

Efficient Simulation of Constant Q Using Coarse-Grained Memory Variables

by Steven M. Day

Abstract Improvements in computing speed have progressively increased the usable bandwidth of seismic wave-field simulations computed with time-stepped numerical schemes (e.g., finite difference, finite element, pseudospectral). As computational bandwidth increases, anelastic losses become increasingly significant for some important applications such as earthquake ground-motion modeling, whole earth seismogram simulation, and exploration seismic profile modeling, and these losses need to be included in the simulations. As bandwidth increases, however, the memory variables necessary to incorporate realistic anelastic losses account for an increasing proportion of total computational storage requirements, a consequence of the broad relaxation spectrum of typical earth materials. To reduce these storage requirements, we introduce a new method in which the memory variables are coarse grained, that is, redistributed in such a way that only a single relaxation time is represented at each node point (and therefore a single memory variable per stress component is required). Guided by a perturbation analysis, we effect this redistribution in such a way that spatial variability of this single relaxation time simulates the full relaxation spectrum. Such coarse graining reduces memory-variable storage requirements by a factor of 8 for 3D problems or a factor of 4 for 2D problems.

In fourth-order finite-difference computations for the 3D acoustic-wave equation, the method simulates frequency-independent Q within a 3% tolerance over 2 decades in frequency, and it is highly accurate and free of artifacts over the entire usable bandwidth of the underlying finite-difference scheme. These results should also hold for the elastodynamic equations. The method is readily generalized to approximate specific frequency-dependent Q models such as power laws or to further reduce memory requirements. In its present implementation, the main limitation of the method is that it generates artifacts at wavelengths equal to 4 grid cell dimensions and shorter, which may, in some limited circumstances, overlap the usable bandwidth of very high-order finite-difference and/or pseudospectral schemes.

Introduction

As computing has become progressively cheaper, it has become possible to use discrete numerical methods such as the finite-difference and finite-element methods to synthesize seismic wave fields over spatial domains that are relatively large compared with the minimum numerically resolvable wavelength. When a seismic pulse propagates a distance many times its dominant wavelength, anelastic attenuation cannot generally be ignored. Even three-dimensional elastodynamic problems can now be solved over domains sufficiently large that attenuation should not be neglected, and the problem of modeling attenuation efficiently and accurately will become increasingly important as computational capabilities advance. Accurate treatment of anelastic attenuation has already become a priority in some important practical applications in three dimensions, includ-

ing, for example, the modeling of earthquake strong ground motion in sedimentary basins (e.g., Frankel and Vidale, 1992; Frankel, 1993; Graves, 1993; Yomogida and Etgen, 1993; Olsen *et al.*, 1995a), synthesis of the elastic response of whole earth models (e.g., Yoon and McMechan, 1995; Igel and Weber, 1995), and the simulation of seismic reflection profiles (e.g., Carcione *et al.*, 1992). Applications such as these now frequently entail three-dimensional computations over domains with dimensions roughly two orders of magnitude larger than the minimum wavelength (e.g., Olsen *et al.*, 1995b).

The most general anelastic law represents the stress at (\mathbf{x}, t) as a convolution integral over the complete strain history at \mathbf{x} , a form of the stress-strain relation that is not tractable for numerical calculations, because it imposes enormous storage and computational requirements. Day and

Minster (1984) developed a framework for incorporating anelastic attenuation laws into time-stepped numerical methods such as the finite-difference method by approximating them using internal, or “memory,” variables. This approach can be viewed in the time domain as the replacement of the convolution operator by a low-order differential operator. Alternatively, it can be viewed in the Laplace transform domain as the approximation of an algebraic transfer function by a pole-zero representation. Each pole then corresponds to a memory variable that evolves according to a first-order differential equation analogous to that governing a standard linear solid. The latter can be time stepped to compute the evolution of the memory variables along with the other field variables such as velocity and stress. This framework has been refined and improved by a number of investigators, who have used a variety of methods to determine efficient pole-zero representations of the stress–strain relationship (e.g., Emmerich and Korn, 1987; Carcione *et al.*, 1988, Witte and Richards, 1990; Krebs and Quiroga-Goode, 1994; Blanch *et al.*, 1995).

While the memory variable framework is effective, it does impose stringent demands for computer memory. To approximate an anelastic quality factor (Q) that is approximately constant over a specified frequency band requires 2 to 3 memory variables per decade of bandwidth per stress component per computational unit cell, even under the most highly optimized memory variable formulations (e.g., Robertsson *et al.*, 1994; Blanch *et al.*, 1995). This requirement becomes particularly onerous in three dimensions, where there are six independent stress components per computational unit cell. The typical 3D elastic (nonattenuative) staggered grid finite-difference method requires nine words of storage per unit cell for the field variables, that is, the velocity and stress components. Anelastic computations in three dimensions, in contrast, would require the addition of roughly 30 memory variables per unit cell to achieve a near-constant Q over a bandwidth of 2 decades. Thus, the total memory requirement would be roughly quadrupled. It is therefore not surprising that, for example, none of the 3D simulations of basin response to earthquakes done to date have included anelastic memory variables.

We present a method that approximates constant Q to very high accuracy, over 2 to 3 decades of frequency, with a single memory variable per stress component per unit cell. The method is most efficient in higher dimensions. In two dimensions, the method achieves a fourfold reduction in memory variable storage requirements (to achieve Q constant over a given bandwidth), compared with the most efficient current methods. In three dimensions, the method provides an eightfold reduction. The operations count associated with the memory variable evolution is reduced in the same proportions.

The method is conceptually simple and easily implemented. Given an N -pole description of the anelastic relaxation function, we distribute the N memory variables over the stress node points of the computational grid, one per

stress component per unit cell, in such a way that spatial averages of the relaxation function remain well approximated. This procedure can be viewed as replacement of the microscopic relaxation structure of the medium with a coarse-grained equivalent.

The following section reviews the operational form of the stress–strain relationship of linear anelasticity, its representation in terms of a relaxation spectrum, and its approximation via first-order differential equations governing a set of memory variables. The subsequent section analyzes the effect of coarse graining the memory variables. The analysis is approximate and is intended to demonstrate the plausibility of the technique and to indicate how the evolutionary equations should be weighted in the coarse-grained medium to achieve the desired Q value. We then test the method with a numerical application to constant Q simulation in three dimensions. Finally, we discuss possible extensions, as well as limitations, of the method.

Anelasticity and Memory Variables

This section reviews the reduction of the 1D, anelastic stress–strain relationship to differential form, resulting in its (approximate) expression in terms of N internal, or “memory,” variables. In higher dimensions, there will be one such relationship for each independent stress component. Adhering to the notation of Day and Minster (1984), we write the 1D stress–strain relation as a convolution integral over the step response M :

$$\sigma(t) = \int_0^t M(t - t') d\varepsilon(t'), \quad (1)$$

where σ is the stress and ε is the strain. We apply the Laplace transform in s -multiplied form, that is,

$$\tilde{\sigma}(s) = s \int_0^{\infty} f(t) e^{-st} dt, \quad (2)$$

to transform the stress–strain relation to its operational form:

$$\bar{\sigma}(s) = \bar{M}(s) \bar{\varepsilon}(s). \quad (3)$$

Note that the operational modulus \bar{M} has the same dimensions as the step response M . The unrelaxed modulus M_u , the relaxed modulus M_R , the relaxation of the modulus δM , and the normalized relaxation function ϕ are given by

$$M_u = M(0) = \bar{M}(\infty), \quad (4)$$

$$M_R = M(\infty) = \bar{M}(0), \quad (5)$$

$$\delta M = M_u - M_R, \quad (6)$$

and

$$M(t) = M_R + \delta M \phi(t). \quad (7)$$

We represent the relaxation function in terms of a relaxation spectrum Φ ,

$$\phi(t) = \int_{-\infty}^{\infty} e^{-t/\tau} \Phi(\ln \tau) d(\ln \tau), \quad (8)$$

resulting in the following integral expression for the operational modulus:

$$\bar{M}(s) = M_u - \delta M \int_{-\infty}^{\infty} \frac{\Phi(\ln \tau)}{s\tau + 1} d(\ln \tau), \quad (9)$$

which is equivalent to equation (8) of Day and Minster (1984). We use (9) and the definition of Q^{-1} (as a function of angular frequency ω),

$$Q^{-1}(\omega) = \frac{\text{Im}[\bar{M}(i\omega)]}{\text{Re}[\bar{M}(i\omega)]}, \quad (10)$$

to get, in the low loss approximation ($Q^{-1} \ll 1$),

$$Q^{-1}(\omega) \approx \frac{\delta M}{M_u} \int_{-\infty}^{\infty} \frac{\omega\tau}{\omega^2\tau^2 + 1} \Phi(\ln \tau) d(\ln \tau). \quad (11)$$

To reformulate (9) in differential form, we approximate the integral in (9) by a numerical quadrature, that is, a sum over N discrete relaxation times τ_i ,

$$\int_{-\infty}^{\infty} \frac{\Phi(\ln \tau)}{s\tau + 1} d(\ln \tau) \approx \sum_{i=1}^N \frac{\lambda_i}{s\tau_i + 1}, \quad (12)$$

where the λ_i are the quadrature weights. Next, we substitute the approximation into (3) and invert the Laplace transform, obtaining

$$\sigma(t) = M_u \left[\varepsilon(t) - \sum_{i=1}^N \zeta_i(t) \right]. \quad (13)$$

The ζ_i , $i = 1, \dots, N$, are then internal, or memory, variables that evolve, respectively, according to the N first-order differential equations,

$$\tau_i \frac{d\zeta_i(t)}{dt} + \zeta_i(t) = \lambda_i \frac{\delta M}{M_u} \varepsilon(t), \quad (14)$$

which are specified by the $2N$ values of τ_i and λ_i . The latter are, respectively, the poles and residues of the operational modulus (Day and Minster, 1984). Each stress component, at each computational node, thus has associated with it the N memory variables ζ_i . The resulting low-loss approximation to Q^{-1} is then

$$Q^{-1}(\omega) \approx \frac{\delta M}{M_u} \sum_{i=1}^N \frac{\lambda_i \omega \tau_i}{\omega^2 \tau_i^2 + 1}. \quad (15)$$

The natural criterion on which to judge the approximation (12) is the degree to which the corresponding approximation (15) to $Q^{-1}(\omega)$ matches the exact value (11). Most seismic applications entail approximating a $Q^{-1}(\omega)$ that is nearly constant over a broad frequency range. To accomplish this, Day and Minster (1984) transformed the integration variable in (9) to τ^{-1} and applied Gaussian quadrature to derive analytical expressions for the poles and residues τ_i and λ_i (and showed that this is equivalent to a Padé approximant of the operational modulus). Subsequent investigators have made substantial improvements by using numerical optimization to determine the τ_i and λ_i , usually leading to roughly uniform distribution of the poles over $\ln(\tau)$ (e.g., Emmerich and Korn, 1987). For practical purposes, about 1 pole per octave is sufficient (Robertsson *et al.*, 1994; Blanch *et al.*, 1995) to approximate a frequency-independent Q function.

Coarse Graining of Memory Variables

Outline of the Method

Our objective is to reduce the number of memory variables to at most 1 per stress component per node, without sacrificing accuracy in representing $Q(\omega)$. We do so by introducing a single, spatially variable, relaxation time $\tau(\mathbf{x})$, together with a weighting function $w(\mathbf{x})$, in place of the spectrum of relaxation times in (8). Then, in place of (13), we write

$$\sigma(\mathbf{x}, t) = M_u [\varepsilon(\mathbf{x}, t) - \zeta(\mathbf{x}, t)], \quad (16)$$

where ζ satisfies

$$\tau(\mathbf{x}) \frac{d\zeta(\mathbf{x}, t)}{dt} + \zeta(\mathbf{x}, t) = w(\mathbf{x}) \frac{\delta M}{M_u} \varepsilon(\mathbf{x}, t). \quad (17)$$

We will show by analysis and numerical examples that this method can be made highly successful when the relaxation time function $\tau(\mathbf{x})$ and weight function $w(\mathbf{x})$ satisfy two conditions. First, $\tau(\mathbf{x})$ and $w(\mathbf{x})$ should be constructed so that their spatial distributions approximate the target relaxation spectrum, in the sense described in the next paragraph. Second, the scale length of their spatial variability should be limited to a few times the computational cell dimension, as shown by the perturbation analysis in the next section.

Note that $\Phi(\ln\tau)$ in equation (8) is non-negative and has unit area. We can thus think of $\Phi(\ln\tau)$ as a probability density function, in which case the integral in (9) defining the operational modulus would be the expected value of $(s\tau + 1)^{-1}$ when $\ln\tau$ is sampled from the distribution Φ . It is natural to specify $\tau(\mathbf{x})$ such that it is analogously distributed, in the following sense. Choose $\tau(\mathbf{x})$ to be a measurable function on domain X , and define the set A by

$$A = \{\mathbf{x} : \ln\tau(\mathbf{x}) < a\}, \tag{18}$$

where a is a real number. We introduce a measure ν , such that the measure of a subset E of X is defined by

$$\nu(E) = \int_E w(\mathbf{x})dV, \tag{19}$$

where $w(\mathbf{x})$ is a Lebesgue-measurable, non-negative function having unit mean value on domain X . This means that

$$\nu(X) = m(X), \tag{20}$$

where m is the Lebesgue measure of X (i.e., the volume of the domain of interest). Then we would specify $\tau(\mathbf{x})$ such that the measure of A , $\nu(A)$, satisfies

$$\frac{\nu(A)}{\nu(X)} = \int_{-\infty}^a \Phi(\ln \tau)d(\ln \tau), \tag{21}$$

for all real a . In other words, given a $w(\mathbf{x})$ satisfying (20), we specify $\tau(\mathbf{x})$ such that, if spatial points were sampled randomly, with probability weighted by $w(\mathbf{x})$, the resulting $\ln\tau$ samples would have distribution Φ . In practice, of course, we will only need to specify w and τ at a discrete set of computational nodes.

There are actually two distinct scales of spatial variation implicit in equation (17). The microscale (subwavelength) variation that we have just discussed, and that is designed to simulate the relaxation spectrum, is indicated explicitly by the notation in (17). Equation (17), however, does not explicitly note the spatial variation of the factor $\delta M/M_u$, though it too may vary spatially. This factor, however, scales the macroscopic Q , and its variation will usually have a scale length larger than that of τ and w .

We will refer to the conventional description of anelasticity, as given by equation (9), as the ‘‘relaxation-spectrum’’ representation. We will call the method based on a heterogeneous relaxation time, as given by equations (16) and (17), the ‘‘coarse-grain’’ representation, in accordance with the physical analogy described below.

Physical Analog

Some motivation for this approach and terminology may be found in the physical origin of the relaxation spectra

of earth materials. The characteristic attenuative behavior of rock results from the fact that rock is a heterogeneous aggregate. Relaxation mechanisms corresponding to a broad range of scale lengths and activation energies are present on the microscopic level (see, e.g., Minster, 1980). When the material is deformed over wavelengths large compared to the scale of the heterogeneities, these relaxation mechanisms are activated concurrently, giving rise to a broad spectrum of relaxation times. Conventionally, these relaxations are subsumed into a constitutive equation describing the mean behavior of the aggregate. Our method, in a sense, reverses this synthesis. We revert to the more fundamental, heterogeneous description of the relaxations, albeit in an artificial, coarser-grained aggregate, with grain size equal to the computational unit cell. In this coarse aggregate, each computational unit cell has only a single relaxation time.

In discrete time-stepped numerical wave propagation computations (e.g., finite difference, finite element, pseudo-spectral), the minimum wavelength that can be computed accurately necessarily exceeds several times the cell size of the grid. Thus, it is plausible that we can coarse grain the medium in such a way that its attenuative properties will be preserved for all wavelengths of computational interest. The following analysis of the coarse-grain representation by means of perturbation theory suggests that this is indeed the case.

Analysis of the Method

In the coarse-grain representation of the medium, the effective attenuation can no longer be regarded as a purely local process with properties derivable from constitutive equations alone. To understand the nature of the approximation, it is necessary to introduce the elastodynamic equations into the analysis. We begin with the well-known analysis, via perturbation theory, of anelastic attenuation in terms of an eigenfunction expansion of the elastodynamic field. We simplify the analysis, without sacrificing any useful insights, by restricting consideration to wave fields in a bounded volume (and by neglecting issues raised by eigenfrequency degeneracy).

In a finite volume V with homogeneous boundary conditions, all solutions to the elastodynamic equations (without anelastic losses) can be written as linear combinations of $\mathbf{u}_n(\mathbf{x})e^{i\omega_n t}$, $n = 1, 2, \dots$, where \mathbf{u}_n is the (normalized) eigenfunction belonging to the eigenvalue ω_n^2 of the operator H^0 given by

$$H_i^0(\mathbf{u}) = \rho^{-1} (C_{ipq}^0 u_{p,q})_{,j} \tag{22}$$

In (22), $\mathbf{u}(\mathbf{x})e^{i\omega t}$ is the displacement field, $\rho(\mathbf{x})$ the density, and $C^0(\mathbf{x})$ the elastic tensor. The eigenvalues are positive and the eigenvectors are orthogonal under the inner product defined by

$$(\mathbf{f}, \mathbf{g}) = \int_V \mathbf{f}(\mathbf{x}) \cdot \mathbf{g}^*(\mathbf{x})\rho(\mathbf{x})dV. \tag{23}$$

The introduction of weak, anelastic attenuation has the effect of adding a small additional, frequency-dependent term H' to H^0 . The new operator H is then

$$H = H^0 + H', \quad (24)$$

where

$$H'_i(\mathbf{u}) = \rho^{-1} (C'_{ijpq} u_{p,q})_{,j}. \quad (25)$$

Each component of C' has the form of the second term on the right-hand side of (9), evaluated at $s = i\omega$. The resulting perturbation to the eigenvalue ω_n^2 can be estimated from perturbation theory. The first-order correction $\delta\omega_n^2$ to the eigenvalue is

$$\delta\omega_n^2 = (H'(\mathbf{u}_n), \mathbf{u}_n), \quad (26)$$

where \mathbf{u}_n is the normalized, unperturbed eigenfunction of H^0 . The temporal Q^{-1} (Aki and Richards, 1980) of the mode, Q_n^{-1} , can be expressed in terms of $\delta\omega_n$, to first order in the components of C' , as

$$\begin{aligned} Q_n^{-1} &= 2\omega_n^{-1} \text{Im}(\delta\omega_n), \\ &= \omega_n^{-2} \text{Im}(H'(\mathbf{u}_n), \mathbf{u}_n). \end{aligned} \quad (27)$$

To simplify the subsequent discussion, we now specialize to the case of an isotropic medium, though this assumption is not essential to the analysis. To begin with, we consider the relaxation-spectrum representation of the medium. We take a normalized relaxation spectrum Φ that is independent of position \mathbf{x} (though this assumption is not essential either). Then we can find H' from (9), replacing M and δM by their bulk and shear modulus equivalents:

$$\begin{aligned} H'_i(\mathbf{u}) &= -\rho^{-1} (\delta\kappa \varepsilon_{kk} \delta_{ij} \\ &+ 2\delta\mu \varepsilon'_{ij})_{,j} \int_{-\infty}^{\infty} \frac{\Phi(\ln \tau)}{1 + i\omega\tau} d(\ln \tau), \end{aligned} \quad (28)$$

where ε is the infinitesimal strain tensor and ε' its deviatoric part, and κ and μ denote, respectively, bulk and shear moduli. Then, substituting (28) in (27), and applying the divergence theorem and the boundary conditions, we obtain

$$\begin{aligned} Q_n^{-1} &= \omega_n^{-2} \int_{-\infty}^{\infty} D_{\omega_n}(\tau) \Phi(\ln \tau) d(\ln \tau) \int_V \left[\left(\frac{\delta\kappa}{\kappa_u} \right) \kappa_u \varepsilon_{(n)ij} \varepsilon_{(n)kk}^* \right. \\ &\quad \left. + \left(\frac{\delta\mu}{\mu_u} \right) 2\mu_u \varepsilon'_{(n)ijk} \varepsilon'_{(n)jk} \right] dV, \end{aligned} \quad (29)$$

where ε_n and ε'_n are strain tensors associated with the eigenvector \mathbf{u}_n , and $D_{\omega}(\tau)$ is defined by

$$D_{\omega}(\tau) \equiv \frac{\tau\omega}{1 + \tau^2\omega^2}. \quad (30)$$

Recalling our reinterpretation of $\Phi(\ln \tau)$ as a probability density function, we can interpret the first integral in (29) as the expected value of $D_{\omega_n}(\tau)$. Denoting this expected value by $\langle D_{\omega_n} \rangle$, and also introducing $E_{n\kappa}$ and $E_{n\mu}$ to denote, respectively, the bulk and shear strain energy densities of mode n , we can write (29) in the form

$$\begin{aligned} Q_n^{-1} &= \omega_n^{-2} \int_V \left(2 \frac{\delta\kappa}{\kappa_u} E_{n\kappa} \langle D_{\omega_n} \rangle \right. \\ &\quad \left. + 4 \frac{\delta\mu}{\mu_u} E_{n\mu} \langle D_{\omega_n} \rangle \right) dV. \end{aligned} \quad (31)$$

Now we are ready to introduce the coarse-grain representation, as given by equations (16) to (21). Therefore, we drop the integral over the relaxation spectrum, that is, $\langle D_{\omega_n} \rangle$, in our representation of H' , and instead allow the relaxation time τ to be a function of position, $\tau(\mathbf{x})$, as described earlier. The result is that $\langle D_{\omega_n} \rangle$ in (31) is replaced by $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$. We denote by \tilde{Q}_n the resulting approximation to the temporal Q of mode n . Then, \tilde{Q}_n^{-1} is

$$\begin{aligned} \tilde{Q}_n^{-1} &= \omega_n^{-2} \int_V \left(2 \frac{\delta\kappa}{\kappa_u} E_{n\kappa} D_{\omega_n}[\tau(\mathbf{x})] \right. \\ &\quad \left. + 4 \frac{\delta\mu}{\mu_u} E_{n\mu} D_{\omega_n}[\tau(\mathbf{x})] \right) w(\mathbf{x}) dV. \end{aligned} \quad (32)$$

We can assess the coarse-grain representation by evaluating the accuracy of (32) as an approximation to (31), as the latter is the relaxation-spectrum representation of the eigenfunction attenuation. The comparison is more informative with the further simplification that $\delta\kappa/\kappa_u$ and $\delta\mu/\mu_u$ are spatially uniform. From its method of construction, we know that the volume average of $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$ is also equal to the expected value $\langle D_{\omega_n} \rangle$ (see equations 18 to 21). We write $\bar{E}_{n\kappa}$ and $\bar{E}_{n\mu}$ for the corresponding volume averages of $E_{n\kappa}$ and $E_{n\mu}$, and (31) simplifies to

$$Q_n^{-1} = 2\omega_n^{-2} V \left(\frac{\delta\kappa}{\kappa_u} \bar{E}_{n\kappa} + \frac{\delta\mu}{\mu_u} \bar{E}_{n\mu} \right) \langle D_{\omega_n} \rangle. \quad (33)$$

The corresponding expression for attenuation of the eigenfunction in the coarse-grain representation is

$$\begin{aligned} \tilde{Q}_n^{-1} &= 2\omega_n^{-2} \left\{ \frac{\delta\kappa}{\kappa_u} \int_V E_{n\kappa}(\mathbf{x}) D_{\omega_n}[\tau(\mathbf{x})] w(\mathbf{x}) dV \right. \\ &\quad \left. + 2 \frac{\delta\mu}{\mu_u} \int_V E_{n\mu}(\mathbf{x}) D_{\omega_n}[\tau(\mathbf{x})] w(\mathbf{x}) dV \right\}, \\ &= Q_n^{-1} + 2\omega_n^{-2} \left(\frac{\delta\kappa}{\kappa_u} C_{\kappa} + 2 \frac{\delta\mu}{\mu_u} C_{\mu} \right). \end{aligned} \quad (34)$$

The C_κ and C_μ in (34) are just the spatial correlations of $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$ with the bulk and shear strain energy densities, respectively:

$$C_M = \int_V [E_{nM}(\mathbf{x}) - \bar{E}_{nM}] \{w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})] - \overline{wD_{\omega_n}}\} dV, \quad (35)$$

where M refers to either κ or μ and $\overline{wD_{\omega_n}}$ is the volume average of $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$. We have used the equality between $\overline{wD_{\omega_n}}$ and $\langle D_{\omega_n} \rangle$ in deriving (34) and (35). Equation 35 states that the error introduced by the coarse-grain representation, relative to the relaxation-spectrum representation, is proportional to the degree of spatial correlation between the modal strain energy density and the function $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$.

We should be able to minimize the correlation of $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$ with all modes of practical interest by making $\tau(\mathbf{x})$ and $w(\mathbf{x})$ periodic in space, with wavelength short compared with the minimum wavelength at which the numerical wave propagation computation (e.g., finite difference, etc.) is considered accurate. We assess this strategy by examining (34) and (35) in the simplified case of plane wave excitation of a uniform, 1D acoustic medium.

In this special case, the eigenfunction has the form

$$u_n(\mathbf{x}) = B \cos(\omega_n x/\alpha) \quad (36)$$

(apart from an irrelevant phase shift), where α is the wave-speed and B is a constant. The deviation of the strain energy density from its mean value of $(1/4) \kappa_u (\omega_n B/\alpha)^2$ is

$$E_{nc}(\mathbf{x}) - \bar{E}_{nk} = -\frac{1}{4} \kappa_u (\omega_n B/\alpha)^2 \cos(2\omega_n x/\alpha), \quad (37)$$

which has $1/2$ the wavelength of the eigenfunction. We will assume that $w(\mathbf{x})$ and $\tau(\mathbf{x})$ are periodic, with period L ; that is, $\tau(x) = \tau(x-mL)$ for integer m , and similarly for w . Then the function $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$ is also periodic with period L and therefore has a Fourier series representation, the zeroth-order term of which is just $\overline{wD_{\omega_n}}$, so that the deviation of $w(\mathbf{x})D_{\omega_n}[\tau(\mathbf{x})]$ from its mean is

$$w(\mathbf{x})D_{\omega_n}[\tau(x)] - \overline{wD_{\omega_n}} = \sum_{j=1}^{\infty} a_j \cos(2\pi jx/L) + \sum_{j=1}^{\infty} b_j \sin(2\pi jx/L). \quad (38)$$

Combining (37) and (38) with (35), we find that the correlation C_κ is only appreciable when the condition

$$j \left(\frac{2\pi\alpha}{\omega_n} \right) = 2L, \quad j \geq 1, \quad (39)$$

is satisfied, where $2\pi\alpha/\omega_n$ is the wavelength. Therefore, the largest value of wavelength for which the error in the coarse-grain representation is significant is $2L$. In other words, the estimated error is only significant when the wavelength is less than or equal to twice the periodicity of the relaxation time function $\tau(x)$. Note, by the way, that (39) is just the Bragg condition for backscatter at normal incidence to the scattering plane.

Implementation

Equation (39) implies a bandwidth limitation in the coarse-grain method. Our strategy for its implementation is to match this short-wavelength cutoff to the short-wavelength limit already present in discrete numerical methods. We make $w(\mathbf{x})$ and $\tau(x)$ periodic, with period equal to twice the computation unit cell dimension. According to (39), this period-two coarse-graining will not result in significant error for wavelengths longer than 4 unit-cell dimensions. For relatively low-order finite-difference methods, including the widely used fourth-order staggered grid methods (e.g., Levander, 1988), wavelengths shorter than 4 unit-cell dimensions suffer strong numerical dispersion and are therefore already subject to large errors. Thus, any spurious effects of period-two coarse-graining should be isolated at wavelengths that are already outside the usable bandwidth of the wave propagation calculations. High-order finite-difference methods and pseudospectral methods are less restricted by numerical dispersion, but suffer other limitations associated with the representation of material interfaces (Witte and Richards, 1990) when the wavelength is shorter than 4 times the cell dimension, and the period-two coarse-grain method may turn out to be satisfactory in those applications as well.

In three dimensions, this two-unit-cell periodicity allows for eight different values of τ at grid nodes. To be specific, if this chosen set of discrete relaxation times is τ_k , $k = 1, \dots, 8$, and unit cells are indexed by integers p, q, r , such that stress nodes are centered at $\mathbf{x} = p\mathbf{a}_1 + q\mathbf{a}_2 + r\mathbf{a}_3$, then

$$\tau(\mathbf{x}) = \tau_{1+p \bmod 2 + 2(q \bmod 2) + 4(r \bmod 2)}, \quad (40)$$

and $w(\mathbf{x})$ has a similar discrete representation

$$w(\mathbf{x}) = w_{1+p \bmod 2 + 2(q \bmod 2) + 4(r \bmod 2)}. \quad (41)$$

Because this approach only permits a discrete set of relaxation times (8 in 3D) to be realized, implementation of criterion (21) requires first developing a discrete approximation to the continuous relaxation spectrum; that is,

$$\Phi(\ln \tau) = \sum_{i=1}^N \lambda_i \delta(\ln \tau - \ln \tau_i), \quad (42)$$

as in (12). This is the same step required in the generation of a conventional memory variable representation, and any

of the optimization methods proposed in the literature (e.g., Emmerich and Korn, 1987; Carcione *et al.*, 1988; Witte and Richards, 1990; Blanch *et al.*, 1995) for generating conventional memory variable representations can be used at this point. Once the continuous spectrum is replaced by the discrete spectrum, the conditions defined by (18) to (21) will be satisfied if the w_i are identified with the quadrature weights λ_i but normalized to have unit volumetric mean. Because the normalization of Φ implies that the λ_i sum to 1, we then have

$$w_i = N\lambda_i, \quad (43)$$

where N is 8 in three dimension or 4 in two dimension. Figure 1 illustrates this period-two scheme.

Eight relaxation times should be sufficient to approximate a frequency-independent $Q(\omega)$ with very high accuracy over roughly 3 decades in frequency. In two dimensions, the two-unit-cell periodicity allows for four distinct values of τ , adequate to approximate frequency-independent $Q(\omega)$ over one and a half decades or more. Compared with a conventional memory variable method, that is, one in which the memory variables are not coarse grained, the coarse-grain method with two-unit-cell periodicity can approximate $Q(\omega)$ with the same accuracy, but reduces the memory requirement by a factor of 8 in three dimension and a factor of 4 in two dimension.

Numerical Example

The above analysis is approximate, and a numerical test is necessary to establish the practical accuracy of the coarse-grain method. In this section, we test its period-two implementation for the case of a plane wave in a uniform, 3D acoustic medium, with Q approximately frequency independent. We solve the 3D acoustic equations using a staggered grid method, fourth order in space and second order in time [Graves (1996) describes the corresponding algorithm for the elastodynamic equations]. Equation (17) is time stepped using the formula

$$\begin{aligned} \zeta(\mathbf{x}, t + \delta t) &= \zeta(\mathbf{x}, t)e^{-\delta t/\tau(\mathbf{x})} \\ &+ w(\mathbf{x}) \frac{\delta\kappa}{\kappa_u} \frac{1}{2} [\varepsilon(\mathbf{x}, t) \\ &+ \varepsilon(\mathbf{x}, t + \delta t)] [1 - e^{-\delta t/\tau(\mathbf{x})}]. \end{aligned} \quad (44)$$

A relaxation spectrum, constant on its support interval $(\ln\tau_m, \ln\tau_M)$, yields a good approximation to a frequency-independent Q for $\tau_m^{-1} \ll \omega \ll \tau_M^{-1}$ [this is a change of notation from Day and Minster (1984) in that τ_m and τ_M correspond to their τ_1 and τ_2]. If ω_0 is a reference frequency

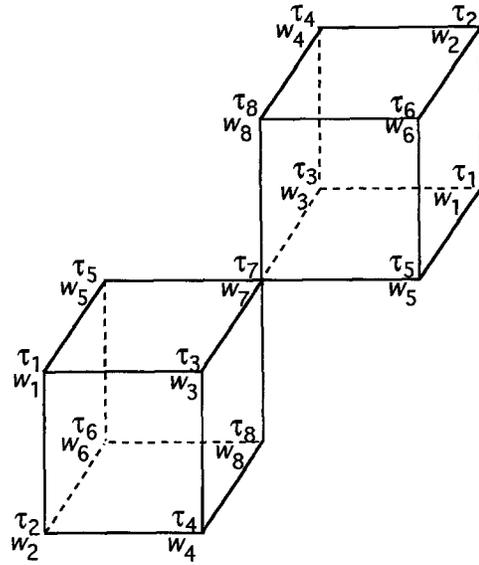


Figure 1. Distribution of relaxation times τ_i and weights w_i in the period-two coarse-grain scheme. Vertices represent computational node points for a given stress component.

near the center of the absorption band, then $Q_0^{-1} \equiv Q^{-1}(\omega_0)$ is given approximately by

$$Q_0^{-1} \approx \frac{\pi\delta\kappa}{2\kappa_u} \left[\ln\left(\frac{\tau_M}{\tau_m}\right) + \frac{\delta\kappa}{\kappa_u} \ln(\omega_0\tau_m) \right]^{-1}. \quad (45)$$

To approximate this absorption band, we take $w(\mathbf{x})$ equal to 1 everywhere and prescribe $\tau(\mathbf{x})$ according to (40), with the $\ln\tau_k$ uniformly spaced over the interval $(\ln\tau_m, \ln\tau_M)$:

$$\ln \tau_k = \ln \tau_m + \frac{2k-1}{16} (\ln \tau_M - \ln \tau_m). \quad (46)$$

The plane wave is given an initial displacement pulse shape $u_0(t)$, where

$$u_0(t) = \frac{2\pi t}{T} e^{-2\pi t/T}, \quad (47)$$

which is a broadband, minimum phase pulse with its spectral corner at frequency $1/T$. In the cases shown below, T is set equal to 10π times the propagation time across the unit cell, so that the corner frequency corresponds to a wavelength of about 30 grid cells. The propagation direction is parallel to one of the lattice vectors of the grid.

We have set the ratio τ_M/τ_m to 10^4 , with the objective of producing a nearly frequency-independent Q for $10\tau_M^{-1} < \omega < 0.1\tau_m^{-1}$. Q_0 is defined at reference frequency ω_0 , which we set equal to the geometrical mean of τ_M^{-1} and τ_m^{-1} . We perform calculations for Q_0 of 100 and 20. The lower cutoff τ_m is set equal to $(2\pi)^{-1}$ times the propagation time across the unit cell.

Figure 2 shows the pulse for the case Q_0 equal to 100, after it has propagated a distance of 26 times the corner wavelength. Figure 2a shows the entire pulse, with the later, low-amplitude parts of the waveform magnified by factors of 100 and 1000. Figure 2b shows the high-amplitude part of the pulse on an expanded timescale. For comparison, both parts of Figure 2 also show the exact solution [based on the frequency-independent Q model of Kjartansson (1979)]. It is evident that the coarse-grain method reproduces the analytical solution with very high accuracy. The small, negative precursor in the numerical solution is a consequence of ordinary numerical dispersion associated with the fourth-order finite-difference method and is unrelated to the anelastic relaxation model.

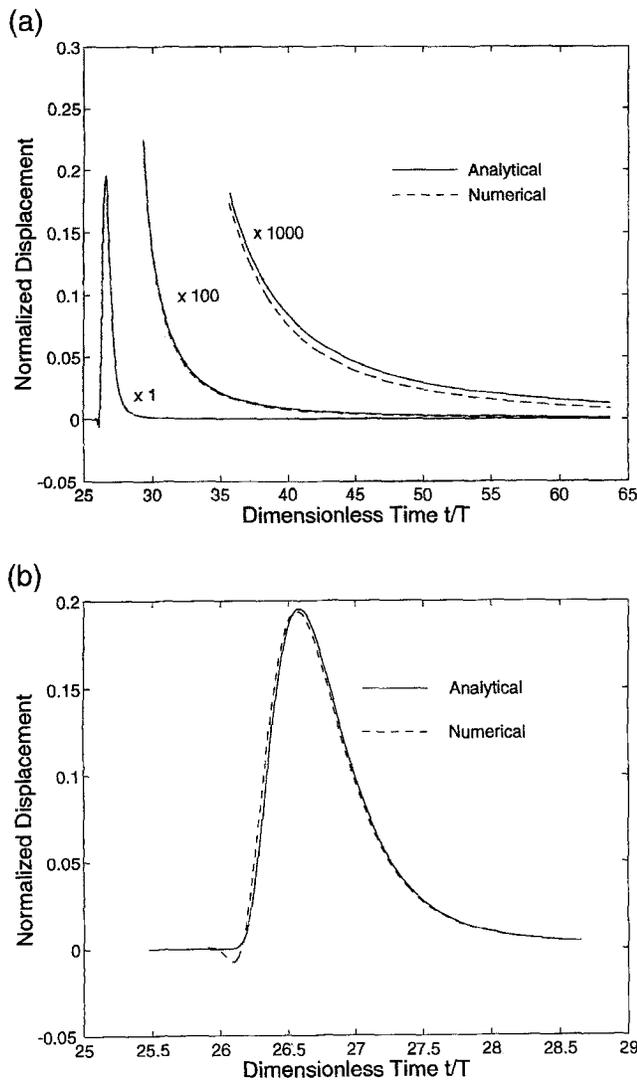


Figure 2. Broadband pulse, for the case of Q_0 equal to 100, after propagating a distance of 26 times its corner wavelength. Part (a) shows the entire pulse, with low-amplitude portions magnified by factors of 100 and 1000. Part (b) shows the initial portion on an expanded timescale. Time has been scaled by the source corner frequency $1/T$.

Figure 3 shows the corresponding results for Q_0 equal to 20. As in the previous case, the numerical and exact solutions are compared in the figure, with the late portion magnified in Figure 3a and the early portion expanded in Figure 3b. The numerical solution reproduces both the attenuation and physical dispersion of the exact solution with very high accuracy. Figure 4 compares the analytical and numerical solutions again, this time after convolution with a Ricker wavelet with dominant frequency f_R equal to twice the corner frequency. Thus, the pulse has propagated 52 times its dominant wavelength. Again, numerical and analytical solutions are virtually indistinguishable, even at late times.

We can determine an apparent Q , denoted by $\hat{Q}(\omega)$, from the numerical solutions. We do this using spectral ratios of the pulse at successive locations; that is,

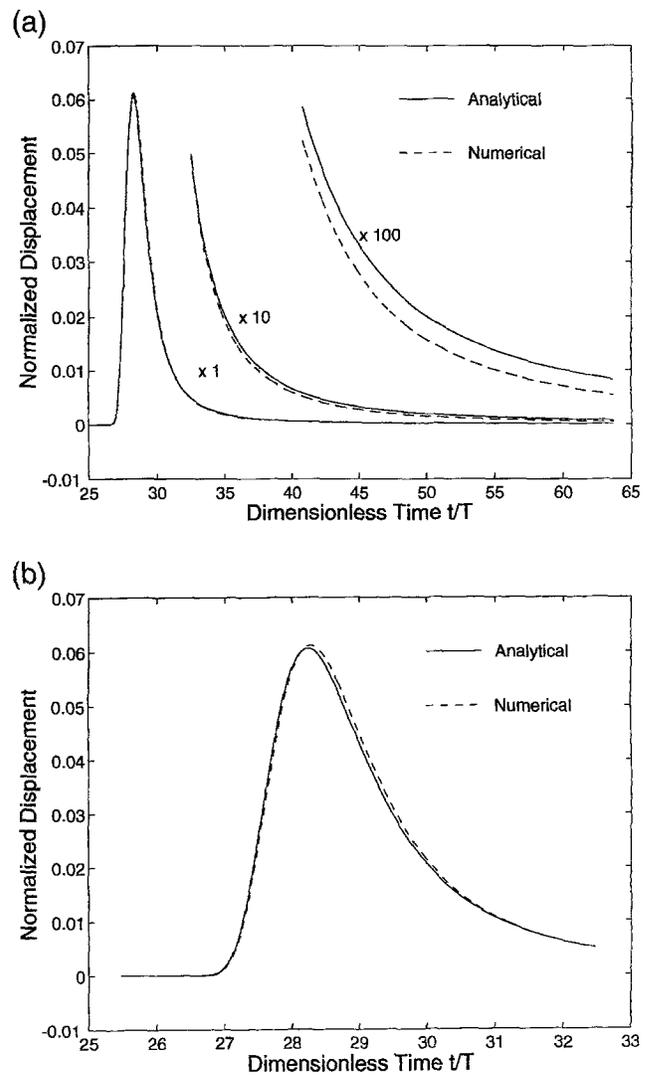


Figure 3. Broadband pulse, for the case of Q_0 equal to 20, after propagating a distance of 26 times its corner wavelength. Part (a) shows the entire pulse, with low-amplitude portions magnified by factors of 10 and 100. Part (b) shows the initial portion on an expanded timescale.

$$\hat{Q}^{-1}(\omega) = \frac{-2 \ln[u(x + \Delta x, \omega)/u(x, \omega)]}{\omega \alpha \Delta x}, \quad (48)$$

where Δx is the propagation distance between locations and α is the wave speed. This operational definition of Q differs from that of equation (10), but at large values of Q , the difference is negligible (see, e.g., O'Connell and Budiansky, 1978). The difference can become significant for small Q , and this difference will become evident in our case when Q_0 equals 20.

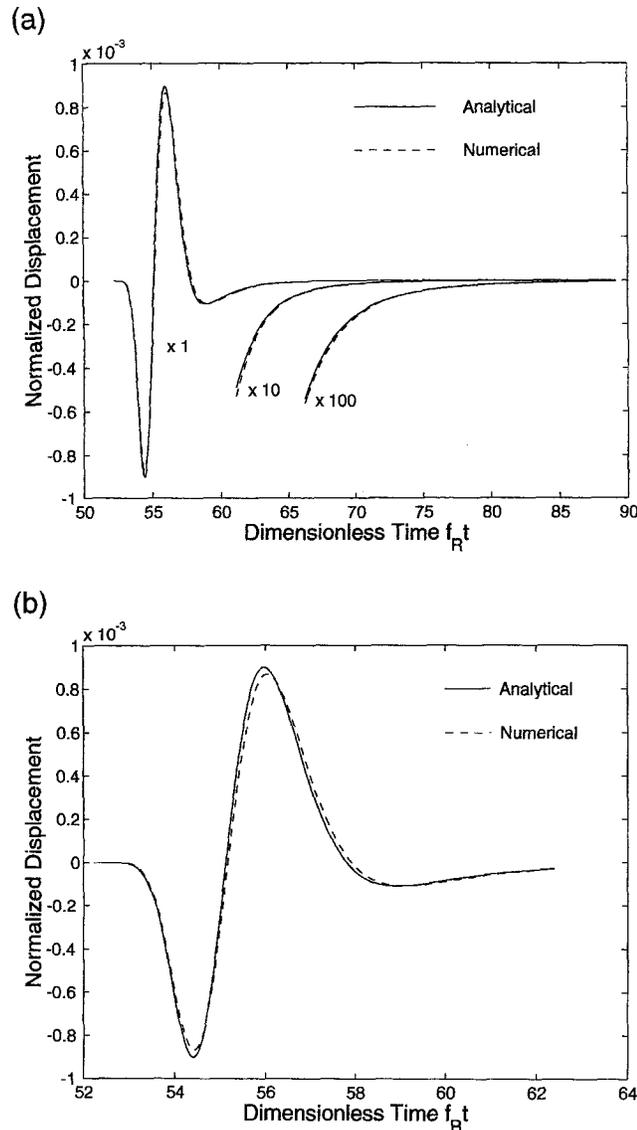


Figure 4. Narrow-band pulse, for the case of Q_0 equal to 20, after it has propagated a distance of 52 times its dominant wavelength. The pulse was constructed by convolving the broadband pulse with a Ricker wavelet. The latter has a dominant frequency f_R equal to twice the corner frequency of the initial broadband pulse. Part (a) shows the entire pulse, with low-amplitude portions magnified by factors of 10 and 100. Part (b) shows the initial portion on an expanded timescale.

Figure 5 shows $\hat{Q}(\omega)$ for the numerical solution in the case $Q_0 = 100$. Four estimates are shown, corresponding to spectral ratios at successive 200 grid-cell intervals. The $\hat{Q}(\omega)$ curves are indistinguishable from one another over the entire range of frequencies from 10^3 grids per wavelength to 4 grids per wavelength. Thus, the coarse-grain method yields a spatially homogeneous attenuative behavior over the full range of usable frequencies present in the finite-difference computations. The variability of the curves at wavelengths longer than 10^3 grid cells is simply a result of the shorter pulse durations available for Fourier analysis at greater distances (due to later arrival times), 10^3 grids per wavelength corresponding to a period equal to the total post-arrival duration of the computation at the most distant location. Likewise, 4 grids per wavelength is well below the practical limit imposed by numerical dispersion in the fourth-order method used here (Levander, 1988).

In addition to indicating a spatially homogeneous behavior, Figure 5 shows that the method is highly successful in producing a frequency-independent $\hat{Q}(\omega)$. The dotted box encloses the usable computational bandwidth and indicates the 6% tolerance interval about a \hat{Q} of 100. For the entire computational bandwidth, \hat{Q} lies within 6% of the target value of 100. Figure 6 shows the corresponding results for $Q_0 = 20$. In this case, the difference in definition between \hat{Q} and Q is significant, and when Q (as defined by equation 10) is frequency independent and equal to 20, \hat{Q} (as defined by equation 48) should be approximately 19, as can be verified analytically using expressions in Kjartansson (1979). As Figure 6 indicates, this target \hat{Q} is again achieved within 6% tolerance over virtually the entire usable computational frequency band. For the 2 decades from 5 grids per wavelength to 500 grids per wavelength, the variation in \hat{Q} is less than $\pm 3\%$ in both the $Q_0 = 20$ and $Q_0 = 100$ cases.

At wavelengths shorter than 4 grid points, as predicted

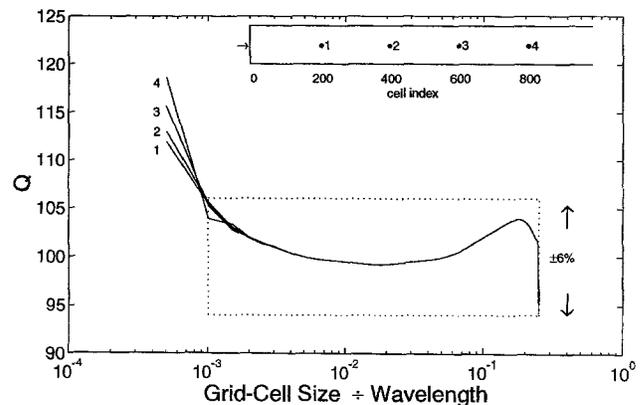


Figure 5. $\hat{Q}(\omega)$ for the numerical solution in the case Q_0 equal to 100. Four estimates are shown, corresponding to spectral ratios at successive 200 grid-cell intervals. The dotted box delimits the usable computational bandwidth and the 6% tolerance interval about a \hat{Q} of 100.

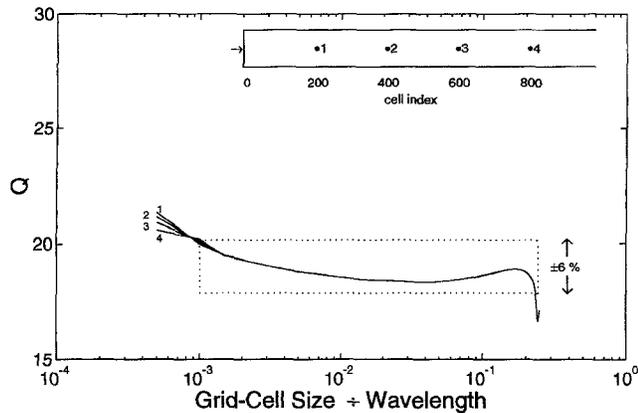


Figure 6. $\hat{Q}(\omega)$ for the numerical solution in the case Q_0 equal to 20. Four estimates are shown, corresponding to spectral ratios at successive 200 grid-cell intervals. The dotted box delimits the usable computational bandwidth and the 6% tolerance interval about a \hat{Q} of 19 (which corresponds to a Q of 20).

by the perturbation analysis, backscatter due to coarse graining adds to the numerical errors that are already present in the form of numerical dispersion. For a low-order finite-difference method (such as the fourth-order method tested here), wavelengths shorter than four grid points are usually considered to lie outside the usable computational bandwidth, as a result of ordinary numerical dispersion alone. Whenever that is the case, the anelastic coarse graining imposes no additional limitation on usable bandwidth. On the other hand, higher-order finite-difference and pseudospectral methods may have sufficient accuracy for some applications at wavelengths shorter than four grid points. When that is the case, the coarse-grain method will not be appropriate.

Discussion

In the numerical example, the period-two coarse-grain method achieved a frequency-independent Q , to high accuracy, over more than 2 decades in frequency, using a single memory variable per stress component per node. This accuracy was achieved without any attempts to optimize the frequency distribution or weighting of the relaxation times in (46). In fact, we did not exploit the flexibility provided by the weight function w at all, simply setting it to 1 in the example.

The weight function can be exploited in several different ways to optimize and/or generalize the coarse-grain method. For example, if further accuracy were desired in the approximation of a frequency-independent Q , any of the optimization methods that have been proposed in the literature for generating conventional memory variable representations could be used in conjunction with period-two coarse graining (e.g., Emmerich and Korn, 1987; Carcione *et al.*, 1988; Witte and Richards, 1990; Blanch *et al.*, 1995). The first step in doing so would be the conventional one: choose

relaxation times τ_i and quadrature weights λ_i (eight of each for three dimension, four of each for two dimension) to optimize the approximation of (11) by (15). Then those relaxation times would be identified with the discrete form of $\tau(\mathbf{x})$ (equation 40); the quadrature weights, after being normalized to unit mean, would be identified with the discrete form of $w(\mathbf{x})$ (equation 41).

A second application of the weight function would be to approximate some other desired frequency dependence of Q , a power law, for example. As in the case of frequency-independent Q , any of the memory variable optimization techniques in the literature could be used to find the eight relaxation times and quadrature weights, which are then identified with the coarse-grain relaxation times and (after normalization) weights.

A third application of the weight function would be to further reduce memory requirements in instances where the accuracy possible with eight relaxation times is redundant. In some seismic exploration problems, for example, a fairly narrow-band representation of Q may be acceptable in practice. Blanch *et al.* (1995) have argued that two properly chosen relaxation times provide sufficient constant- Q bandwidth (Q constant within 8% over roughly one decade in frequency) for many practical purposes, especially given the limited problem sizes that are practical with current computers [Robertsson *et al.* (1994) have even found a single relaxation time to provide an adequate approximation in some narrow-band applications in high-resolution reflection seismology]. In cases where the two-relaxation-time approximation is acceptable, for example, we could assign these two values to τ_1 and τ_2 in (40), assign values to w_1 and w_2 in (41) such that $w_1 + w_2$ equals 8 (in three dimension; or 4 in two dimension), and set the remaining w_i to zero. This strategy would eliminate the need to store any memory variables at all at 75% of the nodes in the grid and only 1 per node at the remaining 25%, thereby reducing the memory variable storage by a factor of 8 (or a factor of 4 in two dimension), relative to a conventional 2 memory variable method. Obviously, then, with the aid of the weight function, any conventional memory variable scheme could be coarse grained to achieve an eightfold memory reduction in three dimension, or fourfold reduction in two dimension.

The above numerical example was restricted to the acoustic-wave equation. There is no apparent reason, however, that the method should not work equally well for the elastodynamic case. In fact, it is for the later case that the method is likely to be most useful, because the presence of six stress components in 3D elastodynamics places very large memory requirements on the conventional memory variable method. Furthermore, the advantages of the coarse-grain method will become even more important as higher computing speeds become available. The speed increases will permit 3D elastodynamic simulations to be computed with ever greater usable bandwidth, and, as computational bandwidth increases, memory variables constitute a pro-

gressively larger proportion of the total computational memory requirement.

We have established the viability of the coarse-grain method for the fourth-order staggered-grid finite-difference method, and it should work equally well with other low-order finite-difference or finite-element methods. It has yet to be established, however, whether the coarse-grain method will work effectively with higher-order finite-difference methods or with pseudospectral methods. The main issue is likely to be whether coarse graining entails any loss of usable bandwidth when applied with those numerical methods. The staggered-grid pseudospectral methods, in particular, may be questionable candidates for the coarse-grain method. In pseudospectral methods, the spatial differentiation operators are exact if the medium is homogeneous (and the wave field spatially periodic and sufficiently bandlimited). In that case, the only source of numerical dispersion arises from approximations associated with the time integration. In the case of the staggered-grid pseudospectral methods, numerical dispersion can be kept very small even at the spatial Nyquist wavenumber, that is, down to only two nodes per wavelength. On the other hand, when discrete interfaces or sharp transitions in wave speed are present in the model, three to four nodes per wavelength are required in order to accurately model the reflection and transmission behavior (Witte and Richards, 1990). In cases where the practical limit of the pseudospectral method is determined to be four nodes per wavelength or greater, the coarse-grain method should be effective, though this conjecture will have to be tested numerically. If so, the method will provide the same degree of memory savings that it provides in the case of low-order finite-difference methods.

Summary

We have developed a new method for introducing anelastic attenuation into time-stepped numerical methods for solving the equations of acoustics and elastodynamics. The method is designed to reduce drastically the memory requirements of anelastic modeling. It takes as its starting point an N -pole (i.e., N relaxation time) approximation to the relaxation spectrum. Then, the N corresponding memory variables are distributed over the stress node points of the computational grid, one per stress component per unit cell. An analysis via perturbation theory shows that this single memory variable method will reproduce the attenuative behavior of the N memory variable model, provided the weights assigned each of the N relaxations (corresponding to the residues in the N -pole representation) are rescaled to have unit volumetric mean, and provided the relaxation times are distributed in a periodic array, with period less than half the minimum wavelength to be modeled. We make N equal to 8 for 3D problems and 4 for 2D problems and distribute the memory variables so that, for each stress component, the relaxation time array has a spatial period, in each direction, equal to twice the unit-cell dimension. This period-two

coarse-grain method is accurate for wavelengths exceeding four unit-cell dimensions.

Accuracy of the method has been verified in numerical examples for the 3D acoustic wave equation. The test problems were solved using a fourth-order staggered-grid finite-difference method. In the numerical examples, a frequency-independent Q was achieved within a 3% tolerance over the 2 decades in frequency from 5 grids per wavelength to 500 grids per wavelength, using only a single memory variable per stress component per node. Straightforward generalizations of the method can be used to approximate specific frequency-dependent Q models such as power laws, or to further reduce memory requirements.

We expect the method to work equally well for low-order finite-difference or finite-element solutions in elastodynamics. For higher-order finite-difference methods, as well as for pseudospectral methods, the utility of the coarse-graining method will depend upon the limiting wavelength at which the underlying numerical method has acceptable accuracy. In those numerical methods in which the usable bandwidth is restricted to wavelengths longer than four unit-cell dimensions, the coarse-grain method should be accurate and efficient.

In those cases where it is appropriate, the method reduces storage requirements for memory variables by a factor of 8 for 3D problems and by a factor of 4 for 2D problems. As computer speeds increase, larger problems, and therefore higher computational bandwidths, will become feasible. As computational bandwidth increases, memory variables will account for an increasing proportion of total storage requirements, putting a further premium on the concision achieved by the coarse-grain method.

Acknowledgments

This work was undertaken while the author was a Visiting Scholar at the Institute of Geophysics and Planetary Physics of the University of California, San Diego. The work owes much to discussions with Bernard Minster. He also pointed out the analogy between the numerical method and the physical origin of the relaxation spectrum. Careful reviews of the manuscript by Bernard Minster and Chris Bradley led to significant improvements. The work was supported in part by grants from the University of California Campus-Laboratory Collaboration Program (Project Number 951322), the Air Force Office of Scientific Research (under Grant Number F49620-94-1-0205), and the Southern California Earthquake Center (SCEC). SCEC is funded by NSF Cooperative Agreement EAR-8920136 and USGS Cooperative Agreements 14-08-0001-A0899 and 1434-HQ-97AG01718. The SCEC Contribution Number is 391.

References

- Aki, K. and P. G. Richards (1980). *Quantitative Seismology, Theory and Methods*, W. H. Freeman and Co., New York.
- Blanch, J. O., J. O. A. Robertsson, and W. W. Symes (1995). Modeling of a constant Q : methodology and algorithm for an efficient and optimally inexpensive viscoelastic technique, *Geophysics* **60**, 176–184.
- Carcione, J. M., D. Kosloff, and R. Kosloff (1988). Wave propagation in a linear viscoacoustic medium, *Geophys. J. R. Astr. Soc.* **93**, 393–407.

- Carcione, J. M., D. Kosloff, A. Behle, and G. Seriani (1992). A spectral scheme for wave propagation simulation in 3-D elastic-anisotropic media, *Geophysics* **57**, 1593–1607.
- Day, S. M. and J. B. Minster (1984). Numerical simulation of attenuated wavefields using a Padé approximant method, *Geophys. J. R. Astr. Soc.* **78**, 105–118.
- Emmerich, H. and M. Korn (1987). Incorporation of attenuation into time-domain computations of seismic wavefields, *Geophysics* **52**, 1252–1264.
- Frankel, A. (1993). Three-dimensional simulations of ground motions in the San Bernardino Valley, California, for hypothetical earthquakes on the San Andreas fault, *Bull. Seism. Soc. Am.* **83**, 1020–1041.
- Frankel, A. and J. Vidale (1992). A three-dimensional simulation of seismic waves in the Santa Clara Valley, California, from a Loma Prieta aftershock, *Bull. Seism. Soc. Am.* **82**, 2045–2074.
- Graves, R. W. (1993). Modeling three-dimensional site response effects in the Marina District basin, San Francisco, California, *Bull. Seism. Soc. Am.* **85**, 1042–1063.
- Graves, R. W. (1996). Simulating seismic wave propagation in 3D elastic media using staggered-grid finite differences, *Bull. Seism. Soc. Am.* **86**, 1091–1106.
- Igel, H. and M. Weber (1995). SH-wave propagation in the whole mantle using high-order finite differences, *Geophys. Res. Lett.* **22**, 731–734.
- Kjartansson, E. (1979). Constant Q wave propagation and attenuation, *J. Geophys. Res.* **84**, 4737–4748.
- Krebs, E. S. and G. Quiroga-Goode (1994). A standard finite-difference scheme for the time-domain computation of anelastic wavefields, *Geophysics* **59**, 290–296.
- Levander, A. R. (1988). Fourth-order finite-difference P-SV seismograms, *Geophysics* **53**, 1425–1436.
- Minster, J. B. (1980). Anelasticity and attenuation, in *Physics of the Earth's Interior, Proc. Enrico Fermi Int. School Phys.*, A. M. Dziewonski and E. Boschi (Editors), North-Holland, Amsterdam.
- O'Connell, R. J. and B. Budiansky (1978). Measures of dissipation in viscoelastic media, *Geophys. Res. Lett.* **5**, 5–8.
- Olsen, K. B., J. C. Pechmann, and G. T. Schuster (1995a). Simulation of 3D elastic wave propagation in the Salt Lake Basin, *Bull. Seism. Soc. Am.* **85**, 1688–1710.
- Olsen, K. B., R. J. Archuleta, and J. R. Matarese (1995b). Three-dimensional simulation of a magnitude 7.75 earthquake on the San Andreas fault, *Science* **270**, 1628–1632.
- Robertsson, J. O. A., J. O. Blanch, and W. W. Symes (1994). Viscoelastic finite-difference modeling, *Geophysics* **59**, 1444–1456.
- Witte, D. C. and P. G. Richards (1990). The pseudospectral method for simulating wave propagation, in *Computational Acoustics, Seismo-Ocean Acoustics and Modeling*, D. Lee, R. Cakmak, and Vichnevetsky (Editors), North-Holland, Amsterdam, 1–18.
- Yomogida, K. and J. T. Etgen (1993). 3D wave propagation in the Los Angeles Basin for the Whittier-Narrows earthquake, *Bull. Seism. Soc. Am.* **83**, 1325–1344.
- Yoon, K.-H. and G. A. McMechan (1995). Simulation of long-period 3-D elastic responses for whole earth models, *Geophys. J. Int.* **120**, 721–730.

Department of Geological Sciences
San Diego State University
San Diego, California 92182

Manuscript received 26 September 1997.