VELOCITY PHASE ENCODED MRI OF GAS FLOW IN THE ACOUSTIC BOUNDARY LAYER

by

Geoffrey Archibald

B.Sc. (Physics), Mount Allison University, 1998

THESIS SUBMITTED IN PARTIAL FULFILLMENT

OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN THE

DEPARTMENT OF PHYSICS

FACULTY OF SCIENCE

APPROVAL

Name: Geoffrey Archibald

Degree: Doctor of Philosophy

Title of Thesis: Velocity Phase Encoded MRI of Gas Flow in the Acoustic Boundary Layer

Examining Committee: Dr. J. Steven Dodge (Chair)

Dr. Michael Hayden, Senior Supervisor
Professor, Department of Physics, SFU

Dr. Jenifer Thewalt, Supervisor
Professor, Department of Physics, SFU

Dr. Igor Herbut, Supervisor
Professor, Department of Physics, SFU

Dr. David Broun, Internal Examiner
Associate Professor, Department of Physics, SFU

By written consultation from Los Alamos

Dr. Gregory Swift, External Examiner
Retired Laboratory Fellow and Guest Scientist,
Los Alamos National Laboratory

Date Approved: August 19, 2011

Partial Copyright Licence

Abstract

This thesis explores the use of magnetic resonance imaging (MRI) to study acoustic oscillations of a gas in a cylindrical tube. It describes experiments performed under conditions where the gas is in the Acoustic Viscous Boundary Layer and its flow is laminar. Velocity maps acquired at discrete phases of the acoustic oscillation are presented, and are compared with thermoacoustic theory. This represents the first time that such information has been obtained using MRI.

An important component of the work reported in this thesis involves the design, construction, and characterization of an acousto-mechanical resonator (AMR). This device can drive oscillatory gas motion and impose density variations at rates that are compatible with MRI data acquisition. To date it has been operated at frequencies ranging from 0.7 Hz to 1.65 Hz and with peak gas displacement amplitudes of up to 2.5 cm. The AMR is based on a modular design intended to permit the study of acoustic flow through a variety of different structures and under a variety of different conditions.

MRI experiments were performed on a mixture of thermally-polarized ^{3}He and O$_2$. The latter is used to increase the ^{3}He longitudinal nuclear relaxation rate T_1^{-1} to a value comparable to the acoustic frequency. In turn, measurements of T_1^{-1} provide a means for determining the precise composition of the gas mixture. Velocity phase-encoding techniques were then used to map acoustic flow fields: A bipolar magnetic field gradient pulse inserted into the imaging sequence stores velocity information in the phase of the complex image data. The MRI pulse sequence is synchronized with the periodic motion of the gas so that the velocity measurement can be performed at discrete and well-defined phases of the acoustic cycle. These non-invasive flow imaging experiments provide information that is complementary to that which can be obtained from other gas velocity probes, and may lead to new opportunities in the study of acoustic devices.

Acknowledgments

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Magnetic resonance imaging (MRI) is a powerful imaging modality and an indispensable tool for medical diagnostics. Decades ago it was already clear that simple MRI images – reflecting contrasting characteristics of different structures and tissues – would provide medical practitioners with valuable new insight. Still, it is almost certain that the capabilities of modern MRI far exceed the vision of those who first contributed to its development. Examples of modern applications range from monitoring brain activity [1, 2] and its chemical makeup [3] to visualizing and tracking blood flow [4, 5, 6]. MRI has also found significant application outside of the medical industry, in fields ranging from materials science [7, 8] to gas exploration [9]. The vast majority of these applications involve situations in which the underlying substance that is being probed is dense, usually a liquid or liquid-rich tissue. In recent years, however, there have been an increasing number of promising NMR- and MRI-based studies involving *gaseous* samples [10]. In most of these experiments, the gas acts as a probe (or in the language of MRI – as a "contrast agent") intended to reveal information about its environment: usually some sort of complex or convoluted structure like the human lung [11, 12, 13]. In a small subset of these experiments, it is the gas itself – or more precisely – its *dynamics* that are of interest. So, for example, it has been demonstrated that one can map out the velocity field associated with gas flowing through various channels [14, 15, 16] and around obstructions [17, 18]. To date, these demonstrations have been limited to unidirectional flow.

In this thesis I describe the first MRI-based study of gas motion in an acoustic field, of which I am aware. That is, a situation in which the direction of gas flow is periodically

reversed. In particular, I focus on a situation in which the gas under study is in what is known as the "Acoustic Viscous Boundary Layer" or AVBL. This corresponds to the volume of gas that is close enough to an immobile solid surface so that its motion is strongly influenced by momentum transfer with the wall. The model system that I describe and study involves long-wavelength acoustic oscillations in a gas confined to a very simple and well defined channel: a long, narrow, cylindrical tube. In a sense, this problem is the acoustic analog of Poiseuille flow through a cylindrical pipe. It is a problem that is very well understood from both theoretical and experimental perspectives.

Despite the simplicity of the model system described in my thesis, the more general problem of gas dynamics in the AVBL is of considerable technological interest. These dynamics are fundamental to the operation of a broad and versatile class of "thermoacoustic" devices, which primarily act as heat engines and refrigerators [19, 20, 21]. I thus anticipate this work will be of interest to researchers in both the MRI and thermoacoustics communities, between which there has not traditionally been much overlap. My goal in presenting this work is neither to develop radically new MRI sequences nor to study a particularly challenging problem in acoustic flow. Rather it is to explore the manner in which one might adapt reasonably well established imaging techniques to a very well-defined acoustic system. In this sense it is a feasibility study, intended to explore the viability and utility of MRI as a probe of acoustic processes.

In the remainder of this introductory chapter I present a brief introduction to the field of thermoacoustics and highlight some of the issues that the thesis will attempt to explore. For context, I then describe some existing experimental probes of oscillatory gas flow. The goal here is simply to identify some benchmarks against which the performance of MRI-based probes will ultimately have to be judged. This is followed by a statement of the experimental challenge, that is, an outline in general terms of the apparatus that I set out to construct and the experiments that I was able to perform. The chapter concludes with a brief outline for the remainder of the thesis.

1.1 Thermoacoustics

Thermoacoustics is a field of study that brings the principles of thermodynamics to bear on the study of acoustic oscillations in a compressible fluid that occur in narrow chan-

nels or near immobile surfaces. The earliest scientific enquiries in this field are connected to the observation of spontaneous sound production in narrow tubes along which intense thermal gradients are imposed [22, 23]. Lord Rayleigh set the groundwork for future theoretical descriptions of the effect more than a century ago [24], but his picture was limited to tubes with large transverse dimensions and ultimately failed to produce accurate predictions. Although others contributed to the field over many decades [25, 26], it was Nikolaus Rott who in 1969 first managed to integrate acoustics and thermodynamics to give a consistent mathematical description of low amplitude acoustic oscillations of a gas in a narrow tube (and other geometries) along which a temperature gradient is imposed [27, 28, 29, 30, 31, 32, 19]. Rott's work was primarily motivated by the observation of what are known as "Taconis oscillations," in which acoustic oscillations are spontaneously produced in narrow tubes in which one end (usually open) is cooled to low temperatures by immersion in a cryogenic liquid while the other (usually closed) remains at room temperature. It was quickly realized, however, that his "thermoacoustic theory" applied to an emerging and potentially important class of acoustic devices including pulse tube [33] and standing-wave refrigerators [34] and the traveling-wave acoustic amplifier [35]. The insight provided by Rott's theory has since proved to be invaluable. At a very fundamental level the framework he established paved the way for some remarkable innovations in acoustic engineering, including development of the thermoacoustic Stirling heat engine [36, 37, 38].

It is well known that sound waves involve gas displacement and pressure/density oscillations, but the fact that temperature variations are also usually involved is often overlooked. By coupling temperature and displacement oscillations (with appropriate phasing) to solid surfaces (such as the wall of a confining tube, or some other intentionally inserted structure) it is possible to convert heat into acoustic energy or to use acoustic power to transport heat from one location to another, including up a thermal gradient. In other words, one is able to harness acoustic oscillations in a gas – appropriately coupled to solid structures that act as thermal reservoirs – to construct heat engines and refrigerators. A key advantage of thermoacoustic heat engines and refrigerators over more conventional devices is that they can often be designed with few or no moving parts, and thus can in principle be very robust against the wear normally encountered with sliding seals. Moreover, the design tolerances required by thermoacoustic devices tend to be less demanding than those required by conventional devices, and the working fluid is typically helium gas which is in-

ert. The greatest commercial success relevant the field of thermoacoustics is the pulse-tube cryocooler [33, 39, 40], but from the standpoint of technical merit and intrinsic potential the somewhat newer class of hybrid thermoacoustic-Stirling heat engines [36, 37] is probably on an equal footing. Both have achieved remarkable thermodynamic efficiencies. These devices use the controlled phasing of acoustic waves to move their working gases through the necessary thermodynamic cycles, eliminating the sometimes complex arrangement of pistons and cams required by conventional heat engines and refrigerators.

In order to achieve high power density, acoustic heat engines and refrigerators typically operate outside of the small amplitude limit or "acoustic approximation" that is inherent in Rott's original formulation of thermoacoustic theory. This presents a design challenge. Non-linearities such as turbulence, entrance effects and mass streaming come into play as acoustic amplitudes increase. These effects introduce parasitic losses that tend to degrade device efficiency relative to theory. In order to account for these losses (and more importantly, to mitigate their effect on efficiency through design, or to harness them for useful purposes) one is often forced to resort to empirical characterizations of complex acoustic phenomena. Turbulence, for example, has long been studied for steady flow through straight pipes, bends, junctions, and transitions involving changes in tube diameter. In each case, these systems exhibit rich and complex behaviour. If one now adds to this the constantly changing motion associated with acoustic flow, the issues one needs to consider become incredibly complex. Impressive progress has been made toward identifying and characterizing their influence through local measurements of acoustic pressure (and in some cases through probes of gas velocity). Still, at an intuitive level, it would be very useful to have a tool that would enable one to peer inside a functioning thermoacoustic device and observe gas motion in an unobtrusive manner. The hope is that this would help to identify and improve understanding of non-linear processes, and possibly take the efficiency of devices to a higher level. The imaging of gases as they undergo acoustic oscillations is an area of research that is actively being pursued, but which is still very much in its infancy. Below I summarize a few experimental techniques that are currently employed.

1.2 Methods for Measuring Oscillatory Gas Flow

Several experimental methods have been developed for studying gas motion: Hot Wire Anemometry (HWA), Laser Doppler Anemometry (LDA), and Particle Image Velocimetry (PIV) as well as various forms of MRI-based probes including Velocity Phase Encoded MRI (VPE-MRI). The first three methods have all been used previously to study thermoacoustic phenomena. This thesis describes the first attempt to apply VPE-MRI to the study of gas flow in the acoustic boundary layer.

HWA uses heat loss from a thin electrically heated wire to infer the velocity of a fluid as it passes by. If the resistance of this wire is temperature dependent, and if the current passing through the wire is adjusted to keep the wire resistance constant, then a measurement of the electrical power dissipated by the wire can be used to infer the rate at which heat is lost to the environment. As long as this heat loss is dominated by forced convection driven by flow, one can infer velocity. This works very well for unidirectional flows, but the interpretation of HWA data is more challenging for oscillatory flows. The problem is that the heated wire ends up repeatedly interacting with its own thermal plume. Huelsz *et al.* have shown that with careful calibration, HWA can indeed be applied to oscillatory flow [41]. However, when these techniques are applied to flow in the AVBL, systematic deviations between theory and experiment are observed [42].

LDA employs light reflected from seed particles that are injected into a fluid stream. Crossed laser beams are focused to a small area, typically around 50 μm in size, forming a linear interference pattern. As a particle passes through this interference pattern, the intensity of reflected light is modulated at a frequency that is proportional to velocity. The point at which the laser beams intersect can then be scanned throughout the sample to map the gas velocity as a function of position. A variety of seed particles can be used in conjunction with gas flows, including smoke, olive oil and water glycerol mixtures. Castrejón-Pita *et al.* have used LDA to observe oscillatory gas motion in the acoustic viscous boundary layer near a wall [43]. Others, including Yazaki *et al.* [38] and Bailliet *et al.* [44], have used LDA to probe oscillatory gas motion in more complex situations, including the flow encountered inside thermoacoustic engines.

PIV employs light reflected from seed particles that are illuminated by a thin sheet of laser light. This reflected light is imaged using a camera that is aimed perpendicular to

the light sheet. A sequence of images is then acquired as the laser light is pulsed. The displacement of particles from one image to the next is then determined and used to infer fluid velocity. Castrejón-Pita *et al.* have compared PIV and LDA as probes of gas motion in the acoustic boundary layer [43]. With a spot size of volume 0.64 mm × 0.075 mm × 0.075 mm, their LDA results showed the greatest accuracy and reproduced theoretical expectations – even at very short distances away from the wall. Others, including Shi *et al.* [45], have used PIV to study more complex problems such as vortex shedding from a thermoacoustic stack in an acoustic flow.

A recent review of laser-based methods for characterizing the motion and temperature of a gas in a thermoacoustic context has been presented by Shi *et al.* [46]. One of the assumptions underlying both the LDA and PIV approaches is that the motion of seed particles reflects that of the gas. Care must be taken to select particles that are large enough to reflect a detectable amount of light and yet are small enough to faithfully follow the flow. LDA and PIV also both require optical access in order to illuminate particles and detect reflected light. Typically this means that the apparatus must be built with optically flat transparent materials or that suitable windows can be inserted near the location of interest. This requirement limits the geometry and design of structures and devices that can be studied. For example, a structure that is commonly used in thermoacoustic Stirling heat engines and pulse tube refrigerators is a "regenerator." This often consists of a stack of fine steel mesh, creating a tortuous path for the gas to travel through. It would be difficult to probe the interior of a regenerator using LDA, PIV, or for that matter HWA.

1.3   MRI and Velocity Phase Encoding

When an atom with nuclear spin is placed in a magnetic field it can undergo transitions between discrete states with different energies, as the orientation of the spin with respect to the field is changed. In doing so, the atom will absorb or emit radiation at a frequency that depends on the strength of the magnetic field. The excitation and subsequent observation of radiation from these nuclei is the basis of Nuclear Magnetic Resonance (NMR). MRI is a form of NMR in which linear magnetic field gradients are applied to a sample, so that the resonant frequency of nuclei becomes a function of position. A measurement of the frequency distribution of the energy radiated by a sample can then be used to determine the

underlying spatial distribution of nuclear spins and their local environment. MRI is largely associated with medical imaging, but a wide range of non-medical applications have also been developed, including the measurement of fluid flow.

Over the last 15 years there has been considerable interest in the use of MRI to image spaces filled with the inert noble gases ^{3}He and ^{129}Xe, as well as hydrocarbon- and fluorine-based molecular gases [10]. Historically, imaging of these types of spaces has been hampered by the significant reduction in signal strength that is encountered when a liquid is replaced with a gas. Effects associated with the more rapid diffusion of atoms in a gas also serve to lower signal strength and degrade image resolution. The disadvantage of lower density can be counteracted through the use of laser optical pumping techniques. These techniques enable one to increase nuclear spin polarizations by several orders of magnitude (a process referred to as "hyperpolarization") and hence to increase signal-to-noise-ratios (SNRs) by a significant margin. Another approach is to use gases made up of molecules that each have many nuclei that can be probed, such as ^{1}H nuclei in propane (C_3H_8) or ^{19}F nuclei in sulfur hexafluoride (SF_6).

One of the simplest MRI-based methods for studying flow involves selective excitation of a well-defined portion of the sample (often referred to as a "slice"), followed by interrogation at a later time that reads back information about the location. This "time of flight" imaging approach has been used previously to study gases [47, 48].

The method explored in this thesis, which is called Velocity Phase Encoded MRI or VPE-MRI, employs additional magnetic field gradients that are rapidly turned on and off ("switched" or "pulsed"). These pulsed field gradients encode information about the velocity of the underlying flow, and so when the nuclear spin response is read out the data can be used to generate velocity vector field images. In effect, the imaging sequence is "sensitized" so that the local sample response is proportional to velocity, to the extent that all of the nuclei in that region are moving in the same direction, at a given instant in time.

VPE-MRI has been used previously to study the flow of hyperpolarized ^{129}Xe [16, 49], hyperpolarized ^{3}He [15], thermally polarized hydrocarbon gases [50], and thermally polarized gases that are rich in fluorine such as SF_6 [18]. It has not previously been used in connection with thermally polarized ^{3}He gas, despite the fact that static images of such gases have been reported [51].

1.4   Experimental Goal and Outline of the Thesis

Channels with widths that range from a few to several times the thickness of the AVBL are commonly found in thermoacoustic devices. Gas motion inside these channels is strongly influenced by viscous interactions with the walls. Thermoacoustic theory can be used to derive analytic expressions for the time and spatial dependence of the gas velocity inside a variety of different channels, including tubes with rectangular, triangular, and cylindrical cross sections.

The goal of the work described in this thesis is to explore the feasibility of using VPE-MRI techniques to probe oscillatory gas motion inside one of these channels; that is, in a well-defined geometry where gas flow is in the AVBL. The particular geometry that will be investigated is a long cylindrical tube. This tube will be referred to as an "acoustic flow tube." The dimensions of this tube, the operating frequency, and even the gas pressure and composition are chosen to facilitate VPE-MRI experiments. The result is a device that operates at frequencies over the range 0.5 to 1.65 Hz and which drives gas motion with peak displacement amplitudes of up to 2.5 cm in a tube with an inner diameter 1.3 cm. While these conditions might seem foreign to someone who is used to working with practical thermoacoustic devices that operate at much higher frequencies (10 to 1000 Hz) and which involve channels with correspondingly smaller transverse dimensions, the principles of similitude can ultimately be brought to bear on this problem. Thus, the apparatus described in this thesis ought to be viewed as an "acoustic wind tunnel." With this in mind, it has been designed so that the cylindrical acoustic flow tube can easily be replaced with other tubes or with other more complex structures. Similarly, even though the particular investigation that is reported here only involves gas displacements, the apparatus has been constructed to permit studies involving superimposed acoustic displacement and pressure/density oscillations.

The remainder of this thesis is arranged as follows. Chapter 2 is devoted to a summary of basic principles and theories underlying the fields of Magnetic Resonance Imaging and Thermoacoustics. The introductions to these fields are far from being exhaustive, but should give readers enough background to understand the rest of the thesis. Chapter 3 describes the design, construction, and instrumentation of the experimental apparatus, including an "Acousto-Mechanical Resonator" that is used to drive gas motion and various aspects of the MRI system. Chapter 4 then discusses the manner in which the various components of

the apparatus were characterized and/or calibrated, prior to the VPE-MRI work. Chapter 5 recounts the manner in which VPE-MRI experiments were implemented, characterized, and ultimately improved as experience was gained over time. Chapter 6 then describes a full set of VPE-MRI data that map gas velocities in the acoustic flow tube through an entire cycle of motion. This chapter also includes brief summaries of similar measurements performed using alternative experimental methods. Finally, an assessment of the VPE-MRI experiments is given in Chapter 7, along with some ideas that might help motivate future work.

Chapter 2

Theoretical Background

The work described in this thesis bridges two fields of study that until now have seen little overlap: magnetic resonance imaging (MRI) and thermoacoustics. In this chapter I will briefly describe the theoretical framework underlying each of these fields. Those interested in reading more about MRI are referred to the excellent monograph by Callaghan [52]. The first half of this chapter deals with MRI (Sec. 2.1). I start with a description of nuclear magnetic resonance (NMR), including nuclear spin dynamics and relaxation. Issues surrounding signal strength and detection are then outlined. Finally, the technique of MRI itself is described, including discussions of slice selection, read gradients, phase encoding, and the concept of velocity phase encoding. This last topic plays a central role in the experiments described in the latter part of this thesis. The second half of this chapter deals with thermoacoustics (Sec. 2.2), in what is known as the small amplitude limit or acoustic approximation. The focus of the discussion is the oscillatory flow of a compressible medium in a cylindrical duct. Again, this topic plays an important role in the latter part of the thesis. For a more in depth discussion of the relevant theory, as well as factors that must be accounted for in the design of the thermoacoustic devices, Swift's book on thermoacoustics is highly recommended [53].

2.1   Magnetic Resonance Imaging

MRI is a subfield of NMR in which gradients in the strength of the longitudinal component of the magnetic field are used for the purpose of determining the spatial extent of various

properties of the imaged object. In an elementary sense, NMR is based on the principle that an atom with non-zero nuclear spin will precess at a well-defined resonant frequency in the presence of a magnetic field. This precession is intrinsically a quantum mechanical phenomenon, but in context of this thesis, the usual semi-classical description of spin dynamics is sufficient. The use of magnetic field gradients and NMR to measure the spatial distribution of nuclear spins was first demonstrated in 1973 by Lauterbur [54] and Mansfield [55]. My description of NMR begins with the concepts of polarization, resonant excitation, and relaxation. This is followed by a discussion of signal strength and signal-to-noise ratio. I then present an overview of MRI including spatial localization, slice selection, echo formation, and k-space mapping. The final topic in this section introduces the technique of velocity-phase-encoded MRI and discusses the role of diffusion.

2.1.1 Nuclear Magnetic Resonance

The nuclei of atoms are composed of protons and neutrons, each of which has an intrinsic spin of 1/2 associated with its angular momentum. Pairs of like nucleons tend to align antiparallel with one another, thus their contributions to the net spin of the nucleus cancel. As a result, atoms with nuclei that have an odd number of protons, an odd number of neutrons, or both, will have a non-zero net nuclear spin I (or angular momentum) of integer or half-integer value. It is these atoms that can be used in NMR and hence MRI. For example, hydrogen-1 (^{1}H) has a single proton and a nuclear spin of 1/2. Helium-3 (^{3}He), on the other hand, has two protons and one neutron. The proton spins cancel leaving one unpaired neutron which gives the nucleus a net spin of 1/2 as well. The magnetic moment of these nuclei is aligned (or anti-aligned) with their angular momentum the ratio of the two quantities is known as the gyromagnetic ratio, denoted by the symbol γ. The gyromagnetic ratio for ^{1}H is $\gamma_H = 26.7519 \times 10^7$ T^{-1}s^{-1} and that for ^{3}He is $\gamma_{He} = -20.38 \times 10^7$ T^{-1}s^{-1} [56]. When placed in a magnetic field of magnitude B_0, the interaction of this dipole moment with the field causes the energy of the nucleus to be shifted by an amount

$$E = -\gamma \hbar B_0 m \tag{2.1}$$

where $\hbar = h/2\pi$ is the reduced Planck's constant and $m = -I, -I+1, \ldots, I-1, I$ is the projection of the nuclear spin along the field direction. For atoms with nuclear spin

$I = 1/2$, there are two energy states: one with the spin aligned with the magnetic field and the other with the spin anti-aligned with the field. The state with the higher energy depends on the sign of the gyromagnetic ratio. Transitions from the lower to the higher state can occur with the absorption of radiation at the resonant angular frequency $\omega = \gamma B_0$, which is known as the Larmor frequency. For ^{1}H in a magnetic field of 1.5 T, which is typical of imaging magnets, the Larmor frequency is approximately 64 MHz, which is in the radio frequency (RF) band of the electromagnetic spectrum. For this reason, the radiation used to manipulate nuclear spins is known simply as "RF" in the NMR and MRI communities.

2.1.2 Polarization

For the most part, NMR experiments are performed on macroscopic samples. The nuclear spin of each atom in the sample will either be aligned or anti-aligned with the magnetic field, and the equilibrium distribution is described by Boltzmann statistics. That is, the ratio of the number of spins aligned anti-parallel to the field N^- and parallel to the field N^+ is

$$\frac{N^-}{N^+} = e^{-\Delta E/k_B T} \tag{2.2}$$

where ΔE is the difference in energy between the two states implied by Eq. (2.1), k_B is the Boltzmann constant, and T is the temperature. The polarization of an ensemble of atoms is defined as

$$P = \frac{N^+ - N^-}{N^+ + N^-} = \tanh(\frac{\Delta E}{2k_B T}) \tag{2.3}$$

which is the ratio of the difference in the two populations divided by the total number of atoms. In thermal equilibrium and for $I = 1/2$ the magnetization is given by

$$\mathbf{M} = \frac{N_s \gamma \hbar}{2} P \, \mathbf{z} \tag{2.4}$$

where N_s is the spin density. The magnetization is a vector quantity that is either aligned or anti-aligned with the background magnetic field depending on the sign of the gyromagnetic ratio. At room temperature the polarization of ^{1}H in a 1.5 T magnetic field is of order 10^{-6}. It is this miniscule population difference that is detected in NMR.

It is exceedingly cumbersome, and in most cases little insight is gained, to work with individual spin dynamics when dealing with macroscopic samples. Instead, a semi-classical

description of the dynamics of ensembles of spins is used. In this approach it is the time dependence and manipulation of the magnetization of **M** that is of interest.

One of the greatest challenges in performing NMR on a gas, compared to a liquid, is the low molecular density. There are many factors that go into determining the total signal strength, but one of the major considerations is the total number of spins. For a fixed sample volume, this depends on the net magnetization of the sample and more specifically the nuclear spin density. In this thesis, experiments involving room temperature ^{3}He at a pressure of 2 atm are described, corresponding to a number density of order 5×10^{25} m^{-3}. As a basis for comparison, the number density of water molecules at room temperature is 3.3×10^{28} m^{-3}, which is nearly 3 orders of magnitude larger. On top of that, there are two hydrogen atoms in each water molecule effectively doubling the signal density. The relatively weak signal available from thermally polarized ^{3}He gas near atmospheric pressure, compared to that obtained in conventional MRI performed on ^{1}H rich tissues and fluids, is a significant hurdle to overcome. Other issues surrounding signal strength and detection are discussed in more detail in Sec. 2.1.5.

2.1.3   Resonant Excitation

The net magnetization vector can be manipulated through the use of alternating magnetic fields applied orthogonal to the background magnetic field at a frequency that is close to the Larmor frequency. In the quantum mechanical description of the state of a nuclear spin, the absorption or emission of radiation induces transitions between energy levels. In the semi-classical description, the alternating magnetic field causes the net magnetization to rotate away from the static field direction.

The net magnetization is the density of magnetic dipole moments in the sample; i.e. $\mathbf{M} = \mu_{\text{net}}/V$ where V is the volume. To begin the mathematical description, we start by equating the torque due to the magnetic field and the rate of change of the angular momentum giving

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} \quad . \tag{2.5}$$

For a static background magnetic field along the z-axis $B_0\mathbf{z}$, the solution for $\mathbf{M}(t)$ involves precession of a constant vector about the magnetic field at the Larmor frequency $\omega_0 = \gamma|\mathbf{B_0}|$. This is analogous to a gyroscope, whose angular momentum vector points along

Figure 2.1: Qualitative trajectory of the magnetization vector during a tipping pulse.

Figure 2.2: Linearly polarized oscillating magnetic field broken into two counter-rotating components.

the axis of rotation, precessing because of the torque caused by the force of gravity acting on the center of mass. The torque in the case of the spin ensemble is due to the applied magnetic field acting on the net magnetization.

If, in addition to the main field $\mathbf{B}_0 = B_0\,\mathbf{z}$, there is a small field B_1 applied in the perpendicular plane and which rotates about $\mathbf{z}$ at angular frequency ω_0, the total magnetic field will be

$$\mathbf{B}(t) = B_1\left(\cos\omega_0 t\,\mathbf{x} - \sin\omega_0 t\,\mathbf{y}\right) + B_0\,\mathbf{z} \quad . \tag{2.6}$$

Equation (2.5) implies that the magnetization will be forced to rotate around the $\mathbf{B}_1$ component of the total field at a frequency $\omega_1 = \gamma B_1$. Since $B_1 \ll B_0$, the precession about $\mathbf{B}_1$ will be much slower than that about $\mathbf{B}_0$, and the magnetization trajectory will trace out a spiral on the surface of a sphere as portrayed qualitatively in Fig. 2.1. Note that the spiral in the figure does not accurately depict the trajectory of the spin which, in a more realistic scenario, would make many more turns around the static field as the direction of $\mathbf{M}$ is canted away from $\mathbf{B}_0$. A short pulse of RF excitation producing a field $\mathbf{B}_1$ for a time t results in the tipping of the sample magnetization away from $\mathbf{B}_0$ field by an angle $\omega_1 t$. As noted earlier, the Larmor frequency is typically in the radio frequency band and therefore these pulses are known as RF tipping pulses. Using radians as a measurement of angle, a pulse that rotates the magnetization from the z-axis into the x-y-plane is called a $\pi/2$ pulse, and one that flips the magnetization is referred to as a π pulse.

In practice, linearly-polarized oscillating magnetic fields are simpler to produce than ro-

tating fields. A linearly-polarized oscillating magnetic field can be expanded into the vector sum of two counter-rotating fields, as shown in Fig. 2.2. Only one of these components rotates in the same sense as the precession of the nuclear magnetization, and thus acts on the spins. In the high field limit where $B_0 \gg B_1$, the counter-rotating component typically has negligible effect.

Greater intuition and ease of calculation can be gained by transforming to a reference frame that is rotating with $\mathbf{B}_1$ at angular frequency ω_1. Using a simple transformation, the effective field in this reference frame can be written

$$\mathbf{B}_{\text{eff}} = (B_0 - \omega_1/\gamma)\,\mathbf{z} + B_1\mathbf{x}' \qquad (2.7)$$

where $\mathbf{x}'$ is the axis in the transverse plane rotating at angular frequency ω_1 along which $\mathbf{B_1}$ is aligned. When $\omega_1 = \omega_0$, the longitudinal component of the field disappears, and the magnetization will precess about $\mathbf{B}_1$. When ω is off resonance the magnetization will precess around $\mathbf{B}_{\text{eff}}$ at some angle between the longitudinal and transverse planes. Off resonance excitation is of importance when analyzing the spectral bandwidth of tipping pulses, and is also useful in manipulating spin systems involving multiple species.

2.1.4 Relaxation

After the application of a $\pi/2$ RF pulse that brings the magnetization into the x-y-plane, Eq. (2.5) implies that $\mathbf{M}$ will simply precess about the z-axis. In the absence of any other interactions, this precession would continue indefinitely. However, inevitable interactions between nuclear spins and the environment, as well as interactions between spins within the ensemble, cause the magnetization to eventually return to its equilibrium state of alignment with the static magnetic field. Although numerous microscopic mechanisms causing nuclear relaxation have been identified, they can be classified in terms of two phenomenological parameters. The first consists of relaxation processes that return the magnetization to its thermal equilibrium state aligned with the background field, and is known as spin-lattice relaxation or longitudinal relaxation. The second consists of processes that lead to the destruction of coherent precession of the magnetization about the z-axis, as the spins come into thermal equilibrium with each other, and is known as spin-spin or transverse relaxation. The characteristic timescales for these relaxation processes are conventionally

denoted T_1 and T_2, respectively. The inverse of these two relaxation times correspond to the phenomenological relaxation rates for the longitudinal and transverse components of **M**.

Longitudinal relaxation involves processes that exchange energy between the spin system and the thermal reservoir of the surrounding environment (the "lattice"). The phenomenological equation that describes the time evolution of the z-component of the magnetization as it returns to thermal equilibrium is

$$\frac{\mathrm{d}M_z}{\mathrm{d}t} = -\frac{M_z - M_0}{T_1} \tag{2.8}$$

with solution

$$M_z(t) = M_z(0)e^{-t/T_1} + M_0(1 - e^{-t/T_1}) \tag{2.9}$$

where T_1 is the longitudinal relaxation time. This time scale can be quite long for pure materials in liquid or gas form and sets a limitation on the rate at which experiments can be performed. For pure deionized water $T_1 > 1$ s is typical, whereas for pure ^{3}He gas it can be much longer, ranging from hours to days with careful preparation of the container [57]. Strongly paramagnetic particles in the fluid or the container walls are particularly effective at enhancing relaxation through collisions. Copper sulfate, which is paramagnetic, is often added to water samples (phantoms) used in evaluating NMR systems specifically to decrease T_1, and thus decrease the amount of time one must wait for thermal equilibrium to be attained. In the experiments described in this thesis, oxygen, which is also paramagnetic, is added to ^{3}He gas for the same purpose. This is discussed further in Sec. 4.4.

As nuclear spins precess in a magnetic field, the coherence of their alignment decreases because they each inevitably experience slight variations in field strength and hence Larmor frequency through interactions with other spins. This leads to the random accumulation of relative phase between spins (dispersion), and hence irreversible degradation of the coherence of the precession. These variations in field strength can be caused by momentary interactions between spins during collisions or by random diffusive motion. The phenomenological equation that describes the decay of the transverse components of the magnetization is

$$\frac{\mathrm{d}M_{x,y}}{\mathrm{d}t} = -\frac{M_{x,y}}{T_2} \tag{2.10}$$

with solution

$$M_{x,y}(t) = M_{x,y}(0)e^{-t/T_2} \quad . \tag{2.11}$$

In this case, T_2 is the spin-spin relaxation time. There is a hard limit that $T_2 \leq T_1$, but typically the spin-spin relaxation time is significantly shorter than the spin-lattice relaxation time. In many situations, the background field will have slight inhomogeneities leading to variations in the Larmor frequency for each spin. This will cause additional dephasing of the spins. The resulting envelope of the detected oscillation remains exponential, but with a shorter decay constant denoted ${T_2}^* < T_2$.

Collectively, when equations (2.5), (2.8) and (2.10) are combined and written in a form appropriate to the rotating frame, they are known as the Bloch equations:

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = \gamma \left(M_y B_0 + M_z B_1 \sin \omega t \right) - \frac{M_x}{T_2} \tag{2.12}$$

$$\frac{\mathrm{d}M_y}{\mathrm{d}t} = \gamma \left(-M_x B_0 + M_z B_1 \cos \omega t \right) - \frac{M_y}{T_2} \tag{2.13}$$

$$\frac{\mathrm{d}M_z}{\mathrm{d}t} = \gamma \left(-M_x B_1 \sin \omega t - M_y B_1 \cos \omega t \right) - \frac{M_z - M_0}{T_1} \quad . \tag{2.14}$$

This set of phenomenological equations describe a great deal of the magnetization dynamics observed in NMR, including the work described in this thesis.

2.1.5 Signal Detection and Noise

It was noted in Sec. 2.1.3 that the sample magnetization could be rotated away from the z-axis through application of an oscillating magnetic field $\mathbf{B}_1$. This field is typically created by an AC current flowing in an RF resonator such as a wire loop, a solenoid, or a birdcage coil [58]. The birdcage coil used for the experiment described in this thesis is discussed in Sec. 3.4. Similarly, the precession of nuclear spins in a magnetic field after a tipping pulse is applied will induce an EMF in the coil. It is this EMF that corresponds to the signal that is detected in conventional NMR and MRI.

Any process that causes a finite magnetization to form in the x-y-plane will subsequently lead to a free induction decay (FID). That is, the Bloch equations (2.12) and (2.13) imply that the transverse components of the magnetization will precess about the z-axis, while the magnitude of both longitudinal and transverse components of $\mathbf{M}$ decay due to spin-lattice and spin-spin relaxation mechanisms. For situations where $T_2 \ll T_1$, the time dependence of the magnetization will then be

$$\mathbf{M}(t) = \left[M_0 \cos \omega_0 t \, \mathbf{x} + M_0 \sin \omega_0 t \, \mathbf{y} \right] \exp(-t/T_2) \quad . \tag{2.15}$$

Figure 2.3: Quadrature detection scheme. The reference oscillator signal is split into two signals, one of which is phase shifted by 90°. The phase sensitive detector (PSD) outputs the difference between the NMR signal and the reference oscillator.

Using complex notation, where the y-component of the magnetization is mapped to the imaginary axis, Eq. (2.15) can be written

$$M_+(t) = M_0 \exp(i\omega_0 t) \exp(-t/T_2) \tag{2.16}$$

where M_+ indicates the transverse component of the magnetization. The EMF that is induced in the coil will thus be an exponentially decaying sinusoid at the Larmor frequency.

The precessing magnetization described by Eq. (2.16) is complex, having real and imaginary components. These components can be separated using the quadrature detection scheme shown in Fig. 2.3. The output of a reference oscillator is set to the Larmor frequency and is divided into two branches, one of which is shifted in phase by 90°. The reference and the phase-shifted signals are mixed with the NMR signal using phase sensitive detectors (PSD). The phase sensitive detector produces the product of the two input signals. This results in a signal that contains terms that oscillate at the sum and difference of the input signal frequencies. The higher frequency "sum" term is removed using a low pass filter leaving the lower frequency "difference" term as the output of the PSD. The output from one of the PSDs is in phase with the reference signal and is labelled the real component. The result from mixing the NMR signal with the phase shifted branch is similar to the real component, but is shifted in phase by 90°. This is the imaginary component. Using this detection method, both the amplitude and the phase of the NMR signal can be extracted.

The size of the EMF induced in the coil depends on a number of factors, including the volume of the sample and the magnitude of the magnetization at thermal equilibrium. Using

Eqs. (2.1), (2.3), and (2.4) and assuming the high temperature limit where $\gamma \hbar B_0 \ll k_B T$, the equilibrium magnetization of a sample of spin-1/2 particles can be written

$$\mathbf{M}_0 = \frac{N_S \gamma^2 \hbar^2 \mathbf{B}_0}{4 k_B T} \quad . \tag{2.17}$$

The signal detected by the coil is related to the field that would be produced by a current flowing through the coil such that

$$S = -\frac{\partial}{\partial t} \int_{V_s} \frac{\mathbf{B}_1(x,y,z)}{i} \cdot \mathbf{M} \, dV \tag{2.18}$$

where the integral is over the volume of the sample [59]. Here, $\mathbf{B}_1(x,y,z)$ is the spatial profile of the field produced by the coil with current i. The magnitude of the signal is then $|S| = \omega_0 K (B_1/i)_{xy} M_0 V_s$, where $(B_1/i)_{xy}$ is the transverse component of $\mathbf{B}_1/i$, and K describes the result of integrating the field profile over the volume of the sample. A good coil design produces large $B_1(x,y,z)$ in the vicinity of the sample for a given current.

The voltage signal received from the precessing spins must compete with the electronic noise of the detection system. To good approximation, the total noise is often white, meaning that the distribution in voltage fluctuations at the spectrometer is nearly Gaussian, and the power spectral density is constant for a given bandwidth over all center frequencies. A measure of the relative strength of the signal compared to the noise is the signal-to-noise-ratio (SNR), which is the magnitude of the signal strength divided by the root-mean-square (RMS) amplitude of the total noise.

There are a number of measures that can be taken to reduce noise in an NMR experiment. The most basic step is the placement of RF shielding around the probe coil or resonator to minimize the influence of external sources of noise. Also low noise amplifiers are used and are placed as close as possible to the coil to boost the signal above any noise that may be picked up along the transmission line.

Once physical noise sources have been minimized, there are further methods that can be employed to improve the SNR. The first is the use of signal averaging: the experiment is repeated N times and the results are added. Since the signal is coherent, the associated EMF is summed (proportional to N). At the same time, the phase of the noise signal is random and the associated EMFs add in quadrature (proportional to $\sqrt{N}$). The net improvement in SNR increases as $\sqrt{N}$. There is a practical limit to how much the SNR can be improved

through signal averaging, since each factor of 2 requires 4 times as many acquisitions. At some point the duration of the experiment becomes prohibitively long and other means of improving the SNR must be explored.

Another method for reducing noise is by the selection of the receiver bandwidth Δf through analog or digital filtering. The RMS voltage from a white noise source is proportional to $\sqrt{\Delta f}$, and therefore, the bandwidth used in signal acquisition must be optimized in order to maximize the SNR. The receiver bandwidth should be large enough to be sensitive to the frequencies of the NMR signal. One component of the apparatus that influences the bandwidth, and hence the noise, is the structure of the RF coil used to generate $\mathbf{B}_1$ fields and detect the signal from precessing nuclear spins. This coil is usually a part of a tuned resonant circuit with a finite bandwidth. The frequency response of a resonator can be characterized by the quality factor Q which, for high Q factors, is defined as $Q = f_0/\Delta f_p$, where f_0 is the resonant frequency, and Δf_p is the full-width-at-half-maximum (FWHM) of the power spectrum. A high Q is desirable in a MRI coil because of the increased response to the small signal from the sample.

Issues associated with SNR play a significant role in the experiments described in this thesis because of the relatively small signal available from thermally polarized ^{3}He gas. At all points in the design of the experiment, care was taken to minimize external sources of noise and to optimize the acquisition parameters.

2.1.6   Spatial Localization

The magnets used in NMR are typically constructed so as to be as homogeneous as possible, so that the entire sample is immersed in a field of uniform strength. This ensures that the frequency at which spins precess is independent of location. The goal of MRI is to map spatial properties of the sample. This is done through the addition of linear gradients to the homogeneous background field, so that the total magnetic field becomes

$$\mathbf{B} = (B_0 + \mathbf{G} \cdot \mathbf{r})\hat{\mathbf{B}}_0 \tag{2.19}$$

where $\mathbf{G}$ is the gradient vector, $\mathbf{r}$ is the position, and $\hat{\mathbf{B}}_0$ is the unit vector aligned with the background magnetic field. There are two configurations of field gradient implied by this equation that are used in MRI: axial gradients, which by convention are directed along

Figure 2.4: Pictorial representation of the (a) longitudinal and (b) transverse magnetic field gradient configurations implied by Eq. (2.19). The arrows point along the z-axis, which is the direction of the background magnetic field. In practice, $\mathbf{G} \cdot \mathbf{r} \ll B_0$.

the z-axis parallel to the main field, and transverse gradients, which are directed along the x- or y-axes perpendicular to the main field. Figure 2.4 illustrates the functional nature of these gradient fields. As a result of Maxwell's equations, the production of a field gradient, as implied by Eq. (2.19), also creates orthogonal gradients terms, known as concomitant gradients. In high field MRI, the contributions of these concomitant gradients to the total field are small and, to a good approximation, may be neglected.

The field described by Eq. (2.19) causes the Larmor frequency within the sample to depend on position:

$$\omega(\mathbf{r}) = \gamma B_0 + \gamma \mathbf{G} \cdot \mathbf{r} \quad . \tag{2.20}$$

The linear dependence of the Larmor frequency on location in the magnetic field gradient forms the basis of MRI. Through the use of pulsed magnetic field gradients, it is possible to create three dimensional images of the sample.

The signal acquired from spins precessing in a magnetic field gradient is simply the superposition of contributions from across the sample. Using exponential notation and ignoring relaxation, the time- and gradient-dependent signal coming from a small volume of the sample may be written

$$dS(\mathbf{G}, t) \propto \rho(\mathbf{r}) dV \exp\left[i\omega(\mathbf{r})t\right] \tag{2.21}$$

where $\rho(\mathbf{r})$ represents the sample density at position $\mathbf{r}$, and $\omega(\mathbf{r}) = \gamma B_0 + \gamma \mathbf{G} \cdot \mathbf{r}$ is the precession angular frequency as a function of position along the field gradient. Neglecting relaxation in this manner is possible when the applied magnetic field gradient dominates the dephasing of spins. Using quadrature detection at the Larmor frequency removes γB_0

from the signal. Under these conditions, the total signal is

$$S(t) = \int \rho(\mathbf{r}) \exp\left[i\gamma\mathbf{G} \cdot \mathbf{r}t\right] \mathrm{d}V \tag{2.22}$$

where the integral is evaluated over the entire volume of the sample and the constant of proportionality has been ignored. This integral represents a Fourier transform (FT),[1] a point that can be made more obvious through the substitution

$$\mathbf{k} = (2\pi)^{-1}\gamma\mathbf{G}t \quad . \tag{2.23}$$

The signal can then be expressed in reciprocal- or k-space notation as

$$S(\mathbf{k}) = \int \rho(\mathbf{r})e^{i2\pi\,\mathbf{k}\cdot\mathbf{r}}\mathrm{d}V \quad . \tag{2.24}$$

Similarly, the inverse Fourier transform of $S(\mathbf{k})$ returns the sample density as a function of position along the gradient:

$$\rho(\mathbf{r}) = \int S(\mathbf{k})e^{-i2\pi\,\mathbf{k}\cdot\mathbf{r}}\mathrm{d}\mathbf{k} \quad . \tag{2.25}$$

Imaging the sample density $\rho(\mathbf{r})$ is reduced to the measurement of $S(\mathbf{k})$ through the application of time dependent linear magnetic field gradients $\mathbf{G}(t)$ during the acquisition. This quantity is then Fourier transformed to reveal $\rho(\mathbf{r})$. The role of magnetic field gradients in the mapping of k-space is described below. The MRI pulse sequence used in this thesis is then described in Sec. 2.1.9.

2.1.7 Slice Selection

Although it is possible to directly acquire three-dimensional data sets spanning much of k-space in order to produce an image of the entire volume, it is frequently more convenient to only excite a thin slice of the sample and generate a pseudo two-dimensional image. This is accomplished through the use of an RF tipping pulse that is transmitted while a magnetic field gradient is applied.

A short pulse of RF radiation at frequency ω will necessarily have a finite spectral bandwidth. The RF pulse can be written as $\psi(t) = A(t)\exp[i(\omega t)]$, where $A(t)$ represents

---

[1]A brief review of Fourier transforms is given in App. A.

Figure 2.5: Sinc RF pulse of Eq. (2.26) and its Fourier transform, which is known as the boxcar or tophat function.

its amplitude or envelope. Through the convolution theorem, the Fourier transform of the product of two functions is equal to the convolution of the Fourier transforms of the individual functions. The frequency spectrum of $\psi(t)$ is then the convolution of the Fourier transform of the envelope function $A(t)$ and the Dirac delta function centered at frequency ω. An RF pulse that is frequently used in MRI is

$$R(t) = R\operatorname{sinc}(at)\cos(\omega t) \tag{2.26}$$

where

$$\operatorname{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{2.27}$$

and thus $R(t)$ is known as a sinc pulse. The frequency spectrum of a sinc pulse has constant amplitude across bandwidth $\Delta\omega = 1/a$ which is centered around ω. This Fourier transform and inverse Fourier transform pair is shown in Fig. 2.5. The constant amplitude and well defined bandwidth of this type of pulse are particularly useful in MRI because of the degree of control it gives in selective excitation.

When an RF pulse is applied at the same time as a linear magnetic field gradient, only the spins whose Larmor frequency lie within the bandwith of the pulse will be excited. For a G_z gradient applied along the z-axis and a sinc pulse of bandwidth Δf, a slice of width $\Delta z = \Delta\omega/\gamma G_z$ is excited. Thinner slice selection is achieved by using longer RF pulses, which have a narrower bandwidth, or by using stronger magnetic field gradients, which cause the spin precession frequency to vary more rapidly as a function of position. For pulses that have an amplitude dependent spectrum, care must be taken to properly account for the dynamics of spins that are only partially tipped away from the z-axis during

excitation.

It should be noted that a $\pi/2$ slice selection RF pulse causes a linear variation in the phase of spins to build up because of the dependence of the Larmor frequency on position along the field gradient. This must be refocused by applying a magnetic field gradient of the opposite polarity and time integral equal to half that of the original gradient pulse [60].

2.1.8 Echoes

The decay of the signal strength during an FID that occurs while a field gradient is applied is due to the dephasing of the spins as they each experience slightly different Larmor precession. Hahn realized that this is an inherently reversible process and demonstrated this fact by generating the first spin-echoes [61]. Figure 2.6 illustrates the formation of a spin echo as viewed in the rotating frame. First a $\pi/2$ RF pulse is applied to tip spins into the transverse plane, producing an FID which decays with time constant T_2^*. This decay is caused by the dephasing of spins via spin-spin interactions and magnetic field inhomogeneities. The distribution of spin orientations across the sample is represented by the shaded arc shown in the figure. At a later time τ, a π RF pulse is applied. This has the effect of inverting the phase of the spins. Immediately after the π RF pulse, the same magnetic field inhomogeneities that caused the spin phases to diverge now refocus the magnetization. This is illustrated in the fifth panel of Fig. 2.6. The signal amplitude then increases as a function of time, peaking at a time 2τ and subsequently decaying back to zero. This peaked signal is a spin echo. The amplitude of the echo will be smaller than the initial signal amplitude of the FID by a factor of $\exp(-2\tau/T_2)$. Recall that T_2 is the time constant for an FID when spin-spin processes are the only ones that contribute to the decay. In the absence of diffusion, the field inhomogeneity effects responsible for making $T_2^* < T_2$ are eliminated.

In Sec. 2.1.6 it was shown that to produce an image of the sample, a map of k-space must be acquired. Movement in k-space is accomplished through the application of magnetic field gradients. The k-vector, defined by Eq. (2.23), depends both on the magnetic field gradient strength and on time. Acquiring a signal during the application of a constant amplitude magnetic field gradient, a process called frequency encoding, will map out a line in k-space. For example, a $\pi/2$ RF pulse followed by a positive x-directed field gradient of constant amplitude will cause k_x to increase linearly at a rate γG_x starting from an

Figure 2.6: Spin-echo sequence in the rotating reference frame. Initially, the magnetization is in thermal equilibrium. A $\pi/2$ RF pulse at time $t = 0$ rotates the magnetization into the transverse plane. As spins precess, they dephase because of spin-spin interactions and field inhomogeneities. A π pulse is applied at time $t = \tau$, flipping the phase of the spins. The reversible processes that previously caused spins to dephase now work to refocus them, forming an echo at $t = 2\tau$.

initial value of zero. The k-vector can be thought of as growing in magnitude along the positive k_x axis. As the gradient is applied, spins dephase according to the local Larmor frequency. Similarly, a negative x-directed gradient will cause the k-vector to increase along the negative k_x axis. The Fourier transform of this line spanning the positive and negative k-axis will produce a projection image of the sample along the direction of the field gradient.

A more efficient method for mapping out a line in k-space spanning the positive and negative axes is the gradient-echo. A gradient-echo is produced by first applying a negative gradient pulse of duration τ (taking the spins to $k_0 = -\gamma G\tau$) followed by a continuous positive gradient such that $k(t) = k_0 + \gamma Gt$. As $k(t)$ approaches 0, the spins will rephase, and an echo will form. The amplitude of the detected signal subsequently decays as the k-vector increases further. The initial gradient pulse in this example is often called a rewind pulse, and the second gradient (applied during data acquisition) is called the read gradient.

An alternative method for mapping the negative and positive k_x axes uses a spin echo. After the $\pi/2$ RF pulse, a positive rewind field gradient is applied, taking the spins to $k_0 = \gamma G_x\tau$. This is followed by a π pulse that inverts the phase of the spins, taking them to $-k_0$. The application of the read field gradient then produces a gradient echo as the k-vector passes through 0. One advantage of the spin-echo over the gradient echo is that the π RF pulse refocuses the spin dephasing caused by background field inhomogeneities. This permits the mapping of plus and minus k-vectors as well as refocusing field inhomogeneities yielding a stronger signal. The gradient-echo is often preferred for rapid imaging sequences or where the absorption of RF power in the subject is an issue.

Figure 2.7: Simple pulse sequence that can be used to acquire a two dimensional map of the signal amplitude in k-space. The Fourier transform will reveal an image of a slice of the sample. One of the possible trajectories through k-space for this pulse sequence is shown in Fig. 2.8.

2.1.9 Mapping k-Space

In order to produce two-dimensional images of a sample, it is necessary to acquire a two-dimensional map of $S(\mathbf{k})$ (which happens in Eq. (2.25)) in k-space. This is done using orthogonal time-dependent magnetic field gradients G_x and G_y to move through k-space (via Eq. (2.23)), while acquiring data. A pulse sequence that accomplishes this task is illustrated in Fig. 2.7. One of the possible trajectories through k-space for this pulse sequence is shown in Fig. 2.8.

The pulse sequence begins with a slice selective $\pi/2$ RF pulse along the z-axis, followed by a refocusing magnetic field gradient pulse with the opposite polarity. A rewind field gradient pulse is then applied along the x-axis, moving k-vector in the positive k_x direction. Next, a short field gradient pulse along the G_y axis, called a phase encoding pulse, is applied. This moves the k-vector in the negative k_y direction, as shown in Fig. 2.8. Immediately afterwards, a π RF pulse inverts the phase of the spins. The net effect is to move the k-vector from (k_x, k_y) to $(-k_x, -k_y)$. The inversion of the k-vector with the π RF pulse also takes advantage of the refocusing of the dephasing, due to background magnetic field inhomogeneities in the spin echo that is produced at $k_x = 0$. Finally, the data is acquired as the read gradient is applied and the k-vector moves in the positive k_x direction. Repeating this sequence multiple times using different amplitudes for the phase encoding

Figure 2.8: One of the possible k-space trajectories for the Fourier imaging pulse sequence shown in Fig. 2.7.

field gradient pulse produces a two-dimensional map of k-space that is Fourier transformed to produce an image of the sample (cf. Eq. (2.25)). This pulse sequence is closely related to one used in the experiments described in this thesis. One should be aware that many other methods for acquiring k-space images exist.

A complete and continuous map of $S(\mathbf{k})$ would yield a perfect image. However, real data sets are discrete and finite in extent. N_x data points are acquired during the frequency encoding phase of the sequence, sampled at regular time intervals $\Delta\tau$ (called the dwell time). This is repeated for N_y phase encoding gradient pulse amplitudes. The extent of k-space that is actually sampled is therefore N_x by N_y points and produces an image data set of the same dimensions.

The area represented by an image is called the field of view (FOV). Using the properties of the Fourier transform, the FOV along the x-axis is

$$\text{FOV}_x = \frac{1}{\Delta\tau\gamma G_x} \tag{2.28}$$

while along the y-axis it is

$$\text{FOV}_y = \frac{N_y}{2\gamma \int G_{y\,\text{max}}\mathrm{d}t} \quad . \tag{2.29}$$

The digital resolution in each dimension can be found by dividing the FOV by the respective number of points. These resolutions define the dimensions of a pixel (picture element) of the image. If the resolution along the z-axis is included, the resulting dimensions define a voxel (a volume picture element). The resolution along the z-axis is usually determined by the thickness established by slice selection. Although less common, it is also possible to acquire full three-dimensional k-space maps, which can then be used to create volume images through a three-dimensional Fourier transform.

The finite area of k-space that is sampled results in the production of truncation artifacts, such as Gibbs ringing, which is manifested by ripples in image intensity near sharp edges. Qualitatively, the k-space map provides a cosine Fourier series representation of the image. For k-vector $\mathbf{k} = (k_x, k_y)$, k_x is the frequency of the cosine function along the x-direction, and similarly, k_y is the frequency along the y-direction. $S(k_x, k_y)$ provides the coefficients for the cosine series. In this mapping, large k-space vectors are responsible for high frequency contributions to the image, whereas small k-vectors contain low frequency information. The origin in k-space (zero frequency component) sets the overall intensity of the image. In truncating the region of k-space that is sampled, some of the high frequency information needed to represent sharp edges is removed. The result is akin to the ripples produced when a square wave is generated from a truncated Fourier series.

The true resolution of an image must be evaluated using the Rayleigh criterion [62], which sets the minimum separation between point spread functions that can be resolved by the eye. The maximum resolution depends on the width of the features that are imaged. This is influenced by a number of factors such as T_2^*, the SNR, the imaging parameters, and diffusion. An MR image can be thought of as the convolution of the FID line shape with the idealized image. The resolution is therefore limited by the FWHM of the FID spectrum, which is given by $\Gamma = (\pi T_2^*)^{-1}$. The optimal resolution of an image is obtained by setting the pixel separation so as to be approximately equal to this width. Diffusion, especially during the read gradient, plays a significant role in limiting image resolution because the signal is attenuated during the acquisition phase. This issue plays a role in experiments described in this thesis and is discussed further in the following section (see Eq. (2.35)).

2.1.10   Velocity Phase Encoding and Diffusion

MRI is a powerful tool for imaging stationary samples. However, it can also be used for quantifying motion. One of the earliest methods used to track fluid motion, called a time-of-flight measurement, involves using slice-selection, or some other selective excitation, to tag a group of spins at one location and then probe them at a later time, so as to infer the displacement [4]. The method used in this thesis and described below is called velocity-phase-encoding (VPE). In this method, pulsed magnetic field gradients are used to store the velocity information in the phase of the image.

The discussion to this point has focused on the use of magnetic field gradients to probe the spatial properties of a sample. These gradients can also be used to directly probe the motion of spins. The effect of magnetic field gradients on the phase of a particle in motion was first presented by Carr and Purcell [63] and later by Hahn [64]. As a spin moves through a region of space in which a magnetic field gradient exists, the Larmor frequency changes as displacements along the gradient direction take place. This is true both for the coherent motion of a spin packet and for individual spins undergoing random diffusive motions. The phase accumulated by a spin (in the rotating frame) moving along a path in a time dependent gradient is given by

$$\phi(t) = \gamma \int_0^t \mathbf{G}(t') \cdot \mathbf{r}(t') \, \mathrm{d}t' \quad . \tag{2.30}$$

If the spin moves at constant velocity, such that $\mathbf{r}(t) = \mathbf{r}_0 + \mathbf{v}t$, then the phase becomes

$$\phi(t) = \gamma \left(\int_0^t \mathbf{G}(t') \cdot \mathbf{r}_0 \, \mathrm{d}t + \int_0^t t' \mathbf{G}(t') \cdot \mathbf{v}(t') \, \mathrm{d}t' \right) \quad . \tag{2.31}$$

The first term represents the total area under the gradient-time graph. The second term represents the first moment of the gradient pulse. We can define a vector from the first moment term

$$\mathbf{p} = \gamma \int_0^t t' \mathbf{G}(t') \, \mathrm{d}t' \tag{2.32}$$

such that it gives the phase shift along the gradient at any position $\Delta\phi(\mathbf{r}) = \mathbf{p} \cdot \mathbf{v}(\mathbf{r})$. Commonly used forms of $\mathbf{G}(t)$ are bipolar with a pulse of trapezoidal, half-sine, or other shape followed by a delay, and then a second pulse with the same shape but opposite polarity. Figure 2.9 shows a bipolar pulse constructed using half-sine shaped lobes. The duration of each lobe is τ and the time between the leading edges of the pulses is Δ. The p-vector for this bipolar pulse is given by

$$\mathbf{p} = \gamma |\mathbf{G}|_{\mathrm{max}} \left(\frac{2\Delta\tau}{\pi} \right) \hat{\mathbf{G}} \tag{2.33}$$

where $\hat{\mathbf{G}}$ is the field gradient direction unit vector. The effect of the bipolar pulse on stationary spins can be understood by looking at Eq. (2.31). The second integral term will disappear because $\mathbf{v}(\mathbf{r}) = 0$, leaving only the first term. After the first lobe of the bipolar pulse, the stationary spins will accumulate a phase of $\phi = \gamma \mathbf{G} \cdot \mathbf{r}\tau$. The second lobe has

Figure 2.9: Bipolar gradient pulse used in velocity phase encoding and diffusion measurements. τ is the length of each half sine lobe and Δ is the time between the leading edges of the pulses.

the same shape but opposite sense, and so it will unwrap the phase acquired during the first lobe. Thus, for stationary spins, the bipolar pulse has no effect. For spins moving at constant velocity along the direction of the field gradient, the first term again goes to zero at the end of the pulse pair, but this time the second term is non-zero. As the spins move, the local field strength changes during the bipolar pulse, and the second lobe no longer undoes the action of the first. As a result, a spin moving at constant velocity will accumulate a phase $\Delta\phi = \mathbf{p} \cdot \mathbf{v}$. Similar to the manner in which the field of view (FOV)was defined, the field of speed (FOS) is defined

$$\text{FOS} = \frac{2\pi}{|p|} \quad .\tag{2.34}$$

Care must be taken to ensure that the maximum velocity expected in the measurement is less than the FOS, so as to avoid phase wrap errors.

The application of a bipolar gradient pulse during an imaging sequence will cause a phase shift that is proportional to velocity, but in order to determine the change in the phase due to motion, a reference is needed. One approach is to acquire two successive images: a reference image acquired with the bipolar pulse turned off and a VPE image with the bipolar pulse turned on. A pixel by pixel phase subtraction of the images yields a map of the velocity in the direction of the gradient.

In the discussion to this point, spins have been assumed to be perfectly stationary, or to be in uniform motion at constant velocity. In a fluid or gas, each particle executes a random walk as it collides with other particles. That is, it undergoes self-diffusion. The displacement of spins has a Gaussian distribution, and the RMS distance travelled by spins in time t will be $\Delta x = \sqrt{2Dt}$, where D is the diffusion coefficient. During the application of a

field gradient, motion will cause a spin to explore regions with different Larmor frequencies resulting in variations in phase. This process is irreversible, and any attempt to undo the accumulated phase of the spin will result in incomplete refocusing. Thus, particle diffusion will attenuate the signal strength [63] during spin- and gradient-echoes sequences. In the presence of a continuous gradient, the signal strength will change with time according to

$$S(t) = S_0 \exp\left(-\frac{1}{3}\gamma^2 G^2 D t^3\right) \quad . \tag{2.35}$$

Stejskal and Tanner demonstrated that the attenuation in signal strength caused by a bipolar field gradient pulse could be used to measure the self-diffusion coefficient D [65]. For the case of isotropic free diffusion during a half-sine bipolar pulse, the attenuation takes the form [66]

$$\ln\left(\frac{S}{S_0}\right) = -D\gamma^2 |\mathbf{G}|_{\mathrm{max}}^2 \tau^2 \left(4\Delta - \tau\right)/\pi^2 \tag{2.36}$$

where D is the diffusion coefficient, and the times Δ and τ are shown in Fig. 2.9 [66].

Diffusion plays an important role in the VPE-MR images described later in this thesis. The irreversible loss of signal during gradient pulses not only degrades SNR, but also limits spatial resolution. One method of decreasing the loss of signal during the application of a field gradient pulse is to make it shorter in time, due to the exponential dependence of the signal strength on $G^2 t^3$ as indicated in Eq. (2.35). This must be balanced against the generation of induced eddy currents, which is described in Sec. 4.8.

2.2 Thermoacoustics

The first calculation of the viscous damping of sound in a tube was reported by Helmholtz, although his results were presented without showing the derivation [67]. Kirkhhoff was the first to present the theory; he also expanded the calculated results to include heat conduction to the walls. A detailed account of this work is given by Rayleigh in his book *Theory of Sound* [24]. Kramers [26] used Kirkhhoff's work in an attempt to explain the spontaneous emission of sound from a capillary tube that is closed at one end as the other end is removed from liquid nitrogen or helium, an effect known as Taconis oscillations. Ultimately, that explanation was deemed unsatisfactory. Nikolaus Rott built on the work of Kramers, and developed a successful and versatile theory that laid the groundwork for the

field of thermoacoustics [27]. In the second half of this chapter, I will present a few aspects of Rott's theory that are relevant to this thesis. In particular I will focus on the oscillatory flow of gas in a cylindrical tube. The experimental work in the later chapters focuses on this system. The notation used here follows that of Swift in his book *Thermoacoustics: A Unifying Perspective for some Engines and Refrigerators* [53]. This book makes extensive use of concepts from thermodynamics, fluid mechanics, and acoustics. For those who are unfamiliar with these topics, I recommend Bejan's *Advanced Engineering Thermodynamics* [68], Landau's *Fluid Mechanics* [69] and Kinsler's *Fundamentals of Acoustics* [70] as excellent sources of complementary information.

2.2.1 Acoustic Approximations

A simple standing-wave thermoacoustic engine is depicted in Fig. 2.10. Also shown is the coordinate system used in this section and a few of the length scales of interest that will be discussed. The device is a half-wavelength acoustic resonator with velocity nodes at the end walls and a pressure node in the center. The wave propagation direction is along the x-axis, and the wavelength is $\lambda = a/f$, where a is the sound speed, and f is the oscillation frequency. The heat exchangers facilitate heat flow into and out of the gas. The thermoacoustic stack is a structure designed to maintain loose thermal contact with the gas. It is the interactions between the oscillating gas and the temperature gradient along the stack between the heat exchangers that gives rise to the work done to the gas. A qualitative description of the operation of this device is given at the end of this section.

We begin by looking at the acoustic aspect of thermoacoustics. In acoustics, it is generally assumed that the time dependent part of all variables is small, and that the amplitude varies as $e^{i\omega_A t}$ about the mean value. A subscript m is used to denote the mean value and a subscript 1 is used to show the complex amplitude of the oscillating component. For example, the gas density ρ and temperature T are written

$$\rho(x, y, z, t) = \rho_m(x) + \Re\left[\rho_1(x, y, z)e^{i\omega_A t}\right], \tag{2.37}$$

$$T(x, y, z, t) = T_m(x) + \Re\left[T_1(x, y, z)e^{i\omega_A t}\right]. \tag{2.38}$$

where $\Re$ indicates the real part of a complex quantity. It is assumed that only the acoustic oscillation can cause gradients perpendicular to the wave propagation direction, and therefore mean values only depend on x. One advantage of this complex notation is that all time

Figure 2.10: Schematic of a simple standing-wave thermoacoustic heat engine.

derivatives can be replaced by $\partial/\partial t \equiv i\omega_A$. In the small amplitude limit, the oscillating part of each parameter is much smaller than its mean value (e.g. $|p_1| \ll p_m$). The acoustic oscillations are assumed to be purely sinusoidal with no mean flow. This implies that the mean pressure has no spatial dependence due to the absence of steady flow. Under these assumptions, the gas displacement ξ and velocity u can be written

$$\xi(x, y, z, t) = \Re\left[\xi_1(x, y, z)e^{i\omega_A t}\right] \quad , \tag{2.39}$$

$$u(x, y, z, t) = \Re\left[u_1(x, y, z)e^{i\omega_A t}\right] \quad . \tag{2.40}$$

The velocity and displacement amplitudes are related by $u_1 = \partial\xi_1/\partial t = i\omega_A \xi_1$. The magnitude of the displacement amplitude is always shorter than the wavelength, but it can often be a significant fraction of the lengths of some of the components used in thermoacoustic devices, such as the stack or heat exchangers described below. The velocity components in the y and z directions are assumed to be much smaller than u_1. Consequently, it is assumed that the gradients in pressure perpendicular to the wave propagation direction are also small, and the pressure can be written

$$p(x, y, z, t) = p_m + \Re\left[p_1(x)e^{i\omega_A t}\right] \quad . \tag{2.41}$$

Fluid properties in thermoacoustics are typically assumed to follow the ideal gas law, and therefore the equations of state for the pressure p and the internal energy per unit mass η are

$$p = \rho RT \tag{2.42}$$

and

$$\eta = \frac{RT}{\gamma_{\text{ideal}} - 1} \tag{2.43}$$

where R is the ideal gas constant divided by the molar mass, and $\gamma_{\text{ideal}} = c_p/c_v$ is the ratio of the isobaric specific heat and the isochoric specific heat. Acoustic waves consist of coupled pressure and displacement oscillations. The ideal gas law implies that variations in pressure give rise to variations in temperature. The speed of sound of an ideal gas is $a = \sqrt{\gamma_{\text{ideal}} RT}$. Other thermodynamic properties of the gas can be derived using Eqs. (2.42) and (2.43).

Viscous, thermal diffusion, and mass diffusion effects are most important in the direction perpendicular to the displacement of the gas, and are characterized by three length scales. The thermal penetration depth is defined as

$$\delta_\kappa = \sqrt{2k/\omega_A \rho c_p} = \sqrt{2\kappa/\omega_A} \tag{2.44}$$

where k is the thermal conductivity of the gas, κ is the thermal diffusivity, and c_p is the isobaric specific heat per unit mass. δ_κ is a measure of the characteristic distance heat can diffuse in a time T_A/π. The temperature of an imaginary gas parcel that is located a distance much greater than δ_κ from the solid boundaries of the device will not be in thermal contact with the walls. For a structure where good thermal contact is desired, such as a heat exchanger, one wants the mean distance between a parcel of gas and the nearest wall to be approximately equal to δ_κ. Thermoacoustic stacks are designed to maintain reasonable thermal contact with the gas: while the gas is in motion little thermal energy is transferred between the gas and the stack but at the turning points of the gas motion, good thermal contact with the walls is established. The viscous penetration depth is

$$\delta_\nu = \sqrt{2\mu/\omega_A \rho} = \sqrt{2\nu/\omega_A} \tag{2.45}$$

where μ and ν are the dynamic and kinematic viscosities, respectively. δ_ν represents the characteristic distance momentum can diffuse in a time of T_A/π. Again, a parcel of gas that is located a distance much greater than δ_ν will move freely without feeling drag from the stationary walls. The acoustic viscous boundary layer (AVBL) thickness is of order δ_ν. Finally, the mass-diffusion penetration depth is

$$\delta_D = \sqrt{2D/\omega_A} \tag{2.46}$$

where D is the mass-diffusion coefficient (cf. Sec. 2.1.10). δ_D represents the characteristic distance an atom or molecule can diffuse in a time T_A/π. Note that the Prandtl number

$$\sigma = \mu c_p/k = (\delta_\nu/\delta_\kappa)^2 \tag{2.47}$$

is a measure of the ratio of viscous to thermal effects. It is of order 2/3 for pure monatomic ideal gases and significantly smaller for some gas mixtures. The Prandtl number is of great importance because viscous effects can lead to significant losses associated with friction between the gas and the walls, reducing the efficiency of a thermoacoustic device. For example, thermoacoustic stacks have pore dimensions of order δ_κ, so for a Prandtl number equal to one, this leads to significant losses. Lowering the Prandtl number will therefore reduce viscous losses. Another relevant dimensionless number is the ratio of thermal to mass-diffusion effects

$$L = k/\rho c_p D = (\delta_\kappa/\delta_D)^2 \quad . \tag{2.48}$$

Like the Prandtl number, L is of order one for ideal gases [71]. Simple kinetic theory for an ideal gas also predicts that σ and L are independent of pressure and temperature [72]. The dimensionless number L is of interest in thermoacoustic mixture separation processes [71], where it plays a role in governing the mass separation flux, and contributes additional losses to the viscous boundary layer because of diffusion.

The importance of the length scales outlined above can be illustrated through a qualitative description of the action of the thermoacoustic heat engine shown in Fig. 2.10. The outer walls form a simple tube that is closed at both ends making this a $\lambda/2$ resonator with velocity nodes at the ends and a pressure node at the center. The thermoacoustic stack is slightly offset from one of the ends and is sandwiched between two heat exchangers. Heat is injected into the system at the hot heat exchanger, which is closest to the end. The other heat exchanger is held at ambient temperature, thereby setting up a temperature gradient along the stack. The stack is made from a material with low thermal conductivity to minimize heat conduction between the heat exchangers. The plate spacing of the stack is designed such that the gas is in good thermal contact with the walls at the turning points, but weak thermal contact while the gas is in motion.

Figure 2.11 shows a cartoon of a parcel of gas undergoing a thermodynamic cycle in the stack of a standing wave thermoacoustic heat engine. It is the temperature gradient along the stack that provides the source of energy and the feedback mechanism needed to drive oscillations in the gas. To begin, imagine that the parcel of gas is in the leftmost position where it is compressed by the acoustic standing wave (recall that the closest end wall is located to the left of this figure) and yet its temperature is lower than that of the adjacent stack material. Heat is transferred from the stack to the gas parcel, causing it to

Figure 2.11: Cartoon showing snapshots of an imaginary parcel of gas undergoing a thermodynamic cycle in the stack of a standing wave thermoacoustic heat engine. The temperature gradient is such that the stack gets progressively hotter as one moves to the left. The closest end wall is located further to the left.

expand. The parcel then moves to the right, away from the end wall, losing thermal contact with the stack in the process. As the parcel approaches the rightmost position the local pressure in the standing wave is reduced, thus allowing it to expand and cool. However, the temperature gradient along the stack is steep enough so that as the gas parcel slows down and comes back into thermal contact with the stack (at the rightmost extreme of its excursion), it finds itself hotter than the local temperature of the wall. As a result, heat is transferred from the gas to the stack, and the parcel volume decreases. At this point the parcel reverses direction and while it is in motion, thermal contact between the gas and the wall is once again reduced until the turning point at the left is reached, and the cycle repeats. Through this cycle the gas undergoes expansion at high pressure and compression at low pressure resulting in net work being done on the gas. All of the gas in the stack undergoes a similar thermodynamic cycle and the net effect is the generation of acoustic power from the flow of heat applied to the hot heat exchanger and rejected by the cold heat exchanger. The displacement amplitude of the gas combined with the temperature changes associated with pressure oscillations sets a lower limit on the minimum temperature gradient ∇T required along the stack, imposed by the hot and cold heat exchangers, to sustain oscillations.

2.2.2 Rott's Thermoacoustics

Prior to Rott, the theory of acoustic oscillations in a tube was limited to dimensions much larger than the AVBL. Kirchhoff and Kramers worked in the limit where the viscous penetration depth δ_ν is much less than the tube radius R. In this limit, viscous and thermal

diffusion effects need not be considered. Rott extended the model of Kramers to situations where $R \sim \delta_\nu$, and discovered that the viscous layer is responsible for many of the thermoacoustic effects that had been observed. Thermoacoustic theory starts with the continuity equation

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{2.49}$$

the momentum equation

$$\rho \left[\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} \right] = -\nabla p + \nabla \cdot \sigma' \tag{2.50}$$

and the general equation of heat transfer

$$\rho T \left(\frac{\partial s}{\partial t} + \mathbf{v} \cdot \nabla s \right) = \nabla \cdot k \nabla T + (\sigma' \cdot \nabla) \cdot \mathbf{v} \tag{2.51}$$

for fluid mechanics [69]. Here s is the entropy per unit mass, $\mathbf{v}$ the velocity of an infinitesimal volume of the fluid, and σ' is the nine-component viscous stress tensor. For most situations, gradients in viscosity can be neglected and —with respect to momentum effects— the flow can be considered as being incompressible. With these approximations, Eq. (2.50) can be simplified to

$$\rho \left[\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} \right] = -\nabla p + \mu \nabla^2 \mathbf{v} \tag{2.52}$$

where μ is the dynamic viscosity. From this starting point, a number of further simplifications are then employed for application to thermoacoustic phenomena.

As mentioned in the previous subsection (Sec. 2.2.1), the time dependent component of all variables is assumed to be small, and therefore products of these small variables can often be neglected. For example, the term $(\mathbf{v} \cdot \nabla)\mathbf{v}$ in the momentum equation (2.52) can be neglected relative to $\partial \mathbf{v}/\partial t$. Further simplifications are possible through noticing that derivatives with respect to x tend to be of order $1/\lambda$ and can be neglected relative to derivatives with respect to y or z which tend to be of order $1/\delta_\nu$ and $1/\delta_\kappa$. Based on these simplifications and substituting $i\omega_A$ for time derivatives, the continuity equation (2.49) can be written

$$i\omega \rho_1 + \frac{\mathrm{d}\rho_m}{\mathrm{d}x} u_1 + \rho_m \nabla \cdot \mathbf{v}_1 = 0 \quad , \tag{2.53}$$

the momentum equation (2.52) becomes

$$i\omega_A \rho_m u_1 = -\frac{\mathrm{d}p_1}{\mathrm{d}x} + \mu \left(\frac{\partial^2 u_1}{\partial y^2} + \frac{\partial^2 u_1}{\partial z^2} \right) \tag{2.54}$$

and the heat-transfer equation (2.51) is

$$i\omega_A \rho_m c_p T_1 + \rho_m c_p \frac{\mathrm{d}T_m}{\mathrm{d}x} u_1 = i\omega_A p_1 + k \left(\frac{\partial^2 T_1}{\partial y^2} + \frac{\partial^2 T_1}{\partial z^2} \right) \quad . \tag{2.55}$$

These are the general equations of thermoacoustics in the low amplitude limit. With his derivations of these equations, Rott was able to quantitatively explain the origin of Taconis oscillations, which was experimentally verified by Yazaki *et al.* [73]. His work was then applied to explain the cooling effect of pulse-tube refrigerators, and inspired the design of new refrigerators [34] and thermoacoustic Stirling heat engines [35, 36, 38, 74].

2.2.3   Ducts

Although Eqs. (2.53), (2.54) and (2.55) are much simpler than Eqs. (2.49), (2.50) and (2.51), greater intuition can be gained through introducing yet further simplifications appropriate to the flow of an inviscid gas in a short segment of tubing or a duct. For tubes of diameter much larger than the thickness of the AVBL, the viscosity of the gas may be neglected. It is useful to define the total volume flow rate

$$U_1(x, t) = \int \int u_1(x, y, z, t) \mathrm{d}y \mathrm{d}z \quad . \tag{2.56}$$

For an inviscid fluid, the velocity is independent of x and y, therefore $U_1 = A u_1$, with A being the area of the duct. Assuming the mean temperature T_m has no x dependence, the relationship between the mean pressure and the mean density under isentropic conditions is $\rho_m a^2 = \gamma p_m$. The variations in pressure and density are then related by $(\mathrm{d}p_1/\mathrm{d}\rho_1)_s = a^2$. For a lossless ideal gas in a duct of constant cross-section, the continuity equation (2.53) can be written

$$p_1 = -\frac{\gamma_{\text{ideal}} p_m}{i\omega_A A \Delta x} \Delta U_1, \tag{2.57}$$

$$= -\frac{1}{i\omega_A C} \Delta U_1 \tag{2.58}$$

(a) Compliance   (b) Inertance

Figure 2.12: Diagrams illustrating the influence of the compliance and inertance of a duct on gas flow. (adapted from Swift [53])

where Δx indicates the length. The compliance C appearing in Eq. (2.58) is defined

$$C = \frac{V}{\gamma_{\text{ideal}}p_m} \qquad (2.59)$$

where the volume $V = A\Delta x$ and $1/\gamma_{\text{ideal}}p_m$ represents the compressibility. The compliance describes the springlike behaviour of the duct (see Fig. 2.12a). A duct with a large volume or compressibility will behave like a softer (more compliant) gas spring.

Using a similar procedure, the momentum equation (2.54) can be rewritten as

$$\Delta p_1 = -i\omega_{\text{A}}\frac{\rho_m\Delta x}{A}U_1 \qquad (2.60)$$

$$= -i\omega_{\text{A}}LU_1 \qquad (2.61)$$

where

$$L = \frac{\rho_m\Delta x}{A} \qquad (2.62)$$

represents the inertance, a quantity that describes the inertial behaviour of the gas in the duct (see Fig. 2.12b). As an acoustic wave propagates through a segment of a duct that acts as a pure compliance, the volume velocity experiences a change in magnitude and phase, represented by ΔU_1 in Fig. 2.12a. For an acoustic wave passing through a segment of a duct that acts as a pure inertance, it is the acoustic pressure amplitude p_1 that is modified.

If we allow a gas to become viscous but restrict the discussion to a duct where δ_ν is much smaller than the hydraulic radius r_h of the tube, it can be shown that the momentum

equation becomes

$$\Delta p_1 = -\frac{i\omega_A \Delta x/A}{1 - (1 - i)\delta_\nu/2r_h}U_1 \tag{2.63}$$

$$= -ZU_1 \tag{2.64}$$

$$= -(i\omega_A L + R_\nu)U_1 \tag{2.65}$$

where

$$R_\nu \approx \frac{\mu\Pi\Delta x}{A^2\delta_\nu} \tag{2.66}$$

represents the viscous resistance, Π is the perimeter of the duct, and $Z = R_\nu + i\omega L$ is the complex flow impedance. An acoustic wave propagating through this duct will experience a pressure amplitude drop much like the voltage drop across a lossy inductor in an AC circuit. For a narrow duct where the inertial mass and therefore L become negligible, the duct will be analogous to a resistor.

Many thermoacoustic devices can be treated as a series of ducts with different diameters. The joining conditions require that all of the gas leaving one section must enter the next, implying that U_1 should be continuous between ducts. The pressure variation p_1 must also be treated similarly because discontinuities in pressure would result in shock waves. Each duct can be characterized by an acoustic impedance in terms of its compliance C, inertance L and viscous resistance R_ν. These elementary thermoacoustic components are directly analogous to AC electric circuit components, allowing conventional methods to be used to determine the frequency response of complex systems. In this regard, the pressure and volume velocity are analogous to AC voltage and current, respectively. A compliance behaves like a capacitor, an inertance like an inductance, and a viscous resistance as a resistor. An acoustic resonator can be modelled and analyzed in the same fashion as an equivalent electronic resonator might be modelled using lumped or distributed circuit elements, including the use of phasor diagrams. This analogy can provide significant intuition, allowing one to manipulate the geometry of a thermoacoustic structure to control the relative phase of the pressure, displacement, and velocity, and thereby move the gas through some desired thermodynamic cycle.

An example of a lumped element AC circuit model of an acoustic resonator is shown in Fig. 2.13. The device on the left is a double Helmholtz resonator, which consists of two spherical bulbs of volume V that act like compliances connected by a narrower duct of area

Figure 2.13: Double Helmholtz resonator and equivalent AC circuit.

A and length Δx that behaves like an inertance. The equivalent AC circuit involving two capacitances C (associated with the compliances) and an inductance L (associated with the inertance) is shown on the right. The combination of capacitances give a total capacitance of $C/2$ and a resonant frequency of

$$\omega = 2\pi f = \sqrt{\frac{2}{LC}} \quad . \tag{2.67}$$

Dimensional analysis of this equation using SI base units for the inertance (kg/m^4) and the compliance (m^{4}s^2/kg) reveals the correct units for ω_{A} (s^{-1}).

The definitions given above for the compliance, inertance, and viscous resistance are often sufficient for modeling the acoustic properties of a device that can be treated as a simple series of ducts. For more complex behaviour that includes heat flow, especially in the presence of a temperature gradient, a full thermoacoustic treatment is required.

2.2.4 Radial Velocity Profile in a Cylindrical Tube

The radial velocity profile of gas in the acoustic boundary layer of a cylindrical tube is central to the work described in this thesis. The viscous interaction of a gas oscillating between widely-spaced parallel plates ($l \gg \delta_\nu$) was first investigated by Rayleigh [24] in the nineteenth century. Arnott's method for calculating the velocity profile for an arbitrary geometry [75] has been used to find solutions for circular and rectangular pores as well as other configurations.

Rott's momentum equation (2.54) is a differential equation for $u_1(y, z)$, with $u_1 = 0$ at the wall. The solution is

$$u_1 = \frac{i}{\omega_{\mathrm{A}}\rho_m} \left[1 - h_\nu(y, z)\right] \frac{dp_1}{dx} \tag{2.68}$$

where h_ν is a complex function that depends on the geometry of the tube. Integrating both sides of Eq. (2.68) over y and z (i.e. the area of the tube) gives the volume velocity U_1. If

the resulting equation is rearranged, one finds

$$\frac{dp_1}{dx} = -\frac{i\omega_{\mathrm{A}}\rho_m}{1-f_\nu}\frac{U_1}{A} \tag{2.69}$$

where f_ν represents the transverse spatial average of $h_\nu(y,z)$. Substituting Eq. (2.69) into Eq. (2.68) yields a useful expression for the gas velocity as a function of position for a duct of arbitrary cross-section

$$u_1(y,z) = \frac{1-h_\nu(y,z)}{1-f_\nu}\frac{U_1}{A} \quad . \tag{2.70}$$

The complex functions $h_\nu(y,z)$ and f_ν have been calculated for a cylindrical geometry as well as parallel plate, rectangular and pin array configurations [53]. Measurements related to the acoustic flow in the parallel plate [76] and pin array [77] configurations have been performed. Of interest here is the result obtained for a uniform circular pore of radius R, for which

$$h_\nu = \frac{J_0[(i-1)r/\delta_\nu]}{J_0[(i-1)R/\delta_\nu]} \tag{2.71}$$

and

$$f_\nu = \frac{2J_1[(i-1)R/\delta_\nu]}{J_0[(i-1)R/\delta_\nu](i-1)R/\delta_\nu} \tag{2.72}$$

where $r = \sqrt{y^2 + z^2}$ and the $J_\alpha(x)$ are Bessel functions of the first kind.

Figure 2.14 shows the real and imaginary parts of u_1A/U_1 from Eq. (2.70) plotted at $t = 0$ for various values of the ratio R/δ_ν. For $R/\delta_\nu \lesssim 2$, all of the gas is in the AVBL. For larger values of R/δ_ν, the gas near the center of the tube is able to move freely. In the limit that the tube radius is much larger than the viscous penetration depth, the viscous interaction with the wall becomes negligible, and the gas velocity is independent of radial position (plug flow). In the experiments described in this thesis, $R/\delta_\nu \approx 4$. In this situation, the motion of the gas near the walls is strongly influenced by viscous effects, while near the center of the tube it is much less so. Using VPE-MRI techniques (Sec. 2.1.10), the velocity profile inside a cylindrical tube is measured. The results yield maps analogous to Fig. 2.14 that are then directly compared to Eqs. (2.70), (2.71) and (2.72).

2.2.5 DELTAEC

DELTAEC is an anacronym for Design Environment for Low-amplitude ThermoAcoustic Energy Conversion [78]. It is a software package that is used for evaluating the performance

Figure 2.14: Theoretical oscillatory velocity profile for gas in a cylinder of radius R. The real and imaginary parts of $u_1 A/U_1$ are shown for several values of the ratio R/δ_ν.

of thermoacoustic devices that is freely available online from the Los Alamos National Laboratory Thermoacoustics Group. DELTAEC numerically integrates in one spatial dimension the equations of Rott's thermoacoustics for a series of segments, such as the ducts, compliances, heat exchangers and thermoacoustic stacks, described above. The software also includes segments that model many other useful functions, such as transducers for conversion between acoustic and electric power and regenerators, which are integral to Stirling engine and pulse-tube refrigerator designs.

Modeling using DELTAEC begins with the user defining the geometrical aspects of the problem as a series of interconnected segments. The gas conditions (mean pressure, mean temperature, frequency, initial volume flow rate, etc.) are also defined or allowed to vary as free parameters as necessary, in order to arrive at a solution. The coupled equations (2.53) and (2.54) for the pressure amplitude p_1 and the volume flow rate U_1 are sequentially integrated for each segment subject to the requirements that p_1, U_1 and T_m are continuous from one segment to the next. Some of the variables are set be varied in order to arrive at a solution that satisfies the equations and the boundary conditions. Whenever heat flow is relevant, Eq. (2.55) is solved simultaneously with (2.53) and (2.54). DELTAEC also provides many capabilities that are useful in the design process, such as finding a series of

solutions as a particular parameter is varied. Any of the outputs of the model can then be plotted against this varied parameter. Empirical approximations for some non-linear effects, such as turbulence and other loss mechanisms, can also be included in models.

The apparatus I built to drive acoustic oscillations in a gas confined to a cylindrical pore is presented in the next chapter (Sec. 3.2). The DELTAEC model that was used to simulate the frequency response of this device is discussed in Sec. 4.5, and a sample output file from the simulation is presented in App. B.

2.2.6   Reynolds, Dean and Womersley numbers

Rott developed his theory for thermoacoustics in the small amplitude limit because he was most interested in the onset conditions for Taconis oscillations. Perhaps surprisingly, the theory works for a broad variety of practical applications. However, as designers push to create high power thermoacoustic devices in compact packages, the small amplitude limit starts breaking down, and non-linear effects begin to manifest themselves. The onset of non-linear effects such as turbulence, mass streaming, and jetting in an oscillating fluid, have been the subject of much recent research.

In fluid mechanics, several dimensionless numbers are conventionally and conveniently used to differentiate between flow regimes, as well as to determine the similtude of flows in different geometries and in fluids with different physical properties. The most well-known of these is the Reynolds number (Re), which provides a measure of the ratio of inertial to viscous forces during steady flow. It is used to delineate the transition from steady laminar to turbulent flow. The definitions of several dimensionless numbers including the Reynolds number are given in Table 2.1; note that V and D represent the characteristic velocity of the fluid and diameter of the pore or channel. For most situations the transition from laminar to turbulent flow occurs near a Reynolds number of about 2000. Other factors such as surface roughness can alter the onset point.

In the experiments described later in this thesis, fluid is forced to flow through channels that change direction, that is, the flow path is bent. At high flow rates there can be a significant pressure drop across these bends. The Dean number (De) is conventionally used for classifying steady flow in a bent tube [79, 80]. It is the product of the Reynolds number, and the square root of the curvature ratio $D/2R$, where R is the mean radius of the bend.

Table 2.1: Dimensionless numbers related to viscous flow in a pipe.

number	symbol	definition	application
Reynolds	Re	$\dfrac{\rho V D}{\mu}$	steady straight flow
Dean	De	$\dfrac{\rho V D}{\mu}\left(\dfrac{D}{2R}\right)^{1/2}$	steady flow curved tubes
Womersley	α	$\dfrac{D}{2}\left(\dfrac{\omega\rho}{\mu}\right)^{1/2}$	pulsatile flow straight tubes

The relevant geometry is shown in Fig. 2.15. For low Dean numbers (De $\ll 1$), where the flow velocity is small or the radius of the bend is large, flow patterns are largely unchanged from those encountered in straight tubes. As the flow rate increases, or as the bend is made tighter, inertial forces begin to disturb the flow and the velocity of the fluid increases toward the outer periphery of the bend. Figure 2.16 shows examples of primary (axial) and secondary (transverse) flow patterns in a bent cylindrical tube calculated for De=96 and De=606. As De increases, the peak velocity of the axial flow moves toward the outside of the bend and the secondary flow rate increases. The onset of turbulence is estimated to take place near De=5000, resulting in significant losses manifested by a pressure drop across the elbow. An excellent review of steady and oscillatory flow in curved pipes has been presented by Berger *et al.* [81].

The fluid motion encountered in the experiments described in this thesis are purely oscillatory with no steady flow component. The Reynolds and Dean numbers therefore have limited applicability, and yet another dimensionless number is required. The Womersley number α is conventionally used to classify pulsatile flow in a straight tube (oscillatory flow superimposed on a steady flow, although the limiting case of pure oscillatory flow is encompassed by the definition). It was originally developed in the context of trying to

Figure 2.15: Dimensions relevant to calculating the Dean number. R is the mean radius of the bend and D is the diameter of the tube.

understand blood flow in the body. The Womersley number is the ratio of the oscillatory inertia force to the shear force, and is used to estimate the effect of oscillation frequency and peak volume velocity on flow characteristics. For $\alpha < 1$, Poiseuille flow is observed at the peak flow velocity, and for $\alpha > 10$, plug-like flow is established.

Figure 2.16: Calculated flow patterns for low and intermediate Dean numbers in a cylindrical tube (adapted from McConalogue and Srivastava [82]). The inner radius of the bend is to the left (denoted I) and the outer radius to the right (denoted O).

The nature of unsteady flow in a bent tube has been studied [83, 84, 85], but little progress has been made in generalizing the results. The Reynolds, Dean, and Womersley numbers, along with the curvature ratio, all play a role in determining the flow regime, and the variety of flow patterns that can be produced is incredibly complex. Progress has been made toward understanding a number of limiting cases, but empirical methods and numerical simulations are critical for the general case.

Chapter 3

Apparatus

The goal of the work presented in this thesis is to measure the motion of gas undergoing acoustic oscillations in the viscous boundary layer, using the techniques of MRI. In this chapter, I describe the design and construction of the apparatus built to perform these experiments. In Sec. 3.2, I describe an acousto-mechanical resonator that is used to drive a column of gas in simple harmonic motion. I also describe some of the instrumentation used to monitor and characterize the resulting flow. In Sec. 3.3, I present the design of the gas handling system used to prepare and manipulate gas mixtures. Finally, in Sec. 3.4, I describe a birdcage coil resonator that was built to enable resonant excitation and detection of nuclear spin precession, and thereby enable MRI experiments to be performed on the oscillating gas.

3.1 Design Motivation

Before presenting the design of the apparatus, some "back of the envelope" estimates are summarized here to motivate choices that were made regarding an appropriate oscillation frequency and amplitude, including an estimate of the viscous boundary layer thickness. These qualitative descriptions will be expanded later in the thesis.

Typical MR imaging acquisition times span the range from milliseconds to seconds and achieve spatial resolutions of order millimeters. The period of the gas motion must be long enough that the imaging sequence can resolve the velocity at several discrete phases of oscillation. Based on these rough numbers, it was decided to choose a nominal operating

47

frequency of 1 Hz with a displacement amplitude of order 2 cm. Pure ^{3}He gas samples can have very long longitudinal nuclear relaxation times [57], ranging from hours to days in length, depending on conditions in the cell. Since T_1 is the time scale for nuclear magnetization to return to thermal equilibrium after a perturbation, the use of pure ^{3}He could result in a situation where one would have to wait an inordinate amount of time between acquisitions. Fortunately, T_1 can be reduced through the addition of a paramagnetic gas such as oxygen [86, 87], thereby decreasing the time needed for the ^{3}He magnetization to relax to thermal equilibrium. In anticipation of the extensive signal averaging required for these imaging experiments (because of the weak signal available from thermally polarized ^{3}He) it was felt that T_1 should be reduced to a value comparable to the period of the acoustic oscillations. This would allow for approximately one acquisition per acoustic cycle. An O_2 partial pressure of around 1 atm is sufficient to bring the ^{3}He nuclear relaxation time T_1 close to 2 s [86]. Having previously performed NMR experiments on cells containing ^{3}He at pressures of around 2 atm with some success [87], it was decided that a similar ^{3}He partial pressure should be used in this experiment as well. Finally, in order for most of the gas to be in the acoustic boundary layer, the viscous penetration depth δ_ν cannot be too much smaller than the radius of the tube. For a mixture of 1 atm of O_2 and 2 atm of ^{3}He, a kinetic theory calculation the viscosity (as outlined in App. C) along with Eq. (2.45) leads to $\delta_\nu \sim 0.2$ cm. This in turn implies that a tube radius R in the range $0.4 - 1.2$ cm (i.e. chosen such that $R/\delta_\nu \approx 2 - 6$) would place a reasonable fraction of the gas in the viscous boundary layer.

3.2   Acousto-mechanical Resonator

The acoustic wavelength in air at 1 Hz is greater than 300 m, making a purely acoustic approach to the design of an oscillator for this experiment impractical. Instead, an acousto-mechanical resonator, shown in Fig. 3.1, was designed. It uses two coupled mercury-filled U-tube resonators as large travel pistons to move the gas in harmonic motion. The system is driven by an electric motor, coupled to a locomotive-style piston that compresses a berylium-copper bellows, which is in turn attached to a compliance (referred to as the piston compliance). The displacement amplitude of the bellows is adjustable, providing one means for controlling the amplitude of the gas oscillations. A 1.5 m tube connects the

Figure 3.1: Schematic diagram of the acousto-mechanical resonator, drawn roughly to scale. The piston assembly is described in Sec. 3.2.1 (see Figs. 3.3 - 3.4), the compliances are described in Sec. 3.2.2 (see Fig. 3.9), the U-tube assembly in Secs. 3.2.3 through 3.2.5 (see Figs. 3.5 - 3.7), and the pressure transducers in Sec. 3.2.6 (see Figs. 3.9 and 3.10). The gas handling system that connects to the air lines and experiment fill lines is presented in Sec. 3.3 (see Figs. 3.11-3.13).

piston/compliance assembly to one of the U-tube resonators. This length of tube serves to distance the motor and other magnetic components of the drive system from the MRI magnet. The role of the compliance is to act as a soft spring, reducing the force exerted on the mercury in the U-tube. The second U-tube resonator is coupled to the first through the gas in a 1.5 m length of tube passing through the MRI magnet. This "acoustic flow tube" has a low compliance and creates a stronger coupling between the two liquid-filled (i.e. U-tube) resonators than that which exists between the drive piston and the closest U-tube resonator. Following the second U-tube resonator, another 1.5 m length of tube leads to a second compliance (referred to as the end compliance) that is identical to the first. This ensures that the entire acousto-mechanical resonator is symmetric about its midpoint, inside the MRI magnet.

Coupled identical harmonic oscillators have two normal modes. In the low frequency mode, the masses move in unison. In the high frequency mode, the masses move in antiphase, such that when one mass is at its positive maximum excursion, the other mass is at its negative maximum. The frequency of the lower mode is equal to the resonant frequency of the isolated oscillators. The frequency of the higher mode relative to that of the lower mode is determined by the strength of the coupling between the resonators. For the acousto-mechanical resonator (AMR) described above, it is the motion of the gas in the acoustic flow tube that is of interest. In the high frequency mode, the pressure variations in the gas are largest, and the displacement is small. In the low frequency mode, it is the gas displacement oscillations that are maximized. In order to study gas motion in the AVBL, the operating point of the AMR was chosen to be near the frequency of the low mode. The frequency response and operating point of the AMR is discussed further in Sec. 4.5.

The tubing used in the construction of the apparatus is primarily $15.9 \, \text{mm OD} \times 9.5 \, \text{mm}$ ID poly(methyl methacrylate) (PMMA) tube. When the acoustic flow tube was first filled with a $^3\text{He-O}_2$ mixture, the pressure in the system was observed to drop at an alarming rate. It was surmised that the PMMA is too permeable to helium gas to be used in this application. Both the horizontal and vertical runs of tubing in the experimental region were replaced with appropriately formed lengths of 15 mm OD $\times$ 13 mm ID soda lime glass. Despite the smaller wall thickness, the leak rate slowed considerably. The details of the leak rate measurement and the improved performance of the soda lime glass are given in Sec. 4.1. Note that the nominal length of the glass acoustic flow tube (from one free
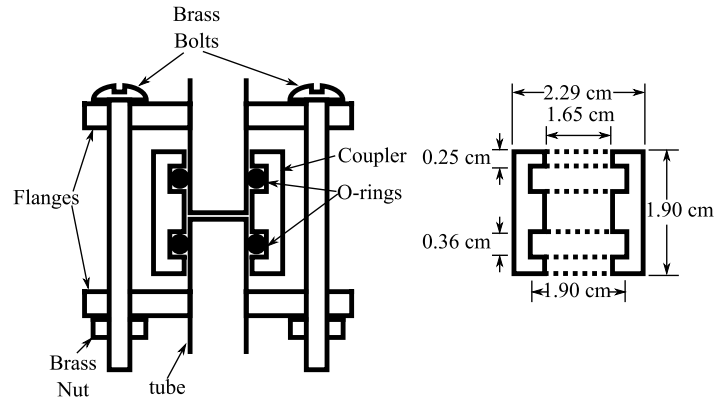
Figure 3.2: Schematic of the coupler with the flanges and bolts in place and a dimensional drawing of the static radial seal coupler. The Buna-N O-rings are compressed between the coupler body and the walls of the glass or plastic tubing. The flanges are epoxied to the tubing and the brass nuts and bolts hold the tubes together against the force of the internal pressure of the gas. The flanges and coupler are made from PMMA plastic.

mercury surface to the other) is 1.48(2) m. The straight run of tube between the two bends is 1.12(1) m long and the bends have a mean radius of curvature of 0.9(2) cm. The acoustic flow tube pressure transducer is located 2.5 cm from the bend (see Fig. 3.1).

The individual sections of the AMR are modular and are connected using the simple static radial seal couplers shown in Fig. 3.2. To hold these joints together against the 3 atm internal pressure, each coupler is spanned by three 1/4-20 brass bolts fastened through external flanges that are glued to the tubes. In the case of the PMMA tube, the flanges are glued in place using Weld-On # 16 Plexiglass Glue Adhesive. In the case of the glass tube, Hysol Tra-Bond F112 epoxy adhesive was used to attach the flanges. This design allows for parts to be easily exchanged in order to perform experiments on different geometries and flow conditions. Prior to filling the system with mercury, it was pressure tested to 5 atm using air.

In the following several subsections, I will describe the components of the AMR in greater detail, moving from right to left in Fig. 3.1. I start with the piston assembly in Sec. 3.2.1, then the compliances in Sec. 3.2.2, and finally the U-tube resonators in Sec. 3.2.3. Next, I describe the system used to adjust the amount of mercury in the U-tubes in Sec. 3.2.4. I then describe the transducers used to measure the mercury displacement

in Sec. 3.2.5 and the pressure amplitude in Sec. 3.2.6. Finally, I describe the gas handling system used to fill the experiment with the gas mixture, as well as to remove the gas and separate the O_2 from ^{3}He in Sec. 3.3.

3.2.1 Piston Assembly

The side and top views of the drive mechanism for the AMR are depicted in Fig. 3.3. A photograph is shown in Fig. 3.4. It is a locomotive-style piston in which the rotational motion of the variable speed DC electric motor is used to produce sinusoidal linear motion. The entire assembly is mounted on a 14 cm × 51 cm × 1.3 cm brass plate fastened to the top of a wooden box that aligns the drive system with the axis of the MRI magnet. The DC electric motor drives a brass cam through a cog-and-chain gearing system. The chain loops around a large drive gear and a smaller driven gear. The latter sits below the brass plate and is attached to the 6.4 mm steel shaft of the cam. The shaft passes through the mounting plate to the brass cam via a cylindrical ball bearing assembly. The connecting rod is attached to the cam through a cylindrical plane bearing, and leads to a universal joint connected to a hardened steel piston rod. This rod connects though a linear bearing to a flange soldered to a berrylium-copper bellows that is, in turn, soldered to the brass end-plate of the piston compliance. The displacement amplitude of the bellows can be adjusted by changing the radial distance from the center of the cam shaft to the rod end connection. The bellows has a mean radius of 2.12 cm, and for the work presented in this thesis, the displacement amplitude of the piston was 0.55 cm. Piston frequencies up to 1.65 Hz can be generated with this mechanism in its current configuration.

The outer diameter of the bellows is 5 cm. With a pressure difference of 2 atm across the face, there is static force of approximately 400 N for the motor to work against. At the point of greatest compression of the bellows, the backlash in the electric motor is released causing the drive system to jerk. To ameliorate this issue, a system of springs was implemented to compensate for the back reaction force exerted on the piston by the gas. Four 28.5 cm springs with spring constants of 13 N/cm each are compressed between 1.27 cm diameter steel bolts mounted near the cam and a flange that moves with the bellows. The bolt mounting plate is threaded to enable easy adjustment of the bolts, and is rigidly fastened to the brass mounting plate of the piston assembly. The bolts are used to adjust the spring
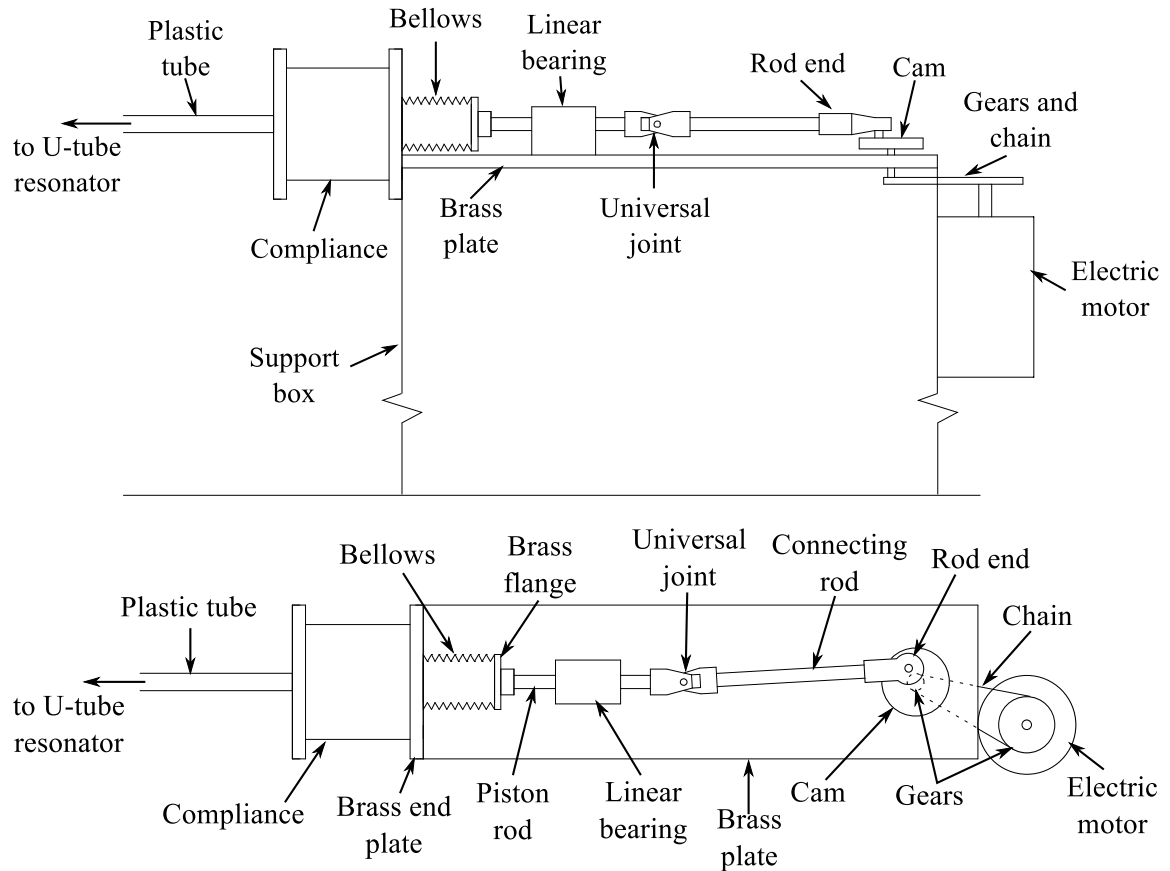
Figure 3.3: Side and top views of the drive piston assembly. The springs used in compensating for force exerted by the internal pressure of the bellows, have been omitted.

compression until the net force on the bellows is zero at the midpoint of the piston travel. The end of each spring is trapped in a hole bored into the face of each bolt. Also mounted in the face of each bolt and running along the axis of each spring is a 0.64 cm diameter aluminum rod that ultimately passes through a clearance hole in the bellows flange. As the piston operates the bellows flange slides on the four aluminum spring center rods but the spring ends move with the flange. The motor must still work against the force arising from variations in spring length, but this has been minimized by making the spring length long compared to the amplitude of the motion. Implementation of the spring loaded counterbalance system was an important component of the system design, as it greatly reduced the higher order harmonic content of the drive.

Figure 3.4: Photograph of the piston assembly highlighting elements of the spring-loaded pressure compensation system. Only one of the four springs is in place so that the center rods can also be seen.

3.2.2   Compliances

The two compliances, located at either end of the AMR, act like soft gas springs that isolate motion of the coupled mercury columns from the environment. A direct consequence is that the coupling between the drive piston and the U-tube resonators is weak (i.e. they are coupled via the piston compliance). The body of each compliance is made from a 9.6 cm length of 9.93 cm ID copper tube sandwiched between two square end plates with 12.7 cm sides. Face seal O-ring grooves are machined into the end plates and 6 symmetrically spaced bolts compress the assembly to contain the internal pressure.

   The construction of the piston compliance and the end compliance differ slightly. One face of the piston compliance is a 1.27 cm thick brass plate with a 4.24 cm diameter hole leading to the bellows. The plate and bellows are soldered together and coupled via this hole. The other face of the piston compliance is made from a 1.27 thick PMMA plate. It is glued using Weld-On # 16 Plexiglass Glue Adhesive to the tube leading to the U-tube resonators. Also attached to the plastic plate is a pressure transducer that is described further in Sec. 3.2.6. Both faces of the end compliance are made from 1.27 cm thick PMMA. The side leading from the U-tube resonators is, again, glued to the PMMA tube. The other end plate is rigged with a pressure transducer mount. A photograph of the end compliance is shown in Fig. 3.9.

In order for the AMR to operate over a wide range of internal pressures, the effective spring constant of the two compliances must be small compared to that of the restoring force caused by gravity acting on the mercury in the U-tubes. The following section, Sec. 3.2.3, describes the mechanical resonance of the U-tubes. As the compliances become stiffer, the restoring force they exert on the mercury increases, raising the resonant frequency. The effective diameter of the compliances is much larger than δ_ν, and so they can functionally be approximated as ducts with compliances $C = 1.75 \times 10^{-9}$ m^3 Pa^{-1}, using Eq. (2.59). Multiplying the acoustic pressure amplitude p_1 (Eq. (2.58)) by the internal area of the PMMA tube $A_t = 7.54 \times 10^{-5}$ m^2 and relating the volume flow rate U_1 to the acoustic displacement amplitude ξ_1 (i.e. $U_1 = i\omega u_1 = i\omega A\xi_1$), the magnitude of the force on the mercury caused by the gas spring can be written

$$|F_1| = \frac{A_t^2}{C}\xi_1 = K\xi_1 \tag{3.1}$$

where K is the effective spring constant of the compliance, which evaluates to 3.2 N/m at a total fill pressure of 3 atm. This represents approximately 10% of the mean spring constant acting on the mercury (see Sec. 4.5), and shifts the resonant frequency upward by a small amount.

3.2.3 U-tube Resonators

A U-tube, as the name implies, is a U-shaped tube with long vertical sections. A volume of liquid moving in a U-tube of constant cross-section will undergo simple harmonic motion with a resonant frequency

$$f = 2\pi\sqrt{2g/l} \tag{3.2}$$

where g is the gravitational constant and l is the linear length of the fluid. A photograph of one of the U-tube resonators used in this work is shown in Fig. 3.5.

The two U-tubes were constructed from 15.9 mm OD $\times$ 9.5 mm ID PMMA tube. The section with the $180°$ bend has an inner diameter of 4.1 cm and was formed by bending the tube around a semi circular jig as it was gently warmed with a heat gun. Care was taken not to over-heat the plastic, which results in bubbling and discoloration. Slow, even heating produced the highest quality bends of constant radius of curvature, but some collapse of the circular cross-section was unavoidable. The effect of the change in cross-sectional area of

Figure 3.5: Photograph of one of the U-tubes after having been removed from the AMR. The mercury is visible in the vertical plastic section on the left. The vertical glass section on the right is wrapped in copper foil for the height measurement system. Couplers and flanges can be seen at the bottom ends of the vertical sections. The mercury height adjustment piston is mounted to the ring stand post behind the U-tube. Nearly all of the mercury has been displaced out of the piston and into the U-tube.

the tube in the vicinity of these bends on the resonant frequency of the U-tube is discussed in Sec. 4.5.

Each $180°$ bend is extended by two 18 cm-long vertical tubes. The vertical sections closest to the acoustic flow tube are made from 15.0 mm OD $\times$ 13.0 mm ID soda lime glass and are wrapped in copper foil. This foil facilitates the mercury height measurement system discussed in Sec. 3.2.5. The other vertical sections are made from PMMA tubing. Individual sections are connected using the cylindrical couplers and bolted flange assembly shown in Fig. 3.2. Each U-tube is attached to an aluminum ring stand using custom non-magnetic clamps. The mercury level adjustment system, described in Sec. 3.2.4, connects to the bottom center of the bends via a 0.64 cm diameter stainless steel tube that is glued in place using Emerson and Cuming Stycast 1266 epoxy.

Mercury was chosen as the working fluid because it is very dense, has a low vapor pressure, is a good conductor and —most importantly— neither the vapour nor the liquid increase the nuclear relaxation of ^{3}He [88]. The high density of the fluid means that the restoring force due to gravity and the large mass will dominate the resonant behaviour of the U-tubes. As was shown in Sec. 3.2.2, the effective spring constant associated with the air in the compliances only has a small effect on the resonant frequency. The low vapour pressure of Hg (0.0012 Torr at $20°$ C [56]) ensures long term stability in the amount of fluid over time, and prevents the gas mixture from becoming polluted. The electrical conductivity of the mercury is exploited for measuring the displacement of the liquid, as will be described in Sec. 3.2.5. Precautions were taken to minimize the potential hazard should the mercury be released from the apparatus. This included building a spill containment system and the installment of a mercury spill kit. Aside from its well known toxicity, mercury has some unique physical properties that make it difficult to work with; most notably, it forms amalgams with most metals. Exceptions include iron, steel, and a few other metals.

3.2.4   Mercury Level Adjustment

Being able to change the amount of mercury in the U-tubes allows for the manipulation of the resonant behaviour of the coupled oscillator system. It also provides a natural means for removing the liquid from the U-tubes altogether in order to exchange parts. Two approaches were taken in the design of the mercury level adjustment system, both of which involved

displacing the liquid from a reservoir.

The first approach involved a large brass valve in which the body acted as the reservoir and the bellows could then be used to displace the fluid. Just prior to filling with mercury, it was realized that the lead-based solder that sealed the bellows to the body of the valve would dissolve in the mercury quite rapidly, and that the copper bellows would also slowly be consumed. Consequently, this particular system was never put into service.

The final design of the mercury level adjustment system is outlined in Fig. 3.6. The cylindrical body is built from the same PMMA tubing material as used for the U-tubes (9.5 mm ID). The plunger is machined from PMMA and uses two series O-rings to seal against the cylinder walls. It is threaded into a brass ball joint that decouples the plunger from the adjustment screw. The screw is made from a length of 1/4-20 brass threaded rod machined to mate with the socket of the ball joint, so that it can raise or lower the plunger as desired. Affixed to the top of the piston is a female threaded brass insert that guides the adjustment screw. A wing nut tightened against another nut provides leverage to turn the screw. Epoxied to the bottom of the piston is a length of flexible 0.64 cm diameter PVC tubing that leads to the U-tube. A short length of stainless steel tube is epoxied to the bottom of the U-tube using Emerson and Cuming Stycast 1266 epoxy. The connection to the PVC tube is made through a stainless steel union.

After the reservoir has been charged with mercury, it can be attached to the U-tube. This is most easily done with the reservoir in a horizontal orientation, and held physically lower than the bottom of the steel attachment tube. Once the seal is made, the mercury can be injected into the U-tube by advancing the plunger with the threaded rod. The piston can then be returned to a vertical orientation and raised to allow the mercury to enter the U-tube. This process can be reversed to remove the liquid from the U-tube, or to detach the reservoir altogether. The linear length of mercury in the U-tube can be estimated by measuring the travel of the plunger as the fluid enters the U-tube. This information can then be used to estimate the resonant frequency of the liquid (cf. Eq. (3.2)).

3.2.5   Mercury Displacement Measurement

The mercury in the U-tubes acts as a liquid piston that is coupled to the gas in the acoustic flow tube. Knowing the position and velocity of the liquid piston as a function of time

Figure 3.6: Cross-sectional view of the mercury height adjustment system.

provides information about the motion of the gas that can be used for comparison with VPE-MR images. In the process of developing a suitable method for measuring the mercury height as a function of time, a number of techniques were explored. The first method involved two thin capacitor plates placed inside the vertical section of a U-tube and electrically insulated from the mercury by a thin layer of epoxy. As the mercury filled the space between the plates, the capacitance would increase. Ultimately, it proved to be too difficult to adequately insulate the mercury from the plates and to reliably feed electrical leads through the walls. The second trial technique used a length of fine stainless steel wire running along the inside of a vertical tube. The electrical resistance of the mercury column was much smaller than that of the steel wire, and forms a parallel high conductance path with the wire. As the mercury level rises, the total resistance measured between the ends of the wire is reduced. This system showed promise at first, but the mercury oxide layer that forms on the free surface of the liquid would also tend to coat the wire, forming an electrically insulating layer. At this point, the readings would become intermittent and

unreliable.

The final design of the system used for displacement measurements is again a capacitive method, as shown in Fig. 3.7. One plate of the capacitor is formed by wrapping copper foil around the vertical glass section of the U-tube. The copper foil is held in place using electrical tape. The mercury itself acts as the second electrode of the capacitor. Electrical contact to the mercury is made through the stainless steel tube that leads from the bottom of the U-tube to the liquid height adjustment system described in Sec. 3.2.4. The theoretical capacitance per unit length of this coaxial capacitor is

$$\frac{C}{L} = \frac{2\pi k_g \varepsilon_0}{\ln(b/a)} \tag{3.3}$$

where k_g is the dielectric constant of the glass (6.32 for soda lime glass [56]), ε_0 the permittivity of free space, and a and b are the inner and outer radii of the glass tube, respectively. The mercury height is then proportional to the capacitance. In practice, the height of the mercury column is determined by measuring the current flow through the capacitor in response to a 0.500 V_{RMS}, 1.000 kHz AC voltage using a Stanford Research Systems SR830 DSP lock-in amplifier. The calibration and evaluation of this system is described in Sec. 4.3.

3.2.6 Pressure Transducers

The gas dynamics in the central acoustic flow tube involve coupled velocity and pressure oscillations. The velocity is directly measured using VPE-MRI, and can be inferred from the mercury displacement. A measurement of the acoustic pressure amplitude provides complementary information to the mercury displacements. The pressure amplitude p_1 is much smaller than the mean pressure p_m making a total pressure measurement insensitive to the small variations. A differential pressure transducer can be used to measure the variations relative to the mean pressure. A special mount was designed that incorporates an acoustic low pass filter to hold the reference port of a two port differential pressure transducer at the mean pressure thereby measuring only p_1.

The acoustic pressure amplitude p_1 is measured at three points along the apparatus. These points are indicated as PT in Fig. 3.1. The transducers are Nova Sensor NPC-1210-005D-3L two-port differential piezoresistive pressure sensors with a nominal range of 5 psi (34000 Pa). A photograph of one of these transducers and a custom-designed mounting

Figure 3.7: Mercury height measurement system. The plastic sections are made from PMMA tube. The vertical section on the left (closest to the MRI magnet) is made from 15 mm OD × 1 mm wall thickness soda lime glass and is wrapped in copper foil. The mercury height is determined by monitoring the capacitance between the copper foil and the liquid.

interface attached to the end compliance is shown in Fig. 3.9. The pressure transducer has two ports, both of which couple to the mounting interface using O-ring static radial seals. The p_1 measurement port hole passes through into the AMR volume. The mean pressure port ends in a small volume. A 0.80 mm diameter hole is drilled perpendicular into each port mount hole and an Upchurch Scientific NanoPort$^{\text{TM}}$ assembly is epoxied to the face of the mount. These are used to connect the ports through a 37 cm length of 0.13 mm ID capillary tube. The capillary tube acts like an acoustic resistance, and the volume around the mean pressure port acts like a compliance forming an acoustic low pass filter. The cutoff frequency of the filter was measured to be 0.08 Hz. At the operating frequency of 1.31 Hz used in the VPE-MRI experiments the pressure amplitude at the reference port is 6% of p_1.

The pressure sensors make use of a resistive bridge network in which two parallel legs of the bridge involve piezoresistive elements. Changes in pressure induce strain on these elements, thereby causing their resistance to change. A constant current source and instru-
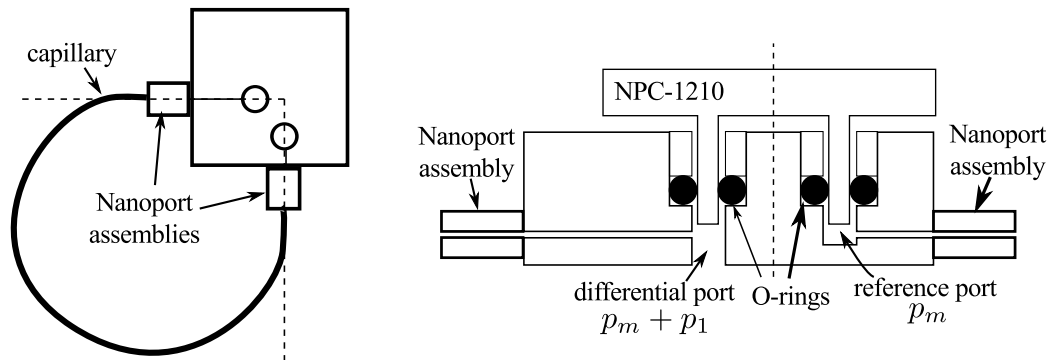
Figure 3.8: (left) Top view of the Nova Sensor NPC-1210 two-port differential pressure transducer mount. The dotted lines indicate the views for the cutaway drawing (right). The reference port is connected to the differential port via capillary sealed to the body using Nanoport assemblies. The capillary combined with the reference port volume form an acoustic low pass filter. The tubes of the pressure transducer join to the mount using static radial seals with Buna-N O-rings. A photograph of a mount and transducer is shown in Fig. 3.9

Figure 3.9: Photograph showing a pressure transducer interface mount on the end compliance. The pressure transducer is near the bottom of the interface mount.

mentation amplifier were constructed for each sensor to monitor changes in sensor resistance and hence pressure. The circuit diagram is shown in Fig. 3.10. The calibration of these pressure transducers is discussed in Sec. 4.2.
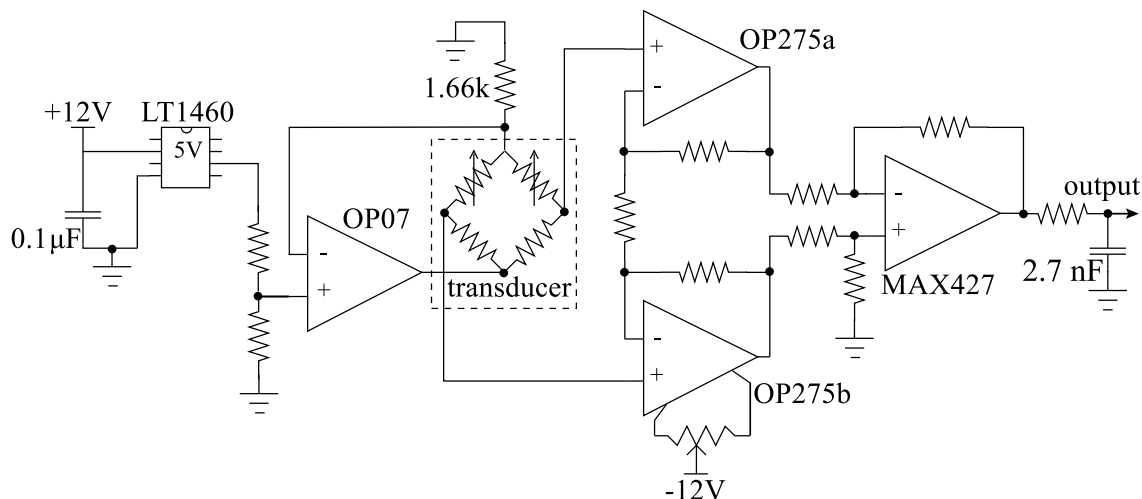
Figure 3.10: Schematic diagram of the pressure transducer circuit. All unmarked resistors are precision wire-wrapped with a nominal value of $100\,\mathrm{k\Omega}\pm1\%$. The opamps have a supply voltage of ±12 V. The transducer is functionally represented by the resistance bridge inside the dashed box, with the variable resistors representing the piezoresistive elements. The LT1460 5 V reference combined with the left-most opamp and the $1.66\,\mathrm{k\Omega}$ resistor form a 1.5 mA constant current source. The three opamps on the right make up an instrumentation amplifier with a signal gain of 3 and a common mode gain of 1. The output is single-ended and can be read by any voltmeter. The low-pass RC filter at the output has a cutoff frequency of ~600 Hz and is intended to remove high frequency noise.

3.3   Gas Handling System

The AMR is filled with two different gas mixtures (see Fig. 3.1). The acoustic flow tube, located between the two U-tubes, is filled with a nominal mixture of 2 atm ^{3}He and 1 atm O_2. The compliances, which bookend the two U-tubes, are filled with 3 atm of dry air. In order to manipulate the composition of the gas mixture in the acoustic flow tube, balance

the overall pressure distribution (to keep the mercury slugs in place), and to add or remove gas from the system, the gas handling system (GHS) shown in Fig. 3.11 was built.
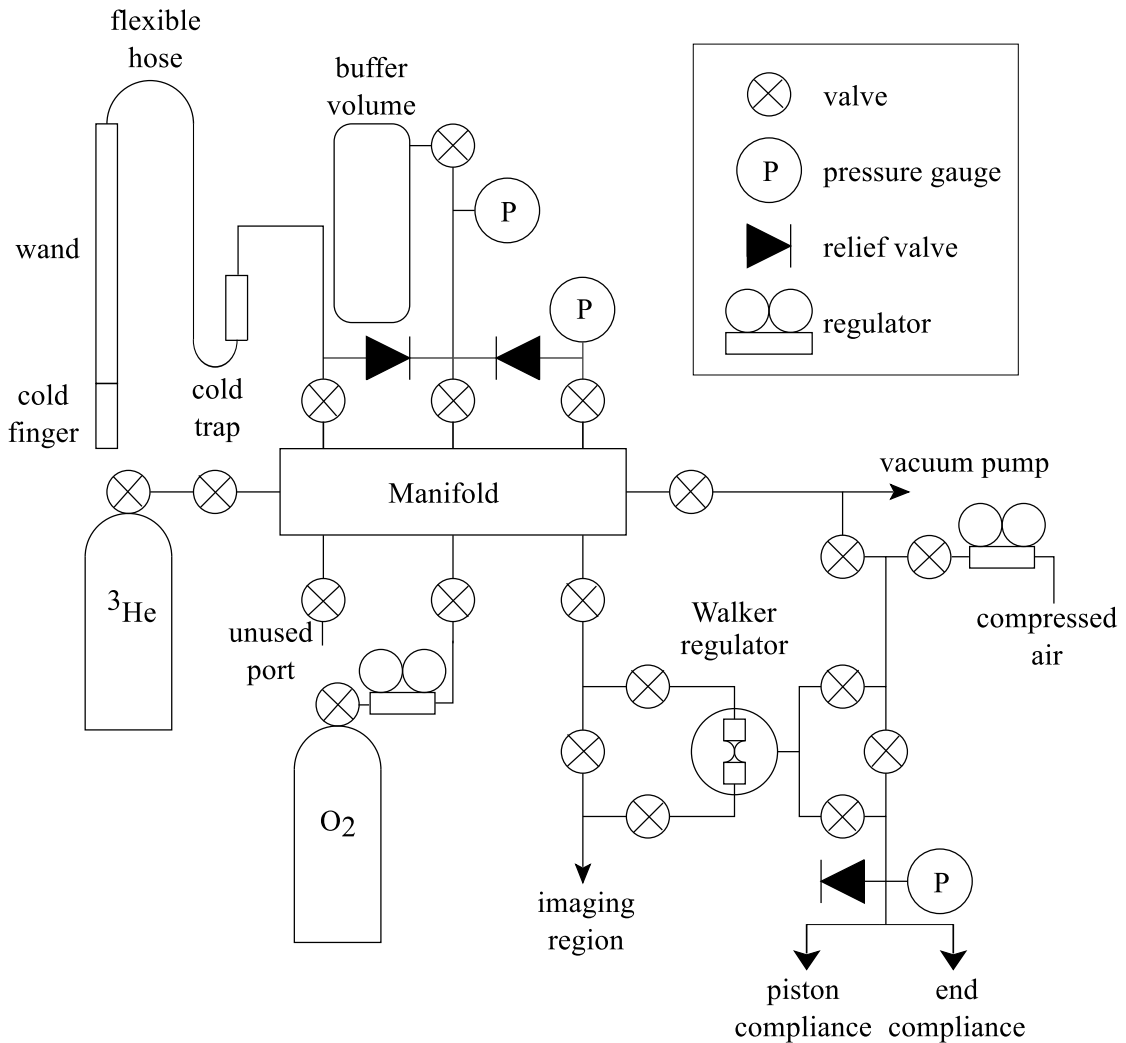
Figure 3.11: Schematic diagram of the gas handling system.

The GHS is composed of two separate sections. One is for working with the ^{3}He and O_2 mixture in acoustic flow tube, and the other is for adjusting the air pressure in the compliances. The gas mixture system is built around a stainless steel manifold with eight valved ports. The various elements, connected to the manifold, include gas bottles for ^{3}He and O_2, a cold finger for cryopumping, a pressure gauge, a pumping line, and fill line leading to the acoustic flow tube. Additionally, provision is made to access a large buffer

volume that is normally coupled to the rest of the system by pressure relief valves set to activate at a pressure difference of 5 atm. These valves protect against accidental over-pressure in the case of rapid boil-off from the cryopump, or a failure of the regulator from the high pressure O_2 bottle. The 5 litre buffer volume captures any gas that escapes, so that the ^{3}He can later be recovered. The compressed air section of the GHS involves a small manifold with connections to a pressure regulated air supply line and a pumping line. The mixture and air sections of the GHS use a common rotary vacuum pump for evacuating the gas. The two sides of the GHS are also indirectly linked, connecting the mixture system and the air system by a Walker regulator, which is a pressure controlled value.

Each of these components and the operating procedure for the GHS are outlined below. The manifold and pressure gauge used to prepare the gas mixture are discussed in Sec. 3.3.1. The cryopumping system is presented in Sec. 3.3.2. The compressed air system and Walker regulator are described in Sec. 3.3.3. Finally, the systems used for filling the AMR with a gas mixture and for recovering the ^{3}He are outlined in Sec. 3.3.4.

3.3.1 Gas Manifold

The manifold provides a central connection between the gas bottles, the cryopump, the rotary vacuum pump, the acoustic flow tube, and a pressure gauge. There is a sixth connection that ends in a KF 16 flange for miscellaneous connections or venting to atmosphere. The pressure gauge is an Ashcroft general purpose digital vacuum-100 psi gauge (model 25D1005PS02LV&100#-BL) with 0.25% full scale accuracy (13 Torr). A second Ashcroft pressure gauge is connected to the compliance fill lines, as shown in Fig. 3.11. The manifold was constructed from a 7.62 cm × 2.54 cm × 2.54 cm block of stainless steel. A 2.4 mm diameter hole was drilled lengthwise through the body and six equally spaced thru-holes were drilled along the width. A 0.637 cm OD stainless steel tube was welded in place at each opening, and connected to a Swagelok® model SS-4h stainless steel valve.

3.3.2 Cryopump

When a substance, such as activated charcoal or Zeolite, is cooled to cryogenic tempera-tures, gas molecules will adsorb to its surface. This process, which is called cryopumping, can be used to evacuate a volume. The gas can later be released by warming the substrate

(or "sorb"). The reversibility of this pumping process is exploited in the GHS to move ^{3}He gas between volumes.

The cryopumping system consists of a cold finger at the end of a wand that can be inserted into a Dewar containing liquid helium, and a U-shaped cold trap that is located between the wand and the manifold. The wand is made from 1.27 cm OD thin-walled stainless steel tube. This material has a low thermal conductivity which helps prevent the transport of heat from room temperature down into the Dewar. An 8.5 cm length of 1.27 cm OD copper tube is brazed to the end of the stainless steel section forming the body of the cold finger. A small packet made from fine copper mesh and containing 3.0 g of aquarium-grade activated charcoal was placed inside the copper tube. Copper has a much higher thermal conductivity than stainless steel, and provides good thermal contact between the charcoal sorb and the liquid helium. A 100 cm length of 1.27 cm mean diameter flexible stainless steel bellows connects the top of the wand to the cold trap. The other end of the bellows is mounted 200 cm above the ground, providing sufficient maneuverability for the insertion of the wand and cold finger into helium Dewars of various sizes. A 76 cm, 0.64 cm OD stainless steel tube connects the flexible bellows to the cold trap.

The pumping rate of the cryopump and the desorption rate as the sorb warms can be controlled by the depth to which the wand and cold finger are inserted into a liquid helium Dewar. This is best done with a Dewar that is relatively empty, having a liquid depth of around 10 cm. As the cold finger is slowly lowered into the Dewar, vapour begins to cool the sorb. The temperature of the sorb decreases as it is inserted further, and maximum cooling is achieved with the cold finger immersed in liquid. The process can be reversed to control the rate that gas desorbs from the charcoal surface.

The function of the cold trap is to remove impurities (especially O_2) from the ^{3}He gas as it is transferred to and from the experimental volume and/or the storage bottle via cryopumping. During the transfer process, gas passes through the cold trap twice, which increases its effectiveness at filtering out the oxygen. A drawing of the U-shaped cold trap is shown in Fig. 3.12. It is constructed of 0.64 cm OD stainless steel with the straight sections being made from thin-walled tube. One of the vertical arms contains an 8.5 cm length of 1.27 cm OD copper tube, similar to that employed in the cold finger. Inside this segment is fine copper mesh packet filled with 3.2 g of Zeolite 5A. Zeolite 5A is a ceramic sorb material with a characteristic pore size of 5 Å that is safe to use with high concentrations of

Figure 3.12: Sketch of the cold trap. All of the stainless steel tube, including the thin-walled sections, have a 0.64 cm OD. The copper section holds a copper mesh packet containing 3.2 g of Zeolite 5A. The U-shaped design facilitates insertion into a Dewar flask filled with liquid nitrogen.

O_2. The amount of sorb used in the cold trap is sufficient for trapping most of the O_2 in the acoustic flow tube as the mixture is removed from the experiment and the ^{3}He is returned to the storage bottle.

The cold finger is cooled to 4 K by inserting it into a liquid helium Dewar and the cold trap is cooled to 77 K by immersing it in a liquid nitrogen Dewar. The binding energy of ^{3}He to the Zeolite surface is relatively small at 77 K therefore very little of it adsorbs to the Zeolite surface. On the other hand, O_2 has a much higher binding energy to the Zeolite and at 77 K of most the gas adsorbs to the surfaces. The binding energy for ^{3}He to the activated charcoal is large enough at 4 K that the gas can be evacuated from the AMR to pressures less than 0.1 Torr. To remove residual water and other impurities adsorbed to the charcoal and Zeolite, both the cold finger and the cold trap were heated to 150° C using heating tape while being evacuated by the rotary pump.

3.3.3 Compressed Air System and Walker Regulator

The two compliances, the piston bellows, and the PMMA tubes connecting the compliances to the U-tubes and connected volumes are filled with dry compressed air (see Fig. 3.1). The compressed air is supplied from a centralized source. The pressure of the air is regulated to approximately 3 atm absolute to mitigate over-pressure hazards and is measured

using another Ashcroft model 25D1005PS02LV&100#-BL pressure gauge. The air manifold is connected to the same rotary vacuum pump that services the mixture manifold (see Fig. 3.11).

A Walker regulator is a pressure dependent valve and can be used to maintain zero pressure difference between two volumes as they are evacuated. The purpose for including this device in the GHS was to keep the ^{3}He/O$_2$ mixture in the acoustic flow tube and the air in the compliance at the same pressure as the gas is removed from the AMR. A schematic of the Walker regulator design is shown in Fig. 3.13. The device consists of a 9 cm diameter $\times$ 6 cm high cylindrical chamber that is connected to the compliant volumes at either end of the acoustic flow tube. Passing through the center of this chamber is the pumping line connected to the acoustic flow tube of the AMR. In the middle of the pumping line are two cylindrical seals that are spanned by a flexible latex finger cot. To operate the device, the cryopump is used to evacuate the acoustic flow tube through the finger cot, and the rotary pump is used to evacuate the compliances. If the air pressure in the compliances is greater than that of the gas mixture in the acoustic flow tube, the finger cot will collapse and the cryopump flow rate will be reduced or stopped. Conversely, if the air pressure in the compliances is less than that of the gas mixture in the acoustic flow tube, the finger cot will inflate and the flow rate to the cryopump will be maximized. In order to monitor the status of the finger cot, the front face of the cylindrical air chamber is made from transparent PMMA. As long as the maximum rate of change of pressure in the compliances is less than that in the acoustic flow tube, the Walker regulator will maintain a small pressure difference between the two volumes. In this way, only one valve (controlling the air pumping rate) needs to be adjusted during the AMR evacuation process.

Unfortunately, the Walker regulator did not perform as well as expected. The volume of the acoustic flow tube is much smaller than that of the compliances. The pressure in the acoustic flow tube therefore changes more rapidly with the addition or removal of gas. When pumping the acoustic flow tube through the Walker regulator the latex valve did not sufficiently restrict the flow. Consequently, the pressure change in the acoustic flow tube was too rapid to adjust the air pumping rate accordingly to maintain the displacement of the mercury. Another issue with the Walker regulator is that the latex from which the finger cot is made becomes brittle and cracks during the span of months between usages. The finger cot therefore requires replacement before each usage. It is simpler to bypass the Walker

Figure 3.13: Schematic of the Walker regulator used in the GHS.

regulator and manually control the pumping of the gas by alternating between removing mixture from the acoustic flow tube tube and air from the compliances.

3.3.4   AMR Filling and ³He Recovery Procedures

To charge the apparatus with 3 atm of gas, care must be taken to maintain a small pressure difference across the mercury in the U-tubes, so that it is not displaced beyond a maximum excursion of 9 cm. Exceeding this level will result in gas being forced from one chamber to the other (compliance to acoustic flow tube or vice versa) through the bottom of the U-tube.

The ³He bottle (Fig. 3.11) has an internal volume of approximately 2500 cm³. When it was first opened, the pressure read 657 Torr. The combined volume of the manifold, acoustic flow tube, fill line, and the room temperature cold finger is about 500 cm³. In order to achieve the desired pressure of 2 atm (1520 Torr) in the acoustic flow tube, the cold finger was first used to extract an appropriate amount of ³He from the storage bottle. This gas was then expanded to the acoustic flow tube. Impurities are removed from the gas during this procedure as it passes through the cold trap. To charge the apparatus with gas, the entire system is first pumped out and then the cold trap temperature is lowered to 77 K by filling the Dewar flask with liquid nitrogen. The wand and cold finger are then lowered into a liquid helium Dewar, and a quantity of gas is cryo-pumped from the ³He bottle. Once the storage bottle and manifold pressure are less than 1 Torr, the valve to the ³He bottle is closed, and the cold finger is slowly withdrawn to release the gas into the manifold. ³He is admitted into the acoustic flow tube until the mercury columns reach their upper-

most travel limit ($\sim \pm 7$ cm) in the visible vertical sections of the U-tubes. Air is then administered to the compliances until the mercury columns approach their lower limits. This procedure is repeated until the pressure in the AMR reaches an absolute pressure of 1520 Torr. The acoustic flow tube is then valved off from the GHS, and the cold finger is once again lowered into the liquid helium Dewar. This extracts the remaining ^{3}He from the manifold and filling tube. With the cold finger still in the Dewar, the valve to the cryopumping system is closed. Next, the regulator on the O_2 bottle is set to just over 2280 Torr in preparation for completing the mixture. A similar metering procedure is then employed to add O_2 from its storage bottle to the acoustic flow tube with air being added to the compliances as necessary until the total pressure reaches approximately 2280 Torr. The mercury resonators are then leveled by ensuring that the ^{3}He-O_2 gas mixture and air pressures are equal, at which point the entire AMR is valved off from the GHS using valves attached to the compliances and to the acoustic flow tube as shown in Fig. 3.1.

The GHS is also used to remove the gas from the AMR and separate the O_2 from the ^{3}He. After filling the acoustic flow tube with the gas mixture, the manifold and fill line contain primarily O_2. This is discarded by evacuating the manifold and fill line with the rotary vacuum pump. The cold trap is then cooled to 77 K by filling the Dewar with liquid nitrogen and the cryopumping wand is inserted into a liquid helium Dewar until it is immersed in liquid. The valve to the acoustic flow tube is then opened, admitting the gas mixture into the manifold. The process used in filling the AMR is then reversed. Air is pumped out of the compliances until the mercury level is displaced by approximately 7 cm. The valve to the cryopump is then opened slightly, to control the pumping rate, until the mercury is displaced to the opposite extent of its travel. This is repeated until the valves to the rotary pump and cryopump can be opened to the system fully, without worry of a large pressure difference across the mercury. Once the pressure is below 1 Torr, the acoustic flow tube fill line is valved off from the manifold. At this point, the gas mixture has passed through the cold trap into the cold finger, which collects most of the O_2. The ^{3}He is then transferred back into the storage bottle. The cold finger is slowly warmed by raising the wand out of the liquid helium Dewar, and the ^{3}He makes a second pass through the cold trap as it enters the manifold and storage bottle volumes. Once the cold finger has warmed to room temperature, the valve to the storage bottle is closed. The cold trap is then warmed, and the manifold and cold finger are pumped out with the rotary pump, discarding

the filtered O_2 and the leftover ^{3}He.

3.4 Birdcage coil

An integral piece of equipment required for MRI is the RF coil used to apply B_1 tipping pulses and to detect the precessing nuclear magnetization. A perfectly homogeneous transverse B_1 field can be produced in an infinitely long conducting cylinder by an axial surface current of amplitude proportional to $\sin\theta$, where θ is the azimuthal angle. The birdcage coil approximates this ideal surface current distribution with a finite number of evenly-distributed axial current rungs, in which the phase of the current varies as $\sin\theta$. Thus, at any instant of time, an approximation to the $\sin\theta$ axial current density distribution is produced. The birdcage coil was first reported by Hayes as an improvement over the saddle coil [58]. A schematic diagram and lumped-element equivalent circuit for a "low-pass" birdcage coil design is shown in Fig. 3.14. This coil involves two conducting azimuthal end rings that are joined by capacitors along N equally spaced axial rungs. The inductance of the end rings and rungs, combined with capacitance of the capacitors, forms a transmission line with $N/2 + 2$ modes. The lowest frequency mode produces the most homogeneous B_1 field and is doubly degenerate with counter-rotating circularly polarized fields. Quadrature drive and reception will break the degeneracy. This reduces the power required to tip the spins by a factor of 2 and increasing the SNR by $\sqrt{2}$. The birdcage coil used in this work has only a single coupling coil, thus producing a linearly-polarized B_1 field normal to the area enclosed by the coupling coil.

The 48 MHz birdcage coil used in this work was built around a former consisting of a 22 cm length of 42 mm OD diameter $\times$ 5 mm wall thickness PVC tube. Top view and end view photographs of the birdcage coil are shown in Figs. 3.15 and 3.16. A freely distributed Java program, called Birdcage Builder [89], was used to estimate the dimensions of the end rings and rungs, as well as the capacitance needed to achieve the desired resonant frequency. The completed coil has 3 cm-wide end rings and six 21.0 cm-long $\times$ 5.2 mm-wide axial rungs, all fabricated from 0.15 mm-thick copper sheet. The capacitors are placed at one end of the resonator and are formed by inserting a 70 μm sheet of Teflon between the end ring and a 10.5 mm-wide by 3 cm long copper pad at the end of each rung. Pressure is then applied to each capacitor via PVC caps that are fastened to the PVC former using 6-32 nylon

Figure 3.14: Schematic diagram of a low-pass birdcage coil (left) showing the geometric configuration of conductors and capacitors. The diagram on the right indicates the lumped element equivalent circuit. Points A and B connect to points C and D, respectively.

6/6 molded bolts. The Birdcage Builder software is limited to performing calculations for a set of predetermined frequencies. The closest value to the desired frequency was 42.56 MHz. The nominal valve of each capacitor for this frequency is 56 pF. Coupling to the birdcage coil is accomplished through a single rectangular loop. The loop is attached to the former using two 1/4-20 nylon 6/6 molded bolts. The spacing between the coupling loop and the birdcage coil is varied to optimize the impedence match between the coil and the transmission line. Another single turn loop is fastened to the opposite side of the birdcage coil. The spacing between this coil and the birdcage is used to tune the frequency of the resonance with a range of approximately 200 kHz.

The birdcage coil was tuned to the ^{3}He Larmor frequency at a magnetic field B_0 of 1.5 T. The impedance of the coil must also be matched to the transmission line for optimal efficiency. To measure the frequency response, the coil is connected to an RF frequency generator set to sweep over a bandwidth of 500 kHz. A directional coupler is placed in line with the signal to measure the reflected signal. An oscilloscope, triggered by the sync signal from the frequency sweep, is then used to monitor the reflected signal. A zero in the reflected signal is seen at resonance under critical coupling conditions. Course tuning of the coil is done by adjusting the capacitances by tightening or loosening the screws that hold the capacitor caps in place. Fine tuning is accomplished by moving the tuning coil closer to or further from the birdcage. The tuning must be accompanied by adjustments to the matching circuit, which are done by moving the coupling coil closer to or further from the birdcage. The frequency response and B_1 field homogeneity of the birdcage coil are

discussed in Sec. 4.6.

Figure 3.15: Top view photograph of the birdcage coil. One end ring is visible on the left. The PVC caps used to apply pressure to the capacitors and thereby adjust the capacitances are visible on the right. The single turn coupling coil can be seen edge-on near the bottom of the image. The tuning coil is similarly visible near the top. Nylon bolts and wingnuts are used to adjust the spacing of these coils relative to the resonator.

In order to fit the birdcage coil around the acoustic flow tube in the imaging region of the magnet, the plastic former was split lengthwise with a slitting saw on a milling machine. The end rings thus form a C shape, allowing the former to open at one end. The copper end rings opposite the opening act as hinges. The gap in the end ring is placed at a point where the azimuthal current is zero and thus the resonance of the coil is not influenced. The birdcage is positioned around the acoustic flow tube, and held in place via a pair of PMMA end caps. The end caps have a 16 mm wide slot cut from the edge to the center. They are glued to one half of the PVC former allowing the structure to open and close. An end view photograph of the open birdcage coil is shown in Fig. 3.16.

Figure 3.16: End view photograph showing the birdcage coil in the open position. The resonator is closed once it has been placed around the acoustic flow tube.

Chapter 4

Characterization of the Apparatus

Much of the apparatus built for this project, and described in Ch. 3, required calibration and characterization. This chapter is organized as follows. It begins with a discussion of the diffusion of helium through the walls of the acoustic flow tube in the imaging region, and modifications that were made to inhibit this effective leak (Sec. 4.1). The discussion then moves to issues related to gas composition, acoustics and mechanical resonances. This is followed by a presentation of the calibration of the pressure and mercury displacement transducers (Secs. 4.2 and 4.3). The determination of relevant thermophysical and transport properties of the ^{3}He-O$_2$ gas mixture is discussed in Sec. 4.4. The frequency response of the AMR and a DELTAEC model that was developed to simulate the behaviour of the system are presented in Sec. 4.5. The second half of the chapter deals with issues related to magnetic resonance. A measurement of the homogeneity of the RF fields produced by the birdcage coil used in subsequent MRI experiments is described in Sec. 4.6. Finally, the shimming of the main B_0 field and the development of pulsed magnetic field gradient driven eddy current compensation techniques are discussed in Secs. 4.7 and 4.8.

4.1   Permeability of the Acoustic Flow Tube

When the acoustic flow tube of the AMR was first filled with a mixture of 2 atm ^{3}He and 1 atm O$_2$, the level of the mercury in the U-tubes was observed to displace toward the middle of the apparatus over the course of a few days. The AMR had previously been pressure tested with air to 5 atm, and no indication of a significant leak had been observed.

Figure 4.1: $\ln(p/p_0)$ as a function of time for acoustic flow tubes made with PMMA and soda lime glass.

The effect was therefore attributed to the diffusion of helium gas through the walls of the acoustic flow tube, which were made using PMMA. To quantify the leak, the gas mixture was removed and the acoustic flow tube was refilled with 2 atm of ^{4}He and the compliances were filled to an equal pressure of air. The pressure in the acoustic flow tube was then logged over a period of a few days. Prior to making each pressure measurement, the amount of the air in the compliances was adjusted until the mercury in the U-tubes returned to zero displacement. The pressure was recorded from the Ashcroft gauge connected to the compliances (see Fig. 3.11). The decrease in the pressure in the acoustic flow tube is the result of gas diffusing through the walls. Soda lime glass is much less permeable to helium and oxygen than PMMA, and thus using this material instead should reduce the leak rate. The acoustic flow tube, including the vertical sections of the U-tubes closest to the MRI magnet, were rebuilt using soda lime glass tube with an inner diameter of 1.30 cm and a wall thickness of 1 mm. Once the acoustic flow tube was replaced with soda lime glass, the leak rate was again tested using the procedure outlined above.

For gas diffusing through the wall of a container of fixed volume, the pressure as a

function of time is

$$p(t) = p_0 e^{-Ct} \tag{4.1}$$

where p_0 is the initial pressure at time $t = 0$ and C is a rate constant related to the permeability of the walls and the geometry of the container [90]. Figure 4.1 shows a plot of $\ln(p/p_0)$ as a function of time for the PMMA and soda lime glass acoustic flow tubes. From the slopes of straight lines fit to the data the time constant for the PMMA tube is $C_{\text{PMMA}} = 5.3 \times 10^{-7} \, \text{s}^{-1}$ and the time constant for the glass tube is $C_g = 0.5 \times 10^{-7} \, \text{s}^{-1}$. The permeability of soda lime glass to helium is $K_g = 7.5 \times 10^{-16}$ std. cm^3/s·mm/cm^2/Torr [91]. The measured leak rate constant for the glass acoustic flow tube C_g was compared to the calculated time constant for helium diffusing through the soda lime glass walls obtained using K_g. The observed leak rate is much higher than what is predicted solely for helium diffusion through glass. This excess is attributed to gas diffusing through remaining components in the vicinity of the acoustic flow tube that are still made from plastic, such as the tube couplers and the pressure transducer mount. Although a significant improvement in performance is evident, a measureable loss of gas occurs each day. In order to keep the internal pressure of the AMR constant, O_2 is added to the acoustic flow tube and air is added to the compliances to maintain constant pressure.

The soda lime glass tubing is 0.03 mm larger in outer diameter than that of the PMMA. This is convenient because it allows the tube couplers to be used with only a slight modification to the inner diameter. The inner diameter of the soda lime glass is 0.35 mm larger than that of the PMMA. This change has a more dramatic effect on the experiment as it alters the velocity profile of the gas. The estimated viscous penetration depth in the gas mixture is of order 0.2 cm, as will be shown Sec. 4.4. The nominal value of the ratio R/δ_ν will therefore increase from 2.4 for the PMMA tube to 3.3 for the glass tube. The largest effect of changing the tube diameter in the vicinity of the mercury was to alter the behaviour of the U-tube resonators. The amplitude of the mercury displacement in the glass tube is reduced relative to that in the PMMA tube by the ratio of their internal cross-sectional areas, as will be discussed in further detail in Sec. 4.3. The influence of this change in diameter on the resonant behaviour of the U-tube resonators is discussed in Sec. 4.5.

4.2 Pressure Transducer Calibration

The pressure transducers described in Sec. 3.2.6 required calibration to determine the correspondence between pressure differences and voltage output. The calibration was done in situ, but without any mercury in the U-tubes. The capillary connecting the two ports of each transducer was removed so that the reference port was held at atmospheric pressure. A 1000 Torr full scale MKS Baratron absolute capacitance manometer was connected to the unused port of the manifold and the valves between the manifold and the AMR were left open. A positive differential pressure calibration was performed by first raising the pressure in the AMR by approximately 200 Torr and then slowly venting this overpressure while recording the voltage outputs of the Baratron gauge and the differential transducers. A negative differential pressure calibration was performed in an analogous manner after first using the rotary vacuum pump to lower the pressure in the AMR by approximately 200 Torr. The pressure calibration data and a typical plot of the dynamic pressure differential transducer output while being driven by the AMR are shown in Fig. 4.2. Linear regressions of the calibration data yield sensitivities of 41.39(6) Torr/V, 43.64(6) Torr/V and 41.96(6) Torr/V for the transducers located near the piston, the acoustic flow tube, and the end compliance, respectively.

4.3 Mercury Displacement Transducer Calibration

The mercury displacement transducer is described in Sec. 3.2.5. It was calibrated with 3 atm of a ^{3}He-O_2 mixture in the acoustic flow tube. The copper foil wrapped around the glass tube precludes a visual (optical) measurement of the mercury level in that segment. Instead, the measurement of the relative height of the mercury column was performed on the vertical plastic branch of the U-tube. The mercury level was manipulated by adjusting the pressure of the air in the compliances. Increasing the pressure moves the liquid toward the acoustic flow tube, and vice-versa. The position of the liquid was measured relative to the bottom coupling flange on the vertical tube. The volume of mercury displaced by a change in level Δy is proportional to the area of the tube. The fluid displacement in the glass tube was inferred from measurements of displacement in the PMMA tube using $\Delta y_g = (A_p/A_g)\Delta y_p = 0.57(1)\Delta y_p$, where the subscript g denotes the glass and p the
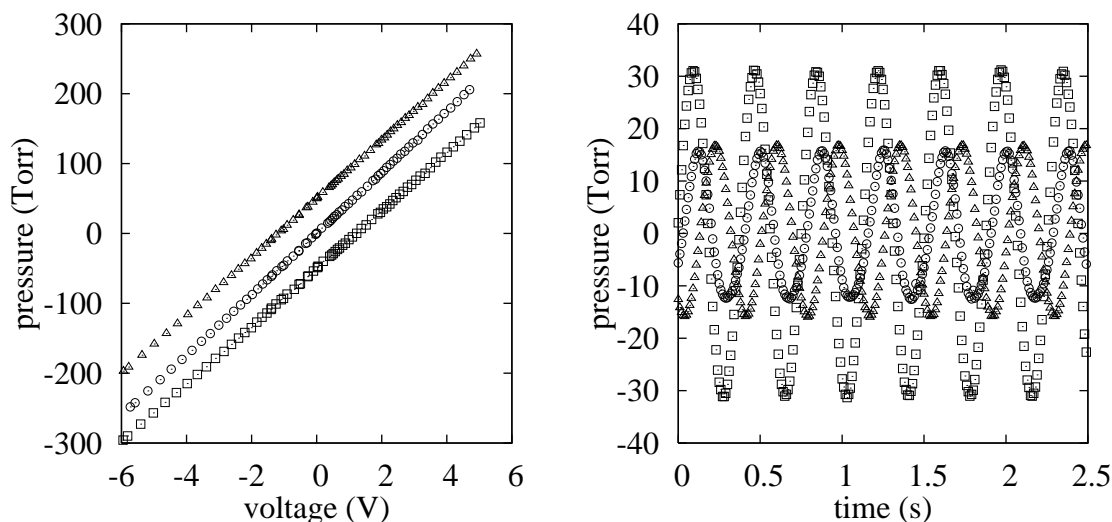
Figure 4.2: Left: Pressure transducer calibration. The data have been vertically separated by 50 Torr for clarity. Right: A typical time-resolved pressure measurement. In both plots the squares are from the transducers near the piston, the circles are from the transducers near the middle and the triangles are from the transducers near the far end.

plastic. The phase offset of the lock-in amplifier was set so that changes in capacitance result in a variation of the real component of the display. The real component of the reading on the lock-in was then recorded along with the level of the mercury. A straight line with an offset was then fit to the data. The slope of this line with an output voltage scaling of a factor of 10 V/500 nA for the lock-in amplifier voltage output were used to arrive at a final calibration of 0.82(3) cm/V for both capacitors. The mercury height data and the best fit lines are shown in Fig. 4.3.

4.4 Gas Properties

To connect the measurement of the velocity profile of the gas with theory, one needs to know the thermophysical properties of the gas mixture. The theoretical velocity profile u_1 (Eq. (2.70)) depends on the viscous penetration depth δ_ν (Eq. (2.45)), which in turn depends on the dynamic viscosity μ and density ρ of the mixture. A comparison of the VPE-MRI measurement was also made against the predictions of a model for the AMR created using

Figure 4.3: Plot of mercury displacement in the vertical glass section of the U-tube as a function of current through the capacitor. Both data sets yield the same slope: 0.0164(5) cm/nA.

DELTAEC, as will be presented in Sec. 4.5. The thermophysical properties of the mixture were input to the model through an external fluid file. This file consists of a series of user defined polynomial coefficients that are used to calculate the density, the specific-heat at constant pressure, the thermal conductivity, the square of the sound speed, and the dynamic viscosity of the gas. Appropriate coefficients were determined using kinetic theory and interpolation of empirical data, as described below.

First, the precise composition of the gas mixture must be determined. Following the procedure described in Sec. 3.3.4, the acoustic flow tube was filled with a mixture of approximately 2 atm of ^{3}He and 1 atm of O_2 at a temperature of 20° C. Although the partial pressures of ^{3}He and O_2 can be determined at the time of filling from known internal volumes and pressures in the the manifold, the pressure in the AMR changes over time as ^{3}He gas diffuses through the walls (cf. Sec. 4.1). It also changes as the ambient temperature of the apparatus changes. Regular adjustments were made to the gas mixture in the acoustic flow tube to maintain the system at a constant pressure. O_2 was added or some of the mixture was removed, and the air pressure in the compliances was adjusted to return the

mercury to equilibrium. The total pressure of the gas in the acoustic flow tube is measured with each adjustment. In order to determine the makeup of the mixture it is sufficient to measure the partial pressure of one of the two components.

As discussed in Sec. 2.1.4, the longitudinal relaxation time T_1 of ^{3}He decreases with the addition of a paramagnetic gas such as O_2 or nitric-oxide (NO). For example, Archibald *et al.* [87] studied the dependence of the relaxation rate, $1/T_1$, of ^{3}He on the concentration of NO. In a field of 1.5 T and at room temperature they found $1/T_1 = 0.0502(3) \times [NO] \mathrm{s}^{-1}$ where [NO] is expressed in amagat.[1] Similarly, Saam *et al.* measured the relaxation of ^{3}He in the presence of O_2 and found $1/T_1 = 0.45(1) \times [O_2] \mathrm{s}^{-1}$, again with $[O_2]$ expressed in amagat [86]. A density of 1 amagat of O_2 would therefore result in a ^{3}He longitudinal relaxation time of 2.22(5) s. If collisions with paramagnetic gas molecules are the dominant relaxation process, then a measurement of T_1 can be used to determine the concentration of O_2 in the mixture. The other relaxation mechanisms that one might consider include dipole-dipole interactions between ^{3}He atoms, collisions with mercury vapour, interactions with the wall, and the diffusion of ^{3}He in magnetic field gradients. The longitudinal relaxation time for pure ^{3}He due only to dipole-dipole interactions within the gas is estimated by Newbury *et al.* to be $1/T_1 = 1.34 \times 10^{-5}[^3\mathrm{He}] \mathrm{s}^{-1}$ [92], which is negligible compared to the effect observed for O_2. Timsit has measured the effect of various metal surfaces on the relaxation of ^{3}He and noted that mercury vapour did not cause any appreciable relaxation over a time scale of order 10^4 s [88]. The relaxation of ^{3}He associated with walls made of various glasses (and with various coatings) has been measured by several groups [88, 57]. Relaxation times in excess of 10^4 s were routinely measured. Finally, diffusion of ^{3}He in magnetic field gradients should be considered. A particle undergoing a random walk in a magnetic field gradient will experience variations in magnetic field strength. If the spectral density of this changing field is non-zero at the Larmor frequency, nuclear spin flips (and hence relaxation) can be induced. Schearer and Walters[2] showed that the relaxation associated with Brownian motion in a magnetic field gradient yields a relaxation rate [94]

$$\frac{1}{T_1} = \frac{2}{3}\gamma^2 G^2 \left\langle U^2 \right\rangle \frac{\tau_c}{\omega_0^2(\omega_0^2\tau_c^2 + 1)} \tag{4.2}$$

---

[1] An amagat is a unit of density equivalent to the density of an ideal gas at atmospheric pressure and $0°$ C (1 amg ≈ 44.6 mol/m^3).

[2] Gamblin independently derived a similar result using a classical random walk argument [93].

where $G = \partial B / \partial x$ is the magnitude of the transverse magnetic field gradient, $\langle U^2 \rangle$ is the mean square molecular velocity, τ_c is the time between collisions, and ω_0 is the mean Larmor frequency. Taking $G = 10$ G/cm as the worst case scenario (gradients near the ends of the acoustic flow tube where it leaves the MRI magnet) yields values of $T_1 \sim 10^4$ s. Even this is likely to be an underestimate as the time required for ^{3}He to diffuse half the length of the acoustic flow tube ($\tau \sim l^2 / D$) is of order 10^4 s as well.

The estimates presented above indicate that O_2 will completely dominate the ^{3}He longitudinal relaxation time, and hence the O_2 density can be inferred from a measurement of T_1. These measurements were regularly made as necessary using the inversion-recovery method. First, the nuclear magnetization is inverted with a π pulse after which it immediately starts to relax back to equilibrium. After a delay τ, a $\pi/2$ pulse is applied to measure the z-component of the magnetization at that time. The inversion-recovery sequence is then repeated for different delay times τ and the resulting signal is compared to the function

$$S(\tau) = S_0(1 - 2e^{-\tau/T_1}) \quad , \tag{4.3}$$

which is simply Eq. (2.9) with $M_z(0) = -M_0$. In practice it is often difficult to achieve perfect π pulses, and hence the inversion of the magnetization is imperfect. This can be remedied to some extent by replacing the coefficient of 2 appearing in Eq. (4.3) by a variable b such that $1 \leq b \leq 2$. Fitting this empirical function to the data typically resulted in a value of $b = 2$ within a 2% uncertainty, indicating that the quality of the π pulses was very good. A sample T_1 inversion-recovery data set is shown in Fig. 4.4 along with a fit to Eq. (4.3). The O_2 density inferred from these data is 1.40(4) amagat (1.50(5) atm at 20° C). With a total gas pressure of 3.013(15) atm, the remaining 1.51(4) atm is associated with ^{3}He.

Once the composition of the gas mixture has been determined, its thermophysical properties can be calculated using kinetic theory and the ideal gas law. Treating the mixture as an ideal gas, the mean molar mass is $M = x_1 M_1 + x_2 M_2$ where x_1 and x_2 are the ^{3}He and O_2 molar fractions and M_1 and M_2 are the respective molecular masses. The mixture density is $\rho_m = M p_m / RT$ where R is the universal gas constant. Similarly, the isobaric heat capacity per mole is $C_p = x_1 C_{p1} + x_2 C_{p2}$ which leads to an isobaric heat capacity per unit mass of $c_p = C_p / M$. The same method is used to calculate the isochoric heat capacity per mole and heat capacity per unit mass. With these quantities in hand, the heat capacity ratio $\gamma_{\text{ideal}} = C_p / C_v$ and the sound speed $a = \sqrt{\gamma_{\text{ideal}} RT / M}$ can then be calculated.

Figure 4.4: Amplitude of the NMR signal as a function of time for an inversion-recovery experiment. The solid line represents a fit to Eq. (4.3) yielding $T_1 = 1.59(4)s$.

The viscosity, thermal conductivity, mass diffusion coefficient and thermal diffusion ratio of the mixture are then calculated using the procedure presented in *Molecular Theory of Gases and Liquids* by Hirschfelder, Curtiss and Bird [72]. The approach is based on the Chapman-Enskog theory using a Lennard-Jones (6-12) potential. The details of the implementation of the transport coefficient calculations used in this work are given in App. C.

4.5 DELTAEC Model and AMR characterization

DELTAEC is an important tool that is often used during the preliminary design of thermoacoustic devices (see Sec. 2.2.5 for a brief introduction). The algorithm is based on Rott's thermoacoustics, which was developed in the low amplitude limit where the time dependent components of pressure, velocity, etc. are small (see Sec. 2.2). Practical devices often operate outside of this limit and consequently the results generated through integration of the appropriate differential equations are often inadequate. A number of empirical techniques have consequently been incorporated into DELTAEC to improve the correspondence between model calculations and the performance of real devices. This is done by inserting

Figure 4.5: Schematic of the DELTAEC model of the AMR.

segments that modify the solution process. For example, at high volume flow rates, changes in gas flow direction are associated with additional pressure drops, known as minor losses, that are not accounted for in linear DELTAEC models. Additional segments can be added to a model that compensate for these minor losses to improve the agreement between the model and measurements used to characterize the device.

One of the goals for building the AMR and attempting to study gas motion using VPE-MRI is to be able to ultimately create acoustic flow conditions that are found in practical devices, but that are not accurately modelled in DELTAEC. To this end, a DELTAEC model for the AMR was developed for comparison with VPE-MRI data.

The model of the AMR comprises eight distinct components. From right to left in Fig. 4.5 these are the bellows, the piston compliance, the piston duct, the piston U-tube resonator, the acoustic flow tube, the end U-tube resonator, the end duct and the end compliance. The bends in the tubing are neglected for the time being. The duct and compliance components are modelled using the DUCT segment. In DELTAEC, a DUCT is simply a tube where the perimeter, the cross-sectional area, and the length are used to define the dimensions. The bellows is also modeled as a duct, in which the average of the inner and outer diameters is used to calculate the cross-sectional area and the perimeter is an effective value set so as to give the correct surface area for the convolutions of the walls.

The U-tube resonators are modelled using the VESPEAKER segment, which simulates a voice coil driver that is in series with the flow. A voice coil behaves like a damped mass and spring resonator with an electromagnetic drive consisting of a resistive coil that moves a permanent magnet. The U-tube resonators can also be modelled as a damped mass and spring resonator which means that they can simulated using the damped harmonic oscillator parameters of the VESPEAKER segment with the electromagnetic drive parameters set to

zero.[3]

Empirical mass, spring constant and damping coefficients for each U-tube were established by measuring the decay of free oscillations of the mercury at atmospheric pressure. First the mixture was removed from the acoustic flow tube and the AMR was brought to atmospheric pressure. The piston duct and acoustic flow tube were disconnected from the top of the piston U-tube. The piston drive was set near the resonant frequency and the piston duct was temporarily connected to excite oscillations in the mercury. The piston duct was then decoupled from the U-tube allow the mercury to oscillate freely. The output of the mercury displacement transducer was recorded for 10 s. A similar procedure was employed to measure the free oscillation decay for the end U-tube resonator.

The equation of motion for a damped mass and spring harmonic oscillator with a linear resistive force can be written

$$\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} + \frac{b}{m}\frac{\mathrm{d}x}{\mathrm{d}t} + \frac{k}{m}x = 0 \tag{4.4}$$

where m is the mass, b the damping coefficient and k is the spring constant. The solution is of the form

$$x(t) = Ae^{-\Gamma t/2}\cos(\omega t + \alpha) + x_0 \tag{4.5}$$

where x_0 is the initial position, A is the amplitude of the motion, $\Gamma = b/m$ and α is the phase of the oscillation. The angular frequency of the motion is given by

$$\omega^2 = \omega_0^2 - \frac{\Gamma^2}{4} = \frac{k}{m} - \frac{b^2}{4m^2} \tag{4.6}$$

where $\omega_0 = \sqrt{k/m}$ is the angular resonant frequency of the oscillator in the absence of damping. The VESPEAKER model requires k, m, and b to define the behaviour of the damped oscillator. Equation (4.5) was fit to the data from each U-tube with A, Γ and α as free parameters. It was immediately evident that there was an additional exponentially decaying offset in the data. That is, the mercury displacement was best fit by an empirical function of the form

$$x(t) = x_0 + A\cos(\omega t + \phi)e^{-\Gamma t/2} + A_1(1 - e^{-\Gamma_1 t/2}) \tag{4.7}$$

---

[3]In practice, it was necessary to assign a non-zero value for the coil resistance to avoid a divide by zero error. The value of this resistance has no effect on the outcome of the model calculations with the other electrical parameters of the VESPEAKER segment set to zero.

(a) Piston U-tube

(b) End U-tube

Figure 4.6: Free decay of mercury oscillations in the piston and end U-tube resonators. The solid line in each plot is the result of fitting Eq. (4.7) to the data. The dashed lines represent the exponentially decaying offset. The parameters extracted from these fits are: (a) $\omega = 7.383(1)$ s^{-1}, $\Gamma = 0.569(2)$ s^{-1}, $\Gamma_1 = 0.35(2)$ s^{-1} (b) $\omega = 7.405(3)$ s^{-1}, $\Gamma = 0.730(5)$ s^{-1}, $\Gamma_1 = 0.63(4)$ s^{-1}

where A_1 is the offset amplitude and Γ_1 the second decay constant. Mercury displacement data for both the piston and end U-tubes are shown in Fig. 4.6, along with fits of Eq. (4.7). The contribution of the exponentially-decaying offset term in Eq. (4.7) is also plotted. Agreement between the fit and the data is very good. The VESPEAKER segment is not capable of modeling the exponentially decaying offset, and so this piece of information is disregarded in the model. Mechanisms that might be responsible for this effect are discussed later in this section.

Two of the variables extracted from the fits of Eq. (4.7) to the data were used to establish corresponding parameters in the VESPEAKER segments of the DELTAEC model: Γ and ω. However, this is not sufficient. A third parameter, either the effective spring constant or the effective mass of the U-tube resonators, needs to be specified. The frequency of the damped oscillation ω depends on the undamped resonant frequency of the system ω_0. An expression for the theoretical value ω_0 can be used to completely determine the VESPEAKER parameters.

(a) Simple U-tube   (b) One wider vertical arm   (c) Consticted U-tube

Figure 4.7: Models of the U-tube resonator. (a) Simple U-tube with uniform cross-sectional area. (b) Uniform cross-sectional area through the bend and one wider vertical section. (c) A constriction in the bend.

One method for analyzing the resonance of an undamped simple harmonic oscillator is to look at the total mechanical energy of the system [95]. The sum of the kinetic and potential energies is

$$E = K + U = \frac{1}{2}mv^2 + \frac{1}{2}kx^2 \quad . \tag{4.8}$$

where v is the velocity of the fluid. Three physical models of the U-tube resonator are shown in Fig. 4.7, in order of increasing complexity. To begin, consider the simple U-tube shown in Fig. 4.7a. It has a uniform cross-sectional area A and the total mass of liquid is given by $m = \rho A L$, where ρ is the fluid density and L is the linear length of liquid. Assuming that all of the fluid moves at the same velocity, the kinetic energy is $K = \frac{1}{2}\rho A L v^2$. The potential energy associated with a displacement x away from equilibrium is $U = g\rho A x^2$, where g is the gravitational constant. From these expressions we can define an effective spring constant $k_{\text{eff}} = 2g\rho A$ and effective mass $m_{\text{eff}} = \rho A L$. The natural resonant frequency of the system is then $\omega_0^2 = 2g/L$.

The analysis is similar for the case shown in Fig. 4.7b, in which one of the two vertical sections of the U-tube involves a segment of length l_g that has a larger area A_g. The remainder of the tube (including the bend) has area A_p and length l_p. The subscripts here are intended to refer to the glass and plastic tubes used in the construction of the real apparatus. Recall that the displacement of liquid is measured in the glass tube. The potential energy

associated with a displacement x is now

$$U = \frac{1}{2}\rho g A_g \left(1 + \frac{A_g}{A_p}\right) x^2 \quad . \tag{4.9}$$

The effective spring constant of this system is then $k_{\text{eff}} = \rho g A_g (1 + A_g/A_p)$. If the amplitude of the displacement is small so that $|x| \ll l_g$, then the volume of liquid in each section remains approximately constant and the kinetic energy can be written

$$K = \frac{1}{2}\rho A_g \left(l_g + l_p \frac{A_g}{A_p}\right) v_g^2 \tag{4.10}$$

giving an effective mass of $m_{\text{eff}} = \rho A_g (l_g + l_p A_g/A_p)$. In practice this is not a good assumption. The maximum excursion of the mercury is comparable to l_g. From this perspective, one is better off using Eq. (4.9) to extract an effective spring constant rather than using Eq. (4.10) to extract an effective mass. The analysis procedure was thus to extract Γ and ω from fits of Eq. (4.7) to the mercury oscillation decay data (e.g. Fig. 4.6) and then rely on tube dimensions to provide an effective spring constant k_{eff} (through Eq. (4.10)). The effective mass was then assigned so as to give the natural resonant frequency implied by Eq. (4.6) (i.e. $\omega_0^2 = k_{\text{eff}}/m_{\text{eff}}$). The effective mass, spring constant, and damping coefficients extracted from this analysis were then used in the DELTAEC VESPEAKER model of the U-tube resonator.

A number of assumptions are buried in the semi-empirical resonator model described above. The first of these is that the plastic section of the U-tube has a constant cross-sectional area. In the process of forming the resonators the tubing collapsed slightly and thus the cross-sectional area in the vicinity of the bend is smaller than the nominal area A_p. This situation is shown in Fig. 4.7c. Under these conditions, the velocity of the liquid through the bend will be higher than in the straight plastic section, increasing the total kinetic energy. This is not a significant issue and in practice is accounted for by using an effective mass. The second assumption is that the displacement amplitude is small. As noted above, this is a poor assumption and the primary motivation for not using Eq. (4.10) to define the effective mass. Accounting for large amplitude displacements leads to the following nonlinear differential equation

$$E = \frac{1}{2}\rho A_g \left[l_p \frac{A_g}{A_p} + l_g + \left(1 - \frac{A_g^2}{A_p^2}\right) x\right] \left(\frac{\mathrm{d}x}{\mathrm{d}t}\right)^2 + \rho g A_g \left(1 + \frac{A_g}{A_p}\right) x^2 \tag{4.11}$$

where now the effective mass in the kinetic energy term is a function of displacement. Another effect that could have a significant impact on the performance of the model is associated with the flow of liquid through the junction between the glass and plastic tubes. As a real fluid flowing through a tube experiences a sudden change in diameter, some kinetic energy is dissipated. For steady flows losses are typically of the form $E \propto Kv^2$ where v is the velocity of the fluid and K is a loss coefficient. In general, the value of K associated with a sudden enlargement of tube diameter is greater than that for a sudden contraction in diameter. For an oscillatory flow, this assymetry results in a mean pressure difference across the junction. This effect has been used previously to build a "gas diode" which in effect leads to a partial rectification of flow driven by the net pressure difference across the constriction [96]. It may play a role in the work described here, in that it may contribute to the exponentially decaying offset observed during the mercury free decay experiments.

The remaining pieces of information that must be entered into the DELTAEC model are the conditions of the gas in the system. There are two fluids used in the AMR. The two compliances and associated volumes contain dry air and the acoustic flow tube in the imaging region is filled with a mixture of ^{3}He and O_2. DELTAEC has a number of "built in" fluids that can be used to construct device models, including a model for air. It also has the capability for accepting "user defined" fluids. A file defining the ^{3}He-O_2 mixture used in this work was created using parameters derived from measurements of mixture composition and calculations of relevant thermophysical and transport properties as outlined in Sec. 4.4. The remaining parameters required by the DELTAEC model include the mean internal pressure, the mean temperature, the drive frequency, and the amplitude of the drive volume velocity, all of which are measured. The drive volume velocity is estimated by calculating $|U_1^{\text{piston}}| = \omega |\xi_1^{\text{piston}}| A^{\text{piston}}$, where all three parameters on the right hand side of this expression are determined from measurements.

The DELTAEC model of the AMR was tested against measurements of mercury displacement under a variety of conditions. The photodiode trigger pulse, pressure transducer, and mercury displacement transducer outputs were recorded for a period of 5 s at various drive frequencies. The time data were then fit to a sinusoidal function to extract the amplitude and phase relative to the trigger pulse of the pressure and mercury displacement oscillations.

The initial characterization of the AMR model was performed at atmospheric pressure,

(a) Piston U-tube   (b) End U-tube
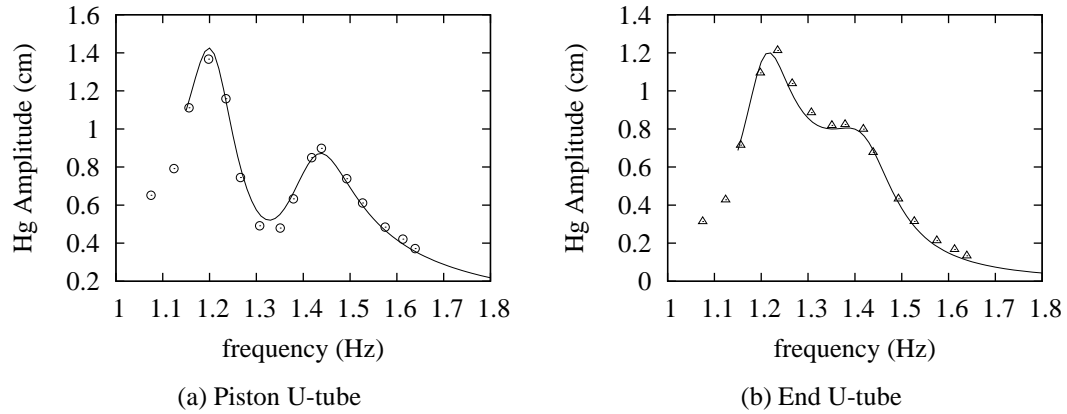
Figure 4.8: Measured (data points) and modelled (curve) frequency response of the AMR when the entire apparatus is filled with air at atmospheric pressure.

and with the entire volume (including the acoustic flow tube) filled with air. The only variable that was adjusted to improve the correspondence between the model and the measured response was the volume displaced by the bellows as it was moved by the piston. The frequency response of the mercury displacement under these conditions is shown in Fig. 4.8, along with predictions of the DELTAEC model. The model had difficulties converging on a solution at frequencies below 1.15 Hz. Here, and in subsequent plots showing the frequency response of the AMR, conversion from measured mercury displacement amplitude ξ_1 to volume flow rate U_1 is obtained through $U_1 = \omega_A A_g \xi_1$ and is typically known to 2%. The mercury filled U-tube resonators form a coupled harmonic oscillator with two natural modes, with the gas in the acoustic flow tube acting as the coupling agent. The low frequency mode is characterized by in-phase motion of the mercury in the two resonators and yields the highest gas displacement amplitudes. The high frequency mode involves antiphase motion of the mercury, producing large amplitude pressure oscillations and small amplitude gas displacement near the center of the acoustic flow tube. These two modes are associated the two peaks in mercury displacement amplitude that are evident in Fig. 4.8.

The internal pressure of the system is approximately 3 atm when the acoustic flow tube is charged with an appropriate gas mixture for MRI. The higher pressure causes the frequency of both natural modes to increase. This is due to the stiffening of the compressibility of the gas, which is inversely proportional to pressure, in the compliances as well

as in the acoustic flow tube. In order to test the response of the system and to compare its behaviour with the DELTAEC model, the frequency response was measured again with the ^{3}He-O$_2$ gas mixture in the acoustic flow tube and a corresponding amount of air in the compliances. The external fluid model of the gas mixture was inserted into the acoustic flow tube section of the AMR in the model. The data and the model calculations are shown in Fig. 4.9. Only the lower frequency mode is still visible in the same frequency window. The maximum of the measured mercury displacement amplitude, and hence the mercury velocity amplitude, shows an increase in going from 1 atm to 3 atm internal pressure. This is the result of stronger coupling between the piston and the U-tubes at higher pressure. It is evident that there are significant additional losses at higher amplitude that are not accounted for in the model, which assumes laminar flow.

(a) Piston U-tube   (b) End U-tube

Figure 4.9: Measured frequency response of the AMR when charged to 3 atm, and the predictions of the DELTAEC model assuming laminar gas flow throughout. The only damping damping included here is associated with the mercury.

The Reynolds (Re), Dean (De) and Womersley (α) numbers are used for classifying flow patterns in tubes for straight steady flow, flow around bends, and pulsatile flow, respectively. A description of these dimensionless numbers can be found in Sec. 2.2.6 and their definitions are given in Table 2.1. The AMR has four 90° elbows that cause the gas flow to transition between horizontal motion (in the piston duct, end duct, and acoustic flow tube) and vertical motion (in the U-tubes); see Fig. 3.1. The two elbows fabricated from PMMA (containing air) have a mean bend radius $R = 2.5$ cm while those fabricated from

glass (containing the ^{3}He-O$_2$ mixture) have $R = 0.9$ cm. Typical Reynolds, Dean, and Womersley numbers for air at 1 atm, for air at 3 atm, and for a ^{3}He-O$_2$ mixture at 3 atm under conditions found in the AMR, are listed in Table 4.1. The Reynolds and Dean numbers are calculated using the measured peak mercury velocity at 1 atm and 3 atm. The Reynolds numbers are all well below the value of order 2000 at which which self-maintaining turbulence is observed in steady pipe flow. The Womersley numbers are in the transitional range between Poiseuille-like flow (α<1) and plug flow (α>10) and thus do not lead one to anticipate unstable flow. In fact, a comparison of the conditions in the acoustic flow tube with the experimental conditions reported by Kurzweg *et al.* [97] strongly support the hypothesis that flow in this region of the apparatus should be laminar. The Dean numbers for gas flow through the elbows, on the other hand, are large enough that significant secondary flows are expected; however they are still well below the typical transition threshold for the onset of turbulence at De~5000 [81]. Here it is worth noting that the ratio of the tube radius to the mean bend radius is of order 0.7 for the glass elbows and 0.2 for the plastic elbows. For ratios this large, deviations from the Dean-like flow patterns illustrated in Sec. 2.2.6 can be expected [81]. This could lead to additional losses as the gas moves through the elbows at peak velocity; however, it is not even remotely possible to reconcile the anticipated magnitude of these losses [98] with the increase in damping that was observed when the mean pressure in the AMR was increased from 1 atm to 3 atm.

Table 4.1: Peak Reynolds, peak Dean, and Womerseley numbers for air at 1 atm and 3 atm, and for a 3 atm 50-50 mixture of ^{3}He-O$_2$ at an AMR operating frequency of 1.3 Hz.

Number	Air 1 atm	Air 3 atm	1.5 atm ^{3}He-1.5 atm O$_2$
Reynolds	230	680	270
Dean	99	370	230
Womerseley	3.5	6.0	6.1

Instead, the mechanism underlying the enhanced damping observed at higher pressure is almost certainly caused by the flow of mercury around the 180° bends in the two U-Tubes. As was the case for gas flow through the 90° elbows, secondary (Dean) flows will be established in the mercury as it oscillates. Taking the mean bend radius R of a U-tube to be 2.25 cm, the inner diameter D of the tube to be 0.95 cm, the acoustic frequency to

be 1.31 Hz and the kinematic viscosity of mercury to be 1×10^{-7} m^2/s, one obtains a peak Dean number of 3400 for a measured peak displacement amplitude of 0.5 cm and 17000 for a measured peak displacement amplitude of 2.5 cm. Here it is important to recall that the mercury displacement amplitudes are measured in the larger diameter glass tube. Thus, mercury dispacement amplitudes in the acrylic bends are larger than measured amplitudes by a factor $A_g/A_p \sim 1.75$ (see Sec. 4.3). These estimates imply a transition from laminar to turbulent flow at De ~ 5000 for a peak measured displacement amplitude of order 0.7 cm at a frequency of 1.31 Hz, corresponding to a peak velocity amplitude of 12 cm/s. The measured peak velocity amplitude of the mercury in Fig. 4.8 is 18 cm/s and in Fig. 4.9 is 37 cm/s, suggesting losses due to turbulent flow in the mercury are to be expected.

The observed non-linear (amplitude dependent) damping was modeled in DELTAEC by introducing MINOR segments for each U-Tube, which introduce pressure drops proportional to U_1^2. The coefficients governing the MINOR segment losses were then adjusted through a trial and error until satisfactory agreement was reached between the model and the data. Caution is required in interpreting the physical meaning of the MINOR loss coefficients. In particular, the MINOR segments should not be viewed as a model for dissipation in the gas; rather they serve as an empirical model for dissipation in the mercury. In principle this could also be modeled by introducing a velocity dependent mechanical resistance in the VESPEAKER segments used to model the U-Tubes.

(a) Piston U-tube   (b) End U-tube

Figure 4.10: Same measured AMR response as shown in Fig. 4.9. Minor loss segments have been added to the DELTAEC model to account for the observed additional damping.

Figure 4.10 shows the measured frequency response of the mercury displacement amplitude at 3 atm along with the predictions of the DELTAEC model after the MINOR loss coefficients were optimized. It is evident that there is a marked improvement in the agreement between the model and the data. The measured phase of the mercury displacement in both U-tubes as well as the results of the model are shown as a function of frequency in Fig. 4.11a. The phase reference for the model was shifted by 0.5 rad to compensate for the timing of the photodiode trigger pulse. Also plotted in this figure is the model prediction for the phase of the gas in the vicinity of the birdcage coil. The VPE-MRI experiments that will be described in Ch. 5 were performed at an acoustic frequency of 1.31 Hz. At this frequency the motion of the mercury in the two U-tubes is nearly in phase. The solid line in Fig. 4.11a shows the DELTAEC model prediction for the phase of the gas velocity in the acoustic flow tube. The model indicates that the gas displacement amplitude and phase can be estimated from the average of the measured mercury displacement versus time data. The results are shown in Fig. 4.11b along with the DELTAEC model prediction for gas displacement near the birdcage coil. Again there is good qualitative agreement between the model and the data.

On average, the displacement amplitude data lie within 10% of the model predictions. The data appear to reach their maximum amplitude at a slightly higher frequency, and reveal a slightly broader peak, both of which are indications that damping is still underestimated in the DELTAEC model. Improving the level of agreement between the model and the data would likely require further characterization of the amplitude dependent losses in the U-Tubes.

4.6 Birdcage Coil RF Field Homogeneity

In MRI, the application of RF tipping pulses and the inductive detection of nuclear spin precession (i.e. signal acquisition) is done through RF coils or resonators tuned to the Larmor frequency. In order for all of the spins in the coil to experience the same tipping pulses, the B_1 field must be homogeneous over the volume of the sample. Under these conditions, the sensitivity of the RF coil to spin precession will also be homogeneous.

The RF field homogeneity inside of the birdcage coil was evaluated using a cavity perturbation method similar to that outlined by Hayden and Hardy [99]. The shift in the res-

(a) Phase
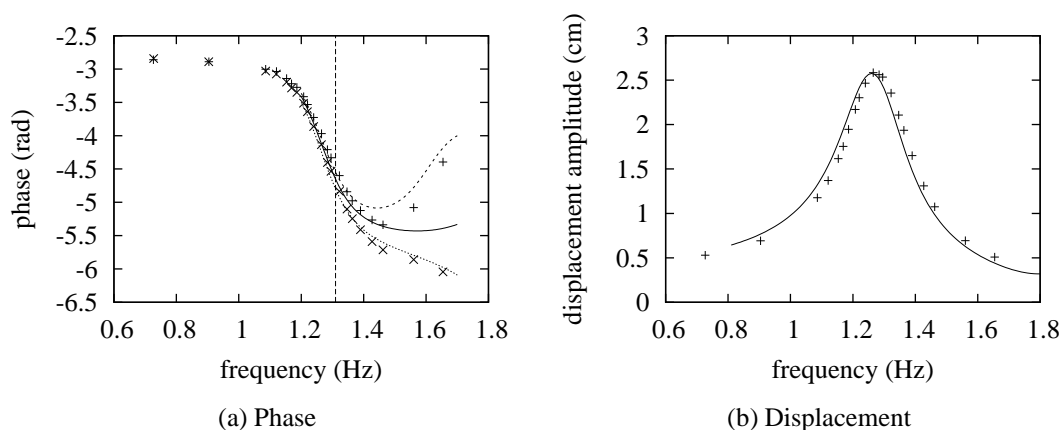
(b) Displacement

Figure 4.11: (a) Phase of the mercury displacement in the (+) piston and (×) end U-tubes. The dashed and dotted lines are the phases from the DELTAEC model for the mercury at the piston and end sides of the acoustic flow tube, respectively. The solid line is the model prediction for the phase of the gas displacement at the location of the birdcage coil. The VPE measurements were performed at a frequency of 1.31 Hz which is marked by the vertical dashed line. (b) The mean displacement amplitude of the mercury in the U-tubes. The line is the DELTAEC model prediction.

onant frequency of a tuned coil is measured as a small sphere of dielectric or conducting material is scanned through the sample volume. A dielectric sphere will perturb the electric field of the RF resonator while a conducting sphere will influence both the electric and magnetic fields. It is therefore possible to directly probe the RF electric field but more challenging to isolate just the magnetic contribution. Fortunately in the quasistatic limit appropriate to the birdcage coil used in this work, the RF electric fields are entirely confined to the capacitors around the end ring. This was confirmed by scanning a 4 mm diameter Teflon sphere, which has a dielectric constant $\eta = 2.1$, through the sample volume; no shift in the resonance of the coil was observed at the level of 1 part in 10^5. The implication is that all subsequent resonant frequency shifts caused by the conducting spheres described below are a result solely of RF magnetic field perturbations.

The resonant frequency of the birdcage coil was measured using the phase-locked-loop (PLL) circuit depicted in Fig. 4.12. This circuit uses feedback to maintain a constant phase shift between the signal provided by the function generator and signal detected by the RF
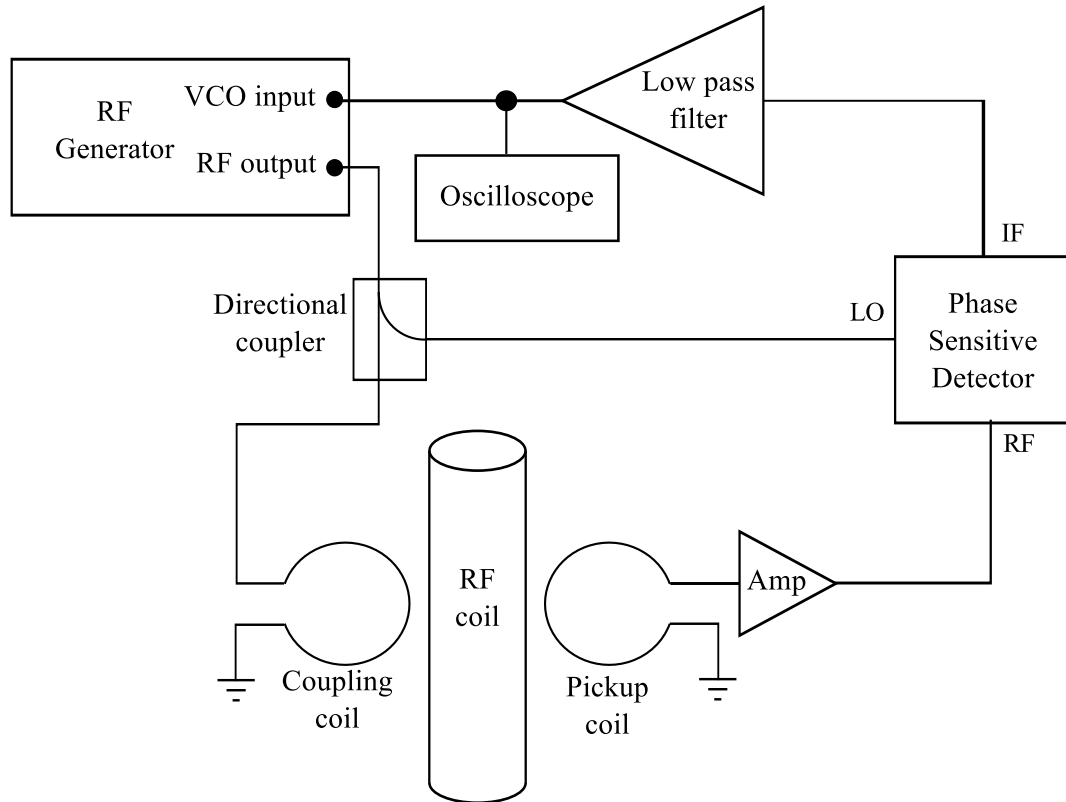
Figure 4.12: Phase locked loop used in the cavity perturbation measurement of the birdcage RF magnetic field homogeneity.

coil. The function generator is operated as a voltage-controlled-oscillator (VCO), in which the output RF frequency depends on the voltage applied to VCO input. The output of the VCO is inductively coupled to the resonator through the coupling loop. A small pickup loop located on the opposite side of the resonator is then used to monitor the RF fields excited in the birdcage coil. The direct coupling (mutual inductance) between the coupling coils and pickup coils is very small. The RF coil behaves like a forced-damped oscillator, and in general there is a phase shift between the driving signal applied to the coupling coil and that received by the pickup coil. When the drive frequency matches the resonance of the RF coil this phase shift is 90°. The phase shift is measured using a Mini-Circuits SRA-2CM phase sensitive detector (PSD). This device mixes two RF signals producing an output that is the product of the two inputs, and which can be decomposed into a superposition of signals at the sum and difference frequencies of the inputs. The component oscillating

at the sum of the two input frequencies is typically at a much higher frequency than the component oscillating at the difference frequency, and can be removed using a low pass filter. For small phase angles the amplitude of the component at the difference frequency is proportional to the phase difference. One of the PSD inputs is derived from the output of the RF generator by way of directional coupler. The other input is the signal from the pickup loop, after being is amplified to the same level as the reference signal. The low pass filter then removes the component of the PSD output at the sum frequency, and the remaining signal at the difference frequency is fed back to the VCO input of the function generator to complete the loop. The net result is that the RF output frequency tracks the resonant frequency of the birdcage coil. Perturbations to this frequency, caused by the insertion of a small conducting sphere are reflected in the VCO frequency output. The VCO input signal, which is proportional to the frequency shift, is measured using an oscilloscope.

Small 4.07(5) mm diameter spheres of Teflon and steel were used to perturb the RF field inside the birdcage coil. The experimental setup is depicted in Fig. 4.13. The beads were attached to a fine nylon thread running through the resonator (parallel to its axis) and over a pulley to a counter weight, providing tension. A set of Vernier calipers with 0.01 mm precision fixed to the lab bench were used to position the bead along the length of the resonator. The VCO input voltage was monitored using a digital oscilloscope set to scan at 1 s/div providing 10 s of recording time. After the bead was inserted and the PLL had stabilized, the oscilloscope was stopped. The change in voltage caused by the insertion of the bead was then measured using the cursors. Measurements were made along the axis of the coil as well as at radial displacements of 1.67(5) cm and 2.93(5) cm.

The resonant frequency of the birdcage coil drifted in time at a rate where some care had to be taken to not compromise the perturbation measurement. The uncertainty caused by this effect was minimized by rapidly inserting the bead into the RF coil and measuring the shift on a timescale of order 1 second. The Vernier calipers were first adjusted to set the bead position and then a bight of thread was withdrawn to remove the bead from the coil, as shown in Fig. 4.13. The bight was then released, allowing the counter weight to pull the bead back into position, at which point the oscilloscope trace was stopped and the change in voltage was measured.

Figure 4.14 shows the measured frequency shift caused by the metal sphere as a function of position. These data directly reflect the relative RF magnetic field amplitude. The

Figure 4.13: Experimental setup used in the cavity perturbation measurement.

homogeneity of this field is very good over an axial distance of 4 cm and out to a diameter of 4.86 cm, a volume over which the measured frequency shift varies by less than 10%.

The frequency response of a resonant circuit such as the birdcage coil can be characterized by a Quality factor. This dimensionless number is measured by finding the ratio of the resonant frequency and the full-width-at-half-maximum (FWHM) of the power spectrum. The loaded Quality factor of the birdcage coil, isolated from conducting surfaces, was measured to be 440.

4.7 B_0 Field Homogeneity

In order to obtain artifact-free MRI images, the static magnetic field B_0 must be as homogeneous as possible. This ensures that all of the spins in the sample have the same Larmor frequency. Imaging magnets typically employ a superconducting solenoid carrying a persistent current to produce this field. Often, it is not intrinsically homogeneous at the level required for MRI. In order to improve homogeneity the magnet must be shimmed. That is, small spatially-inhomogeneous magnetic fields are added to B_0 in order modify the overall profile of the field. The term "shim" originates from the use of small pieces of ferromagnetic sheet metal placed in the field to improve the homogeneity of electromagnets. Modern magnets are typically shimmed using specially-designed coil sets that produce linear and non-linear magnetic field gradients that can be independently controlled by adjusting the

Figure 4.14: Frequency shift of the birdcage coil as a function of position during a cavity perturbation measurement performed using a conducting sphere, as described in Sec. 3.4. The resonant frequency was 48.351 MHz. The different datasets reflect measurements performed at various radial distances from the axis of the coil.

appropriate currents.

The imaging magnet used in this work is a 30-cm-bore Oxford Instruments Horizontal Superconducting Magnet System 80/300 operated at 1.5 T. The imaging system consists of a Tecmag Apollo console controlled by NTNMR version 2.4.21 software. The x, y, and z linear magnetic field gradients are driven using Techron 8607 gradient amplifiers that are controlled via NTNMR through an 18-bit DAC. The imaging magnet is equipped with superconducting shims that were set and left in persistent mode when the magnet was first cooled. Inserted into the bore of the magnet is a set of room temperature (i.e. resistive) shim coils that are driven by an array of constant current sources. The current delivered to each coil is independently set by adjusting a 10 turn potentiometer.

The homogeneity of the magnetic field can be characterized by measuring the FID from a sample. This can be done either in the time domain or, by performing a Fourier transform (FT), in the frequency domain. In the time domain, an improvement in shimming leads to a larger initial signal amplitude and a slower decay of the signal over time. In the frequency

domain, improvements are characterized by an increase in the area under the curve and a narrowing of the linewidth. In fact, the linewidth characterizes the distribution of Larmor frequencies over the sample.

Figure 4.15: Tube coil used to image the water phantom.

The initial round of magnetic field shimming was performed using an 88 μmol/cm^3 CuSO$_4$-water solution as a sample or a "phantom." The CuSO$_4$ was added to lower the ^{1}H nuclear relaxation time T_1 to about 0.15 s. At this point, the AMR had already been assembled and pressure tested in situ, prior to energizing the magnet. Consequently the acoustic flow tube obstructed the axis of the magnet. In order to place the water sample near the center of the magnet a special tuned tube coil was built. As shown in Fig. 4.15, this resonator fits around the acoustic flow tube. The cylindrical former for the tube coil has a pair of holes perpendicular to its axis and has been cut in half lengthwise so that it can be assembled around the the acoustic flow tube. The coil itself is made with copper sheet wrapped around the former creating a single turn inductor. The two ends of this inductor are brought close together, with only a thin Teflon sheet between them, forming a capacitive junction. The area and spacing of these capacitor plates is then adjusted to set the resonant frequency to 64 MHz (the ^{1}H Larmor frequency in a 1.5 T field). At resonance, this coil produces a $\mathbf{B}_1$ field along its axis. Two 10 ml phantoms were placed in the tube coil, one above and the other below the acoustic flow tube. Once in place, the resonator and phantoms could be translated along and rotated about the z-axis of the magnet with ease.

The superconducting shims and, initially, the room temperature shims (see Table 4.1) were adjusted through a trial and error process in which the FID from the water phantom was used as an indicator. NTNMR provides a numerical indicator called 'Shim Units' that

Table 4.2: List of the shims used to improve the homogeneity of the overall magnetic field. Function indicates the spatial dependence of the z-component of the small shim fields that are superimposed on $\mathbf{B}_0$. The x- and y-components of these shim fields are assumed to have negligible effect on the overall field, which is true as long as B_0 is large. Gradient Order is the largest exponent in the function describing the field. Interaction Order indicates the relative sequence in which the shims must be adjusted. When a higher interaction order shim is adjusted, all of the shims of lower interaction order must subsequently be adjusted in order to optimize field homogeneity.

Common Shim Name	Function	Gradient Order	Interaction Order
X	x	1	0
Y	y	1	0
Z	z	1	0
Z^2	$2z^2 - (x^2 + y^2)$	2	1
ZX	zx	2	2
ZY	zy	2	2
2XY	xy	2	1
X^2-Y^2	$x^2 - y^2$	2	1
Z^3	$z[2z^2 - 3(x^2 + y^2)]$	3	2
Z^4	$8z^2[z^2 - 3(x^2 + y^2)] + 3(x^2 + y^2)^2$	4	2

is the sum of the magnitude points in the data. This is essentially a numerical integral of the acquired signal magnitude. A larger number indicates that the initial magnitude of the FID is larger, or that the decay time is longer, or both, all of which indicate an improvement in average field homogeneity regardless of the line shape. First the x, y, and z shims were optimized. One of the second-order shims was then adjusted and the x, y, and z shims were re-optimized in accordance with the interaction order outlined in Table 4.2. If an improvement was observed, this procedure was then repeated for each of the same higher order shim until no further gains were made. This procedure was then repeated for each of the higher-order shims until no further improvement seemed possible. After the initial round of shimming with the water phantom at 64 MHZ the ^{3}He signal was found

at 48.6 MHz using the birdcage coil and additional shimming adjustments were made to optimize the signal coming from the gas volume. At this point, the linewidth of the ^{3}He signal in the frequency domain seemed limited to a real component FWHM of 30 Hz.

To make further improvements in B_0 field homogeneity, the orthogonalization and optimization routine developed by Michal was employed [100]. The spatial dependence of the magnetic field gradients used in shimming is based on spherical harmonics. These shim fields should ideally be orthogonal to one another with minimal interaction. That is, the total field in the vicinity of the sample is of the form

$$\mathbf{B} = \mathbf{B}_0 + \sim_{i=0}^{\infty} a_i g_i \qquad (4.12)$$

where the functions g_i describe the spatial dependence of the shim fields. In practice, the shim fields are not perfectly orthogonal. This situation is exacerbated by the magnetic susceptibility of the sample which can alter the field profile of both B_0 and the shims themselves. The procedure presented by Michal produces a set of composite shim fields that are orthogonal. A composite shim is composed of a number of shims whose control currents are adjusted proportionately. The procedure implemented in this work, which differs in some respects to the original prescription, is presented below.

The technique is based on numerically integrating

$$I = \int S(\omega)\omega^2 \mathrm{d}\omega \qquad (4.13)$$

where $S(\omega)$ is the line shape data from the Fourier transform of the FID. If the maximum is located at $S(0)$, this integral gives the second moment of the spectrum. I varies parabolically with the strength of the applied shim field. As the current flowing through a shim coil moves away from its optimum value, the width of the spectrum increases. The weighting factor ω^2 in the integrand increases the magnitude of I for a broader spectrum. The bottom of the parabola formed by plotting I as a function of shim current represents the optimum shim value for any one shim control. In practice only three spectra yielding three measurements of I are required. One spectrum is acquired at an initial shim value d_0 producing a second moment I^0. Two more spectra are acquired with the shim current increased and decreased by a symmetric amount d relative to d_0 yielding second moments I^+ (for $d + d_0$) and I^- (for $d - d_0$). The new optimum shim setting is then

$$a = \frac{I^- - I^+}{4dA} \qquad (4.14)$$

where A is the quadratic coefficient given by

$$A = \frac{I^+ + I^- - 2I^0}{4d^2} \quad . \tag{4.15}$$

The next shim field that is to be adjusted can be orthogonalized relative to the previous shim using the Gram-Schmidt orthonormalization procedure [101]. This is done by applying currents to two shim coils simultaneously with offsets d_i and d_j and calculating the second moment for the resulting spectra on the basis of Eq. (4.13). Once the four second moment integrals have been calculated, the interaction coefficient of the gradients is given by

$$C_{ij}(i \neq j) = \frac{I^{++} - I^{+-} - I^{-+} + I^{--}}{8d_i d_j} \tag{4.16}$$

where the plus and minus signs indicate the polarity of the current offsets d_i and d_j applied to shims i and j. For example, the spectrum for the I^{++} term is acquired with both shims set to the high current setting. The orthogonal composite shim is then given by

$$g' = g_j - C_{ij}g_i \quad . \tag{4.17}$$

This composite shim can then be optimized by acquiring spectra with symmetric current as described above for a single shim using Eqs. (4.13), (4.14), and (4.15). A third shim can then be orthonormalized to the composite shim, the new composite shim that is created can be optimized, and so on, until all of the available shim controls have been adjusted. In principle, one complete cycle through all of the available shim controls should produce an optimally-homogeneous field but in practice further improvements can be achieved with two or more passes. After applying several iterations of this method the ^{3}He line width in the acoustic flow tube was reduced from a FWHM of 30 Hz to about 5 Hz. A sample spectrum is shown in Fig. 4.16.

4.8 Pulsed Magnetic Field Gradients

MRI pulse sequences require the application of linear magnetic field gradient pulses in order to measure the spatial distribution of nuclear spins in the sample. The shape and timing of magnetic field gradient pulses are set in NTNMR using a "shape table" and a parameter defining the duration of the pulse τ. The shape table consists of a list of n numerical values

Figure 4.16: Typical ^{3}He FID lineshape at 48.6 MHz after performing the Michal shimming procedure. The line has a FWHM of 5.1 Hz.

that range from -100 to 100 and that are proportional to the amplitude of the gradient pulse. The time between points is τ/n. Another parameter called the "gradient amplitude", which ranges between -100 and 100, is used to specify the total amplitude of the pulse.

An Oxford Instruments 265/235 Gradient Coil System was inserted into the bore of the room temperature shim cylinder. The x, y, and z linear gradient set of this insert were used in this work. The inner diameter of this insert is 23.5 cm and the x, y, and z gradient coils are nominally calibrated to produce 0.12174 G/cm/A, 0.11431 G/cm/A, and 0.11511 G/cm/A respectively. The location of the isocenter of the magnetic field gradients is also specified relative to the ends of the former by the manufacturer. The Techron 8607 Gradient Amplifiers used to drive currents through the gradient coils each have a current monitor output with a calibration of 20 A/V. In the absence of induced eddy currents, a measurement of the current monitor output as a function of time could be used to calculate gradient strengths. In practice, this at best provides a rough estimate. An accurate in situ measurement of the magnetic field gradients produced at the location of the gas sample was performed using a set of custom gradiometer coils. The MR imaging of a water phantom of known geometry was used to calibrate gradient field strengths.

As magnetic field gradients are switched on and off (pulsed) in the magnet, it is possi-

ble for eddy currents in conducting materials located near the coils to build up and decay. These eddy currents alter the gradient fields to which the sample is exposed. Consequently, the strength and time dependence of applied field gradients cannot always be directly inferred from a measurement of applied currents. The conducting structures that surround the linear gradient coil set include the cylindrical metal bore of the magnet itself, the room temperature shim coil set, and a copper RF Faraday shield that is situated on the inner surface of the gradient coil insert. The room temperature gradient shim set of the MRI magnet includes x, y and z linear gradient coils that were not used, and whose current leads were left unconnected. The strength of induced eddy currents depends on the time rate of change of the magnetic field $\mathrm{d}B/\mathrm{d}t$. A magnetic field gradient pulse with a square envelope will induce large eddy currents because of the large $\mathrm{d}B/\mathrm{d}t$ at the leading and trailing edges of the pulse. From this perspective, it is desirable to apply long, low amplitude pulses. To acquire images rapidly and to improve SNR by reducing effects related to T_2 dephasing, it is necessary to shorten the duration of gradient pulses which in turn increases the amplitude of induced eddy currents. A compromise between these two goals must be reached.

The shape of the gradient pulse that is used in a pulse sequence depends on its intended function (see Sec. 2.1). The two field gradient shapes employed in the VPE-MRI sequence that will be described in Sec. 5.1 are the trapezoid and the half-sine. One method for reducing eddy currents relative to those induced by square pulses is to employ pulse shapes that expose conducting surfaces to lower values of $\mathrm{d}B/\mathrm{d}t$. For example, frequency encoding along the x-axis during data acquisition is often accomplished using a trapezoidal field gradient envelope. A linear ramp brings the amplitude of the gradient up to the desired value, at which point the current is held constant. After the signal has been acquired from the sample another linear ramp brings the amplitude of the field gradient back to zero. Any eddy currents that circulate during the acquisition phase will alter the nuclear spin precession and create image artifacts. To reduce the influence of eddy currents on precession, the $\mathrm{d}B/\mathrm{d}t$ of the ramp can be reduced simply by increasing the amount of time taken to reach the maximum. The role of read-rewind, phase-encoding, and velocity-phase-encoding field gradients is to induce a net phase shift in the precession of the nuclear magnetization. Therefore, it is the time integral of the pulse that is of interest. A variety of shapes are used to generate these types of pulses; in this work the half-sine field gradient pulse shape was used. This gradient pulse mimics the positive lobe of the sine function. Although $\mathrm{d}B/\mathrm{d}t$

is large near the leading and trailing edges of the pulse, it quickly decreases in magnitude with time and is zero when the gradient is at its peak amplitude. The peak rate of change is significantly less than a square pulse of similar area and duration.

Figure 4.17: Longitudinal and transverse gradiometer coils. The shaded areas represent the locations of the coils. The coils are connected in series and counterwound so that the voltage across either pair is the difference in induced EMF. The parameter s is the distance between the coils, which is used in Eq (4.18) to calculate the gradient strength.

Two gradiometer coil sets were used to measure the magnetic field gradients. One set of coils was used to calibrate the z-gradient and the other set was used to calibrate the x- and y-gradients. The physical arrangement of the coils in these gradiometers is shown in Fig. 4.17. The longitudinal gradiometer consists of two identical coaxial coils. The coils each have $N = 65$ turns of wire, an area $A = 4.52(5)$ cm^2, and they are separated by a distance $s = 3.54(5)$ cm. When the z-axis magnetic field gradient is pulsed the EMF $\mathcal{E} = NA\mathrm{d}B/\mathrm{d}t$ induced in the two coils will be different, because of the separation that exists between them. If the coils are connected in series but counterwound with respect to one another then the voltage measured across the circuit will be the difference $\Delta\mathcal{E}$ between the two EMFs. From a measurement of $\Delta\mathcal{E}$ the field gradient amplitude can be calculated:

$$G = \int \frac{\Delta\mathcal{E}}{NAs}dt \quad .$$

(4.18)

The transverse gradiometer operates in a similar manner. Its two coils are mounted so that they lie in the same plane. The center to center separation of the two coils is $s = 5.52(5)$ cm; they are otherwise identical to the longitudinal gradiometer coils.

These gradiometers were used to measure the time dependent amplitude of the magnetic field gradient pulses employed in the VPE-MRI sequence shown in Fig. 5.1. A special NT-

NMR sequence was written to repeat a gradient pulse 128 times, using the same shape table, timing, and field gradient amplitude as used in the imaging sequence itself. Both trapezoid and half-sine pulse shapes were characterized for the x field gradient. The transverse gradiometer output was connected to a PAR 113 pre-amp to amplify the net EMF across the coils. This pre-amp has a 100 MΩ input impedance and provided a gain of 100. The pre-amp output was monitored using a Tektronix TDS 7054 DPO digital oscilloscope. After averaging, the oscilloscope was used to numerically integrate the EMF and display the result as a function of time. This integrated signal is proportional to the field gradient strength. The raw data were also transferred to a computer where they were numerically integrated and analyzed using the Origin 8 SRV v8.0951 software package.

Figure 4.18: Ringing in the z field gradient, observed in response to a 1 ms duration half sine pulse.

First, the spatial linearity of the magnetic field gradients was studied. While one of the gradients was pulsed repeatedly and the response was monitored on the oscilloscope, the gradiometer was translated through the imaging region near the isocenter of the field gradient coils. The x and y field gradient maps revealed deviations from linearity of less than 1% within a 9 cm radius of the isocenter. As soon as the first observations of the z field gradient were made it was evident that a decaying oscillation occurred at the end

of each pulse. This was surprising, because the x and y field gradients did not exhibit this behaviour. An example of these oscillations can be seen in Fig. 4.18. The ringing looks like an exponentially-decaying sinusoid with a frequency of 1.1085(3) kHz and a time constant of 1.318(3) ms. The phase and amplitude of the ringing varied as the gradiometer was translated along the z-axis, but the phase remained constant for any z-position as the gradiometer was moved in the x-y plane. The exact origin of this ringing is unknown.[4]

One method for dealing with eddy currents is to modify the shape of the current pulse sent to the amplifiers so as to achieve the desired magnetic field gradient pulse shape at the sample location. The modification of the current pulse shape is known as preemphasis. Typically the build up and decay of eddy currents are exponential in time, but multiple time constants can be involved depending on the number and type of eddy current sources. NTNMR has a built in preemphasis function that can be used to modify the field gradient pulse shape tables before the waveform is sent to the amplifiers. For a gradient pulse that is described in time by $G(t)$, the preemphasis function produces an output

$$G_{\text{out}}(t) = G(t) \left[A_0 + A_1 e^{-t/\tau_1} + A_2 e^{-t/\tau_2} + A_3 e^{-t/\tau_3} \right] \quad . \tag{4.19}$$

The NTNMR preemphasis function was applied to both the x and y field gradients. Because the acoustic flow tube obstructs the axis of the magnet the center of the transverse gradiometer was positioned 5 cm below and 4 cm to the side (along +x) of the tube, and was aligned with the center of the z field gradient. Preemphasis adjustments were performed while observing gradient pulses identical to those used in the VPE-MRI pulse sequence that will be described in Sec. 5.1. The rewind pulse for the x field gradient, the phase encoding pulse along y, and each lobe of the VPE gradient along z all have a half-sine shape and a 1 ms duration. The field gradient pulse shape was observed using the numerically integrated gradiometer signal produced by the oscilloscope. The dominant time constant was estimated by finding the time interval required for the tail of the signal to drop by a factor of e^{-1} relative to its amplitude at the end of the applied pulse. The amplitude parameter associated with this time constant was then adjusted by trial and error until the

---

[4]The origin of the ringing was unknown as of the date this thesis was sent to the external examiner. Additional experiments were performed between the thesis submission and examination dates. These experiments revealed that the observed ringing is related to acoustic modes in the acoustic flow tube and not due to eddy currents. That work is described in an addendum to the thesis in Appendix D. The discussion of the ringing in terms of the eddy current hypothesis is retained throughout the thesis.

(a) Without preemphasis
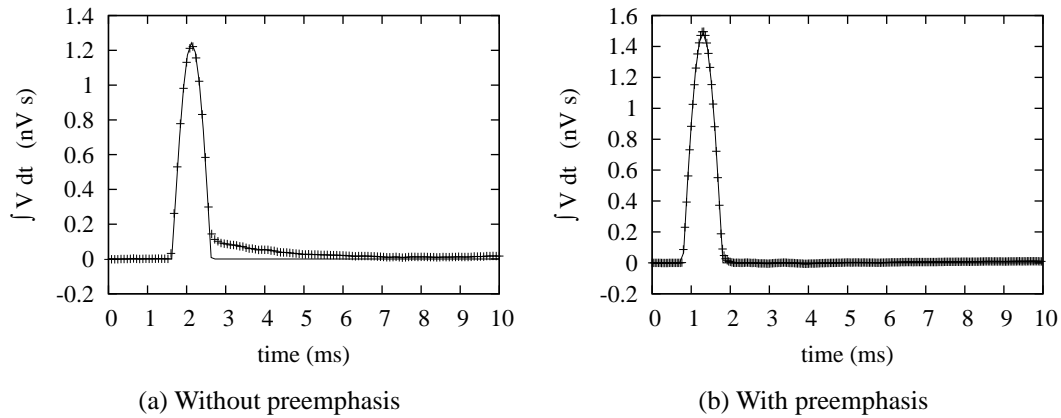
(b) With preemphasis

Figure 4.19: Measured time integral of the gradiometer voltage during a y magnetic field gradient pulse (a) without preemphasis and (b) with preemphasis. The solid lines show the ideal shape of the 1 ms half-sine pulse.

lineshape could no longer be improved. This process was then repeated for two more time constants. An example illustrating the improvement made to the y magnetic field gradient is shown in Fig. 4.19. The field gradient pulse produced using preemphasis compensation is indiscernible from the ideal pulse shape represent by the solid line.

The z magnetic field gradient required special treatment to compensate for the oscillations that occur at the end of the half-sine pulse. A special shape table was created that counteracts the ringing as well as the exponentially decaying eddy currents. This shape table was created using the Origin software package. The voltage signal from the gradiometer was transferred from the oscilloscope to a computer and then loaded into Origin. The data were then numerically integrated and scaled so that the maximum amplitude had a value of 100, corresponding to the maximum input value allowed in a shape table. The idealized half-sine shape was then subtracted from the signal leaving a residual associated with eddy currents. A function of the form

$$S(t) = A_1 e^{-t/\tau_1} \sin \omega t + \phi + A_0 e^{-t/\tau_0} \tag{4.20}$$

was fit to the oscillating residual. The parameters A_1 and τ_1 represent the initial amplitude and time constant of the decaying oscillation with frequency ω and phase ϕ, and the parameters A_0 and τ_0 represent the initial amplitude and time constant of the dominant

exponentially-damped eddy current. The phase parameter in Eq. (4.20) was then shifted by 90° to create a decaying oscillation that leads the induced oscillation. This phase-shifted function was then added to the idealized half-sine shape table. The magnetic field gradient pulse that was produced using this new shape table was then measured using the gradiometer and the process repeated. In Fig. 4.18 the amplitude of the ringing component of the signal is 15% of the maximum amplitude of the ideal half-sine pulse and has a time constant of 1.3 ms. The shape table used in the pulse sequence and a bipolar field gradient pulse measured using the gradiometer are shown in Fig. 4.20. With the compensated half-sine shape table the maximum deviation after the end of the pulse as measured using the gradiometer is reduced to approximately 2% of the maximum pulse amplitude.

(a) Compensated half-sine shape table   (b) Compensated bipolar field gradient pulse
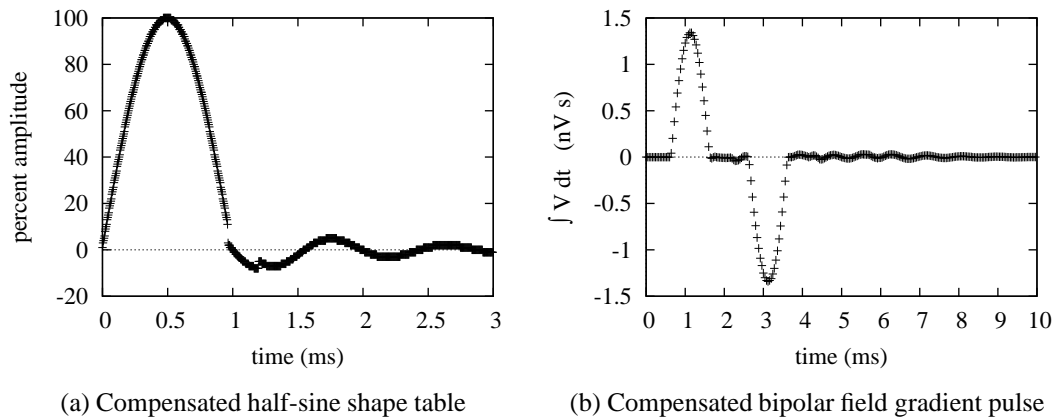
Figure 4.20: (a) Plot of the shape table used to null the ringing in the half-sine z field gradient pulse. (b) The measured time integral of the gradiometer voltage during z magnetic field gradient pulses.

4.9 Calibration of Magnetic Field Gradients

Quantitative analysis of VPE-MRI experiments requires accurate calibration of the magnetic field gradients used in the pulse sequences. Although the calibration of the amplitude of the magnetic field gradients could be performed using the gradiometer along with Eq. (4.18), MRI gives a more direct measure of how the field gradients affect the images

Figure 4.21: The inner dimensions of the water phantom used in calibrating the magnetic field gradients.

that are acquired. This includes the scaling of the image as well as any image artifacts due to inhomogeneities in $\mathbf{B}_0$. After the VPE-MRI experiments were completed the gas mixture was removed from the acoustic flow tube and the AMR was disassembled to allow unfettered access to the center of the MRI magnet.

The images used in the calibration of the magnetic field gradients were acquired using a phantom filled with an 8 μmol/cm^3 Cu-SO$_4$-water solution inside of a 5.0 cm diameter, 5.0 cm long tube resonator tuned to 63.57 MHz. The phantom inner dimensions are shown in Fig. 4.21. The imaging sequence was the same spin echo sequence as was used for ^{3}He, and will be described in Sec. 5.1. The read-rewind and phase-encoding field gradient pulses involved the same 1 ms half-sine shape. The other timings were varied to produce an echo time T_E, as shown in Fig. 5.1, of 52.4 ms. The mapping of k-space is 256 by 256 points with a SNR at the center of the echo with $k_y = 0$ of approximately 2500. No slice selection was used in acquiring the images which are therefore two dimensional projections of the water distribution in the phantom. The read gradient was applied along the x-axis, which is in the horizontal direction perpendicular to the bore of the magnet. Two images were acquired: one with the phase encoding gradient applied along the y-axis and the other with the phase encoding in the z-direction, which is parallel to the bore of the magnet. The FOV in the read and phase encoding directions were calculated based on the size of the water phantom in the images. The magnetic field gradient strengths were then calculated using Eqs. (2.28) and (2.29).

The x-y projection image is shown in Fig. 4.22a. As can be seen there are some imaging artifacts in the phase encoding (vertical) direction. The distortion is primarily a skew effect

(a) x-y projection

(b) x-z projection

Figure 4.22: Projection image of the water phantom in the x-y plane with 256×256 points. The dimensions of the phantom in the image are given in points, which can be scaled to find the image FOV from the physical dimensions of the phantom specified in Fig. 4.21.

where the top of the image is shifted to the left relative to the bottom, parallel to the x-axis. This can be corrected by a shear mapping algorithm. The x-z projection image shown in Fig. 4.22b reveals further evidence of the skew image artifact in the x-y plane. This can be seen as a shift of the smaller diameter top of the the phantom to the left relative to the bottom. This is consistent with the shift seen in the x-y image. These data reveal the nature of residual inhomogeneities in the B_0 magnetic field.

Chapter 5

VPE-MRI of Acoustic Motion in a Gas

The motion of gas in the acoustic viscous boundary layer has been studied since the mid-1860s [67, 25, 24]. The publication of Rott's theory in 1969 [27], explaining the rich diversity of thermoacoustic phenomena, brought about a resurgence of research around this topic. Lately there has been another push to develop imaging modalities capable of capturing details of gas motion under conditions found in thermoacoustic devices. These techniques, which were briefly noted in Ch. 1, include hot-wire anemometry [42], laser-Doppler anemometry [43, 44], and particle-image velocimetry [43].

MRI has been used previously to acquire 2D images of thermally polarized ^{3}He mixed with O_2 [51] in a rat lung. It has also been used in conjunction with hyperpolarized ^{3}He to produce velocity maps reflecting – amongst other things – the steady flow of gas in a tube and the inspiration of gas into the lungs of a human subject [15]. In this chapter I will describe time- and spatially-resolved measurements of gas velocities in the viscous boundary layer (AVBL) obtained using VPE-MRI on thermally polarized ^{3}He mixed with O_2. The apparatus used to perform these experiments was described in Chs. 3 and 4.

The coordinate system used in this chapter follows the MRI convention in which slice selection is performed along the z-axis, phase encoding is along the y-axis, and frequency encoding is along the x-axis. The acoustic flow tube in the imaging region is therefore parallel to the z-axis in this scheme, and the velocity of the gas will be a function of x and y. This contrasts with the convention used in thermoacoustic theory (e.g. Sec. 2.2) which has the acoustic wave propagating in the x-direction.

5.1 Pulse Sequence

There are a seemingly limitless number of pulse sequences used in MRI, each with its own advantages and disadvantages. The VPE-MRI pulse sequence used in this work to produce gas velocity maps is pictured in Fig. 5.1. It is a variation of the spin echo imaging sequence outlined in Sec. 2.1.9. The application of a bipolar VPE pulse, shown as a dotted line in the figure, introduces a velocity-dependent phase shift in the precession of the ^{3}He nuclear spin, as described in Sec. 2.1.10. The images produced by this sequence are 2D snapshots of the gas velocity.

Figure 5.1: VPE-MR imaging sequence used in this work. The external trigger is used for synchronizing image acquisition with gas motion. The delay T_ϕ sets the phase of the image relative to the acoustic cycle of the gas. The echo time is shown as T_E, and is twice the time between the two RF pulses. The bipolar VPE field gradient pulse is indicated by a dotted line. The width τ of the pulse lobes and the time Δ between the leading edges of the pulses, along with the maximum amplitude of the field gradient, determine the FOS given by Eqs. (2.33) and (2.34).

Note that slice selection is not used in this pulse sequence. Slice selection would lead to a reduction of signal amplitude both because it would involve excitation of a smaller

volume of spins and also because of irreversible losses (attenuation) associated with the diffusion of ^{3}He during application of the necessary (additional) field gradients. Instead, there is an effective slice selection imposed by the axial extent of the RF fields produced by the birdcage resonator (about 4 cm; see Fig. 4.14). The alignment of the acoustic flow tube in the imaging region, with the geometric axis of the magnet, was measured to be within $1°$. The motion of gas in this tube is symmetric about the center of the magnet and has very little z-dependence, as confirmed by modeling. It will thus simply be assumed that the velocity is uniform along the axial extent of each voxel. Based on the Womerseley number and Reynolds number calculated using the peak gas velocity during each cycle of the oscillation (Table 4.1), the flow inside this tube should be laminar [97]. Consequently, there should be no coherent motion of the gas in the transverse (x and y) directions and a measurement of the z-component of the velocity alone should suffice.

The imaging sequence begins with a 1 ms 5-lobe truncated sinc $\pi/2$ RF pulse that tips the spins into the transverse plane. Next, read-rewind and phase-encoding magnetic field gradient pulses are applied, along with the first lobe of the bipolar velocity-phase-encoding (VPE) field gradient pulse. These gradients are applied simultaneously to avoid unnecessary dephasing by intrinsic T_2. The field gradient pulses have a half-sine shape and a duration $\tau = 1$ ms. The half-sine shape was chosen to limit the excitation of eddy currents, as discussed previously in Sec. 4.8. The negative lobe of the bipolar VPE pulse is applied after a time delay Δ, as measured from the rising edge of the first lobe of the VPE pulse. Next, a 1 ms 5-lobe sinc π RF pulse prepares the spin echo by inverting the phase of the nuclear magnetization. The read field gradient is then ramped up, again to limit the excitation of eddy currents, and the spin echo is recorded during the constant plateau. The relative strength of the read-rewind field gradient pulse and the read gradient are chosen so as to form an echo at the same time as the spin echo caused by the π RF pulse at time T_E. This sequence is repeated with different phase encoding field gradient strengths to produce a two dimensional map of k-space. The advantage of using this sequence over a gradient-recalled echo sequence is that some residual B_0 field inhomogeneities are refocused.

The full imaging sequence is performed twice. One image is acquired precisely as described above and, in this thesis, is referred to as the VPE image. A second image is taken with the bipolar pulses turned off. This provides a phase reference and is thus called the reference image. The phase of the complex data from the reference image is subtracted

from the phase of the corresponding data from the VPE image. The velocity of the gas is proportional to the resulting phase difference by a factor of FOS/2π, where the FOS is found using Eqs. (2.33) and (2.34). Recall that the image field of speed (FOS) is determined by the pulse time τ, the delay time Δ, and the amplitude of the VPE pulse. The attenuation of the signal caused by diffusion during the bipolar field gradient pulse is described by Eq. (2.36).

The initiation of the pulse sequence is synchronized with the motion of the gas in the acoustic flow tube using the external trigger input of the Apollo console. The synchronization signal is generated by a photodiode emitter-receiver pair that is periodically interrupted by a small piece of copper foil attached to the piston rod. This signal alternates between 0 V and 5 V with a period equal to that of the piston. The rising edge of the photodiode signal triggers a Stanford Research Systems model DS345 function and arbitrary waveform generator that is programmed to output a 1 ms TTL pulse. The waveform generator output is fed to the Apollo console external trigger input where it is used to launch the pulse sequence. It was found that the waveform generator was triggered more predictably by the photodiode signal than was the spectrometer. After the synchronization pulse is received there is a delay T_ϕ before the initial $\pi/2$ RF pulse is generated. This sets the phase of the acoustic cycle of the gas at which the image is taken. A delay was also added to the end of the entire sequence so that it would be repeated every second acoustic cycle. The relative phase between the gas motion and the photodiode signal depends on the frequency of the AMR drive, as discussed in Sec. 4.5.

5.2 Uncertainty in the Measured Velocity

The heterodyne detection scheme used in this experiment provides complex data having real and imaginary components. To compare velocity images with theory, it is necessary to estimate the uncertainty in the data. The velocity in each pixel of the image is calculated from the phase difference between the VPE image and the reference image, multiplied by a factor of FOS/2π. The uncertainty in the velocity will therefore depend on the uncertainty in the phase of the images.

The phase at each point in the image is

$$\phi = \tan^{-1}\left(\frac{I_{re}}{I_{im}}\right) \tag{5.1}$$

where I_{re} and I_{im} are the real and imaginary pixel intensities, respectively. Assuming white noise, the uncertainty in the data is characterized by the standard deviation of fluctuations in the real $\sigma(I_{re})$ and imaginary $\sigma(I_{im})$ channels acquired when no signal is present. An examination of the noise in each channel of the spectrometer used in this work reveals that they are independent and well-balanced (i.e. exhibit the same noise level). The noise can therefore be characterized by a single standard deviation $\sigma(I_0) = \sigma(I_{re}) = \sigma(I_{im})$. By a propagation of errors calculation the uncertainty in the phase calculated using Eq. (5.1) is [102]

$$\sigma(\phi) = \frac{\sigma(I_0)}{|I|} \tag{5.2}$$

where $|I|$ is the magnitude of the signal. The uncertainty in the measured velocity will then be

$$\sigma(u_1) = \frac{\text{FOS}}{2\pi}\sqrt{\sigma(\phi_{\text{ref}})^2 + \sigma(\phi_{\text{VPE}})^2} \tag{5.3}$$

where $\sigma(\phi_{\text{ref}})$ and $\sigma(\phi_{\text{VPE}})$ represent standard deviations in the phase noise of the reference and VPE images, respectively. The magnitude of the signal in the VPE images is much smaller than that of the reference images because of the signal attenuation caused by ^{3}He diffusion during the bipolar field gradient pulse, but the standard deviation of the amplitude noise in both images is equal. Consequently, the noise in the phase signal will be higher in the VPE images than in the reference images. The standard deviation of the amplitude noise $\sigma(I_0)$ was calculated for each image using data that comes from far away from the ^{3}He signal (i.e. outside of the acoustic flow tube). These measurements confirmed that $\sigma(I_{re}) = \sigma(I_{im})$. Typical velocity uncertainties range from 2 to 5 cm/s while peak measured velocities are as high as 30 cm/s.

5.3 Preliminary Velocity Maps

The goal of this study was to measure the velocity of gas in the AVBL as a function of transverse position and time (i.e. $u_1(x, y, t)$). The time dependence here is sinusoidal and

can be separated from the spatial dependence such that

$$u_1(x, y, t) = u_1(x, y)e^{i\omega_A t + \phi_0} \tag{5.4}$$

where ω_A is the angular frequency of the oscillatory motion and ϕ_0 is the initial phase. The spatial dependence of the gas velocity in a cylindrical geometry depends on the radial distance from the tube axis $r = \sqrt{x^2 + y^2}$, and is given by Eqs. (2.70), (2.71), and (2.72). This function is complex, and involves a phase shift of the gas motion near the cylinder wall relative to that along the axis. It is the real component of Eq. (5.4) that is reflected in the measured velocity maps. Each map represents a snapshot at one phase $\phi_A(t) = \omega_A t + \phi_0$ in the acoustic cycle. Here, the phase ϕ_0 is measured relative to the photodiode trigger used to synchronize the pulse sequence with the AMR drive. The time dependence of the velocity can therefore be studied by acquiring images at different delay times spanning one period of the acoustic cycle.

5.3.1 Data Acquisition

The first complete set of velocity phase encoded (VPE) images were taken over a 7 day period during which data were acquired on 5 out of the 7 days. Starting each day, the AMR was set in motion at a frequency of 1.31 Hz (period of 763 ms). The photodiode trigger and the piston U-tube mercury displacement signals were monitored on an oscilloscope. The drive frequency tended to vary slowly with time. Any change in this frequency would cause changes in the trigger period, the mercury displacement amplitude, and the relative phase between the two traces. Small manual adjustments were made to the frequency control to maintain the drive at 1.31 Hz. Before launching the imaging sequence, the birdcage coil was tuned and the π RF pulse was calibrated by adjusting the power level until a null signal was observed in the FID. The $\pi/2$ RF pulse amplitude was then set by looking for the maximum FID signal. The first image that was acquired each day had the velocity-phase-encoding pulses turned off so as to produce a reference image. This was followed by a succession of velocity-phase-encoded images taken at 40 ms delay increments. The reference image and the VPE images were Fourier transformed and the phase of the signal in each pixel was calculated. The phase of the reference image was subtracted, pixel-by-pixel, from each of the VPE images. The resulting datasets were multiplied by the VPE

scaling factor FOS/2π to produce velocity maps. The uncertainty in the velocity associated with each pixel was calculated using Eq. (5.3). This process was repeated each day until velocity maps spanning the acoustic cycle had been acquired.

Table 5.1: VPE-MRI pulse sequence parameters and the data acquisition order for the preliminary data set. N_p is the number of phase encoding points, N_r the number of read points, and N_A the number of averages. The delay T_ϕ is the time between the AMR trigger and the $\pi/2$ pulse.

(a) Imaging parameters

parameter	value
T_E	10.0 ms
T_π	6.1 ms
τ	1.0 ms
Δ	3.0 ms
G_p max	0.419 G/cm
N_p	32
G_r	0.10 G/cm
G_{VPE} max	1.67 G/cm
N_r	128
T_R	1.527 s
N_A	64
Δf	$\pm$2500 Hz

(b) Acquisition order

Day	$T_\phi(ms)$
1st	317, 357, 397, 437, 477
2nd	517, 557, 597
3rd	637, 677, 717, 277
4th	237
7th	197, 157, 117, 77, 37

The measured image delay times and the order in which data were acquired, along with the imaging parameters used for this set of experiments, are summarized in Table 5.1. The resulting pixel dimensions are 0.121(3) cm, along the frequency-encoding direction (x), by 0.58(2) cm, along the phase encoding-direction (y). The FOS was 96(3) cm/s mapping phase shifts $[-\pi, \pi]$ to [-48,48] cm/s. The uncertainties in pixel dimensions and the FOS are based on the calibration of the field gradients as described in Sec. 4.9. The image acquisition time was 52 minutes.

Sample images reflecting the magnitude (as opposed to the phase) of reference and VPE

(a) Reference   (b) VPE
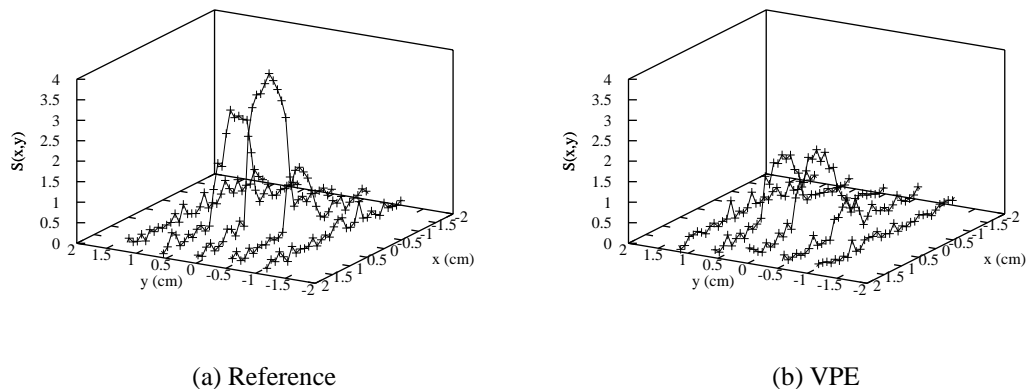
Figure 5.2: Sample images reflecting the magnitude of the detected ^{3}He precession signal for reference and VPE datasets. The FOV has been reduced to show details near the acoustic flow tube. The magnitude of the VPE field gradient pulse was 1.67 G/cm.

datasets are shown in Fig. 5.2. The average SNR of this reference image is 21 and that of the VPE image is 11. The decrease in the VPE image SNR relative to the reference image is a consequence of ^{3}He diffusion while the (bipolar) VPE field gradient is applied, which results in an irreversible loss of coherent spin precession.

5.3.2   Analysis

The measured velocity maps were compared to the theoretical prediction for a $1.30(5)$ cm diameter cylindrical tube with an acoustic frequency of 1.31 Hz and a volume flow rate $U_1 = 26.5$ cm^3/s. The viscous penetration depth was calculated to be $\delta_\nu = 0.15(2)$ cm for the mixture of 1.5 atm ^{3}He and 1.5 atm O$_2$ that was used in the experiments. The mixture composition was determined from a measurement of the ^{3}He nuclear relaxation time T_1 performed on the first day of the data run.

The theoretical velocity profile $u_1(r)$ is given by Eqs. (2.70), (2.71), and (2.72). A program (Gnuplot v. 4.4) using an implementation of the non-linear least-squares Marquardt-Levenberg curve fitting algorithm was used to fit the theoretical expression for the velocity profile as a function of time $u_1(r)\exp(i(\omega_A t + \phi_0))$ to the data. The Bessel functions $J_0(x)$ and $J_1(x)$ in Eqs. (2.71) and (2.72) have complex arguments but the built-in Bessel func-

tions used in the fitting program could only handle real arguments. Instead, user-defined functions were employed based on the Taylor series expansions of the complex Bessel functions around $x = 0$ given by [103]

$$J_0^{\text{fit}}(r) = \sum_{m=0}^{10} \frac{(-1)^m}{(m!)^2} \left(\frac{(i-1)r/\delta_\nu}{2} \right)^{2m} \tag{5.5}$$

$$J_1^{\text{fit}}(r) = \sum_{m=0}^{10} \frac{(-1)^m}{m!(m+1)!} \left(\frac{(i-1)r/\delta_\nu}{2} \right)^{2m+1} . \tag{5.6}$$

The complete fitting function used the following equations:

$$h_\nu^{\text{fit}}(r) = \frac{J_0^{\text{fit}}(r)}{J_0^{\text{fit}}(R)} \tag{5.7}$$

$$f_\nu^{\text{fit}} = \frac{2 J_1^{\text{fit}}(R)}{J_0^{\text{fit}}(R)(i-1)R/\delta_\nu} \tag{5.8}$$

$$u_1^{\text{fit}}(r, \phi) = \Re \left(\frac{1 - h_\nu^{\text{fit}}(r)}{1 - f_\nu^{\text{fit}}} \langle u_1 \rangle e^{i\phi} \right) \tag{5.9}$$

where R is the radius of the cylinder and $\langle u_1 \rangle = U_1/A$ is the spatial average of the velocity amplitude, which is equal to the volume flow rate divided by the cylinder area A. The phase of the acoustic cycle is

$$\phi = \omega_A t + \phi_0 \tag{5.10}$$

where $\omega_A = 2\pi f_A$ is the angular frequency of the acoustic oscillation and ϕ_0 is the initial phase. The fitting function Eq. (5.9) as well as the real part of Eq. (5.4), as calculated using GNU Octave v.3.2.3, are plotted in Fig. 5.3 for $\phi = 0$ and $R/\delta_\nu = 0.65/0.15 = 4.3$. Also plotted is the residual, which is calculated by taking the difference between the two functions. The largest deviations are of order one part in 10^6 of the peak velocity. The smallest relative uncertainties in the velocity measurements are typically around 10%; therefore the truncated series given above are more than sufficient for the fitting process. The computation times for these fits are only a few seconds so there is little advantage to using fewer terms.

The pixel dimensions for the velocity maps are 0.12 cm along the x-direction and 0.58 cm along the y-direction. The radius of the acoustic flow tube is $R = 0.65$ cm. A sketch showing an overlay of the data grid on the sample is shown in Fig. 5.4. The alignment of the grid with the sample is such that one line of data along the x-axis lies within

Figure 5.3: A plot of $u_1(r/R)/\langle u_1 \rangle$ versus the non-dimensional radius r/R using the approximation given by Eq. (5.9) (solid line) and a calculation performed using Octave (+). Also plotted is the difference between these functions ($\times$), with the scale for the latter shown on the right vertical axis.

one millimeter of the axis of the tube. The adjacent lines of data pass close to the tube wall. They contain weaker signals leading to larger uncertainties in velocities. The initial attempt at fitting Eq. (5.9) to the data was done using only the central line of data along the x-axis to take advantage of the stronger signal. A mean velocity amplitude $\langle u_1 \rangle = 20.0$ cm/s was used in the fit. This value was taken from the mercury displacement measurement at the acoustic frequency $f_A = 1.31(1)$ Hz. The viscous penetration depth $\delta_\nu = 0.15$ cm was calculated for the gas mixture of 1.5 atm ^{3}He and 1.5 atm O_2. This left ϕ as the only free parameter determined in the fit.

Figure 5.5 shows the phase extracted from fits of Eq. (5.9) to the data as a function of the delay time T_ϕ divided by the acoustic period T_A. The velocity maps were acquired in 40 ms delay increments starting from $T_\phi = 37$ ms and ending at 717 ms. From Eq. (5.10) the phase of each velocity map is expected to be $\phi(T_\phi) = 2\pi f_A T_\phi + \phi_0$, which is shown as a solid line in the figure. The initial phase ϕ_0 is therefore the phase of the gas velocity at the time of the trigger signal $T_\phi = 0$. At $f_A = 1.31(1)$ Hz, the DELTAEC model of the AMR predicts the gas volume flow rate phase to be $-1.5(2)$ rad.

Figure 5.4: Arrangement of pixels relative to the 1.30 cm diameter acoustic flow tube for images acquired during the preliminary study. The pixel centers are marked by a dot and the dashed lines mark their boundaries. The lines of pixels in the y-direction are slightly offset from the axis of the cylinder.

Figure 5.5: Phase ϕ extracted from fits of Eq. (5.9) to the velocity maps as a function of delay time T_ϕ divided by the acoustic period T_A. The solid line has a slope of 2π and the y-intercept is taken from the phase of the velocity relative to the AMR drive trigger signal as measured during an independent measurement of the AMR frequency response. The data acquisition order is shown in Table 5.1b.

The data shown in Fig. 5.5 shows two distinct data groupings that parallel the expected

linear relationship (note that they are grouped by day of acquisition, as indicated by symbol type). The data acquired during the first four days show the expected trend in phase, although their absolute values are offset by a substantial factor. There is then a jump in phase of approximately 0.8 radians between the fourth day and the seventh day of the acquisition period. This reflects some kind of uncontrolled systematic change in experimental parameters.

As mentioned in Sec. 5.3.1, the frequency of the acoustic oscillations tended to slowly drift away from the 1.31 Hz nominal target over the course of an acquisition. To maintain the acoustic frequency at a constant value, manual corrections were made to the drive frequency control. The period of the photodiode trigger signal and the amplitude and phase of the piston U-tube mercury displacement transducer output were monitored on an oscilloscope as indicators of the drive frequency. With these manual corrections, the maximum deviation in mercury displacement amplitude and drive frequency were estimated to be 10% and 1%, respectively. The phase of the gas displacement in the acoustic flow tube is shown in Fig. 4.11, as a function of drive frequency. At 1.31 Hz, the rate at which the gas displacement phase (in the acoustic flow tube) changes relative to the drive frequency is $\Delta\phi/\Delta f = -10$ rad Hz^{-1}. The phase of the gas velocity should exhibit the same frequency dependence. A variation in drive frequency by 0.01 Hz thus results in a 0.1 rad shift in the gas velocity phase.

One important parameter that was not monitored during the seven day data acquisition period was the internal pressure of the AMR. Over this time period, gas diffusing through the walls of the apparatus would cause a decrease in pressure each day. Based on the effective leak rate measurement presented in Sec. 4.1 the pressure in the system could have dropped by up to 2% over the seven day span. Also, changes in the ambient temperature of the apparatus would be accompanied by adiabatic pressure changes. On hot sunny days, the room temperature increases throughout the morning, peaking in the early afternoon before cooling again in the evening. Temperature changes of up to 4 C have been observed under these conditions. On cooler days, the ambient temperature is much more stable at 20(1) C. To improve the repeatability of the conditions under which gas motion occurs from day to day, the amount of gas in the AMR must be adjusted to maintain the system at constant internal pressure.

The preliminary round of experiments also showed that the tuning of the birdcage coil

tended to drift, and that these drifts are strongly correlated with temperature. As the tuning of the RF coil shifts away from the Larmor precession frequency, its coupling to the external circuitry degrades and some of the incident RF power is reflected back to the transmitter. This situation results in the production of a weaker B_1 field and hence an imperfect π pulse. This in turn is manifested by the appearance of an additional FID signal at the beginning of the data acquisition period. Note that detuning of the resonator also degrades the amplitude of the detected NMR signal, as viewed by the spectrometer.

The sensitivity of the resonant frequency of the birdcage to changes in temperature was measured to be of order 37 kHz/K (or 770 ppm/K). A likely explanation for this sensitivity is related to the fact that the tuning capacitors between the end ring and the axial rungs are sandwiched between two pieces of PVC that are compressed by a series of nylon bolts. The thermal expansion coefficient for PVC is typically greater than that of nylon, and thus the tendency is for the capacitor plates to be squeezed closer together (raising the resonant frequency) as the temperature increases. Given the nominal thickness of the dielectric (70 μm) and the distance over which unconstrained differential contraction can occur (3 mm), a differential thermal contraction of order 4×10^{-5} K^{-1} is sufficient to produce the observed effect. This value is consistent with typical thermal expansion coefficients for the two materials [104].

Irrespective of the mechanism underlying this temperature sensitivity, the magnitude of the effect highlights a serious deficiency in the design of the experimental apparatus. The loaded Quality factor of the critically coupled coil is of order 440, and so a temperature change of order 1 C is sufficient to shift the resonance by an amount comparable to its half-width-at-half-maximum (HWHM). It is likely that temperature changes of this magnitude were incurred during the acquisition of data. Acquisitions that showed a large FID signal prior to the echo were aborted and the coil was retuned. The shift in the resonant frequency observed prior to retuning was typically of order the HWHM. In addition to its influence on the quality of tipping pulses and the amplitude of the detected signal, changes in the tuning of the coil influence the phase of the NMR signal recorded by the spectrometer. Temperature changes that occur between the acquisition of reference and VPE images (or during the acquisition of any one image) thus introduce systematic errors into the inferred velocity field. No alterations were made to the birdcage between the first- and second-round of experiments; the temperature of the room was simply monitored and experiments were

performed during time periods when it was stable.

5.3.3   Summary

The preliminary measurements indicated that VPE-MRI can be used to image the gas veloc-
ity in the AVBL as a function of space and time. Reducing the size of the pixel dimension
along the y-direction is necessary to enable the mapping of the velocity in two dimensions.
The jump in the phase of the measured velocity relative to the trigger pulse observed be-
tween the last two days of data acquisition was most likely due to a combination of factors.
The pressure in the AMR most certainly decreased significantly over that time span as gas
diffused through the surfaces of the apparatus. This factor, combined with a change in am-
bient temperature, could have shifted the velocity amplitude and phase of the gas in the
acoustic flow tube at the operating frequency of 1.31 Hz by enough to explain the discon-
tinuity shown in Fig. 5.5. The pressure in the AMR must be adjusted each day to ensure
that the drive frequency and the gas velocity amplitude and phase is consistent. Care must
also be take to ensure that the tuning of the birdcage coil is stable. Shifts in the resonant
frequency during or between acquisitions will introduce systematic errors to the velocity
measurements.

5.4   Improved Measurement

The experiments described above are promising. They suggest that a reasonable measure-
ment of gas velocities in the AVBL ought to be possible. With the experience gained from
this initial exercise in hand, I introduced some improvements to the experimental method.
The timing, field gradient pulse strengths and averaging schemes of the pulse sequence
were altered to decrease pixel dimension along the y-axis, while maintaining or improving
the SNR in the velocity maps. A routine to maintain the pressure in the AMR was intro-
duced. Finally, the tuning of the birdcage coil was monitored more carefully to minimize
systematic effects introduced by changes in its resonant frequency.

5.4.1 Data Acquisition

The pulse sequence parameters used in this measurement as well as those of the preliminary experiment are summarized in Table 5.2. The changes to the read-gradient strength, acquisition bandwidth, and phase-encoding gradient pulse resulted in pixels of dimensions 0.18(2) cm × 0.39(3) cm in the x- and y-directions, respectively. The amplitude of the VPE gradient pulse was decreased to give a FOS of 144(3) cm/s mapping phase shifts $[-\pi, \pi]$ to $[-72, 72]$ cm/s. In an effort to improve the SNR, the number of averages N_A was doubled while the number of phase encoding points N_p was cut in half. The total image acquisition time for the revised sequence was 78 min. The gas in the imaging region was removed and the ^{3}He was extracted. The system was then recharged with a mixture of 1.7(1) atm ^{3}He and 1.3(1) atm O_2, determined from the total gas pressure and a measurement of $T_1 = 1.8(1)$ s. The viscous penetration depth under these conditions, and at an acoustic frequency of 1.31 Hz, is calculated to be $\delta_\nu = 0.17(1)$ cm.

Table 5.2: Comparison of the VPE-MRI parameters used in the two experiments. The information from Table 5.1a is duplicated here in the preliminary dataset column.

parameter	preliminary dataset	improved dataset
T_E	10.0 ms	16.6 ms
T_π	6.1 ms	8.3 ms
τ	1.0 ms	1.0 ms
Δ	3.0 ms	3.0 ms
G_p max	0.419 G/cm	1.26 G/cm
N_p	32	16
G_r	0.10 G/cm	0.067 G/cm
G_{VPE} max	1.67 G/cm	1.12 G/cm
N_r	128	32
T_R	1.527 s	2.290 s
N_A	64	128
Δf	± 2500 Hz	± 625 Hz

Each day the total mean pressure of the gas in the system was brought to 2,277 Torr. This was done to ensure that the response of the AMR to the 1.31 Hz drive was consistent from day-to-day. The birdcage coil was then tuned to the Larmor frequency and monitored until the resonance was stable for at least an hour. The π ($\pi/2$) pulses were then calibrated by adjusting the RF power until the first zero (maximum) in the FID amplitude was found. Eleven velocity maps were acquired at ten evenly spaced phases of the acoustic cycle with increments in the delay time set at $\Delta T_\phi = 76.3$ ms. A reference image was acquired prior to each VPE image. The increased image acquisition time as well as the extra reference image acquisitions meant that only one or two velocity maps were produced per day. The data acquisition period lasted for two weeks.

Before the time-series data were Fourier transformed, the number of points along the frequency encoding dimension was truncated to 32. This is discussed further in the following section. Sample reference and VPE images reflecting the magnitude of the signal are shown in Fig. 5.6. The average SNR in these reference and VPE magnitude images are 21 and 18, respectively. The bipolar field gradient used in the VPE pulse is 2/3 the amplitude of that used in the preliminary work. The loss of signal strength in the VPE image compared to the reference image is therefore less substantial than in the previous round of measurements. The average SNR of the velocity calculated using Eq. (5.3) is comparable for the two sets of imaging parameters.

5.4.2 Truncation and Asymmetric Sampling of k-space

The field gradient strengths and pulse sequence timing listed in Table 5.2 were intended to decrease the y-dimension of each pixel. The peak of the spin echo produced by this sequence occurs 7.2 ms into the 102.4 ms read data acquisition time. This asymmetric acquisition window is arranged to minimize signal losses caused by T_2^* dephasing and additional losses associated with the diffusion of ^{3}He during the frequency-encoding field gradient. Figure 5.7a shows the magnitude of the signal acquired during frequency encoding but without the phase encoding pulse. The "calculated" line indicates the magnitude of the signal that is expected for the inverse Fourier transform of a 1.3-cm-diameter cylinder and a read gradient amplitude of 0.067 G/cm. Also shown on this plot are curves representing the contributions to signal amplitude decay through T_2^* losses and ^{3}He diffusion

(a) Reference image   (b) VPE image
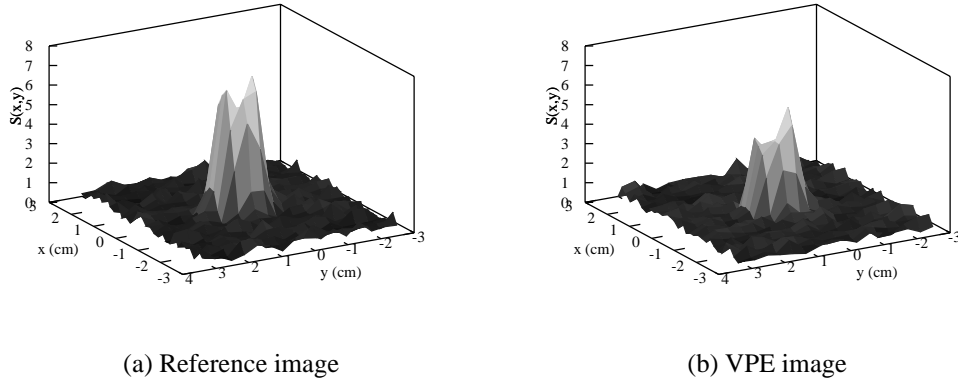
Figure 5.6: Sample images reflecting the magnitude of the detected ^{3}He precession signal for reference and VPE datasets. The magnitude of the VPE field gradient pulse was 1.12 G/cm.

during the frequency-encoding gradient. Diffusion during the read magnetic field gradient dominates the signal losses. The time at which the echo appears was chosen so that full-width of the central maximum of the signal is recorded along with part of its first side lobe in the $-k_x$ direction. Delaying the echo any further in time would result in significant diffusion related signal losses. Approximately 25 ms after the beginning of signal acquisition, the signal amplitude is completely buried in noise and contributes little information. The dataset was therefore truncated from 128 points to 32 points along the read direction before the Fourier transform was performed. The measured and calculated magnitude of the time-series signal are shown in Fig. 5.7b. The correspondence between the two is reassuring. No adjustable parameter other than a scaling factor for the amplitude is involved. This speaks to the accuracy of the field gradient calibrations.

The map of k-space that is acquired using this pulse sequence is discrete, finite and asymmetric. This will result in image artifacts. To develop some intuition as to the types of artifacts that might be expected, the 32 by 16 point k-space mapping was simulated using the Fourier transform of the cylinder function (see Eq. (A.9) in App. A). The discrete Fourier transform was then used to calculate the ideal image one would expect in Cartesian space. The resulting k-space map and corresponding Cartesian space magnitude image are shown in Fig. 5.8. This model does not incorporate attenuation of the signal due to $T_2{}^*$
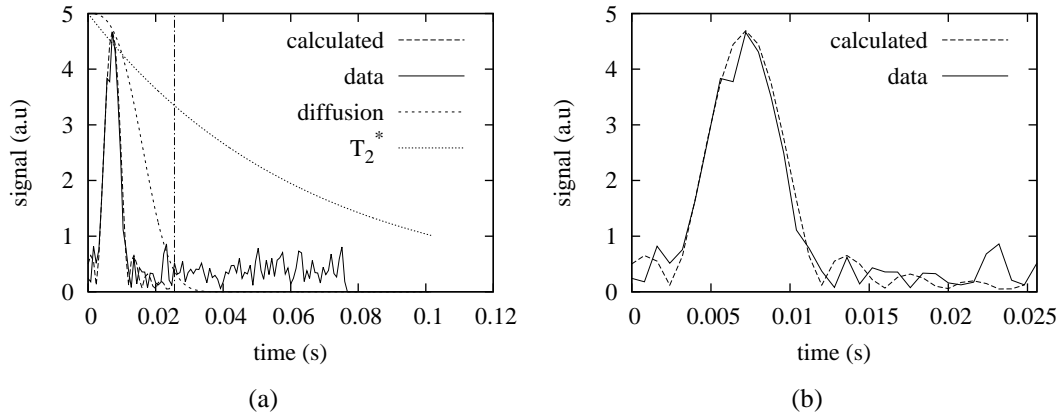
Figure 5.7: Magnitude of a signal acquired during the read gradient without phase encoding. Panel (a) shows magnitude data along with the time-dependence expected from the Fourier transform of the cylinder function in the absence of decay processes (Eq. (A.9)). It also shows the signal attenuation one would expect to observe if the decay was dominated by T_2^* (Eq. (2.11)) or diffusion during the read gradient (Eq. (2.35)). The vertical line marks the point at which the data were truncated. Panel (b) shows the same echo data after truncation to 32 points.

dephasing or diffusion of ^{3}He in the read gradient. The simulation of the truncation of the k-space map predicts two signal peaks in the image magnitude along the x-axis. This Gibbs ringing artifact is clearly visible in the sample reference image and to a lesser extent in the VPE image shown in Fig. 5.6. The simulation also predicts broadening, with signal being observed in pixels located just outside the acoustic flow tube. The strength of the signal in these pixels is of order 20% that of the maximum signal strength in the image.

5.4.3 Evaluation

An evaluation of the velocity maps acquired using the improved procedure was performed in a manner analogous to that used during the preliminary work. The time series data were Fourier transformed, the reference image phase was subtracted from the VPE image phase and the result was scaled by the FOS/2π. Equation (5.9) was then fit to the line of velocity measurements passing closest to the center of the tube with the phase ϕ as the

(a)   (b)

Figure 5.8: (a) Simulation of a two dimensional MRI k-space map for a 1.3-cm-diameter cylindrical sample using the imaging parameters summarized in Table 5.2. (b) Fourier transform of the simulated dataset.

only free parameter. A mean velocity amplitude $\langle u_1 \rangle = 20.0$ cm/s and viscous penetration depth $\delta_\nu = 0.17$ cm were used in the fit. The results are shown in Fig. 5.9 along with a line indicating the expected trend. The y-intercept of this line is set to -1.5 rad based a measurement of the phase of the mean gas velocity relative to the photodiode trigger, as outlined in Sec. 4.5. The velocity maps were acquired in order of increasing T_ϕ. The maps represented by points at $T_\phi/T_A = 0$ and 1 are separated in phase by 2π and therefore image the velocity at the same phase in the acoustic cycle. These images were acquired on the first and last days, respectively, of the two week data collection interval. The difference in the extracted phase between these images is 0.11(6) rad. This suggests that the reproducibility of the experiment has been improved relative to the preliminary round of measurements.

5.4.4 Summary

A set of eleven velocity maps were acquired at ten equally-spaced phases of the acoustic cycle. The initial evaluation, performed by fitting Eq. (5.9) to the velocity maps to determine ϕ, shows the expected sinusoidal dependence of gas velocity on the delay time T_ϕ between the photodiode trigger pulse and the launch of the imaging sequence. The image pixel dimensions are 0.18(2) cm $\times$ 0.39(3) cm which provides sufficient resolution to produce two

Figure 5.9: The phase ϕ extracted from fits of Eq. (5.9) to the velocity maps as a function of delay time T_ϕ normalized to the acoustic period T_A. The solid line has a slope of 2π and the y-intercept is taken from the phase of the velocity relative to the AMR drive trigger signal as measured during an independent measurement of the AMR frequency response.

dimensional maps of the flow field. The average uncertainty in the velocity measurements remains unchanged from the preliminary work. The repeatability of the flow conditions in the acoustic flow tube were improved by maintaining the AMR at constant pressure. The temperature dependence of the birdcage coil resonant frequency remains as a significant factor limiting reproducibility. A quantitative analysis of the data from the second round of image acquisition is presented in Ch. 6.

Chapter 6

Results and Discussion

This chapter presents a more detailed analysis of the data from the second round of VPE-MRI experiments outlined in the previous chapter. It amounts to a study of the velocity profile for a gas undergoing acoustic oscillations in a cylindrical tube. The conditions for this experiment were such that much of the gas was in the viscous boundary layer. Velocity maps were acquired at 10 evenly-spaced phases of the acoustic cycle and are compared to the theoretical velocity profile expected for a cylindrical tube, as summarized in Ch. 2 (Sec. 2.2.4). The experimental results and method are then compared with AVBL velocity measurements as performed by Huelsz *et al.* using Hot Wire Anemometry (HWA) [42], and by Castrejón-Pita *et al.* using Laser Doppler Velocimetry (LDA) and Particle Image Velocimetry (PIV) [43].

6.1   Velocity Maps

Velocity maps were produced using the spin-echo pulse sequence outlined in Sec. 5.1 and the procedure described in Sec. 5.4.1. The acousto-mechanical resonator was adjusted to drive spatially-averaged peak displacement oscillations of 2.4 cm in the acoustic flow tube at a frequency of $1.31(1)$ Hz (acoustic period $T_A = 763(6)$ ms). The peak volume flow rate was thus $U_1 = 26.5(7)$ cm^3/s. The phase of the acoustic motion relative to the photodiode trigger pulse was $-1.50(8)$ rad, as determined from the acousto-mechanical resonator characterization described in Sec. 4.5. The gas mixture was composed of 1.7 atm ^{3}He and 1.3 atm O_2 and so at this frequency the viscous penetration depth $\delta_\nu = 0.17(1)$ cm, as

133

calculated from kinetic theory (see Sec. 4.4). This corresponds to 26% of the acoustic flow tube radius $R = 0.65(1)$ cm. Alternatively, $R/\delta_\nu = 3.8$ (cf. Fig. 2.14). The delay time $T'_\phi = T_\phi + T_{\text{VPE}}$ between the photodiode trigger pulse and the midpoint of the bipolar VPE gradient pulse was sequentially incremented by $T_A/10 = 76.3$ ms to acquire images at 10 equally-spaced phases of the acoustic cycle. These velocity maps have a nominal resolution of 0.179 cm $\times$ 0.385 cm in the x- and y-directions, respectively as set by k-space sampling.

In addition to this, the characteristic distance for ^{3}He diffusion in any one direction over the time interval between the initial $\pi/2$ pulse and the peak of the echo (ie. T_E) is $\delta_{\text{TE}} = \sqrt{2DT_E} = 0.09$ cm, where the diffusion coefficient $D = 0.25$ cm^2/s was calculated from kinetic theory [72]. Close to the wall the effective value of δ_{TE} is reduced because random motion of the gas is somewhat restricted. The transverse diffusion of ^{3}He during the VPE measurement tends to smear out the radial velocity profile. The spatial resolution imposed by this effect is of order $2\delta_{\text{TE}} = 0.18$ cm (or 28% of the tube radius) which is the same as the nominal image resolution along the x-direction (or frequency encoding direction).

Figure 6.1 shows the alignment of the pixel grid relative to the acoustic flow tube. The center of the tube was located from the image data by manually centering a 0.65 cm radius circle on the magnitude images. The estimated uncertainty in the location of the image center is 0.02 cm in the x-direction and 0.03 cm in the y-direction. Velocity data associated with pixels centered at a radial distance greater than 0.5 cm were cut from the dataset. Although somewhat arbitrary, this mask effectively eliminates data located within a distance $2\delta_{TE} \approx \delta_\nu$ of the walls. Velocity data from outside of this cutoff (but still located within the acoustic flow tube) showed significant scatter. To aid in visualization, the datasets for each two-dimentional velocity map have been divided into three horizontal lines. The line that passes through the center of the acoustic flow tube is labelled y_0. The lines displaced upwards and downwards by one pixel (0.39 cm in the y-direction) are labelled y_+ and y_-, respectively.

In some sense, the velocity associated with each pixel represents an average over a finite area. Differences between the pixel averaged velocity and pixel center velocity are expected in regions where the curvature in the velocity profile is large or near walls where the pixel area extends outside the acoustic flow tube. To examine the extent of these deviations, $u_1(x, y)$ (i.e. Eq. (5.9)) was numerically averaged over the area of each pixel as shown in

Figure 6.1: Location of pixel grid relative to the acoustic flow tube. The outer circle represents the 0.65 cm radius tube. The inner circle represents a 0.5 cm radius mask outside of which data were ignored during the analysis. The velocity map data are grouped into three lines along the x-axis. These are labelled y_+, y_0, and y_-, as shown.

Fig. 6.1. For pixels included in the analysis (lying within the inner gray circle of Fig. 6.1), deviations between the pixel averaged velocity and the pixel center velocity were less than 5%. This is comparable to the uncertainty in measured velocities calculated using Eq. (5.3) and much smaller than the observed scatter in the data. It was thus decided to ignore this difference and simply use the pixel center velocity calculated using Eq. (5.9). A more robust treatment of the problem might involve fitting data to the model in k-space before performing the inverse Fourier transform to obtain an image.

The 10 velocity maps are shown in Fig. 6.2. The lines of data corresponding to y_+ and y_- have been displayed as being offset from the y_0 data by 25 cm/s up and down, respectively. The maps are labelled by the delay time T'_ϕ between the photodiode trigger pulse and the midpoint of the bipolar VPE pulse. The approximation to the velocity profile Eq. (5.9) was fit to the entire set of velocity maps simultaneously with $\langle u_1 \rangle$, δ_ν and ϕ_0 as free parameters. Since the VPE-MRI sequence measures the average z-velocity of the ^{3}He over the duration of the bipolar velocity phase encoding gradient $(\Delta + \tau)$, the phase ϕ for each velocity map is given by $\phi = \omega_A T'_\phi + \phi_0$. The results from the fit are $\langle u_1 \rangle = 19.9(5)$ cm/s, $\delta_\nu = 0.22(2)$ cm, and $\phi_0 = -1.50(2)$. The best fit curves for the velocity profile approximation Eq. (5.9) are also plotted in Fig. 6.2 for the three lines of data. There is a respectable correspondence between the data and the predicted amplitude and phase of the velocity, as functions of both position and time. Most notably, the phase of the gas velocity near the wall leads that of the gas near the center of the tube. This is most evident along the y_0 (central) line of data in panels (a), (f), and (k), of Fig. 6.2 where the sense

of the velocity near the walls is opposite that near the center of the acoustic flow tube. This striking observation is clear evidence of the fact that the gas motion is indeed strongly influenced by the presence of the walls, as one would expect.

Table 6.1: Comparison of the free parameters extracted from the fit of Eq. (5.9) to the velocity maps shown in Fig. 6.2 with values determined independently from the AMR characterization measurements outlined in Ch. 4.

fitting parameter	characterization measurement	fitting result
$\langle u_1 \rangle$	20.2(5) cm/s	19.9(5) cm/s
ϕ_0	-1.50(8)	-1.50(2)
δ_ν	0.17(1)	0.22(2)

Table 6.1 summarizes the parameters determined from the fit and those determined from independent characterization measurements described in Ch. 4. The mean velocity amplitude and phase measurements are in very good agreement with one another (less than one standard deviation). However, there is some disagreement between the viscous penetration depth extracted from the VPE-MRI data and the value calculated based on measured gas composition and calculated viscosity. Here the discrepancy is more than two standard deviations. There are at least two factors that might contribute to this discrepancy. First, the calculation used to determine δ_ν from kinetic theory was based on a model for a binary mixture of gases. The calculation of the viscosity by this method should be considered to be an approximation [72]. If the mass diffusion coefficient is known for the mixture then a calculation of the viscosity based on D_{12} is preferred. No such correction was applied to obtain the calculated values and therefore the calculated viscous penetration depth δ_ν reported to this point should be taken as an approximation. A more concrete value for the viscosity of the ^{3}He-O_2 mixture would certainly benefit the interpretation of the data. A second factor that might influence the value of the viscous penetration depth extracted from the VPE-MRI data is spatial resolution. That is, the combined effect of finite pixel size, transverse gas diffusion during echo formation (T_E), and the lack of data near the walls (because of the mask) might tend to smear the data out along r, tending to increase the apparent value of δ_ν. Again, this would have to be studied properly before drawing any

strong conclusions. This issue is revisited in the following section.

The individual velocity maps are snapshots of the two dimensional velocity profile at an instant in time with overall time resolution of order $T_E/2 = 8.3$ ms or 1% of an acoustic cycle. To evaluate the use of this VPE-MRI method in the context of yet more complicated flow patterns it is worthwhile comparing individual maps with predicted velocity profiles one-by-one. In examining the various velocity maps shown in Fig. 6.2, panels (d), (e), (h), (i), and (k) stand out as showing good agreement with the predicted velocity profile. Most of the data on these maps lie within one standard deviation of the global best-fit curve with the exception of one or two notable outliers. The poorest agreement is evident in Fig. 6.2, panels (c) and (j) where most of the data deviates significantly from the prediction by as much as 10 cm/s or $\langle u_1 \rangle /2$. During acquisition of these particular velocity maps it was noted that the tuning of the birdcage coil was drifting excessively. This problem is reflected in the time series data through the appearance of a small FID at the begining of the acquisition phase but prior to the anticipated echo (i.e. during frequency encoding). This is a symptom of a poor RF π pulse resulting from a change in coil tuning as discussed in Sec. 5.3.2. The data in panels (a), (b), (f), and (g) of Fig. 6.2 show good agreement with the predicted velocity profile across the diameter of the tube (i.e. along the line labelled y_0 in Fig. 6.1) but exhibit some significant deviations along the lines y_+ and y_-. In situations where velocity measurements are in good agreement with the theoretical prediction, the uncertainty in individual velocity measurements range from 2 to 5 cm/s and the standard deviation between measured and nominal predicted velocities is 5.8 cm/s. The uncertainty in position associated with identification of the center of the tube is 0.02 cm in the x-direction and 0.03 in the y-direction. More significant is the uncertainty in spatial resolution imposed by transverse diffusion of gas during the VPE-MRI echo formation process; i.e., a length scale of order $\delta_{\text{TE}} = 0.09$ cm. Velocity measurements performed within 1.5 mm of the tube wall showed significant scatter and were thus not included in Fig. 6.2. The average uncertainty in individual velocity measurements for this excluded subset of the data is 9 cm/s. The most significant issue with the quality and repeatability of the velocity maps is related to the stability of the birdcage coil tuning. Because of the temperature dependence of the coil resonant frequency discussed in Sec. 5.3.2, the tipping angle produced by RF pulses can change significantly in response to a temperature change of just 1 C. This will influence signal strength and the phase of the nuclear precession signal seen by the spectrometer.

(a) $T'_\phi = 4.3$ ms

(b) $T'_\phi = 80.6$ ms

(c) $T'_\phi = 157.0$ ms

(d) $T'_\phi = 233.3$ ms

(e) $T'_\phi = 309.6$ ms

(f) $T'_\phi = 386.0$ ms
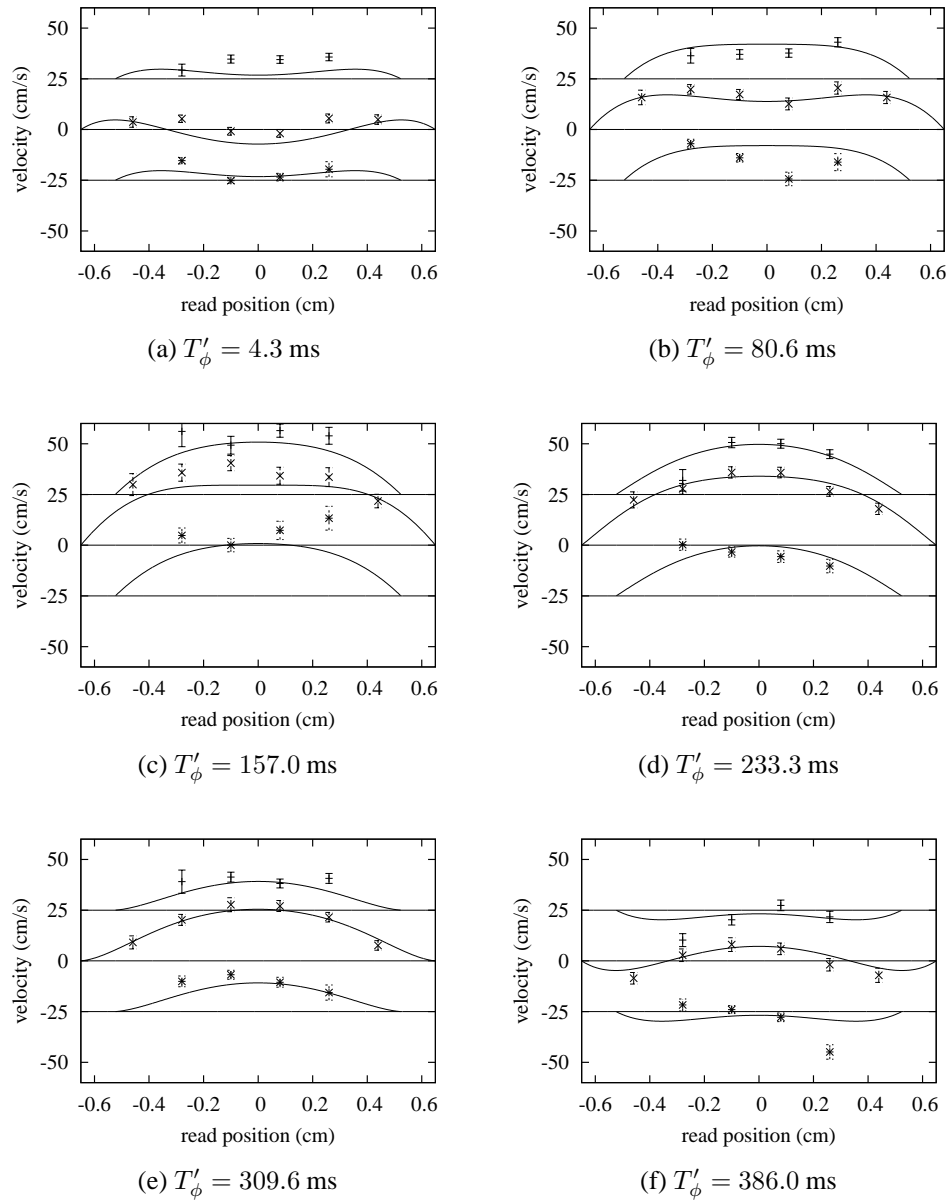
Figure 6.2: Velocity maps for a gas oscillating in a cylindrical tube of radius $R = 0.65$ cm. The curves represent Eq. 5.9 with the best fit parameters given in Table 6.1. The central set of data in each plot corresponds to the central line of pixels shown in Fig. 6.1 and labelled y_0. The other two sets of data correspond to the lines of pixels y_+ and y_-; velocities for these datasets have been offset by ± 25 cm/s, respectively. Note the time $T'_\phi = T_\phi + T_{\text{VPE}}$. This figure is continued on the following page.

(g) $T'_\phi = 462.3$ ms

(h) $T'_\phi = 538.7$ ms

(i) $T'_\phi = 615.0$ ms

(j) $T'_\phi = 691.3$ ms

(k) $T'_\phi = 767.7$ ms
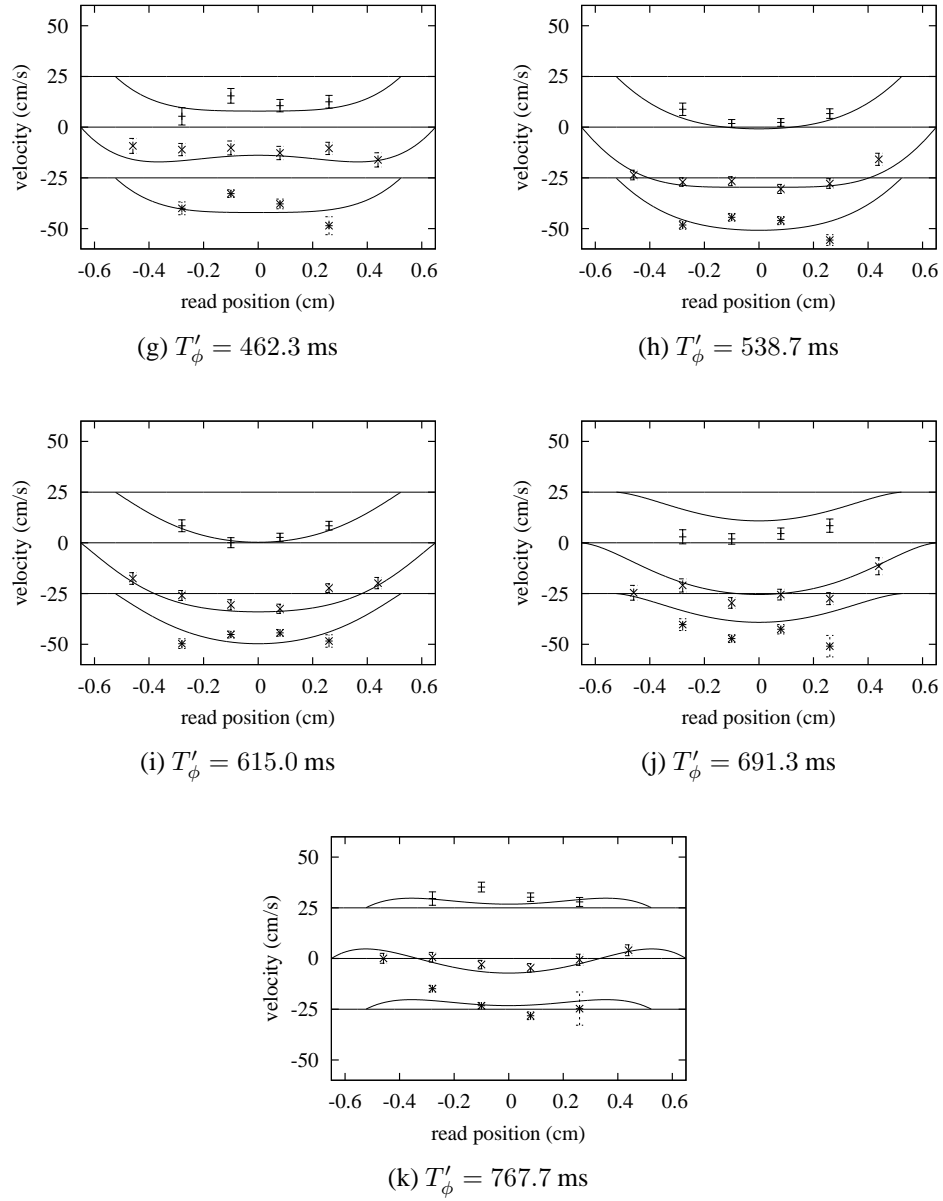
Figure 6.2: (continued)

6.2 Magnitude and Phase of the Velocity in Each Pixel

Each velocity map shows the velocity of the gas as a function of position at a time T_ϕ' relative to the synchronization pulse. An alternative to the global fit of Eq. (5.9) to the data as presented above is to plot the gas velocity in each pixel as a function of time. These velocities will be functions of the form

$$u_1(r_p, T_\phi) = u_1 \cos[\omega_A T_\phi' + \phi] \qquad (6.1)$$

where u_1 and ϕ are parameters representing the local amplitude and phase. The central line of velocity data (acquired from the pixels labelled y_0 in Fig. 6.1) are plotted in this format in Fig. 6.3. Also shown are fits of Eq. (6.1) to the data with u_{1p} and ϕ as free parameters. This figure reveals that the measured velocity does indeed exhibit the expected sinusoidal time dependence, and that the peak amplitude decreases toward the tube wall as the phase advances. This is consistent with observations made in the context of Fig. 6.2 (i.e. the velocity maps) where it was seen that the velocity of the gas near the wall leads that of the gas near the center, and that the velocity amplitude decreases to zero at the wall. The scatter in the velocity data relative to the nominal predicted distribution increases with distance from the center of the tube. There are two notable outliers in each panel of Fig. 6.3. These occur at $T_\phi' = 157.0$ ms and 691.3 ms, corresponding to the velocity maps shown in panels (c) and (j). As noted previously, the data associated with these two velocity maps are suspect.

Figure 6.4 shows the normalized velocity amplitude and phase, extracted from fits of Eq. (6.1) to the data as a function of pixel position $r = \sqrt{x^2 + y^2}$ normalized to tube radius R. The velocity amplitude u_1 is normalized by dividing by the mean velocity amplitude $\langle u_1 \rangle$ extracted from the global fit (i.e. Table 6.1). The phase is given in terms of the shift relative to the average phase of the data from the two pixels closest to the center of the tube. Presented in this form, the data are useful for comparing velocity profiles for tubes of different diameters and for different peak velocities but similar ratios of R/δ_ν. As before, the spatial resolution associated with each datum is limited by diffusion in the transverse direction, and is of order $\delta_{\text{TE}} = 0.09$ cm. Normalized to the tube radius, this uncertainty becomes 0.14. Also shown in Fig. 6.4 are the nominal (i.e. theoretical) velocity distributions given by Eq. (2.70) for two different viscous penetrations depths. The curves for $\delta_\nu = 0.17$ cm correspond to the result expected when this parameter is calculated from

(a) $x = -0.45$ cm

(b) $x = -0.27$ cm

(c) $x = -0.09$ cm

(d) $x = 0.09$ cm

(e) $x = 0.27$ cm

(f) $x = 0.45$ cm

Figure 6.3: Velocity measured at each pixel location along the line labelled y_0 in Fig. 6.1 as a function of delay time T'_ϕ. The labels for x essentially indicate the distance from the tube center for this particular line of data. The curves represent Eq. (6.1) with u_1 and ϕ being adjusted for each panel to give the best fit.

kinetic theory based on a mixture of 1.7 atm ^{3}He and 1.3 atm O_2 being driven at 1.31 Hz. The curves for $\delta_\nu = 0.22$ cm were obtained by fitting the normalized velocity data to Eq. (2.70). The same result is obtained as was found when a global fit of Eq. (5.9) to the velocity map was performed, as described in Sec. 6.1. Comparing the nominal (i.e. best fit theoretical) velocity distributions with the data it seems that the fit of Eq. (5.9) is strongly influenced by data near the center of the image. There is very little difference between the two curves in the vicinity of the cluster of points in the range $0.6 < r/R < 0.7$. The dataset includes two points at $r/R = 0.4$ that show better agreement with the velocity amplitude expected for $\delta_\nu = 0.17$. At the same time there is stark disagreement between this curve and the velocity amplitude data for the two points at $r/R = 0.14$. This discrepancy is resolved by allowing δ_ν to increase to 0.22 cm. Notably, all four of the central velocity measurements are in general associated with well-behaved data sets (see Fig. 6.3 panels (b), (c), (d), and (e)). In this light, there is little evidence to suggest that one pair or the other ought somehow to be discounted. As was pointed out previously, calculation of δ_ν from kinetic theory as it was performed is an approximation. A proper determination of the viscosity of the ^{3}He-O_2 mixture from empirical data would enable a more critical distinction between these curves to be made. At the same time there are clearly a number of experimental issues (such as temperature stability of the birdcage) that have yet to be resolved.

Finally, some useful assessments of the VPE-MRI measurement technique can be inferred from Figs. 6.3 and 6.4. The uncertainties in u_1 increase with radius, but reveal an instrumental pixel-by-pixel velocity amplitude resolution of order 12% near the center of the tube (increasing to 22% over the range $0.6 < r/R < 0.7$). These figures are comparable to the scatter between individual measurements and the nominal velocity distribution if one accepts the best fit viscous penetration depth $\delta_\nu = 0.22$ cm. Similarly one can evaluate the distrubution of measured acoustic phases relative to the nominal acoustic phase (the abscissae in Fig. 6.3). This reveals a pixel-by-pixel phase resolution of order 0.12 rad near the center of the tube (increasing to 0.16 rad for $0.6 < r/R < 0.7$)

6.3 Comparison of AVBL Measurements

The study of gas oscillations in the AVBL is of importance to a variety of nonlinear thermoacoustic and acoustic phenomena, such as streaming, which have important technologi-
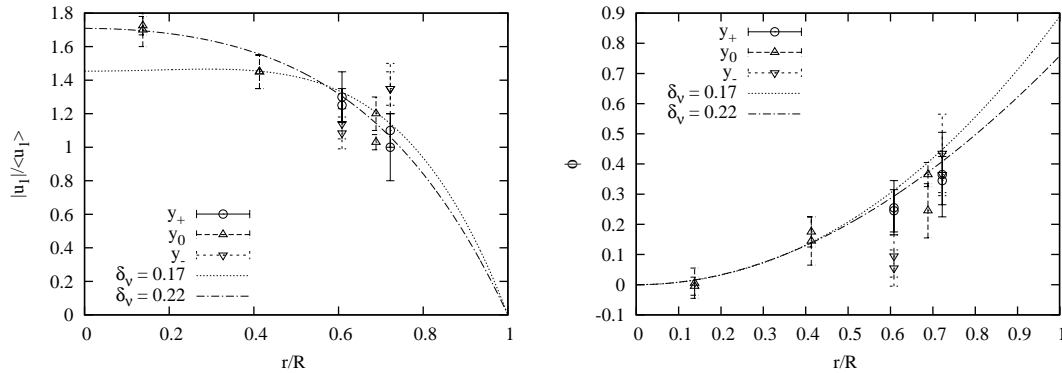
Figure 6.4: (left) Normalized velocity amplitude as a function of the normalized radial position. (right) Phase offset as a function of normalized radial position. The velocity amplitude and phase at each pixel location was found by fitting Eq. (6.1) to the data as illustrated in Fig. 6.3. The curves represent the normalized velocity amplitude and phase given by Eq. (5.9) for two values of the viscous penetration depth δ_ν. The value $\delta_\nu = 0.17$ cm is calculated from kinetic theory. The value $\delta_\nu = 0.22$ cm is obtained through fitting Eq. (5.9) to the data shown in the amplitude plot (left). The same value is obtained when a global fit to all of the velocity data is performed as described in Sec. 6.1

cal implications related to the performance of heat engines, heat pumps, refrigerators, and other devices. Below, I assess the results of my VPE-MRI investigation of flow in the AVBL with similar measurements performed using Hot-Wire Anemometry, Laser-Doppler Anemometry and Particle Image Velocimetry.

6.3.1 HWA, LDA, PIV Measurements of Gas Velocity in the AVBL

Huelsz *et al.* studied the velocity amplitude and phase of air in the acoustic viscous boundary layer (AVBL) using hot-wire-anemometry (HWA) [42]. In HWA, gas velocities are inferred from the convective heat transferred away from a heated wire maintained at constant temperature as gas flows past. Huelsz *et al.* employed a quarter-wavelength standing wave acoustic resonator with a rectangular internal cross-section of 9.8 cm $\times$ 5.4 cm. Resonators with three different internal lengths were studied, corresponding to operating frequencies of 35(1) Hz, 46(1) Hz, and 130(1) Hz. The viscous penetration depths δ_ν at these frequencies

were 0.38 mm, 0.33 mm, and 0.20 mm, respectively. Acoustic oscillations were driven by a loudspeaker and the other end of the resonator was closed by a massive cap. The pressure amplitude p_1 at the location of the end cap was measured using a microphone. Typically in HWA, a steady flow is superimposed on the oscillatory motion to prevent the heated gas from periodicaly returning to the heated wire. The HWA probe used by Huelsz *et al.* was specially calibrated for use with oscillatory flow without superimposed steady flow [41].

To find the theoretical gas velocity as a function of distance from a wall, their resonator was treated as a pair of parallel plates in the widely-spaced limit. That is, the distance between plates l is much larger than δ_ν. If the distance from the wall is normalized to the viscous penetration depth (i.e. $y = y^*/\delta_\nu$), time is normalized to the angular period of the oscillation (i.e. $t = t^*\omega$) and velocity is normalized to its value far from the wall (i.e. $u = u^*/u_\infty$) then the non-dimensional velocity is given by [24]

$$u = \{1 - \exp[-(1+i)y]\}e^{it} \tag{6.2}$$

where the superscript $*$ indicates dimensional quantities.

Figure 6.5 shows the normalized velocity amplitude u_a and phase α as a function of non-dimensional distance from the wall as measured by HWA at 130 Hz by Huelsz *et al.*, along with the theoretical prediction given by Eq. 6.2. The measured velocity amplitude exhibits the expected trend, increasing from zero at the wall and reaching a plateau for $y > 6$. However, the data do not reveal the small peak in amplitude expected near $y = 2$. The measured velocity phase shows good agreement for $y > 0.5$. Discrepancies between the measured velocity amplitude and the prediction of Eq. (6.2) is attributed to interaction of the flow around the HWA probe and the wall. The authors hypothesized that the technique used in this measurement could be used to calibrate HWA probes for use near walls. Their data seem to reveal that at some level, HWA is intrinsically an invasive measurement technique.

Castrejón-Pita *et al.* repeated the study reported by Huelsz *et al.* using laser-Doppler-anemometry (LDA) and particle-image-velocimetry (PIV) [43]. A resonator, with an internal rectangular cross-section of 9.8 cm $\times$ 5.4 cm and lengths of 0.6 m and 1.2 m (corresponding to 114.5(5) Hz and 68.5(5) Hz), was employed. The apparatus was built using Plexiglas to enable optical access. The resonator was filled with air at atmospheric pressure and an ambient temperature of 23 C, resulting in viscous penetration depths of $\delta_\nu = 0.27$ mm and 0.21 mm, for 68.5 Hz and 114.5 Hz, respectively. The pressure am-

plitude and phase, measured with a microphone, were used to determine the mean velocity and phase of the oscillating gas at the center of the resonator, assuming adiabatic conditions. The mean velocity and phase were then used to normalize the LDA and PIV gas velocity amplitude and phase measurements.
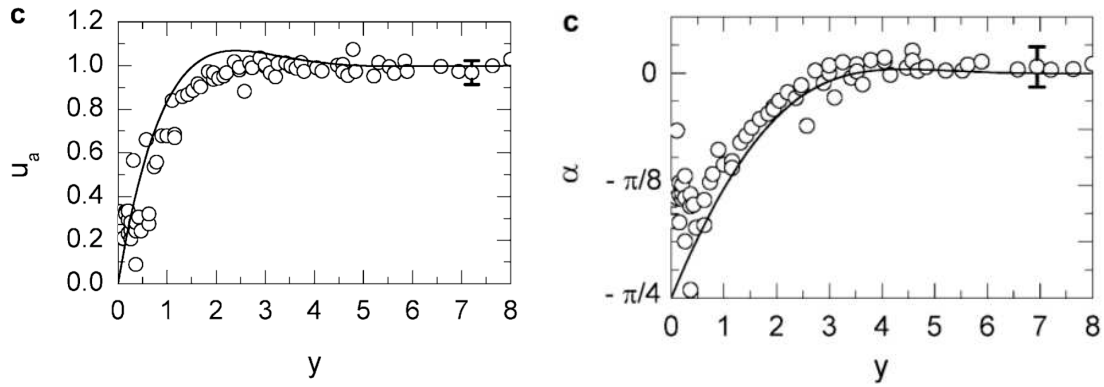
Figure 6.5: Non-dimensional velocity amplitude $u_a = |u|/|u_\infty|$ and phase α as a function of non-dimensional distance from the wall $y = y^*/\delta_\nu$, as measured using HWA. The acoustic frequency is 130 Hz and $\delta_\nu = 0.20$ mm. (adapted from Huelsz *et al.* [42])

LDA relies on crossed laser beams to create an interference pattern at the velocity measurement location. The fluid is seeded with small particles, such as smoke or fine mist, that are suspended and entrained in the flow. As the seed particles pass through the crossed laser beams, the scattered light intensity (detected using a photomultiplier tube) is modulated by the interference pattern. The frequency of the modulation is used to infer seed particle velocities. Scanning the point at which the laser beams intersect through the region of interest produces a velocity map. The LDA system employed by Castrejón-Pita *et al.* had a measurement volume of dimensions 0.64 mm $\times$ 0.075 mm $\times$ 0.075 mm and a positioning precision of 1 μm. The fluid was seeded with smoke particles with sizes of order 1 μm.

Figure 6.6 shows the normalized gas velocity amplitude and phase as a function of nondimensional distance from the wall as measured by LDA by Castrejón-Pita *et al.*, along with the theoretical prediction given by Eq. (6.2) [43]. The LDA measurement of the velocity amplitude and phase reveal much better agreement with Eq. (6.2) than do the HWA data of Huelsz *et al.*. This in turn indicates that the discrepancies observed by Huelsz *et al.* were related to the technique and not a failure of the theory. The LDA velocity measurements

were made at distances as close as 0.01 mm away from the waveguide wall, with only a slight increase in the noise. This was made possible through allowing the interrogation volume formed by the crossed laser beams to partially enter the transparent Plexiglas wall. The slight peak in velocity amplitude near $y = 2$ was successfully resolved at both drive frequencies. For $y \geq 1$ the measured amplitudes are within 1% of the theoretical prediction.

Figure 6.6: LDA velocity measurements of a gas in the AVBL. (adapted from Castrejón-Pita *et al.* [43])

PIV also uses laser light scattered from seed particles entrained in fluid flow to measure gas velocity. A sheet of laser light illuminates a thin slice of the fluid. Seed particles in the slice scatter this light, which is then captured as a time series of images using a camera. The data in each image are windowed into small interrogation regions. The displacement of particles between successive images is then used to infer velocity. The displacement is found by cross-correllating each window of one image with the corresponding window of the subsequent image. That is, the pixel intensity data in the two images are multiplied together as the first image is shifted relative to the second. A peak in the correlation function or "image product" occurs when the image displacement matches the mean particle displacement. By performing cross-correlations in the x- and y-directions, a two dimensional map of fluid velocity is obtained. In the setup employed by Castrejón-Pita *et al.*, the sheet of laser light had a thickness of 1 mm. The windowing procedure resulted in an interrogation volume of

0.124 mm $\times$ 0.031 mm $\times$ 1 mm.

Figure 6.7 shows the normalized gas velocity amplitude u_a and phase α as a function of nondimensional distance from the wall as measured using PIV by Castrejón-Pita *et al.* along with the theoretical prediction given by Eq. (6.2) [43]. Again, the PIV data reveal good agreement with theory. For $y \geq 1$ the measured velocity amplitudes are within 3% of the prediction. However, the technique is limited to measuring the velocity at distances greater than 0.04 mm from the wall, which in this case corresponds to $0.15\delta_\nu$.
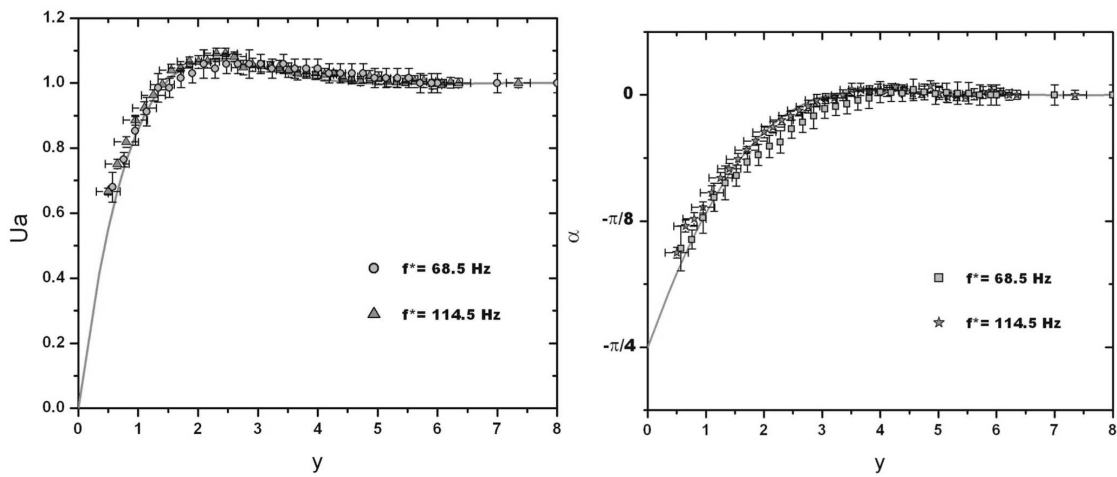
Figure 6.7: PIV velocity measurements of a gas in the AVBL.

6.4 Summary

Four experimental methods for measuring oscillating gas velocities in the acoustic viscous boundary layer have been summarized in this chapter. Implementation of the first method, Velocity-Phase-Encoded MRI, was the primary focus of this thesis. Although the intended demonstration was quite successful, it should be reasonably obvious that the state-of-the-art for the other three experimental methods – Hot Wire Anemometry, Laser Doppler Anemometry, and Particle Image Velocimetry – is much more refined. In some sense the AVBL velocity profile data acquired with the prototype VPE-MRI system is comparable to the HWA data of Huelsz et al. [42]; the velocity amplitude resolution of both techniques is of order 10 to 30%, and their spatial resolutions are certainly of order $\delta_\nu/2$ or better.

It is even reasonably clear that the HWA technique is invasive, which is not the case for VPE-MRI. On the other hand, the VPE-MRI technique, as it was implemented here, very obviously suffers from the influence of some uncontrolled systematic issues that lead to intermittent erratic responses. In this sense, the reliability and reproducibility of data from the VPE-MRI probe are not yet in the same league as HWA. Moreover, it is not yet understood why data from the prototype VPE-MRI system exhibits excessively large scatter close to the wall, and whether or not this represents a fundamental limitation of the probe. Neither HWA nor VPE-MRI (as it was implemented) yield AVBL data of the same quality as PIV or LDA. The latter in particular sets the current benchmark against which it seems appropriate to compare other techniques.

Chapter 7

Conclusion

The primary objective set out at the beginning of this thesis was to investigate the feasibility of using Velocity-Phase-Encoded MRI to probe oscillatory gas motion in the Acoustic Viscous Boundary Layer. That objective has been met. An Acousto-Mechanical Resonator was designed, constructed, and characterized, and was then used to drive large amplitude, low frequency, oscillatory gas motion in a long, narrow, cylindrical tube. The diameter ($2R = 1.3$ cm) of this Acoustic Flow Tube was chosen so as to ensure that a large fraction of the oscillating gas would be in the AVBL. The tube was filled to 3 atm with a ^{3}He-O$_2$ mixture, and was driven by the AMR at a frequency of 1.3 Hz. Under these conditions, and at room temperature, the nominal viscous penetration depth δ_ν for this binary gas mixture was calculated (from kinetic theory) to be 0.17 cm, hence the ratio $R/\delta_\nu = 3.8$. The role of the ^{3}He in the experiment is to act as an NMR-detectable gas. The role of the O$_2$, which is strongly paramagnetic, is to induce longitudinal nuclear spin relaxation of the ^{3}He at a rate comparable to the acoustic frequency. This permits rapid repolarization of the nuclear spin system in the background magnetic field, and enables one to launch a new image acquisition sequence for approximately every other acoustic cycle.

VPE-MRI experiments were successfully performed on thermally-polarized ^{3}He gas in a static magnetic field of 1.5 Tesla using a 48 MHz birdcage resonator designed to fit around the acoustic flow tube. Velocity sensitization along the axis of the tube was accomplished using a bipolar pair of half-sine field gradient pulses synchronized with the mechanical phase of the AMR drive. Pairs of images were then collected; one in which the VPE gradients were activated and the other in which they were turned off, acting as a phase

reference. Phase differences between these images were calculated pixel by pixel, and then mapped onto an absolute velocity scale. Each image involved 16 phase-encoding steps and individual VPE echo data were sampled 32 times at 1250 Hz, ultimately producing a 32 by 16 grid of discrete pixels with dimensions 0.179 cm by 0.385 cm. Each VPE sample was sampled and averaged 128 times, resulting in an image acquisition time of 78 minutes. Of the 512 pixels associated with each image, 14 were associated with pixels that overlapped the acoustic flow tube.

Images of absolute ^{3}He gas velocities were obtained, and reasonable agreement with theoretical expectations was observed. Some parameters, such as the spatial average of the velocity and its phase relative to the mechanical drive were captured by the VPE data with high accuracy (a few percent, or a few degrees). Others, such as the apparent viscous penetration depth deviated significantly (more than two standard deviations) from the value nominally expected from kinetic theory (albeit an expectation based on an approximation that may not be justified). Many velocity measurements associated with individual pixels showed quite reasonable correspondence with expectations; the standard deviation between measurement and theory increased as a function of radius, varying from approximately 15% near the center of the tube to about 25% for $0.6 < r/R < 0.7$. At the same time, other velocity measurements exhibited unusually large deviations from theory (several standard deviations). It is suspected that these outliers are produced in response to uncontrolled changes in experimental parameters. A known culprit in this respect is the tuning of the birdcage resonator used to apply RF tipping pulses, which is excessively sensitive to changes in temperature. The formation of spin echoes, which is central to the formation of VPE MR images, is almost certainly influenced when the resonator is detuned from the nuclear Larmor precession frequency. In some instances, such as for pixels located within approximately δ_ν of the walls, the scatter in measured velocities was inexplicably high. The mechanism responsible for these fluctuations was not identified, and ultimately data from pixels located close to the walls were excluded from analysis.

To some extent, the pixel size for the VPE-MRI experiments described in this thesis can be decreased, through appropriate k-space sampling. In particular, the only real penalty for increasing the pixel density along the phase encoding axis (the y-axis in Fig.6.1) is the total image acquisition time. However, there are additional intrinsic limitations to spatial resolution; in particular, diffusion during the encoding and readout procedure can result

in the exchange of atoms between different velocity classes, particularly at distances of order δ_ν from walls where gradients in velocity are large. For the VPE-MRI experiments described in Chapter 6 the characteristic length scale for diffusion during echo formation (i.e. $T_E = 16.6$ ms) was 0.09 cm, or approximately $\delta_\nu/2$ at the operating frequency of the AMR. Some improvement in resolution could be obtained through reducing T_E (by perhaps a factor of 2 or 3); otherwise, one would have to increase δ_ν (for example, through lowering the AMR drive frequency).

When compared to other experimental methods that have been used to study oscillatory gas flow in the AVBL – Hot Wire Anemometry, Laser Doppler Anemometry, and Particle Image Velocimetry – it is clear that VPE-MRI, as implemented in this thesis, is in need of further development. In some respects the sensitivity of VPE-MRI appears similar to that of HWA, but uncontrolled systematic factors such as the tuning of the birdcage resonator make it challenging to undertake a proper comparison at the moment. The other two experimental techniques, and in particular LDA, seem to outperform VPE-MRI and HWA insofar as AVBL applications are concerned.

There is no obvious reason why the stability and reproducibility of the VPE-MRI probe couldn't be improved quite substantially. The first step would obviously be to redesign the birdcage resonator with an eye to both temperature and mechanical stability. The stability of the mechanical resonator could also be improved, perhaps through the use of feedback. Improved AMR frequency stability would help to ensure consistent phasing between imaging sequences and the acoustic motion. Another modification that would be simple to implement would be to replace the elbows located on either side of both U-tubes with larger-radius bends. There is some evidence that this would reduce damping of the AMR, but it would also reduce the likelihood of distorting the velocity field in the acoustic flow tube. Yet another modification one might consider would be to improve the homogeneity of the static magnetic field $\mathbf{B}_0$ over the resonator volume. This would certainly help to improve resolution through narrowing of the ^{3}He linewidth, but would likely imply using a more modern superconducting imaging magnet.

Finally, a decision was made early on in this project to work with thermally polarized ^{3}He gas. This choice shaped several aspects of the existing experimental design, but one could imagine modifications to permit the introduction of hyperpolarized gases or the use of gases with many spin-bearing nuclei [105]. In both cases the motivation would be

to improve SNR, and thereby enhance image quality and/or decrease image acquisition times. While the goal of improving SNR has obvious merit, there could be some significant challenges associated with these alternatives. For example, it is difficult to replenish high levels of nuclear polarization from an external source without imposing a steady flow of gas, which then limits the types of problems that can be studied. Similarly, molecules like propane or SF_6 that each have many NMR detectable nuclei tend to exhibit much shorter longitudinal nuclear relaxation times in the gas phase than ^{3}He. If the goal is to maintain nuclear spin coherence over one acoustic period, higher frequencies are necessary, implying a thinner AVBL, and placing greater demands on spatial resolution.

Eventually, one could start to think about applications where VPE-MRI might reveal information that is not accessible via HWA, LDA, and PIV. For example, apart from material constraints (the need to exclude conducting and magnetic materials), VPE-MRI may well be the least invasive of the four velocity probes. Thus, one might consider trying to apply VPE-MRI to the study of acoustic flow in complex geometries. This might include studies of oscillating flow through orifices and/or tortuous channels, around objects, or at interfaces between different structures, such as between a heat-exchanger and a thermoacoustic stack or regenerator. These are all situations that are encountered in technological applications of acoustics; in some cases they are difficult to model with computational fluid dynamics techniques and in most cases the geometry presents challenges for optical access. Here it is worth noting that a variety of MR imaging techniques have already been used to examine the steady flow of gas through tubes (e.g. Poiseuille flow of propane [48]), through orifices (e.g. laminar flow of hyperpolarized ^{129}Xe [16]), around obstructions (e.g. laminar flow of hyperpolarized ^{129}Xe [49] and turbulent flow of thermally polarized SF_6 [18]), and in human lungs (inspiration of hyperpolarized ^{3}He [14, 15]). The modular design of the AMR is intended to permit replacement of the existing acoustic flow tube with more complex structures, in effect becoming an acoustic wind tunnel.

A comparison of these other MR-based probes of gas velocity with the work described in this thesis and other experimental methods reveals an important distinction. In principle, MR-based probes provide access to gas composition, while HWA, LDA, and PIV do not. This feature has not yet been exploited as a probe of acoustic dynamics, but it was used in this thesis as an *in-situ* monitor of gas composition. That is, a situation was contrived in which the ^{3}He longitudinal nuclear relaxation rate T_1^{-1} was dominated by interactions

with paramagnetic O_2 molecules; since this interaction is well characterized [86, 87], measurements of T_1^{-1} and total gas pressure were sufficient to determine the concentration of both constituents. One could imagine trying to exploit this sensitivity as a probe of acoustic mixture separation [71, 106]. A more esoteric version of this experiment might involve trying to investigate the phenomenon of spin-sorting [107], in which forces associated with magnetic field gradients can be used to enhance nuclear polarization above thermal equilibrium levels. More generally, NMR is an excellent non-invasive probe of diffusion. One could imagine integrating and correlating time-resolved diffusion-sensitized maps of gas diffusion in acoustic flow-fields with corresponding velocity-sensitized maps. Thus, experiments in well-defined geometries (analogous to those reported here) would be expected to reveal two characteristic length scales (δ_ν and δ_D) instead of one. Note that – as part of the investigation that led to this thesis – an attempt was made to directly measure the ^{3}He diffusion coefficient in ^{3}He-O_2 mixtures. This measurement was ultimately thwarted by systematic effects related to the temperature sensitivity of the birdcage resonator, but in principle it could provide valuable information. Finally, it is noteworthy that pulsed-field-gradient MRI sequences can be sensitized to gas motion in a variety of different ways. The method employed here – VPE MRI – sensitized the nuclear response to gas velocity. Modifications of this sequence enables one to sensitize images to gas displacement, velocity, acceleration, or even higher moments of the gas displacement.

An interesting observation hints at a peculiarity of the VPE-MRI experiment as it was implemented. As noted in Chapter 4, an unusual level of ringing was observed in the axial (z-directed) pulsed magnetic field gradients. This effect, which was attributed to some unfortunate resonant excitation of eddy currents, was not observed when transverse gradients were applied. It required special attention to the design of pulse shapes for gradients applied in the axial direction, but was ultimately brought under control. The acoustic flow tube was recently removed from the bore of the imaging magnet to facilitate a check of gradient strength calibrations. Quite unexpectedly, no ringing was observed in any of the pulsed field gradients. This change in behaviour was investigated after the thesis was completed but before the exmination date. A summary of this study is presented in App. D. It was discovered that mechanical vibrations driven by the z-gradient excite axial acoustic standing wave modes of the gas, which is confined by the two mercury pistons. The full implication of this effect on the outcome of the VPE-MRI experiments has not been evaluated,

but it is clearly something one would want to avoid in the future.

One of the unique features of the work described in this thesis is the focus on quantifying periodic gas motion, as opposed to steady flow. Under suitable conditions, however, acoustic fields can drive streaming currents (time-independent gas flow patterns). These streaming currents are notorious in the sense that they often materialize (unexpectedly) inside thermoacoustic heat engines and refrigerators, and (unless they are controlled or manipulated through design) tend to naturally degrade thermodynamic performance. In principle VPE-MRI could be used to reveal the presence of streaming currents as a time average asymmetry in a velocity field. Here it is noteworthy that Mastikhen *et al.* [108, 109] have used a variety of ^{1}H pulsed-field-gradient (motion-sensitized) MRI techniques to study 790 to 835 Hz acoustic standing waves in large diameter (10 to 20 cm) tubes filled with propane. Their experiments do not directly reveal the time and spatial distribution of the acoustic velocity field, but instead map out one-dimensional time-averaged velocity distributions averaged over transverse dimensions comparable to the tube diameter. Their data reveal structure in the velocity distribution that reflects the underlying standing wave pattern (in which $\lambda \approx 29$ cm), and which can be understood as a manifestation of Raleigh streaming.

Appendix A

Fourier Transform

At the heart of NMR spectroscopy and MRI lies the Fourier transform which is used to map the time domain signal from the acquisition to the frequency domain where analysis typically takes place. One of the many ways of writing the Fourier transform is

$$F(\mathbf{r}) = \int_{-\infty}^{\infty} f(\mathbf{k}) e^{i2\pi \mathbf{k} \cdot \mathbf{r}} \mathrm{d}\mathbf{k} \quad . \tag{A.1}$$

The inverse Fourier transform is then

$$f(\mathbf{k}) = \int_{-\infty}^{\infty} F(\mathbf{r}) e^{i2\pi \mathbf{k} \cdot \mathbf{r}} \mathrm{d}\mathbf{r} \quad . \tag{A.2}$$

For the discrete time and voltage data set acquired in a MRI experiment it is necessary to perform a discrete Fourier transform (DFT) or its inverse which are defined by

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N} kn} \quad k = 0, \ldots, N-1 \tag{A.3}$$

and the inverse discrete Fourier transform (iDFT) is then

$$x_n = \frac{1}{N} \sum_{k=0}^{N-1} X_k e^{\frac{2\pi i}{N} kn} \quad n = 0, \ldots, N-1 \quad . \tag{A.4}$$

The brute force approach to performing a DFT can be computationally instensive but there are a number of more efficient algorithms that take advantage of symmetries. For data sets that have 2^n points the number of calculations can be reduced significantly. The most famous of these algorithms is known as the fast Fourier transform or FFT [110].

Two Fourier transform pairs that are used in this thesis are given below with $\mathcal{F}$ indicating the Fourier transform operator. The rectangular function defined by

$$\text{rect}(x) = \begin{cases} |x| < \frac{1}{2} : a \\ |x| = \frac{1}{2} : 0.5 \\ |x| > \frac{1}{2} : 0 \end{cases} \tag{A.5}$$

has a Fourier transform

$$\mathcal{F}[\text{rect}(x)](k) = \frac{1}{a}\text{sinc}\left(\frac{k}{a}\right) \tag{A.6}$$

where

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \quad . \tag{A.7}$$

The Cylinder function

$$C(x, y) = \begin{cases} 1 & : \sqrt{x^2 + y^2} \leq a \\ 0 & : \sqrt{x^2 + y^2} > a \end{cases} \tag{A.8}$$

has a Fourier transform

$$\mathcal{F}\left[C(x, y)\right](k_x, k_y) = 2\pi a^2 \frac{J_1(\sqrt{k_x^2 + k_y^2}a)}{\sqrt{k_x^2 + k_y^2}a} \tag{A.9}$$

where $J_n(x)$ is a Bessel function of the first kind.

The convolution theorem states that the Fourier transform of the product of two integrable functions is equal to the convolution of the Fourier transforms of the individual functions. This is expressed as

$$\mathcal{F}[f \cdot g] = \mathcal{F}[f] * \mathcal{F}[g] \tag{A.10}$$

where the $*$ indicates the convolution operation. The convolution theorem is used in determining the bandwidth of RF pulses from the pulse envelope function.

Appendix B

DeltaEC model for the AM resonator

```
Acousto-mechanical resonator
!Created@11:11:59  25-Aug-2011 with DeltaEC version 6.2b3 under darwin,

!------------------------------- 0 -------------------------------
BEGIN
 3.0400E+5  a Mean P Pa
    1.3100  b Freq   Hz
  293.00    c TBeg   K
-3431.0     d |p|    Pa        G
  111.91    e Ph(p)  deg       G
 1.0270E-4  f |U|    m^3/s
    0.0000  g Ph(U)  deg
air          Gas type
!------------------------------- 1 -------------------------------
RPN piston volume flow rate
    0.0000 a G or T                5.2678E-5                    A
0b 0.0000064 * 2 * 3.14159 * =U1
!------------------------------- 2 -------------------------------
SURFACE start
 1.4150E-3  a Area   m^2            3431.0     A |p|    Pa
                                   -68.088     B Ph(p)  deg
                                   5.2669E-5   C |U|    m^3/s
                                   2.4850E-2   D Ph(U)  deg
                                   6.5749E-2   E Htot   W
ideal          Solid type          3.3682E-2   F Edot   W
!------------------------------- 3 -------------------------------
DUCT bellows
sameas   2a a Area   m^2            3431.0     A |p|    Pa
    0.5080  b Perim  m             -68.088     B Ph(p)  deg
 7.4900E-2  c Length m             4.5369E-5   C |U|    m^3/s
                                   -2.9186     D Ph(U)  deg
                                   6.5749E-2   E Htot   W
```

```
ideal           Solid type          3.2684E-2  F Edot   W
!------------------------------ 4 ------------------------------
DUCT piston compliance
 7.7490E-3  a Area   m^2     Mstr    3431.0     A |p|    Pa
    0.31205 b Perim  m       4a      -68.089    B Ph(p)  deg
 9.6000E-2  c Length m               2.0475E-5  C |U|    m^3/s
                                     -93.259    D Ph(U)  deg
                                     6.5749E-2  E Htot   W
ideal           Solid type          3.1790E-2  F Edot   W
!------------------------------ 5 ------------------------------
DUCT piston duct
 7.5400E-5  a Area   m^2     Mstr    3418.9     A |p|    Pa
 3.0536E-2  b Perim  m       5a      -68.263    B Ph(p)  deg
    1.4600  c Length m               2.4497E-5  C |U|    m^3/s
                                     -111.31    D Ph(U)  deg
                                     6.5749E-2  E Htot   W
ideal           Solid type          3.0601E-2  F Edot   W
!------------------------------ 6 ------------------------------
MINOR piston U-tube minor losses
sameas    5a a Area   m^2            3081.9     A |p|    Pa
 3000.0      b K+                    -62.086    B Ph(p)  deg
sameas    6b c K-                    2.4497E-5  C |U|    m^3/s
                                     -111.31    D Ph(U)  deg
                                     6.5749E-2  E Htot   W
                                     2.4651E-2  F Edot   W

!------------------------------ 7 ------------------------------
VESPEAKER piston U-tube
sameas    9a a Area   m^2            1319.0     A |p|    Pa
    1.0000  b R       ohms           -98.323    B Ph(p)  deg
    0.0000  c L       H              2.4495E-5  C |U|    m^3/s
    0.0000  d BLProd  T-m            -111.32    D Ph(U)  deg
    0.9190  e M       kg             6.5749E-2  E Htot   W
   50.200   f K       N/m            1.5741E-2  F Edot   W
    0.5230  g Rm      N-s/m              0.0000 G WorkIn W
    0.0000  h |V|     V                  0.0000 H Volts  V
    0.0000  i Ph(V)   deg                0.0000 I Amps   A
                                         0.0000 J Ph(Ze) deg
                                     2163.4     K |Px|   Pa
ideal           Solid type           139.04     L Ph(Px) deg
!------------------------------ 8 ------------------------------
BEGIN 3He-O2 gas mixture
sameas    0a a Mean P Pa
sameas    0b b Freq   Hz
sameas    0c c TBeg   K
sameas    7A d |p|    Pa
sameas    7B e Ph(p)  deg
sameas    7C f |U|    m^3/s
sameas    7D g Ph(U)  deg
```

```
heo2.tpf          Gas type
!------------------------------ 9 ------------------------------
DUCT acoustic flow tube up to the BC coil
 1.3273E-4  a Area   m^2      Mstr     1317.8    A |p|   Pa
 4.0841E-2  b Perim  m        9a        -98.407  B Ph(p) deg
    0.6000  c Length m                 2.4811E-5 C |U|   m^3/s
                                       -116.46   D Ph(U) deg
                                       1.5741E-2 E Htot   W
ideal          Solid type             1.5544E-2 F Edot   W
!------------------------------ 10 -----------------------------
DUCT acoustic flow tube from the BC coil
sameas  9a a Area   m^2      Mstr     1315.7    A |p|   Pa
 4.0840E-2  b Perim  m        10a       -98.524  B Ph(p) deg
    0.8800  c Length m                 2.5615E-5 C |U|   m^3/s
                                       -123.66   D Ph(U) deg
                                       1.5741E-2 E Htot   W
ideal          Solid type             1.5255E-2 F Edot   W
!------------------------------ 11 -----------------------------
MINOR end U-tube minor losses
sameas  10a a Area   m^2               1151.6    A |p|   Pa
 6000.0     b K+                        -94.628  B Ph(p) deg
sameas  11b c K-                       2.5615E-5 C |U|   m^3/s
                                       -123.66   D Ph(U) deg
                                       1.5741E-2 E Htot   W
                                       1.2896E-2 F Edot   W
!------------------------------ 12 -----------------------------
VESPEAKER end U-tube
 1.3273E-4  a Area   m^2               1500.2    A |p|   Pa
sameas  7b b R       ohms               147.81   B Ph(p) deg
sameas  7c c L       H                 2.5614E-5 C |U|   m^3/s
sameas  7d d BLProd  T-m               -123.66   D Ph(U) deg
    0.9130  e M       kg               1.5741E-2 E Htot   W
   50.200   f K       N/m              4.9259E-4 F Edot   W
    0.6660  g Rm      N-s/m               0.0000 G WorkIn W
    0.0000  h |V|     V                   0.0000 H Volts  V
    0.0000  i Ph(V)   deg                 0.0000 I Amps   A
                                          0.0000 J Ph(Ze) deg
                                       2275.0    K |Px|   Pa
ideal          Solid type              121.15    L Ph(Px) deg
!------------------------------ 13 -----------------------------
BEGIN air
sameas   0a a Mean P Pa
sameas   0b b Freq   Hz
sameas   0c c TBeg   K
sameas  12A d |p|    Pa
sameas  12B e Ph(p)  deg
sameas  12C f |U|    m^3/s
sameas  12D g Ph(U)  deg
```

```
air          Gas type
!----------------------------- 14 -------------------------------
DUCT end duct
sameas   5a a Area   m^2       Mstr      1516.8     A |p|    Pa
 3.0536E-2  b Perim  m         14a        147.64    B Ph(p)  deg
    1.4600  c Length m                    2.2054E-5 C |U|    m^3/s
                                         -123.12    D Ph(U)  deg
                                          4.9259E-4 E Htot   W
ideal           Solid type                2.1998E-4 F Edot   W
!----------------------------- 15 -------------------------------
DUCT end compliance
sameas   4a a Area   m^2       Mstr      1516.8     A |p|    Pa
    0.31205 b Perim  m         15a        147.64    B Ph(p)  deg
sameas   4c c Length m                    5.9610E-8 C |U|    m^3/s
                                          147.64    D Ph(U)  deg
                                          4.9259E-4 E Htot   W
ideal           Solid type                4.5208E-5 F Edot   W
!----------------------------- 16 -------------------------------
SURFACE end surface
sameas   4a a Area   m^2                  1516.8     A |p|    Pa
                                          147.64    B Ph(p)  deg
                                          1.0072E-15 C |U|    m^3/s
                                           61.836   D Ph(U)  deg
                                          4.9259E-4 E Htot   W
ideal           Solid type                5.5936E-14 F Edot   W
!----------------------------- 17 -------------------------------
HARDEND End
    0.0000 a R(1/z)          =17G       1516.8     A |p|    Pa
    0.0000 b I(1/z)          =17H        147.64    B Ph(p)  deg
    0.0000 c Htot   W                    1.0072E-15 C |U|    m^3/s
                                          61.836   D Ph(U)  deg
                                         4.9259E-4 E Htot   W
                                         5.5936E-14 F Edot   W
                                         7.7830E-15 G R(1/z)
                                        -1.0600E-13 H I(1/z)
! The restart information below was generated by a previous run
! and will be used by DeltaEC the next time it opens this file.
guessz  0d   0e
xprecn  3.2397E-2   2.4512E-3
targs   17a  17b
mstr-slave 6 4 -2 5 -2 9 -2 10 -2 14 -2 15 -2
! Plot start, end, and step values.  May be edited if you wish.
! Outer Loop:                        | Inner Loop .
```

Appendix C

^{3}He-O$_2$ External Fluid File

The DELTAEC model of the acousto-mechanical resonator presented in App. B relies on an external fluid file to define the thermophysical properties of the ^{3}He-O$_2$ mixture in the acoustic flow tube. External fluid filenames for pure gases or gas mixtures with constant partial pressures must end with a <.tpf> extension, and are known simply as tpf files. Each gas property is defined in terms of an expression of the form

$$\text{gas property} = C_0 + C_1 \frac{p_m}{T_m + p_m C_2} + C_3 T_m + C_4 T_m^2 + C_5 T_m^{C_6} + C_7 p_m^2 T_m^{C_8} + C_9 p_m \quad \text{(C.1)}$$

where T_m and p_m represent the mean temperature and pressure in the DELTAEC model. The tpf file consists of five lines of space-delimited coefficients, each corresponding to C_0, $C_1,\ldots$, C_9 for a particular property. Trailing omitted coefficients are automatically set to zero. The order in which the thermophysical properties must be specified is mean density ρ_m, isobaric specific heat capacity c_p, thermal conductivity k, square of the sound speed a^2, and dynamic viscosity μ. Comment lines may be inserted into the file using a leading exclamation point. A sample tpf file corresponding to the primary ^{3}He-O$_2$ gas mixture used in this work is listed below (referred to as heo2.tpf in App. B). Various calculated thermophysical properties of this gas mixture (including some that are not required for the tpf file) are then given for reference in Table C.1.

```
! heo2.tpf
! External fluid; 1.7atm 3He, 1.3atm O2 mixture.
! Density, rho (kg/m^3):
```

```
0. 0.001872
! Isobaric heat capacity, cp (J/kg/K):
1574.
! Thermal conductivity, k0 (W/m/K):
0. 0. 0. 0. 0. 0.01419 0.279
! Square speed of sound, a^2 (m^2/s^2):
0. 0. 0. 816.0
! Viscosity, mu (kg/s/m):
0. 0. 0. 0. 0. 3.40e-7 0.725
```

The transport properties of the gas mixture were calculated using the procedure described in Chapters 7 and 8 of *Molecular Theory of Gases and Liquids* by Hirschfelder, Curtiss and Bird [72], which is based on the Chapman-Enskog theory. In principle, the Chapman-Enskog theory is suitable for monatomic gases and gas mixtures. It can, however, be extended to obtain the viscosity and diffusion coefficients for pure polyatomic gases and for mixtures of gases containing polyatomic molecules, such as O_2. In calculating the thermal conductivity of a polyatomic gas, the Eucken correction factor must be applied to account for the energy stored in the internal molecular degrees of freedom.

The calculation proceeds by first assigning to each constituent of the mixture (i.e. ³He and O_2) appropriate Lennard-Jones (6-12) pair interaction potential parameters (a collision diameter σ and well depth ϵ). These are taken from Table I-A in App. A of Ref. [72], and allow one to define reduced temperatures $T^* = kT/\epsilon$ for ³He-³He and O_2-O_2 pairs. The specific values used in this work are $\sigma = 2.57$ Å and $\epsilon/k = 10.22$ K for ³He and $\sigma = 3.43$ Å and $\epsilon/k = 113$ K for O_2, where k is the Boltzmann constant. These values are appropriate to the temperature range 80 K $< T <$ 300 K. Note that the values assigned to ³He-³He interactions are in fact those specified in Ref. [72] for ⁴He-⁴He interactions. The arithmetic mean of the two collision diameters and the geometric mean of the two well depths are then used as effective Lennard-Jones interaction parameters for ³He-O_2 collisions, which in turn enable one to define a third reduced temperature. With these reduced temperatures in hand, appropriate temperature-dependent collision integrals $\Omega^{(l,s)\star}$ (for $(l,s) = (1,1)$, $(1,2)$, $(1,3)$, and $(2,2)$) are extracted from Table I-M in App. A of Ref. [72], where they are listed for $0.3 \leq T^* \leq 400$. Interpolation was accomplished by

fitting tabulated values for each collision integral to an expression of the form

$$\Omega^{(l,s)\star} = y_0 + A_1 e^{(T^*/C_1)} + A_2 e^{(T^*/C_2)} + A_3 e^{(T^*/C_3)} + A_4 e^{(T^*/C_4)} \tag{C.2}$$

over the range $0.85 \leq T^* \leq 100$.

The dynamic viscosity μ of the ^{3}He and O$_2$ mixture was calculated using Eqs. (8.2-18), (8.2-21) and (8.2-22) of Ref. [72]. The thermal conductivity was similarly calculated using Eqs. (8.2-31), (8.2-33), (8.2-35) and (8.2-36). Here, Eq. (8.2-33) includes the Eucken correction factor needed to account for the fact that O$_2$ is polyatomic. Values for this correction factor were extracted from Table 8.4-10 of Ref. [72]. Finally, the mass diffusion coefficient D and thermal diffusion ratio k_T were calculated using Eqs. (8.2-44) and (8.2-50), respectively.

Table C.1: Calculated thermophysical properties of a mixture of 1.7 atm ^{3}He and 1.3 atm O$_2$ at T=293 K.

property	value
ρ	1.94 kg/m^3
c_p	1574 J/kg/K
k_0	0.0692 W/m/K
a	489 m/s
μ	2.089$\times$10^{-5} kg/s/m
D	0.25 cm^2/s
k_T	-0.12

Calculated values of viscosity and thermal conductivity for pure ^{3}He were compared to similar calculations performed by Slaman and Aziz [111]; agreement to within 1% was observed over the temperature range of 250 to 330 K. Similarly, calculated values of viscosity and thermal conductivity for pure O$_2$ were compared to those obtained by Boushehri *et al.* [112] and Roder [113], respectively; again agreement was observed at the 1% level in both cases. The viscosity and thermal conductivity of the ^{3}He-O$_2$ mixture were then calculated at five temperatures spanning the range 250 to 330 K. These calculated values were fit to

$$f(T) = C_5 T^{C_6} \tag{C.3}$$

in order to extract the coefficients C_5 and C_6 in Eq. (C.1). This simple parameterization reproduces the calculated transport properties to better than 1% over the temperature range 250-300 K.

Appendix D

Addendum

Section 4.8 describes the observation of exponentially decaying oscillations in the z magnetic field gradient following a 1 ms duration, half-sine shaped gradient pulse. No similar oscillations were observed in response to x- or y-directed magnetic field gradient pulses. Throughout the thesis it is assumed that these oscillations are somehow associated with eddy currents in some conducting structure or shim coils of the MRI magnet. After submission of the thesis to the examining committee, a series of additional experiments were performed to search for the source of these oscillations, and are summarized here.

The ^{3}He-O$_2$ gas mixture was removed from the acoustic flow tube, and this volume was back-filled with air to 1 atm pressure. The pulsed magnetic field gradient calibrations using the gradiometer described in Sec. 4.8 were then repeated, with the goal of confirming the magnetic field gradient calibration performed using the water phantom. To my surprise, the ringing that had previously been observed in response to the pulsed z-directed magnetic field gradient had disappeared. The only difference between the two experiments was the change in gas composition in the acoustic flow tube (from a 3 atm mixture of ^{3}He and O$_2$ to 0.96 atm of air). At this point it was surmised that the ringing must have been related to an unfortunate mechanically driven acoustic oscillation in the gas confined to the acoustic flow tube, rather than eddy currents in the magnet. That is, the linear length of gas trapped between the two free surfaces of the (massive) mercury slugs is 1.48 m, forming a half-wavelength acoustic resonator. Changing the gas composition changes the sound speed, and hence the frequencies at which this column resonates.

To test this hypothesis, the acoustic flow tube pressure transducer signal was monitored

with a lockin amplifier as the z magnetic field gradient was driven with a continuous sine-wave excitation. The frequency of the excitation was scanned over the range 100 to 1200 Hz in search of resonant responses in the pressure transducer signal. The acoustic flow tube was first filled with pure O_2 for which the sound speed at room temperature is 326 m/s and the fundamental acoustic mode is expected to be at 110 Hz. It was then filled with 1 atm of O_2 and 2 atm of ^{4}He for which the sound speed at room temperature is calculated to be 532 m/s and the fundamental acoustic mode is expected to be at 180 Hz. The fundamental and its first two overtones were clearly identified as high Q resonances (of order 30) for both gas mixtures. A number of mechanical resonances in the magnet and gradient coil structure were also observed; these did not shift in frequency when the gas mixture was changed. Acoustic resonances from modes near 1 kHz were also observed for both gas mixtures.

The coupling between the applied magnetic field gradients and the acoustic modes of the flow tube is mechanical, that is, momentary forces are exerted on the magnet structure as pulsed currents flow through the gradient coils. These vibrations are in turn mechanically coupled to the acoustic flow tube, exciting resonant modes of the gas. Reducing the mechanical coupling between the magnet and the acoustic flow tube decreases the amplitude of the resonant excitations. The symmetry of the z magnetic field gradient is such that it drives larger amplitude vibrations in the axial direction than do the x or y field gradient coils, and therefore couples strongly to acoustic modes in the confined gas.

In this light, it is believed that the 1 ms duration bipolar field gradients used for velocity phase encoding coupled strongly to an acoustic mode in the ^{3}He-O_2 gas mixture, causing an *apparent* decaying oscillation in the associated magnetic field gradient. Replacing the gas with air changed the sound speed and shifted the resonance far enough to effectively eliminate the coupling. This coincidence is unfortunate in the sense that the shape of the VPE gradients was engineered (as described in Sec. 4.8) to eliminate ringing in the gradiometer signal. This procedure was almost certainly unnecessary. At the same time, it is the first moment of the bipolar VPE gradients that matters for velocity determination. On this basis, I estimate that the modulation introduced into the bipolar gradient pulses will result in an overestimation of the FOS (see Eq. (2.34)) by approximately 3%.

This observation highlights the importance of making sure that the acoustic flow tube is mechanically isolated from the imaging magnet. It also suggests that the timing of magnetic

field gradient pulses could be chosen to minimize their spectral content in the vicinity of known acoustic modes. Viewed in another light, the fact that acoustic modes in the confined gas can be excited and detected suggests that one could try to exploit measurements of sound speed and damping as yet another probe of gas mixture properties.

Bibliography

[1] K. K. Kwong, J. W. Belliveau, D. A. Chesler, I.E. Goldberg, R.M. Weisskoff, B.P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen, R. Turner, H. Cheng, T. J. Brady, and B. R. Rosen. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences of the United States of America*, 89:5675–5679, 1992.

[2] P. A. Bandettini, A. Jesmanowicz, E. C. Wong, and J. S. Hyde. Processing strategies for time-course data sets in functional MRI of the human brain. *Magnetic Resonance in Medicine*, 30:161–173, 1993.

[3] A. MacKay, K. Whittall, J. Adler, D. Li, D. Paty, and D Graeb. In-vivo visualization of myelin water in brain by magnetic-resonance. *Magnetic Resonance in Medicine*, 31(6):673–677, 1994.

[4] J. R. Singer. Blood flow rates by nuclear magnetic resonance measurements. *Science*, 130(3389):1652–1653, 1959.

[5] D. G. Nishimura. Time-of-flight MR angiography. *Magnetic Resonance in Medicine*, 14:194–201, 1990.

[6] D. L. Dumoulin, S. P. Souza, M. F. Walker, and W. Wagle. Three-dimensional phase contrast angiography. *Magnetic Resonance in Medicine*, 9, 1989.

[7] F. A. Bovey, R. E. Cais, J. W. Jelinski, F. C. Schilling, W. H. Starnes, and A. E. Tonelli. Structural and dynamic characterization of polymers by C-13 and F-19 NMR. *Advances in Chemistry Series*, 203:421–439, 1983.

[8] S. D. Beyea, S. L. Codd, D. O. Kuethe, and E. Fukushima. Studies of porous media by thermally polarized gas NMR: current status. *Magnetic Resonance Imaging*, 21:201–205, 2003.

[9] B. Sun and K.-J. Dunn. A global inversion method for multi-dimensional NMR logging. *Journal of Magnetic Resonance*, 172(1):152–160, 2005.

[10] B. M. Goodson. Nuclear magnetic resonance of laser-polarized noble gases in molecules, materials, and organisms. *Journal of Magnetic Resonance*, 155:157–216, 2002.

[11] M. S. Albert, D. Cates, G, B. Driehuys, W. Happer, B. Saam, C. S. Springer, and A. Wishnia. Biological magnetic-resonance-imaging using laser polarized Xe-129. *Nature*, 370:199–201, 1994.

[12] H. Middleton, R. D. Black, B. Saam, G.D. Cates, G. P. Cofer, R. Guenther, W. Happer, L. W. Hedlund, G. A. Johnson, K. Juvan, and J. Swartz. MR-imaging with hyperpolarized He-3 gas. *Magnetic Resonance in Medicine*, 33:271–275, 1995.

[13] H. E. Moller, X. J. Chen, B. Saam, K. D. Hagspiel, G. A. Johnson, T. A. Altes, E. E. de Lange, and H. U. Kauczor. MRI of the lungs using hyperpolarized noble gases. *Magnetic Resonance in Medicine*, 47:1029–1051, 2002.

[14] L. de Rochefort. *Imagerie dynamique et vélocimétrie IRM des gaz hyperpolarisés*. PhD thesis, Université Paris-Sud XI, 2006.

[15] L. de Rochefort, X. Maître, R. Fodil, L. Vial, B. Louis, D. Isabey, C. Croce, L. Darrasse, B. Apiou, G. Caillibotte, J. Bittoun, and E. Durand. Phase-contrast velocimetry with hyperpolarized ^{3}He for in vitro and in vivo characterization of airflow. *Magnetic Resonance in Medicine*, 55:1318–1325, 2006.

[16] E. Brunner, M. Haake, L. Kaiser, A. Pines, and J. A. Reimer. Gas flow MRI using circulating laser-polarized ^{129}Xe. *Journal of Magnetic Resonance*, 138:155–159, 1999.

[17] Song-I Han, Kimberly L. Pierce, and Alexander Pines. NMR velocity mapping of gas flow around solid objects. *Phys. Rev. E*, 74(1):016302, Jul 2006.

[18] B. Newling, C. C. Poirier, Y. Zhi, J. A. Rioux, A. J. Coristine, D. Roach, and B. J. Balcom. Velocity imaging of highly turbulent gas flow. *Physical Review Letters*, 93(15):154503, 2004.

[19] N. Rott. Thermoacoustics. *Advances in Applied Mechanics*, 20:135–175, 1980.

[20] G. W. Swift. Thermoacoustic engines and refrigerators. *Physics Today*, 48:22–28, 1995.

[21] G. W. Swift. Thermoacoustic engines. *Journal of the Acoustical Society of America*, 84:1145–1180, 1988.

[22] K. T. Feldman. Review of the literature on Rijke thermoacoustic phenomena. *Journal of Sound and Vibration*, 7:83–89, 1968.

[23] K. T. Feldman. Review of the literature on Sondhauss thermoacoustic phenomena. *Journal of Sound and Vibration*, 7:71–82, 1968.

[24] Lord Rayleigh. *The Theory of Sound*, volume 2. Dover, New York, 1945.

[25] G. Kirchhoff. Über den Einfluss der Wärmeleitung in einem Gase auf die Schallbewegung. *Poggendorffs Annelen*, 134:177, 1868.

[26] H. A. Kramers. Vibrations in a gas column. *Physica*, 75:971, 1949.

[27] N. Rott. Damped and thermally driven acoustic oscillations in wide and narrow tubes. *Zeitschrift fur Angewandte Mathematik und Physik*, 20:230–243, 1969.

[28] N. Rott. Thermally driven acoustic oscillations. Part II: Stability limit for helium. *Zeitschrift fur Angewandte Mathematik und Physik*, 24:54–72, 1973.

[29] N. Rott. Thermally driven acoustic oscillations, Part III: Second-order heat flux. *Zeitschrift fur Angewandte Mathematik und Physik*, 26:43–49, 1975.

[30] N. Rott. Thermally driven acoustic oscillations, Part IV: Tubes with variable cross-section. *Zeitschrift fur Angewandte Mathematik und Physik*, 27:197–224, 1976.

[31] N. Rott. Thermally driven acoustic oscillations, Part V: Gas-liquid oscillations. *Zeitschrift fur Angewandte Mathematik und Physik*, 27:325–334, 1976.

[32] U. A. Müller and N. Rott. Thermally driven acoustic oscillations, Part VI: Excitation and power. *Zeitschrift fur Angewandte Mathematik und Physik*, 34:609–626, 1983.

[33] R. Radebaugh. A review of pulse tube refrigeration. *Advances in Cryogenic Engineering*, 35:1191–1205, 1990.

[34] T. Hofler. *Thermoacoustic refrigerator design and performance.* PhD thesis, Physics Department, University of California, San Diego, 1986.

[35] P. H. Ceperley. A pistonless Stirling engine: The traveling wave heat engine. *Journal of the Acoustical Society of America*, 66(5):1508–1513, 1979.

[36] S. Backhaus and G. W. Swift. A thermoacoustic Stirling heat engine. *Nature*, 399:335–338, 1999.

[37] S. Backhaus and G. W. Swift. A thermoacoustic-Stirling heat engine: Detailed study. *Journal of the Acoustical Society of America*, 107(6):3148–3166, 2000.

[38] T. Yazaki, A. Iwata, T. Maekawa, and A. Tominaga. Traveling wave thermoacoustic engine in a looped tube. *Physical Review Letters*, 81(15):3128–3131, 1998.

[39] A. T. A. M. de Waele, P. P. Seijaert, and J. Gijzen. Thermodynamical aspects of pulse tubes. *Cryogenics*, 37:313–324, 1998.

[40] A. T. A. M. de Waele, P. P. Seijaert, and J. J. Koning. Thermodynamical aspects of pulse tubes II. *Cryogenics*, 38:329–335, 1998.

[41] G. Huelsz and F. López-Alquicira. Hot-wire anemometry in acoustic waves. *Experiments in Fluids*, 30:283–285, 2001.

[42] G. Huelsz, F. López-Alquicira, and E. Ramos. Velocity measurements in the oscillatory boundary layer produced by acoustic waves. *Experiments in Fluids*, 32(6):612–15, 2002.

[43] J. R. Castrejón-Pita, A. A. Castrejón-Pita, G. Huelsz, and R. Tovar. Experimental demonstration of the Rayleigh acoustic viscous boundary layer theory. *Physical Review E*, 73(3):36601–1–5, 2006.

[44] H. Bailliet, P. Lotton, M. Bruneau, V. Gusev, J.C. Valière, and B. Gazengel. Acoustic power flow measurement in a thermoacoustic resonator by means of laser doppler anemometry (L.D.A.) and microphonic measurement. *Applied Acoustics*, 60:1–11, 2000.

[45] L. Shi, Z. Yu, and A. J. Jaworski. Vortex shedding flow patterns and their transitions in oscillatory flows past parallel-plate thermoacoustic stacks. *Experimental Thermal and Fluid Science*, 34:954–965, 2010.

[46] L. Shi, Z. Yu, and A. J. Jaworski. Application of laser-based instrumentation for measurement of time-resolved temperature and velocity fields in the thermoacoustic system. *International Journal of Thermal Sciences*, 49:1688–1701, 2010.

[47] J. Granwehr, E. Harel, S. Han, S. Garcia, A. Pines, P. N. Sen, and Y. Q. Song. Time-of-flight flow imaging using NMR remote detection. *Physical Review Letters*, 95(7):075503, 2005.

[48] I. V. Koptyug, A. V. Matveev, and S. A. Altobelli. NMR studies of hydrocarbon gas flow and dispersion. *Applied Magnetic Resonance*, 22:187–200, 2002.

[49] S.-I. Han, K. L. Pierce, and A. Pines. NMR velocity mapping of gas flow around solid objects. *Physical Review E*, 74:016302, 2006.

[50] I. V. Koptyug, S. A. Altobelli, E. Fukushima, A. V. Matveev, and R. Z. Sagdeev. Thermally polarized H NMR microimaging studies of liquid and gas flow in monolithic catalysts. *Journal of Magnetic Resonance*, 147:36–42, 2000.

[51] F. Kober, B. Koenigsberg, V. Belle, M. Viallon, J. L. Leviel, A. Delon, A. Ziegler, and M. Décorps. NMR imaging of thermally polarized helium-3 gas. *Journal of Magnetic Resonance*, 138:308–312, 1999.

[52] P. T. Callaghan. *Principles of Nuclear Magnetic Resonance Microscopy*. Oxford University Press Inc., 1991.

[53] G. W. Swift. *Thermoacoustics: A Unifying Perspective for some Engines and Refrigerators*. Acoustical Society of America, 2002.

[54] P. C. Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 242:190–191, 1973.

[55] P. Mansfield and P. K. Grannell. NMR 'diffraction' in solids? *Journal of Physics C: Solid State Physics*, 6:L422–L426, 1973.

[56] R. C. Weast and M. J. Astle, editors. *CRC Handbook of Chemistry and Physics*. CRC Press, Inc., Boca Raton, Florida, 59th edition, 1978.

[57] W. Heil, H. Humblot, E. Otten, M. Schafer, R. Sarkau, and M. Leduc. Very long nuclear relaxation times of spin polarized helium 3 in metal coated cells. *Physics Letters A*, 201:337–343, 1995.

[58] C. E. Hayes, W. A. Edelstein, J. F. Schenck, O. M. Mueller, and M. Eash. An efficient, highly homogeneous radio frequency coil for whole-body NMR imaging at 1.5T. *Journal of Magnetic Resonance*, 63:622–628, 1985.

[59] D. Hoult and R. Richards. Signal-to-noise ratio of nuclear magnetic-resonance experiment. *Journal of Magnetic Resonance*, 24(1):71–85, 1976.

[60] D. R. Bailes and D. J. Bryant. NMR imaging. *Contemporary Physics*, 25(5):441–475, 1984.

[61] E. L. Hahn. Spin echoes. *Physical Review*, 80(4):580–594, 1950.

[62] Lord Rayleigh. Investigations in optics, with special reference to the spectroscope: Resolving, or separating, power of optical instruments. *Philosophical Magazine*, 8:261–274, 1879.

[63] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Physical Review*, 94(3):630–638, 1954.

[64] E. L. Hahn. Detection of sea-water motion by nuclear precession. *Journal of Geophysical Research*, 65(2):776–777, 1960.

[65] E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1):288–292, 1965.

[66] W. S. Price and P. W. Kuchel. Effect of nonrectangular field gradient pulses in the Stejskal and Tanner (diffusion) pulse sequence. *Journal of Magnetic Resonance*, 94:133–139, 1991.

[67] H. L. F. Helmholtz. Über den einfluss der reibung in der luft auf die schallbewegung. *Verhandlungen des Natuhisorisch-Medizinshcen Vereins zu Heidelberg*, 3:16–20, 1865. presented February 27, 1863.

[68] A. Bejan. *Advanced Engineering Thermodynamics*. Wiley, second edition, 1997.

[69] L. D. Landau and E. M. Lifshitz. *Fluid Mechanics*. Pergamon, 1982.

[70] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders. *Fundamentals of Acoustics*. John Wiley and Sons, 4^{th} edition, 1999.

[71] G. W. Swift and P. S. Spoor. Thermal diffusion and mixture separation in the acoustic boundary layer. *Journal of the Acoustical Society of America*, 106:1794–1800, 1999.

[72] J. O. Hirschfelder, C. F. Curtiss, and R. B. Bird. *Molecular Theory of Gases and Liquids*. Wiley, New York, 1964.

[73] T. Yazaki, A. Tominaga, and Y. Narahara. Experiments on thermally driven acoustic oscillations of gaseous helium. *Journal of Low Temperature Physics*, 41(1/2):45–60, 1980.

[74] P. H. Ceperley. Gain and efficiency of a short traveling wave heat engine. *Journal of the Acoustical Society of America*, 77(3):1239–1244, 1985.

[75] W. P. Arnott, H. E. Bass, and R. Raspet. General formulation of thermoacoustics for stacks having arbitrarily shaped pore cross-sections. *Journal of the Acoustical Society of America*, 90:3228–3237, 1991.

[76] A. J. Jaworski, X. Mao, X. Mao, and Z. Yu. Entrance effects in the channels of the parallel plate stack in oscillatory flow conditions. *Experimental Thermal and Fluid Science*, 33:495–502, 2009.

[77] M. E. Hayden and G. W. Swift. Thermoacoustic relaxation in a pin-array stack. *Journal of the Acoustical Society of America*, 102(5):2714–2722, 1997.

[78] W. Ward and G. W. Swift. Design environment for low-amplitude thermoacoustic engines. *Journal of the Acoustical Society of America*, 95(6):3671–3672, 1994.

[79] W. Dean. Note on the notion of fluid in a curved pipe. *Philosophical Magazine*, 4(20):208–223, 1927.

[80] W. Dean. The stream-line motion of fluid in a curved pipe. *Philosophical Magazine*, 5(30):673–695, 1928.

[81] S. A. Berger, L. Talbot, and L.-S. Yao. Flow in curved pipes. In M. Van Dyke and J. V. Wehausen, editors, *Annual Review of Fluid Mechanics*, volume 15, pages 461–512. Annual Reviews Inc., 1983.

[82] D. J. McConalogue and R. S. Srivastava. Motion of fluid in a curved tube. *Procedings of the Royal Society of London Series A*, 307:37–53, 1968.

[83] J. R. Olson and G. W. Swift. Energy dissipation in oscillating flow through straight and coiled pipes. *Journal of the Acoustical Society of America*, 100(4):2123–2131, 1996.

[84] W. H. Lyne. Unsteady viscous flow in a curved pipe. *Journal Of Fluid Mechanics*, 45:13–31, 1971.

[85] K. Sudo, M. Sumida, and R. Yamane. Secondary motion of fully-developed oscillatory flow in a curved pipe. *Journal Of Fluid Mechanics*, 237:189–208, 1992.

[86] B. Saam, W. Happer, and H. Middleton. Nuclear relaxation of ^{3}He in the presence of O_2. *Physical Review A*, 52:862, 1995.

[87] G. Archibald, E. Brief, C. Lei, T. Pausak, and M. E. Hayden. Nuclear relaxation of ^{3}He in the presence of NO. *Physical Review A*, 73(2):022721, 2006.

[88] R. S. Timsit, J. M. Daniels, and A.D. May. Nuclear relaxation of ^{3}He gas on various solid surfaces. *Canadian Journal of Physics*, 49(5):560–575, 1971.

[89] Penn State Hershey College of Medicine Center for NMR Research. Software: Birdcage builder. `http://pennstatehershey.org/web/nmrlab/resources/software/birdcagebuilder`, July 2011.

[90] J. Crank. *The Mathematics of Diffusion*. Oxford University Press, 1975.

[91] F. J. Norton. Helium diffusion through glass. *Journal of the American Ceramic Society*, 36(3):90–96, 1953.

[92] N. R. Newbury, A. S. Barton, G. D. Cates, W. Happer, and H. Middleton. Gaseous ^{3}He magnetic dipolar spin relaxation. *Physical Review A*, 48(6):4411–4420, Dec 1993.

[93] R. L. Gamblin and T. R. Carver. Polarization and relaxation processes in ^{3}He gas. *Physical Review*, 138(4A):A946–A960, May 1965.

[94] L. D. Schearer and G. K. Walters. Nuclear spin-lattice relaxation in the presence of magnetic-field gradients. *Physical Review*, 139(5A):A1398–A1402, Aug 1965.

[95] A. P. French. *Vibrations and Waves*. W. W. Norton & Company, New York, 1971.

[96] A. Petculescu and L. A. Wilen. Oscillatory flow in jet pumps: Nonlinear effects and minor losses. *Journal of the Acoustical Society of America*, 113(3):1282–1292, 2003.

[97] U. H. Kurzweg, E. R. Lindgren, and B. Lothrop. Onset of turbulence in oscillating flow at low Womersley number. *Physics Of Fluids A*, 1:1972–1975, 1989.

[98] I. E. Idelchik. *Handbook of Hydraulic Resistance*. Begell House, New York, 3rd edition, 1994.

[99] M. E. Hayden and W. N. Hardy. Technique for measuring magnetic filling factors with applications to cryogenic magnetic resonance experiments. *Review of Scientific Instruments*, 67:1905–1911, 1996.

[100] C. A. Michal. Magnetic field homogeneity: A new approach to orthogonalizing and optimizing shim gradients. *Journal of Magnetic Resonance*, 185:110–117, 2007.

[101] G. Stran. *Linear Algebra and is Applications*. Harcourt Brace Jovanovich Inc., San Diego, third edition, 1988.

[102] T. E. Conturo and Smith G. D. Signal-to-noise in phase angle reconstruction: Dynamic range extension using phase reference offsets. *Magnetic Resonance in Medicine*, 15:420–437, 1990.

[103] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Dover, New York, 1965.

[104] H. F. Mark, N. G. Gaylord, and N. M. Bikales, editors. *Encyclopedia of Polymer Science and Technology*, volume 13. Wiley, New York, 1970.

[105] B. Newling. Gas flow measurements by NMR. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 52:31–48, 2007.

[106] G. W. Swift and D. A. Geller. Continuous thermoacoustic mixture separation. *Journal of the Acoustical Society of America*, 120:2648–2657, 2006.

[107] Y. V. Chang, S. E. Haywood, J. C. Woods, and M. S. Conradi. Spin sorting: apparent longitudinal relaxation without spin transitions. *Chemical Physics Letters*, 437:126–131, 2007.

[108] I. Mastikhin, B. Newling, and S. Culligan. MRI detection of acoustic streaming in gases. In A. A. Atchley, V. W. Sparrow, and R. M. Keolian, editors, *Innovations in Nonlinear Acoustics: 17th International Symposium on Nonlinear Acoustics*, volume 838, pages 469–471. AIP Conference Proceedings, 2006.

[109] I. V. Mastikhin and B. Newling. MRI of fluids in strong acoustic fields. In S. L. Gould and J. D. Seymour, editors, *Magnetic Resonance Spectroscopy*. Wiley-VCH Verlag Gmbh & Co., Weinheim Germany, 2009.

[110] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C (2nd ed.): The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992.

[111] M. J. Slaman and R. A. Aziz. Accurate transport properties and second virial coefficients for helium based on a state-of-the art interatomic potential. *International Journal of Thermophysics*, 12:837–854, 1991.

[112] A. Boushehri, J. Bzowski, J. Kestin, and E. A. Mason. Equilibrium and transport properties of eleven polyatomic gases at low density. *Journal of Physical and Chemical Reference Data*, 16(3):445–466, 1987.

[113] H. M. Roder. The thermal conductivity of oxygen. *Journal Of Research of the National Bureau of Standards*, 87(4):279–346, 1982.