

Resonant Ultrasound Spectroscopy: theory and application

Brian J. Zadler,¹ Jérôme H. L. Le Rousseau,² John A. Scales¹ and Martin L. Smith³

¹Physical Acoustics Laboratory and Center for Wave Phenomena, Department of Geophysics, Colorado School of Mines, Golden, CO, USA

²Laboratoire d'Analyse, Topologie et Probabilités – CNRS UMR6632, Centre de Mathématiques et Informatique, Université de Provence (Aix-Marseille I), France

³New England Research, White River Junction, VT, USA

Accepted 2003 June 27. Received 2003 June 19; in original form 2002 November 19

SUMMARY

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). We use synchronous detection methods 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 normal-mode frequencies. We have successfully applied this technique to a variety of isotropic and anisotropic samples, both man-made and natural. In this paper we will show in detail the procedure applied to a cylindrical core of Elberton granite. By means of a statistical fit of the measured normal modes and an independent laser ultrasonic measurement, the granite core was inferred to have orthorhombic symmetry. A 10 per cent P -wave anisotropy was measured in the plane perpendicular to the core axis.

Key words: acoustic spectroscopy, normal modes, resonance, rock physics.

1 INTRODUCTION

1.1 Laboratory measurements of elastic properties

Mechanical methods for measuring the elastic properties of laboratory specimens are divided into three types: quasi-static, resonance, and time-of-flight. 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 (Batzle & Wang 1992), subject the sample to deformations that are 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 10^3 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 behaviour 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 data set 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 principle to measure arbitrary anisotropy and it is often reasonable in practice to do so. Under favourable 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 centimeters 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 comfortably easy to visualize.

1.2 Resurgent resonance

Resonance techniques have been used to measure elastic properties in the laboratory since at least the 1920s (McSkimmin 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. See (Lucet *et al.* 1991) for an extensive discussion of resonance measurements for long, thin samples.

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 (1968) 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 & 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 greatly 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 some 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 (Maynard 1996), 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, usually ≥ 500 . 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 in a given range that 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.

The biggest obstacle to 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 & Sarrao (1997). 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 min. For a 10 cm 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 THEORY AND COMPUTATION

The theory of the small oscillations of linear systems is a well-developed part of classical mechanics. Goldstein (Goldstein 1981) provides an extensive introduction to the topic and notes that its origins lay in Routh's 1877 Adam's Prize Essay. Lord Rayleigh's landmark *The Theory of Sound* (Rayleigh 1896, 1945 re-issue), also published in 1877, developed many of the results central to the topic. He also examined applications of small oscillation theory to strings, plates, bars, and other elastic systems.

More recent advances relevant to laboratory-scale resonance measurements are covered in Migliori & Sarrao (1997), including sufficient discussion of computational and laboratory aspects to begin using resonance measurements. A much different modern reference, Dahlen & Tromp (1998), provides a deep and *exhaustive* investigation of elastic resonances with particular emphasis on the nearly-spherical case so important to planetary geophysics.

In this section we will summarize the principal theoretical and numerical apparatus we need for our purposes:

- (i) *variational approximations for resonances*—the basic formulation of resonance as a variational calculation and Visscher's simple, approximate basis that makes the necessary computations fast enough and very flexible
- (ii) *elastic energy and damping*—expressions that tell us how energy in a particular resonance is partitioned between compression and shear, how its attenuation is caused by the respective damping of compressional and shear waves, and how the mode's resonant frequency is sensitive to the elastic properties of the medium
- (iii) *excitation calculations*—how to compute the expected relative excitation of individual resonances during a particular resonance experiment.

3.1 Variational approximations for resonance

Let \mathcal{V} be an isolated body, that is, one bounded by a closed, stress-free surface. Let C_{ijkl} be its elastic stiffness tensor and let ρ be its density; both quantities may vary with position in \mathcal{V} .

Let ω be a non-negative real number, and $\mathbf{u}(\mathbf{r})$ be a real-valued function of position \mathbf{r} in \mathcal{V} . Then the combination $\{\omega, \mathbf{u}\}$ is a *free oscillation* or *resonance* if the real-valued displacement field

$$\mathbf{s}(\mathbf{r}, t) = \Re\{\mathbf{u}(\mathbf{r})e^{i\omega t}\} \quad (1)$$

satisfies the elastic equations of motion in \mathcal{V} and the stress-free boundary condition on its surface.

There is no *prima facie* reason to believe that such solutions must exist. In fact, for a given \mathcal{V} and virtually all possible values of ω there will not exist any \mathbf{u} which leads to a solution of the equations of motion and boundary conditions. It is a deep, well-known, and important result of mechanics that a countably infinite set of such solutions must exist; it is a miracle that there are practical ways to compute them.

To proceed, we need the notions of the kinetic and potential energy associated with a resonance. The potential energy E_p associated with the displacement field \mathbf{u} is given by the strain energy (Aki & Richards 1980)

$$E_p = \frac{1}{2} \int_{\mathcal{V}} C_{ijkl} \partial_j u_i \partial_l u_k dV, \quad (2)$$

where u_i , $i = 1, 2, 3$ are the Cartesian components of \mathbf{u} and we are using the summation convention for repeated indices. The corresponding kinetic energy E_k is given by

$$E_k = \omega^2 K \quad (3)$$

where

$$K = \frac{1}{2} \int_{\mathcal{V}} \rho u_i u_i dV. \quad (4)$$

We can obviously compute E_p and E_k for any ω and \mathbf{u} , resonance or otherwise. Classical mechanics comes to our aid by assuring us that (Dahlen & Tromp 1998, Section 4.1.3) the quantity

$$\mathcal{I} = \omega^2 K - E_p \quad (5)$$

is *stationary* if and only if ω and \mathbf{u} are a resonance of \mathcal{V} . (By stationary we mean that the value of \mathcal{I} does not change if we replace \mathbf{u} by $\mathbf{u} + \delta\mathbf{u}$ where $\delta\mathbf{u}$ is any small displacement which obeys the stress-free boundary conditions.) Stationarity, happily, turns out to be the key to a powerful approximation technique, called the *Rayleigh–Ritz* method.

Our next step is to turn the function \mathbf{u} into something we can get a grip on; we do this by representing it in a prescribed basis. Let $\{\phi_\lambda(\mathbf{r}), \lambda = 1, \dots, N\}$ be a set of N specified functions over \mathcal{V} . By *specified* we mean that they are some set which we have chosen and thus know how to evaluate numerically. We then represent each of the components of \mathbf{u} by

$$u_i = a_{i,\lambda} \phi_\lambda. \quad (6)$$

So now in order to determine \mathbf{u} we must specify the $3N$ numbers $a_{i,\lambda}$. In particular, we must determine them so that the resultant value of \mathcal{I} is stationary.

If we insert the representation 6 into eqs (2) through (4), we can express the stationary quantity 5 as

$$\mathcal{I} = \omega^2 \alpha \cdot \mathbf{K} \cdot \alpha - \alpha \cdot \mathbf{E} \cdot \alpha \quad (7)$$

where α is a vector comprising the juxtaposed components of $a_{i,\lambda}$. Stationarity asserts that \mathcal{I} will be unchanged to first order under any small perturbation in the values of α which are consistent with the stress-free boundary conditions. This requirement leads to the conclusion that $\{\omega^2, \alpha\}$ are a stationary solution if and only if they satisfy

$$\omega^2 \mathbf{K} \cdot \alpha = \mathbf{E} \cdot \alpha. \quad (8)$$

Eq. (8) is a standard form for the *generalized symmetric eigenvalue problem* (see Wilkinson & Reinsch 1971) and is readily solved

with modern, freely-available software packages such as Lapack (Anderson *et al.* 1999). Both \mathbf{K} and \mathbf{E} are real and symmetric, and \mathbf{K} is positive definite. For any energetically-stable material, the elastic energy matrix \mathbf{E} will be non-negative definite. These conditions guarantee that the eigenvalues of 8 will be non-negative. In fact there will be six zero eigenvalues, corresponding to three degrees of rigid-body translation and three degrees of rigid-body rotation, and all of the rest will be positive, corresponding to elastic resonances.

The devil in Rayleigh–Ritz methods often lies in the details of the basis functions, $\phi_\lambda(\mathbf{r})$. The first studies used trigonometric functions (Holland 1968) and orthogonal polynomials (Demarest 1969) tailored to a rectangular sample shape. These choices promised good numerical efficiency, in the one case because the trigonometric functions mimicked the basic wave-like nature of the elastic displacement fields and in the other case because the orthogonal polynomials were known to have good numerical stability. On the other hand they required substantial effort to formulate and programme the various integrals, and the effort would be needed afresh for each new sample geometry and symmetry.

These choices were appropriate in an environment in which algebraic and programming effort was regarded as *cheaper* than computer time. The major contribution of Visscher *et al.* (1991) was the discovery that the simplest possible choice of ϕ_λ , namely monomials of the form

$$\phi_\lambda = x^{\eta(\lambda)} y^{\zeta(\lambda)} z^{\xi(\lambda)}, \quad (9)$$

where η , ζ , and ξ are positive integers, could provide a basis that was both numerically adequate (given the great advances in computing power) and very flexible and convenient to implement for a wide variety of shapes \mathcal{V} and symmetries. In fact, with a little attention to scaling (Maynard & So 1991) it is numerically much more than adequate.

We still have to choose the set of exponents $\{\eta, \zeta, \xi\}$. The most common practice is to choose some positive integer N_p and use the set of all values satisfying

$$\eta + \zeta + \xi \leq N_p \quad (10)$$

The larger N_p , of course, the more accurate the results. For a given N_p there are

$$3 \binom{N_p + 3}{3} = (N_p + 1)(N_p + 2)(N_p + 3)/2 \quad (11)$$

basis elements counting all three components of \mathbf{u} . The size of the problem, the length of α , grows as N_p^3 . Since the computational effort required to solve 8 scales as the cube of the length of α , that effort grows as N_p^9 .

An alternative scheme is to restrict the individual ranges of the three indices to be proportional to some measure of sample size in the dimensions. This approach tries to roughly spread the resolution of the basis around in a way that recognizes the wavelike nature of the solution, and is probably the better choice. Our experience is that both are satisfactory.

Our task now is to compute the components of \mathbf{K} and \mathbf{E} , and then to solve the resultant generalized eigenvalue problem for its eigenvalues and eigenvectors. Each eigenvalue is the squared angular frequency of a resonance and its eigenvector comprises the coefficients of the basis functions for the resonance displacement field. To properly formulate this we have to sort out some bookkeeping issues.

Recall that α is a vector containing all of the coefficients for all of the components of \mathbf{u} . Each element in α is one of the elements from

the vector of coefficients \mathbf{a} for one of the cartesian components of \mathbf{u} . Use λ as an index into α . For each value of λ , we must keep track of

- $i(\lambda)$ the Cartesian component of \mathbf{u} to which α_λ applies
- $\eta(\lambda)$ the exponent of x in ϕ_λ
- $\zeta(\lambda)$ the exponent of y in ϕ_λ
- $\xi(\lambda)$ the exponent of z in ϕ_λ .

Then

$$K = \frac{1}{2} \int_{\mathcal{V}} \rho u_i u_i dv \quad (12)$$

$$K_{\lambda\lambda'} = \frac{1}{2} \int_{\mathcal{V}} \rho \alpha_\lambda x^{\eta(\lambda)} y^{\zeta(\lambda)} z^{\xi(\lambda)} \delta_{i(\lambda)i(\lambda')} \quad (13)$$

$$\times \alpha_{\lambda'} x^{\eta(\lambda')} y^{\zeta(\lambda')} z^{\xi(\lambda')} dv \quad (14)$$

$$= \frac{1}{2} \int_{\mathcal{V}} \rho x^{\eta(\lambda)+\eta(\lambda')} y^{\zeta(\lambda)+\zeta(\lambda')} z^{\xi(\lambda)+\xi(\lambda')} dv \quad (15)$$

There is an analogous but messier expansion for \mathbf{E} . In both cases the results look messier than they are. Each consists of a bookkeeping hair-shirt wrapped around integrals of the form

$$\int_{\mathcal{V}} f(x, y, z) x^l y^m z^n dv \quad (16)$$

where $f(x, y, z)$ is either density or one of the Cartesian components of the elastic stiffness tensor. In any region $S \subseteq \mathcal{V}$ in which $f(x, y, z)$ can be written as a sum of monomial basis functions, a condition that includes the nearly universal case of being constant in S , all of the integrals we need are of the form

$$\mathcal{I}_S(l, m, n) = \int_S x^l y^m z^n dv \quad (17)$$

Usually, S is \mathcal{V} . In Appendix B we give explicit results for \mathcal{I}_S for the cases of a rectangular parallelepiped, an ellipsoidal cylinder, a triaxial ellipsoid, and a right-circular cylindrical wedge.

Eq. (8) can be further simplified for its numerical evaluation. In the orthorhombic and higher symmetry case, we can order matrices \mathbf{K} and \mathbf{E} 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.

3.2 Elastic energy and damping

We shall need to know how perturbations in the sample's material properties, ρ and, particularly, C , affect its resonance frequencies. The results of this calculation are also useful in the special case that perturbations in C are *imaginary*, in which instance the results tell us how anelasticity in the sample contribute to the total dissipation or Q of each resonance.

In order to make this computation we must appeal to Rayleigh's Principle, famously stated in Article 88 of Rayleigh (1896, 1945 re-issue), which tells us how to compute the perturbed eigenvalues of a perturbed hermitian operator without computing the perturbed eigenvectors. (This is really just standard first-order perturbation theory for hermitian operators—see for instance chapter 10 of Mathews & Walker 1965—but it's more fun to get it from Rayleigh.) This result says that if we perturb the matrices \mathbf{K} and \mathbf{E} in 8 thus

$$\mathbf{K} \rightarrow \mathbf{K} + \delta\mathbf{K} \quad (18)$$

$$\mathbf{E} \rightarrow \mathbf{E} + \delta\mathbf{E}, \quad (19)$$

and we represent the perturbation in the squared frequency of the i th resonance by

$$\omega_i^2 \rightarrow \omega_i^2 + \delta\omega_i^2 \quad (20)$$

then we can compute $\delta\omega_i^2$ from

$$\delta\omega_i^2 = \frac{\alpha^{(i)} \cdot (\delta\mathbf{E} - \omega_i^2 \delta\mathbf{K}) \cdot \alpha^{(i)}}{\alpha^{(i)} \cdot \mathbf{K} \cdot \alpha^{(i)}} \quad (21)$$

where $\alpha^{(i)}$ is the eigenvector associated with ω_i^2 . In eq. (21) $\delta\mathbf{K}$ and $\delta\mathbf{E}$ are prescribed in advance, ω_i^2 and $\alpha^{(i)}$ are assumed available from prior computation, and $\delta\omega_i^2$ is what we want to compute. This result is correct through terms of first order in $\delta\mathbf{K}$ and $\delta\mathbf{E}$.

Eq. (21) *looks* almost self-evident, as though simple differentiation of 8 would do the job, but the remarkable property of 21 is that it does not contain terms in $\delta\alpha^{(i)}$, the perturbation in the eigenvector. (Straightforward differentiation would produce those terms and the thrust of Rayleigh's principle is that in hermitian cases the terms in $\delta\alpha^{(i)}$ must sum to zero.) This leads to a result which is much more tractable and powerful analytically and vastly more efficient computationally.

3.3 Excitation calculations

In order to assess which resonances will be observed in a particular measurement it is *very* helpful to be able to *predict* the response of a sample during a resonance experiment theoretically. Happily, it is straightforward to use the resonance eigenvalues and eigenvectors to make that prediction.

In a resonance experiment, we apply a sinusoidal excitation to some point on the sample, measure its response at some other point, and repeat the process for many frequencies. The calculation we will perform here is the computation of the displacement response of a sample to a sinusoidal point force with a specified polarization and applied at a specified site. The calculation below hinges on the fact that the resonance displacement eigenvectors,

$$v^{(i)}(\mathbf{r}) = \alpha_\lambda^{(i)} \phi_\lambda(\mathbf{r}), \quad \text{for } i = 1, \dots \quad (22)$$

form a complete basis for the space of all finite elastic displacements in \mathcal{V} .

Let

$$\mathbf{f}(\mathbf{r}, t) = \mathbf{f}_0 \delta(\mathbf{r} - \mathbf{r}_s) e^{i\omega t} \quad (23)$$

where \mathbf{f}_0 is a constant vector, \mathbf{r}_s is the location of the applied point force, and ω is its angular frequency. If $\mathbf{s}(\mathbf{r}, t)$ is the response of \mathcal{V} to \mathbf{f} , then we expect \mathbf{s} will also be sinusoidal with frequency ω .

$$\mathbf{s}(\mathbf{r}, t) = \mathbf{s}(\mathbf{r}) e^{i\omega t} \quad (24)$$

Because $v^{(i)}$ are complete, we can find coefficients γ_i such that

$$\mathbf{s}(\mathbf{r}) = \gamma_i v^{(i)}. \quad (25)$$

A standard derivation (Dahlen & Tromp 1998) leads to

$$\gamma_i(\omega) = \frac{v_i(\mathbf{r}_s) \cdot \mathbf{f}_0}{\omega^2 - \omega_i^2}, \quad (26)$$

and thus

$$\mathbf{s}(\mathbf{r}) = \frac{v_i(\mathbf{r}_s) \cdot \mathbf{f}_0}{\omega^2 - \omega_i^2} v^{(i)}(\mathbf{r}). \quad (27)$$

Using (27) we can compute the response at any point \mathbf{r} as a function of the frequency ω of the applied force.

4 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.

4.1 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

$$F = \sum_i w_i (f_i^{(p)} - f_i^{(m)})^2, \quad (28)$$

where $f_i^{(p)}$ are the computed frequencies and $f_i^{(m)}$ 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 (Fletcher 1980). 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 $j(v) = F(m_i + v h_i)$, we find a new model $m_{i+1} = m_i + v_i h_i$. Then the new search direction is given by:

$$h_{i+1} = g_{i+1} + \gamma_i h_i, \quad (29)$$

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

$$\gamma_i = \frac{g_{i+1}^t g_{i+1}}{g_i^t g_i}. \quad (30)$$

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 (30) was derived from that particular case. For further references we refer to Fletcher (1980) and Press *et al.* (1986).

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

$$\partial_p F = \sum_i w_i 2 \partial_p f_i^{(p)} (f_i^{(p)} - f_i^{(m)}), \quad (31)$$

and differentiating eq. (8) with respect to p , composing by a^t on the left and using Γ and E symmetries yield (Migliori & Sarrao 1997)

$$\partial_p(\omega^2) = a^t (\partial_p \Gamma - \omega^2 \partial_p E) a, \quad (32)$$

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 (Migliori & Sarrao 1997).

To perform the non-linear conjugate gradient algorithm we use the C++ COOOL library (CWP Object Oriented Optimization library) (Deng *et al.* 1996a,b). 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 eq. (29) and a new line search is performed.

5 EXPERIMENT DESIGN FOR CYLINDRICAL SAMPLES

In this section we'll use some of the tools we developed above to examine the properties of the elastic resonances of a cylindrical sample, and look at the implications for experiment design. 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*:

density	2.5	gm cm ⁻³
V_P	3.000	km s ⁻¹
V_S	1.400	km s ⁻¹

5.1 General resonance types

In the presence of perfect axisymmetry, that is, a system axisymmetric in both its shape and the distribution and symmetry of its material properties, we can show (but we won't do it here) that the particle displacement for any mode must vary with the cylindrical coordinate θ as $\sin(n\theta)$ for some integer value of n . That result plus some additional symmetry arguments can be used to conclude that every mode of such a sample must fall into one of three classes:

Torsional axisymmetric pure shear motions consisting of rigid rotations of rings of material around the sample axis. The frequencies of these modes depends entirely 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 $\sin(n\theta)$. The axisymmetric modes above represent energy traveling purely up and down the cylinder axis; these non-axisymmetric modes travel along paths that are tilted with respect to that axis.

Although these classes breakdown 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.

5.2 Mode shapes

We have computed the theoretical resonances for a generic soft rock sample with dimensions

height	3.099	cm
radius	0.635	cm

This sample is about 2.5 times as tall as it is wide. Fig. 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 order of atomic radii.)

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 apparent midriff bulge in the figure is an artifact of the plotting algorithm.) 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.

5.3 Experiment design

5.3.1 Transducer location

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 Q values decrease, however, it becomes quickly difficult, and soon impossible, to determine all or even most resonances below some upper limit.

Fig. 2 shows the theoretical resonance response for our elongate generic soft rock sample featured above using eq. (27). The computation assumes that the source is a point force centered on the sample's end and that the receiver is a displacement sensor at the other end. We extended our above results in a straightforward way to account for damping in the sample, and we show results for three assumed values of compressional and shear Q. The three cases represent, by geological standards, very high, moderate, and low (but still plausible) Q regimes.

It's clear that we cannot depend upon observing a complete mode catalogue in the presence of even moderate attenuation. The good news, which is very important, is that we can use these computed results to determine which resonances *are* observable for a given source and receiver combination.

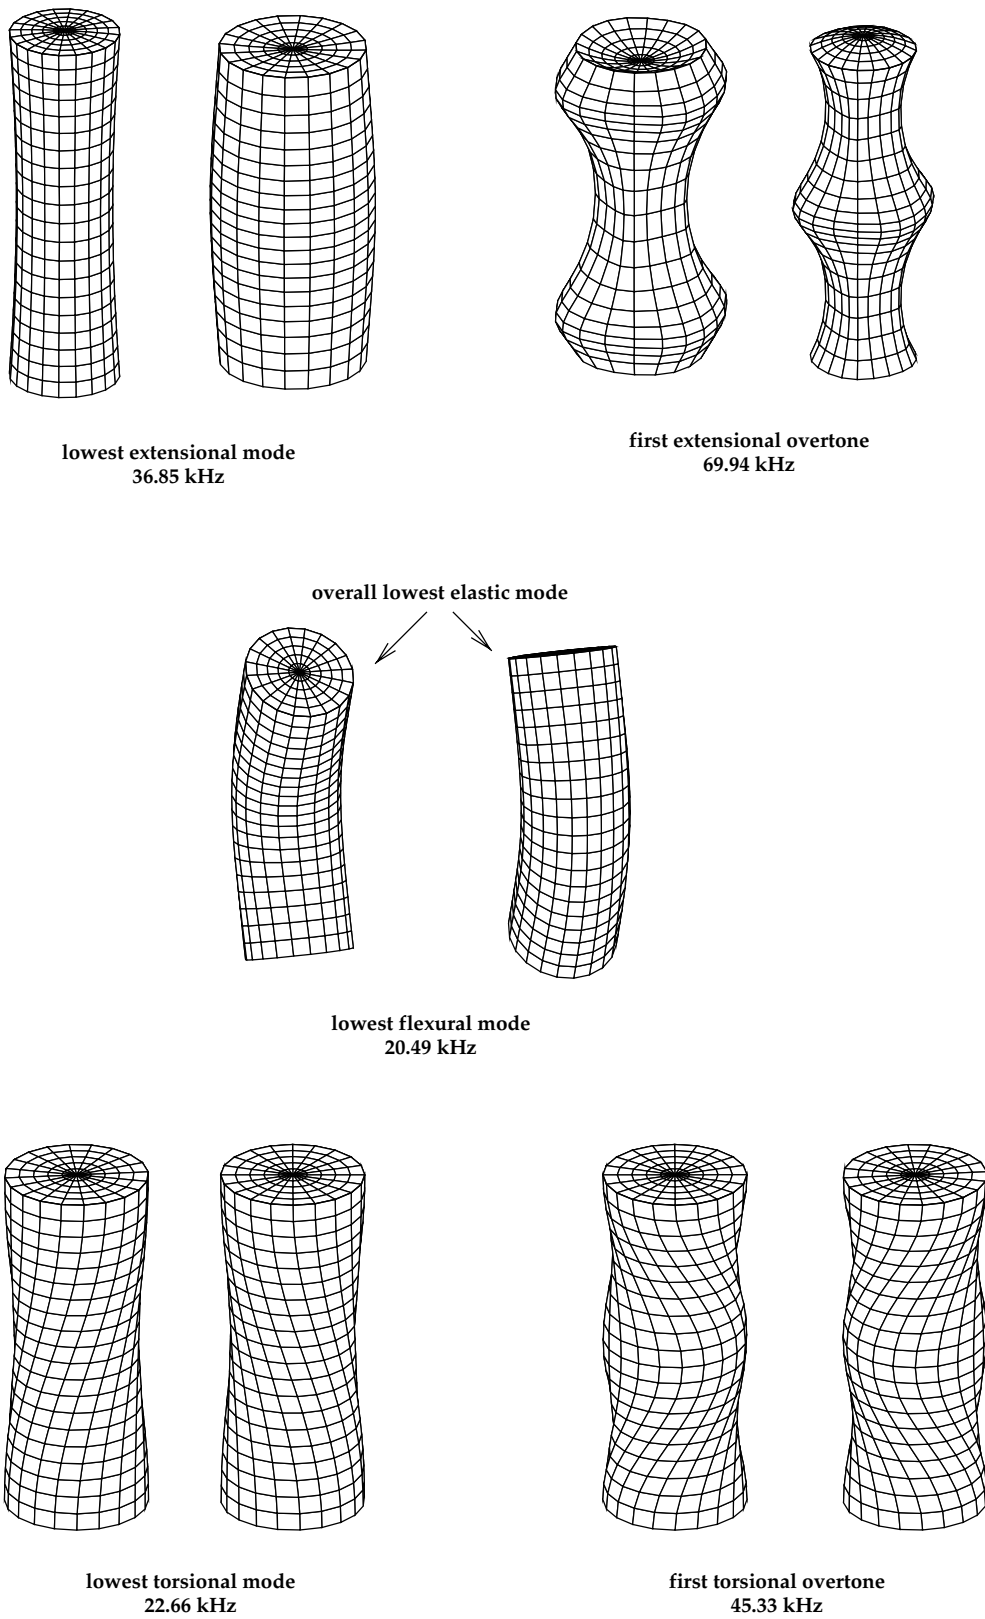


Figure 1. Surface particle displacements of several characteristic modes of the generic soft rock model.




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 eq. (33), for the lowest-frequency modes of the generic soft rock model.

5.3.2 Resonance sensitivities

Fig. 3 shows the relative sensitivity of each mode in Fig. 2 to changes in compressional velocity *versus* shear velocity for our generic soft rock sample. The quantity shown is

$$R_p = \frac{\partial_a f}{\partial_a f + \partial_b f}, \quad (33)$$

where $\partial_a f$ and $\partial_b f$ are the sensitivities of each eigenfrequency to perturbations in V_p and V_s , respectively, computed from Rayleigh's principle, eq. (21). R_p is a measure of the relative importance of the sample's compressional and shear speeds in determining a given mode's frequency. Both $\partial_a f$ and $\partial_b f$ are non-negative (see Rayleigh 1896, 1945 re-issue on 'the march of periods') and so $0 \leq R_p \leq 1$.

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 clearly many modes for which $R_p \ll 1$.




Figure 4. R_p defined in eq. (33) for the generic soft rock model with an inset drawing showing the sample's proportions. The principal plot is the same as that in Fig. 3 but covering a broader frequency range.

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 when V_p is held constant. The elastic cylinder has no modes for which this is true (although it may become asymptotically true for certain families of modes as frequency becomes infinite). It is generally easier to estimate V_s than V_p with resonance methods, so 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 of R_p in Fig. 3 are all depressingly 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 . The results are for a fairly elongate specimen, and this insensitivity to V_p for the lower-frequency modes of such a system is characteristic.

Fig. 4 shows R_p out to about 300 kHz. In this figure we group the modes according to the value of a numerical attribute, L , for each resonance which approximates the value of the *quantum number* n used to classify modes earlier into axisymmetric ($n = 0$) and non-axisymmetric ($n > 0$) families. L is actually the integral of a somewhat messy quadratic function of the displacement eigenfunction which was chosen to yield stable values that converged to $L = n$ when the model was truly axisymmetric.

L is a kind of *quasi-quantum number* which was constructed specifically to function in a cylindrical geometry. We have found it to be a very useful tool in two important functions. First, it partitions the initially-undifferentiated modes enumerated by the eigenvalue calculation into dynamically-related families for study, a role that is most evident in the partition between axisymmetric ($L = 0$) and non-axisymmetric ($L \neq 0$) modes. (Note that an axisymmetric *model* has both axisymmetric and non-axisymmetric *modes*.) This classification scheme is also key to making the inverse calculation more robust because it allows us to track resonances by dynamic identity instead of relying on the rapidly changing location of individual modes in a sorted list of eigenfrequencies. In the inverse calculation we identify a mode by its value of L and by its position in the sequence of modes having that value of L .

Our result is specific to nearly-axisymmetric cylinders but the quantum numbers of various types exist for most highly symmetric systems. It should with modest effort be possible to use this technique in other geometries.

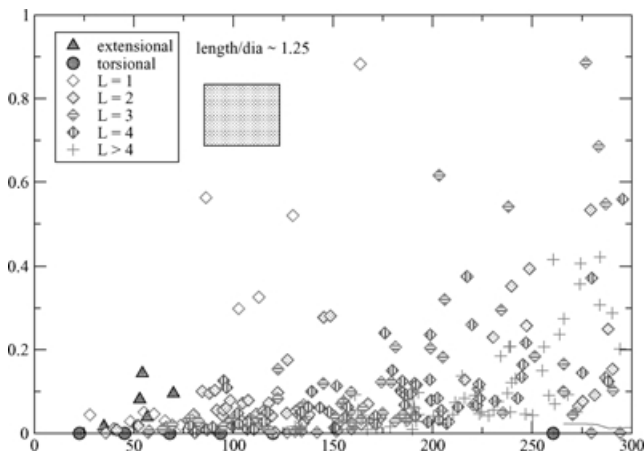


Figure 5. R_p defined in eq. (33) for the generic soft rock model with an inset drawing showing the sample's proportions. Similar to Fig. 4 (preceding), but computed for a model with the same length and *twice* the diameter as the generic soft rock model used for the preceding figure.




Figure 6. R_p defined in eq. (33) for the generic soft rock model with an inset drawing showing the sample's proportions. Similar to Fig. 4 (preceding), but computed for a model with the same length and *four* times the diameter as the generic soft rock model used in Fig. 4.

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 will surely 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. These obstacles are probably insuperable for moderate- Q materials.

We can do better without going to higher frequency, however, by changing the sample geometry. Fig. 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.

Fig. 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




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.

mentioned above, the increased crowding quickly complicates mode identification problems, particular in the presence of dissipation.

5.4 Selective excitation

If we could selectively observe a single, arbitrary mode, we would have achieved resonance paradise. We don't know how to do that but we can look at the role transducer location plays in determining which resonances are observed. Fig. 2 showed that moderate amounts of attenuation could 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 favour particular modes.

Fig. 7 shows the computed response for the our standard sample 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.)

Fig. 8 reproduces the two curves of Fig. 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 tell 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.

6 IDENTIFICATION AND ESTIMATION OF NORMAL MODE FREQUENCIES

6.1 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 (Breit & Wigner 1936). 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.

6.2 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 Fig. 9.

The model we used to fit the amplitudes of one or multiple modes is due to Briet and Wigner (Breit & Wigner 1936):




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

$$A(f) = B_0 + B_1(f - f_0) + \sum_n \frac{C_n + D_n(f - f_0)}{(f - f_n)^2 + \frac{1}{4}\Gamma_n^2} \quad (34)$$

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 affects 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 favour 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 the authors’ knowledge, there is no fool-proof solution for this problem. However, if repeated measurements 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.

6.3 Normal mode identification

There are a number of reasons why a measured spectrum might not be mapped in a reasonable fashion to a modelled 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 modelled 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. Fig. 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 done on a cylinder of Elberton granite in Section 8.2. Although remounting the sample multiple times to estimate error is quite feasible for bench-top measurements, it is too time consuming for measurements involving complex sample preparation.

6.4 Resonance inversion

Missing a normal mode during measurement can be fatal to the inversion if a large 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.* eq. 28). 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 small 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.

7 DATA ACQUISITION

A function generator (Stanford DS345) sends a 10 V (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. Fig. 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 3 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 EXPERIMENTAL RESULTS

We have applied RUS to a variety of samples, both man-made and natural, and both isotropic and anisotropic. Here, for simplicity we focus on two particular samples to illustrate the method. Macor is machinable glass ceramic mainly used in electronic applications as a thermal insulator. The density and modulus of elasticity are similar to that of aluminum. The natural material we chose to show here is Elberton granite, quarried in the state of Georgia, and is predominantly composed of feldspar, and quartz. Previous research has shown that some granites are anisotropic (Douglass & Voight 1969). This seemed an ideal material to test for low symmetries.

8.1 Benchtop macor

A macor cylinder 3.095 cm long and 1.274 cm in diameter was measured in vacuum after being dried. Fig. 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:

f_{obs} (kHz)	f_{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.




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.

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

	V_P (m s ⁻¹)	V_S (m s ⁻¹)
resonance	5655	3220
ultrasonic	5580	3162

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 conservative guess.

There is 1.3 per cent discrepancy between the ultrasonic and resonance values for V_P and a 1.8 per cent discrepancy for V_S . Since the resonance data depend only weakly on V_P , the nearly 2 per cent 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 accurately. However, in the 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.

8.2 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 (Douglass & Voight 1969). Previous work suggests that to be able to accurately measure anisotropy in a sample we would need at least 20 consecutive resonant frequencies (Ulrich *et al.* 2002). We studied a 71 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 Fig. 11. The laser ultrasonic data in Fig. 12 was collected on a similar core. 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 Fig. 13. Frequency uncertainty ranged from 101 Hz to 559 Hz with a mean uncertainty of 246 Hz. Although the amplitudes in Fig. 13 seemed to vary greatly for specific modes, the Q values of the modes were fairly stable. In Fig. 14, the Q's range from approximately 100 to 200, with an average uncertainty of 41.




Figure 11. Spectrum of a cylindrical granite core (2.527 cm diameter by 7.15 cm long). A sweep from 20 to 100 kHz with a 5Hz step was applied.

Table 2. Comparison between the observed ($f^{(obs)}$) and predicted ($f^{(pre)}$) normal mode frequencies with uncertainties for an Elberton granite core. $\sigma^{(obs)}$ is the uncertainty in σ from repeated measurements, $(\frac{f^{(pre)} - f^{(obs)}}{\sigma^{(obs)}})^2$ is the relative contribution of each $f^{(obs)}$ to the total χ^2 , and *freq. #* is how we labelled the identified frequencies.

$f^{(obs)}$ (Hz)	$f^{(pre)}$ (Hz)	$\sigma^{(obs)}$ (Hz)	$(\frac{f^{(pre)} - f^{(obs)}}{\sigma^{(obs)}})^2$	freq. #
15878	15506	79	1.17	1
16026	15793	69	0.60	2
32564	32906	69	1.29	5
33701	34627	31	46.99	6
50031	51316	29	102.98	8
59560	59125	22	20.49	10
62017	62240	27	3.39	11
64331	63586	40	18.67	12
69799	70819	73	10.29	13
71461	71319	769	0.00	14
72553	72440	470	0.00	15
75746	75410	395	0.04	16
76353	76314	402	0.00	17
78331	77323	406	0.34	18
78969	78990	401	0.00	19
79658	79706	406	0.00	20
81338	82327	1084	0.04	21
84621	84717	360	0.00	22
85449	85152	528	0.02	23
86685	85477	331	0.70	24
87228	87594	194	0.19	25

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 207 (as opposed to 500 for a hexagonal model), most of the variance being associated with 5 modes. The comparison between predicted and measured frequencies can be found in Table 2. The inverted model is $c_{11} = 80.7 \pm 1.7$ GPa, $c_{22} =$




Figure 12. Laser ultrasonic measurement of acoustic waves in 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. The arrival after 10 μs is the direct P -wave arrival travelling straight through the sample. The 30 μs arrivals are the counter-rotating surface waves. For more details see (Scales & Malcolm 2003).




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 246 Hz. This uncertainty does include error from misidentification of resonance peaks during fitting.




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

68.1 ± 1.3 GPa, $c_{33} = 74.5 \pm 1.5$ GPa, $c_{23} = 25.6 \pm 1.3$ GPa, $c_{13} = 42.0 \pm 2.2$ GPa, $c_{12} = 37.5 \pm 0.1$ GPa, $c_{44} = 19.89 \pm 0.02$ GPa, $c_{55} = 28.16 \pm 0.01$ GPa and $c_{66} = 28.66 \pm 0.01$ GPa (where the error bars come from the inversion of all 14 data sets). In the plane perpendicular to the symmetry axis, this yields qP wave speeds of 5534 ± 57 m s^{-1} and 5082 ± 50 m s^{-1} in the c_{11} and c_{22} directions, respectively. This nearly 10 per cent P -wave anisotropy is consistent with the ultrasonic measurements shown in Fig. 12 (Scales & Malcolm 2003).

In the anisotropic symmetry plane the S wave speeds are 3298 ± 3 m s^{-1} for the vertical polarization, 2747 ± 16 m s^{-1} and 3269 ± 8 m s^{-1} 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 (Douglass & Voight 1969). The way it is extracted and cut in quarries takes advantage of that anisotropy. This ensured that the anisotropic symmetry axis coincides with the geometrical symmetry axis of the sample. Otherwise, we would have to invert for the crystallographic axis as well.

Of course, going to monoclinic symmetry would reduce our χ^2 , but we have independent measurements indicating weak inhomogeneity at low frequencies. Fig. 12 shows clear evidence of P -wave anisotropy on the order of 10 per cent. This anisotropy is largely due to stress-relaxation micro-fracturing that occurs when the rock is quarried. Since the surface waves all travel around the same path, they are not influenced by the anisotropy in this measurement. However, one can also see splitting or de-phasing of the counter-rotating surface waves caused by heterogeneities along the path. The splitting corresponds to velocity heterogeneity of around 5 per cent. We conclude then, that introducing more degrees of freedom is likely not justified by the data.

9 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.

The method may be applied to samples having small-scale heterogeneities. We have also analyzed the sensitivity of the resulting elastic moduli to transducer position, sample aspect ratio and frequency of the measurement. We have, so far, only studied clean, relatively 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).

ACKNOWLEDGMENTS

We thank Lydia H. Deng for her help with the COOL 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), the members of the consortium project on Seismic Imaging of Complex Structures at the Center for Wave Phenomena and the sponsors of the Fluids III Consortium at the Center for Rock Abuse, Colorado School of Mines.

REFERENCES

- Aki, K. & Richards, P.G., 1980. *Quantitative Seismology*, chap. 2, Freeman and company, San Francisco.
- Anderson, E. *et al.*, 1999. *LAPACK Users' Guide*, 3rd edn, Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Backus, G., 1970. A geometrical picture of anisotropic elastic tensors, *Rev. of Geophys. and Space Phys.*, **8**, 663–371.
- Batzle, M. & Wang, A., 1992. Seismic properties of pore fluids, *Geophys.*, **57**, 1396–1408.
- Breit, G. & Wigner, E., 1936. Capture of slow neutrons, *Physical Review*, **49**, 519.
- Dahlen, F.A. & Tromp, J., 1998. *Theoretical Global Seismology*, Princeton University Press.
- Demarest, H.J., 1969. Cube-resonance method to determine the elastic constants of solids, *J. acoust. Soc. Am.*, **49**, 768–775.
- Deng, H., Gouveia, W. & Scales, J., 1996a. Inverse methods, in *An Object-Oriented Toolbox for Studying Optimization Problems*, pp. 320–330, eds Jacobson, K.M.B.H. & Sibani, P., Springer-Verlag.
- Deng, H., Gouveia, W. & Scales, J., 1996b. The CWP object-oriented optimization library, *The Leading Edge*, **15**, 365–369.
- Douglass, P. & Voight, B., 1969. Anisotropy of granites: a reflection of microscopic fabric, *Geotechnique*, **19**, 376.
- Fletcher, R., 1980. *Practical Methods of Optimization*, Vol. 1, chap. 4, John Wiley & Sons, INC., New York.
- Goldstein, H., 1981. *Classical Mechanics*, 2nd edn, chap. 4-4, Addison-Wesley Publishing Company, Inc., pp. 143–148.
- Gradshteyn, I. & Ryzhik, I., 1965. *Table of Integrals, Series, and Products*, chap. 2.51, Academic Press, New York, pp. 130–131.
- Holland, R., 1968. Resonant properties of piezoelectric ceramic parallelpiped, *J. acoust. Soc. Am.*, **43**, 988–997.
- Lucet, N., Rasolofosaon, P. & Zinszer, B., 1991. Sonic properties of rocks under confining pressure using the resonant bar technique, *J. acoust. Soc. Am.*, **89**, 980–990.
- Mathews, J. & Walker, R.L., 1965. *Mathematical Methods of Physics*, W. A. Benjamin, Inc.
- Maynard, J., 1996. Resonant ultrasound spectroscopy, *Physics Today*, **49**, 26–31.
- Maynard, J.D. & So, J.-H., 1991. Thin film characterization using resonant ultrasound spectroscopy, in *Proceedings of the Resonance Meeting, 1999*, ed. Furr, E.A., National Center for Physical Acoustics.
- McSkimmin, H.J., 1964. Ultrasonic methods for measuring the mechanical properties of liquids and solids, in *Physical Acoustics*, Vol. 1A, ed. Mason, W.P., chap. 4, Academic Press, New York.

- Migliori, A. & Sarrao, J.L., 1997. *Resonant Ultrasound Spectroscopy: Applications to Physics, Materials Measurements, and Non-destructive Evaluation*, John Wiley & Sons, INC., New York.
- Press, W., Flannery, B., Teukolsky, S. & Vetterling, W., 1986. *Numerical recipes*, Cambridge University Press, Cambridge.
- Rayleigh, John William Strutt, B., 1896, 1945 re-issue. *The Theory of Sound*, 2nd edn, Vol. 1, New York: Dover Publications.
- Scales, J. & Malcolm, A., 2003. Laser characterization of ultrasonic wave propagation in random media, *Phys. Rev. E*, **67**, 04 6618.
- Thomsen, L., 1986. Weak elastic anisotropy, *Geophysics*, **51**, 1954–1966.
- Ulrich, T., McCall, K. & Guyer, R., 2002. Determination of elastic moduli of rock samples using resonant ultrasound spectroscopy, *J. acoust. Soc. Am.*, **111**, 1667–1674.
- Visscher, W., Migliori, A., Bell, T. & Reinert, R., 1991. On the normal modes of free vibration of inhomogeneous and anisotropic elastic objects, *J. acoust. Soc. Am.*, **90**, 2154–2162.
- Wilkinson, J.H. & Reinsch, C., 1971. *Linear Algebra*, 1st edn, Springer-Verlag, New York.

APPENDIX A: SYMMETRIES AND BLOCK DIAGONALIZATION

An elastic medium is characterized by its stiffness tensor (C_{ijkl}). With the so-called Voigt notation (Thomsen 1986), one can represent the medium by a 6×6 matrix, (c_{ij}) in accordance with

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

If one restricts 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 is 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 K [see eqs (3) and (4)]. Reordering yields simplifications that can be exploited for numerical efficiency. Consider the entry

$$K_{\lambda i \lambda' k} = C_{ijkl} \int_V \partial_j \phi_\lambda \partial_l \phi'_{\lambda'} dV, \quad (\text{A1})$$

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

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

$$B x^{l+l'-1} y^{m+m'-1} z^{n+n'},$$

where B is some constant. Since we integrate with symmetric limits if such a term is non zero then

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

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

$$B x^{l+l'-1} y^{m+m'} z^{n+n'-1},$$

and if such a term is non zero then

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

If now $i = 1$ and $k = 1$, then the only non-zero terms in the sum in eq. (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

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

We obtain similar results for the cases $(i, k) = (2, 2)$ and $(i, k) = (3, 3)$. There are 8 choices for the parities of \mathbf{l} , \mathbf{m} , and \mathbf{n} . The previous relations give an equivalence relation on the set

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

with 8 equivalence classes. If one takes two representatives of two distinct equivalence classes the associated entry in K is zero. Therefore, by grouping the representative of the same classes we decompose the matrix K into 8 diagonal blocks.

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

For axis definitions, refer to Fig. A1. In the case of symmetries lower than orthorhombic a block diagonalization of the matrices K 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.




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

Note that the rotation matrix involves three additional parameters. These can be chosen to be Euler angles as in Goldstein (1981). 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 cylindrical samples with a circular base the axes of symmetry can only be inverted for up to a rotation around the sample symmetry axis (Backus 1970).

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 proposes 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 dimensions the action of a harmonic tensor $H^{(q)}$ of order q can be written as

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

where \mathbf{r} is the position, $\mathbf{a}^{(v)}$, $1 \leq v \leq q$ 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 $\mathbf{a}^{(v)}$ e.g. $H^{(q)}$ is rotation invariant with respect to some axis if and only if the set of directions $\mathbf{a}^{(v)}$ 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 (17), the case of a rectangular parallelepiped is easiest. We know that if one of \mathbf{l} , \mathbf{m} or \mathbf{n} is odd then \mathcal{I}_V vanishes, otherwise we have

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

B2 Ellipsoidal cylinder

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

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

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 (Gradshteyn & Ryzhik 1965)

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

where

$$(2a)!! = 2a(2a-2)(2a-4)\dots(2) = 2^a a!,$$

$$(2a+1)!! = (2a+1)(2a-1)(2a-3)\dots(1), \text{ and}$$

$$0!! = (-1)!! = 1.$$

which gives

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

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

B3 Triaxial ellipsoid

Here also, \mathcal{I}_V (eq. 17) is only non-zero for \mathbf{l} , \mathbf{m} and \mathbf{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

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

The second integral is equal to (Gradshteyn & Ryzhik 1965)

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

whereas the first is (Gradshteyn & Ryzhik 1965)

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

Therefore, we obtain

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

Note that the form of \mathcal{I}_V shows the spheroidal symmetry of the system. Other geometrical shapes and values of \mathcal{I}_V are given in Visscher *et al.* (1991).

B4 Right-circular cylindrical wedge

This shape is a bite out of a right-circular cylinder which is bounded between planes at heights z_1 and z_2 , between planes at the constant angles θ_1 and θ_2 , and between the cylindrical surfaces of constant radius r_1 and r_2 . This in isolation is a rather bizarre sample shape, but the wedge is a very useful and general component volume for building systems more complex than a single, homogeneous cylinder.

Incorporating the domain limits above yields

$$\mathcal{I}_S = \int_S (r \sin(\theta))^l (r \cos(\theta))^m z^n dv \quad (\text{B1})$$

$$= \int_{r_1}^{r_2} r^{l+m+1} dr \quad (\text{B2})$$

$$\times \int_{\theta_1}^{\theta_2} \sin^l(\theta) \cos^m(\theta) d\theta \int_{z_1}^{z_2} z^n dz \quad (\text{B3})$$

$$= \frac{r^{l+m+2}}{l+m+2} \Big|_{r_1}^{r_2} \frac{z^{n+1}}{n+1} \Big|_{z_1}^{z_2} \mathcal{T}(\theta_1, \theta_2; l, m) \quad (\text{B4})$$

where

$$\mathcal{T}(\theta_1, \theta_2; l, m) = \int_{\theta_1}^{\theta_2} \sin^l(\theta) \cos^m(\theta) d\theta \quad (\text{B5})$$

\mathcal{T} can be computed by a recursive procedure implied in eq. (2.5.10) of (Gradshteyn & Ryzhik 1965).