp. 1313–1324. DOI: 10.1063/1.1706527.

Lighthill, M. J. (1952). "On sound generated aerodynamically I. General theory". In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 211.1107, pp. 564–587. DOI: 10.1098/rspa.1952.0060.

– (1954). "On sound generated aerodynamically II. Turbulence as a source of sound". In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 222.1148, pp. 1–32. DOI: 10.1098/rspa.1954.0049.

MacLean, M., Dufrene, A., Wadhams, T., and Holden, M. (2010). "Numerical and experimental characterization of high enthalpy flow in an expansion tunnel facility". In: *48th AIAA 2019 Aerospace Sciences Meeting*. Orlando, FL. DOI: 10.2514/6.2010-1562.

Mai, C. L. and Bowersox, R. D. (2014). "Effect of a normal shock wave on freestream total pressure fluctuations in a low-density Mach 6 flow". In: *44th AIAA Fluid Dynamics Conference*. Atlanta, GA. DOI: 10.2514/6.2014-2641.

Masutti, D., Spinosa, E., Chazot, O., and Carbonaro, M. (2012). "Disturbance level characterization of a hypersonic blowdown facility". In: *AIAA Journal* 50.12, pp. 2720–2730. DOI: 10.2514/1.J051502.

Merzkirch, W. (1987). *Techniques of flow visualization*. AGARDograph. AGARD-AG 302. Neuilly-sur-Seine, France: AGARD.

Mikusheva, A. (2013). *Stationarity, lag operator, ARMA, and covariance structure*. 14.384 Time Series Analysis. MIT OpenCourseWare.

Milonni, P. W. and Eberly, J. H. (2010). *Laser physics*. John Wiley & Sons, Ltd. DOI: 10.1002/9780470409718.

Mirels, H. (1963). "Test time in low-pressure shock tubes". In: *Physics of Fluids* 6.9, pp. 1201–1214. DOI: 10.1063/1.1706887.

– (1964). "Shock tube test time limitation due to turbulent wall boundary layer". In: *AIAA Journal* 2.1, pp. 84–93. DOI: 10.2514/3.2218.

Morkovin, M. V. (1957). "On transition experiments at moderate supersonic speeds". In: *Journal of the Aeronautical Sciences* 24.7, pp. 480–486. DOI: 10.2514/8.3887.

– (1959). "On supersonic wind tunnels with low free-stream disturbances". In: *Journal of Applied Mechanics* 26.3, pp. 319–324.

Nocera, F. (2004). "LIGO laser intensity noise suppression". In: *Classical and Quantum Gravity* 21.5, S481–S485. DOI: 10.1088/0264-9381/21/5/014.

Norfleet, G. D., Lacey Jr, J. J., and Whitfield, J. D. (1966). *Results of an experimental ivestigation of the performance of an expansion tube*. Tech. rep. Arnold Engineering Development Center, Arnold AFB TN.

Osgood, B. (2007). *The Fourier transform and its applications*. Lecture notes for EE261. Stanford University.

Parker, R., Wakeman, T., Holden, M., and MacLean, M. (2006). "Measuring nitric oxide freestream concentration using quantum cascade lasers at CUBRC". In: *44th AIAA Aerospace Sciences Meeting and Exhibit*. Reno, NV. DOI: 10.2514/6.2006-926.

Parziale, N. J., Rabinovitch, J., Blanquart, G., Hornung, H. G., and Shepherd, J. E. (2013a). "Proposed vertical expansion tunnel". In: *AIAA Journal* 51.12, pp. 2792–2799. DOI: 10.2514/1.J052389.

Parziale, N. J., Shepherd, J. E., and Hornung, H. G. (2012). “Reflected shock tunnel noise measurement by focused differential interferometry”. In: *42nd AIAA Fluid Dynamics Conference and Exhibit*. New Orleans, LA. DOI: 10.2514/6.2012-3261.

– (2013b). “Differential interferometric measurement of instability in a hypervelocity boundary layer”. In: *AIAA Journal* 51.3, pp. 750–754. DOI: 10.2514/1.J052013.

– (2014). “Free-stream density perturbations in a reflected-shock tunnel”. In: *Experiments in Fluids* 55.2. DOI: 10.1007/s00348-014-1665-0.

– (2015). “Observations of hypervelocity boundary-layer instability”. In: *Journal of Fluid Mechanics* 781, pp. 87–112. DOI: 10.1017/jfm.2015.489.

Parziale, N. J. (2013). “Slender-body hypervelocity boundary-layer instability”. Ph.D. thesis. California Institute of Technology.

Pate, S. R. (1980). “Effects of wind tunnel disturbances on boundary-layer transition with emphasis on radiated noise: a review”. In: *11th Aerodynamic Testing Conference*. Colorado Springs, CO: American Institute of Aeronautics and Astronautics.

Paull, A. and Stalker, R. J. (1992). “Test flow disturbances in an expansion tube”. In: *Journal of Fluid Mechanics* 245, pp. 493–521. DOI: 10.1017/S0022112092000569.

Prasad, S. (2017). “Extended Taylor frozen-flow hypothesis and statistics of optical phase in aero-optics”. In: *Journal of the Optical Society of America A* 34.6, pp. 931–941. DOI: 10.1364/JOSAA.34.000931.

Rollins, J., Ottaway, D., Zucker, M., Weiss, R., and Abbott, R. (2004). “Solid-state laser intensity stabilization at the 10^{-8} level”. In: *Optics Letters* 29.16, pp. 1876–1878. DOI: 10.1364/OL.29.001876.

Rufer, S. and Berridge, D. (2012). “Pressure fluctuation measurements in the NASA Langley 20-inch Mach 6 wind tunnel”. In: *42nd AIAA Fluid Dynamics Conference and Exhibit*. New Orleans, LA. DOI: 10.2514/6.2012-3262.

Salyer, T., Collicott, S., and Schneider, S. P. (2000). “Feedback stabilized laser differential interferometry for supersonic blunt body receptivity experiments”. In: *AIAA 38th Aerospace Sciences Meeting and Exhibit*. Reno, NV. DOI: 10.2514/6.2000-416.

Schilden, T., Schröder, W., Ali, S. R. C., Schreyer, A.-M., Wu, J., and Radespiel, R. (2016). “Analysis of acoustic and entropy disturbances in a hypersonic wind tunnel”. In: *Physics of Fluids* 28.5, p. 056104. DOI: 10.1063/1.4948435.

Schmidt, B. (1976). “The shock wave curvature close to the shock tube wall”. In: *Archives of Mechanics* 28.5-6, pp. 809–815.

Schmidt, B. E. and Shepherd, J. E. (2015). “Analysis of focused laser differential interferometry”. In: *Applied Optics* 54.28, pp. 8459–8472. DOI: 10.1364/AO.54.008459.

Schneider, S. P. (2001). "Effects of high-speed tunnel noise on laminar-turbulent transition". In: *Journal of Spacecraft and Rockets* 38.3, pp. 323–333. DOI: 10.2514/2.3705.

– (2008). "Development of hypersonic quiet tunnels". In: *Journal of Spacecraft and Rockets* 45.4, pp. 641–664. DOI: 10.2514/1.34489.

Settles, G. S. and Fulghum, M. R. (2016). "The focusing laser differential interferometer, an instrument for localized turbulence measurements in refractive flows". In: *Journal of Fluids Engineering* 138.10, p. 101402. DOI: 10.1115/1.4033960.

Sharma, M. (2010). "Post-shock thermochemistry in hypervelocity CO_2 and air flow". Ph.D. thesis. University of Illinois at Urbana-Champaign.

Shinn, J. L. and Miller III, C. G. (1978). *Experimental perfect-gas study of expansion-tube flow characteristics*. NASA Technical Paper TP-1317. Langley Research Center, Hampton, VA.

Smeets, G. and George, A. (1973). *Laser-differential interferometer applications in gas dynamics*. Tech. rep. 28/73. French-German Research Institute of Saint-Louis.

Soref, R. A. and McMahon, D. H. (1966). "Optical design of wollaston-prism digital light deflectors". In: *Applied Optics* 5.3, p. 425. DOI: 10.1364/AO.5.000425.

Spurk, J. H. (1965). *Design, operation and preliminary results of the BRL expansion tube*. Tech. rep. AD0630065. Ballistic Research Labs, Aberdeen Proving Ground, MD.

Stainback, P. C. and Wagner, R. D. (1972). "A comparison of disturbance levels measured in hypersonic tunnels using a hot-wire anemometer and a pitot pressure probe". In: *7th Aerodynamic Testing Conference*. Palo Alto, CA: American Institute of Aeronautics and Astronautics.

The MathWorks, Inc. (2019). *MATLAB (R2019a)*. Version 9.6.0. Natick, MA.

Thompson, P. A., Strock, T. W., and Lim, D. S. (1983). "Estimate of shock thickness based on entropy production". In: *Physics of Fluids* 26.1, pp. 48–49. DOI: 10.1063/1.863982.

Trimpi, R. L. (1962). "A Preliminary Study of a New Device for Producing High-Enthalpy, Short-Duration Gas Flows". In: Denver, Colorado: NASA Langley Research Center.

Vaseghi, S. V. (2000). "Spectral subtraction". In: *Advanced digital signal processing and noise reduction*. 2nd ed. John Wiley & Sons Ltd, pp. 333–354.

Virtanen, P. et al. (2020). "SciPy 1.0: fundamental algorithms for scientific computing in Python". In: *Nature Methods* 17, pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

Wagner, A., Schülein, E., Petervari, R., Hannemann, K., Ali, S. R. C., Cerminara, A., and Sandham, N. D. (2018). "Combined free-stream disturbance measurements and receptivity studies in hypersonic wind tunnels by means of a slender wedge probe and DNS". In: *Journal of Fluid Mechanics* 842, pp. 495–531. DOI: 10.1017/jfm.2018.132.

Wagner, R. D. (1971). "Hot wire measurements of freestream and shock layer disturbances". In: *AIAA Journal* 9.12, pp. 2468–2470. DOI: 10.2514/3.6537.

Welch, P. D. (1967). "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms". In: *IEEE Transactions on Audio and Electroacoustics* 15.2, pp. 70–73. DOI: 10.1109/TAU.1967.1161901.

Wheaton, B. M., Berridge, D. C., Wolf, T. D., Araya, D. B., Stevens, R. T., McGrath, B. E., Kemp, B. L., and Adamczak, D. W. (2021). "Final design of the boundary layer transition (BOLT) flight experiment". In: *Journal of Spacecraft and Rockets* 58.1, pp. 6–17. DOI: 10.2514/1.A34809.

Wolfram Research, Inc. (2019). *Mathematica*. Version 12.0. Champaign, IL.

Yanes, N. J. (2020). "Ultraviolet radiation of hypervelocity stagnation flows and shock/boundary-layer interactions". Ph.D. thesis. California Institute of Technology.

Appendix A

FLDI PLANE WAVE RESPONSE DERIVATION

This appendix gives more detail for how Eq. (5.21) was obtained from the evaluation of the ray-tracing equation Eq. (5.20) when the refractive index field is defined as a propagating plane wave per Eq. (5.17).

Expanding the dot product in Eq. (5.17):

$$n(\underline{x}, t) = A \cos\left(k_x x + k_y y + k_z z - \omega t\right) \tag{A.1}$$

Although $\{k_x, k_y, k_z, \omega\}$ are not independent variables, being related to each other by Eqs. (5.18) and (5.19), they are independent of the integration variables $(\tilde{r}, \theta, z)$ and so we can treat them as independent constants during the integrations. Note also that ωt and φ are independent of everything else.

Denoting n_1 as n_+ and n_2 as n_-, apply the co-ordinate transforms of Eq. (4.6) to convert $n_\pm(x, y, z)$ to $n_\pm = n_i(\tilde{r}, \theta, z)$:

$$n_\pm(\underline{x}, t) = A \cos\left[k_x\left(\tilde{r}w\cos\theta \pm \Delta x/2\right) + k_y \tilde{r} w \sin\theta + k_z z - \omega t\right] \tag{A.2}$$

Perform the integrations in the order indicated by Eq. (5.20), i.e. integrate in θ, then $\tilde{r}$, then subtract, and finally integrate in z. Showing the individual steps, first:

$$\begin{aligned}\Theta_\pm(\tilde{r}, z, t) &\equiv \frac{2\pi A}{\lambda_L} \cdot \int_0^{2\pi} n_\pm(\tilde{r}, \theta, z, t)\mathrm{d}\theta \\ &= \frac{2\pi A}{\lambda_L} \cdot 2\pi J_0\left(\sqrt{k_x^2 + k_y^2}\tilde{r}w\right)\cos\left(\pm\frac{k_x \Delta x}{2} + k_z z - \omega t\right)\end{aligned} \tag{A.3}$$

where J_0 is the Bessel function of the first kind and of order 0. Then:

$$\begin{aligned}R_\pm(z, t) &\equiv \frac{2\pi A}{\lambda_L} \cdot \int_0^{\infty} \tilde{r}\exp\left(-2\tilde{r}^2\right)\Theta_\pm(\tilde{r}, z, t)\mathrm{d}\tilde{r} \\ &= \frac{2\pi A}{\lambda_L} \cdot \exp\left(-\frac{1}{8}\left[k_x^2 + k_y^2\right]w^2(z)\right)\cos\left(\pm\frac{k_x \Delta x}{2} + k_z z - \omega t\right)\end{aligned} \tag{A.4}$$

Lastly:

$$
\begin{aligned}
\Delta\Phi(t) &= \frac{2\pi A}{\lambda_L} \cdot \int_{-L}^{L} \left[R_+(z,t) - R_-(z,t)\right] \mathrm{d}z \\
&= \frac{2\pi A}{\lambda_L} \cdot \frac{2\sqrt{2}\pi^{3/2} w_0}{\sqrt{k_x^2 + k_y^2}\lambda_L} \cdot \exp\left(-\frac{w_0^2}{8}\left[k_x^2 + k_y^2 + \frac{16\pi^2 k_z^2}{\left(k_x^2 + k_y^2\right)\lambda_L^2}\right]\right) \\
&\cdot \mathrm{i}\left\{\mathrm{erfi}(\bar{\xi}) - \mathrm{erfi}(\xi)\right\} \cdot \sin\left(\frac{k_x \Delta x}{2}\right) \cdot \sin(\omega t)
\end{aligned}
\tag{A.5}
$$

where the overbar denotes conjugacy, and:

$$
\xi \equiv \frac{4\pi^2 k_z w_0^2 + \mathrm{i}\left(k_x^2 + k_y^2\right) L\lambda_L^2}{2\sqrt{2}\pi\sqrt{k_x^2 + k_y^2}\lambda_L w_0} = B + \mathrm{i}C \tag{A.6}
$$

$\mathrm{erfi}(z) \equiv -\mathrm{i}\,\mathrm{erf}(\mathrm{i}z)$, and like $\mathrm{erf}(z)$, it is an odd function. Using this, along with the the conjugation property of erf:

$$
\begin{aligned}
\mathrm{erf}(\bar{z}) &= \overline{\mathrm{erf}(z)} \\
\Rightarrow \mathrm{erf}(z) + \mathrm{erf}(\bar{z}) &= \mathrm{erf}(z) + \overline{\mathrm{erf}(z)} \\
&= \Re\left[\mathrm{erf}(z)\right] \cancel{+\Im\left[\mathrm{erf}(z)\right]} + \Re\left[\mathrm{erf}(z)\right] \cancel{-\Im\left[\mathrm{erf}(z)\right]} \\
&= 2\Re\left[\mathrm{erf}(z)\right]
\end{aligned}
\tag{A.7}
$$

(where $\Re$ and $\Im$ designate real and imaginary parts, respectively) we find:

$$
\mathrm{i}\left\{\mathrm{erfi}(\bar{\xi}) - \mathrm{erfi}(\xi)\right\} = 2\Re\left[\mathrm{erf}(\Xi)\right] \tag{A.8}
$$

where:

$$
\Xi \equiv C + \mathrm{i}B \tag{A.9}
$$

So the final result Eq. (5.21) is recovered:

$$\Delta\Phi(t) = \frac{2\pi A}{\lambda_L} \cdot \frac{4\sqrt{2}\pi^{3/2} w_0}{\sqrt{k_x^2 + k_y^2}\lambda_L} \cdot \sin\left(\frac{k_x \Delta x}{2}\right) \cdot \sin\left(\omega t - \varphi\right)$$
$$\cdot \exp\left(-\frac{w_0^2}{8}\left[k_x^2 + k_y^2 + \frac{16\pi^2 k_z^2}{\left(k_x^2 + k_y^2\right)\lambda_L^2}\right]\right)$$
$$\cdot \Re\left[\operatorname{erf}\left(\frac{\left(k_x^2 + k_y^2\right) L\lambda_L^2 + \mathrm{i} \cdot 4\pi^2 k_z w_0^2}{2\sqrt{2}\pi\sqrt{k_x^2 + k_y^2}\lambda_L w_0}\right)\right] \tag{A.10}$$

Appendix B

SPECTRAL CONVERSION CONSIDERATIONS

For the work in Chapter 6, spectra of various quantities (represented using PSD) need to be interconverted several times. When applying the different transfer functions, care needs to be taken that the quantities are in the right form and scaled correctly. This appendix gives more information on these minor details, which are nevertheless important for getting quantitatively-correct outputs.

Consider a general time-varying quantity q, decomposed into mean and fluctuating parts:

$$q = \bar{q} + q' \tag{B.1}$$

where by definition $\overline{q'} = 0$. The first point to note is that $PSD\{q\} = PSD\{q'\}$, except perhaps at $f = 0$, depending on the particular implementation of the algorithm. The second point is that $PSD\{q'\} \propto (q')^2$ in a frequency-wise sense. Due to the latter fact, the relationship between absolute and relative spectra is:

$$PSD\{q'/\bar{q}\} = \frac{1}{\bar{q}^2} \times PSD\{q'\} \tag{B.2}$$

If q has dimensions Q, then $[PSD\{q'\}] = \mathrm{Q}^2\mathrm{Hz}^{-1}$ while $[PSD\{q'/\bar{q}\}] = \mathrm{Hz}^{-1}$. The particular quantities that appear in Chapter 6 are the refractive index n, density ρ, and pressure p. These are related by the Gladstone-Dale and isentropic relations:

$$n = K\rho + 1 \tag{B.3a}$$

$$\frac{p'}{\overline{p}} = \gamma\frac{\rho'}{\overline{\rho}} \tag{B.3b}$$

Substituting $\rho = \bar{\rho} + \rho'$ into Eq. (B.3a) and grouping gives:

$$\begin{aligned} n &= K(\bar{\rho} + \rho') + 1 \\ &= (K\bar{\rho} + 1) + (K\rho') \\ &= \bar{n} + n' \end{aligned} \tag{B.4}$$

So then:

$$PSD\{n'\} = \frac{1}{K^2} \times PSD\{\rho'\} \tag{B.5}$$

i.e. the +1 term is not used for the conversion of the fluctuating part; this can easily be verified by synthesizing a $\rho(t)$ signal, converted it to $n(t)$, then comparing the spectra using Welch's method.

Because $\bar{\rho}$, $\bar{p}$, and γ are all constants, Eq. (B.3b) simply yields:

$$PSD\{p'\} = \left(\frac{\gamma\bar{p}}{\bar{\rho}}\right)^2 \times PSD\{\rho'\} \tag{B.6}$$

Appendix C

FLDI BEAM CLIPPING

One important application of FLDI is making measurements of disturbances propagating in supersonic or hypersonic boundary layers. This has been performed on both conical models (Benitez et al., 2020; Parziale et al., 2015) and flat walls (Ceruzzi et al., 2020). Probing at varying heights within a boundary layer can necessitate the FLDI foci approaching closely to the solid surface. Due to the convex surface, it is possible to get very small stand-off distances on a cone without cut-off, but flat geometries will often intersect the beam away from the foci, where the beam diameter is larger. Alternative beam geometries have been proposed to avoid this issue (Houpt and Leonov, 2018, 2019), but here the effect of clipping portions of the standard circular FLDI beam configuration is explored.

To implement clipping in the code, consider the integrand of Eq. (4.1), which can be alternatively expressed as:

$$I_0(\xi,\eta)\sin\left[\frac{2\pi}{\lambda_L}\cdot \mathrm{OPD}(\xi,\eta)\right] = I_0(\xi,\eta)\sin\left[\Delta\phi(\xi,\eta)\right] \tag{C.1}$$

where $\Delta\phi$ is the phase difference between the recombining ray pair at (ξ,η). This is not a well-defined quantity when the rays do not reach the detector, as a non-existent ray does not have a phase. However, the integrand is weighted by the intensity distribution. Regardless of phase, if neither ray in a pair reaches the detector, there can physically be no contribution to the total intensity and thus $I_0 = 0$ at that point. Rather than modifying I_0, it is more efficient in this implementation* to set $\Delta\phi = 0$ at a point, which has the same numerical outcome. Wherever the beam intersects a solid object, the values of $N_\pm$ are set to `NaN`. This causes the ray integrals $J_\pm$ to also evaluate to `NaN` even if only a single point along the ray is blocked. A masking operation then replaces `NaN` with 0 so that $\sin[\Delta\phi(\xi,\eta)] = 0$ as desired. The edge case where only one ray in a pair is clipped is not addressed, because rays are only separated by $\Delta x \sim \mathcal{O}(100\,\mu\mathrm{m})$, so in practice it is highly likely that either both or neither of the rays in a pair will be clipped.

*See discussion in Implementation subsection of Section 4.1

As a test case, a propagating planar wavepacket of the form

$$n'(x) = A \exp\left[\frac{(x - x_0)^2}{l^2}\right] \cos\left[\frac{2\pi}{\lambda}(x - x_0)\right] \tag{C.2}$$

is simulated with various types of beam clipping. $n'(x) = n(x) - n_\infty$ where n_∞ is the undisturbed freestream value. A is the amplitude of the envelope, l is the $1/e$ half-width, x_0 is the offset of the envelope peak, and λ is the wavelength. For simplicity, the wavepacket is modeled as non-dispersive, i.e. the waveform and envelope propagate together at a fixed velocity. The wave has a top-hat distribution in z, extending to $z = \pm 5$ cm about the focal plane. The magnitudes of the constants in Eq. (C.2) are chosen to be similar to those used in Schmidt and Shepherd (2015), which itself models a Mack-mode wavepacket in a hypersonic boundary layer. This gives $A = 3 \times 10^{-7}$, $k = 2\pi/\lambda = 2\,\mathrm{mm}^{-1}$, $l = 20$ mm. The wave propagates in x with velocity $u = 3.5\,\mathrm{km\,s}^{-1}$. This is achieved by updating x_0 at each timestep. The step size $\Delta t = 4.5$ ns is chosen to give a Nyquist frequency ratio of 100.

A flat plate of infinite x-extent is placed at two different heights below the foci: $y_p = 1$ mm and 2 mm (Fig. C.1). This corresponds to the beam first intersecting the plate at $z \approx 3$ cm and 6 cm, respectively. Three different clipping scenarios are used: symmetric ($-\infty < z_p < +\infty$), pitch-side asymmetric ($-\infty < z_p < 0$), and catch-side asymmetric ($0 < z_p < +\infty$).

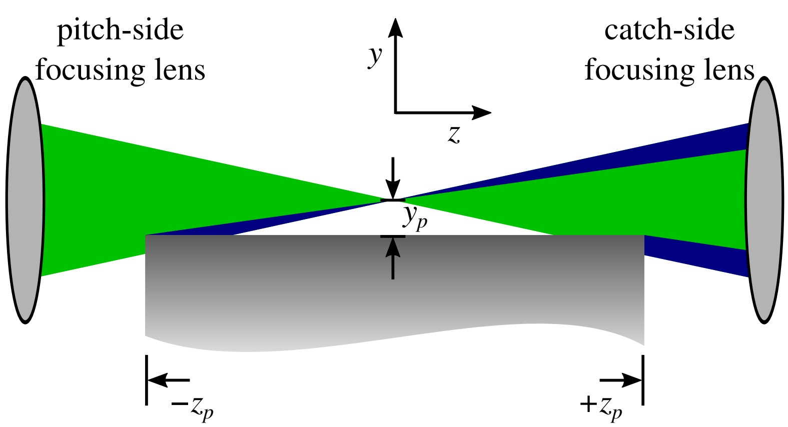

Figure C.1: Schematic of beam clipping by symmetrically-located flat plate. Occluded portions of beams shown in blue.

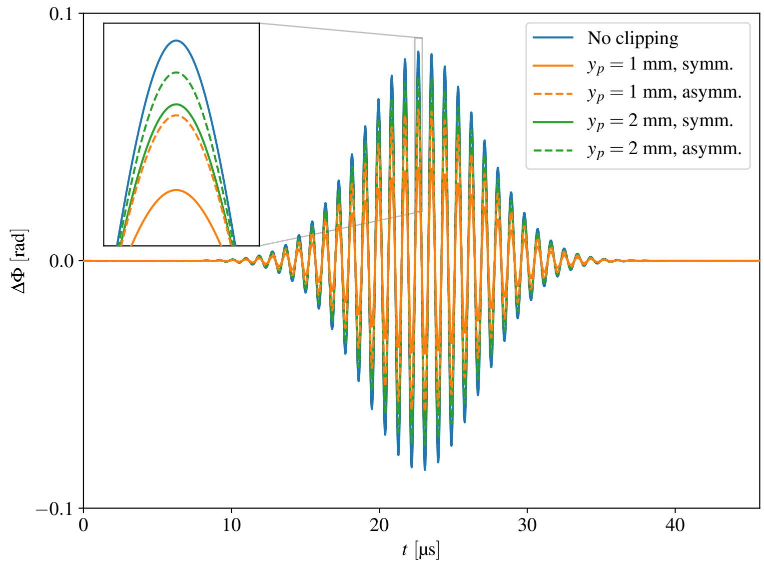

Figure C.2: FLDI response to propagating wavepacket with varying amounts of beam clipping. Detail of central peak shown in inset.

The results (Fig. C.2) show that clipping causes a reduction in signal magnitude only. There is no alteration of frequency or phase information. Increased amounts of clipping give increased reductions in signal; one-sided (asymmetric) clipping gives an identical response regardless of which side the clipping occurs on, and two-sided (symmetric) clipping causes precisely twice the signal reduction as the corresponding one-sided clipping, i.e. the reductions are additive.

This outcome is perhaps obvious given that the test flow is uniform in both y and z. In most use-cases of FLDI in wind tunnels, the flow is indeed likely to be uniform in z, at least in a mean sense and in the core flow where the FLDI is most sensitive. However, boundary layers have varying properties in the y-direction, and some clipping geometries could skew the contributions to the integrated FLDI response. The instrument sensitivity varies in the y direction too, due to the $I_0(\xi,\eta)$ term in Eq. (4.1): as long the clipping is not severe enough to impinge on the central peak of the Gaussian intensity profile, the overall response should not be altered substantially. These results indicate that for many applications, beam clipping should not qualitatively alter the signal, although consideration should be given to intensity attenuation and possible signal-to-noise ratio issues.

Appendix D

FLDI NOISE REDUCTION

One of the leading applications for FLDI is the measurement of complex fluctuating density fields, such as in turbulence. These signals are stochastic and broad-band, and are often quite weak in ground testing facilities where densities can be low as a by-product of reaching hypervelocity conditions. Therefore, it is important to understand the possible sources of experimental noise, and mitigate their effects if possible.

Some types of FLDI experiments can simply rely on averaging out noise without having to address its origin. For example, two of the validation cases that are discussed in Section 4.2 (the laminar jet and the ultrasound beam) were capable of steady continuous operation, allowing for long record lengths. Even though the ultrasound beam was a dynamic field, it had a dominant single-frequency component that was known *a priori* — allowing noise to be discarded by spectral filtering. Some hypersonic facilities can operate in a quasi-continuous manner (e.g. blowdown tunnels that might have test times on the order of minutes); these allow for long record lengths compared to the turbulent timescales of interest, and thereby statistical noise reduction procedures, as discussed later. More challenging conditions are encountered in impulse facilities, where test times are very short, and densities low. Additionally, the impulsive mechanism is often quite violent, leading to strong structural motion that can influence the FLDI signal.

D.1 Previous Work

The first uses of FLDI on impulse facilities were by Parziale et al. (2012, 2013b, 2014, 2015), all on T5. An upper bound on spectral noise was found in their frequency band of interest by using the portion of the signal preceding the start of the flow. Vibrations from the driver piston arrive ahead of the flow, and it was assumed these do not become significantly stronger during test time. When searching for acoustic instability wavepackets, a two-tailed hypothesis test was used to show statistically-significant differences between these intermittent wavepackets and the "background" test-time flow. When measuring freestream flow noise spectra, no attempt was made to address non-flow noise sources, but this was unimportant because the primary goal was to look for spectral trends versus reservoir enthalpy.

It was noted that discrete spikes of radio-frequency noise were observed even when the tunnel was not in operation.

In their treatment of noise concerns, Settles and Fulghum (2016) appear to assume that the main potential noise source is electronic, i.e. in the detection and acquisition instruments. To mitigate this, they use an alternative, more complex single FLDI setup that uses a Berek compensator, a polarizing beamsplitter, and separate photodetectors for each of the two polarized beams — as opposed to the “canonical” configuration used in this work, where the beams are recombined and interfered on a single photodetector. Their arrangement allows for rejection of uncorrelated noise between the two channels, because actual turbulent signals will be correlated. Other suggested measures were to use a more powerful laser source in combination with electronic shielding to improve the SNR.

D.2 Common-Mode Noise Reduction

The 200 mW diode laser [Spectra-Physics EXLSR-532-200-CDRH] used on all FLDI systems throughout this work is far more powerful than the 0.8 mW HeNe laser used by Fulghum (2014) and Settles and Fulghum (2016), so electronic noise was not expected to be a significant issue. However, benchtop testing and preliminary experiments on HET revealed that the no-flow FLDI signal was significantly noisy, often with fluctuations of the same order of magnitude as flow features. Initially, various photodetectors, cables, shielding methods, and oscilloscopes were trialed, yet the noise spectrum remained similar between tests. This implied that the laser itself (the only constant component) was the main source of the noise.

Stronger evidence was provided when the single FLDI was modified to a double FLDI. When similar no-flow measurements were taken with the DFLDI, the noise signals were very similar between the two detectors, not just in spectral space but even in the temporal signals. This would not be the case if the noise were due to, for example, dark current noise in the independent photodetectors. A beamsplitter was added to the DFLDI immediately after the laser aperture to divert 10% of the light to a third photodetector, before it passed through the rest of the optical train. The noise signal here was also similar overall to the pair of photodetectors, although with some discrepancies. This may be due to complicated interference effects between all the back-reflected light off each optic in the system, or from airborne dust particles that affect the main beam path but not the diverted beam.

The nature of this noise was not uniform. Sometimes the fluctuation levels were low,

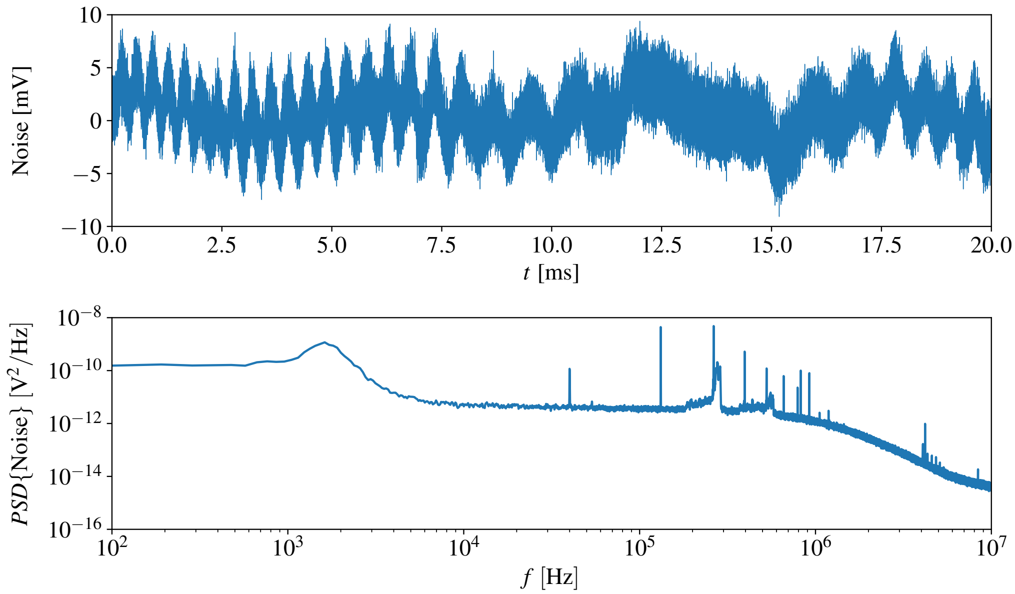

Figure D.1: Typical laser noise signal shown in the temporal (top) and frequency (bottom) domains.

followed by a higher-amplitude wave-train of finite duration and distinct frequency, before settling back to the baseline again. This was postulated to be attributable to some form of cyclic laser instability. The spectra of these noise signals consistently showed a peak around 1700 Hz, and the aforementioned high-amplitude waves had periods corresponding with this frequency, superimposed with much higher-frequency components (Fig. D.1).

These noise fluctuations had typical amplitudes of a few mV, whereas the full range of the FLDI over $0 < \Delta\Phi < \pi$ rad corresponds to a voltage range of about $0 < V < 2.5$ V, and the maximum output of the photodetectors is 5 V when terminated at 50 Ω. The laser manufacturer specifications claim a noise of < 0.2% RMS over 20 Hz–20 MHz, and in fact is advertised as having "exceptionally low optical noise" with interferometry as an application. This RMS value can be computed by integrating the power spectral density of the relative intensity over the given bandwidth. 0.13% RMS noise was obtained, within the specification. So in fact the laser noise is indeed quite low in absolute terms, and furthermore largely contained to a particular bandwidth. However, at low-density conditions, especially in HET, this can still be of a similar magnitude to actual flow signals.

The optical configuration of the DFLDI constructed in this work is such that the polarization directions of the two beam pairs are mirrored. Thus, a given flow feature

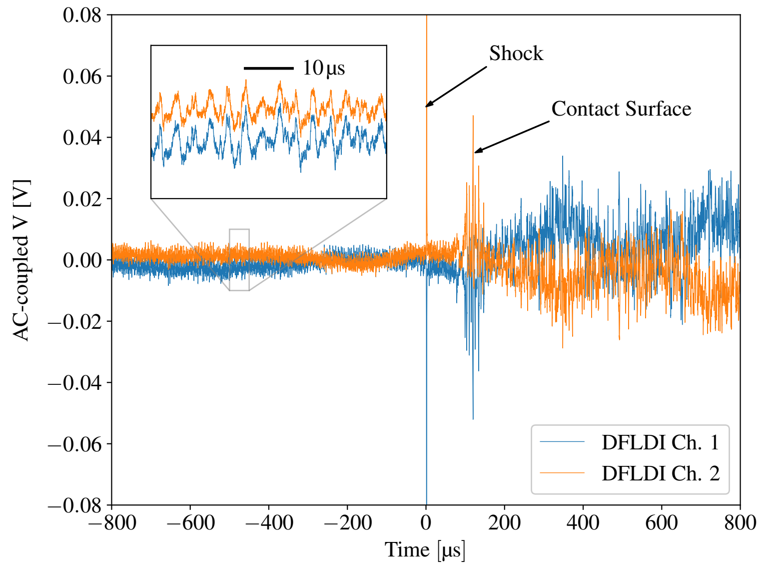

Figure D.2: An example DFLDI signal pair, spanning the starting process and test time of HET. The laser intensity noise during no-flow conditions shows strong correlation, with fluctuations having the same sign (see inset detail). Conversely, flow features give signals of equal magnitude but opposite sign.

actually causes a response of $+\Delta\Phi$ at one detector, and $-\Delta\Phi$ at the other, while laser intensity noise causes the voltage at each detector to move in the same direction. In this way, when the two signals are overlaid, fluctuations due to actual flow features are qualitatively distinguishable from noise fluctuations, as shown in Fig. D.2. This observation offers a starting point for temporally-based noise reduction algorithms, via common-mode rejection.

Consider how laser intensity fluctuations affect the apparent FLDI response. Beginning with a noiseless system, the voltage is given by:

$$V(t) = \frac{(2\epsilon - 1)kI_0}{2} \sin\left[\Delta\Phi(t) + \Delta\Phi_0\right] + \frac{kI_0}{2} = A \sin\left[\Delta\Phi(t) + \Delta\Phi_0\right] + B \quad \text{(D.1)}$$

I_0 is the full intensity of the beam, k is the gain of the detection system, and $0 < \epsilon \leq 1$ is the contrast ratio. $\Delta\Phi_0$ is the initial phase offset of the interferometer, and is experimentally determined along with the constants A and B during the calibration process, i.e. k and ϵ do not need to be determined directly. Note this is just a re-expression of Eq. (3.17).

Next, intensity noise is introduced, such that the mean intensity is still I_0:

$$I(t) = \alpha(t) I_0 \tag{D.2}$$

where:

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T \alpha(t)dt = 1 \tag{D.3}$$

Replacing I_0 with $I(t)$ in Eq. (D.1) yields:

$$V^*(t) = \alpha(t)V(t) \tag{D.4}$$

From here on, the superscript asterisk will denote noise-corrupted signals, and the unmodified symbol will represent the "true" signal. Eq. (D.4) implies that the noise is multiplicative, and that the amplitude of the fluctuations will scale with the underlying mean voltage. This can be easily tested experimentally. In order to get sufficient voltage resolution during low-density experiments, the FLDI photodetector signal is usually recorded in AC-coupled mode; splitting the signal into an additional DC-coupled channel allows the mean voltage V_0 to be recorded alongside the finely-resolved noise fluctuations. $V_0 = A \sin \Delta\Phi_0 + B$ can be controlled by changing $\Delta\Phi_0$; this is done by slightly translating the final Wollaston prism/polarizer combination in the same way as during the regular calibration process detailed in Section 3.4. Fig. D.3 shows the AC- and DC-coupled voltages for a range of $\Delta\Phi_0$ (the transient spikes occur while the optics are being moved). $V'_{RMS} = \mathrm{RMS}(V_{AC})$ and $V_0 = \mathrm{mean}(V_{DC})$ are then extracted and plotted, giving a linear relationship as predicted.

Ultimately, the goal is to extract a denoised $\Delta\Phi(t)$ from the noisy $V^*(t)$. As an intermediate step towards a solution, $V(t)$ is considered as the composite function $g(t)$:

$$V(t) = f\left[\Delta\Phi(t)\right] \equiv g(t) \tag{D.5}$$

where the function f is given by Eq. (D.1). If the assumption is made that both DFLDI stations have the same response function, with the downstream signal lagging that of the upstream signal by a time delay τ, then the two recorded noisy signals are given by:

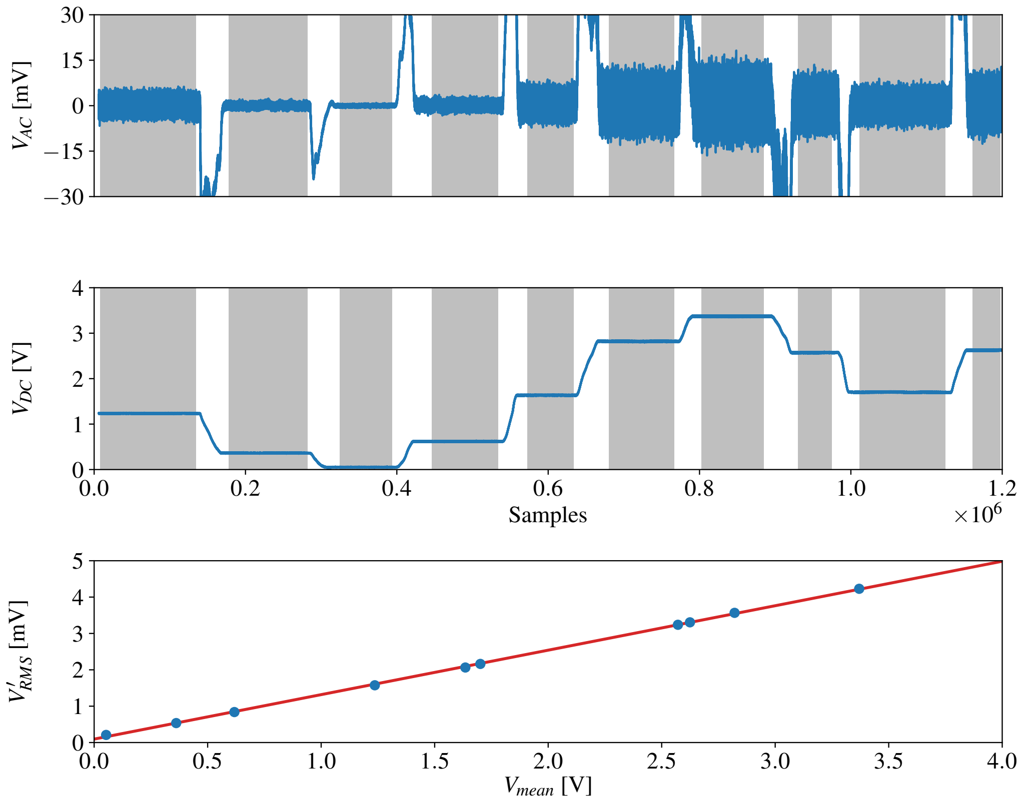

Figure D.3: Experimental data showing that laser intensity noise has a multiplicative nature. Grey bands are portions of the data used to compute the values at each point in the bottom plot.

$$V_1^*(t) = \alpha(t)g(t) \tag{D.6a}$$

$$V_2^*(t) = \alpha(t)g(t-\tau) \tag{D.6b}$$

In the more general case, each station will have a different response function due to variations between photodetectors, in which case:

$$V_1^*(t) = \alpha(t)f_1\left[\Delta\Phi(t)\right] \tag{D.7a}$$

$$V_2^*(t) = \alpha(t)f_2\left[\Delta\Phi(t-\tau)\right] \tag{D.7b}$$

Note that in either case, there is the additional assumption that τ is constant, i.e. that the flow velocity is constant during the sample period of interest. First, a solution is sought for the more idealized model given by Eq. (D.6). The problem statement is: given measured signals $V_1^*(t)$ and $V_2^*(t)$ corrupted by a common multiplicative noise $\alpha(t)$, can the underlying function $g(t)$ be recovered? This is an interesting problem because both the noise and signal are common to both measurements, but the noise propagates through the symmetrical optical system at the speed of light, giving essentially no time delay, whereas the relatively slower flow signal does have a delay.

The common noise can immediately be eliminated by taking the ratio of the measured voltages:

$$r(t) \equiv \frac{V_1^*(t)}{V_2^*(t)} = \frac{g(t)}{g(t-\tau)} \tag{D.8}$$

To make subsequent analysis simpler, Eq. (D.8) is transformed by taking logs. This is valid because the raw voltage signals are both always positive.

$$s(t) = m(t) - m(t-\tau) \qquad \text{where} \begin{cases} s \equiv \log r \\ m \equiv \log g \end{cases} \tag{D.9}$$

The problem is now: with known function $s(t)$ and constant τ, can a unique function $m(t)$ that satisfies Eq. (D.9) be determined? The first attempted approach to tackling this problem was to apply Fourier transforms. Using standard identities:

$$S(\omega) = (1 - \mathrm{e}^{-i\omega\tau})M(\omega) \qquad \text{where} \begin{cases} S(\omega) \equiv \mathcal{F}\{s(t)\} \\ M(\omega) \equiv \mathcal{F}\{m(t)\} \end{cases} \tag{D.10}$$

$$\longrightarrow \quad m(t) = \mathcal{F}^{-1}\{M(\omega)\} = \mathcal{F}^{-1}\left\{\frac{S(\omega)}{1 - \mathrm{e}^{-i\omega\tau}}\right\} \tag{D.11}$$

Problems were encountered actually trying to perform this inversion, however, because the denominator goes to zero periodically, giving rise to an infinite number of simple poles evenly spaced along the real axis. There are methods for inverting this class of Fourier transform that involve deforming the integration contour around each pole (Inverarity, 2003), but these generally require knowing $S(\omega)$ as a continuous analytical function so that the residuals can be computed at each pole. In this case, $s(t)$ originates from discretely-sampled experimental data, so $S(\omega)$ is only known at discrete values of ω, via some numerical method such as FFT. In theory, estimates for the residuals could be obtained by fitting a complex surface to $S(\omega)$ and then numerically computing contour integrals, but this is likely to risk significant errors.

A different approach was instead taken, by applying the theory of lag operators (Mikusheva, 2013). These arise in time-series analysis, where given a discrete time-series x, the lag operator L shifts x by one index:

$$Lx_i = x_{i-1} \tag{D.12}$$

Lag operators can be raised to integer powers of either sign, and lag polynomials can be formed via linear combinations, e.g.:

$$L^2 x_i = x_{i-2} \tag{D.13a}$$

$$a(L) = \sum_{i=0}^{N} a_i L^i \tag{D.13b}$$

Much of the machinery for polynomials of complex variables can be applied to these lag polynomials, including inversion. Since these data are discrete time-series, Eq. (D.9) can be recast as a lag polynomial:

$$s_i = (1 - L^j) m_i \tag{D.14a}$$

$$\rightarrow \quad m_i = (1 - L^j)^{-1} s_i \tag{D.14b}$$

where j is the discrete index number corresponding to the continuous time lag τ.

MATLAB can invert and apply lag operators using the `lagop` class in the Econometrics toolbox (The MathWorks, Inc., 2019). However, there are restrictions on which lag polynomials can be inverted: the inverse function is an infinite series, and for a numerically truncated inverse to be stable, the coefficients must decrease with higher terms. This gives the requirement that the roots of the lag polynomial must lie outside the unit circle in the complex plane. Here, the lag polynomial $(1 - L^j)$ has all its roots on the unit circle, and hence its inverse does not converge. This is possibly another manifestation of the problem that prevented the solution of Eq. (D.11), as the term $e^{-i\omega\tau}$ also describes the unit circle.

This observation motivates a pertubational approach: if the roots can be moved slightly off the unit circle, both the Fourier transform and lag polynomial methods become tractable. Modifying Eq. (D.8) with an exponent κ perturbed from unity:

$$r_\kappa(t) \equiv \frac{V_1^*(t)}{\left[V_2^*(t)\right]^\kappa} = \frac{\alpha(t)g(t)}{\alpha^\kappa(t)g^\kappa(t-\tau)} \approx \frac{h(t)}{h^\kappa(t-\tau)} \tag{D.15a}$$

$$\rightarrow \quad s_\kappa(t) \approx m(t) - \kappa m(t-\tau) \tag{D.15b}$$

For $\kappa \neq 1$, the noise term $\alpha(t)$ no longer perfectly cancels out, introducing some error. However, $|\alpha^{1-\kappa}| << |\alpha|$ as $\kappa \rightarrow 1$, leading to strong noise reduction, although not total noise elimination. This yields the two approximate solution methods (Fourier transforms and lag polynomials, respectively):

$$\left.\begin{aligned} m(t) &\approx \mathcal{F}^{-1}\left\{\frac{S(\omega)}{1-\kappa \mathrm{e}^{-i\omega\tau}}\right\} \\ m_i &\approx (1-\kappa L^j)^{-1} s_i \end{aligned}\right\} \kappa = 1-\epsilon \,,\ |\epsilon| << 1 \tag{D.16}$$

Both methods were applied in MATLAB, trialing various values of κ. A synthetic Gaussian pulse signal was generated for $g(t)$ and corrupted by $\alpha(t)$ drawn from a uniform distribution. Fig. D.4 shows the results: both techniques reduce the amplitude of the noise, but the lag operator method performs much better, with the denoised signal indistinguishable from the original underlying signal at the scale shown here, for $\kappa = 0.999$. The Fourier transform method appears to introduce some spurious frequencies into the signal that were not originally present.

The perturbed lag operator method works well for this idealized scenario. Next, non-common-mode noise $\beta_i(t)$ is introduced such that:

$$V_1^*(t) = \beta_1(t)\alpha(t)h(t) \tag{D.17a}$$

$$V_2^*(t) = \beta_2(t)\alpha(t)h(t-\tau) \tag{D.17b}$$

From observing experimental data, it is known that the majority of the noise is common-mode, so $\beta_i(t)$ should have smaller fluctuations. Both α and β_i are drawn from distributions with unity mean, so for synthetic signals $1-a \leq \alpha \leq 1+a$ and

$1 - b \leq \beta_i \leq 1 + b$ where $b < a$. The same denoising procedures as above were repeated, with $a/b = 20$, with results shown in Fig. D.5; the performance is poor. Interestingly, there seems to be some denoising for intermediate values of κ, but as $\kappa \to 1$, the lag operator signal becomes more noisy than the original corrupted signal, while the Fourier transform signal blows up. This undesirable behavior is probably due to the uncorrelated nature of the non-common noise, which does not (almost) cancel out during the division process.

Even disregarding the issue of non-common-mode noise, there are other impediments to applying these techniques to real data. The assumptions of Eq. (D.6) were used in the derivations, but in reality Eq. (D.7) is more applicable. This equation can be expressed more explicitly as:

$$V_1^*(t) = \alpha(t)\,(A_1 \sin[\Delta\Phi(t) + \Delta\Phi_0] + B_1) \tag{D.18a}$$

$$V_2^*(t) = \alpha(t)\,(A_2 \sin[\Delta\Phi(t-\tau) + \Delta\Phi_0] + B_2) \tag{D.18b}$$

with $A_1 \neq A_2$ and $B_1 \neq B_2$. The above mathematical rearrangements only work when these calibration constants are equal.

The most pressing issue is that τ is not necessarily constant. For a sample of the flow signal, a constant τ can be computed by cross-correlation. If the true signal is of significantly larger magnitude than the noise, the correlation will yield an accurate τ even when applied to the noisy signal (which is all we have). However, if the velocity is fluctuating during the sample, then this τ only represents an average time delay between the two stations, and this then leads to errors when trying to reconstruct the original signal, because information may be taken from the wrong relative positions in each of the noisy signals.

In summary, although the methods here show potential for the elimination of common-mode laser intensity noise, they are inhibited by some of the complicating factors of real signals. There may exist more advanced techniques that look for similar oscillations between each signal of the DFLDI and then perform some form of demodulation, but we are not currently aware of these if so.

D.3 Statistical Noise Reduction

Instead of performing noise reduction in the temporal domain as discussed above, the frequency domain can be used. The laser intensity noise appears to be consistently

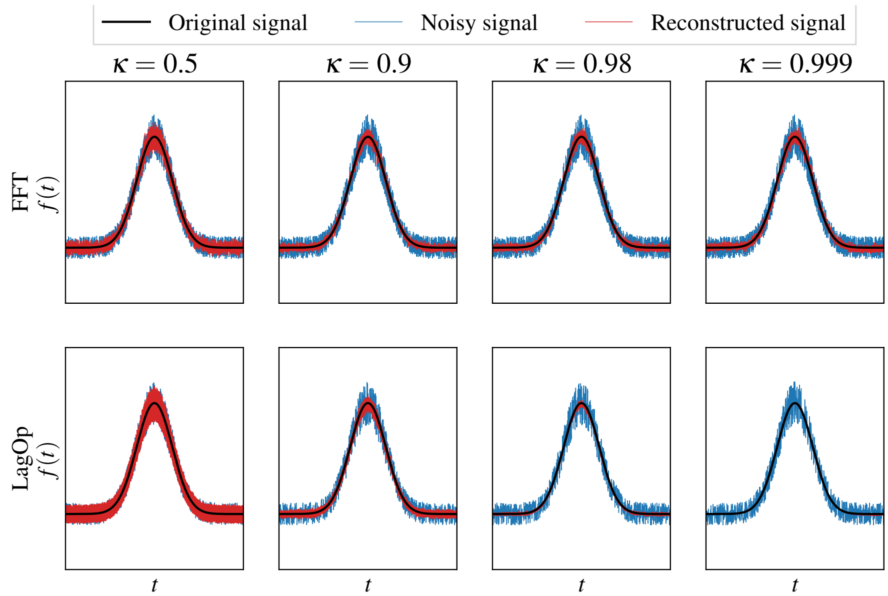

Figure D.4: Comparison of FFT (top) and lag operator (bottom) noise reduction algorithms for various values of κ. Only common-mode multiplicative noise is present.

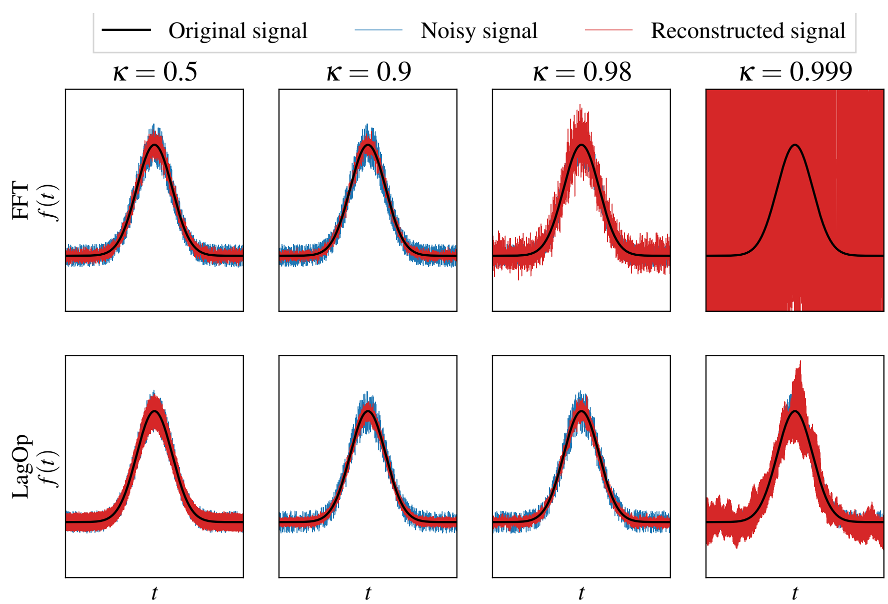

Figure D.5: The same comparison of noise removal algorithms as Fig. D.4, but with both common- and non-common-mode multiplicative noise present.

drawn from the same statistical distribution, which can be measured. Spectral subtraction of additive noise signals is a standard technique (Vaseghi, 2000), e.g. where the signal $x(t)$ is corrupted by noise $n(t)$:

$$y(t) = x(t) + n(t) \tag{D.19a}$$

$$\xrightarrow{\mathcal{F}} \quad Y(\omega) = X(\omega) + N(\omega) \tag{D.19b}$$

Estimates of the original signal can be obtained by subtracting an estimate of the noise spectrum:

$$|\tilde{X}(\omega)|^b = |Y(\omega)|^b - \overline{|N(\omega)|^b} \tag{D.20}$$

where the tilde designates an estimated quantity, and $b = 1$ or $= 2$ for a magnitude or power spectrum, respectively. Note that Eq. (D.20) assumes the signal and noise are uncorrelated, i.e. that the cross-spectral terms drop out. It also holds that:

$$\mathbb{E}\left[|\tilde{X}(\omega)|^2\right] = \mathbb{E}\left[|X(\omega)|^2\right] \tag{D.21}$$

so power spectral densities (PSDs) can be directly subtracted; PSDs being the form usually used when processing data in the frequency domain, e.g. via Welch's method. Note that the above technique applies when the signal of interest is stationary in time. This is, in theory, the case when we are interested in the spectral properties of a steady facility test time. Another factor that can affect this process is that the magnitude of the average noise spectrum may exceed the instantaneous noisy signal at some frequencies where the signal-to-noise ratio is low. In these cases, some form of rectification has to be applied in order to avoid nonsensical negative spectral values.

To apply this to the multiplicative noise of an FLDI signal, logs are taken:

$$PSD\left\{\log[V^*(t)]\right\} = PSD\left\{\log[V(t)]\right\} + PSD\left\{\log[\alpha(t)]\right\} \tag{D.22}$$

The quantity of interest is a spectral estimate of $\Delta\Phi$, which is related to $V(t)$ as shown in Eq. (D.1). If there exists a linear time-invariant transfer function between two signals, then it holds that:

$$y(t) = h(t) * x(t) \tag{D.23a}$$

$$Y(\omega) = H(\omega)X(\omega) \tag{D.23b}$$

$$PSD\{y(t)\} = |H(\omega)|^2 PSD\{x(t)\} \tag{D.23c}$$

However, because Eq. (D.1) is non-linear time-invariant, it is linearized by performing a Taylor expansion about $\Delta\Phi = \Delta\Phi_0$:

$$\begin{aligned}
\log[V(t)] &= \log\left[A\sin\left(\Delta\Phi(t) + \Delta\Phi_0\right) + B\right] \\
&\approx \log[A\sin(\Delta\Phi_0) + B] + \frac{A\cos(\Delta\Phi_0)}{A\sin(\Delta\Phi_0) + B}\Delta\Phi(t) + O\left(\Delta\Phi^2(t)\right) \\
&\approx a_0 + a_1\Delta\Phi(t)
\end{aligned} \tag{D.24}$$

When $|\Delta\Phi(t) - \Delta\Phi_0| \leq 0.1$ rad the error in this linearization is less than 5%. These magnitudes of $\Delta\Phi(t)$ are typical for low-density facilities. Eq. (D.24) then yields:

$$PSD\{V(t)\} \approx |a_0|^2 PSD\{\mathbb{1}\} + |a_1|^2 PSD\{\Delta\Phi(t)\} \tag{D.25}$$

The first term on the right-hand side ideally only contributes a delta function at $\omega = 0$; in practical numerical implementations there will be some spectral leakage, but this will only affect comparatively low frequencies—and the typical short test times of impulse facilities mean that frequencies below $1/\Delta t \sim 1$ kHz cannot be resolved.

Combining Eqs. (D.22) and (D.25) gives the final result:

$$PSD\left\{\widetilde{\Delta\Phi}(t)\right\} \approx \left|\frac{1}{a_1}\right|^2 \left(PSD\left\{\log[V^*(t)]\right\} - PSD\left\{\log[\alpha(t)]\right\}\right) \tag{D.26}$$

Recall that $V^*(t)$ is the raw measured voltage, $a_1 = f(A, B, \Delta\Phi_0)$ where the arguments are all known calibration constants, and $\alpha(t)$ is obtained from separate no-flow measurements. $\widetilde{\Delta\Phi}(t)$ refers to the estimate of the denoised FLDI response, while $\Delta\Phi^*(t)$ refers to the response computed without the application of any denoising process.

Note that the results of this section do not apply only to double FLDI, but equally to single FLDI because they do not rely on common-mode rejection in the temporal domain.

Appendix E

HET RUN CONDITIONS

Table E.1 lists all HET shots performed with the optical arms installed, as used for the results of Chapter 7. The tabulated diaphragm thicknesses refer to the primary aluminum diaphragm; all shots used the same 8.5 μm mylar secondary diaphragms. Note that p_5 is extrapolated using the measured leak rate from the final value recorded prior to firing.

The flow parameters are computed using LETS in real-gas mode; $Re_{m,7}$ relies on transport properties from Cantera.

Table E.1: List of HET run conditions used for FLDI freestream noise campaign.

Shot	Driver	Driven	Accel.	Dia. [mm]	p_4 [kPa]	p_1 [kPa]	p_5 [mTorr]	a_3/a_2	M_7	h_{07} [MJ kg^{-1}]	$Re_{m,7}$ [$\times 10^6$ m^{-1}]
1844	He	Air	Air	1.3	2310	1.28	630.8	0.561	3.899	4.729	1.193
1845	He	Air	Air	1.3	2448	1.86	483.9	0.604	4.450	4.844	1.410
1846	He	Air	Air	1.3	2332	2.61	380.6	0.662	4.977	4.787	1.624
1847	He	Air	Air	1.3	2333	3.55	313.9	0.712	5.473	4.739	1.860
1848	He	Air	Air	1.3	2248	4.72	260.2	0.770	5.972	4.613	2.098
1849	He	Air	Air	1.3	2304	6.15	219.1	0.817	6.479	4.553	2.349
1850	He	Air	Air	1.3	2466	7.86	187.5	0.853	6.987	4.543	2.610
1851	He	Air	Air	1.3	2473	9.90	163.3	0.903	7.471	4.414	2.889
1852	He	Air	Air	1.3	2447	1.86	480.9	0.604	4.456	4.849	1.407
1853	He	Air	Air	1.3	2416	3.55	314.2	0.706	5.478	4.789	1.861
1854	He	Air	Air	1.3	2343	6.15	220.6	0.814	6.473	4.574	2.355
1855	He	Air	Air	1.3	2658	9.90	164.9	0.887	7.474	4.524	2.898
1856	He	CO_2	Air	1.6	3687	1.38	521.0	0.769	5.382	4.946	1.838
1857	He	CO_2	Air	1.6	3803	1.80	459.7	0.811	5.722	4.960	1.987
1858	He	CO_2	Air	1.6	3436	2.90	377.2	0.929	6.298	4.718	2.294

Continued on next page. . .

Table E.1 continued from previous page

1859	He	CO_2	Air	1.6	3827	4.42	319.4	1.001	6.863	4.713	2.620
1860	He	CO_2	Air	1.6	3515	6.44	271.1	1.118	7.363	4.424	2.928
1861	He	CO_2	Air	1.6	3416	9.12	235.5	1.224	7.843	4.185	3.284
1862	He	Air	Air	0.8	1345	1.70	539.9	0.681	4.186	4.062	1.350
1863	He	Air	Air	0.8	1279	2.26	439.8	0.739	4.606	3.978	1.526
1864	He	Air	Air	0.8	1400	3.78	313.7	0.819	5.439	3.980	1.920
1866	He	Air	Air	1.6	3820	1.00	996.5	0.474	3.391	5.023	1.171
1867	He	Air	Air	1.6	3734	2.00	998.2	0.556	3.891	4.758	1.874
1868	He	Air	Air	1.6	3803	4.00	1001.5	0.652	4.438	4.456	2.971
1869	He	Air	Air	1.6	3748	8.00	999.9	0.773	5.028	4.026	4.672
1870	He	Ar	Air	1.6	3781	1.00	1004.5	0.339	2.845	5.108	0.616
1871	He	Ar	Air	1.6	3676	2.00	1000.7	0.412	3.448	4.708	1.153
1872	He	Ar	Air	1.6	3664	4.00	1002.4	0.499	4.131	4.252	2.115
1873	He	Ar	Air	1.6	3789	8.00	1004.0	0.603	4.906	3.765	3.819
1874	He	CO_2	Air	1.6	3607	2.00	999.2	0.842	5.050	4.279	2.984
1875	He	CO_2	Air	1.6	3916	1.00	998.1	0.707	4.557	4.646	2.076
1876	He	CO_2	Air	1.6	3899	4.00	999.7	0.973	5.589	4.019	4.273
1877	He	CO_2	Air	1.6	3814	8.00	1000.4	1.155	6.114	3.575	6.076

Continued on next page. . .

Table E.1 continued from previous page

1878	He	Air	Air	1.3	2427	1.00	997.8	0.524	3.342	4.534	1.152
1879	He	Air	Air	1.3	2615	2.00	999.1	0.605	3.845	4.339	1.852
1880	He	Air	Air	1.3	2554	4.00	1000.3	0.717	4.377	3.960	2.934
1881	He	Air	Air	1.3	2600	8.00	999.9	0.846	4.955	3.556	4.624
1882	He	Ar	Air	1.3	2539	1.00	999.6	0.377	2.836	4.626	0.619
1883	He	Ar	Air	1.3	2819	2.00	998.3	0.443	3.437	4.365	1.157
1884	He	Ar	Air	1.3	2520	8.00	999.8	0.679	4.864	3.229	3.837
1885	He	CO_2	Air	1.3	2642	8.00	1000.0	1.261	6.007	3.140	6.001
1886	He	CO_2	Air	1.3	2562	1.00	998.5	0.777	4.491	4.198	2.040
1887	He	CO_2	Air	0.8	1263	1.00	1002.1	0.915	4.365	3.469	1.973
1888	He	CO_2	Air	0.8	1218	2.00	1000.0	1.089	4.823	3.097	2.834
1889	He	CO_2	Air	0.8	1261	4.00	998.8	1.276	5.296	2.755	4.056
1890	He	CO_2	Air	0.8	1234	8.00	998.8	1.511	5.742	2.326	5.798
1891	CO_2	CO_2	Air	0.8	1287	1.00	997.3	0.382	3.385	0.913	1.327
1892	CO_2	CO_2	Air	0.8	1287	2.00	997.5	0.428	3.801	0.881	2.014
1893	CO_2	CO_2	Air	0.8	1245	4.00	998.2	0.483	4.221	0.834	3.031
1894	CO_2	CO_2	Air	0.8	1296	8.00	1000.6	0.536	4.668	0.795	4.608
1895	CO_2	Ar	Air	0.8	1265	1.00	999.1	0.230	2.486	0.959	0.580

Continued on next page. . .

Table E.1 continued from previous page

1896	CO_2	Ar	Air	0.8	1221	2.00	997.2	0.266	2.991	0.895	1.073
1897	CO_2	Ar	Air	0.8	1264	4.00	997.2	0.304	3.562	0.842	1.963
1898	CO_2	Ar	Air	0.8	1215	8.00	997.8	0.353	4.177	0.760	3.566
1899	CO_2	Air	Air	0.8	1278	1.00	998.5	0.268	2.580	0.977	0.775
1900	CO_2	Air	Air	0.8	1262	2.00	999.1	0.302	3.006	0.939	1.283
1901	CO_2	Air	Air	0.8	1117	4.00	1000.8	0.348	3.446	0.871	2.050
1902	CO_2	Air	Air	0.8	1225	8.00	1001.4	0.386	3.955	0.848	3.274
1903	Ar	Air	Air	0.8	1228	1.00	999.2	0.230	2.390	0.690	0.669
1904	Ar	Air	Air	0.8	1260	2.00	997.9	0.265	2.824	0.696	1.133
1905	Ar	Air	Air	0.8	1223	4.00	999.9	0.308	3.284	0.691	1.859
1906	Ar	Air	Air	0.8	1253	8.00	998.3	0.353	3.787	0.686	2.988
1907	Ar	Ar	Air	0.8	1228	1.00	999.2	0.208	2.370	0.686	0.545
1908	Ar	Ar	Air	0.8	1290	2.00	998.1	0.240	2.873	0.679	1.020
1909	Ar	Ar	Air	0.8	1262	4.00	997.0	0.283	3.430	0.654	1.877
1910	Ar	Ar	Air	0.8	1292	8.00	1002.1	0.331	4.045	0.624	3.453

Concluded.

Appendix F

HET OPTICAL ARM DESIGN

F.1 Design Rationale

F.1.1 Cutter Geometry

The purpose of the optical arms is to shield the outer parts of the FLDI beams from the complex flow-field established when the diffracting transmitted shock reflects off the window cavities. Simulations show that this field is restricted to the region outside the "test rhombus", which in the case of HET's circular tube exit, is the conical volume bounded by the Mach wave from the tube lip, where the flow is nominally uniform (neglecting boundary layer effects).

The "cookie cutters" are intended to isolate a portion of this uniform core flow as cleanly as possible. They have a sharp leading edge with a wedge angle chosen so that the oblique shock wave will remain attached for all freestream Mach numbers of interest. The inner faces of the cutters should be as parallel as possible with the freestream, so that only weak Mach waves will be present in the core flow. To ensure that only the core flow is isolated, the leading edges need to lie within the original Mach cone. These main geometrical considerations are illustrated in Fig. F.1.

The angle of the wedge is a trade-off: smaller angles allow the shock to stay attached to lower M, but also necessitate the wedge being longer, thinner, more difficult to machine, and more susceptible to damage from diaphragm debris. An angle of $\theta = 25°$ guarantees the shock remaining attached down to $M = 3$, even for the worst-case $\gamma = 5/3$ (monatomic test gases); for diatomics at $\gamma = 7/5$, the limit is $M \approx 2.2$. Unfortunately, this precludes the arms from being used with HET operating in shock-tube mode, as there are no shock-tube conditions exceeding $M = 2.2$ for any pressure or gas combination within HET's operational limits.

The width of the cutters transverse to the flow direction is largely bounded by two factors: the diameter of the cylindrical arms they attach to, and the diameter of the tube. The cutters need to be able to penetrate partially into the tube without colliding with it; the bevels on the outer edges of the cutter give a little more clearance for this.

The other components of the arm are all made from aluminum to reduce material

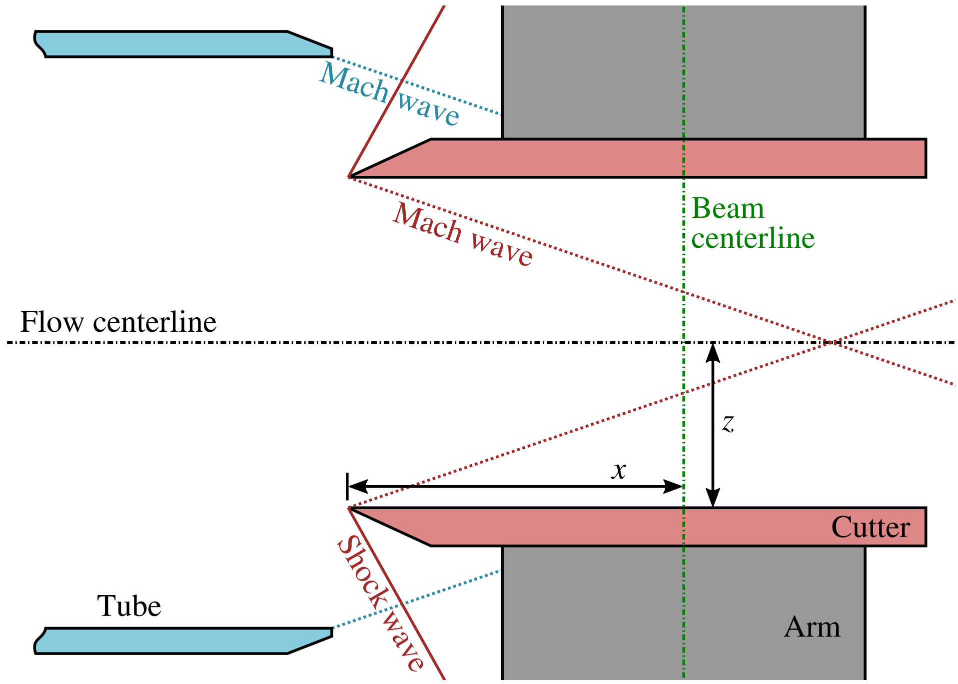

Figure F.1: Schematic of main wave processes for HET optical arms.

and manufacturing expense; the cutters themselves are made of A2 tool steel, in order to better hold a sharp edge, and to resist particulate damage.

The nominal location of the FLDI foci is at the center of the core flow, i.e. the intersection of the flow and beam centerlines in Fig. F.1, at co-ordinates (x, z) relative to the leading edge. The intersection of the Mach waves from the leading edges should lie downstream of the foci as illustrated, so that the flow interrogated by the FLDI is either undisturbed, or at most processed by a single Mach wave. Some simple trigonometry yields the following requirement:

$$x < z \times \sqrt{M^2 - 1} \tag{F.1}$$

Since z is fixed by other geometrical constraints at 32 mm, then for the worst-case $M = 3$, $x < 101$ mm is required; the chosen dimension of 88.9 mm is within this.

The confined space within the HET test section means that the beam profiler cannot be used once the arms are installed. Hence measurements of the foci size and location are made without the arms; initially the foci are located on the flow centerline as intended. However, the relatively thick windows of the arm act as weak lenses, causing the foci to shift down-beam slightly. Although it is difficult to accurately locate the foci without the profiler, they were approximately measured as being

18 mm off the pitch-side window. This means the foci still lie in undisturbed flow for $M \gtrsim 5.0$; in any case it should not matter due to being a weak Mach wave. Bench-top testing using similar windows showed that although the foci position is displaced, the foci size and separation is not altered to within the precision of the beam profiler.

F.1.2 Arm Geometry

To reduce costs, the arm tube itself used hollow cylindrical stock, chosen with a suitable internal diameter that allowed the interior to be left unmachined. A detachable flange of similar dimensions to the standard HET windows is clamped in place by the window flange.

To avoid costly welding to vacuum standards, the arm attaches to both the flange and cutter via bolts. The main downside of this approach is that the cutter bolts must be inserted from the face that touches the core flow. To minimize disturbances, the bolts were countersunk to below the surface datum, then back-filled with a low-viscosity self-leveling silicone. This gives a smooth surface finish, and while the silicone proved durable over the course of many HET shots, it is also easy to remove from the bolt heads if the arms need disassembling.

The wall thickness of the arm tube and the choice of bolts were verified by considering the worst-case loading during facility operation. The peak dynamic pressure, $\max\left(p_{dyn}\right)$, during a shot was used as the basis for this. Because the acceleration section is much lower density, it is reasonable to assume that $\max\left(p_{dyn}\right)$ for an expansion-tube shot never exceeds $\max\left(p_{dyn}\right)$ for the equivalent shock-tube condition (i.e. the same driver and driven states, but no secondary diaphgram). The theoretical $p_{dyn}(t)$ for a shock tube was computed, and found to compare well with pitot data. It was further determined that $\max\left(p_{dyn}\right)$ never exceeds the original driver pressure p_4. To apply static-load analysis to highly dynamic loads, the equivalent static load was taken to be $2p_{dyn}$, which leads to the very conservative target loading of $2p_4 \approx 6.6\,\text{MPa}$ for the thickest primary diaphragm.

This pressure was assumed to act on the entire projected area of the arm tube, rather than just the portion exposed to the flow. A simple finite-element analysis performed in Solidworks predicted negligible bending and peak stresses well below yield, even for this contrived worst-case scenario. For the bolts, the projected area of the cutter was used to compute the shear stress on each bolt, and the exposed rear area was used to compute the tensile stress of the cutter trying to pull away from the arm (this latter scenario will never actually experience a pressure differential anything close to $2p_4$, so the safety margin is very high).

Another limitation of the confined space is that the cutters do not fit through the test section window openings, so the clamped flange is first installed, then the pre-assembled remainder of the optical arms is brought through the front of the test section and bolted into place.

One small but useful feature is the pair of orthogonal grooves scribed onto the outer face of the clamped flange. These are used as alignment targets, in combination with a self-leveling laser cross-hair projector that has proven a valuable tool for optical alignment within our research group.

F.1.3 Window & Seal Design

While the design of the metal parts, bolts and O-rings for metal-to-metal contact all followed standard procedures, the tolerancing and mounting of the glass window was somewhat more involved. The window is sandwiched between two O-rings to provide a vacuum seal and to avoid any metal-to-glass contact; it has a shoulder to allow a flush surface facing the core flow. The design methodology took into account the plate bending and compressive stresses on the window, following the calculations and failure curves given for a similar window used in LIGO (Coyne, 2011). The window clamp is flush with the back surface of the cutter at the design condition, so it is impossible to exceed the design stress on the window by over-tightening the screws.

While the O-rings between metal parts (Items 10–12 in the assembly drawing) use standard grooves for axial seals, the grooves for the window seals (Item 9) are shallower than standard to prevent metal-to-glass contact. A tolerance stack-up was performed involving both grooves and their O-rings, the window, and the window recess in the cutter; even for the worst combinations of tolerances, the glass cannot touch either metal part. Furthermore, both O-rings remain within their allowable squeeze ranges, and the compressive forces do not exceed safe limits for the glass. This applies both for the static assembly, and under vacuum loading (where there is a pressure differential of approximately 1 atm across the window, causing one O-ring to further compress while the other relaxes). The nominal condition has the outer face of the window aligned with the cutter surface; the radial tolerances are tight to give a minimal annular gap. The off-nominal condition is designed to favor a slightly recessed window rather than protruding, as it is thought that this will disturb the flow less. All seals, including the ones touching glass, use vacuum grease. The windows are uncoated silica, so it is not harmful to clean any excess grease off using solvents.

Initial vacuum testing showed that the optical arm assembly actually gave a better overall leak rate than with the standard HET windows installed. The windows also endured the rapid decompression experienced during the regular pump-down procedure, as well as the impulsive loading of the shock waves and processed driver gases. After more than 70 shots, there was no sign of damage to the windows, nor were they even dirty. The standard HET windows tend to become dirty and pitted over time due to impact from secondary diaphragm particulates. It is thought that the cutters deflected these particulates into the outer flow, because the upstream-facing portions of the arm tubes (made of soft aluminum) became substantially pitted after several shots (Fig. F.2). However, the steel cutter arms showed no sign of degradation at the end of the campaign. As previously mentioned, the silicone "caps" covering the recessed bolts were not damaged either.

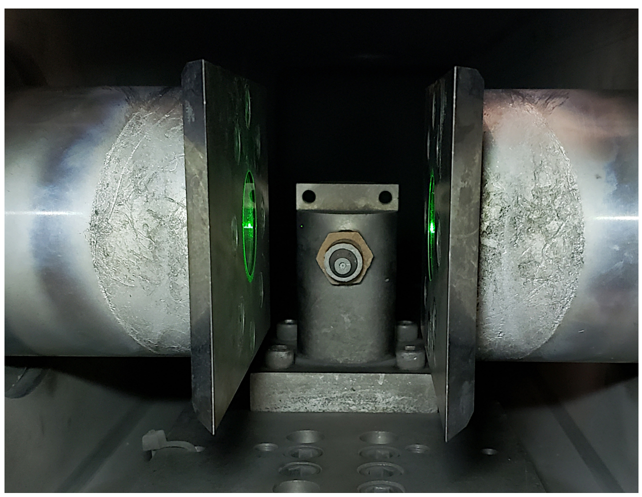

Figure F.2: Photo of the installed optical arms, looking downstream. Pitting on the arm tubes corresponds to the projection of the HET tube exit.

F.2 Engineering Drawings

The following pages contain the engineering drawings of all the custom-made components of the optical arm assembly (standard fasteners and O-rings are called out in these drawings). These are:

1. Arm tube
2. Clamped flange
3. Cookie cutter
4. Window clamp
5. Window

The four metal components were manufactured by H&S Enterprises [Monrovia, CA] while the glass window was made by Esco Optics [Oak Ridge, NJ].

An overview drawing of the full optical arm assembly is also included. For views of the assembly installed into HET, please see Fig. 2.8.

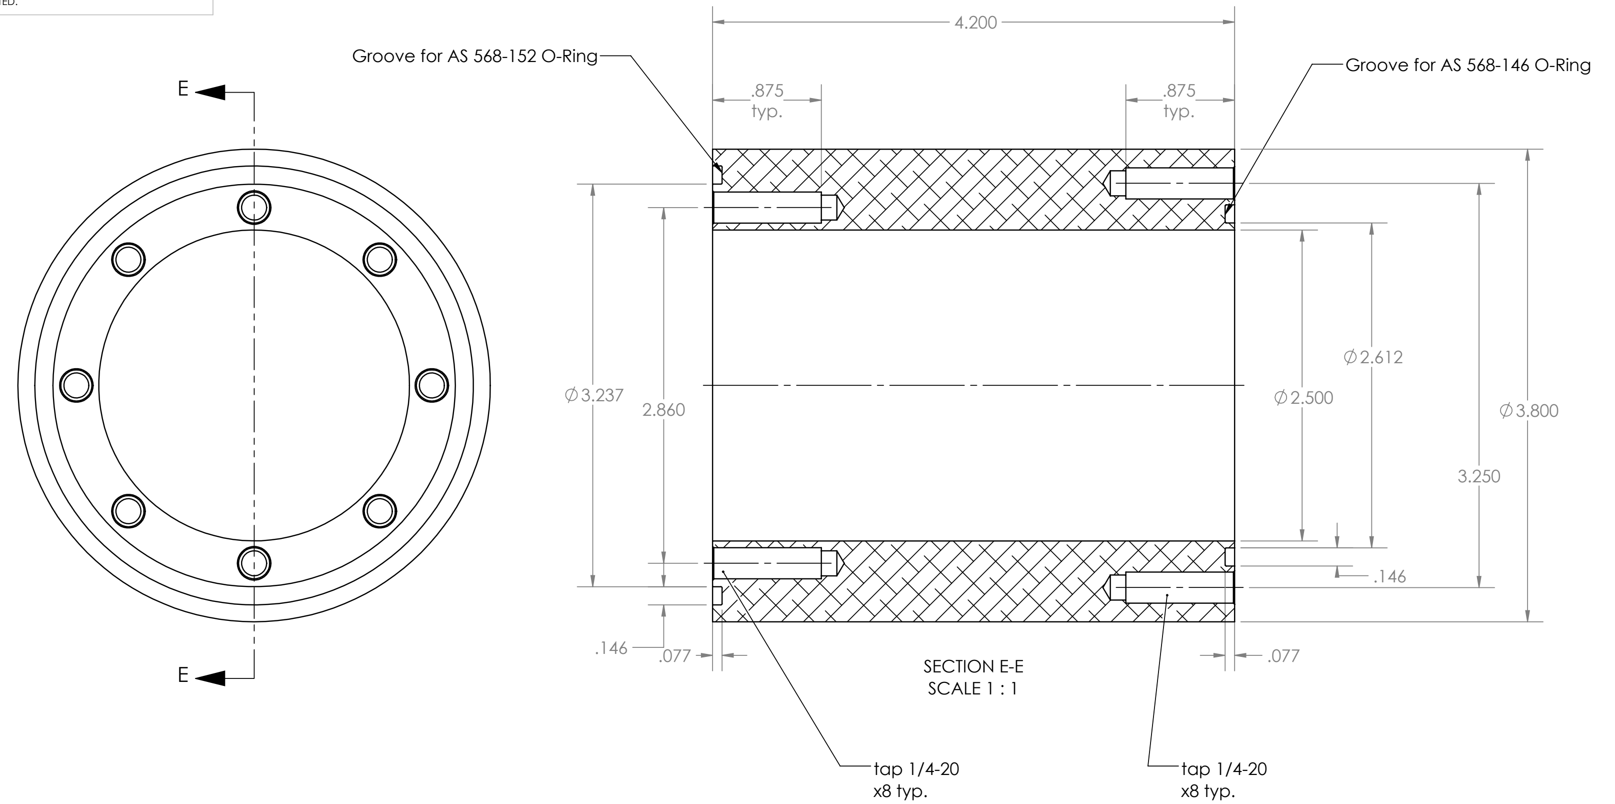

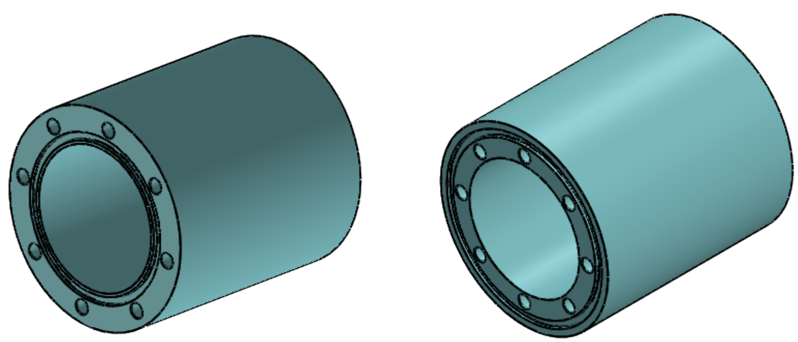

UNLESS OTHERWISE SPECIFIED DIMENSIONS ARE IN INCHES TOLERANCES ARE: FRACTIONS 1/16 DECIMALS ±0.005 ANGLES 1°	CAD GENERATED DRAWING, DO NOT MANUALLY UPDATE	NAME Arm tube
MATERIAL 6061 Alloy	APPROVALS / DATE DRAWN Joel Lawson CHECKED Bahram Valiferdowsi	**Austin Laboratory Caltech, Pasadena**
WEIGHT	RESP ENG	
FINISH	MFG ENG	SIZE B DWG. NO. REV.
QUANTITY two pieces	QUAL ENG	SCALE CAD FILE: SHEET 1 OF 1

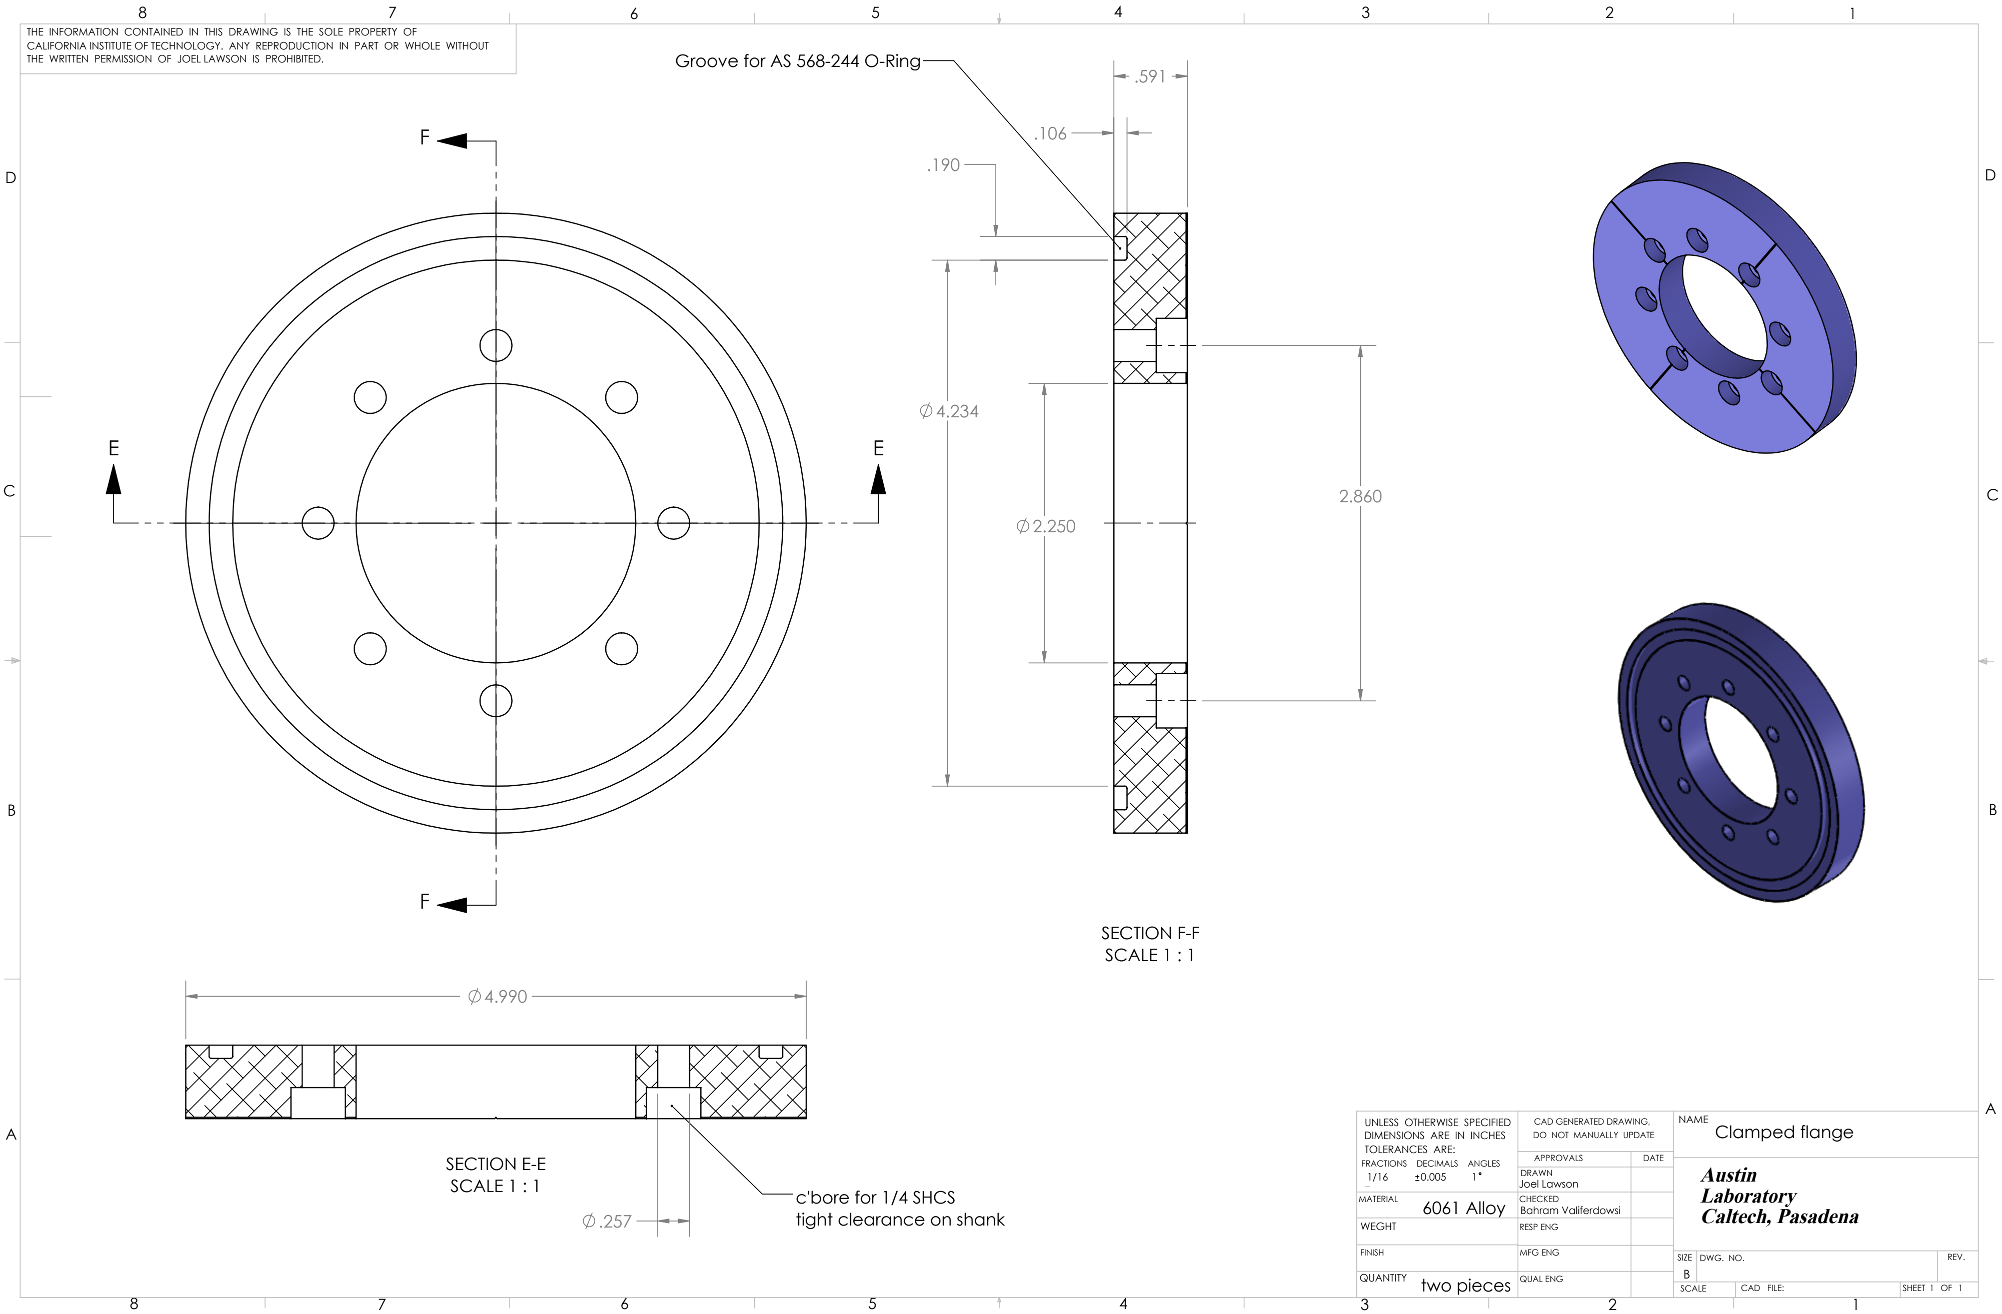
THE INFORMATION CONTAINED IN THIS DRAWING IS THE SOLE PROPERTY OF CALIFORNIA INSTITUTE OF TECHNOLOGY. ANY REPRODUCTION IN PART OR WHOLE WITHOUT THE WRITTEN PERMISSION OF JOEL LAWSON IS PROHIBITED.
Groove for AS 568-244 O-Ring
.591
.106
.190
Ø4.234
Ø2.250
2.860
F
F
E
E
SECTION F-F
SCALE 1 : 1
Ø4.990
SECTION E-E
SCALE 1 : 1
Ø.257
c'bore for 1/4 SHCS
tight clearance on shank
UNLESS OTHERWISE SPECIFIED
DIMENSIONS ARE IN INCHES
TOLERANCES ARE:
FRACTIONS 1/16
DECIMALS ±0.005
ANGLES 1°
CAD GENERATED DRAWING, DO NOT MANUALLY UPDATE
APPROVALS
DATE
DRAWN Joel Lawson
CHECKED Bahram Valiferdowsi
RESP ENG
MFG ENG
QUAL ENG
MATERIAL 6061 Alloy
WEGHT
FINISH
QUANTITY two pieces
NAME Clamped flange
Austin
Laboratory
Caltech, Pasadena
SIZE B
DWG. NO.
REV.
SCALE
CAD FILE:
SHEET 1 OF 1

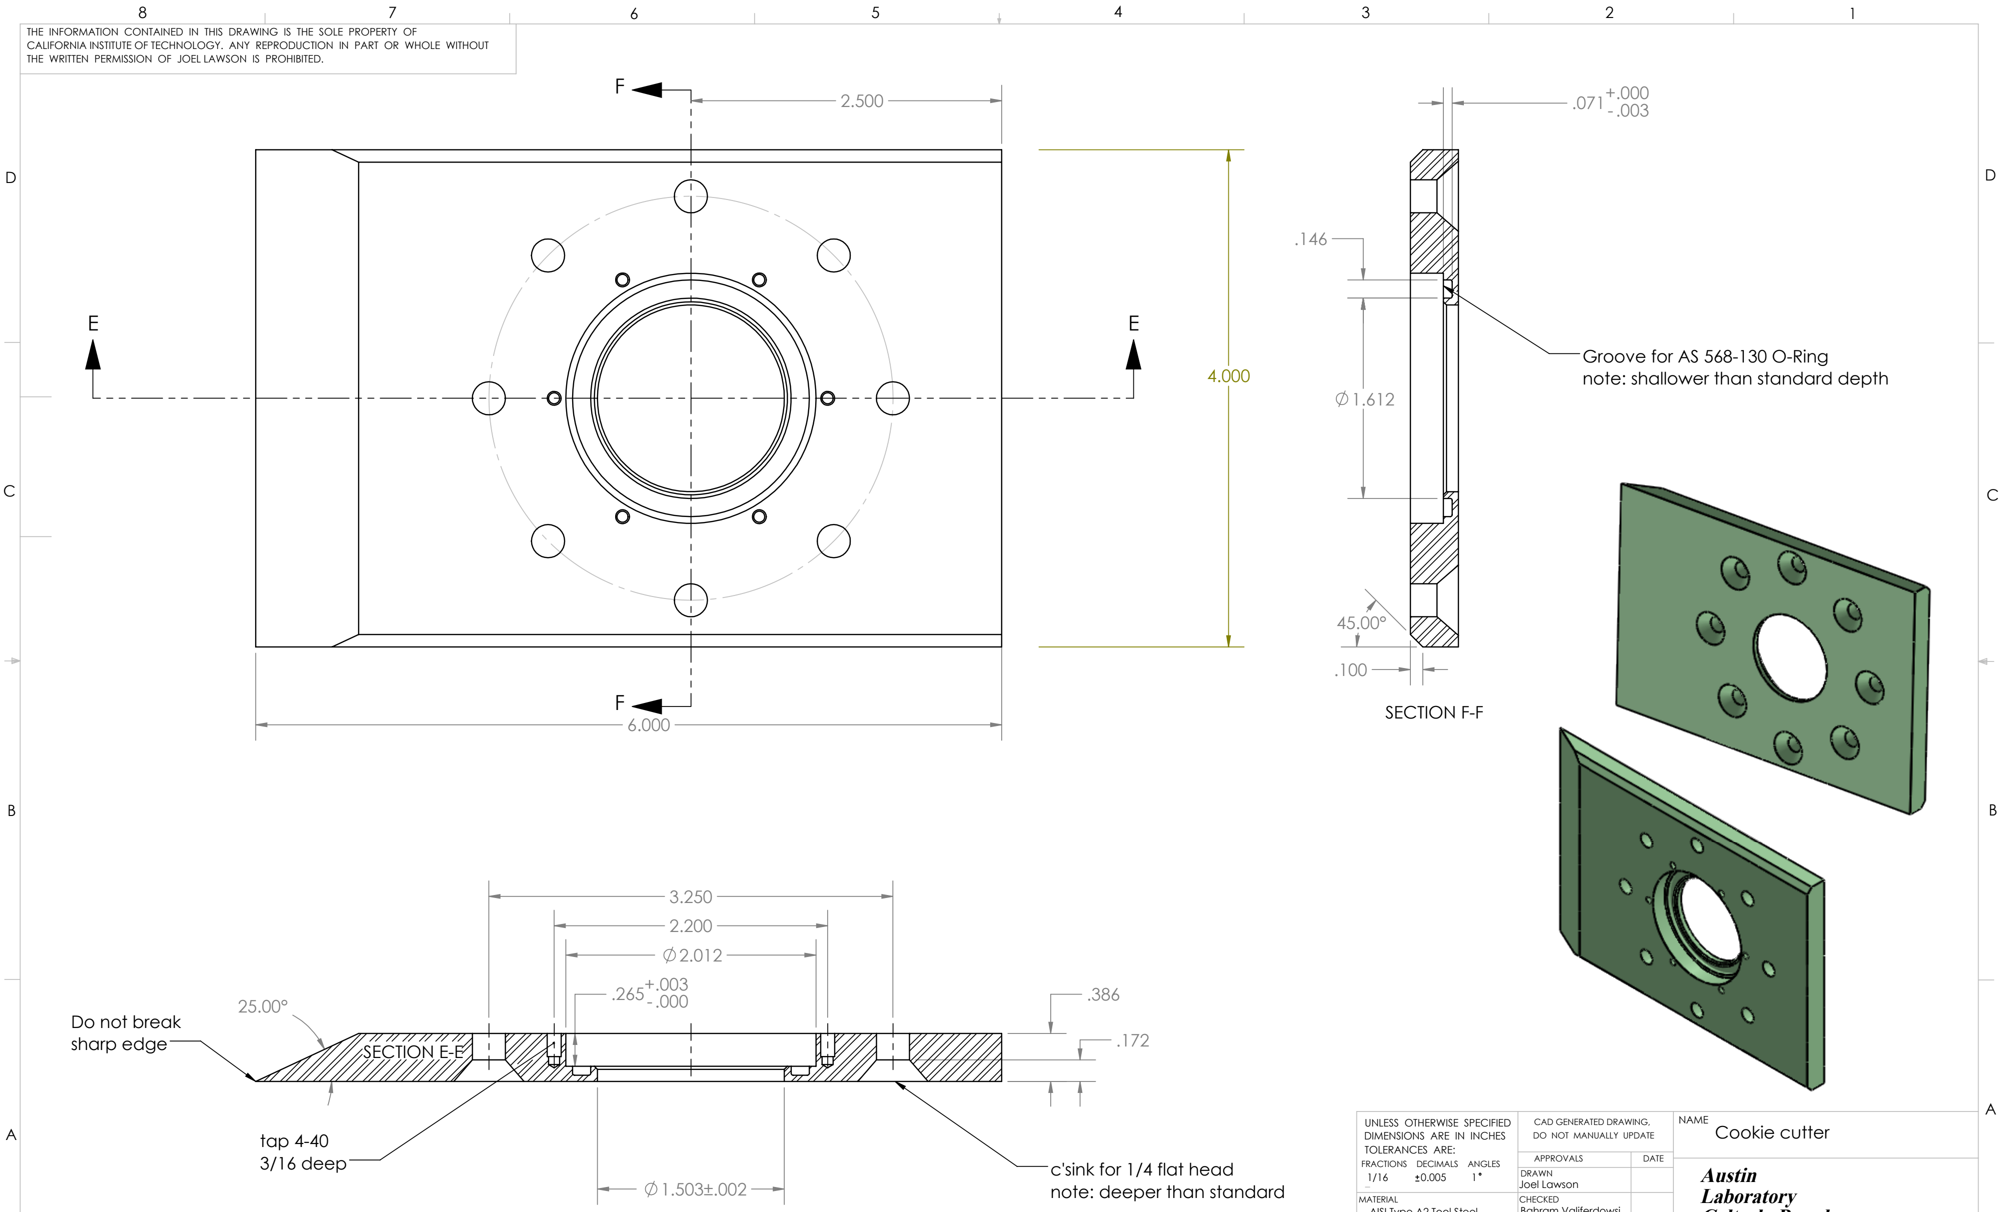

UNLESS OTHERWISE SPECIFIED DIMENSIONS ARE IN INCHES TOLERANCES ARE: FRACTIONS 1/16 DECIMALS ±0.005 ANGLES 1°	CAD GENERATED DRAWING, DO NOT MANUALLY UPDATE	NAME Cookie cutter
	APPROVALS / DATE	
	DRAWN Joel Lawson	**Austin Laboratory Caltech, Pasadena**
MATERIAL AISI Type A2 Tool Steel	CHECKED Bahram Valiferdowsi	
WEGHT 2.031 Lb	RESP ENG	
FINISH -	MFG ENG	SIZE B DWG. NO. REV.
QUANTITY two pieces	QUAL ENG	SCALE CAD FILE: SHEET 1 OF 1

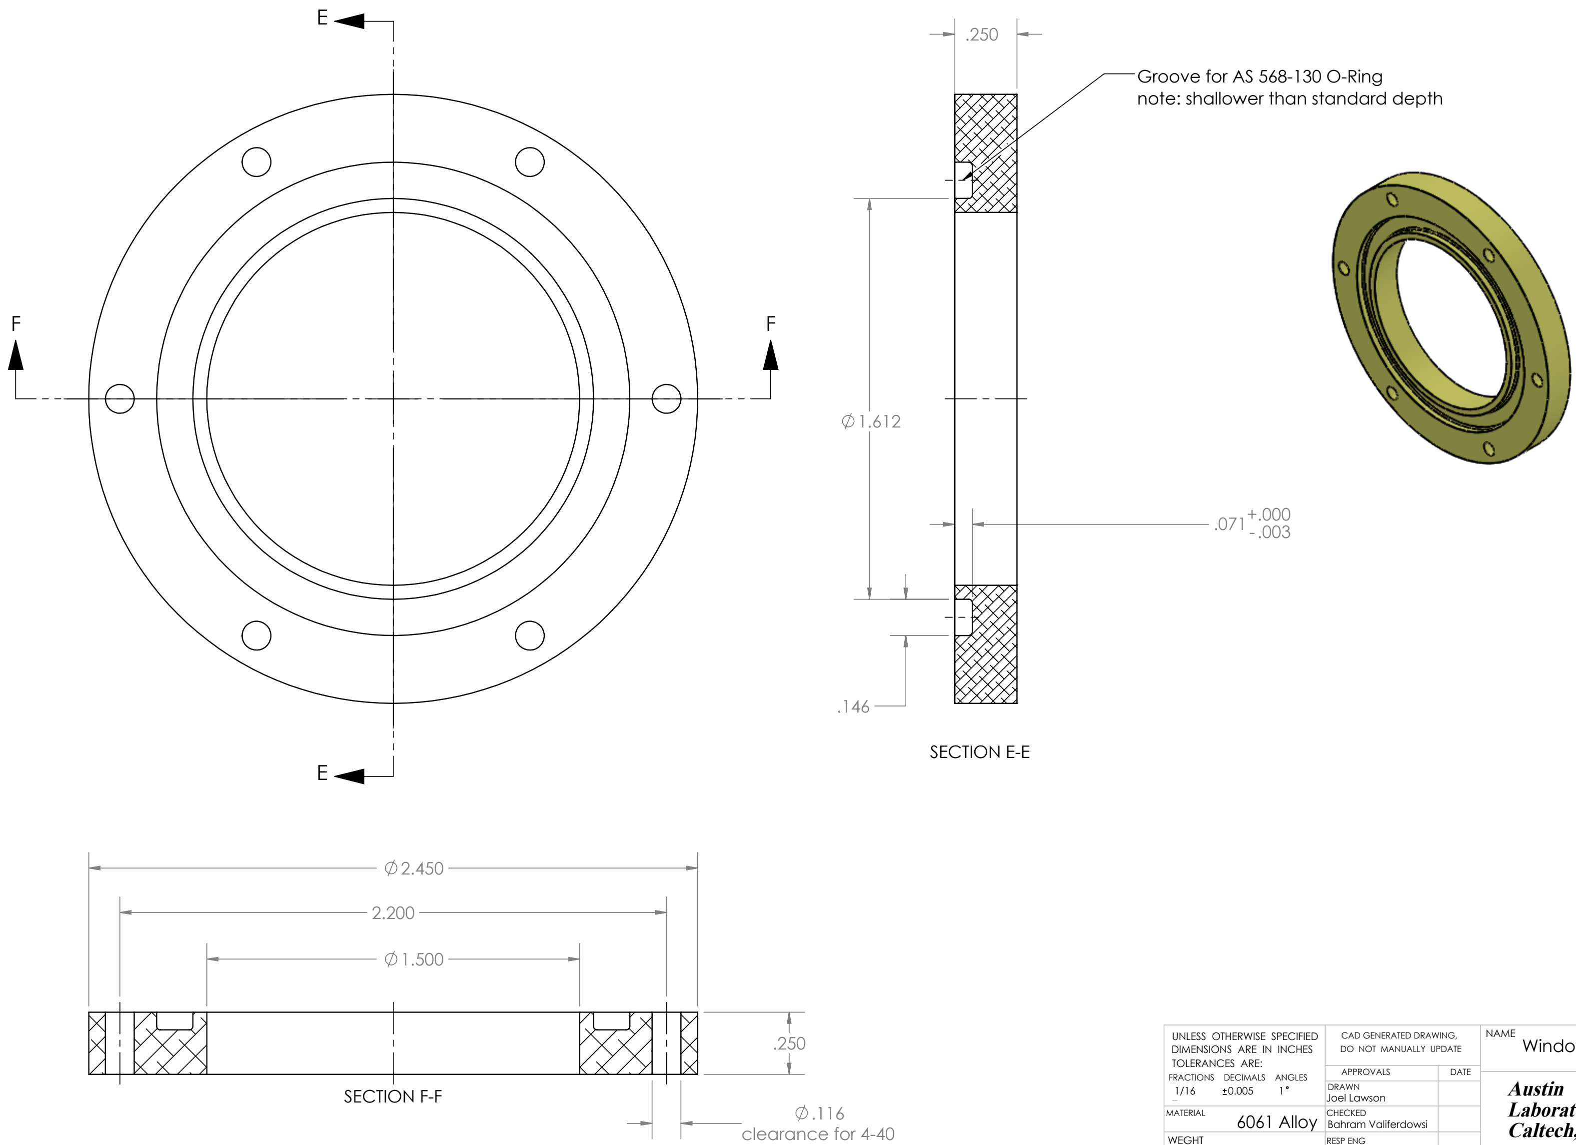

UNLESS OTHERWISE SPECIFIED DIMENSIONS ARE IN INCHES TOLERANCES ARE: FRACTIONS 1/16 DECIMALS ±0.005 ANGLES 1°	CAD GENERATED DRAWING, DO NOT MANUALLY UPDATE	NAME Window clamp
	APPROVALS / DATE	
	DRAWN Joel Lawson	***Austin Laboratory Caltech, Pasadena***
MATERIAL 6061 Alloy	CHECKED Bahram Valiferdowsi	
WEIGHT	RESP ENG	
FINISH	MFG ENG	SIZE B DWG. NO. REV.
QUANTITY two pieces	QUAL ENG	SCALE CAD FILE: SHEET 1 OF 1

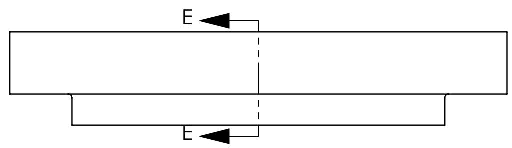

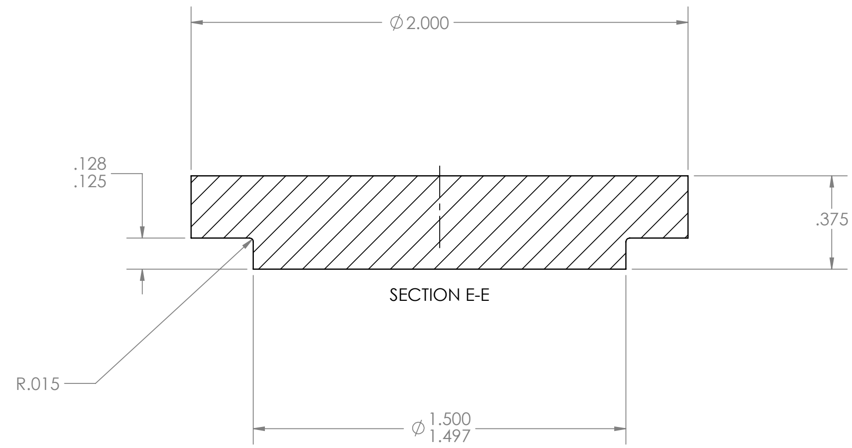

SECTION E-E

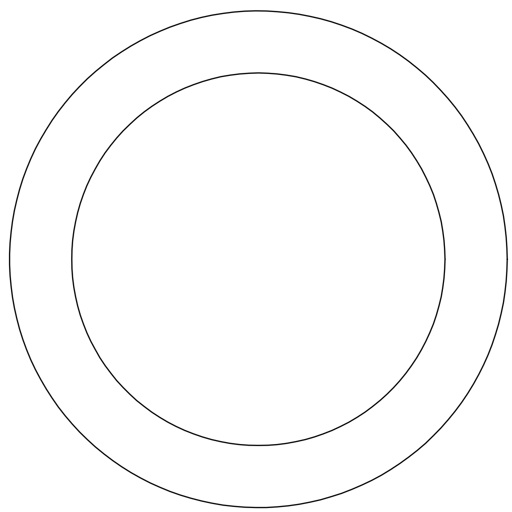

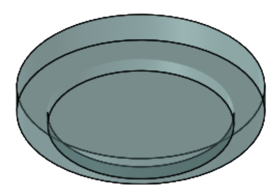

UNLESS OTHERWISE SPECIFIED DIMENSIONS ARE IN INCHES TOLERANCES ARE: FRACTIONS 1/16 DECIMALS ±0.005 ANGLES 1°	CAD GENERATED DRAWING, DO NOT MANUALLY UPDATE	NAME Window
	APPROVALS / DATE	*Austin Laboratory Caltech, Pasadena*
	DRAWN Joel Lawson	
MATERIAL	CHECKED Bahram Valiferdowsi	
WEGHT	RESP ENG	
FINISH	MFG ENG	SIZE B / DWG. NO. / REV.
QUANTITY	QUAL ENG	SCALE / CAD FILE: / SHEET 1 OF 1

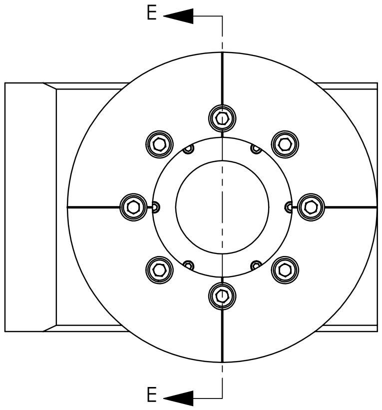

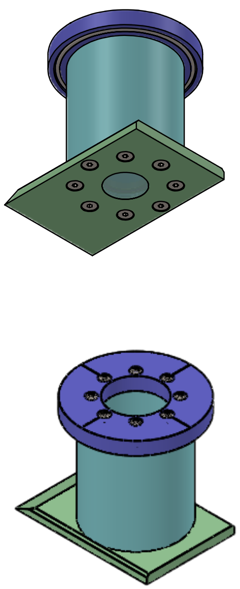

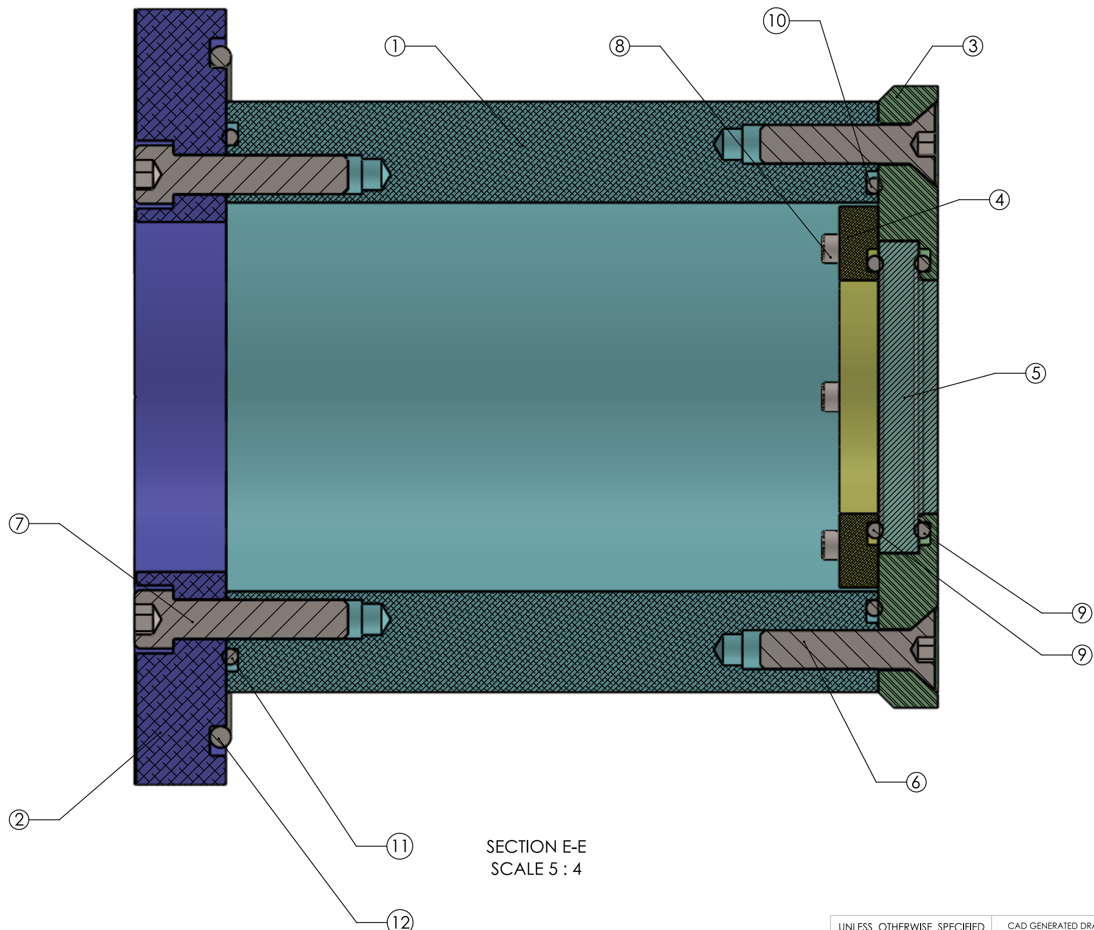

SECTION E-E
SCALE 5 : 4

ITEM NO.	DESCRIPTION	QTY.
1	Arm tube	2
2	Clamped flange	2
3	Cookie cutter	2
4	Window clamp	2
5	Window	2
6	Flat head screw, 1/4"-20, 1-1/8" long	8
7	Socket head screw, 1/4"-20, 1-1/8" long	8
8	Socket head screw, 4-40, 3/8" long	6
9	Buna-N O-ring, AS568-130	2
10	Buna-N O-ring, AS568-146	1
11	Buna-N O-ring, AS568-152	1
12	Buna-N O-ring, AS568-244	1

UNLESS OTHERWISE SPECIFIED DIMENSIONS ARE IN INCHES TOLERANCES ARE:
FRACTIONS 1/16 DECIMALS ±0.005 ANGLES 1°

CAD GENERATED DRAWING, DO NOT MANUALLY UPDATE

APPROVALS		DATE
DRAWN	Joel Lawson	
CHECKED	Bahram Valiferdowsi	
RESP ENG		
MFG ENG		
QUAL ENG		

MATERIAL

WEGHT

FINISH

QUANTITY

NAME Optical arm assembly

Austin Laboratory Caltech, Pasadena

SIZE B DWG. NO. REV.

SCALE CAD FILE: SHEET 1 OF 1