DENSITY OF STATES OF ELASTIC WAVES IN A

STRONGLY SCATTERING POROUS "MESOGLASS"

by

William Kurt Hildebrand

A Thesis submitted to the Faculty of Graduate Studies

in partial fulfilment of the requirements of the degree of

Master of Science

Department of Physics and Astronomy

University of Manitoba

Winnipeg, Canada

August 2009

Abstract

The density of states of elastic waves in a porous amorphous "mesoglass" has been measured in the strong-scattering regime. Samples were constructed by sintering glass beads percolated on a random lattice. This structure was investigated via x-ray tomography, and fractal behaviour was observed with fractal dimension $D = 2.6$. Using sufficiently small samples, the individual modes of vibration could be resolved and counted in the Fourier transform of each transmitted ultrasonic pulse. A statistical treatment of the data, designed to account for the possibility of missing modes, was developed, yielding a robust method for measuring the density of states. In the strong-scattering regime, the data are in good agreement with a simple model based on mode conservation, though the density of states significantly exceeds the predictions of the Debye approximation at low frequencies. At intermediate frequencies, an average density of states of $\mathcal{D}_l = 47.1 \pm 0.3$ MHz^{-1} mm^{-3} was found, with a frequency dependence of $f^{0.01 \pm 0.04}$.

Acknowledgements

First and foremost, I would like to thank Professor John H. Page for his excellent direction and attentive supervision of this thesis. His patience, support and attention to detail have been an integral part of my education, and I am proud to have worked with him.

I would also especially like to thank Laura Cobus, without whose help this project may not have been completed.

My co-workers in the Ultrasonics Research Lab, along with Gilles Roy, were invaluable in day-to-day problem solving, and made for a pleasant working environment.

Thanks to Johan van Lierop and Derek Oliver for their participation on my thesis committee.

Finally, sincere thanks to my wonderful friends and family, who always managed to keep me smiling.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

The density of states is the number of normal modes that exist at a given energy (or frequency), and is a fundamental property of any system. In electronic systems, the density of states refers to the energy levels that can be occupied by the electrons, and this determines the nature of the system – conductors, semi-conductors and insulators all depend on the density of states for their conduction properties. In quantum systems, the density of states also determines the transition probability between states via Fermi's golden rule. The photon density of states (for electromagnetic waves) strongly influences the behaviour of photonic crystals, for example, and helps to explain phenomena such as blackbody radiation. In solid-state physics, the thermal properties of most systems are primarily determined by the phonon density of states. The density of states is, by nature, a statistical quantity, used to glean the properties of the entire system when the individual modes are too many or too complex to be accounted for individually. In highly disordered systems, the density of states becomes increasingly difficult to derive from first principles, motivating new studies that can investigate this ubiquitous physical quantity directly.

In the past thirty years, there has been a renewed interest in wave transport in strongly scattering random media. A huge variety of strong-scattering systems are being studied: light scattering in "milky" colloids, microwaves in resonant cavities, matter waves in Bose-Einstein condensates and ultra-cold atoms with random potentials, plasmons on disordered surfaces, seismic waves in the earth, and of

course, elastic and acoustic waves[1] in many different types of porous media, which are the focus of our research group. Several books have been written providing excellent in-depth descriptions of the theory of wave transport in strongly scattering media, including some experimental results [Sheng 1990, 1995; Van Tiggelen and Skipetrov 2003; Akkermans and Montambaux 2007], and applications have been found, such as *diffusing wave spectroscopy* (DWS) for light, and its acoustic and seismic counterparts, diffusing *acoustic* wave spectroscopy (DAWS) and *coda wave interferometry* (CWI). There has been much discussion lately about Anderson localization in strongly scattering systems [Lagendijk *et al.* 2009], and given the recent successful observation of this phenomenon in an three-dimensional elastic network [Hu *et al.* 2008], characteristics of elastic wave transport in such systems are of considerable interest.

The principles of wave transport in strongly scattering systems are applicable for many types of waves and systems, spanning a huge range of length scales – from the sub-nanometre world of quantum mechanics to the kilometre-scale physics of seismology. Because of this, a better understanding of wave transport in strong-scattering media can provide a wealth of information to many disciplines, and the density of states is a fundamental part of this picture.

---

[1] Acoustic waves refer to the longitudinal (pressure) waves that occur in both fluid and solid systems; these are encountered daily in the form of sounds we hear, but also extend to the subsonic and ultrasonic ranges. Elastic waves also include transverse (shear) waves, which occur over a wide range of frequencies as well, but are only found in solids (and some very viscous fluids).

Wave transport in strongly scattering systems is often treated using the diffusion approximation, wherein the average transport of energy occurs diffusively, much like the familiar examples of heat flow in a solid or molecular diffusion in a gas. In this model, the phase of the wave is assumed to be sufficiently randomized by scattering that interference effects can be ignored in the ensemble average. This approximation works rather well to describe many features of multiply scattered waves, perhaps surprisingly, since interference is such an important aspect of wave transport. In media where the diffusion approximation is applicable, the density of states may exhibit anomalous behaviour, which is still not well-understood.

Percolation theory [Stauffer 1985] provides a useful tool for modelling disordered systems. One feature of percolated systems (above the percolation threshold) is that they exhibit fractal behaviour[1] over a range of length scales. There has been much discussion on the density of states in fractals and percolated structures [Nakayama *et al.* 1994]; however, some debate still remains, and relatively few experimental measurements have been made.

Percolation also suggests a practical method of constructing a model disordered system. By analogy to the random network of atoms that form a glass, a "mesoglass" is constructed by sintering glass spheres that are "percolated" on a random lattice by mixing with transitory iron spheres (which are removed after the sintering process). Thus, along with being a highly porous system which strongly scatters

---

[1] That is, the mass of the system scales with system size to a non-integer power, c.f. § 2.7.1.

elastic waves over a range of frequencies, this medium also serves as a mesoscopic model for an atomic glass.

The advantages of using elastic waves in a system such as this to study wave transport are many. The mesoglass is constructed so that the range of frequencies over which strong scattering occurs is easily accessible via piezoelectric transducers, and the field (including phase information) can be easily measured and recorded using an oscilloscope. The length scales in this mesoglass are large enough that the sample construction can be well-controlled, and the structure can be viewed with optical microscopy and x-ray tomography; while still being small enough that experiments can be done on a table-top. Most importantly for this thesis, the samples can be made small enough to resolve and count the individual modes of vibration, and thereby measure the density of states directly.

Wave transport in these mesoglasses has been studied extensively by our research group. Ultrasonic attenuation and velocity was studied in a similar system (though somewhat smaller in length scale), and the dependence of sample properties on the porosity of the material helped to lay much of the groundwork for my experiments [Pachet 1990; Schriemer *et al.* 1996]. In larger slabs of a medium comparable to that used my experiments, measurements of the longitudinal phase and group velocities, as well as the scattering strength (kl_s) (wave vector times mean free path), were obtained from the ballistic (unscattered and/or forward scattered) portion of an ultrasonic pulse transmitted through the medium (shown in Figure 1.1 on the following page).
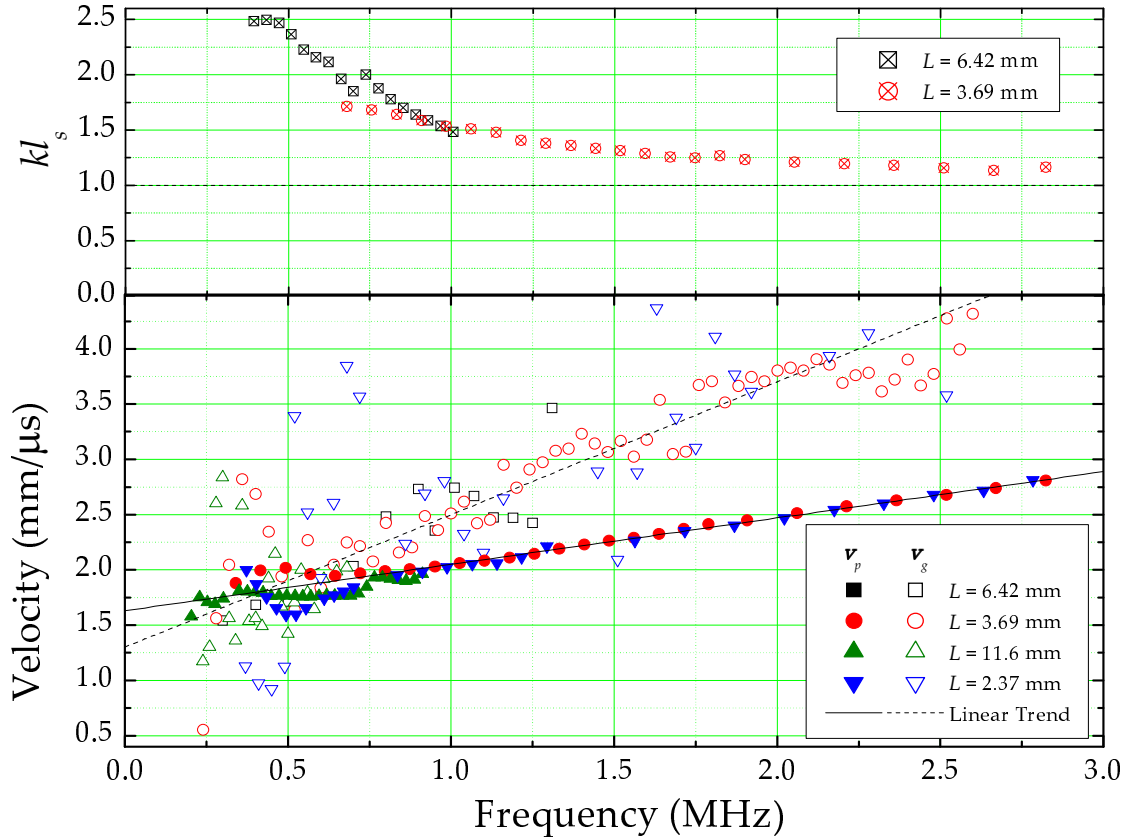
Figure 1.1: Ballistic Wave Transport Properties – The scattering strength (kl_s) and ballistic longitudinal phase (v_p) and group (v_g) velocities are shown for samples of various thicknesses (L). The samples were constructed by Russell Holmes and measurements were made by James Beck (also for Figure 1.2 on the following page).

Additionally, time-of-flight profiles for the multiply-scattered pulse transmitted through this system were used to obtain measurements of the diffusion coefficient (Figure 1.2 on the following page). These previous experiments serve as an excellent basis for my own, and provide valuable information about wave transport in strongly scattering media; however, the density of states is an important property that has remained unknown for these systems, and hence is the subject of my investigation.

Figure 1.2: Diffusion Coefficient Measurements – The diffusion coefficient (D_B), measured from the time-of-flight profiles (inset example) for two similarly prepared samples (of thickness L) are shown. A plateau is found in the strong-scattering regime. The low-frequency crossover to this regime corresponds to the average radius of the largest pores in the medium, and is estimated from results in this thesis (c.f. § 2.3.2 and § 4.3).

In my work, the vibrational modes of the samples are excited by an ultrasonic pulse, and the resulting multiply scattered signal, transmitted through the sample is recorded. If the samples are small enough (since the density of states is proportional to the volume of the sample), peaks in the Fourier transform, corresponding to the individual modes of vibration, will be sufficiently well-separated to be resolved and counted. Thus, the density of states can be directly measured and averaged over many samples to determine the density of vibrational states for elastic waves in this medium.

7

Chapter 2

Theory

2.1 Overview

In my experiments, the density of vibrational states of a porous "mesoglass" was investigated. It is important to understand the qualitatively different mechanisms of wave transport that occur on different length scales within the medium, governed by the relationship between the wavelength (λ) and the size of the pores and constituent particles. These length scales correspond to different frequency regimes, and the behaviour of the density of states in the frequency regime where the scattering is very strong – when the wavelength is comparable to the size of the inhomogeneities – is of particular interest. In this regime, the effect of the disorder is most profound, and anomalous behaviour of the density of states is expected.

In this chapter, predictions are made for the behaviour of the wave transport and the density of states at the various length scales/frequency regimes in the medium. Some existing models for the density of states in a strongly scattering medium are discussed, and the nature of vibrational modes in such systems is examined.

2.2 Debye Model for the Density of States

A simple model for the density of states in three dimensions is obtained by modelling the sample as a homogenous, isotropic closed system, giving rise to standing waves inside the sample. Thus, only waves with wave vector k such that there is an integer number of half-wavelengths contained in the sample are allowed. This gives rise to a finite number of allowed k values in the range $k \to k + dk$, or a density of states in k-space, given by

$$\mathcal{D}_D(k)dk = \frac{Vk^2 dk}{2\pi^2},\tag{2.1}$$

where $\mathcal{D}_D$ is the Debye density of states and V is the volume of the system. From the definitions of phase and group velocity (v_p and v_g), and accounting for all possible polarizations (the subscripts l and t refer to longitudinal and transverse polarizations), the density of states can be written as a function of frequency:

$$\mathcal{D}_D(f)df = 4\pi V f^2 \left(\frac{1}{v_{p,l}^2 v_{g,l}} + \frac{2}{v_{p,t}^2 v_{g,t}} \right) df.\tag{2.2}$$

If velocity is independent of frequency, the density of states grows as f^2. To obtain the correct total number of modes, an upper cut-off frequency must be chosen, such that the total integrated density of states is three times the number of particles in the system (c.f. § 2.5). This cut-off is known as the Debye frequency (f_D), and approximately corresponds, for example, to the smallest allowable wavelength as set by the lattice spacing in a crystal. The Debye model describes the low-frequency density of states fairly well in homogenous and crystalline materials, but breaks down when the scattering is strong and ballistic wave transport no longer occurs.

2.3   Length Scales in a "Mesoglass"

At large length scales, corresponding low frequencies, the ultrasonic wavelength is much greater than the size of the inhomogeneities, and the sample behaves as a uniform medium. In this "effective medium" regime, the wave transport will be determined by the average properties of the sintered network, and will be primarily ballistic (with lower, "effective medium" velocities). At intermediate length scales, the wavelength is comparable to the pore sizes in the medium, the scattering is becomes very strong. In this "strong-scattering" frequency regime, the energy propagates diffusively through the system. At short length scales, corresponding to the high frequency, "bulk material" regime, the wavelength is much smaller than the constituent particles of the system. In this regime, ballistic wave transport occurs within the size of the beads, according to the velocities of the bulk glass.

The medium being investigated consists of a sintered disordered network of glass beads (c.f. Chapter 3 and Chapter 4), which can be thought of as a mesoscopic model of an atomic glass, or mesoglass. Both the size of the beads from which the medium is constructed and the size of the pores in the medium play a role in defining the important length scales in the medium. Note that the crossovers between the various length scales are not abrupt, and the behaviour of the wave transport is expected to vary gradually from one length scale to the next.

2.3.1 Bulk Material Regime

When the wavelength is much smaller than the beads from which the sample is constructed, the wave transport is expected to behave as it would in the bulk material of the beads. The phase and group velocities, as well as the density of states should correspond to that of the bead material in bulk. Because the size of the scatterers (pores) is much larger than the wavelength, a simple ray-tracing picture should accurately describe the wave transport. The bulk crossover length (Ξ) is defined here as the transition point to this regime from the larger wavelength, strong scattering regime, and is estimated to occur when the wavelength is just small enough for standing waves *within one bead* ($\lambda/2 \approx 2a$):

$$\Xi \approx 4a, \tag{2.3}$$

where a is the average bead radius. The high-frequency crossover corresponding to this length scale ($\lambda \leq \Xi$) is defined as $f_\Xi = v_{p,bulk}/\Xi$. However, since the bulk medium is a solid, transverse and longitudinal polarizations should both be considered, and may have different velocities. The bulk crossover frequency (f_Ξ) is then calculated from the energy density-weighted average phase velocity $\langle v_p \rangle_E$. Since the energy density depends on the density of states (c.f. Equation 2.2), the equipartition of energy for one longitudinal and two transverse modes gives:

$$\langle v_p \rangle_E = \left(\frac{1}{v_{p,l}v_{g,l}} + \frac{2}{v_{p,t}v_{g,t}} \right) \Bigg/ \left(\frac{1}{v_{p,l}^2 v_{g,l}} + \frac{2}{v_{p,l}^2 v_{g,l}} \right), \tag{2.4}$$

$$f_\Xi = \frac{\langle v_{p,bulk} \rangle_E}{\Xi}. \tag{2.5}$$

12

2.3.2 Effective Medium Regime

When the wavelength is much larger than the size of the inhomogeneities, the elastic wave behaves as if it were in a uniform medium, with properties influenced by the average porosity and connectivity of the network. In this regime, one expects to observe Rayleigh scattering, which is proportional to λ^{-4}, so the scattering will be weak and the transported wave will be dominated by the unscattered or forward-scattered component. The wave transport will be primarily ballistic, and will occur with the effective medium phase and group velocities. The transition between this regime and the strong scattering, intermediate frequency regime is given by the *percolation correlation length* (ξ), which is equal to the average radius of the largest pores in the medium (c.f. § 2.7). The corresponding low-frequency (effective medium) crossover is given by $k\xi = 1$ [Aharony *et al.* 1987; Schriemer *et al.* 1996], which yields (analogous to Equations 2.4 and 2.5):

$$f_\xi = \frac{\left\langle V_{p,eff} \right\rangle_E}{2\pi\xi}. \tag{2.6}$$

2.3.3 Strong Scattering Regime

At intermediate length scales ($\xi > \lambda > \Xi$), the wavelength is comparable to the sizes of the pores, and the scattering is expected to be very strong. The wave transport is expected to occur diffusively (c.f. § 2.4), and the waves may even be localized (c.f. § 2.7.2 and footnote page 24). The material may also exhibit fractal characteristics (c.f. § 2.7). Because of the complex nature of wave transport in this regime, an investigation of the density of states in the corresponding frequency range ($\mathcal{D}\left(f_\xi < f < f_\Xi\right)$) is of considerable interest. The remainder of this chapter deals with the details of wave transport in this strong scattering regime, the fractal nature of the mesoglass (and its effect on wave transport) and the expected behaviour of the density of states.

2.4   Diffusive Wave Propagation

In the simplest case (which serves to illustrate the essential physics of the system), the propagation of classical scalar waves through any linear, isotropic and uniform-density medium is described by the wave equation:

$$\nabla^2 \psi(r,t) = \frac{\frac{\partial^2}{\partial t^2} \psi(r,t)}{v_p(r)},$$

$$(2.7)$$

where ψ is the wave function (representing pressure in the case of acoustic waves) and $v_p(r)$ is the phase velocity as a function of position (which may vary due to inhomogeneities in the medium). In the absence of absorption, the intensity of the ballistic wave within the medium decays exponentially with the distance travelled x, according to $I(x) = I_0 \exp[-x/l_s]$, where l_s is the scattering mean free path. Given the appropriate boundary conditions, including all details of the internal microstructure, wave propagation in any system can be completely described, in principle, by solving the wave equation. In a complex medium such as an amorphous material, however, it may be exceedingly complex or impossible to solve for the wave function, even if these conditions can be determined. For example, if the scattering becomes very strong, the phase relationship between multiply scattered components of the wave may be intractable.

The *diffusion approximation* assumes that the phase of the wave is sufficiently randomized by scattering that the phase information and interference effects can be ignored in the ensemble average. The phonons can be thought of as undergoing a random walk through the sample, and the propagation of the energy density is

approximated as a diffusive process. The equation governing diffusive energy flow is:

$$\frac{\partial}{\partial t}U(r,t) = D_B \nabla^2 U(r,t), \tag{2.8}$$

where U is the energy density and the parameter D_B is the diffusion coefficient. In three dimensions, the diffusion coefficient is given by:

$$D_B = \frac{1}{3}v_E l^*, \tag{2.9}$$

where v_E is the average local velocity of energy transport, or "energy velocity" and l^* is the average distance travelled by the waves until their direction is randomized, called the transport mean free path. Note that in an isotropic medium, $l_s = l^*$, and $v_g \approx v_E$, in the absence of scattering resonances [Schriemer *et al.* 1997].

While in the effective medium regime the wave will be primarily unscattered or singly scattered (c.f. § 2.3.2), an elastic wave propagating diffusively will be multiply scattered within the sample, effectively taking very long paths through the sample before finally emerging. One expects, then, that the energy contained in a pulse incident on the sample will diffuse out over very long times (compared to the length of the input pulse) if the scattering is strong. Furthermore, in a diffusive system, the energy will be shared equally between all mode polarizations [Weaver 1982; Trégourès and van Tiggelen 2002], allowing for the excitation of all polarizations of vibrational modes, regardless of the polarization of the input pulse.

2.5   Modes and Mode Conservation in Random Structures

The frequencies and line shapes of the normal modes in any system are determined by the size, shape, internal structure and material properties of the system in question. To get an idea of what the modes might look like in a complex system, a simple model was constructed. The model is a one-dimensional system containing three driven coupled harmonic oscillators, with only nearest-neighbour coupling. The properties of each harmonic oscillator (mass, damping, coupling and driving force) were chosen at random. One of the results of these simulations, showing the output amplitude and phase derivative with respect to driving frequency, as measured at each oscillator, is shown in Figure 2.1 below.
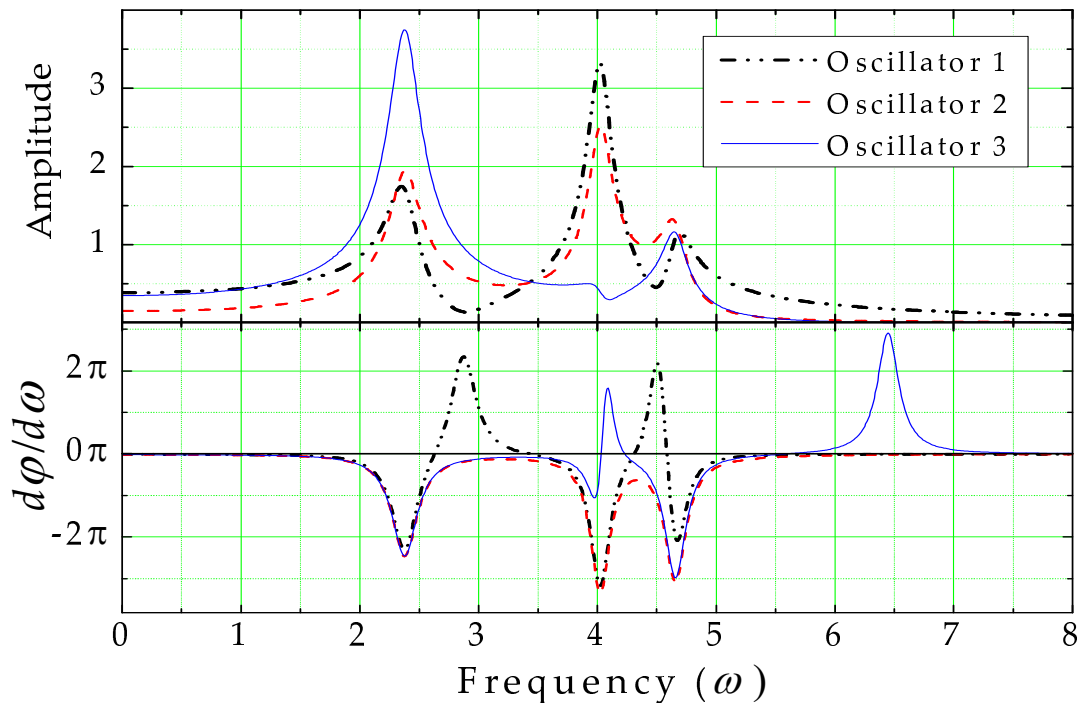
Figure 2.1: Coupled Harmonic Oscillators in 1-D – The response at each of the coupled harmonic oscillators is shown as a function of driving frequency. The units of frequency and amplitude are arbitrary.

The mode shapes and phase response of this coupled system are quite different from the familiar result for a system of uncoupled oscillators, where the peaks are Lorentzian and there is a phase shift of $-\pi$ at each resonance. Clearly, fitting a multi-peaked function (such as three Lorentzians) will not be very successful in locating these resonances (note especially the central peak for the third oscillator in the above figure). Furthermore, though it has been suggested that the density of states can be measured from the cumulative phase, or phase derivative [Sebbah *et al.* 1997], it seems that this would yield strange and inconsistent results for the above simulation (for example, the meaning of the peak far from resonance in $d\varphi/d\omega$ – around $\omega \approx 6.5$ – is not understood). Because of these issues[1], a more empirical method of locating the modes was used (c.f. § 6.5).

If the waves propagate diffusively in some frequency/wavelength regime, the normal modes are expected to appear at random frequencies [Weaver 1989; Mehta 2004][2], and the modal structure of such a system can be characterized by measuring an average number of modes in a given frequency range, i.e. the density of states. Regardless of their distribution, the total number of modes in a given system must be equal to the number of degrees of freedom of the system. In a

---

[1] The measured phase also tends to be somewhat noisy in our experiments, and therefore provides little useful information.

[2] The eigenfrequencies of the system are not strictly random (as they can be in, say, a quantum system), but are determined by the microscopic details of the sample structure, similarly to how the eigenvalues of a Gaussian random matrix are exactly determined by the values of the matrix. A very small change in the sample structure (or matrix values) will cause the eigenmodes to shift unpredictably, and the system is in this sense actually chaotic.

system composed of N_P particles, the total number of modes must be equal to $3N_P$ (3 translational degrees of freedom per free particle, or 3 possible polarizations per normal mode). This condition holds regardless of system size, shape and structure, and is known as *conservation of modes*. Thus, despite the random nature of the modes, the integrated density of states can be known independently from the wave transport behaviour of the medium.

2.6 Volume Scaling

The density of states is an intrinsic and fundamental property of any material (or type of material). In practice, in order to measure the density of states directly, the samples need to be sufficiently small for the individual modes to be resolved. Furthermore, the samples investigated will not all be identical in microscopic detail or size, and the data from several samples must be combined in order to measure the ensemble average properties of the material.

In the Debye approximation, the density of states does indeed scale linearly with volume. Furthermore, the total *integrated* density of states must scale linearly with volume, simply by conservation of modes (c.f. § 2.5), provided, of course, that the samples are statistically similar.

There may also be surface modes that exist in these samples, and the number of these modes should scale with sample surface area, rather than volume. This is, however, not a concern, since the total surface area of the sample should also scale linearly with sample volume, provided that the size of the sample is much larger than the average pore size of the sample ($V >> \xi^3$), since the surface area to volume ratio is set primarily by the pore size and volume fraction of the sample[1]. The *external* surface area is negligible by comparison, and will only affect very-low-frequency modes ($f << f_\xi$).

---

[1] This is analogous to the idea of a unit cell in a crystalline solid. The ratio of surface area to volume is set by that of the unit cell, and is independent of the number of unit cells in the sample.

One expects, therefore, that the density of states should scale linearly with volume, provided the wavelength is much smaller than the sample size ($\lambda << \sqrt[3]{V}$). Fluctuations in the density of states are still expected, even for samples of the same volume, but this simply reflects the random nature of the precise resonant frequencies in strongly scattering, amorphous samples. The measured density of states for each sample should be normalized by sample volume, and the average volume-normalized density of states should then be an intrinsic property of the medium.

2.7   Percolation Theory and Fractal Structures

A simple model that has been used extensively to study the vibrational dynamics of random systems is known as percolation theory. One common type of percolated system consists of a lattice of sites, randomly occupied with probability q[1]. In order to obtain a continuous network (solid), q must be above a certain *percolation threshold* (q_c). For $q > q_c$, these systems are known to exhibit fractal, or self-similar, geometry at short length scales, and to be homogenous at long length scales. The transition between these two regions occurs at the percolation correlation length, ξ, which is equal to the average radius of the largest pores in the medium [Stauffer 1985][2]. Two important quantities that characterize a fractal structure are the fractal (Hausdorff) dimension D, and the fracton (spectral) dimension $\tilde{d}$. Because the samples are constructed by a site-percolation method on a random lattice (c.f. § 3.1), these quantities are expected to be relevant to the analysis of the experiments.

2.7.1   The Fractal (Hausdorff) Dimension *D*

The Hausdorff dimension of any object can be defined as the exponent by which the mass (m) scales with respect to system size (L) [Hausdorff 1919; Mandelbrot 1982], such that

$$m \propto L^{D}. \tag{2.10}$$

---

[1] q is used for the percolation probability (instead of the usual p) to avoid confusion with the mode-counting probability, for which p is used (c.f. § 6.4).
[2] This is the characteristic length of the percolation system, and all the properties of the system should scale with ξ [Schriemer et al. 1996].

This is intuitively understood for objects with integer dimensionality, such as the mass of a rod scaling linearly with length, or the mass of a cube proportional to L^3. An object may exhibit different Hausdorff dimensions at different length scales, such as in Mandelbrot's famous example of the ball of yarn, wherein the object is a point (zero-dimensional) at very long length scales, a sphere (three-dimensional) when viewed more closely, a string (one-dimensional) when viewed even more closely, and so on [Mandelbrot 1982]. Fractals are characterized by having a non-integer, or fractional, dimensionality (this is, in fact, the origin of the word *fractal*), and as in the example just given, the fractal dimension of an object can, itself, be function of length scale.

2.7.2 The Fracton (Spectral) Dimension $\tilde{d}$

Vibrational excitations on a fractal structure have been referred to as *fractons*, and the corresponding *fracton dimension* is defined by the frequency scaling of the density of states [Alexander and Orbach 1982], where

$$\mathcal{D}(f) \propto f^{(\tilde{d}-1)}. \qquad (2.11)$$

For homogenous, non-fractal structures, the fractal, fracton and Euclidian dimensions are all equal, i.e. $D = \tilde{d} = d$, and thus, for a three-dimensional structure, the standard results of $m \propto L^3$ and $\mathcal{D}(f) \propto f^2$ are recovered.

Several predictions have been made for the value of $\tilde{d}$ in a fractal structure. Based on a scalar elasticity model on a percolated network [Born and Huang 1954], a value of $\tilde{d} = 4/3$ was conjectured by Alexander and Orbach [1982], along with the further

conjecture that all vibrational modes on a fractal will be localized[1]. This is consistent [Aharony *et al.* 1987] with the *Ioffe-Regel condition*, which states that localization will occur when [Ioffe and Regel 1960]

$$kl_s < 1, \qquad\qquad (2.12)$$

where k is the wave vector and l_s is the scattering mean free path.

Conversely, a "bond-bending" model (vector elasticity) for vibrational modes on a percolating network predicts that $\tilde{d} = 0.89$ [Kantor and Webman 1984; Webman and Grest 1985; Pachet 1990]. Furthermore, it has been shown that the Ioffe-Regel criterion does not necessarily imply localization for acoustic or elastic waves [Sheng *et al.* 1994; Allen *et al.* 1999; Page *et al.* 2004]. A direct measurement of the density of states in a fractal structure may provide additional insight into these ideas.

---

[1] This refers to strong (Anderson) localization. Though there has been much recent interest in this topic, the theory of localization is beyond the scope of this thesis. An excellent introduction to the subject and review of many important contributions thereto has been published recently [Lagendijk *et al.* 2009], and further information can be found at http://www.andersonlocalization.com/canonal50.php.

2.8 Level Repulsion and Statistics

If the frequency of a resonant mode in the strong scattering regime is random, the probability of observing a degenerate, or nearly degenerate, mode must be considered. Similar systems to that of this thesis have been observed to exhibit level repulsion, which is a reduced probability of observing a mode nearby in frequency to another mode [Weaver 1989]. It has been shown that this result is also predicted by the application of random matrix theory to eigenfrequencies of chaotic systems [Mehta 2004]. This level repulsion is not, however, a Coulomb-type repulsion, but is statistical in nature. This means that the level repulsion does not depend on the *absolute* frequency difference of adjacent modes, but rather, adjacent modes are repelled *relative* to the average mode spacing. This means that the level spacing statistics of samples of different volumes should be statistically similar when normalized by the average mode spacing (which should be inversely proportional to the density of states and therefore the sample volume). As long as the samples are sufficiently small to resolve adjacent modes, the probability of encountering a (near-) degeneracy is very small, and can be ignored safely.

Level repulsion is expected to break down in the case of Anderson localization [Sade *et al.* 2005]. Though localization is predicted by some of the fractal models presented in the previous section, no evidence of Anderson localization has been observed in similar samples in the frequency range studied [Page *et al.* 2004], and level repulsion should persist in the samples studied here. Further investigation into the level spacing statistics may be of interest.

2.9 Estimating the Density of States in a Mesoglass

The mesoglass constructed for these experiments can be thought of as a practical realization of a percolated system. The preceding arguments of this chapter can be used to develop a simple model to estimate the density of states in, and nearby to, the strong-scattering regime. This model was originally proposed to attempt to explain the anomalous Kapitza resistance for sintered metal heat exchangers [Maliepaard *et al.* 1985].

As discussed, in the low-frequency regime ($f << f_\xi$), the density of states function should behave as in an effective homogenous medium described by the Debye model (§ 2.1). At intermediate frequencies, the medium is fractal on the length scale of the elastic wavelength, and the density of states is expected to depend on the fracton dimension. In a practical percolated system like a mesoglass, the bulk medium (glass) must be considered along with the percolated structure, and the internal vibrational modes of the glass beads must also be included (c.f. § 2.3.1).

Above the bulk crossover frequency (f_Ξ), a Debye-like behaviour of the density of states is expected; however, the surface area of the percolated structure is very large, and the contribution of surface waves must be included in the density of states. With this contribution, the density of states in the high-frequency (bulk material) regime is given by [Dupuis *et al.* 1960; Weaver 1989]:

$$\mathcal{D}_\Xi(f)df = 4\pi\phi V f^2\left(\frac{1}{v_l^3}+\frac{2}{v_t^3}\right)df + \frac{\pi S f}{2v_l^2}\left(\frac{2-3(v_l/v_t)^2+3(v_l/v_t)^4}{(v_l/v_t)^2-1}\right), \qquad (2.13)$$

26

where S is the total surface area and ϕ is the volume fraction of the sample (such that ϕV is the *occupied* volume of the sample, because the modes of the bulk medium are being considered), and where the dispersion in the bulk material (glass, in my samples) is negligible ($v_p = v_g$).

The density of states in the high-frequency regime is likely not significantly affected by the structure at long length scales. Because of this, a conservation of modes argument can be used to estimate the density of states in the strong scattering regime. Because the number of total modes must be conserved and the number of modes at *high* frequencies is given by Equation 2.13 (cut off by the appropriate Debye frequency), the total number of modes in the *remaining* frequency range must *also* be given by the integral of $\mathcal{D}_\Xi (f)$ up to the high-frequency cut-off f_Ξ, since these low- and intermediate-frequency modes can only rearrange (i.e. additional modes are not created, c.f. § 2.5). That is,

$$N_{total} = \int_0^{f_D} \mathcal{D}_D(f)\,df = \int_0^{f_D} \mathcal{D}_\Xi(f)\,df = \int_0^{f_\Xi} \mathcal{D}_\Xi(f)\,df + \int_{f_\Xi}^{f_D} \mathcal{D}_\Xi(f)\,df \, , \qquad (2.14)$$

where N_{total} is the total number of modes, given by $3N_P$, where N_P is the number of *atoms* in the sample. Since the total value of the lower-frequency integral (second term from the right) above should be independent of the structure of the sample, *any* function that describes the density of states must satisfy the condition

$$\int_0^{f_\Xi} \mathcal{D}(f)\,df = \int_0^{f_\Xi} \mathcal{D}_\Xi(f)\,df \, . \qquad (2.15)$$

In the low-frequency regime, the density of states should be well-described by the Debye model of Equation 2.2, and this can be integrated up to f_ξ to obtain the total number of low-frequency modes. The remaining density of states in the intermediate frequency regime can then be estimated by subtracting the number of effective medium modes from the total number of modes up to the high-frequency cut-off (from Equation 2.15). If it is assumed, for simplicity, that the number of modes is constant, the following estimate is obtained:

$$\mathcal{D}_I \cong \frac{\int_0^{f_\Xi} \mathcal{D}_\Xi(f)\,df - \int_0^{f_\xi} \mathcal{D}_D(f)\,df}{f_\Xi - f_\xi}, \qquad (2.16)$$

where the I subscript refers to the intermediate frequency range. Note that assuming a constant density of states in the strong-scattering regime would also suggest a spectral dimension of $\tilde{d}=1$ (according to Equation 2.11), which is comparable to the results obtained for the scalar elasticity $(\tilde{d}=4/3)$ and bond-bending $(\tilde{d}=0.89)$ models, so this assumption may be reasonable (c.f. § 2.7.2).

A similar estimate could be made based on a mode-conservation argument wherein the constituent particles (glass beads) are treated as perfectly rigid spheres, and only the degrees of freedom of these particles are considered. Again, the idea behind this is that the high-frequency modes that occur within each bead will be unaffected by the way in which the beads are put together. In this case, the total number of modes of the system *below the high-frequency cut-off* (f_Ξ) is simply given by three times the number of beads from which the sample is made (N_B), and Equation 2.16 becomes:

$$\mathcal{D}_I \cong \frac{3N_B - \int_0^{f_\xi} \mathcal{D}_D(f)\,df}{f_\Xi - f_\xi}.$$ (2.17)

The number of beads in a given sample can be estimated easily from the average bead size and volume fraction in the mesoglass.

Bear in mind that the above expressions are intended to give rough estimates for the density of states in the strong-scattering regime, and the parameters required to calculate numerical values from them may also be somewhat ill-defined. Nevertheless, they should allow for a reasonable estimate of the average density of states in the range of interest to be made, based on the properties of the medium.

Chapter 3

Sample Preparation

3.1 Overview

In order to study the vibrational modes of strongly scattering materials, a medium is required that will strongly scatter elastic waves in the frequency range accessible to via piezoelectric ultrasonic transducers. Such a medium can be constructed by sintering a mixture of glass and iron beads, and subsequently removing the iron through etching. Thus, a highly porous, amorphous mesostructure, or "mesoglass," is created, wherein the length scale of the inhomogeneities is comparable to the ultrasonic wavelength in the medium. Furthermore, the sample is constructed in an analogous way to site percolation on a random lattice, with percolation probability $q = 0.5$, suggesting that the material may be fractal over a range of length scales (c.f. § 2.7). Since the samples must be sufficiently small to allow the individual resonant modes to be resolved and counted, the material is cut into small pieces. These small, amorphous samples allow the density of states to be measured directly in the strong scattering regime.

In this chapter, the details of the sintering, etching and cutting processes, and how the samples were thereby constructed are presented.

Before proceeding further, I wish to express my thanks to Russell Holmes, who had created the large slab sample[1], from which my samples were cut, before my research for this thesis was begun.

---

[1] For the experiments described in this thesis, the slab labelled "Sample 9" was used.

3.2   Material Preparation

The medium was constructed using glass beads with diameters of 106 – 149 μm (mean radius $a \approx$ 64 μm). The glass beads used were *2227 Spheriglass Solid Glass Spheres*, obtained from Potters Industries Inc., sieved to the above range of sizes. The beads are made of *A-Glass*, which is a soda-lime silica glass with a density of 2.5 g/cm^3. The longitudinal and transverse velocities of elastic waves in this glass are $v_l = 5.6$ mm/μs and $v_t = 3.4$ mm/μs [Page *et al.* 1996][1]. The glass beads were mixed with iron beads of the same size range in a volume ratio of 1:1. Considering that the packing fraction of random close-packed monodisperse spheres is approximately $\phi \approx 0.64$ [Torquato *et al.* 2000], this medium is expected to have a volume fraction of about $\phi \approx 0.32$ after the transitory iron beads have been removed through etching.

The glass and iron beads were mixed in the presence of ethanol in order to prevent the particles from sticking together due to the accumulation of static charges. Once the glass and iron beads were thoroughly mixed, the mixture was poured into a mould and firmly packed. Once the mixture had dried, the cover was secured to the mould and the material was then ready to be sintered.

Figure 3.1: Sample Mould – Expanded view one of the moulds used to construct the mesoglass. The thickness of the middle spacer piece determines the thickness of the sample. The dimensions the mould used were 10 cm by 10 cm, with a thickness of 2.38 mm.

[1] The glass beads used in this reference were made of the same material.

3.3 Sintering

In constructing the mesoglass, the aim of the sintering process was to heat the mixture sufficiently to cause the particles to fuse together at their contact points, without entirely melting the material (This is described as *Initial Stage Solid State Sintering* [Kang 2005]. Essentially, the traditional sintering process is being stopped before significant densification occurs). To accomplish this, the material was heated to near its softening point and allowed to "bake" for a fixed amount of time. The material was then removed from the heat source and allowed to cool before the glass had melted entirely (note that the melting point of the iron beads is substantially higher than that of glass, so there is no risk of having the iron melt at these temperatures). There are several concerns that must be considered in the sintering process, as follows:

1. The material must not be heated too quickly. Since the thermal conductivity of the mixture is rather low, slow heating is required for an even temperature distribution within the material.

2. The material must not get too hot, or remain at a high temperature for too long. This could cause the glass to flow too much and thereby destroy the structure of the sample.

3. The material must not be cooled too quickly. This helps to reduce strains in the material induced during sintering, and due to thermal contraction of the glass, iron and surrounding mould during cooling. Ideally, the sample material should be annealed for several hours to relieve any residual strain.

4. The sintering should be done in a nitrogen atmosphere. This will prevent the iron beads from rusting, which could induce larger strains in the sintered glass structure as well as making the iron particles difficult to remove later by etching.

In order to properly sinter the material with the above concerns in mind, the apparatus shown in Figure 3.2 was used, following the procedure described below.

Figure 3.2: Diagram of the Sintering Apparatus – The material to be sintered is contained in the mould, which is placed inside the steel enclosure. The temperature of the material is controlled by moving the enclosure into and out of the furnace. To evacuate the enclosure, valve V3 is closed and V1 is opened. Nitrogen is then introduced by closing V1 and opening V2, and allowed to flow through the enclosure throughout sintering by opening V3.

The mould containing the glass/iron mixture was placed in an airtight steel enclosure. The enclosure was evacuated, in order to remove the air. Nitrogen gas was then introduced, and was allowed to flow through this container throughout the sintering process. This enclosure could be drawn into and out of the furnace, which was kept at a constant temperature of about 810°C. By moving the enclosure into and out of the furnace, and by adjusting the flow rate of the nitrogen gas, the

temperature of the material could be controlled. This was necessary because the temperature of the furnace could not be changed quickly enough to properly sinter the material. The temperature of the material throughout the sintering process is shown in Figure 3.3 below.

Figure 3.3: Sintering Temperature History – The temperature of the material throughout the sintering process is shown, as measured by a thermocouple in contact with the mould. The two softening points shown represent an approximate temperature range for sintering the glass. Note that the material may have been heated and cooled too quickly to produce an ideal sinter.

The quality of the sintered material is very sensitive to the temperature and duration of the sintering process. The sintering temperature of the material can range from approximately the dilatometric softening point (log viscosity ≈ 8-9 Pa s)

to the Littleton softening point (log viscosity = 6.6 Pa s), which is from about 660°C to 730°C for the glass used. To remove strains in the material, the material should be held for several hours in the temperature range of the annealing point to the strain point, around 530°C to 560°C for the glass used [Shelby 2005]. A wide range of temperature/duration combinations are possible to achieve sintering within these temperature bounds. However, despite this broad working temperature range, a change in sintering temperature of only a few degrees or a change in duration of only a few minutes can ruin the material by melting the glass sufficiently to start flowing, or by having insufficient bonds between beads to hold the material together. A sintering duration of about one hour was chosen for convenience (sufficient to achieve good temperature uniformity during sintering, without having the process take an extremely long time), and the appropriate sintering temperature was worked out through trial and error.

Several slabs of this type of material were made for use in other experiments, for which the goal was to measure the ultrasonic wave diffusion coefficient [Page *et al.* 2004]. Because only small samples were required for my mode counting experiments, one of the slabs that broke while being removed from the mould was used to make my samples. This breakage may have been caused by heating or cooling the material too quickly. As seen in Figure 3.3, the material was heated somewhat rapidly, and may have also been cooled to below the strain point too quickly. One of the pieces of this broken slab was subsequently etched and used to make my samples.

3.4 Etching and Cutting

Once the material had been properly sintered, the iron was removed by etching. The material was placed in a 3:1 solution (by volume) of water to HCl ($\approx$12 M), and the solution was gently heated and stirred. As the iron was dissolving, the solution became cloudy, and needed to be changed periodically. After several days of etching, the solution remained clear after changing, indicating that the etching process was complete. The material was then thoroughly flushed with water and allowed to dry, and an open network of connected glass beads remained.

Because the glass and iron beads were originally mixed in a 1:1 ratio, both the glass *and* iron form a continuous network in the medium. It is, therefore, very unlikely that any iron will remain in the medium after etching due to being trapped by complete enclosure by glass. This was also supported by the results of x-ray tomography (c.f. Chapter 4), which showed virtually no evidence of residual iron, even though iron beads would be clearly visible, due to their high x-ray scattering contrast.

A piece of the slab sample was then mounted to an aluminum plate using *Apiezon Wax W* mounting wax. The sample and wax were heated to a temperature of about 130°C, at which the wax flows quite well. An abundance of wax was used in hopes that the wax would flow into the pores of the sample and strengthen it during the cutting process. The aluminum plate was then mounted on a diamond wire saw (shown in Figure 3.4 on the following page), which slowly made a cut into the slab.

Figure 3.4: Diamond Wire Saw – This was used to cut the material into small samples. The thin cutting wire can be seen crossing the mounting plate.

The diamond wire saw was used since it is an excellent method of sectioning small samples, minimizing the possibility of inducing additional strains or otherwise damaging the samples.

A grid was cut into a piece of the slab in order to obtain small cubes. This grid was made with irregular spacing between the cuts in order to produce samples with a range of sizes, and some pieces were cut again to obtain even smaller samples. The aluminum plate was then heated in order to melt the wax and remove the cubes. The cubes were placed in a vial of toluene (a good solvent for the mounting wax used), and this was placed in an ultrasonic bath. The toluene was continually changed until it would remain clear after being placed with the samples in the bath, indicating that the wax was removed completely. The samples were then placed under a heat lamp and allowed to dry thoroughly.

The samples ranged in size (L) from around 0.8 – 1.8 mm. Each sample was given a two-letter name by which it could be identified, and the samples used (along with their volumes) are listed in Table A.1 in Appendix A.

3.5 Summary – Sample Properties

Photographs of one of the larger samples are shown in Figure 3.5 below. The close-up photographs were taken using a digital camera connected to the eyepiece of a viewing microscope. The sample appears orange because of the red background on which it was photographed.

Figure 3.5: Sintered Glass Bead Sample – A close-up photograph of the material surface (top) and photographs of some cut samples (bottom) are shown. The porous structure of the material is clearly visible. The cut samples are shown beside recognizable objects for scale.

Because of the difficulties in lighting and optical contrast when photographing glass, x-ray tomography was used to get a better measurement of the mesostructure of the samples. Three-dimensional reconstructions of the samples could then be created, yielding images such as Figure 3.6 and Figure 3.7 below and on the following page. More reconstructed images, as well as the details of the x-ray tomography and the quantitative results obtained thereby, are shown and discussed in Chapter 4: Sample Analysis by X-Ray Tomography.

Figure 3.6: X-ray Tomography 3-D Reconstruction of a Sample – The structure of one of the samples (BR) can be clearly seen. The sample appears to be somewhat over-sintered, as none of the original glass bead structure remains visible. Many impressions left by the iron beads can be seen.

From the x-ray tomography images, it is clear that the samples are somewhat over-sintered. Rather than a structure of interconnected glass beads fused together at their interfaces, a more continuous glass structure exists, and it appears as if the glass may have flowed too much during the sintering process. These samples are therefore less like an ideal random percolated structure; however, a very porous open glass network remains, which is highly inhomogeneous at the length scales of interest. A large piece of the material was left uncut (shown in Figure 3.7 below) so that the average properties of the medium could be investigated.

≈2 mm

Figure 3.7: X-ray Tomography 3-D Reconstruction of Medium (Slab) –
The general porous structure of the sintered glass bead network can be seen.

From this piece, the density of the mesoglass was found to be $\rho = 0.83 \pm 0.02$ g/cm³. This means the samples have an occupied volume fraction of $\phi = 0.332 \pm 0.007$, which is close to the predicted value of $\phi = 0.32$ (c.f. § 3.2). The volume fraction may be slightly higher than predicted due to the polydispersity of the beads (since polydisperse spheres can pack more tightly than monodisperse spheres) and/or over-melting (densification) during sintering.

The samples were then weighed individually, and their volumes were calculated from their masses. The samples used were found to have volumes ranging from about 0.6 mm³ up to about 5.8 mm³ (c.f. Table A.1).

A summary of the sample properties is shown in Table 3.1 below.

Bulk Material	Soda-Lime Silica Glass
Bulk Material Density	2.5 g/cm³
Bulk Material Longitudinal Velocity ($v_{l,bulk}$)	5.6 mm/µs
Bulk Material Transverse Velocity ($v_{t,bulk}$)	3.4 mm/
Bead Size ($2a$)	106 – 149 µm
Mixing Ratio (Percolation Probability)	1:1 ($q = 0.5$)
Volume Fraction (ϕ)	0.332 ± 0.007
Sample Density (ρ)	0.83 ± 0.02 g/cm³
Sample Volume	0.6 – 5.8 mm³
Effective Medium Phase Velocity[1] ($v_{p,l,eff}$)	1.63 + 0.42 f
Effective Medium Group Velocity[1] ($v_{g,l,eff}$)	1.3 + 1.2 f
Effective Medium Longitudinal/Transverse Velocity Ratio[2] (v_l / v_t)	$\sqrt{2}$

Table 3.1: Summary of Sample Properties

[1] C.f. Figure 1.1.

[2] Estimated from experiments in a similar system [Schriemer *et al.* 1996].

Chapter 4

Sample Analysis by X-Ray Tomography

4.1 Overview

As mentioned in the previous chapter, the mesoglass was also investigated using x-ray tomography. This allowed for a more accurate picture of the structure of the samples, and provided an independent measurement of some of the characteristic parameters of the medium.

Most notably, the crossovers between important length scales in the material were measured; from the sub-bead-size bulk glass scale, to the intermediate fractal-like scale, to the bulk "effective medium" scale. Furthermore, the fractal dimension of the medium (c.f. § 2.7.1) was found.

These x-ray tomography experiments are considered to be complimentary to the primary elastic mode counting experiments discussed in this thesis; therefore, this chapter is intended to be somewhat self contained.

4.2   Tomography Method and Image Acquisition

A diagram showing the basic set-up used in x-ray tomography is shown in Figure 4.1 below.
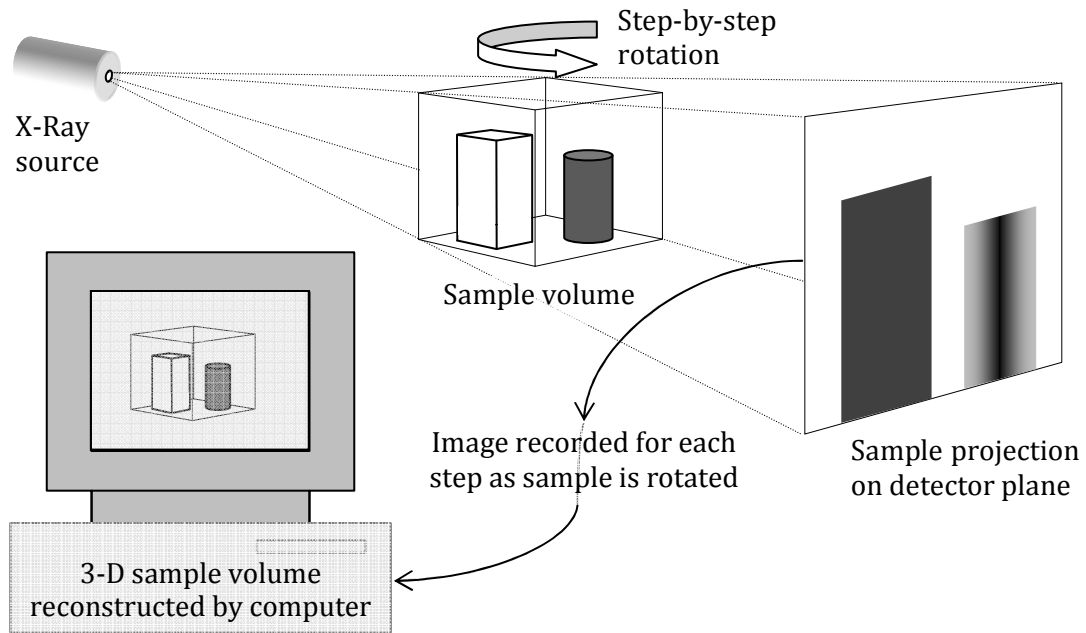
Figure 4.1: Basic X-Ray Tomography Set-Up – The x-ray transmission through the sample is recorded for many angles, and the 3-D sample volume is reconstructed by a computer.

A wide beam of x-rays is transmitted through the sample to a detector plane, creating a two-dimensional projection image of the three-dimensional sample. The intensity (brightness) at a given point on this projection is inversely proportional to the integrated x-ray attenuation (primarily due to scattering) by the sample from the source to this point. Thus, a projection of the entire internal structure of the sample in a particular orientation, based on x-ray attenuation strength, is obtained.

By rotating the sample and recording the projections at many angles, the three-dimensional structure of the sample can be computed and observed.

An example of three projections (in different orientations) of one of the samples (DY) is shown in Figure 4.2 below.

Figure 4.2: X-Ray Projections of Sample – These images show the x-ray transmission through the sample volume in different orientations. The glass shows good x-ray scattering contrast. From these images (at many angles) the entire 3-D sample volume can be reconstructed.

The x-ray tomography was done using an *Xradia MicroXCT* 3-D x-ray transmission microscope. Images were acquired over a 180° range of angles, in 1° steps. The x-ray source voltage was set to 40 kV (x-ray energy < 40 keV), and the x-ray intensity, image dwell time and magnification were all adjusted to suit the sample being measured.

The reconstructed sample volume can be viewed in several ways. The Xradia software allows one to view two-dimensional "slices" of the sample, and can also produce various types of three-dimensional renderings of the sample, which can be rotated and truncated in order to observe the sample volume from any side, or even

"see inside" the sample volume. Examples of a two-dimensional slice and an SSD[1]-rendered image are shown in Figure 4.3 below.

200 μm

Figure 4.3: Reconstructed Images from X-Ray Tomography – A 2-D slice (left) and 3-D SSD rendering (right) of the sample (DY) are shown. The scale bar at the bottom left applies to both images. These reconstructions are obtained from the same image set shown in Figure 4.2.

The reconstruction was performed by the Xradia software, but the details of the reconstruction algorithm are not available (for information on common image reconstruction techniques used in computed tomography, see [Herman 1980]).

---

[1] Surface Shaded Display: This type of rendering creates the appearance of a solid surface (with an apparent light source). The surface is defined by a volume of voxels with similar reconstructed intensity, and therefore creates a good representation of a structure with a uniform, high-contrast x-ray scattering strength. VRT, or Volume Rendering Technique was used for Figure 3.7, and allows for better visualization when the image is somewhat noisy.

The Xradia reconstruction software is also able to remove some artifacts commonly present in computed tomography, such as ring and beam-hardening artifacts. Ring artifacts occur when the gain of one or more detector pixels drifts during image acquisition, and can be identified by the appearance of rings in two-dimensional slice images, perpendicular to the axis of sample rotation. Beam hardening occurs due to the polychromaticity of the x-ray beam. Since the x-ray attenuation of the material is energy dependant, the attenuation of a particular location in the sample may vary with direction, since it may have already passed through more or less sample before reaching that point. This gives rise to a non-linearity in x-ray attenuation with respect to sample thickness [Herman 1980]. Beam hardening artifacts can appear as streaks in the two-dimensional reconstructions, and can also worsen contrast and produce noise.

With care taken to properly acquire and reconstruct x-ray tomographic images, high quality two- and three-dimensional images of the samples were obtained, as seen in this and the previous chapter. These images were then used to make both qualitative and quantitative observations about the structure of the samples.

4.3 Fractal Dimension and Box Counting

The two-dimensional images obtained by x-ray tomography can be used to identify and characterize the three important length scales in the medium, as discussed in Section 2.3. By using a "box-counting" technique [Liebovitch and Toth 1989, Huang *et al.* 1994], the fractal dimension (c.f. § 2.7) of the images can be measured. In this technique, an image of size L is first segmented into "occupied" and "unoccupied" regions. The image is then divided into squares of length ε, and the number of boxes containing at least one occupied pixel ($N(\varepsilon)$) is counted as in Figure 4.4 below.

Figure 4.4: Fractal Box Counting Example – The image of size L is divided into occupied (green, dashed) and unoccupied (red, solid) boxes of size ε. This image has not been segmented, and is shown for illustration only.

The power law by which the number of occupied boxes decreases as the box size increases gives the fractal dimension of the image[1]. Since the image is a two-dimensional cross section of the material in three-dimensional space, the fractal dimension of the material (D) is the fractal dimension measured from the image plus one [Schriemer *et al.* 1996], that is,

$$N(\varepsilon) \propto \left(\frac{L}{\varepsilon}\right)^{D-1},$$ (4.1)

$$D = \frac{d \log N(\varepsilon)}{d \log(1/\varepsilon)} + 1.$$ (4.2)

The fractal dimension can therefore be easily inferred by plotting $N(\varepsilon)$ vs. ε on a log-log plot and measuring the slope in the region of interest, where D = (1 – slope).

Since segmented images are required for a box-counting analysis, care must be taken to obtain high-quality tomographic reconstructions, as free from noise or other artifacts as possible.

---

[1] For example, in a bulk material (D = 3), the entire 2-D cross-section is be occupied. The total number of boxes covering the image decreases as ε^{-2}, and since every box is occupied, $N(\varepsilon)$ follows the same power law. D = 2 is measured for the image, indicating D = 3 for the bulk sample.

4.4 Results

Qualitative observations based on these tomographic reconstructions have already been discussed in Section 3.5. For the fractal box counting analysis, the images were processed and analysed using ImageJ – an open source image processing program (freely available at http://rsbweb.nih.gov/ij/index.html). ImageJ has a built-in fractal box counting function, and macros can easily be created in order to process large batches of images. The results of the fractal box counting are shown in Figure 4.5 below.

Figure 4.5: Fractal Box Counting Results – Results are shown for three samples and three different levels of magnification. Many more samples and tomographic reconstructions were also analyzed and found to be consistent with the above results, and omitted only for clarity of display. The power-law fits and crossover points are estimated by eye.

From the fractal box counting analysis, the important length scales in the medium can be identified, and the fractal dimension can be measured in the intermediate regime. The gradual nature of the crossovers makes it difficult to precisely define a fitting region to measure the power-law dependences in Figure 4.5 and because of the effects of finite pixel and image size, it was difficult to assign a quantitative uncertainty to the data. For these reasons, the fits were done by eye, and the values measured are approximate. Nevertheless, a good understanding of the nature of the samples can be obtained from these data. The power-law fits with reasonable uncertainty estimates yield $D = 2.60 \pm 0.05$, $\Xi_F = 11 \pm 4$ µm and $\xi = 240 \pm 80$ µm. (Here the subscript F on the bulk crossover length Ξ_F indicates this crossover as measured by fractal box counting. C.f. Figure 4.6 and related discussion.) Previous measurements on similarly constructed samples show the percolation correlation length to be about 7 bead radii, or $\xi \approx 7a$ [Schriemer et al. 1996]. In my samples, a value of $\xi \approx 4a$ is found, which may be another indication of over-sintering (c.f. § 3.5).

As discussed in Section 2.7, the percolation correlation length (ξ) should represent the average radius of the largest pores in the medium. The exact physical meaning of the measured bulk crossover length (Ξ_F) is somewhat less clear. Perhaps it can simply be said that below the measured Ξ_F, the material behaves as the bulk glass and that beyond this length, the fractal nature of the material begins to become important. This result is somewhat surprising, as one might expect the bulk crossover length to be related to the constituent bead radius as in Equation 2.3, or $\Xi \approx 4a \approx 250$ µm. To get a better idea of how these length scales relate to the sample

52

features, boxes corresponding to the crossover lengths are shown with a two-dimensional reconstruction of one of the samples in Figure 4.6 below.

Figure 4.6: **Illustration of Crossover Lengths** – Boxes corresponding to the crossover lengths are shown with one of the 2-D reconstructions, to compare with sample features. (Note that the grey area is occupied by glass, while the black area is unoccupied.)

The bulk crossover length measured via fractal box counting appears to correspond with something akin to the average size of the smallest feature of the medium. This length will not correspond to the bulk crossover frequency as per Equation 2.5, as the elastic properties of the bulk medium must clearly begin to dominate on a length scale comparable to the bead size, and this should be considered in any subsequent analysis.

When segmenting the greyscale images obtained, the correct threshold must be chosen to separate the occupied and unoccupied pixels. This was done as objectively as possible, by simply choosing the minimum in the image histogram as the image threshold. An example of one of the greyscale images along with its histogram is shown in Figure 4.7 below.

Figure 4.7: Reconstructed Image with Histogram – This image is a good candidate for segmentation, as shown by the two-peaked histogram of the greyscale values (x-axis on histogram, where the greyscale value is shown in the bar at the bottom). The threshold of such images was chosen to be at the approximate minimum between the two peaks. The resultant occupation fraction is about 35%, reasonably consistent with the measured volume fraction of $\phi = 0.332$.

This quantitative method of image segmentation yields good agreement between the occupation percentage and the previously measured volume fraction, suggesting that the image is indeed being segmented correctly.

Chapter 5

Mode Counting Experiments

5.1 Overview

In the experiments presented in this thesis, the density of states of sintered glass bead samples was measured using a direct mode-counting technique, which is described in this chapter. In order to identify the modes of the samples, the resonances of the sample had to be excited and detected. This was done by placing the sample in direct contact with two ultrasonic transducers, where only one or at most a few points on the sample surface were in contact with each transducer. A short ultrasonic pulse was produced by one of the transducers to excite the vibrational modes of the sample. A second transducer in contact with the sample was used to record the transmitted elastic wave, and the Fourier transform of the recorded signal was used to identify and count the modes.

5.2   Transducer Holders

The experimental set-up is shown in Figure 5.1 below. The transducer holders are an adaptation of an apparatus previously used to investigate ultrasonic velocity and attenuation in similar samples [Pachet 1990].

Figure 5.1: Sample Stage – The sample is in contact with a generating and receiving transducer at a few points. The transducer spacing is controlled by springs and micrometers in order to make light contact with the sample. The dashed red box on the top left diagram is the region shown in the bottom photograph, where the diameter of the transducers is about 7 mm.

The micrometers were used to adjust the spacing of the two transducers, such that the transducers made light contact with the sample (the effect of contact pressure on the acquired spectra will be discussed in § 6.3.2). The micrometer "feet" sat on glass plates affixed to the bottom holder, to isolate the two transducers electrically. Low-frequency ultrasonic waves can be transmitted through the holder itself and picked up by the receiving transducer. Because of this, for experiments in the lowest frequency range investigated, the glass was replaced with soft plastic, in order to damp these vibrations. A comparison of the low-frequency signals transmitted through the holders under these conditions is shown in Figure 5.2 on the following page.

The plastic adequately damps the low-frequency vibrations. Note that some care must be taken to choose the correct damping material, as too soft a material may cause some instability between the top and bottom transducers, which can present a hazard to the sample. Additionally, closed-cell materials are also not useful, as they tend to expand when placed under vacuum, changing the transducer separation. Polypropylene was found to be a suitable plastic for this application.

Figure 5.2: Sample Holder Resonance – The signal (top) from the generating transducer is transferred to the receiving transducer through the sample holder, and the spectrum (bottom) shows several strong low-frequency modes. Replacing the glass "foot pads" with plastic pads significantly reduces the transmitted signal. The signals shown were acquired under vacuum with no sample present. The initial RF pulse (c.f. § 5.3.3.4: RF Pick-Up) is not shown. Both of the spectra on the bottom have been normalized by the maximum peak intensity for the signal acquired with the glass plates in place.

5.3   Signal Path and Electronics

In order to measure the density of states over a broad range of frequencies, several pairs of generating/receiving transducers were used, as each transducer has a limited bandwidth about a specified central frequency. The generating and receiving transducers were always chosen to have a common central frequency, in order to maximize efficiency. A short, tuned pulse was sent through the sample, exciting a broad range of frequencies. The transmitted signal was picked up by the receiving transducer and sent to a digitizing oscilloscope. To minimize the effects of digitizing (bit) noise and electronic noise, the pulse was repeated and averaged. A basic representation of the signal path is shown in Figure 5.3 below.

Figure 5.3: Block Diagram of Basic Signal Path – The signal path is shown with arrows. The two transducers were actually in contact with the sample; the diagram is meant to show signal propagation via elastic waves.

The following sub-sections deal with each step along the signal path in the order of propagation. Since these experiments were performed over the course of several years, the data acquisition procedure and electronics used were modified as the experiments were improved. Note that since RF signals were being used throughout these experiments, input, output and cable impedances had to be matched. Cable lengths were also minimized to reduce noise pick-up and signal reflection effects.

5.3.1 Signal Generation

The input signal was generated using an *Agilent 33220A Arbitrary Waveform Generator* (AWG). To generate a large bandwidth pulse in the frequency range of the transducer pair being used, the output from the AWG was chosen to be a single sine wave oscillation, at the central frequency (f_c) of the transducers. A pulse repetition time of 10 ms was used in these experiments, in order to ensure that any noise or echoes from the previous pulse had died off before the next pulse was produced. The signal produced by the AWG and the corresponding Fourier transform are shown in Figure 5.4 below.
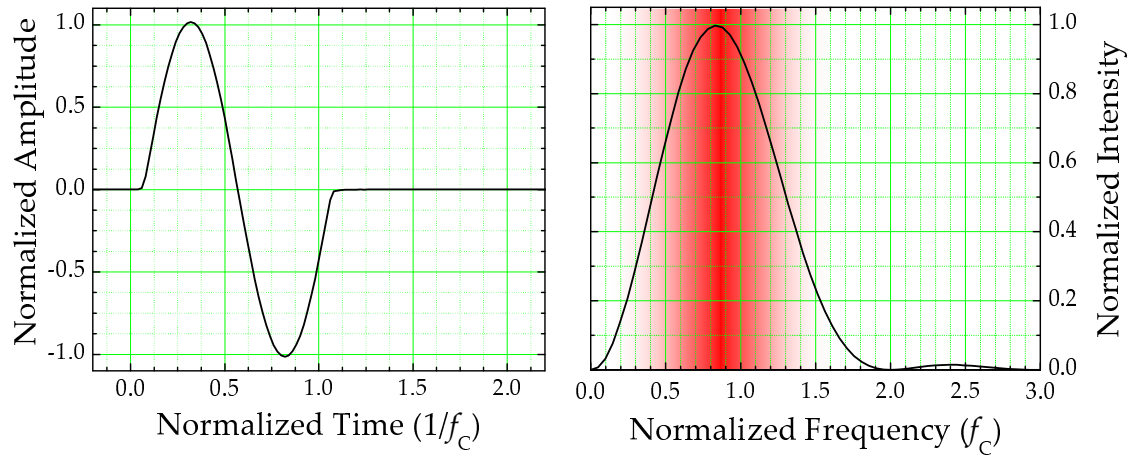
Figure 5.4: Input Pulse – The input signal (left) as produced by the AWG and its Fourier transform (right). The central frequencies (f_c) of the transducers used were 0.5, 1.0, 2.25, 3.5 and 5.0 MHz (c.f. § 5.3.3.1). The AWG output amplitude used was either 50 or 100 mV, depending on the generating transducer used. The region shown in red (shaded) is approximately the frequency range used to obtain mode-counting data in my experiments. Both graphs have been normalized by their respective maxima. The peak frequency in the Fourier transform is shifted below the frequency of the sine wave used to produce the single oscillation, in agreement with the analytic result.

5.3.2 Signal Amplification

The signal was amplified using an *Amplifier Research Model 250L* broadband RF amplifier. The input and output impedances of this amplifier are specified to be 50 Ω, but the input impedance appeared to be matched properly only when the gain was set to maximum (otherwise, cable reflections were observed). The amplifier is quite robust, and is capable of operation without damage using any input or load impedance, and can be used in pulsed or continuous-wave mode. Though a pulse was sought, significant distortion was introduced when using a gating/blanking signal, therefore, the amplifier was used in continuous wave mode, and the pulsing was accomplished by the signal generation electronics, as described in the previous section. The amplifier is operated at its maximum gain setting, which yields a gain of approximately 62 dB. The output is shown in Figure 5.5 below.

Figure 5.5: Amplifier Output – The pulses sent to the generating transducers (left) and their Fourier transforms (right) are shown. The larger pulses were used for the transducer pairs with f_c = 500 kHz and 3.5 MHz). The approximate region of interest is shown in red. The spectra on the right are normalized by their maxima.

The signals on the previous page were acquired by connecting the amplifier to a 50 Ω resistive load, and using an impedance-matching splitter to tap off a small, proportional voltage which could be recorded by an oscilloscope. Note that the voltage applied to the generating transducer is approximately double for the 500 kHz and 3.5 MHz ranges. The experiments performed at these frequencies were done later, after it was found that a higher voltage could be applied to the transducers safely. Since the amplifier must be operated at its maximum gain setting for proper impedance matching, the output voltage applied to the generating transducer was controlled by the AWG output voltage.

5.3.3 Transducers

5.3.3.1 Design and Construction

In these experiments, piezoelectric immersion transducers are used for generation and detection of ultrasound. The transducers are designed for underwater usage, and are therefore acoustically impedance-matched to water. These transducers are chosen for their broadband response, rugged construction and availability. Diagrams (obtained from the manufacturer) of the internal construction of the transducers are shown in Figure 5.6 below.

Figure 5.6: Transducer Construction – Elements of the appropriate thickness and acoustic impedance are selected in order to maximize the efficiency, signal fidelity and bandwidth in a specific frequency range [Panametrics 2006].

The active element in these transducers is made from lead zirconate titanate (PZT), which is a high acoustic impedance piezoelectric material. The PZT has a thickness

of one half-wavelength at the nominal central frequency of the transducer, such that constructive interference produces a strong signal in the desired frequency range. The transducer backing is a high-density, highly attenuative material that is acoustically impedance matched to the PZT element. This prevents "ringing" in the transducer, and gives the transducer a good broadband response. The wear plate on the front of the transducer is designed to provide a better impedance match between the transducer element and water, as well as protecting the piezoelectric element. The wear plate is designed to have a thickness of one quarter-wavelength, in order to minimize losses due to back-reflection. It is worthwhile to note that in the experiments conducted for this thesis, the transducers were not used in water but in direct contact with the samples, and therefore were no longer matched for proper acoustic impedance. This may have affected the performance and frequency response of the transducers.

Each transducer has a nominal central frequency (f_c) which corresponds to the peak frequency response when used in immersion. In this thesis, transducer pairs are referred to by this nominal central frequency in order to indicate the approximate frequency range being studied; however, this does not necessarily reflect the peak frequency response of the transducers when used as in these experiments.

The central frequencies (f_c) of the transducers used were 0.5, 1.0, 2.25, 3.5 and 5.0 MHz, with element sizes ranging from ¼" to 1".

5.3.3.2   Reference Waveforms

Reference waveforms were recorded from each transducer pair, in order to observe the signal used to excite the sample. These references were taken by aligning the transducers with a small air gap in between. The references were taken through air (rather than through water or a solid material) because the very low acoustical impedance of air was thought to provide a load on the transducers similar to the experimental conditions, wherein the load on the transducers was mostly vacuum (zero impedance) with a very small area of the transducer face in contact with the sample. The reference waveforms are shown in Figure 5.7 below.
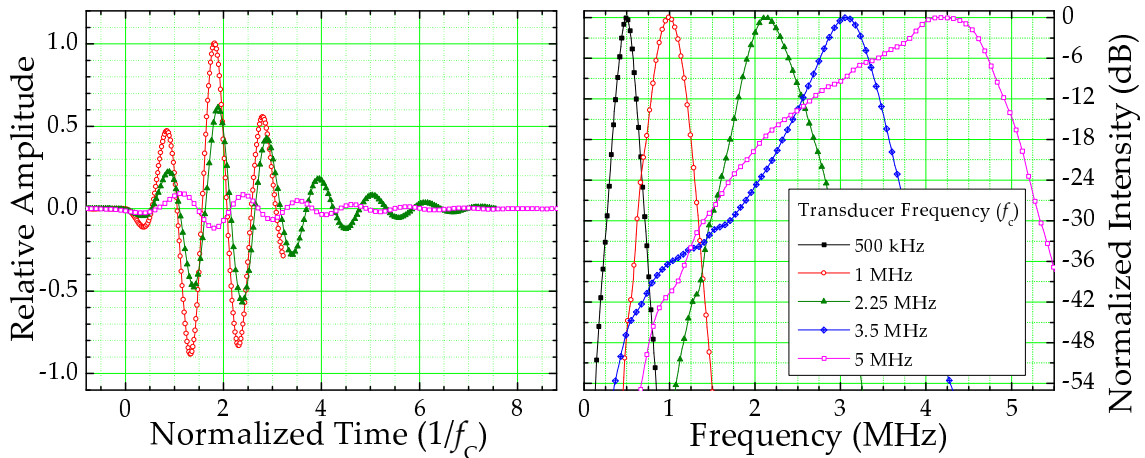
Figure 5.7: Transducer Output – Reference pulses through air (left) and their Fourier transforms (right) are shown. The magnitude decreases with frequency, due to attenuation in air and reduced transducer efficiency. The overlap in bandwidth between transducers allows the density of states to be measured across the entire frequency range. Some signals have been omitted from the left graph for clarity, but were qualitatively similar. The spectra on the right have been normalized by their respective maxima.

The reference signals through air become weaker with increasing frequency. The increased attenuation of air and reduced transducer efficiency at higher frequencies decrease the pulse amplitudes when moving to higher frequency ranges, and shift the central frequency of the pulse to below the transducers' central frequency. Because of this reduced efficiency and the large acoustic impedance mismatch (for air references and when a sample is present), experiments became increasingly difficult at higher frequencies, where signal-to-noise ratios were drastically reduced.

The bandwidth of the transducers in air (given a single-oscillation input pulse – bandwidth ≈ 100%) was approximately 15-30%. The test forms provided by the manufacturer show the transducer bandwidth to be about 65-90%. This suggests that the inefficiency due to improper acoustical impedance matching narrows the transducer bandwidth, in addition to reducing the amplitude of the pulse. Nevertheless, several transducer pairs were used with sufficiently close central frequencies and sufficiently wide bandwidths to allow observation of signals through the samples across a broad frequency range without gaps.

5.3.3.3   Electrical Impedance

The electrical impedance of the transducers was also measured using an *Agilent 4294A Precision Impedance Analyzer*, and the results are shown in Figure 5.8 on the following page. Though the nominal impedance of all the transducers is 50 Ω, the measured impedance varied greatly from this value, spanning several orders of magnitude. The amplifier and receiver are designed for 50 Ω loads, so this electrical impedance mismatch may account for some of the pulse distortion. Notice also that

the impedance is generally closer to 50 Ω at lower frequencies, and this may partially explain the enhanced response at low frequencies, as shown in Figure 5.7. Note that there are two different traces for each central frequency, corresponding to the generating and receiving transducers. Due to availability, these transducers were not necessarily matched in size and make/model, and hence the electrical impedance of transducers of the same central frequency is not necessarily the same.

Figure 5.8: Transducer Electrical Impedance – The electrical impedance of the transducers used is shown. Note the large deviations from the nominal value of 50 Ω.

5.3.3.4 RF Pick-Up

In these experiments, the generating and receiving transducers were in very close proximity (≈ 1 mm) in order to be in physical contact with the small samples. Each transducer acts as a weak antenna, and a signal is transmitted electromagnetically to the receiving transducer, which was observable at early times in the recorded waveforms. As is often the case with RF pick-up, this signal was very sensitive to small changes in the configuration. Factors such as transducer position (separation and parallelism), axial rotation and ground connection between generator, receiver and vacuum chamber all affected the shape, magnitude and duration of the RF pulse significantly. These factors were adjusted to minimize the pick-up signal; however, the RF pulse often had significant amplitude in my mode counting experiments, and care was taken to minimize its effect on the acquired spectra (c.f. § 6.2.1). Two sample RF pulses are shown in Figure 5.9 below.

Figure 5.9: RF Pick-Up Signal – Two examples of RF pick-up signals are shown. Note the differences in frequency, duration and amplitude. The RF pick-up became increasingly problematic at higher frequencies.

5.3.4 Receiving Electronics

A receiver was required to amplify the signal from the receiving transducer to a sufficient level to be efficiently captured by a digitizing oscilloscope (minimizing bit noise). Initially, a *Matec Model 605 Receiver* was used in conjunction with a *Matec Model 253 Broadband Preamp* and a *Telonic 8143A Variable Attenuator*. This receiver/preamp combination did not allow for any gain adjustment, so the attenuator was used to adjust the magnitude of the received signal to minimize digitization noise while avoiding clipping of the waveform. In later experiments, a *Ritec Model BR-640A Broadband Receiver* with a *Ritec Model PAS-0.1-20 Broadband Preamp* was used. The Ritec receiver has adjustable gain and a better signal-to-noise ratio than the Matec receiving electronics. The Ritec broadband preamp was eventually abandoned, as the receiver alone was found to provide sufficient gain without additional noise. Approximately 50 dB of gain was required of the receiving electronics.

High- and low-pass filters were also used to reduce noise outside the frequency range of interest. The Ritec receiver has the advantage of switchable high- and low-pass filters, eliminating the need for external filters.

5.3.5 Signal Acquisition

The signals were acquired using a digitizing oscilloscope, so that they could be averaged and downloaded for analysis. Initially, a *Tektronix TDS 544A Digitizing Oscilloscope* was used to capture and average the signals. This oscilloscope has a resolution of 8 bits, and can perform up to 10,000 averages. In later experiments (especially at higher frequencies, where more averaging and higher resolution was required), a *Gage Compuscope CS14200* digital oscilloscope card was used with *Gagescope Professional* software to acquire the signals. The Gage card has 12-bit resolution, and the number of averages that can be done internally on the card is dependent on the record length required.

The averaging of the Gage card is much faster than that of the Tektronix scope, and the Gagescope software enables the automatic saving of acquired signals. Furthermore, the Tektronix stores the averaged waveform in 16 bit, while the Gage can store the averaged waveform in 32-bit (if "Co-Add" mode is used). Despite the relative difficulty of using the Gagescope (especially involving issues with proper wave scaling and data format conversion), the increased speed and precision made it the superior choice for signal acquisition, especially when dealing with very small signals, or when large amounts of averaging were required. The Gagescope also allowed for clock synchronization with the signal generation electronics, eliminating any possible phase jitter.

5.4 Vacuum System

The experiments were performed under vacuum in order to eliminate the signal propagating through air, as well as to reduce the damping within the sample (since acoustic waves cannot propagate in a vacuum, eliminating the air surrounding the sample eliminates any possible coupling of elastic energy out of the sample except through the contact points with the transducers). It was observed that only a moderate vacuum was required to accomplish these goals (≈ 100 mTorr), and a simple rotary vacuum pump was sufficient. The separation of the transducers was adjusted until contact is established with the sample, and the entire apparatus was placed inside the vacuum chamber, shown in Figure 5.10 below.

Figure 5.10: Vacuum Chamber – The transducers and holders can be seen through the chamber window. Electrical pass-throughs and vacuum connections are also visible.

The effect of performing the experiment under vacuum is shown in Figure 5.11 below. Note that with even a moderate vacuum, the air signal is no longer visible, and the scattered signal through the sample persists for much longer times.
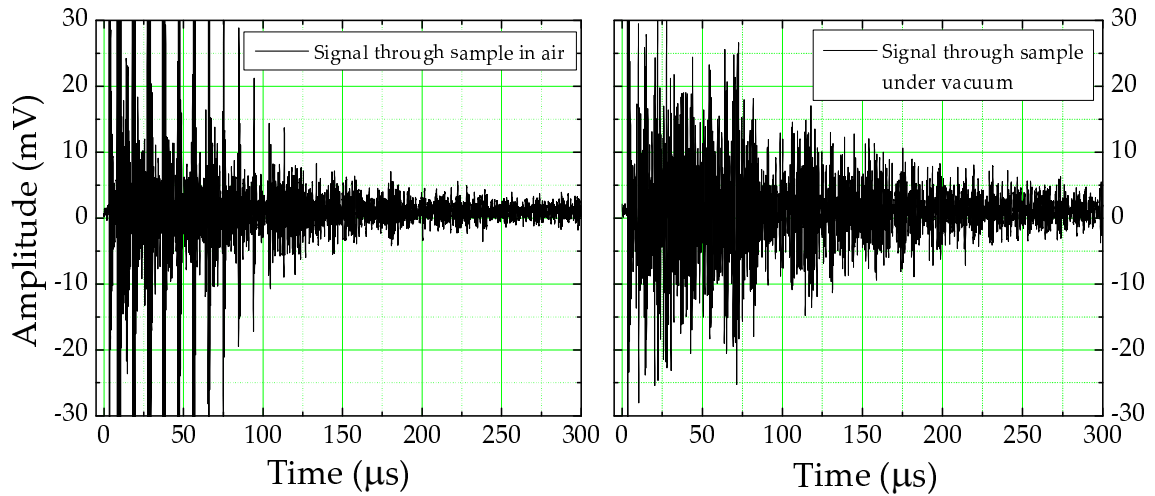
Figure 5.11: Effect of Vacuum on Signal Transmitted through Sample – When air is removed from between the transducers and from the sample pores, the reflections through air disappear and the magnitude and duration of the signal through the sample increases, due to decreased damping. These signals were averaged 1000 times, and the signal through the sample under vacuum was taken at a pressure of 200 mTorr.

Chapter 6

Data and Analysis

6.1 Overview

When the elastic modes of a sample are excited by a short ultrasonic pulse, peaks are found in the Fourier transform of the transmitted signal, indicating the normal modes of vibration of the system. If the sample is small enough, the peaks corresponding to the individual modes will be sufficiently well-separated to be resolved. Once these modes are located, they can be counted and binned to yield the density of states. An example of a signal through one of the samples and its fast Fourier transform (FFT) are shown in Figure 6.1 and Figure 6.2 below.

Figure 6.1: Signal Transmitted through Sample – The signal persists for a long time, showing evidence of multiple scattering. The signal has been averaged 10,000 times.
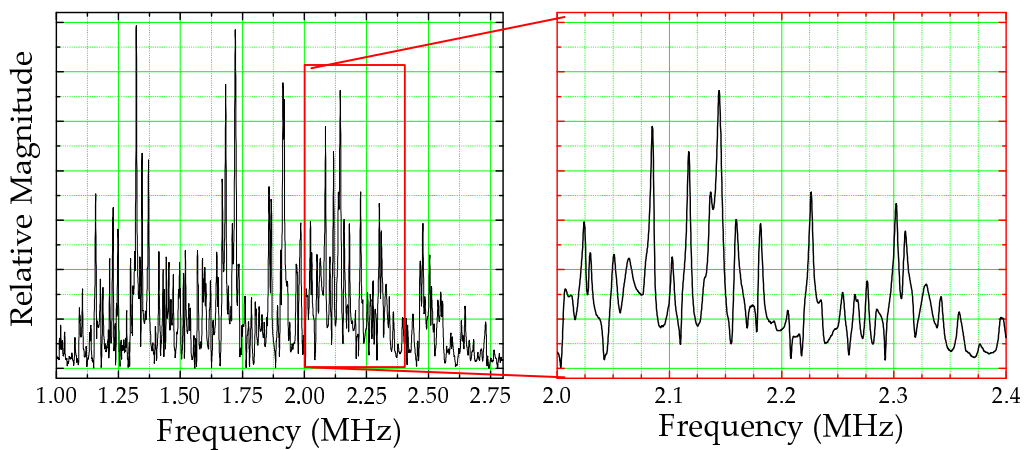
Figure 6.2: Fourier Transform of Transmitted Signal – The peaks corresponding to resonant modes of the sample are clearly resolvable.

From the above example, one might expect that a simple counting and binning procedure (as well as averaging over many samples, in order to measure the average properties of the medium) is all that is required to properly measure the density of states. One must ensure, however, that *all* the modes in a given sample have been accounted for. For example, in one given orientation of a sample, certain modes may only be very weakly coupled to one or both of the transducers, and therefore may not be observed in the FFT. Despite the unlikelihood of degeneracies, or near-degeneracies, due to level repulsion (c.f. § 2.8), adjacent modes may still not be sufficiently separated to be resolvable, if the density of states or peak width is too great.

Much of this chapter is dedicated to addressing the above concerns, wherein a statistical approach to mode counting is shown to be necessary, and is therefore developed.

6.2 Signal Processing Concerns

In order to maximize the usefulness of the acquired spectra, care was taken to properly process the signal to minimize distortion and artifacts. Proper truncation, zero-padding (oversampling) and windowing were essential to getting the best possible spectra for subsequent analysis.

6.2.1 Truncation and RF Pulse Removal

In the signals acquired in my experiments, an electromagnetic RF pulse was always present (c.f. § 5.3.3.4: RF Pick-Up) before, or at the beginning of, the actual elastic signal. Electronic noise was, of course, present through the entire length of the acquired waveforms, but dominated at long times, when the elastic signal had mostly died off. It was therefore useful to truncate the beginning and end of the time domain record, so that the region used for the FFT had the minimum amount of noise and the maximum elastic signal. The appropriate truncation point for the RF pulse removal was easily estimated by eye[1]. A baseline amount of electronic noise (post-averaging) was measured by recording a short pre-trigger region with each signal. The signal was truncated approximately when the signal-to-noise ratio (after averaging) fell to unity. These features are shown in Figure 6.3 on the following page.

---

[1]Often the RF pulse was greater in amplitude than the acoustic signal. Since this portion of the signal was truncated, the waveform was freely allowed to clip in this region, thereby maximizing the useful resolution of the digitizing oscilloscope for the acoustic signal.

Figure 6.3: Truncation Details – The important features used to determine the truncation points are shown. The original signal is shown in the upper-left, the truncated signal in the bottom-right. In the upper-right panel, the RF pulse begins to die off before the bulk of the elastic signal arrives, making this a good choice for truncation. In the bottom-left, the pre-trigger noise level is used to determine when the signal-to-noise ratio is approximately unity, indicating a good region to truncate to reduce electronic noise contributions.

The effect that the truncation has on the spectrum can be seen in Figure 6.4 on the following page. Note the decreased baseline magnitude (especially at high frequencies) and increased smoothness.

Figure 6.4: Spectrum Showing Effects of Signal Truncation – The spectrum from the truncated waveform (blue, dashed) shows a significant improvement in terms of smoothness and reduced baseline noise over the spectrum from the original waveform (black, solid). The magenta and red boxes show the zoomed regions. Note the presence of side-lobe oscillations in the bottom left spectrum (red box).

The figure on the previous page shows that proper truncation can help make the peaks more easily identifiable by reducing the noise that may confuse the resolvability of closely-spaced or small-amplitude peaks.

6.2.2 FFT Padding and Windowing

When spectra are obtained from real data, one must deal appropriately with the inherently finite record length and sampling rate. The finite sampling rate can cause aliasing; however, an analog low-pass filter (with frequency less than the sampling rate) was used to avoid this possibility, as well as to minimize high-frequency electronic noise. Padding and windowing are two common ways of minimizing the effects of finite record length.

Padding simply increases the record length of the waveform by adding zeros at the end of the record. This, in turn, increases the frequency resolution of the FFT, allowing for a more accurate observation of peak shapes and other sample features. Because the FFT algorithm works more efficiently when the record length is an integer power of 2, the signals were all padded to contain 2^{N+1} points, where 2^N is the nearest integer power of 2 greater than the record length of the sample.

A signal with a finite duration is analytically equivalent to multiplying an infinite signal with a "top-hat" or "boxcar" function. The boxcar function is defined as

$$B(t, \Delta t) = \begin{cases} 1, & 0 \leq t \leq \Delta t \\ 0, & \text{otherwise} \end{cases} \tag{6.1}$$

for a signal of duration Δt. If an *ideal* (infinite duration) signal is given by $G(t)$, the measured Fourier transform will be convolved with a sinc function. That is, if

$$g(f) \equiv \mathcal{F}\big[G(t)\big] \tag{6.2}$$

(where $\mathcal{F}$ denotes the Fourier transform operation), then

$$b(f) \equiv \mathcal{F}\big[B(t, \Delta t)\big] = \Delta t \, \text{sinc}(\pi f \Delta t), \tag{6.3}$$

and from the convolution theorem,

$$g'(f) \equiv \mathcal{F}\big[G'(t) \equiv G(t) \times B(t)\big] = g(f) * b(f), \tag{6.4}$$

where $g'(f)$ is the measured Fourier transform. Thus, a finite record length Δt will cause oscillations in the spectrum adjacent to each peak (known as "side lobes" or "spectral leakage"), with spacing

$$\Delta f = \frac{1}{\Delta t}. \tag{6.5}$$

This means, for example, with a record length of approximately 300 μs (as in Figure 6.3, bottom right), the spectrum from the truncated waveform (Figure 6.4 – blue, dashed) should show oscillations with a spacing of approximately 3.3 kHz, and this is clearly seen at around 3.10 MHz in the graph on the bottom left in Figure 6.4.

To minimize the spectral leakage (and thereby maximize resolvability), the original time-domain signal was multiplied by a window function, which smoothly approaches zero at the beginning and end of the record. There are many different types of window functions commonly used, and often spectral leakage is reduced at the cost of increased peak width; Therefore, the window function should be chosen based on the application [Harris 1978][1]. For the analysis of my experiments, a

---

[1] A good introductory discussion on the use of windows, as well as some other concepts found in this section, can be found at http://www.steema.com/FFTProp/FFTProperties/FFTProperties.htm.

"-74dB 4-term Blackman-Harris" window was chosen as it provided very good side lobe reduction and minimal peak broadening. The effect of windowing can be seen in Figure 6.5 below.
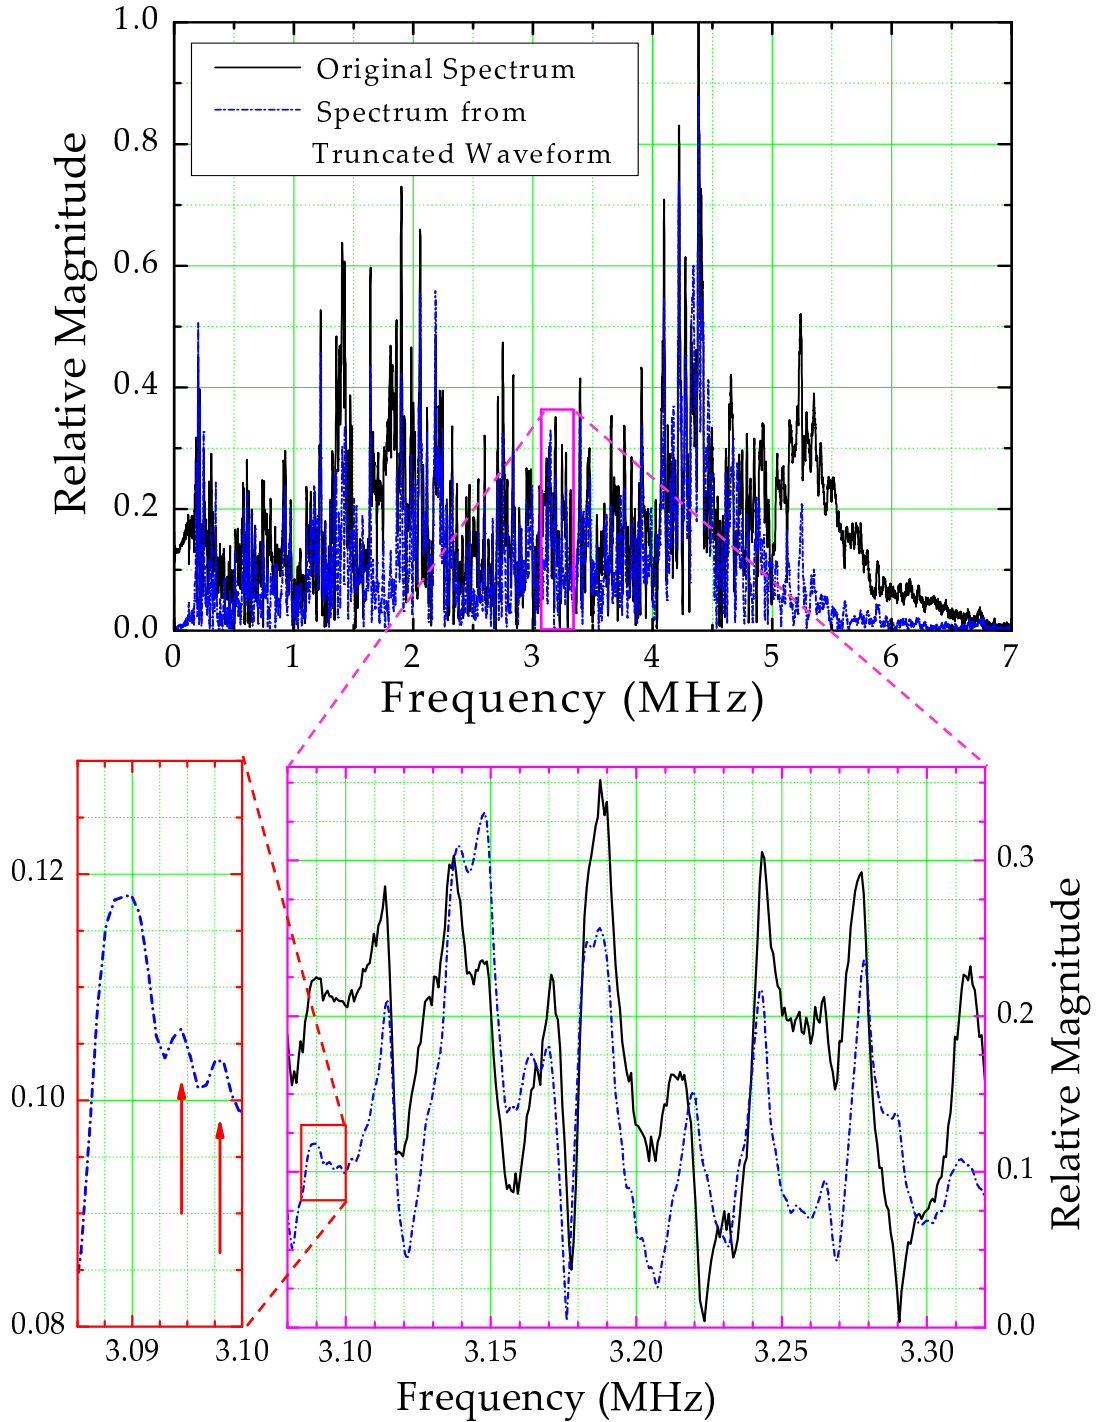

Figure 6.5: Spectrum Showing Effects of Windowing – The spectrum from the windowed waveform (red, dashed) shows a significant improvement over the spectrum from the original waveform (blue solid). Notice the increased resolvability and lack of side lobe oscillations, despite the reduced overall magnitude.

The spectrum from the windowed signal shows increased resolvability and no noticeable side-lobe oscillations. Any short-lived, highly damped modes that may exist will be given less weight, as the window significantly reduces the amplitude of the signal at early times. This may actually be beneficial, as the strong damping of these short-lived modes might indicate that they are greatly affected by coupling with the transducers, and may not well represent the actual modes of the sample.

6.3 Spectrum Reproducibility

As indicated in the overview of this chapter (§ 6.1), one would like to be sure that all the modes present in a particular sample have been accounted for. If the spectra acquired for a particular sample are the same (or reasonably close) under varying experimental conditions, one could have some confidence that all of the normal modes were indeed being seen. In the following section, the effects of various experimental conditions on the reproducibility of the spectra will be investigated.

6.3.1 Sample Reorientation

When the samples are mounted between the generating and receiving transducers, there are several experimental factors that may affect the transmitted signal. Though the mesoglass studied is indeed isotropic, the points on the sample surface that are in contact with the transducers may differ under different orientations, and this may affect the coupling to different modes and/or affect their damping. For example, the nodes in the speckle pattern on the surface of the sample will appear at different frequencies for different positions, implying both that some modes will likely be missed for any given orientation and that the sample can cause different modes to be missed. Also, the transducer displacement is not uniform across its surface, so the sample's position on the transducer surface can also affect the transmitted signal. An example showing the spectra acquired from the same sample in two different orientations is shown in Figure 6.6 on the following page.

Some features of the spectra persist under reorientation, but several peaks appear to be shifted or missing from one spectrum to the other. Furthermore, identifying

common features between such spectra becomes increasingly difficult as the density of states increases (e.g., for larger samples). Thus, the mode counts from a single spectrum are insufficient to determine the actual density of states for a given sample.
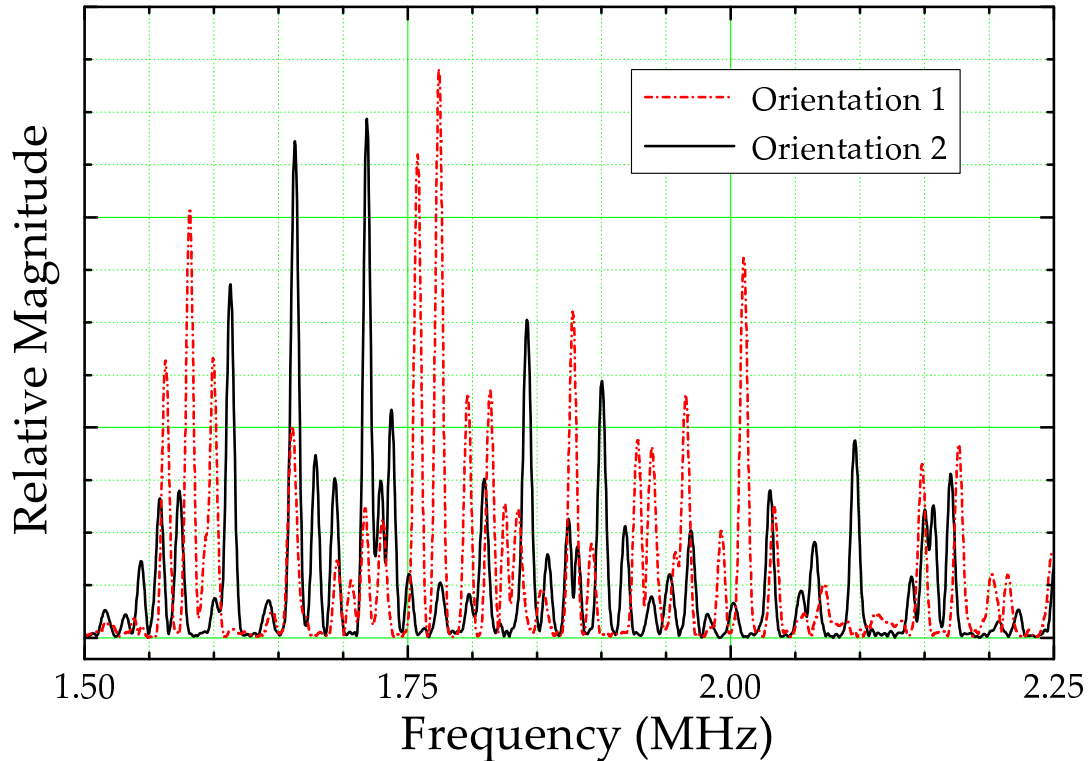
Figure 6.6: Effect of Sample Reorientation – The spectra from the same sample in two different orientations (that is, with different points in contact with the transducers) are shown. Though the spectra look quite similar at some frequencies, the location and number of modes is not consistent.

6.3.2 Contact Pressure

The contact pressure between the transducers and the sample also had a significant effect on the signals. Increased contact pressure caused the modes to broaden significantly, decreasing the resolvability of the spectra. Because of this, the samples were mounted with the minimum contact pressure required to produce an

observable signal[1]. Furthermore, changing the contact pressure caused the modes

to shift in frequency unpredictably, making it difficult to exactly reproduce a signal,

even if the sample and orientation are the same. The effect of contact pressure on

the acquired spectra is shown in Figure 6.7 below.

Figure 6.7: Effect of Contact Pressure – Increasing the contact pressure
causes the modes to broaden and shift (and increase in magnitude). Some of
the modes can no longer be resolved when the pressure is increased (upper
trace). Note that these data were not windowed, so that the peak widths and
shapes could be more meaningfully compared.

---

[1] Properly mounting the sample often required some trial and error – the true signal
was sometimes not observable until the air signal was removed by closing the
vacuum chamber and pumping down!

Again, because this experimental condition had a large effect on the resolvability and location of the modes, a single spectrum was insufficient to determine the density of states for a particular sample.

6.3.3   Vacuum Pressure

Though only a moderate vacuum was required to remove the air signal, the vacuum pressure did have some effect on the acquired spectra, as shown in Figure 6.8 below.
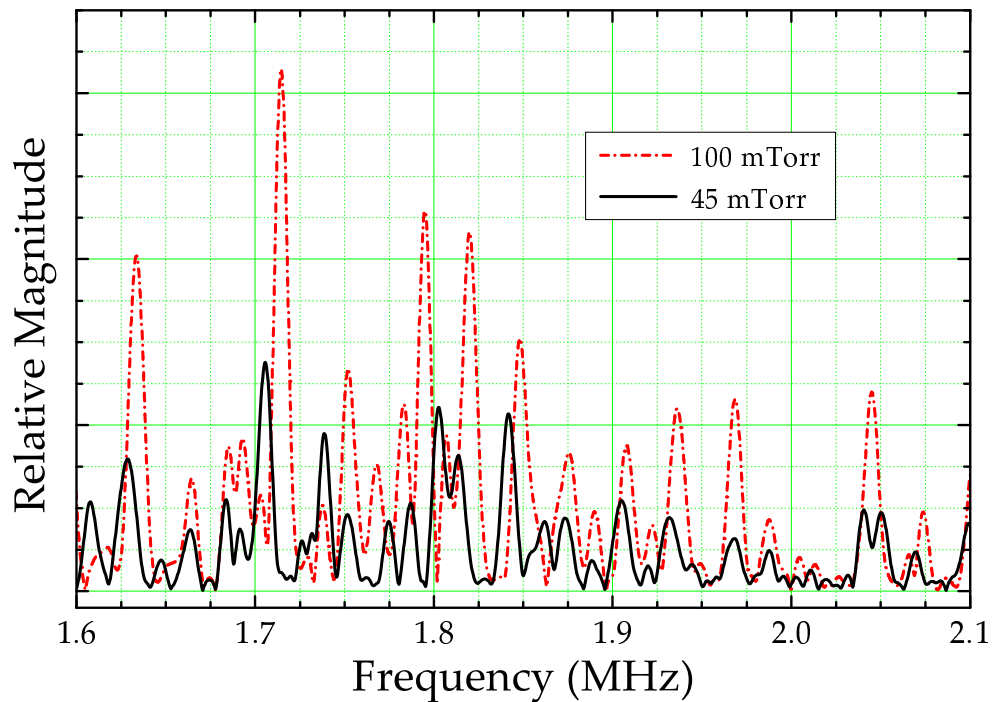
Figure 6.8: Effect of Vacuum Pressure – The spectra for the same sample mounting under different vacuum pressures are shown.   Though both spectra are of similar quality, the peak locations and magnitudes appear to depend on the vacuum pressure.

The spectra shown in the above figure were taken from the same sample mounting; the higher pressure signal was simply acquired first, and then the chamber was further evacuated.   While neither spectrum shows a clear advantage in terms of

87

resolvability, the peak amplitudes and positions are still somewhat different. This again shows the large effect that the experimental conditions can have on the details of the spectra, further illustrating the need for a statistical approach to analysing the mode-counting experiments.

6.4 Statistical Mode Counting

Based on what is shown in the previous sections, one cannot be confident that any single spectrum displays all the modes of a sample in the frequency range of interest. Not only do some modes not appear in some spectra, but the locations (frequencies) of the modes also change with experimental conditions, making unambiguous identification (and therefore easy enumeration) of the modes impossible, even with multiple trials. Instead, the true density of states must be statistically inferred from the distribution of modes counted.

Assume, first of all, that in any one data set for a particular sample, there is a probability p of detecting any one of the n total modes in the frequency range observed. For each mode, there are only two possible outcomes: either the mode is detected or it is not. Therefore, the binomial distribution should be used as a parent distribution for the number of modes observed. That is, the probability P of detecting x of the n total modes is given by [Bevington and Robinson 1992]:

$$P(x;n,p) = \frac{n!\,p^{x}(1-p)^{(n-x)}}{x!(n-x)!}.$$ (6.6)

The mean (μ) of the distribution is simply given by

$$\mu \equiv \langle x \rangle = np,$$ (6.7)

and the standard deviation (σ) of the distribution is given by

$$\sigma = \sqrt{np(1-p)}.$$ (6.8)

If several spectra are acquired, corresponding to different orientations and contact points of the same sample, μ and σ can be measured for the mode counts in each set

89

of measurements. From these, the total number of modes (n) (as well as the probability of finding a mode (p)) can be calculated by solving Equations 6.7 and 6.8, yielding

$$n = \frac{\mu^2}{\mu - \sigma^2} \, , \tag{6.9}$$

$$p = \frac{\mu - \sigma^2}{\mu} \, . \tag{6.10}$$

If there are N trials (for a particular sample), the uncertainty in the mean and standard deviation are given by

$$\Delta_\mu = \frac{\sigma}{\sqrt{N}} \, , \tag{6.11}$$

$$\Delta_\sigma = \frac{\sigma}{\sqrt{2N}} \, . \tag{6.12}$$

Using standard propagation of uncertainties, this yields the following for p and n:

$$\Delta_p = \frac{\sigma^2}{\mu^2} \sqrt{\frac{\left(\sigma^2 + 2\mu^2\right)}{N}} \, , \tag{6.13}$$

$$\Delta_n = \frac{\mu\sigma}{\left(\sigma^2 - \mu\right)^2} \sqrt{\left(\left(\mu - 2\sigma^2\right)^2 + 2\mu^2\sigma^2\right)\Big/N} \, . \tag{6.14}$$

This gives an estimate (with a measurable uncertainty) of the true number of modes in a given frequency range *for a particular sample*. The data from all the samples then needs to be combined in order to get an idea of the average properties of the medium. Note that the uncertainty in the number of modes given by Equation 6.14 represents only the uncertainty in measuring the number of modes for a given sample; however, the samples are also inherently different from one to the next.

That is, *even if the samples could be measured perfectly*, there would still be statistical fluctuations from one sample to the next. To account for this, a factor of $\sqrt{n}$ should be included[1] in the uncertainty in n, so Equation 6.14 becomes

$$\Delta_n = \sqrt{n + \frac{\mu^2\sigma^2\left(\left(\mu - 2\sigma^2\right)^2 + 2\mu^2\sigma^2\right)}{N\left(\sigma^2 - \mu\right)^4}}.$$
(6.15)

To combine the data from several samples, the results from each sample must be normalized by volume (since the density of states is expected to be proportional thereto; c.f. § 2.6) and the ensemble average can be calculated, weighted by the inverse of the uncertainties. The average probability (p) of observing a mode in a given frequency range can also be found, though no volume-normalization is required.

The volume-normalized peak counts are given by

$$\eta \equiv \frac{n}{V} = \frac{n\rho}{m} = \frac{\mu^2\rho}{m\left(\mu - \sigma^2\right)},$$
(6.16)

with uncertainty from standard error propagation and Equation 6.15:

$$\frac{\delta^2}{\eta^2} \equiv \left(\frac{\Delta_n}{n}\right)^2 + \left(\frac{\Delta_\rho}{\rho}\right)^2 + \left(\frac{\Delta_m}{m}\right)^2 = 1 + \frac{\sigma^2\left(\left(\mu - 2\sigma^2\right)^2 + 2\mu^2\sigma^2\right)}{N\mu^2\left(\mu - \sigma^2\right)^2} + \left(\frac{\Delta_\rho}{\rho}\right)^2 + \left(\frac{\Delta_m}{m}\right)^2.$$
(6.17)

The weighted average and its uncertainty are then obtained in the usual way, i.e.,

---

[1] This is the appropriate uncertainty for counting experiments, wherein the data represent the number of observed events per unit interval [Bevington and Robinson 1992]. In the experiments of this thesis, the events are the presence of modes, and the unit interval is the sample mass and frequency-bin width.

$$\bar{\eta} = \frac{\sum\left(\eta_i/\delta_i^2\right)}{\sum\left(1/\delta_i^2\right)}, \tag{6.18}$$

$$\bar{\delta}^2 = \frac{1}{\sum\left(1/\delta_i^2\right)}. \tag{6.19}$$

The average probability (p) of observing modes is also calculated, but using an *un-weighted* average, in order to get an idea of how likely the modes are to be counted, independent of the measurement uncertainties.

The above calculations are done for each frequency bin independently and the resultant set of $\bar{\eta}$, scaled by bin width, yields the density of states.

6.4.1 Treatment of Empty Bins

Before combining the data from various samples, one notable exception to the above formulae should be dealt with. For sufficiently small samples and sufficiently narrow frequency bins, there may be some cases where *no* modes appear in a given bin for all trials of that particular sample. This would yield $\mu = \sigma = 0$, and leave n, p, Δ_n and Δ_p indeterminate (from Equations 6.9, 6.10, 6.13 and 6.15). One might consider setting $n = \Delta_n = 0$ for this case, but this would essentially give the zero-value infinite weight in the weighted sum of Equation 6.18, and still leave p and Δ_p undefined. This means that a density of states of zero would be "measured" at a particular frequency, even though there is a statistically significant probability of finding a mode at that frequency in other samples. This should be accounted for, so that it might be properly included in the weighted average.

For any one trial, with no a priori knowledge of the uncertainty, it is reasonable to assume Poisson statistics, since one is simply counting modes[1]. To estimate the expected mean and uncertainty in a zero-measurement, the probability of observing $\mu = 0$ (in a single measurement) when the *actual* value is $\nu \neq 0$[2] must be found. This is given by the Poisson distribution

$$P_P(\mu; \nu) \equiv \frac{\nu^{\mu}}{\mu!} e^{-\nu} \tag{6.20}$$

evaluated at $\mu = 0$:

$$P_P(0; \nu) = e^{-\nu}. \tag{6.21}$$

The probability of observing zero *every time* if N measurements are made is then given by

$$\left(P_P(0; \nu)\right)^N = e^{-\nu N}. \tag{6.22}$$

Summing over all possible ν, the probability of consistently observing zero for N measurements with an unknown $\nu \neq 0$ is given by

$$P(0; N, \nu \neq 0) \equiv \sum_{\nu=1}^{\infty} e^{-\nu N} = \frac{1}{e^N - 1}. \tag{6.23}$$

The probability that the actual value is zero $(\nu = 0)$ is then simply given by subtracting the above quantity from unity:

$$P(0; N, \nu = 0) = 1 - P(0; N, \nu \neq 0) = 1 - \frac{1}{e^N - 1}. \tag{6.24}$$

---

[1] That is, if only one trial was conducted and μ modes were counted, it would be reasonable to assume an uncertainty of $\sqrt{\mu}$.

[2] ν is used as the mean value of the Poisson distribution, to avoid confusion with the measured mean μ.

The best estimate for n is then given by the expected value of v (the "actual" value) when zero is consistently measured:

$$n(0) = \langle v \rangle = \sum_{v=1}^{\infty} v \left(P_P(0;v) \right)^N = \sum_{v=1}^{\infty} v e^{-vN} = \frac{e^N}{\left(e^N - 1 \right)^2}. \qquad (6.25)$$

The best estimate for Δ_n, should be given by the expected deviation of v, so the second moment of the distribution is used to obtain the RMS uncertainty for a measured value of zero:

$$\Delta_n^2(0) = \langle v^2 \rangle = \sum_{v=1}^{\infty} v^2 \left(P_P(0;v) \right)^N = \sum_{v=1}^{\infty} v^2 e^{-vN} = \frac{e^N \left(e^N + 1 \right)}{\left(e^N - 1 \right)^3} \qquad (6.26)$$

Therefore, when no modes are observed in a frequency-bin for a particular sample, a reasonable value and uncertainty can be assigned to n based on the number of trials performed. For large N, Equations 6.25 and 6.26 both tend to e^{-N}, which intuitively seems to be a sensible value and weighting factor in Equation 6.18, for the case of a zero-measurement.

The most probable value for p could also be calculated based on the above assumption of Poisson statistics. However, that assumption was only intended to give a rough idea of how to assign an appropriate value to n in the case of a zero-measurement, in order that it might be sensibly included in the average. There is no obvious reason to infer a particular value for p, given that one simply does not know if there are *actually* no modes in the given frequency range, or if they have simply all been missed. For this reason, in the case of a zero-measurement, a uniform probability distribution for p is simply assumed, yielding:

$$p(0) = \Delta_p(0) = 0.5. \qquad (6.27)$$

6.5 Peak Finding

One of the most challenging aspects of the data analysis for my experiments was to accurately and objectively identify the peaks in the spectra amongst the background noise. The peak shape, width and height were all found to vary greatly, as is typical for many coupled harmonic oscillators with different masses and degrees of coupling and damping (c.f. § 2.5). A method of peak-finding was eventually developed wherein a "confidence level" was established for each local maximum in the spectra, based on the relative height to the adjacent local minima[1]. The peaks were then accepted or rejected using a threshold for this confidence level. The value for the threshold was set by the properties of the individual spectra, as well as by allowing this parameter to vary and searching for the result that yielded the most self-consistent set of results for the system. Though there is still some subjectivity involved in this method, it was found to yield reasonable and consistent results. Furthermore, there should be some leniency in choosing the threshold, since the analysis already assumes a finite probability $(1-p)$ of missing modes (c.f. § 6.4). The details of the peak-finding process are discussed in the following two sections.

6.5.1 Noise Rejection and Thresholding Method

After the signal was processed (as discussed in § 6.2), the peaks were located using the following steps:

---

[1] The advantage of windowing can be clearly seen by considering this method of analysis in light of the example of Figure 6.5.

1. The spectrum was smoothed using adjacent-point averaging, in order to eliminate small point-to-point fluctuations. Usually ± 4 points were used to compute the average, but up to ±6 points were sometimes used, depending on the noise level and resolution of the spectrum.

2. All of the local minima and maxima were located, by looking for changes in sign of the two-point derivative.

3. A "confidence level" was established for each maximum by taking the sum of the differences between the amplitude of the maximum and adjacent minima.

4. The average confidence level for the entire spectrum was calculated, and all maxima with a confidence level below a given minimum percentage ("noise threshold") of the average were discarded. This percentage level could be input by the user, but was 0.2% by default, though the values used ranged from about 0.1% to about 2%. (The goal of this stage was to ignore all local maxima that were *definitely* noise). This parameter was usually easily established by looking at the spectrum, and rarely needed to be varied.

5. The average confidence level was recalculated with only the remaining peaks.

6. All peaks with a confidence level greater than the "peak threshold" were counted. This threshold was given as a percentage of the new average confidence level, and its chosen value was found to vary widely (from around 0% to 60%) based on the quality of the spectrum as well as the variance in peak heights. It is this parameter that was adjusted when searching for the appropriate threshold to give the most reasonable and self-consistent results for calculating the density of states.

This "double-rejection" method of peak finding was found to be more robust to variable peak densities and amplitudes in the spectra.

6.5.2 Interactive Threshold Choosing

In order to choose the correct threshold for identifying peaks, a process was used wherein the effects of different threshold values on the number of peaks found and the calculated value for n could be easily assessed. In this process, the peak threshold parameter was varied, and the resultant values for n, Δ_n, μ and σ were graphed as a function of threshold value (usually this was done for a wider range of frequencies than the bin width used for the final calculations, in order to obtain better statistics – the possible frequency dependence of the density is of less concern at this stage of analysis). The motivation for this was to find a range of values for the peak threshold where the measured value of n is locally unaffected by small changes in the threshold value (suggesting that the modes are being found probabilistically[1], and the threshold value is appropriate), and where the uncertainty Δ_n is relatively small. After choosing the threshold (for a particular sample), the spectra with the "counted" peaks marked were viewed, to verify their accuracy. If the modes were obviously miscounted, a new threshold was chosen before moving on to the next sample, and if one of the trials looked significantly noisier than the others, it was excluded from the analysis. The chosen threshold is then used for all trials of the sample in question, in order to keep the analysis as consistent as possible.

Once the appropriate parameters were chosen, the peak counts were binned and analysed according to Section 6.4, yielding the density of states.

---

[1] According to § 6.4, the number of modes calculated (n) should be locally independent of the probability of finding a mode (p), which is affected by the peak threshold chosen.

Chapter 7

Results and Discussion

7.1 Overview

In this chapter, the results of my mode-counting experiments are presented. Some important features of the acquired spectra and analysis are noted, and the measured mode-counting probability is shown. The density of states and fracton dimension is shown for the range of frequencies measured, and quantitative estimates based on the models presented in Chapter 2 are given. The results are also compared to other theoretical predictions and measurements made in similar systems.

7.2   Notes on Acquired Spectra

Before looking at the results of the measurements in greater detail, it is worthwhile to note some observed features of the acquired signals. First, the spectra acquired were heavily weighted toward the low-frequency side of the transducer response, and this effect became more pronounced when using higher frequency transducer pairs. This effect may be due partially to the electrical impedance mismatch of the transducers (c.f. § 5.3.3.3), and a qualitatively similar low-frequency bias was found when measuring reference signals through air (c.f. § 5.3.3.2); however, this bias was even greater when the actual experiments were conducted (that is, in a vacuum, with the transducers in contact with the sample). This low-frequency bias is likely due to the acoustical impedance mismatch between the transducer and the load (vacuum/sample). The usable frequency range of each transducer pair is listed in Table 7.1 below.

The spectra acquired from the different transducer pairs also show differences in terms of peak width. The width of the peaks in the acquired spectra determines the resolvability of the modes, and thereby limits the measureable un-normalized density of states.   Approximate average peak widths and quality factors were measured for each transducer pair, and the results are shown in Table 7.1 below.

Transducer Frequency (f_c) (MHz)	**0.5**	**1.0**	**2.25**	**3.5**	**5.0**
Usable Frequency Range (MHz)	0.3–0.7	0.5–1.3	1.0–2.6	0.5–3.5	2.0–4.0
Approximate Peak Width (kHz)	2.0	4.0	5.5	15	10
Approximate Average Q	250	225	325	135	300

Table 7.1: Properties of Acquired Spectra for Each Transducer Pair

There is a general trend of increased width and low-frequency bias for higher frequency transducers. The 3.5 MHz pair shows an abnormally large range of useable frequencies and the peaks are especially wide; this may be related to the fact that these transducers also had a significantly larger surface than the others. Spectra showing a narrow range of frequencies (100 kHz) for each transducer pair are shown for reference in Figure 7.1 below.
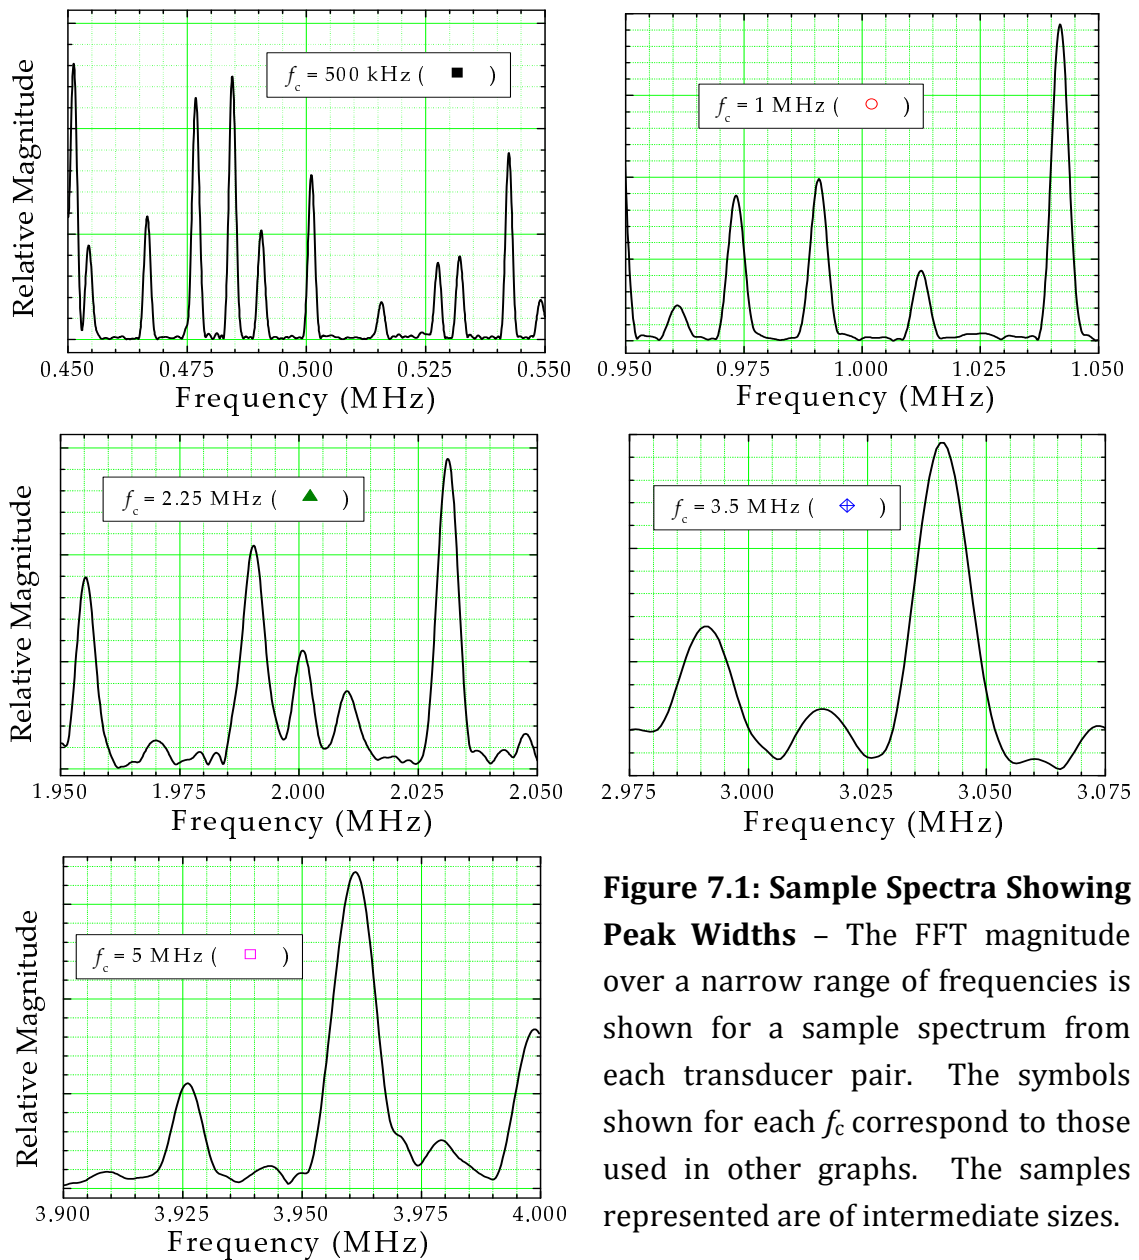
Figure 7.1: Sample Spectra Showing Peak Widths – The FFT magnitude over a narrow range of frequencies is shown for a sample spectrum from each transducer pair. The symbols shown for each f_c correspond to those used in other graphs. The samples represented are of intermediate sizes.

7.3   Peak Thresholding Values

As discussed in § 6.5.2, the dependence of n, Δ_n , μ and σ on the "peak threshold" (c.f. § 6.5.1) was plotted, and a region was found where the calculated number of modes n is relatively independent of the number of peaks counted μ, and where Δ_n is reasonably small.   The appropriate threshold was chosen near the onset of this region (to avoid missing modes unnecessarily), and the results of the peak-finding program with this threshold were viewed and verified.   One example is shown in Figure 7.2 on the following page.

The presence of the plateau indicates that the statistical mode counting method is working as expected, and that the missed peaks are being properly accounted for. The results of the peak-finding program agree quite well with the peaks that one might identify by eye, with perhaps a few modes missing.

The threshold level for each sample was chosen using the method described in Section 6.5.2. The appropriate threshold was not always quite as easy to identify as shown in the example of Figure 7.2.   The plateau region was often smaller, and sometimes several plateau regions could be seen. In cases such as this, the results of the peak-finding program were used to help choose an appropriate value.   It is worthwhile to note that the number of modes calculated was quite robust, as long as the chosen threshold was reasonable according to the stated criteria.
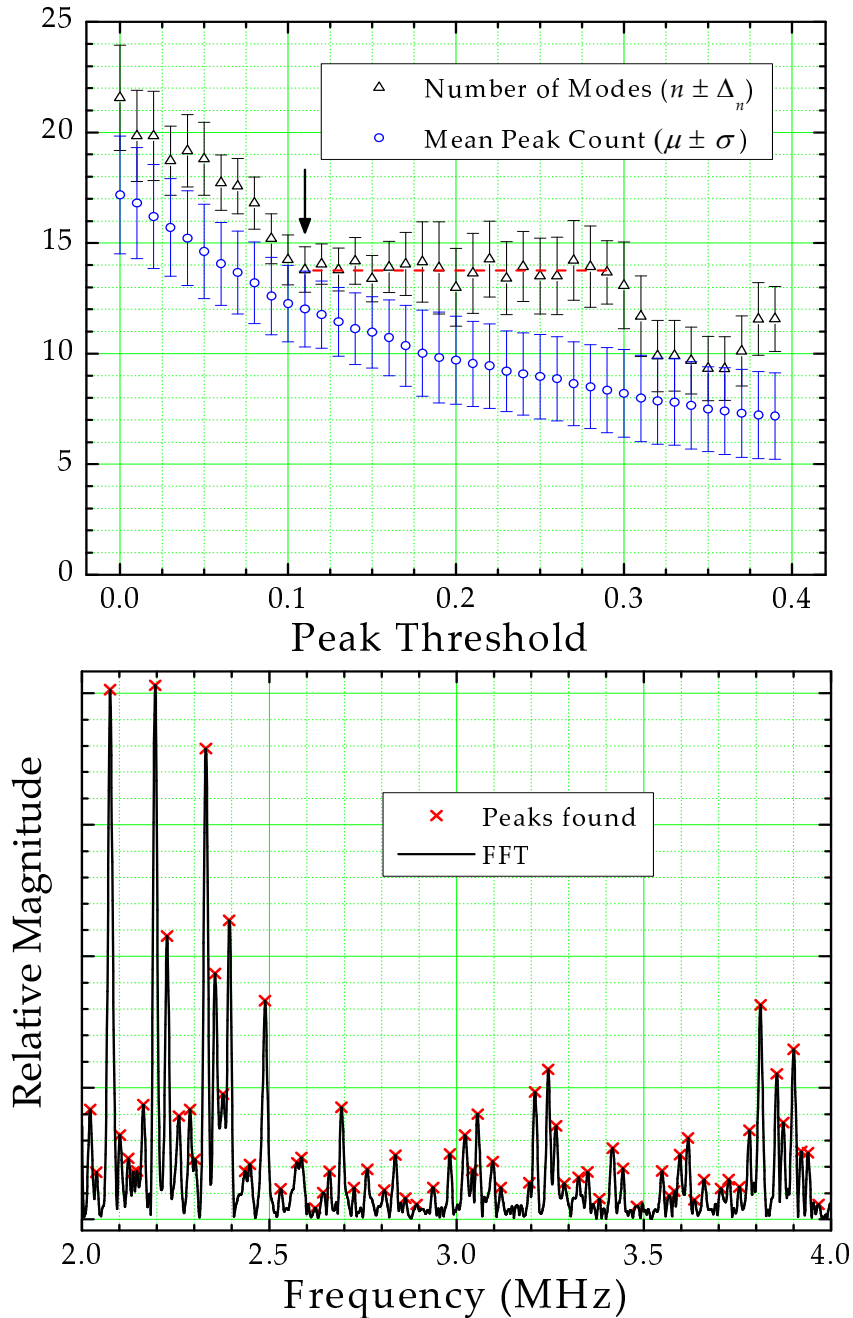
Figure 7.2: Choosing the Appropriate Threshold for Peak Finding – The mean number of peaks counted (μ) and the calculated number of modes (n) are shown as a function of peak threshold (left). A plateau in n (while μ is decreasing) is observed, signifying the correct range for the peak threshold (red line shown as a guide to the eye). The threshold chosen is indicated by the arrow, and the results of the peak-finding program for this value are shown (right). Note that though some modes may be missed, they are accounted for by the method of analysis.

7.4 Volume Linearity

As discussed in Section 2.6, the density of states is expected to be proportional to sample volume. To verify this, the calculated number of modes (*n*) was plotted against sample volume for each frequency range. One example is shown in Figure 7.3 below.

Figure 7.3: Volume Scaling – The density of states is plotted against sample volume for the f_c = 5 MHz range. The data are not well-represented by the straight line fit.

One can clearly see that the measured density of states does not appear to scale linearly with volume. This can be understood by looking closely at the spectra obtained from the larger samples. In these spectra, no "dead" frequencies exist.

That is, all the modes are immediately adjacent to each other, and this suggests that there may in fact be a significant amount of overlap and the mode spacing is smaller than their resolvability. To illustrate this point, spectra obtained in the f_c = 5 MHz range from one small and one large sample are shown in Figure 7.4 below.

Figure 7.4: Comparison of Spectra from a Large and Small Sample – The larger sample may not be sufficiently small for the density of states to be measured directly. This is partially evidenced by the lack of any "dead" frequencies – each mode is immediately adjacent to the next. The peaks in the small sample, however, are easily identifiable. In this frequency range (f_c = 5 MHz), the peak width in the larger sample is greater than the average mode spacing, and the sample was not included in subsequent analysis.

This was further evidenced by the fact that even when a peak threshold of zero was chosen (c.f. § 6.5), the volume-normalized density of states was still significantly lower than that of the smaller volume samples. No other explanation of why the density of states might *not* scale with volume (c.f. § 2.6) is known for this system, so the conclusion that this lack of linearity must simply be due to a saturation in the measureable density of states seems reasonable.

To correct for this artifact, the larger samples were removed from the analysis one-by-one (beginning with the largest) until the remaining samples were well-described by a straight line on the volume scaling plots. Though this procedure is somewhat ad-hoc, the appropriate samples to include/exclude were quite easy to identify, as in Figure 7.3 above. Furthermore, this leads to a consistent result for the density of states in the overlapping frequency ranges from different transducer pairs, as seen in Figure 7.6. This correction was only necessary for the two highest frequency ranges studied, where the modes were found to be significantly wider (c.f. § 7.1).

The corrected volume-scaling plots show reasonably good agreement with a straight line fit for all frequency ranges, and are shown in Figure 7.5 on the following page. Note that the slope of the linear fits gives the average normalized density of states for each frequency range.
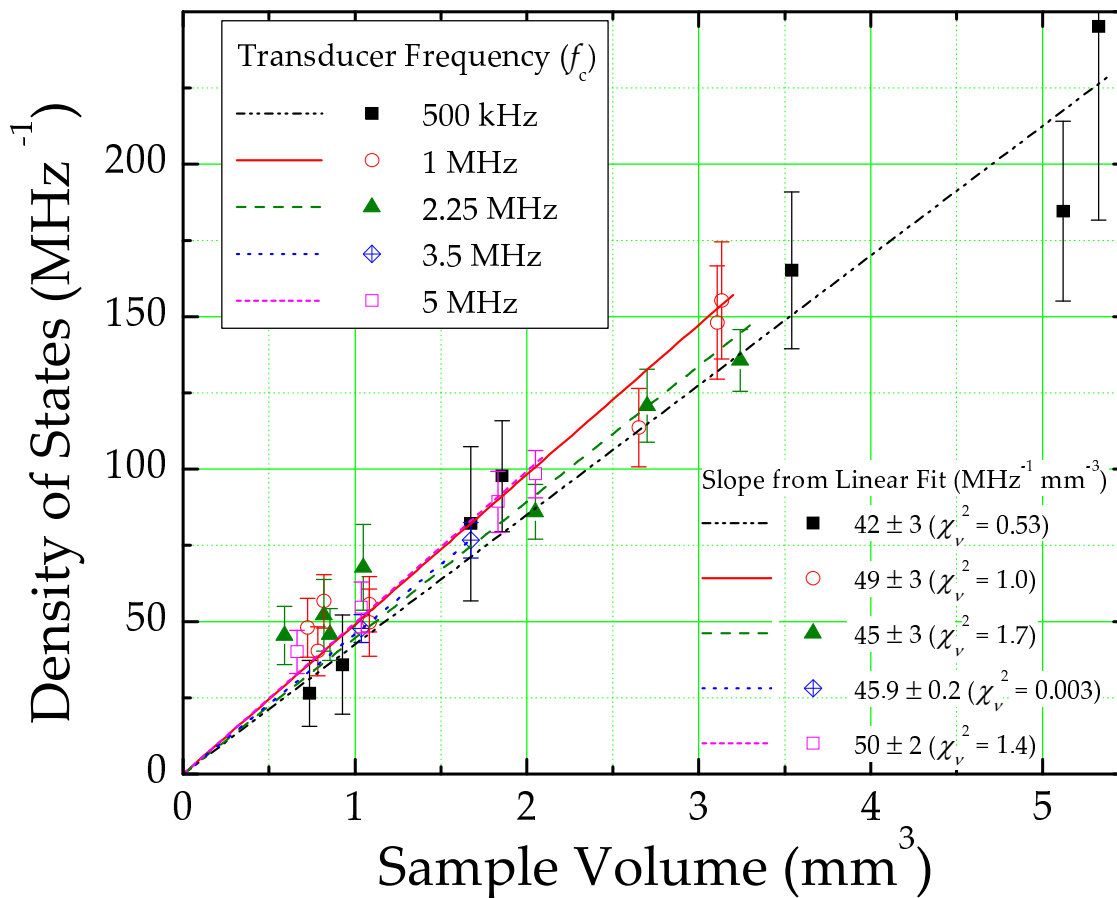
Figure 7.5: Corrected Volume Scaling – The density of states vs. sample volume is shown for all frequency ranges, with some of the larger samples removed as discussed. The slope of the linear fit gives the average normalized density of states for each frequency range. The data are well-described by the linear fits.

7.5 Mode-Counting Probability

The analysis discussed in Section 6.4 gives an estimation of the average probability (p) of observing any given mode, via Equations 6.10 and 6.13. This probability was found to vary widely, depending on the mass of the sample and transducer pair used, and was often smaller near the edges of the useable transducer bandwidth; this variance far outweighed any regular dependence on frequency that may have been present. Though this is an encouraging indication that the statistical mode-counting analysis developed is robust (since n is independent of p within reasonable limits, c.f. § 6.4), it does not allow for a very precise estimate of the true number of modes present based on the number of modes counted in *one* experiment. Still, the (un-weighted) average of p and its variance can be used as a rough estimate of the final results, based on a single measurement in a similar type of experiment. A value of $p = 0.7 \pm 0.2$ is measured, where the uncertainty shown is given by the variance of p, which is equal to the average of the measured uncertainties Δ_p.

Recall that the larger samples at higher frequencies were excluded from the final analysis. The value of p obtained can only be used as a rough estimate of the mode-finding probability if the samples are sufficiently small that the number of modes counted is not significantly limited by the resolvability of the modes.

7.6 Measured Density of States

The density of states was measured over a broad range of frequencies by the methods presented in this thesis, and the results are shown in Figure 7.6 below.

Figure 7.6: Measured Density of States – The normalized density of states over the range of frequencies measured is shown. The result is nearly constant, with a mean value of $\mathcal{D} = 46.6 \pm 0.2$ MHz^{-1} mm^{-3}.

Note that that agreement between measurements in overlapping frequency ranges from different transducer pairs is very good. These measurements are the result of direct mode-counting experiments done on 28 samples, each in 4 – 14 orientations (c.f. Table A.1). The error bars shown are statistically significant, and result from a careful treatment of the uncertainties in the experiments (c.f. § 6.4).

7.6.1 Comparison with Estimates

In order to compare the experimentally measured results with the predictions discussed in Chapter 2, the correct inputs to the given expressions in the context of these particular samples must be known, and many of the relevant sample properties can be found in Table 3.1.

For Equations 2.13 and 2.16, the total surface area (S) of the samples must be known, and can be estimated as follows. Based on the observed structure of the mesoglass (via x-ray tomography, c.f. Chapter 4), the medium might be *roughly* modelled by a network of interconnected "tubes" or cylinders, whose radius is approximately equal to the average bead radius (a). The surface area of the tube (S_T) would be given by

$$S_T = 2\pi a L, \tag{7.1}$$

where L is the length of the tube. The volume of a tube (V_T) is given by

$$V_T = \pi a^2 L. \tag{7.2}$$

Thus, the total surface area of a sample can be estimated from the volume, via

$$S = \frac{S_T}{V_T} \phi V = \frac{2\phi V}{a}, \tag{7.3}$$

where the factor ϕV represents the occupied volume of the sample.

The bulk crossover frequency (f_Ξ) also enters into the estimate of Equation 2.16. As discussed in Section 4.4, this crossover frequency may be better estimated by Equations 2.3 and 2.5 than by the smaller fractal bulk crossover length measured via box counting, and therefore, the values given by these equations are used.

The estimated density of states is shown in Figure 7.7 below, with measured results included for comparison.
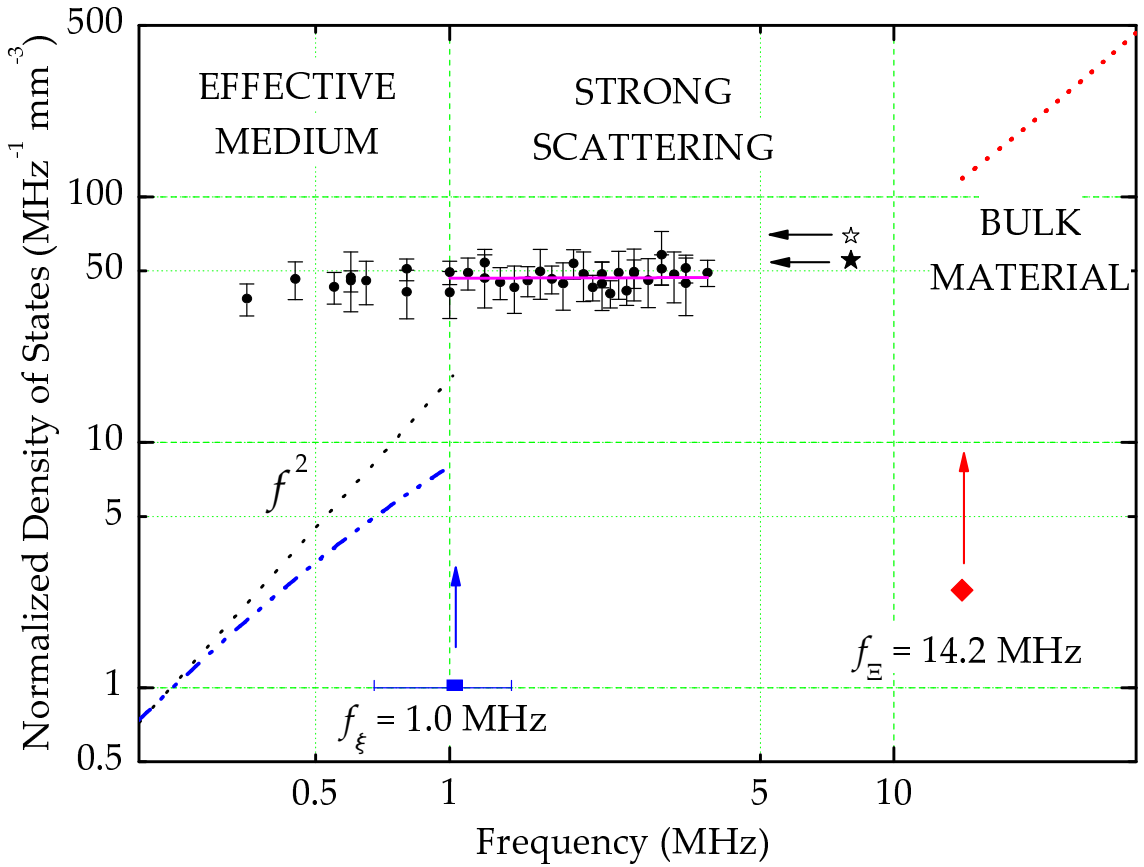
- Measured Density of States
- Power Law fit to data - $f^{0.01 \pm 0.04}$ ($\chi_v^2 = 0.32$)
- ★ Estimated Average in Intermediate Regime = 55 MHz^{-1} mm^{-3} (Eq. 2.16)
- ☆ Estimated Average in Intermediate Regime = 69 MHz^{-1} mm^{-3} (Eq. 2.17)
- Debye Effective Medium Prediction (from Ballistic Velocities) (Eq. 2.2)
- Prediction for Bulk Medium with Surface Modes (Eq. 2.13)

Figure 7.7: Comparison of Measured Density of States with Estimates – The measured density of states at low frequencies is significantly greater than the prediction of the Debye model. In the strong scattering regime, the density of states agrees fairly well with the estimated value. The mean density of states measured in this range is $\mathcal{D}_l = 47.1 \pm 0.3$ MHz^{-1} mm^{-3}, and a power-law fit yields $f^{0.01 \pm 0.04}$, or a fracton dimension of $\bar{d} = 1.01 \pm 0.04$. The uncertainty shown in f_ξ is from the estimated uncertainty in ξ (c.f. § 4.4).

The graph on the previous page shows the predictions for the density of states as discussed in Chapter 2. The low-frequency estimate is given by the Debye model (Equation 2.2), with the frequency-dependent velocities shown in Figure 1.1, and values listed in Table 3.1. The high-frequency estimate is given by a Debye-like model with the inclusion of a surface term (Equation 2.13), with the surface area given by Equation 7.3, and the crossover frequencies are derived from the measured percolation correlation length and the constituent particle radii, according to Equations 2.6 and 2.5.

In the intermediate frequency (strong scattering) regime, the density of states nearly constant and has a measured mean value of $\mathcal{D}_l = 47.1 \pm 0.3$ MHz^{-1} mm^{-3}. This is quite similar to the predicted values of $\mathcal{D}_l = 55$ MHz^{-1} mm^{-3} (from conserving the total number of modes, as in Equation 2.16) and $\mathcal{D}_l = 69$ MHz^{-1} mm^{-3} (from the "rigid bead model, as in Equation 2.17). In fact, the relative agreement of the two predicted values (based on somewhat different assumptions – c.f. § 2.9) suggests that the picture used to estimate the behaviour of the density of states in this regime is relatively self-consistent and robust. A power-law fit to the data in this range yields an $f^{0.01 \pm 0.04}$ dependency, giving a fracton dimension of $\tilde{d} = 1.01 \pm 0.04$ for a fractal (percolation) model applied to this mesoglass (c.f. § 2.7).

The low-frequency estimate is given by the Debye model (Equation 2.2), with the frequency-dependent velocities shown in Figure 1.1, and values listed in Table 3.1. The measured density of states in this low-frequency (effective medium) regime is significantly higher than the Debye prediction, and only a slight downward trend at

low frequencies *may* be inferred from the first point[1]. Since the Debye model only predicts the density of ballistic modes, the enhanced density of states in this regime is consistent with the existence of diffusive modes even at these low frequencies. Even in this range, the acquired waveforms were found to persist for very long times (observable signals lasting up to 1 ms), showing evidence of multiple scattering. From the scattering strength (kl_s) measured for this mesoglass (Figure 1.1, top), the mean free scattering time can be calculated from the equation $\tau_s = l_s/v = kl_s/2\pi f$, and is of the order of 1 µs at 500 kHz. A comparison of the record length with the mean free scattering time shows that multiple scattering is occurring due to the inhomogeneities in the sample, and that the long duration of the waveform is not simply the result of reverberant scattering from the sample surfaces. Though the samples may be too small for the diffusion approximation to be applicable to the data measured at these frequencies, modes of a diffusive nature will exist in larger samples. It is expected that at even lower frequencies, the density of states must agree with the Debye prediction, and this may be an interesting future experimental pursuit.

There is still a large range of frequencies within the strong scattering regime where the density of states has not been measured, and the crossover to the high-

---

[1] A careful observation of the spectra at low frequencies suggests that there is also a slight possibility that some of the resonances of the transducer holders (c.f. § 5.2) are still showing up in the frequency range of the first bin, despite efforts made to avoid this. This would mean that *too many* modes are being measured at low frequencies, though this is unlikely for the first bin, and *very* unlikely for any others.

frequency, bulk material regime has not been observed. These may be interesting

subjects for further experiments.

7.6.2 Comparison with Other Theories and Measurements

The estimate of the density of states is based on a model proposed to explain the

thermal properties of sintered metal powders [Rutherford *et al.* 1984; Maliepaard *et*

al. 1985]. This simple model invokes an f^0 dependence for the density of states, in

agreement with my experiments. This result also agrees with a the measurement of

$\tilde{d} = 1.04 \pm 0.05$ for the fracton dimension in similar glass bead sinters, though this

value was obtained through a different method [Schriemer *et al.* 1996]. Comparison

with the fractal models discussed in Section 2.7.2 gives better agreement with the

"bond-bending" model prediction of $\tilde{d} = 0.89$ than the scalar elasticity prediction of

$\tilde{d} = 4/3$. A numerical simulation of a percolation system on a cubic lattice

[Sheng and Zhou 1991] may show qualitative agreement with my measurements at

intermediate frequencies. Predictions and measurements for the density of states in

an atomic glass may also provide a useful comparison to the results obtained for this

mesoglass.

All of these comparisons warrant some discussion. First, the agreement of the

fracton dimension with that of other similarly constructed glass-bead sinters is not

surprising, but is an encouraging result, considering the very different technique

used in that measurement, and the different bead sizes used in those two sets of

experiments. Second, the density of states model developed for the sintered metal

powders also shows good agreement with the measured result. Preparation of the

metal powder sinters differs from the method used here in that no transitory material is used for "percolation", but rather, the volume fraction of the sinter is simply determined by packing the powder more or less densely before sintering. Nevertheless, the model was applied with some success over a wide range of volume fractions (comparable to those of the present experiments), and gives the correct frequency dependence of the density of states for the mesoglass. Thus, this model may additionally be somewhat applicable to percolation-type systems. Note also that the model of Maliepaard *et al.* simply estimates the effective medium crossover when the wavelength is around ten particle diameters ($\lambda = 20a$), for sinters with the same volume fraction as the mesoglass of this thesis. Estimating the crossover frequency in this way yields (for the lengths and velocities in the mesoglass) a crossover frequency of $f_\xi = 1.28$ MHz, relatively similar to the result obtained from Equation 2.6 and the fractal box counting analysis of Section 4.3, which yields $f_\xi = 1.03$ MHz.

Comparison of the medium studied here with percolation models should be done with some caution, as the sample no longer appears as a percolated structure when observed with x-ray tomography (c.f. Chapter 4); the individual beads are not simply joined at the edges, but the glass has flowed together to form a continuous network. Nevertheless, the occupied volume of the sample is still primarily determined by the locations of the glass beads in the glass/iron mixture *before* sintering (an excellent example of a percolation system on a random lattice), and the material does show fractal behaviour (i.e. non-integer Hausdorff dimension, c.f. § 2.7.1) over more than a decade of length scales (shown in Figure 4.5).

The prediction of scalar elasticity model does not require a fractal structure, but can be derived directly from the Ioffe-Regel condition [Aharony *et al.* 1987]; however, it has been shown that the Ioffe-Regel criterion for localization does not necessarily hold for elastic waves [Sheng *et al.* 1994; Allen *et al.* 1999; Page *et al.* 2004], and though the scattering is strong, there is no evidence so far of localization in the medium used here (via previously measured time-of-flight profiles, c.f. Figure 1.2, inset). The percolation simulation of Sheng and Zhou [1991] is also based on scalar elasticity, and has some qualitative agreement, but shows a very "bumpy" appearance for the density of states at intermediate frequencies, possibly indicating that a simple power law and crossover may not be sufficient to describe the density of states in this range, even for the modestly sized systems of the simulation. (Each simulation in this reference contains $79^3 q \approx 250,000$ particles, for percolation probability $q = 0.55$. The samples studied in the present experiments contain between about 120 – 1500 glass beads each before sintering.) Furthermore, scalar elasticity cannot strictly apply to a system like a porous sinter, where the vector displacements within the medium are important [Webman and Grest 1985], so perhaps it is not surprising that good agreement is not found with these models.

The bond-bending model is expected to be a better representation of a sinter or mesoglass system, and has made excellent predictions for other elastic properties in sintered metal powders [Maliepaard *et al.* 1985]. Despite this, agreement with the predicted value of $\bar{d} = 0.89$ is not found; this may be due to the stiffening of bond-bending modes due to over-sintering. The measurements of the density of states in the present experiments are not, however, sufficiently precise to rule out this value,

and further measurements at higher frequencies are needed to examine this question.

The results of a measurement and simulation of the density of states in amorphous silicon [Allen *et al.* 1999] shows little similarity to the plateau found in my experiments. Notably, a peak in the density of states is found near the low-frequency crossover from ballistic modes, a characteristic found in density of states measurements for many glasses. Such a peak is not seen in present measurements, nor is it seen in simulations of fractal models [Nakayama *et al.* 1994]. This may suggest that the density of states is sensitive to differences in the microscopic details of the elasticity in the mesoglass compared with the particular atomic glass considered in the simulation.

The low-frequency results of the density of states measurements were somewhat surprising, and the enhanced density of states may be due to the presence of diffusive modes, even in this low-frequency regime. Similar results were obtained when the density of states was measured in an aluminum foam, by a similar method [Lobkis and Weaver 2001], where it was concluded that an effective medium model did not well-represent the behaviour of the density of states. No model was found that predicted this behaviour, and this feature continues to be of considerable interest.

Chapter 8

Conclusions

From the Fourier transform of an ultrasonic pulse transmitted through a sufficiently small sample, the individual modes of vibration are separated well enough to be resolved and counted. By doing so, the density of states in a porous mesoglass has been measured directly in the strong-scattering frequency regime.

Using glass beads, samples were constructed for these experiments based on a percolation model for a random lattice. Though the bead structure was no longer apparent after sintering, the medium showed fractal behaviour over a length scale spanning more than a decade: approximately 11 – 240 µm. The fractal dimension measured in this range was $D = 2.60 \pm 0.05$.

The density of states was nearly constant in the frequency range studied, and a mean value of $\mathcal{D} = 46.6 \pm 0.2$ MHz^{-1} mm^{-3} was found. The measured results were fit to a power law in the intermediate frequency regime, and a frequency dependence of $f^{0.01 \pm 0.04}$ was found, indicating a fracton dimension of $\tilde{d} = 1.01 \pm 0.04$. The average probability of observing a mode present in the sample was $p = 0.7 \pm 0.2$.

A simple model was presented to estimate the density of states across three distinct frequency regimes, with two variations on the model (Equations 2.16 and 2.17) in the intermediate range. The measured the frequency dependence of the density of states and its average value are similar to the estimates of the models at intermediate frequencies. In the low frequency regime, the measured density of states was significantly higher than the value predicted by the Debye and effective medium approximations – a surprising result that may be attributed to the presence of diffusive modes at low frequencies. The plateau of the intermediate frequency

regime extends well below the estimated low-frequency crossover, and shows little (if any) convincing evidence of decreasing at the lowest measured frequencies, though it is thought that a crossover to Debye-like behaviour must eventually occur. Measurements also were not performed at sufficiently high frequencies to observe the crossover to the bulk medium (glass) regime. Additional experiments at both ends of the spectrum may be of interest.

At intermediate frequencies, comparison with the results of other measurements and simulations shows a strong similarity with the estimates made for the density of states postulated for sintered metal powders, and the measured result may be consistent with a fractal model based on vector elasticity ("bond-bending"). The predictions of the scalar elasticity model and the measurements and simulations for amorphous silicon both do not agree with the measured results for the mesoglass studied. At low frequencies, no model was found the predicted the observed behaviour, though the results are similar to those obtained for the density of states measured in an aluminum foam.

Thus, the density of states in a highly porous, amorphous mesoglass has been investigated, and this system exhibits some characteristics similar to other sinters and fractal structures, though care should be taken in comparing this mesoglass to atomic glasses. A clear departure from the simple Debye model is seen in both the effective medium and strong scattering regimes, and additional insight is needed in order to explain the behaviour of the density of states.

The behaviour of the density of states in the frequency regime associated with the fractal length scale may be part of a consistent picture for wave transport in a fractal. The scattering strength (kl_s, Figure 1.1), diffusion coefficient (D_B, Figure 1.2) and density of states (Figure 7.6) all are nearly independent of frequency, and this may reflect the self-similar nature of the fractal structure. However, similar behaviour has been found in an aluminum foam, which is not believed to have fractal structure. Both of these observations suggested the need for a more detailed, fundamental and unifying microscopic model for the density of states in strongly scattering, amorphous systems.

Perhaps most importantly, this work has established a robust methodology for directly measuring the density of states of elastic waves in a strong scattering system. The method of sample construction presented illustrates a practical way to construct a highly porous random network with percolation and fractal characteristics. X-ray tomography has also proved to be a useful tool in probing the structure of these materials, both qualitatively and quantitatively, and thus identifying the important length scales in the medium. Through the careful experiments described in Chapter 5 and the new statistical method presented in Chapter 6, the density of states can be measured in a reliable way, even when some of the vibrational modes may not be observed. To my knowledge, no other experiments to date have *directly* measured the density of states in an elastic network with the same degree of rigour, and the hope is that these experiments will provide a good foundation for further research on this important topic.

Works Cited

Aharony, A., S. Alexander, O. Entin-Wohlman, and R. Orbach. "Scattering of fractons, the Ioffe-Regel criterion, and the (4/3) conjecture." *Physical Review Letters* **58**, 132 (1987).

Akkermans, E., and G. Montambaux. *Mesoscopic Physics of Electrons and Photons.* (Cambridge University Press, 2007).

Alexander, S., and R. Orbach. "Density of states on fractals: 'fractons'." *Journal de Physique Lettres* **43**, 625 (1982).

Allen, P., J. Feldman, J. Fabian, and F. Wooten. "Diffusons, locons and propagons: character of atomic vibrations in amorphous Si." *Philosophical Magazine Part B* **79**, 1715 (1999).

Bevington, P. R., and D. K. Robinson. *Data Reduction and Error Analysis for the Physical Sciences.* 2nd ed. (McGraw-Hill, 1992).

Born, M., and K. Huang. *Dynamical Theory of Crystal Lattices.* (Clarendon Press, 1954).

Dupuis, M., R. Mazo, and L. Onsager. "Surface specific heat of an isotropic solid at low temperatures." *The Journal of Chemical Physics* **33**, 1452 (1960).

Harris, F. J. "On the use of windows for harmonic analysis with the discrete Fourier transform." *Proceedings of the IEEE* **66**, 51 (1978).

Hausdorff, F. "Dimension undausseres mass." *Mathematische Annalen* **79**, 157 (1919).

Herman, G. T. *Image Reconstruction from Projections: The Fundamentals of Computerized Tomography.* (Academic Press, 1980).

Hu, H., A. Strybulevych, J. Page, S. Skipetrov, and B. van Tiggelen. "Localization of ultrasound in a three-dimensional elastic network." *Nature Physics* **4**, 945 (2008).

Ioffe, A., and A. Regel. "Non-crystalline, amorphous and liquid electronic semiconductors." *Progress in Semiconductors* **4**, 237 (1960).

Kang, S.-J. L. *Sintering: Densification, Grain Growth, and Microstructure.* (Elsevier Butterworth-Heinemann, 2005).

Kantor, Y., and I. Webman. "Elastic properties of random percolating systems." *Physical Review Letters* **52**, 1891 (1984).

Lagendijk, A., B. van Tiggelen, and D. S. Wiersma. "Fifty years of Anderson localization." *Physics Today* **62**, 24 (2009).

Lobkis, O., and R. Weaver. "Mode counts in an aluminum foam." *The Journal of the Acoustical Society of America* **109**, 2636 (2001).

Maliepaard, M. C., J. H. Page, J. P. Harrison, and R. J. Stubbs. "Ultrasonic study of the vibrational modes of sintered metal powders." *Physical Review B* **32**, 6261 (1985).

Mandelbrot, B. B. *The Fractal Geometry of Nature*. (W.H. Freeman, 1982).

Mehta, M. *Random Matrices*. (Elsevier, 2004).

Nakayama, T., K. Yakubo, and R. Orbach. "Dynamical properties of fractal networks: scaling, numerical simulations, and physical realizations." *Reviews of Modern Physics* **66**, 381 (1994).

Pachet, N. G. *Ultrasonic velocity and attenuation in a glass-bead percolation system*. Thesis (M.Sc.), University of Manitoba (1990).

Page, J., P. Sheng, H. Schriemer, I. Jones, X. Jing, and D. Weitz. "Group velocity in strongly scattering media." *Science* **271**, 634 (1996).

Page, J. H., W. K. Hildebrand, J. Beck, R. Holmes, and J. Bobowski. "Phonons in porous media at intermediate frequencies." *Physica Status Solidi (c)* **1**, 2925 (2004).

Panametrics. "Ultrasonic transducers." (Olympus NDT, 2006). http://www.olympus-ims.com/data/File/panametrics/panametrics-UT.en.pdf

Rutherford, A., J. Harrison, and M. Stott. "Heat transfer between liquid 3 He and sintered metal heat exchangers." *Journal of Low Temperature Physics* **55**, 157 (1984).

Sade, M., T. Kalisky, S. Havlin, and R. Berkovits. "Localization transition on complex networks via spectral statistics." *Physical Review E* **72**, 066123 (2005).

Schriemer, H., M. Cowan, J. Page, P. Sheng, Z. Liu, and D. Weitz. "Energy velocity of diffusing waves in strongly scattering media." *Physical Review Letters* **79**, 3166 (1997).

Schriemer, H. P., N. G. Pachet, and J. H. Page. "Ultrasonic investigation of the vibrational modes of a sintered glass-bead percolation system." *Waves in Random Media* **6**, 361 (1996).

Sebbah, P., O. Legrand, B. Van Tiggelen, and A. Genack. "Statistics of the cumulative phase of microwave radiation in random media." *Physical Review E* **56**, 3619 (1997).

Shelby, J. E. *Introduction to Glass Science and Technology*. (Royal Society of Chemistry, 2005).

Sheng, P., ed. *Scattering and Localization of Classical Waves in Random Media*. (World Scientific, 1990).

Sheng, P., and M. Zhou. "Heat conductivity of amorphous solids: simulation results on model structures." *Science* **253**, 539 (1991).

Sheng, P., M. Zhou, and Z. Zhang. "Phonon transport in strong-scattering media." *Physical Review Letters* **72**, 234 (1994).

Sheng, P. *Introduction to Wave Scattering, Localization, and Mesoscopic Phenomena*. (Academic Press, 1995).

Stauffer, D. *Introduction to Percolation Theory*. (Taylor & Francis, 1985).

Torquato, S., T. M. Truskett, and P. G. Debenedetti. "Is random close packing of spheres well defined?" *Physical Review Letters* **84**, 2064 (2000).

Trégourès, N. P., and B. A. van Tiggelen. "Generalized diffusion equation for multiple scattered elastic waves." *Waves in Random and Complex Media* **12**, 21 (2002).

Van Tiggelen, B., and S. Skipetrov, eds. *Wave Scattering in Complex Media: From Theory to Applications*. (Springer, 2003).

Weaver, R. "On diffuse waves in solid media." *The Journal of the Acoustical Society of America* **71**, 1608 (1982).

Weaver, R. L. "Spectral statistics in elastodynamics." *The Journal of the Acoustical Society of America* **85**, 1005 (1989).

Webman, I., and G. Grest. "Dynamical behavior of fractal structures." *Physical Review B* **31**, 1689 (1985).

Appendix A: Samples Used

The table below lists the samples used for the experiments described in this thesis. Each sample was given a two-letter name for easy identification. Sample volumes were calculated from the measured mass of each sample and the measured density of a large piece of the sintered material. The frequency range(s) for which each sample was used is also included for convenience. An asterisk next to the sample name in the table below indicates that the sample was omitted from the analysis, as per Section 7.4.

Name	Volume (mm³)	Uncertainty (mm³)	Frequency Range (f_c) (MHz)	Number of Trials
AD*	3.25	0.12	5.0	8
AE*	4.87	0.11	3.5	10
AF	3.24	0.08	2.25	9
AG	3.54	0.09	0.5	8
AJ*	3.20	0.08	5.0	5
AM*	3.48	0.09	3.5	7
AP	2.05	0.06	2.25, 5.0	10, 8
AU	3.11	0.08	1.0	8
AX	5.81	0.13	(tomography only)	–
BD	0.73	0.05	0.5	5
BH*	4.78	0.11	5.0	9
BK	1.86	0.06	0.5	7
BL	0.66	0.05	5.0	9
BM	2.70	0.07	2.25	9
BP	2.65	0.11	1.0	8
BR	1.53	0.06	(tomography only)	–

CA	1.83	0.06	5.0	11
CF*	3.13	0.12	5.0	6
CH	3.13	0.12	1.0	6
CL	1.67	0.06	3.5	7
CP	1.08	0.10	1.0	6
CQ	1.67	0.06	0.5	7
DA	1.08	0.10	1.0	14
DB	0.72	0.10	1.0	7
DD	0.78	0.10	1.0	13
DE	0.82	0.05	1.0	9
DG	0.86	0.05	2.25	14
DH	1.05	0.05	2.25	10
DK	0.82	0.05	2.25	9
DM	0.66	0.05	1.0	4
DS	0.59	0.05	2.25	7
DW	0.93	0.05	0.5	8
DX	1.04	0.06	3.5, 5.0	9, 10
DY	0.36	0.05	(tomography only)	–
EB	5.12	0.11	0.5	5
EI	5.33	0.12	0.5	8

Table A.1: List of Samples Used for Experiments Described in This Thesis

Appendix B: Table of Variables

The following table of variables is intended to be a quick reference guide for the reader. Frequently occurring variables and subscripts, those which show up in more than one section, and those of importance are listed, in approximate alphabetical order.

Symbol	Description
a	average glass bead radius (≈ 64 µm)
$bulk$	subscript – referring to bulk medium (glass) properties
d	Euclidian dimension
D	fractal dimension
D	subscript – referring to the Debye model
D_B	diffusion coefficient (Boltzmann)
$\bar{d}$	fracton dimension
$\mathcal{D}$	density of states
Δ_x	uncertainty in x
E	subscript – referring to energy velocity
eff	subscript – referring to effective medium properties
ε	box size (for fractal analysis)
F	subscript – referring to fractal
f_c	transducer central frequency
ϕ	volume fraction
g	subscript – referring to group velocity
I	subscript – referring to the intermediate frequency regime
k	wave vector ($2\pi/\lambda$)
l	mean free path (l^*: transport m.f.p.; l_s: scattering m.f.p.)
l	subscript – referring to longitudinal polarization
λ	wavelength

m	mass
μ	mean number of modes counted
n	number of modes that exist in a small frequency range for one sample
N	number of measurements performed
N_P	number of particles in the system
$N(\varepsilon)$	number of occupied boxes (for fractal analysis)
ν	"actual" number of modes (when zero are counted)
p	probability of observing an existing mode
p	subscript – referring to phase velocity
q	percolation probability
q_c	percolation threshold
ρ	sample density
σ	standard deviation of number of modes counted
t	subscript – referring to transverse polarization
v	velocity
V	volume
ξ	percolation correlation length
ξ	subscript – referring to effective medium (low-frequency) crossover
Ξ	bulk crossover length
Ξ	subscript – referring to bulk (high-frequency) crossover

Table B.1: Common Variables and Subscripts Used in This Thesis