Resonant ultrasound spectroscopy: Theory and application

Brian Zadler[1], Jérôme H.L. Le Rousseau[1], John A. Scales[1] & Martin L. Smith[2]

[1] *Physical Acoustics Laboratory and Center for Wave Phenomena, Department of Geophysics, Colorado School of Mines, Golden, Colorado*

[2] *New England Research, White River Junction, Vermont*

ABSTRACT

Resonant Ultrasound Spectroscopy (RUS) uses normal modes of elastic bodies to infer material properties such as elastic moduli and Q. In principle, the complete elastic tensor can be inferred from a single measurement. For centimeter-sized samples RUS fills an experimental gap between low-frequency stress-strain methods (quasi-static up to a few kHz) and ultrasonic time-delay methods (hundreds of kHz to GHz). Synchronous detection methods are used to measure the resonance spectra of homogeneous rock samples. These spectra are then fit interactively with a model to extract the normal-mode frequencies and Q factors. Inversion is performed by fitting the the normal-mode frequencies. We have developed a forward model to compute resonant frequencies and displacements, based on an approximation that reduces the calculation to a generalized eigenvalue equation that can be solved numerically. The results of this calculation, normal mode frequencies and displacements, depend both on properties of the medium and its geometry. For the inversion we start with a high symmetry model (isotropic) and successively lower the symmetry of the model until a reasonable fit with the data is achieved. This ensures that inferred anisotropic parameters are statistically significant.

1 INTRODUCTION

1.1 Laboratory Measurements of Elastic Properties

All mechanical methods for measuring the elastic properties of laboratory specimens are divided into three types: quasi-static, resonance, and time-of-flight (if we are willing to accept modest intellectual casualties). Those types reflect the ratio of the wavelength of the mechanical signals used to the size of the specimen under test.

Quasi-static methods, such as cyclic loading (@warning Citation 'batzle' on page 1 undefined), subject the sample to deformations that are very slow compared to any of its natural mechanical resonances, so that the sample is close to mechanical equilibrium at all times during the test. By measuring both the applied stresses as well as the strains induced we hope to infer the sample's elastic compliances; if we think of stress, strain, and compliance as complex functions of frequency, we can see that the phase of the compliance (which follows from the phase shift between stress and strain) tells us about sample anelasticity. Cyclic loading measurements can be made from frequencies of the order of 10^2 Hz (above which the typical apparatus tends to become animated) down to zero frequency and have the geophysically-attractive property that they can be made at frequencies which overlap those of exploration (at least) seismology. The most serious limitations of this technique are the difficulty of accurately accounting for the complex mechanical behavior of the measurement apparatus and of making accurate, low-noise strain measurements at low frequencies.

Resonance techniques measure the frequencies of the specimen's elastic resonances, or free oscillations. These frequencies reflect the size, shape, and elastic composition of the sample; each corresponds to a particular bundle of bouncing, interconverting traveling waves which conspire to exactly repeat at intervals of $1/f$, where f is the resonance frequency. Given a sufficient

dataset of observed resonance frequencies we can make useful inferences about the sample's properties. The experimental apparatus is simple and undemanding and easily adaptable to a wide range of sample sizes and shapes. For typical specimen sizes the frequency of measurement ranges from about 5×10^3 Hz on up, appreciably higher than seismic frequencies but still within the range of borehole acoustics. Interpreting the observed frequencies in terms of sample composition is complex but can be automated and is well within the reach of small, consumer computers. It is possible in principal to measure arbitrary anisotropy and it is often reasonable in practice to do so. Under favorable conditions this approach can provide very good attenuation estimates. On the other hand the method loses quite a bit of its power when the sample is strongly coupled to its surroundings, as is typically the case for measurements under stress conditions at depth. Because the raw observations are curves of signal amplitude *versus* frequency with possibly complex features and because the intermediate observables, the resonance frequencies, are themselves not coupled to sample elastic properties in a simple way, this technique offers the greatest potential for confusion and delusional misinterpretations.

Time-of-flight methods measure the transit time of various elastic wave types across the specimen and sometimes alterations in the wave's shape due to attenuation and other phenomena inherent to the sample. These measurements assume wavelengths which are short compared to sample dimensions. For typical specimen lengths of the order of a few inches that implies frequencies above about 5×10^5 Hz and is this method's principal shortcoming from a geophysical point of view. It has a number of appealing strengths: We can usually arrange the measurement geometry so that it is little affected by conditions at the sample surface; this isolation is very important for accommodating imposed conditions of pressure and saturation. The observed signals are readily analyzed and the physics involved is comfortingly easy to visualize.

1.2 Resurgent Resonance

Resonance techniques have been used to measure elastic properties in the laboratory since at least the 1920s (McSkimin, 1964). They were confined to samples with rather unusual geometries, such as bars with length to diameter ratios of ten or more, because those were the only geometries for which it was practical to solve the associated forward resonance problem. Later researchers, starting with Francis Birch's pre-World War II unpublished measurements, exploited spherical samples of isotropic materials. The forward problem for this case is computationally harder than for thin rods but generally within the reach of tables of special functions and determination. Successful numerical use of variational approaches to the forward problem by Holland (1967) and Demarest (1969) were not enough to break the field free of its sample geometry limitations, probably because computational resources were so limited and the bases used in these earliest efforts were essentially useful only for rectangular geometries. Migliori and Sarrao (1997) review this process and observe that the liberating step was the discovery by Visscher et al (1991) that a very simple basis for the variational calculation led to an accurate and extremely flexible approach to computing the elastic resonances of anisotropic elastic bodies of arbitrary shape.

The ability of this approach to handle a wide variety of geometries and symmetries in a unified way has added significantly to the appeal of elastic resonance methods. It has also led to some extent to a tendency to regard the individual resonances of a particular geometry as indistinguishable and equally useful bits of incremental information. Combined with advances in computing speed, we are now able to expend more effort in pursuit of numerical results. We also have access to a large variety of algorithms and languages that we can exploit with only modest effort.

1.3 Issues with Geological Materials

Effective resonance measurements require

(i) accurate measurements of the sample's resonances,

(ii) a good numerical model of the resonator that ties the object's elastic properties to its resonances, and

(iii) a way to tell the *predicted* resonance to which an *observed* resonance corresponds.

Although *(i)* may require considerable ingenuity in a hostile environment or other challenging circumstances , when we deal with geological materials we usually find that *(ii)* and *(iii)* pose the greater challenge.

Common resonance measurement practice is aimed at specimens that are homogeneous in both composition and crystallographic symmetry and that have high to very-high Q's . Homogeneity means that we can accurately model the measurement with a small number of spatially-constant elastic parameters. Low attenuation often implies that we can observe every, or nearly every, resonance the sample is capable of and that we can easily map them onto the set of predicted resonances.

Geological materials (rocks) are frequently spatially inhomogeneous at scales comparable to and smaller than the centimeter scale of laboratory specimens. It's not unusual to find both gradations and sharp discontinuities in specimens and it is not uncommon that these variations are difficult to discern visually. The specimens themselves have usually been extracted from a larger body, such as the Earth, through mechanically aggressive means and suffer various degrees of peripheral damage: after coring, for example, a sample may microfracture as a result of stress relaxation.

First, we deal with sample sizes on the order of a few centimeters, not millimeters. We choose this scale because it is a common sample size for low frequency *quasi-static* measurements. Second, rocks tend to have unknown symmetries which are usually assumed known *a priori*. For example, during coring procedures, the rock matrix will undergo relaxation causing aligned microfractures resulting in anisotropy. Heterogeneity is present in all rocks. Since our acoustic spectroscopy techniques assume homogeneous media, this will introduce some uncertainty.

The biggest deterrent from using RUS on rocks is the low quality factors. Observation of normal modes in a dry, competent rock such as Elberton granite is simpler since the Q of any given mode will likely be on the order of several hundred. However, in a softer and more porous rock like Berea sandstone, the observed modes often have Q values of 100 or less, making it much more difficult to identify them without first using a forward modeling code. Drying techniques can bring this number up into the hundreds, but measurements of dry rocks may not be of great interest except as starting points for confining pressure or saturation experiments.

2 PLAN OF THIS PAPER

The goal of this paper is to introduce the theory of *resonant ultrasound spectroscopy* (RUS) and its application to the characterization of the elastic moduli of homogeneous rock samples. The theory we outline is nearly identical to that found in the book of Migliori and Sarrao (@warning Citation 'mi97' on page 3 undefined). We also step through the inversion procedure, as well as the process of extracting frequencies and Q values from resonance spectra.

Upon excitation, elastic bodies isolated in a free space oscillate at discrete frequencies. These normal modes can be observed and measured. The geometry and size of the system strongly influence the range of frequencies at which such normal modes occur. The Earth, for example, has a fundamental mode with a period of about 36 minutes. For a 10 centimeter piece of sandstone, the fundamental mode has a frequency of about 10 kHz. The elastic moduli also influence the resonance frequencies: the stiffer the sample, the higher the frequencies. The idea behind resonant ultrasound spectroscopy (RUS) is to use these dependencies to infer elastic properties or shape parameters of samples from a suite of measured resonance frequencies. In order to perform an inversion we must have a way of predicting these frequencies for an arbitrary elastic body.

The forward modeling allows the design of an inversion procedure using predicted and measured normal mode frequencies. We seek a model that predicts the data in a least-squares sense. We design an objective function and minimize it using non-linear optimization methods. The inversion usually starts using an isotropic model. If a proper fit cannot be achieved we lower the symmetry of the model used. Using a model with too many free parameters we may however end up fitting noise rather than just fitting the data. χ^2 tests give us the statistical means to decide on the proper anisotropic model to use.

We start with a nontechnical discussion of the elastic modes of homogeneous cylinders. Our algorithms allow us to apply RUS to a variety of regular geometries, but the cylinder is a useful specific case. Then we show how the calculation of the normal modes of an elastic body can be cast as a generalized eigenvalue problem via a Rayleigh-Ritz method. We solve this eigenvalue problem using standard numerical tools. The inverse problem is then formulated for the components of the elastic tensor as a least-squares fitting of the predicted to measured normal-mode frequencies. Next we discuss the problem of measuring the resonance spectrum of a sample and extracting the frequencies from this spectrum. Finally we show application of the RUS method to rock samples, both isotropic and anisotropic, at room conditions and under gas confining pressure. Software implementing all aspects of the calculation (spectra fitting and least squares inversion) is freely available from our web site, listed at the end.

3 ELASTIC MODES OF A CYLINDRICAL SAMPLE

In order to get a sense of the basic scheme of the mode world, consider the elastic resonances of a homogeneous right-circular cylindrical sample of an isotropic material. To be specific, we assume values for density and elastic properties that are about that for shale and refer to the material as *generic soft rock*. Specifically, its properties are

height	3.099	cm
radius	0.635	cm
density	2.5	gm/cm^3
V_P	3.000	km/sec
V_S	1.400	km/sec

3.1 General Categories of Modes

Our cylindrical sample has perfect axisymmetry. Because of that symmetry, we can show (but we won't do it here) the particle displacement for any mode must vary with the cylindrical coordinate θ as $\sin(n\theta)$ for some integer value of n. That fact plus some additional symmetry arguments can be used to show that every mode of our sample must fall into one of three classes:

torsional axisymmetric pure shear motions consisting of rigid rotations of rings of material around the axis. The frequencies of these modes depends wholly upon the sample's shear velocity.

extensional axisymmetric mixtures of compression and shear motions. At low frequencies, when wavelengths are long compared to the sample's diameter, these modes are essentially an axial compression coupled to a radial expansion, their frequencies are then *almost* wholly dependent on the sample's bar velocity, $V_E = \sqrt{\frac{E}{\rho}}$ and the modes are called *bar* modes. At higher frequencies the modes become more complex and it would be more appropriate to call them simply *non-torsional*.

non-axisymmetric *or* flexural all modes for which $n \neq 0$ in $\sin(n\theta)$ (above) are in this class. The axisymmetric modes represent energy traveling purely up and down the cylinder axis; the non-axisymmetric modes travel along paths that are tilted with respect to that axis.

Although these classes break down when axisymmetry is destroyed (as is almost always the case in a real experiment), they are still very useful in "nearly-symmetric" circumstances, and even in "not-very-symmetric" ones.

The flexural modes occur in pairs, called *doublets*, both members of which have the same resonance frequency. This obscure distinction is important when the modes are perturbed by small, non-symmetric changes in the sample's composition or shape: perhaps a small crack develops in one side of a specimen. In that circumstance the resonance frequencies are slightly shifted by the perturbation. For an axisymmetric mode this shift simply moves the modes peak a little bit in frequency. The two members of a doublet, however, may experience different shifts and what appeared to be a single peak in the unperturbed sample may become two distinct peaks in the perturbed sample.

3.2 Mode Shapes

Before we get into the exact mechanics of calculating the displacement field for a generic object, it is beneficial to picture the possible deformations for a common right cylinder. Figure 1 shows the particle motion of the sample surface for several types of modes at two extreme points of the cycle. (The deformation is greatly exaggerated: under normal circumstances in the laboratory, particle motions are of the same order as atomic diameters.) The top row of figures shows the mode shapes for the two lowest extensional modes. The motion in both cases is clearly axisymmetric, and has both axial and radial components. The bottom row of figures shows the two lowest torsional modes which are also axisymmetric, but here the particle motion is entirely azimuthal, that is, in the local direction of the θ coordinate. The middle figures show the shape of the first flexural mode, which is also the lowest-frequency mode of the system. The mode's "bending" shape is clearly non-axisymmetric.

4 DISPLACEMENT CHARACTERIZATION

Next we turn to the problem of calculating the resonances for an arbitrary elastic body. More rigorous studies of body resonances have been performed in the past (@warning Citation 'strutt' on page 4 undefined; @warning Citation 'tromp' on page 4 undefined; @warning Citation 'vi91' on page 4 undefined; @warning Citation 'mi97' on page 4 undefined). Consider an isolated purely elastic body (no attenuation) with a volume V. Its kinetic energy E_k is (@warning Citation 'ak80' on page 4 undefined; @warning Citation 'mi97' on page 4 undefined)

$$E_k = \int_V \tfrac{1}{2}\, \rho\, \partial_t u_i\, \partial_t u_i\, dV\ ,$$

where u_i, $1 \leq i \leq 3$, is the displacement field in a Lagrangian system of coordinates for R^3 and ρ is the density. The summation convention for repeated indices is applied. The potential energy E_p is given by the strain energy (@warning Citation 'ak80' on page 4 undefined)

$$E_p = \int_V \tfrac{1}{2}\, C_{ijkl}\, \partial_j u_i\, \partial_l u_k\, dV\ ,$$

where (C_{ijkl}) is the elastic stiffness tensor.

The Lagrangian L of the system can then be written as

$$\begin{aligned} L = E_k - E_p \quad &= \quad \int_V \tfrac{1}{2}\, \rho\, \partial_t u_i\, \partial_t u_i\, dV \\ &- \quad \int_V \tfrac{1}{2}\, C_{ijkl}\, \partial_j u_i\, \partial_l u_k\, dV\ . \end{aligned}$$

We want to characterize displacement fields that naturally occur in terms of normal modes in the absence of driving force. Such displacement fields may have been initiated by a driving force but we assume this force is not present anymore and the transient response is negligible. Applying Hamilton's principle we get two equations

$$\rho\, \ddot{u}_i - \partial_j(C_{ijkl}\, \partial_l u_k) = 0\ , \quad 1 \leq i \leq 3\ ,$$

which is the elastic wave equation (with no right-hand side) and

$$(C_{ijkl}\, \partial_l u_k\, n_j)\, |_S = 0\ , \quad 1 \leq i \leq 3\ , \tag{1}$$

which is the boundary condition at the free surface S.

In the frequency domain, the homogeneous elastic wave equation becomes

$$\rho\, \omega^2 u_i + \partial_j(C_{ijkl}\, \partial_l u_k) = 0\ , \quad 1 \leq i \leq 3\ . \tag{2}$$

The system of equations (2) and (1) has solutions only for a set of frequencies corresponding to the resonant frequencies of our sample. As we shall see later, they

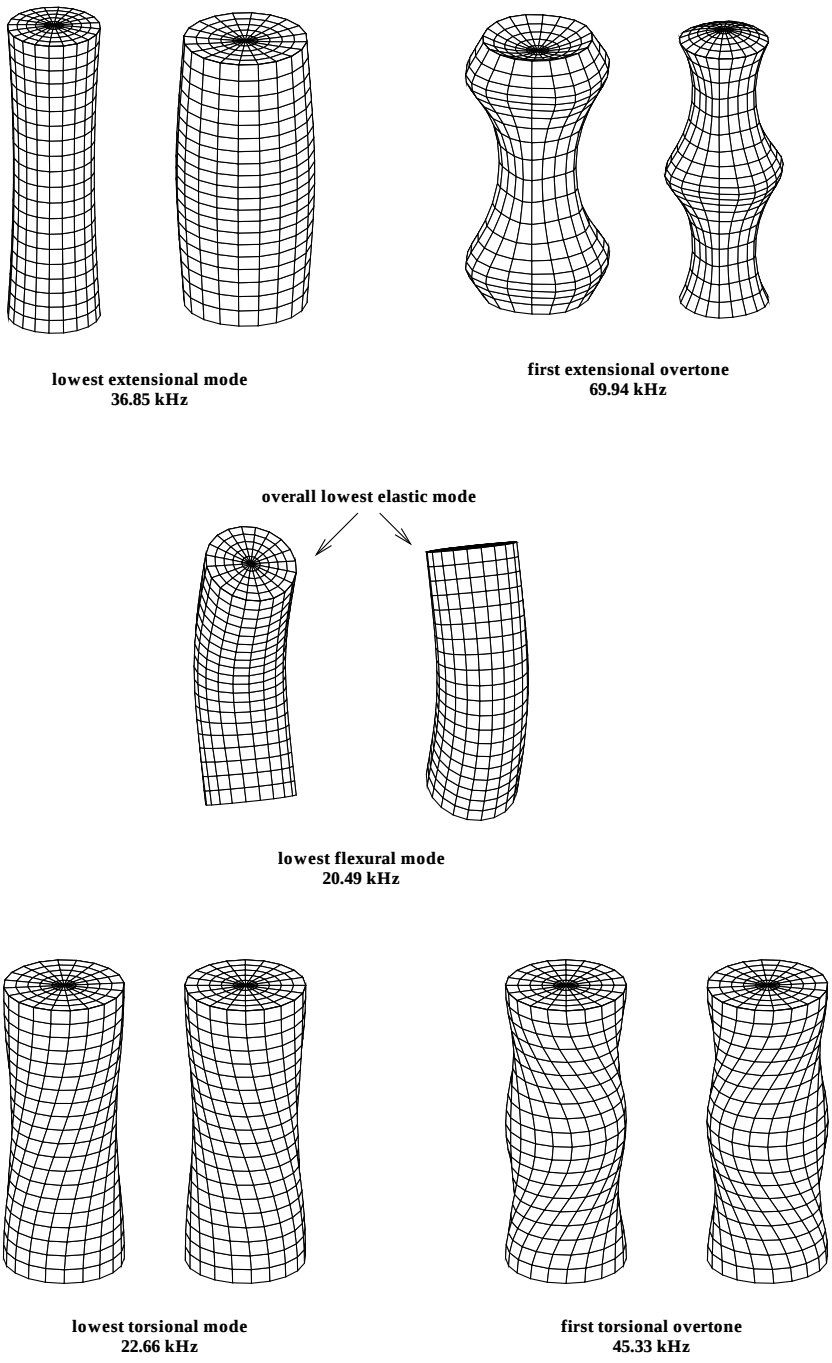

Figure 1. Surface particle displacements of several characteristic modes of the generic soft rock model

depend both on the shape and the elastic properties of our body.

If we now define a new function T as

$$T = \int_V \left\{ \tfrac{1}{2}\, \rho\, \omega^2 u_i u_i - \tfrac{1}{2}\, C_{ijkl}\, \partial_j u_i\, \partial_l u_k \right\} dV \ ,$$

we then have (@warning Citation 'mi97' on page 5 undefined)

$$\begin{aligned} \delta T &= \frac{\partial T}{\partial \alpha} \delta\alpha \\ &= \int_V \left\{ \rho\, \omega^2 u_i\, \delta u_i \ - C_{ijkl}\, \partial_l u_k\, \partial_j \delta u_i \right\} dV = 0 \end{aligned} \quad (3)$$

If we choose a basis we can decompose any component u_i, $1 \le i \le 3$, of our displacement field into this ba-

sis. With such a decomposition equation (3) can then be viewed as a generalized eigenvalue problem (see also next section). Solving this eigenvalue problem, we find a set of frequencies (eigenvalues) for which the system has a non-zero solution. These are the normal modes. The displacement field can then be decomposed in a Fourier series scheme into the basis of the normal modes.

5 NUMERICAL APPROACH OF THE GENERALIZED EIGENPROBLEM

To perform numerical approximations of the normal modes we follow closely Migliori and Sarrao (@warning Citation 'mi97' on page 6 undefined). We restrict ourselves to polynomials of degree less than or equal to N, i.e., the subspace generated by the monomials $\phi_\lambda = x^{\mathsf{l}}\, y^{\mathsf{m}}\, z^{\mathsf{n}}$ for which

$$\mathsf{l} + \mathsf{m} + \mathsf{n} < N \ .$$

We can then write for $1 \leq i \leq 3$ and λ finite

$$u_i = \sum_{\lambda \in \Lambda} a_{i,\lambda} \phi_\lambda \ ,$$

with $\phi_\lambda = x^{\mathsf{l}}\, y^{\mathsf{m}}\, z^{\mathsf{n}}$.

We expect to obtain an approximation good enough for our purpose with N large enough (see Sec. 10). This leads to $\binom{N+3}{3} = (N+1)(N+2)(N+3)/6$ choices of l, m and n for each component of u. Hence, the dimension of the problem is now $R = 3 \times (N+1)(N+2)(N+3)/6$. Having restricted ourselves to this finite dimensional subspace, we can write T in the approximate form

$$\begin{aligned} T &= \int_V \tfrac{1}{2}\, \rho\, \omega^2\, a_{i,\lambda}\, a_{i,\lambda'}\, \phi_\lambda\, \phi_{\lambda'} dV \\ &- \int_V \tfrac{1}{2}\, C_{ijkl}\, a_{i,\lambda}\, a_{k,\lambda'}\, \partial_j \phi_\lambda\, \partial_l \phi_{\lambda'}\, dV \ , \end{aligned}$$

which can be written in a matrix form

$$T = \tfrac{1}{2}\, \omega^2\, a^t\, E\, a - \tfrac{1}{2}\, a^t\, \Gamma\, a \ .$$

Both matrices E and Γ are of dimension R. They have the forms

$$E_{\lambda i \lambda' k} = \delta_{ik}\, \rho \int_V \phi_\lambda\, \phi'_\lambda\, dV \ , \tag{4}$$

and

$$\Gamma_{\lambda i \lambda' k} = C_{ijkl} \int_V \partial_j \phi_\lambda\, \partial_l \phi'_\lambda\, dV \ , \tag{5}$$

where we have assumed that (C_{ijkl}) and ρ are constant. This is the major assumption of this section.

To search for an extremum of T, i.e., an approximate normal mode, we set the differential of T, evaluated at a, to zero

$$d_a T(a) = \omega^2 E a - \Gamma a = 0 \ , \tag{6}$$

which is a generalized eigenvalue problem. The matrices E and Γ have remarkable properties that we can exploit: they are real symmetric. Furthermore, E is positive definite. System (6) can be furthermore simplified for its numerical evaluation. In the orthorhombic and higher symmetry case, we can order matrices E and Γ in such a way to make them block diagonal. There are then 8 diagonal blocks. Instead of (numerically) solving for a large eigenvalue problem, we now solve for 8 smaller ones. See A for further details.

To solve this generalized eigenvalue system we use standard computer libraries such as *Lapack* (@warning Citation 'laug' on page 6 undefined). The eigenvalues and eigenvectors are both dependent on the elastic properties and the geometry of the specified body according to equation (4) and equation (5). The geometry dependency is found in the general value of the following type of integrals:

$$\mathcal{I}_V = \int_V x^{\mathsf{l}} y^{\mathsf{m}} z^{\mathsf{n}} dV \ . \tag{7}$$

If we can compute integral (7) we can compute both E and Γ and can numerically solve for the (approximate) normal modes. Here we consider three different shapes: rectangular parallelepipeds, cylinders, spheroids. The dimensions of those shapes are described in Figure A1. In all those cases, we use the center of symmetry of the shape as the coordinate origin.

The calculations of the integral (7) for spheres, cylinders and rectangular parallelepipeds are carried out in Appendix B.

6 THE INVERSE PROBLEM

So far, we have dealt with a forward modeling problem, i.e., knowing the characteristics of an elastic body we compute its normal modes. We would like to solve the inverse problem: starting from the measured eigenfrequencies, can we infer the elastic parameters? Here, our model could have different variables: we can invert for the dimensions of the body (we still have to assume a type of shape though). We can also invert for the elastic moduli or the crystallographic axis. We shall now focus on the latter topic.

6.1 Objective function and non-linear optimization algorithms

To estimate how good a given model is, i.e., how well it can predict the data, we make use of an the following objective function

$$F = \sum_i w_i (f_i^{(p)} - f_i^{(m)})^2 \ , \tag{8}$$

where $f^{(p)}$ are the computed frequencies and $f^{(m)}$ are the measured ones. We minimize the difference between

predicted and measured frequencies in a least square sense. The weights w_i characterize the confidence we have in the measurements. We minimize the objective function, up to some tolerance, over the space of all feasible models. In a neighborhood of the an extremum this is a quadratic minimization problem. However, given the presence of noise we say a model fits the data if F is less than some tolerance, which can be quantified via a χ^2 type criterion.

An efficient way to minimize the objective function is to apply a conjugate gradient method (@warning Citation 'fl80' on page 7 undefined). We use a non-linear form of conjugate-gradient which amounts to making repeated quadratic approximations. One advantage of such a method is that it does not require the knowledge of the Hessian of the objective function. A sequence of search directions, h_i, is constructed and line searches are performed in order to minimize the objective function along such directions. The procedure is as follows: assume you have a search direction h_i and a model m_i. Let $g_i = -\nabla F(m_i)$. By minimizing the objective function, F, along the h_i direction, i.e., minimizing $j(\upsilon) = F(m_i + \upsilon h_i)$, we find a new model $m_{i+1} = m_i + \upsilon_i h_i$. Then the new search direction is given by:

$$h_{i+1} = g_{i+1} + \gamma_i h_i \ , \tag{9}$$

with $g_{i+1} = -\nabla F(m_{i+1})$ and

$$\gamma_i = \frac{{g_{i+1}}^t g_{i+1}}{{g_i}^t g_i} \ . \tag{10}$$

The method is initiated by taking the steepest descent direction as the starting search direction, i.e., $h_0 = g_0 = -\nabla F(m_0)$. In the case of a true quadratic form such an algorithm yields convergence in a finite number of steps, in exact arithmetic. Formula (10) was derived from that particular case. For further references we refer to Fletcher (@warning Citation 'fl80' on page 7 undefined) and Press *et al.* (@warning Citation 'pr86' on page 7 undefined).

We can obtain an exact expression for the gradient of the objective function. The partial derivative of the objective function with respect to a particular parameter p is

$$\partial_p F = \sum_i w_i 2 \partial_p f_i^{(p)} (f_i^{(p)} - f_i^{(m)}) \ ,$$

and differentiating equation (6) with respect to p , composing by a^t on the left and using Γ and E symmetries yield (@warning Citation 'mi97' on page 7 undefined)

$$\partial_p(\omega^2) = a^t \left(\partial_p \Gamma - \omega^2 \partial_p E\right) a \ ,$$

where we made use of the normalization $a^t E a = 1$. Hence, knowing the eigenvector a as well as the eigenfrequency ω allows us to compute the gradient of the objective function. Knowledge of the gradient is a key parameter for convergence and efficiency. Knowledge of the Hessian would allow more efficient optimization methods, e.g., Newton methods, but would require some significant additional computational effort while computing the eigenvectors (@warning Citation 'mi97' on page 7 undefined).

To perform the non-linear conjugate gradient algorithm we use the C++ COOOL library (CWP Object Oriented Optimization library) (@warning Citation 'de96-a' on page 7 undefined; @warning Citation 'de96-b' on page 7 undefined). A code does the forward modeling and computes both the objective function values and its gradient using the measured resonant frequencies. Using this result, COOOL can perform the line search along the direction chosen and find the best step to use. Starting from this point, with the value of the new gradient a new direction is chosen as described in equation (9) and a new line search is performed.

7 EXPERIMENT DESIGN

In typical RUS measurements, we try to measure all of the resonances below some upper limit. Having a complete set of resonances assures us that we have extracted all of the available information and significantly simplifies the inverse calculation. Most important, we can confidently match each observed mode with a computed one, a step that is essential to interpreting the data. As Qs decrease, however, it becomes increasingly difficult, and eventually impossible, to determine all (or even most) resonances.

Figure 2 shows the theoretical resonance response for the generic soft rock sample (defined in Section 3) using edge-contacting pinducers, for three assumed values of compressional and shear Q. The three cases represent, by geologic standards, very high, moderate, and low (but still plausible) Q regimes. It's clear that we cannot depend upon observing a complete mode catalog in the presence of even moderate attenuation.

7.1 Sample Shape and Mode Sensitivities

Figure 3 shows the relative sensitivity of each mode to changes in compressional velocity for our generic soft rock defined in Section 3. The quantity shown is

$$R_P = \frac{\partial_a f}{\partial_a f + \partial_b f}, \tag{11}$$

where $\partial_a f$ and $\partial_b f$ are related to the sensitivities of the mode to V_P and V_S, respectively. R_P is a measure of the relative importance of the sample's compressional and shear speeds in determining the given mode's frequency. Both $\partial_a f$ and $\partial_b f$ are non-negative (see (@warning Citation 'strutt' on page 7 undefined)) and so $0 \leq R_P \leq 1$.

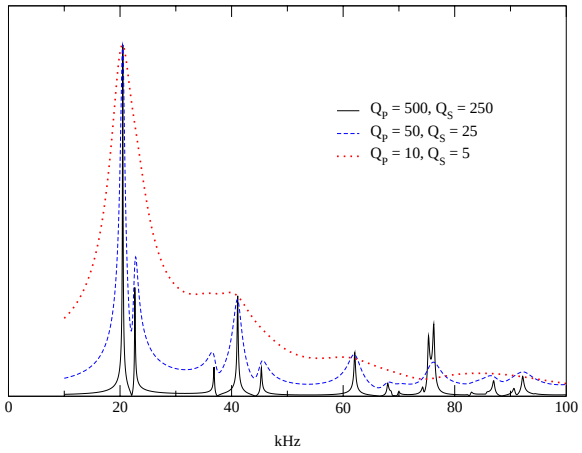

Figure 2. The theoretical resonance response of a generic soft rock sample for high, moderate, and low values of compressional and shear Q.

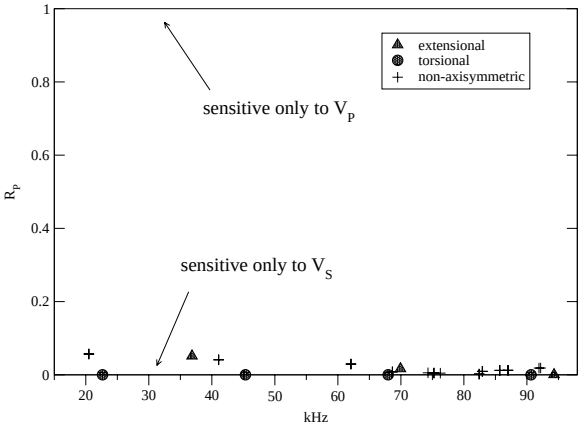

Figure 3. Relative sensitivity to compressional and shear velocity, R_P, defined in equation (11), for the lowest-frequency modes of the generic soft rock model

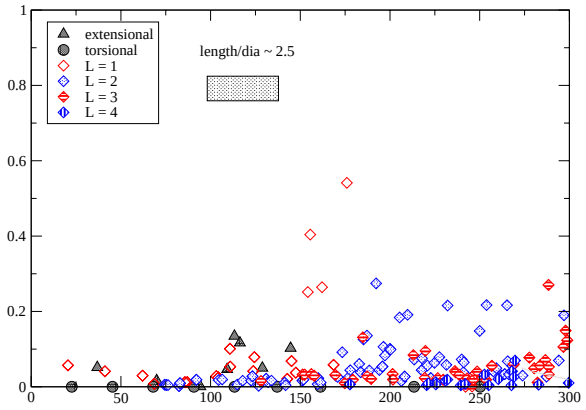

Figure 4. R_P defined in equation (11) for the generic soft rock model with an inset drawing showing the sample's proportions. The principal plot is the same as that in Figure 3 but covering a broader frequency range.

If $R_P = 0$ for some particular mode, then the resonance frequency of that mode is completely insensitive to changes in the sample's V_P. The sample motion is purely shear and the frequency of the mode is exactly proportional to the sample's shear velocity. The axisymmetric torsional modes are the only modes for which the above are exactly true, although there are many modes for which $R_P \ll 1$.

If, on the other hand, $R_P = 1$ for some mode, then the resonance frequency of that mode is completely insensitive to changes in the sample's V_S. The elastic cylinder has no modes for which this is true, although it probably becomes asymptotically true for certain families of modes as frequency becomes infinite. It turns out that we have an interest in looking for modes at reasonable frequencies for which R_P is as large as possible.

The values in figure 3 are all remarkably close to 0. None of these modes is more than slightly sensitive to V_P and we can't expect to get robust estimates of V_P. At a given measurement error level, the errors in V_P will be twenty or more times those in V_S.

Figure 4 shows R_P out to about 300 kHz. The first mode substantially sensitive to V_P is a flexural mode at about 150 kHz, for which $R_P > 0.4$. Note that if we wanted to use this mode in a measurement we would have to extend the upper frequency range of our measurements. Going to higher frequencies may bring additional problems: the number of resonances increases rapidly with frequency and mode identification problems are likely to increase, particularly in the presence of substantial attenuation.

We can do better without going to higher frequency, however, by changing the sample geometry. Figure 5 shows R_P for a model identical to the one used above *except* that it has twice the diameter. In this case the first P-sensitive mode is a flexural mode at about 80 kHz, about a factor of two lower in frequency than for the original sample geometry.

Figure 6 shows R_P for a model greater by yet another factor of two in diameter. In this case the first P-sensitive mode is at about 60 kHz. But notice that the *number* of modes in any particular frequency band is increasing rapidly as we make the sample wider. As we mentioned above, the increased crowding quickly complicates mode identification problems, particular in the presence of dissipation.

7.2 Excitation and Observation

In this section we'll look at a few numerical results that show how transducer location plays a central role in determining which resonances are observed. Figure 2 showed that moderate amounts of attenuation could

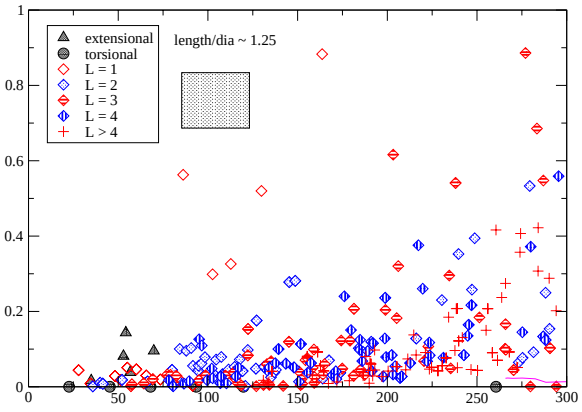

Figure 5. R_P defined in equation (11) for the generic soft rock model with an inset drawing showing the sample's proportions. Similar to Figure 4 (preceding), but computed for a model with the same length and *twice* the diameter as the generic soft rock model used for the preceding figure.

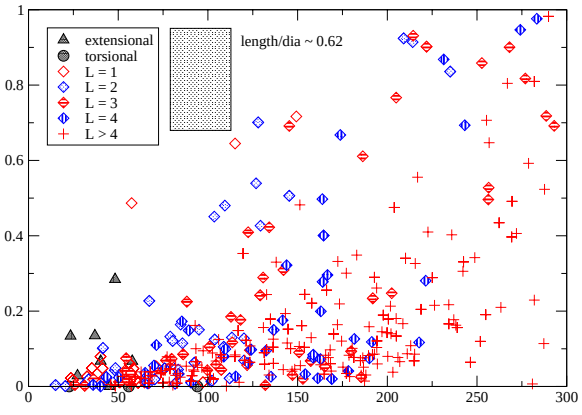

Figure 6. R_P defined in equation (11) for the generic soft rock model with an inset drawing showing the sample's proportions. Similar to Figure 4 (preceding), but computed for a model with the same length and *four* times the diameter as the generic soft rock model used in figure 4.

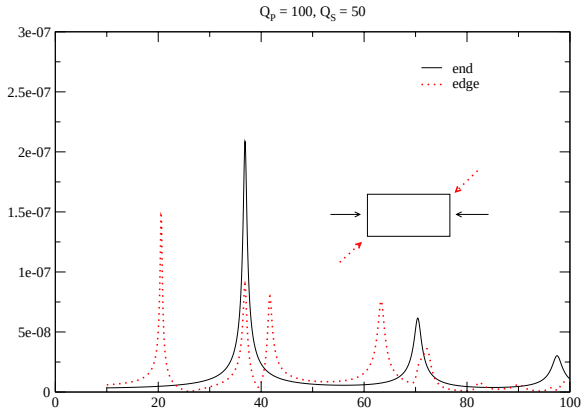

Figure 7. The theoretical resonance response of our generic soft rock sample for the edge and end mounted transducer configurations.

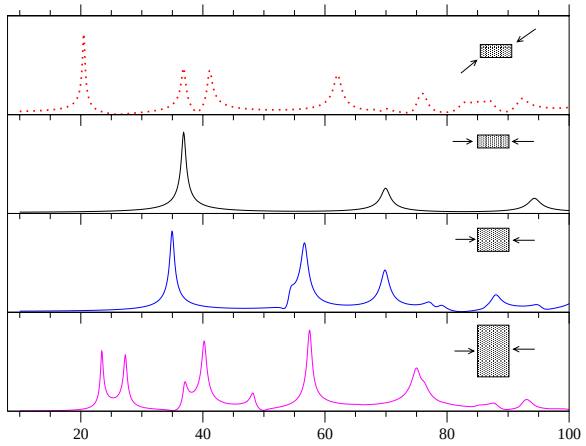

Figure 8. The theoretical resonance response of a generic soft rock sample for several transducer arrangements and sample shapes

make it impossible to discern all of the resonances in a particular frequency interval. A possible tool for solving this problem would be the use of transducer arrangements that favor particular modes.

Figure 7 shows the computed response for the generic soft rock sample (Section 3) for two transducer arrangements for moderate values of Q. The important point to note is that there are modes that appear prominently in the *edge* spectrum that are missing from the *end* spectrum. The great simplicity of the latter curve is one of the strengths of the resonant bar technique.

Figure 8 reproduces the two curves of figure 7 in the top two plots and adds two more plots showing the effects of making the sample, respectively, twice and four times greater in radius. As the sample becomes wider, the spectrum undergoes two important changes. First, bar-wave approximation breaks down, as we can see by the loss between the second and third plots of the simple regularly-spaced peaks of the narrow sample; this change is good because some of the changed mode types are more sensitive to V_P than the bar modes were. Hand-in-hand with that, however, is much greater spectrum complexity, as seen most dramatically in the bottom curve. In the presence of attenuation, spectral complexity makes it much harder to interpret the observed spectrum in terms of specific modes.

At this point, we think these results tells us that developing the ability to get V_P and V_S in a single experiment will probably involve a balance between sample aspect ratio and somewhat more sophisticated transducers. We do not yet have a global solution to the resonance experiment design problem.

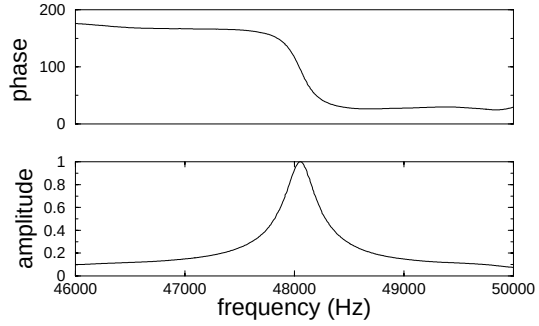

Figure 9. Amplitude and phase of an isolated mode in granite.

7.3 Data acquisition

A function generator (Stanford DS345) sends a 10 Volt (peak-peak) swept sine wave to the source transducer. This signal is detected synchronously with a DSP lock-in amplifier (Stanford SR850), which digitizes the input and reference signal with 18 bit precision, directly measuring the in-phase and quadrature components. Amplitude and phase are calculated from these. Figure 9 shows a typical amplitude/phase response for an isolated mode measured in a granite core. We see the π phase shift associated with the arctangent function as the frequency passes through the resonance.

To limit loss of energy and control the humidity in the bench top measurements we place the sample in a vacuum chamber held at approximately 160 mbar at room temperature, 23 C. In the theoretical derivations of section 4 we did not implement boundary and radiation conditions. With the use of the vacuum chamber in the experiments we can stick to those assumptions. The vacuum chamber also yields a quieter environment for the experiment. For porous media it is well known that even a single monolayer of water adsorbed onto the pore space of the sample can have a dramatic influence on the moduli.

8 RESONANT ULTRASOUND SPECTROSCOPY VS. RESONANT BAR

Before moving into the experimental observations of this paper, it will be useful to mention the differences between RUS and resonant bar techniques. The most obvious difference is sample shape. For resonant bar measurements, the sample is a "bar" with a length-to-diameter ratio of 10 or more. For RUS, the sample can be any shape that can be modeled, though the most common are spheres, cylinders and parallelepipeds. The scale of the objects can also different. The full frequency range available for resonant bar techniques is similar to RUS (5 kHz - 200 kHz), however with a sufficiently long resonant bar of rock (approximately 30 cm), normal mode frequencies in the seismic frequency range (0-3000 Hz) can be measured. A specific mode catalog is excited and studied for the sample (@warning Citation 'lucet91' on page 10 undefined) (torsional and/or extensional as described in Section 3.1). In RUS, all modes are allowed to be driven and measured for a more complete catalog which can give a measure of anisotropy. This can be both good and bad. Experimentally, it may beneficial to mount a sample on its edges to generate more modes and reduce losses due to system coupling. (For rectangular samples, mounting at a corner of three sides eliminates the possibility of accidental degeneracy.) On the other hand, *in-situ* measurements are more difficult for RUS since they involve strong coupling to a pressurized medium.

9 IDENTIFICATION AND ESTIMATION OF NORMAL MODE FREQUENCIES

9.1 Mode *identification* vs. mode *estimation*

Before delving further, it is necessary to make a distinction between mode *estimation* and mode *identification*. Mode *estimation* is the process of extracting a set of parameters $\{p_i\}$ from a given resonance peak such that the peak could be reconstructed at a later time. This is done interactively by fitting a model to a peak or set of peaks. The equation should include at least the peak frequency, Q, amplitude and some combination of background terms. In our case, the Breit-Wigner model is used (@warning Citation 'breit-wigner' on page 10 undefined). This does **not** associate the observed resonance peak with the appropriate member(s) of the eigenfrequency spectrum of the forward model. We term that mode *identification*, which is also an interactive process. Identifying modes should not be a problem for nicely homogeneous, well-made samples. However, for samples with heterogeneity or low Q, poorly cut samples or those with spectra containing a high eigenfrequency density, mode *identification* can be difficult or impossible.

9.2 Normal mode estimation

The eigenfrequency of a mode is not exactly at the maximum of the corresponding peak in the spectrum. Any one mode is assumed to have Lorentzian shape, and multiple overlapping modes are considered to be a superposition of Lorentzians. An example of this shape can be seen in Figure 9.

The model we used to fit the amplitudes of one or multiple modes is due to Briet and Wigner (@warning Citation 'breit-wigner' on page 10 undefined):

$$A(f) = B_0 + B_1(f - f_0) + \sum_n^N \frac{C_n + D_n(f - f_0)}{(f - f_n)^2 + \frac{1}{4}\Gamma_n^2} \quad (12)$$

where the displacement amplitude A is a function of frequency. The parameters B_0 and B_1 describe a constant and linear background, respectively. The shape of each mode is then constructed via 4 parameters: amplitude C_n, skewness D_n, eigenfrequency f_n and full width at half max Γ_n. Therefore, a simultaneous fit of n resonance peaks will have $4*n+2$ parameters. A non-linear least squares fit of the data is then performed. Since the fit is sensitive to the initial parameters, they need to be chosen carefully and accurately. This reduces the chance of divergence in the fit and greatly reduces the computation time.

In an ideal resonance spectrum, the observed modes would be spaced apart such that each mode appears isolated. In addition, each mode would be fit one at a time. Unfortunately, this is not always possible due to degeneracy and the overlapping modes. Every mode in the measured spectrum effects every other mode to some extent, the size of the effect being directly proportional to the frequency spacing between adjacent modes. This is mostly a problem when one peak dominates its neighboring peak in amplitude to the point where you see a small 'hill' on the side of a 'mountain'. The fitting procedure will tend to favor the larger peak, so the estimates of the parameters of the smaller peak will have a higher uncertainty.

We put a quantitative number on this effect by putting less (or zero) weight on the small peak during the inversion process if we're unsure of its existence in the forward model. One way to be fairly certain the peak is a true resonance is the following method. *If* there are a number of prominent peaks (at least 4) before the one in question, the inversion can be performed on the set of prominent peaks with the questionable peak ending the set. Giving a weight of one to all peaks except the last will result in a fit to the first 4 'high confidence' peaks, and the last 'zero-weight' peak will be effectively *guessed* at by the inversion. If the frequency from the inversion seems reasonably close to the observed (but questionable) peak, then it *may* be a true resonance. This is an iterative game that can be played under the right conditions, but to this author's knowledge, there is no fool-proof solution for this problem. However, if repeated measurments are possible, the less prominent peaks may become greatly enhanced in other data sets, giving rise to an increased confidence level in their true existence.

9.3 Normal mode identification

There are a number of reasons why a measured spectrum cannot be mapped in a reasonable fashion to a modeled spectrum. It may be that the proposed model is not able to reflect the observed measurement. For example, trying to fit an isotropic model to a sample with cubic symmetry. It may be that during mode *identification*, a number of modes were identified incorrectly due to a low quality factor. It may also be that the sample itself is poorly cut or heterogeneous to the point that it cannot be modeled as homogeneous, in which case either a heterogeneous model or a new sample is needed. In all but the last case, a little more care and more numerical effort can fix the problem.

As mentioned earlier, the lower the Q of the specimen, the more difficult it is to identify consecutive peaks in any particular spectrum. This is often the case with soft rocks. Figure 2 showed this effect best. In high Q materials the peaks are relatively narrow and splitting of degenerate peaks is easier to detect than in lower Q materials. However, if a rock is inhomogeneous, the degenerate peaks may split enough so that both (or all) are observable. In this case, the inhomogeneity saves you during identification, but will likely come back to bite you in the inversion uncertainty.

Another source of difficulty in eigenfrequency identification comes from the physical coupling of the sample to the apparatus. Slight differences in the positioning of the sample between the transducers along with the force applied by the transducers contribute to this uncertainty. This error can be estimated by completely mounting and unmounting the same sample several times. This was performed on a cylinder of Elberton granite in Section 10.2. Although remounting the sample multiple times to estimate error is quite feasible for benchtop measurements, it is too time consuming for measurements involving complex sample preparation.

9.4 Resonance Inversion

Missing a normal mode during measurement can be fatal to the inversion if a shift between the measured and computed frequencies appears. The optimization may converge to the wrong model. To avoid this problem, we use sets of peaks we are confident in to perform the inversion. We can also use the weight (inverse data standard deviation) w_i in the objective function (cf. equation (8)). Lower weight would be associated with measured frequencies in which we have low confidence; although we will directly estimate the variance in the measured eigenfrequencies later.

The first iterations of the inversion are done with low order polynomials (N=5-7). These approximations can rapidly predict the first few mode frequencies accurately but fail to accurately predict higher modes. This allows one to obtain a rough idea of the true model. Higher order polynomials considerably increase the computation time but yield a better resolution using information of higher normal modes. Along with the different iterations, new information becomes new *a priori* information for the next iterations. One usually starts with an isotropic model. If a proper fit cannot be achieved we lower the symmetry of the model used: cubic, then hexagonal, etc.

The following results show that we can achieve a

model that fits our data. Uncertainties can arise from the polynomial fitting, from uncertainties in the data, from heterogeneities in the samples, and inaccuracies in the specification of the sample geometry (e.g., angles between faces are not exactly 90°, faces not perfectly flat, surface can have some roughness that scatters waves). Finding shifts between the predicted values with the inverted model and the actual measured frequencies is not surprising. What we need to know is if a model fits the data in a statistically significant way. We implement this in two ways. First, In our non-linear conjugate gradient method we assume to have found an answer when the gradient is lower than a certain threshold. Deciding on this threshold depends on the confidence we have in our data and what noise level we have. Second, we use a χ^2 test to know if a proper fit has been achieved.

In using the χ^2 test, it is necessary to know the uncertainty σ_i in frequency for each observed eigenfrequency f_i^o. To obtain these values, we assume that only physical mounting of the sample contributes to the uncertainty. Each time a sample is mounted, the eigenfrequency shifts slightly introducing some error. By completely mounting and unmounting a sample N times, we get N realizations of each spectrum. Then for each mode i, we calculate a mean eigenfrequency f_i and standard deviation σ_i. It is from these values, f_i and σ_i, that we will later calculate how statistically significant our inversion procedure is.

10 EXPERIMENTAL RESULTS

10.1 Benchtop Macor

A macor cylinder 3.095 cm long and 1.274 cm in diameter was measured in vacuum after being dried. Figure 10 shows the spectrum obtained. Five resonances were extracted from the data along with their corresponding Q values:

$f(kHz)$	Q
45.496	426
45.860	335
52.116	311
81.643	442
91.838	380

The first two peaks are the split members of a doublet. All five peaks were used in the inversion. Since macor is known to be an elastically stiff material with moduli similar to aluminum, we fit an isotropic model to the data. The observed frequencies and those predicted by the inverted model are:

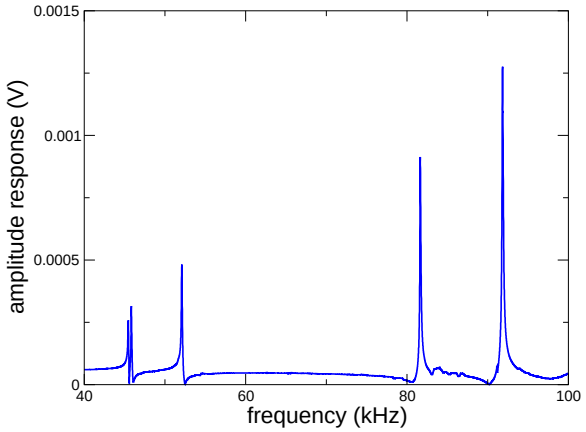

Figure 10. Benchtop sweep of a macor cylinder (3.095 cm × 1.274 cm × 1.274 cm). A sweep from 20 to 100 kHz was applied.

	$V_P(m/s)$	$V_S(m/s)$
resonance:	5655	3220
ultrasonic	5580	3162

Table 1. Observed macor elastic speeds from resonance and ultrasonic time-of-flight measurements.

$f_{\text{obs}}(kHz)$	$f_{\text{final}}(kHz)$	relative error
45.496	45.370	0.0028
45.860	45.455	0.0088
52.116	52.026	0.0017
81.643	81.957	-0.0038
91.838	91.826	-0.0001

The inverted model elastic velocities as well as the results of ultrasonic measurements on a different specimen are listed in Table 1. Assuming a 50 Hz error in measurement, the observed frequencies were fit with a χ^2 of 4.5. Since the sample was not remounted and measured multiple times, the 50 Hz error is a very rough guess. However, from experience, the quality of the data leads us to believe that 50 Hz is a safe guess, considering a 200 Hz error would give a χ^2 of 0.25 which is undoubtedly optimistic.

There is 1.3% discrepancy between the ultrasonic and resonance values for V_P and a 1.8% discrepancy for V_S. Since the resonance data depend only weakly on V_P, the nearly 2% discrepancy for V_S is likely a better measure of the difference in the two measurements.

We therefore conclude that for isotropic materials, we can find V_P and V_S quite accurately. However, in next section we attempt the same scheme for Elberton granite which is anisotropic in order to test our ability to measure and invert for lower symmetry materials.

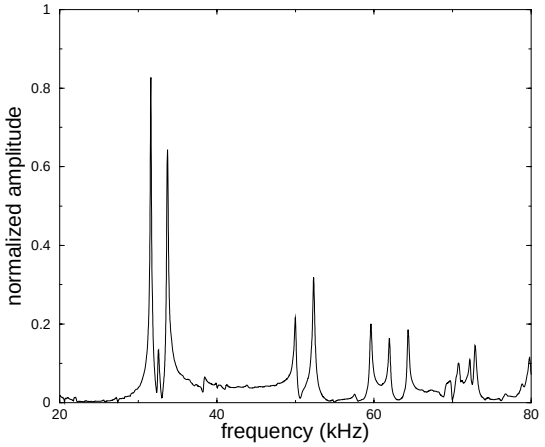

Figure 11. Spectrum of a cylindrical granite core (2.27 cm diameter by 7.15 cm long). A sweep from 20 to 100 kHz with a 5Hz step was applied.

10.2 Granite cylinder

The anisotropy of particular granites due to microfractures has been known for many years and is a good test for our RUS methods (@warning Citation 'granite' on page 12 undefined). Previous work suggests that to be able to accurately measure anisotropy in a sample we would need at least 20 consecutive resonant frequencies (@warning Citation 'ulrich02' on page 12 undefined). We studied a 75 mm by 25 mm cylindrical core of Elberton granite and were able to measure and identify the first 25 normal modes. An example of a measured spectrum can be seen in Figure 11. In order to obtain the uncertainties in frequency for our χ^2, the granite sample was completely mounted and unmounted 14 times. The mean values of the eigenfrequencies and their corresponding uncertainties are shown in Table 2. The differences in frequency and amplitude can be seen in Figure 13. Frequency uncertainty ranged from 22 Hz to 801 Hz with a mean uncertainty of 258 Hz. Although the amplitudes in Figure 13 seemed to vary greatly for specific modes, the Q values of the modes were fairly stable. In Figure 14, the Q's range from approximately 100 to 200, with an average uncertainty of 41.

A first attempt to match the first six normal modes with an isotropic model was unsuccessful. Larger frequencies could not be fitted at all. We then tried anisotropic models with lower and lower symmetries. Finally with the use of an orthorhombic model (with 9 independent parameters, c_{11}, c_{22}, c_{33}, c_{23}, c_{13}, c_{12}, c_{44}, c_{66}, c_{55}) we can achieve a fit of 25 consecutively measured resonances with a χ^2 of 37. The resulting χ^2 from fitting all 25 peaks with each possible symmetry is:

$f^{(obs)}$ (MHz)	$f^{(pre)}$ (MHz)	$\sigma^{(obs)}$ (Hz)	freq. #
0.031490	0.031490	50	4
0.032380	0.031680	58	5
0.033620	0.032540	29	6
0.049670	0.050080	31	8
0.052280	0.052230	43	9
0.059500	0.059400	22	10
0.061970	0.061960	27	11
0.064380	0.065220	39	12
0.069760	0.070050	75	13
0.071860	0.071800	801	14
0.072330	0.072310	327	15
0.075700	0.075750	376	16
0.076020	0.076070	411	17
0.078200	0.078300	396	18
0.079200	0.078640	401	19
0.079200	0.079256	406	20
0.081080	0.080978	789	21
0.084800	0.084833	393	22
0.085600	0.085088	575	23
0.086800	0.087480	394	24
0.087500	0.087512	569	25

Table 2. Comparison between the observed ($f^{(obs)}$) and predicted ($f^{(pre)}$) normal mode frequencies with uncertainties for an Elberton granite core.

symmetry	*# parameters*	χ^2
isotropic	2	3486
cubic	3	4089
hexagonal	5	1793
tetragonal	6	496
orthorhombic	9	37

The comparison between predicted and measured frequencies can be found in Table 2. The inverted model is $c_{11} = 67.87 \pm 1.14$GPa, $c_{22} = 81.94 \pm 0.70$GPa, $c_{33} = 81.83 \pm 0.69$GPa, $c_{23} = 27.15 \pm 1.10$GPa, $c_{13} = 28.95 \pm 0.86$GPa, $c_{12} = 39.76 \pm 0.14$GPa, $c_{44} = 23.72 \pm 0.02$GPa, $c_{55} = 29.16 \pm 0.01$GPa and $c_{66} = 28.67 \pm 0.01$GPa. In the plane perpendicular to the symmetry axis, this yields qP wave speeds of 5003±21 m/s and 4556±38 m/s in the c_{22} and c_{11} directions, respectively. This nearly 10% P-wave anisotropy is consistent with the ultrasonic measurements performed by

Malcolm and Scales (@warning Citation 'malcolm:02' on page 13 undefined) shown in Figure 12.

In the anisotropic symmetry plane the S wave speeds are 2986±5 m/s for the vertical polarization, 2694±10 m/s and 2961±5 m/s for the horizontal polarization in the vertical and horizontal directions, respectively. For orthorhombic symmetry, c_{44} and c_{55} decouple and are not the same as shown by their fit values. Also note the difference in values for the other decoupled coefficients: c_{12} and c_{23}, c_{11} and c_{22}. Here our χ^2 test was used as a tool to decide the symmetry of the sample. Granite is known to be anisotropic. The way granite is extracted and cut in quarries takes advantage of that anisotropy. This ensured, fortunately, that for such a sample we have the anisotropic symmetry axis coinciding with the geometrical symmetry axis of the sample. This makes our assumptions valid in this case. Otherwise, we would have to invert for the crystallographic axis as well.

It must be mentioned that it is generally accepted that 5 resonances are needed for each elastic parameter being fit. This being the case, that means for orthorhombic symmetry, we would need 45 (consecutive) observable resonances. For nearly any rock, this is all but impossible. Of course, going to monoclinic symmetry would reduce our χ^2 as well, but we have independent measurements telling us to stop at orthorhombic. Data by Malcolm and Scales (@warning Citation 'malcolm:02' on page 14 undefined) in Figure 12 shows that there is evidence of anisotropy in the sample (roughly 10%), but it is not due simply to inhomogeneity. By observing the surface waves traveling around the core in opposite directions, it can be said that the paths of the counter-rotating waves are the same, therefore the observed splitting in the arrivals of the two waves at some point must be due to inhomogeneity. This was measured to be 5%. We conclude then, that while we are certain it is not necessary to reduce symmetry below orthorhombic, it is accepted that a portion of the error in the χ^2 is attributed to the sample being inhomogeneous.

11 CONCLUSION

RUS offers a promising technique for characterizing the elastic moduli of rocks. It fills a gap in existing spectrum of measurement techniques between low-frequency stress-strain and ultrasonic delay-time measurements. RUS allows one to infer the complete elastic tensor from a single measurement, without having to machine a sample according to the assumed anisotropic symmetry.

To implement RUS, we first derived the forward model to compute the frequencies and shapes of the normal modes. The forward model is greatly simplified with a polynomial approximation, which allows for numerical evaluations. Further symmetry assumptions yield a speed-up of the algorithm.

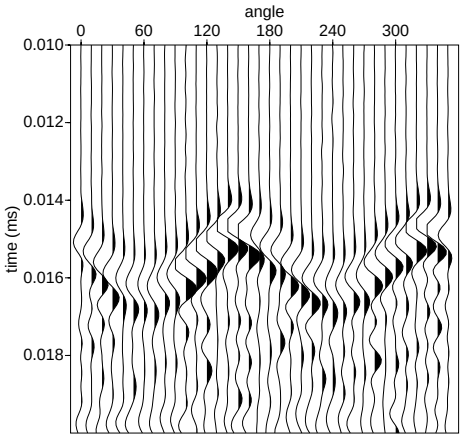

Figure 12. Laser ultrasonic measurement of P-waves traveling straight across a 5.5 cm diameter granite core. A pulsed IR laser in the thermoelastic regime is used to excite elastic waves which are measured with a laser Doppler vibrometer. The source and detector positions are antipodal on a line through the middle of the sample. The sample is rotated 10 degrees between measurements. For more details see (@warning Citation 'malcolm:02' on page 14 undefined).

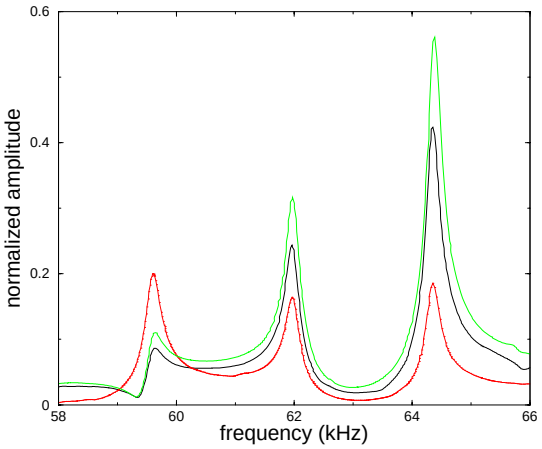

Figure 13. A set of 3 resonance peaks from 3 of 14 Elberton granite data sets. Amplitudes varied with each set due to mounting. The average uncertainty in eigenfrequency from 14 complete mountings and remountings was $258Hz$. This uncertainty does include error from misidentification of resonance peaks during fitting.

RUS methods have been used extensively in the materials science community to evaluate high-Q crystals of small sizes. The method then yields an efficient way to invert for elastic properties. In this paper we show that this method may be applied to larger objects and also objects that have small-scale heterogeneities such as rocks. We have, so far, only studied clean, relatively

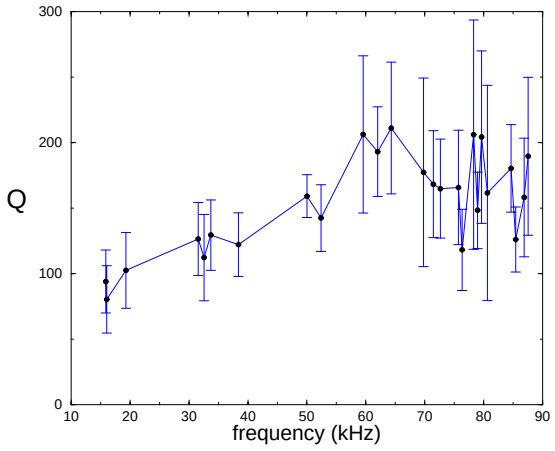

Figure 14. Q values for Elberton granite have an increasing trend over the measured range. Higher frequencies had sligtly higher uncertainties, or more variation in estimated Q from spectrum to spectrum.

homogeneous rock samples, using normal modes whose wavelengths are large compared to the granularity of the sample. For samples with large-scale heterogeneity, it will presumably be necessary to perform a preliminary analysis (e.g., tomography) to characterize the sample. We have also analyzed the sensitivity of the resulting elastic moduli to transducer position, sample aspect ratio and frequency of the measurement. We have found no globally optimal solution to the problem of experiment design. In practice it is important to know which modes most influence the parameters one is interested in and how best to excite those modes by appropriate choice of sample shape and transducer position.

ACKNOWLEDGMENT

We thank Lydia H. Deng for her help with the COOOL library, Mike Batzle for his help in the sample preparation and Gary R. Olhoeft for valuable comments. This work was partially supported by the National Science Foundation (EAR-0111804), and the sponsors of the Fluids III Consortium at the Center for Rock Abuse, Colorado School of Mines.

APPENDIX A: SYMMETRIES AND BLOCK DIAGONALIZATION

An elastic medium is characterized by its stiffness tensor (C_{ijkl}). With the so-called Voigt notation (@warning Citation 'th86' on page 15 undefined), one can represent the medium by a 6×6 matrix, (c_{ij}) in accordance with

$$\begin{array}{lccccccc} ij \text{ or } kl: & 11 & 22 & 33 & 32=23 & 31=13 & 12=21 \,. \\ & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\ & 1 & 2 & 3 & 4 & 5 & 6 \end{array}$$

If one restrict oneself to orthorhombic or higher symmetries the non-zero entries of the matrix $(c_{ij})$are given by

$$\begin{pmatrix} c_{11} & c_{12} & c_{13} & & & \\ c_{12} & c_{22} & c_{23} & & & \\ c_{13} & c_{23} & c_{33} & & & \\ & & & c_{44} & & \\ & & & & c_{55} & \\ & & & & & c_{66} \end{pmatrix} .$$

The corresponding symmetries are isotropic, cubic, hexagonal, tetragonal, and orthorhombic. Note that we also assume that the symmetry axis are aligned with the coordinate axis.

We can exploit the vanishing entries and symmetries of the resulting stiffness tensor C_{ijkl} while reordering the matrices E and Γ [see equations (4) and (5)]. Reordering yields simplifications that can be exploited for numerical efficiency. Consider the entry

$$\Gamma_{\lambda i \lambda' k} = C_{ijkl} \int_V \partial_j \phi_\lambda \, \partial_l \phi'_\lambda \, dV \,, \tag{A1}$$

where $\phi_\lambda = x^{\mathsf{l}} \, y^{\mathsf{m}} \, z^{\mathsf{n}}$ and $\phi'_\lambda = x^{\mathsf{l}'} \, y^{\mathsf{m}'} \, z^{\mathsf{n}'}$.

Assume first the $i = 1$ and $k = 2$. Then the only non-zero terms in the sum in equation (A1) are for $(j, l) = (1, 2)$ and $(j, l) = (2, 1)$. In both cases the integrand is of the form

$$K \, x^{\mathsf{l}+\mathsf{l}'-1} \, y^{\mathsf{m}+\mathsf{m}'-1} \, z^{\mathsf{n}+\mathsf{n}'} \,.$$

Since we integrate with symmetric limits if such a term is non zero then

$$\begin{aligned} \mathsf{l} \text{ even (odd)} &\implies \mathsf{l}' \text{ odd (even)} \,, \\ \mathsf{m} \text{ even (odd)} &\implies \mathsf{m}' \text{ odd (even)} \,, \\ \mathsf{n} \text{ even (odd)} &\implies \mathsf{n}' \text{ even (odd)} \,. \end{aligned}$$

If now $i = 1$ and $k = 3$, then the only non-zero terms in the sum in in equation (A1) are for $(j, l) = (1, 3)$ and $(j, l) = (3, 1)$. In both cases the integrand is of the form

$$K \, x^{\mathsf{l}+\mathsf{l}'-1} \, y^{\mathsf{m}+\mathsf{m}'} \, z^{\mathsf{n}+\mathsf{n}'-1} \,,$$

and if such a term is non zero then

$$\begin{aligned} \mathsf{l} \text{ even (odd)} &\implies \mathsf{l}' \text{ odd (even)} \,, \\ \mathsf{m} \text{ even (odd)} &\implies \mathsf{m}' \text{ even (odd)} \,, \\ \mathsf{n} \text{ even (odd)} &\implies \mathsf{n}' \text{ odd (even)} \,. \end{aligned}$$

If now $i = 1$ and $k = 1$, then the only non-zero terms in the sum in in equation (A1) are for $(j, l) = (1, 1)$, $(j, l) = (2, 2)$, and $(j, l) = (3, 3)$. In these cases to obtain non-zero terms we require

$$\begin{aligned} \mathsf{l} \text{ even (odd)} &\implies \mathsf{l}' \text{ odd (even)} \,, \\ \mathsf{m} \text{ even (odd)} &\implies \mathsf{m}' \text{ even (odd)} \,, \\ \mathsf{n} \text{ even (odd)} &\implies \mathsf{n}' \text{ odd (even)} \,. \end{aligned}$$

We obtain similar results for the cases $(i, k) = (2, 2)$ and

$(i,k) = (3,3)$. There are 8 choices for the parities of l, m, and n. The previous relations give an equivalence relation on the set

$$\{(i,\lambda) \mid 1 \leq i \leq 3\ ,\ \lambda = (\mathsf{l},\mathsf{m},\mathsf{n}) \text{ with } \mathsf{l}+\mathsf{m}+\mathsf{n} < N\}\ ,$$

with 8 equivalence classes. If one takes two representatives of two distinct equivalence classes the associated entry in Γ is zero. Therefore, by grouping the representative of the same classes we decompose the matrix Γ into 8 diagonal blocks. It is straightforward to check that this yield a similar decomposition for the matrix E because of the Kronecker symbol in equation (4).

With such a block diagonal decomposition, instead of solving one large generalized eigenvalue problem, we solve eight smaller ones. Since the computational complexity of the algorithm used is proportional to R^3 (R is the dimension of the problem) the gain in efficiency is obvious.

A1 Low symmetry

In the case of symmetries lower than orthorhombic a block diagonalization of the matrices Γ and E does not apply. This simplification similarly breaks down when symmetry axes do not align with coordinates axes (i.e., sample faces). In such a case call (R_{ij}), $i,j = 1,2,3$, the rotation matrix transforming coordinates from the system associated to the symmetry axes to that associated to the sample faces.. The stiffness tensor (c_{ijkl}) transforms according to $c_{i'j'k'l'} = R_{ii'}R_{jj'}R_{kk'}R_{ll'}c_{ijkl}$. Such a transformation breaks the structure of the matrix in the sample-face coordinate system. This yields a higher computational complexity for the modeling operation.

Note that The rotation matrix involves three additional parameters. These can be chosen to be Euler angle as in (@warning Citation 'gold' on page 16 undefined). To determine symmetry axes one has to invert for these three additional parameters. Note also that the determination of symmetry axes is restricted to samples with distinct face orientation. In the case of a spherical sample the additional parameters cannot be inverted for. In the case of a cylindrical samples with a circular base the axes of symmetry can only be inverted for up to a rotation around the sample symmetry axis (@warning Citation 'backus' on page 16 undefined).

The inversion procedure we have followed assumes some a priori knowledge of the symmetry axis for the anisotropy of the sample. Another approach that one could follow is to invert for the general elastic tensor, i.e. the 21 elastic constants. In such a case inversion requires knowledge of a large number of resonance frequencies. Assuming that one can invert the 21 elastic parameters, one would then want to know what type of anisotropic symmetry the sample exhibits. In other words the question is: Can one guess the anisotropic symmetry with the elastic tensor given in any orthogonal coordinate system? Backus (page 649) propose one strategy to answer this question. His approach is based on the representation of the elastic tensor with harmonic tensors. The elastic tensor, C_{ijkl} is first decomposed into two tensors $C = S + A$, one symmetric S_{ijkl}, one asymmetric, A_{ijkl}. S can be represented with the help of three harmonic tensors, one of order 4, one of order 2, one of order 0. S can be represented with the help of two harmonic tensors, one of order 2, one of order 0. In three dimension the action of a harmonic tensor $H^{(q)}$ of order q can be written as

$$H^{(q)}(\mathbf{r}) = A \prod_{1 \leq \nu \leq q} \mathbf{a}^{(\nu)} \cdot \mathbf{r} + r^2 P^{(q-2)}(\mathbf{r})\ , \tag{A2}$$

where $\mathbf{r}$ is the position, $\mathbf{a}^{(\nu)}$, $1 \leq \nu \leq q$ is a *unique* set of directions, A is a scalar, $r = |\mathbf{r}|$, and $P^{(q-2)}$ is a homogeneous polynomial of order $q-2$. Backus shows that the symmetry properties of $H^{(q)}$ coincide with that of the set of directions $\mathbf{a}^{(\nu))}$, e.g. $H^{(q)}$ is rotation invariant with respect to some axis if and only if the set of directions $\mathbf{a}^{(\nu)}$ is so too. Hence one can decompose the five Harmonic tensors according to (A2) and obtain five 'bouquets' of directions. The common symmetry of these 'bouquets' is then that of the elastic tensor C.

In our inversion framework, one would thus have the geometrical mean to obtain the symmetry axis. Because of uncertainties in the data and hence in the inverted components of the elastic tensor one would only be able to make a guess. One can then however transform the elastic tensor C into the system of guessed symmetry axis. In such a system some entries of the tensor should vanish, at least within the uncertainties. Such an observation can then confirm the symmetry type of the elastic medium. With an anisotropic symmetry in mind, one could now proceed with the main method discussed in this paper and having fewer parameters to invert for, sharpen the estimates of the elastic constants.

APPENDIX B: EVALUATION OF POLYNOMIAL INTEGRALS

B1 Rectangular parallelepiped

As far as computing integral (7), the case of a rectangular parallelepiped is easiest. We know that if one of l, m or n is odd then $\mathcal{I}_V$ vanishes, otherwise we have

$$\mathcal{I}_V = \frac{8\, d_1^{\mathsf{l}+1} d_2^{\mathsf{m}+1} d_3^{\mathsf{n}+1}}{(\mathsf{l}+1)(\mathsf{m}+1)(\mathsf{n}+1)}\ .$$

B2 Cylinder

Again, if either l, m or n is odd then $\mathcal{I}_V$ is zero. If they are all even, then $\mathcal{I}_V$ is equal to

$$\mathcal{I}_V = \frac{2 d_1^{\mathsf{l}+1} d_2^{\mathsf{m}+1} d_3^{\mathsf{n}+1}}{\mathsf{n}+1} \int_0^1 r^{\mathsf{l}+\mathsf{m}+2}\, dr \int_0^{2\pi} \cos^{\mathsf{l}}\theta\ \sin^{\mathsf{m}}\theta\, d\theta$$

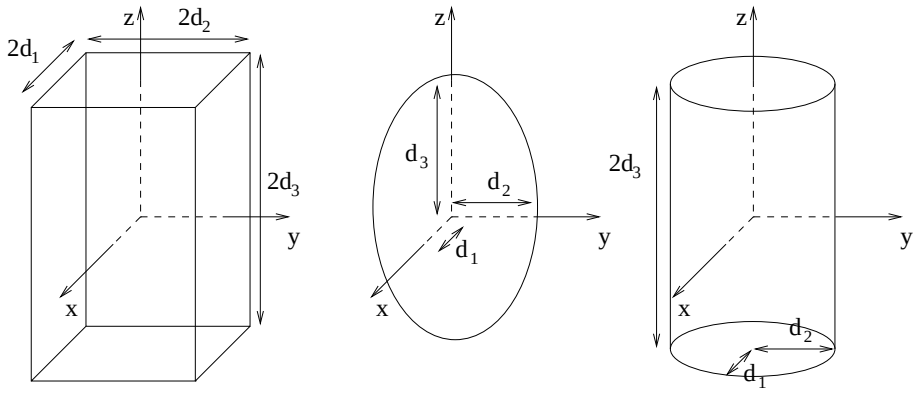

Figure A1. Dimension associated with the different shapes used in our work: rectangular parallelepipeds, spheroids, cylinders.

after integrating over z and applying the change of variables, $x = d_1 r\cos\theta$, $y = d_2 r\sin\theta$. The second integral is the integral of powers of trigonometric functions over a period and is equal to (@warning Citation 'gr65' on page 16 undefined)

$$2\pi\,\frac{(\mathsf{l}-1)!!\,(\mathsf{m}-1)!!}{(\mathsf{l}+\mathsf{m})!!}\ ,$$

where

$$\begin{aligned}&(2a)!! = 2a(2a-2)(2a-4)...(2) = 2^a a!,\\ &(2a+1)!! = (2a+1)(2a-1)(2a-3)...(1),\ \text{and}\\ &0!! = (-1)!! = 1\ .\end{aligned}$$

which gives

$$\mathcal{I}_V = 4\pi\,\frac{d_1^{\mathsf{l}+1} d_2^{\mathsf{m}+1} d_3^{\mathsf{n}+1}}{(\mathsf{n}+1)}\,\frac{(\mathsf{l}-1)!!\,(\mathsf{m}-1)!!}{(\mathsf{l}+\mathsf{m}+2)!!}\ .$$

Note that the form of $\mathcal{I}_V$ shows the cylindrical symmetry of the system.

B3 Spheroid

Here also, $\mathcal{I}_V$ (equation (7)) is only non-zero for l, m and n even. In the spheroid case, we apply the change of variables,$x = d_1 r\sin\theta\,\cos\phi$, $y = d_2 r\sin\theta\,\sin\phi$, $z = d_3 r\cos\theta$. This yields the integral

$$\begin{aligned}\mathcal{I}_V &= d_1^{\mathsf{l}+1} d_2^{\mathsf{m}+1} d_3^{\mathsf{n}+1}\\ &\quad\int_0^1 r^{\mathsf{l}+\mathsf{m}+\mathsf{n}+2}\,dr\int_0^{\pi}\sin^{\mathsf{l}+\mathsf{m}+1}\theta\,\cos^{\mathsf{n}}\theta\,d\theta\int_0^{2\pi}\cos^{\mathsf{l}}\phi\,\sin^{\mathsf{m}}\phi\,d\phi\\ &= \frac{d_1^{\mathsf{l}+1} d_2^{\mathsf{m}+1} d_3^{\mathsf{n}+1}}{\mathsf{l}+\mathsf{m}+\mathsf{n}+3}\int_0^{\pi}\sin^{\mathsf{l}+\mathsf{m}+1}\theta\,\cos^{\mathsf{n}}\theta\,d\theta\int_0^{2\pi}\cos^{\mathsf{l}}\phi\,\sin^{\mathsf{m}}\phi\,d\phi\ .\end{aligned}$$

The second integral is equal to (@warning Citation 'gr65' on page 17 undefined)

$$2\pi\,\frac{(\mathsf{l}-1)!!\,(\mathsf{m}-1)!!}{(\mathsf{l}+\mathsf{m})!!}\ ,$$

whereas the first is (@warning Citation 'gr65' on page 17 undefined)

$$2\,\frac{(\mathsf{l}+\mathsf{m})!!\,(\mathsf{n}-1)!!}{(\mathsf{l}+\mathsf{m}+\mathsf{n}+1)!!}\ .$$

Therefore, we obtain

$$\mathcal{I}_V = 4\pi\,d_1^{\mathsf{l}+1} d_2^{\mathsf{m}+1} d_3^{\mathsf{n}+1}\,\frac{(\mathsf{l}-1)!!\,(\mathsf{m}-1)!!\,(\mathsf{n}-1)!!}{(\mathsf{l}+\mathsf{m}+\mathsf{n}+3)!!}\ .$$

Note that the form of $\mathcal{I}_V$ shows the spheroidal symmetry of the system. Other geometrical shapes and values of I_V are given in Vissher et al. (@warning Citation 'vi91' on page 17 undefined).