

Model order reduction for seismic waveform modelling: inspiration from normal modes

R. Hawkins¹, M. H. Khalid,² K. Smetana³ and J. Trampert¹

¹*Department of Earth Sciences, Utrecht University, 3584 CS Utrecht, Netherlands. E-mail: rhys.hawkins@gmail.com*

²*Department of Applied Mathematics, University of Twente, 7522 NB Enschede, Netherlands*

³*Department of Mathematical Sciences, Stevens Institute of Technology, Hoboken, NJ 07030, USA*

Accepted 2023 May 3. Received 2023 April 3; in original form 2022 November 16

SUMMARY

The computational cost of full waveform simulation in seismological contexts is known to be expensive and generally requires large clusters of computers working in parallel. Although there have been many methods proposed over recent years to reduce this burden, in this work, we focus on a particular method called model order reduction (MOR) whereby a full waveform system of equations is projected onto a lower dimensional space to reduce computational and memory requirements at the cost of introducing approximation errors. In this paper, inspired by normal mode (NM) theory, we use the eigenmodes of the seismic wave equation to span this lower dimensional space. From this we argue that NM theory can be seen as an early form of MOR. Using this as inspiration, we demonstrate how free body oscillations and a form of Petrov–Galerkin projection can be applied in regional scale problems utilizing recent advanced eigensolvers to create a MOR scheme. We also demonstrate how this can be applied to inverse problems. We further conjecture that MOR will have an important role to play in future full waveform applications, particularly those of a time-critical nature such as seismic hazard monitoring.

Key words: Computational seismology; Theoretical seismology.

1 INTRODUCTION

The use of full waveform modelling as a means for interpreting observed seismograms and for understanding the composition of the Earth has been a boon for many decades for both global seismology (Komatitsch *et al.* 2002; Tromp *et al.* 2008; Fichtner *et al.* 2009a) and exploration scale problems (Tarantola 1984; Virieux & Operto 2009). A key problem with full waveform modelling is its high computational demand. An unfortunate side effect of this computational expense is that many potential applications, where full waveform approaches are necessary or would be preferable, are presently simply infeasible, for example, time-critical monitoring applications.

For this reason, alternative approaches are often used to approximate the full wavefield. Arguably the most common being ray theory which takes the limit at infinite frequency in homogeneous media as an approximation and reduces inference to travel time type problems (but with reflections, transmissions, geometric spreading also possible). A common criticism of ray theoretical approaches is that as a consequence of the infinite frequency approximation, the sensitivity of a ray is infinitely thin. Finite frequency approximations (Dahlen *et al.* 2000; Hung *et al.* 2001) attempt to overcome this shortcoming to generate sensitivity kernels based upon paraxial ray theory approximations of scattering.

By contrast, full waveform approaches attempt to model the underlying physics of viscoelastic wave propagation and avoid some limitations of ray and finite frequency approaches, that is, effects such as diffraction, multiple scattering and multiple arrivals are accounted for. Over the years, direct integration techniques for modelling the various physics of wave propagation have evolved and improved (Virieux *et al.* 2011). Early successes for modelling of wave propagation accurately involved finite difference approaches with acoustic approximations (Tarantola 1984) or a staggered grid approach, whereby wavefield variables and material parameters are offset by half of the spatial grid distance (Virieux 1986; Levander 1988).

More recently in the global seismology community, the spectral element approach has proved popular, owing to spectral convergence properties of the underlying basis and a diagonal mass matrix enabling efficient numerical time integration with explicit schemes. Key community codes that allow full waveform modelling and inversion are the SpecFEM family of 2-D, 3-D Cartesian and Global codes (Komatitsch & Tromp 2002a, b; Peter *et al.* 2011) and Salvus (Afanasyev *et al.* 2019).

While the spectral element method is well-suited to modelling of waveform propagation, it remains computationally expensive and there have been a number of methods used to compute accurate waveforms using clever approaches. One example is to use the near

spherical symmetry of the Earth to calculate synthetic seismograms in a 2-D spherically symmetric media at a dramatically reduced cost (Nissen-Meyer *et al.* 2014). Another approach is the use of adaptive remeshing to improve the runtime of forward and adjoint simulations (van Driel *et al.* 2020; Thrastarson *et al.* 2020).

In model order reduction (MOR) methods (Rozza *et al.* 2008; Benner *et al.* 2016, 2020b; Quarteroni *et al.* 2016; Hesthaven *et al.* 2016), the physical solutions for any parameter or frequency value are approximated by a linear combination of numerical solutions of the physical equations for few well-chosen parameters. This then only requires the solution of a very small (non)linear system of equations (typically of size 100–200) for every new parameter. MOR has been successfully applied, for example for structural and multibody dynamics, fluid dynamics and acoustics (Benner *et al.* 2020a). Although more difficult, recently there have been several developments in the application of MOR to the wave equation (Glas *et al.* 2020; Bigoni & Hesthaven 2020; Henning *et al.* 2022).

In many respects, these MOR techniques share similarities with perhaps one of the first methods used for full waveform modelling of long period seismograms, the so called normal mode (NM) theory. NMs are the free oscillations of the Earth and its theoretical underpinnings for seismological contexts began in the middle of last century (Pekeris & Jarosch 1958) and have developed and applied through to recent times (Gilbert 1970; Park & Gilbert 1986; Woodhouse 1988; Al-Attar *et al.* 2012). Within seismology, they have a long history of being used to infer the internal structure of the Earth at long scale lengths and they are claimed to be the best method for resolving the Earth's density structure.

The aim of this paper is to first demonstrate that traditional NM theory can in fact be seen as an early form of MOR. Secondly, that simulations of seismic wave propagation at regional scales, that is for small parts of the Earth, can be modelled using an analogous free oscillation approach. With the incorporation of absorbing boundaries and the treatment of rigid body modes we can rapidly and accurately calculate synthetic seismograms in a regional context using modal summation, analogous to typical NM summation for seismograms on the whole Earth. This creates a bridge between NM theory and MOR, in that NM theory is in essence a MOR method using Galerkin projection and surmise that eigenmodes of a free body oscillation should be seen as a benchmark to be improved upon in any MOR technique for seismic applications. This is particularly true for the global case, where the basis functions used are spherical harmonics calculated on a spherical symmetric, non-rotating, elastic isotropic (SNREI) model. Finally, we demonstrate how this approach can be more generally used in rapid calculations of regional seismograms for inverse problems and how they could be used to bring about new opportunities for time critical full waveform applications.

2 BACKGROUND THEORY IN A NUTSHELL

2.1 Model order reduction

In recent years we have witnessed the rapid growth of numerical simulation of physical processes at various scales, of which full waveform modelling is one example. The increasing applicability of such high fidelity physical modelling is due predominantly to a combination of improved algorithms and rapidly increasing computational resources. Nonetheless, in parallel with these developments, efforts over the last three decades or more have expanded

on the concept of MOR for approximating high fidelity physical modelling, primarily for applications that require either rapid calculation for time-critical problems, or large numbers of repeated simulations calculating similar related problems.

For the purposes of this paper, we give a brief conceptual introduction to MOR, which uses linear projection onto a reduced dimensional subspace of the underlying high fidelity space, and refer for further details to the text books (Benner *et al.* 2016, 2020b; Quarteroni *et al.* 2016; Hesthaven *et al.* 2016).

Using some spatial discretization scheme, the wave equation can be reduced to the following matrix form in the time domain

$$\mathbf{M}\dot{\mathbf{s}} + \mathbf{C}\dot{\mathbf{s}} + \mathbf{K}\mathbf{s} = \mathbf{f}, \quad (1)$$

with \mathbf{s} the time varying 3-D displacement vector of dimension (n), \mathbf{f} the time dependent seismic excitation and \mathbf{M} , \mathbf{C} and \mathbf{K} square matrices of dimension ($n \times n$).

The key idea in MOR is to exploit that while \mathbf{s} will generally belong to a high-dimensional space, properties of the solutions of eq. (1) will allow us to approximate \mathbf{s} very accurately in a low-dimension subspace spanned by appropriate basis vectors. This can be achieved by constructing a reduced basis orthogonal matrix \mathbf{V} of dimension ($n \times r$) and $r \ll n$, with which we can approximate the true displacement using

$$\mathbf{s} \approx \mathbf{V}\tilde{\mathbf{s}}, \quad (2)$$

and then rewrite eq. (1) as

$$\mathbf{M}\mathbf{V}\dot{\tilde{\mathbf{s}}} + \mathbf{C}\mathbf{V}\dot{\tilde{\mathbf{s}}} + \mathbf{K}\mathbf{V}\tilde{\mathbf{s}} \approx \mathbf{f}. \quad (3)$$

Upon multiplying with \mathbf{V}^T (transpose of \mathbf{V}) from the left, we obtain the reduced problem

$$\mathbf{V}^T\mathbf{M}\mathbf{V}\dot{\tilde{\mathbf{s}}} + \mathbf{V}^T\mathbf{C}\mathbf{V}\dot{\tilde{\mathbf{s}}} + \mathbf{V}^T\mathbf{K}\mathbf{V}\tilde{\mathbf{s}} \approx \mathbf{V}^T\mathbf{f}. \quad (4)$$

We note that instead of multiplying with \mathbf{V}^T from the left, which is called a Galerkin approximation, one can instead multiply from the left with a different matrix \mathbf{W}^T that helps with stabilization and which is then called a Petrov–Galerkin approximation. The final result is a reduced dimension matrix equation with terms such as $\mathbf{W}^T\mathbf{M}\mathbf{V}$ having dimension $r \times r$ and this process is illustrated in cartoon form in Fig. 1.

In MOR, a very popular way to construct the matrix \mathbf{V} is the so-called proper orthogonal decomposition (POD; Berkooz *et al.* 1993; Kunisch & Volkwein 2001; Sirovich 1987) where we compute the specific solution of eq. (1) or its counterpart in the frequency domain for well-chosen time or frequency and/or parameter values and store all these solutions in the columns of a matrix \mathbf{S} . Subsequently, we perform a singular value decomposition of the matrix and select the singular vectors belonging to the k largest singular values. POD is in effect a principal component analysis (PCA) of the solution vectors, in other words, we are trying to find the ‘principal components’ of \mathbf{S} or the vectors that best represent all vectors in the columns of \mathbf{S} . One way to decide on the number k , and thus the size of the linear system of equations (eq. 4), is to require that the k th singular value divided by the first and largest singular value lies below a certain tolerance. Usually, we observe a rapid decay of the singular values. However, for transport-dominated and wave-type problems with discontinuities, rapid decay of singular values may not occur requiring either non-linear approximations, exploiting properties of the physical system or a careful adaptation of the solution vectors in the matrix \mathbf{S} (e.g. Ohlberger & Rave 2013; Iollo & Lombardi 2014; Welper 2017, 2020; Reiss *et al.* 2018; Nair & Balajewicz 2019; Taddei 2020).

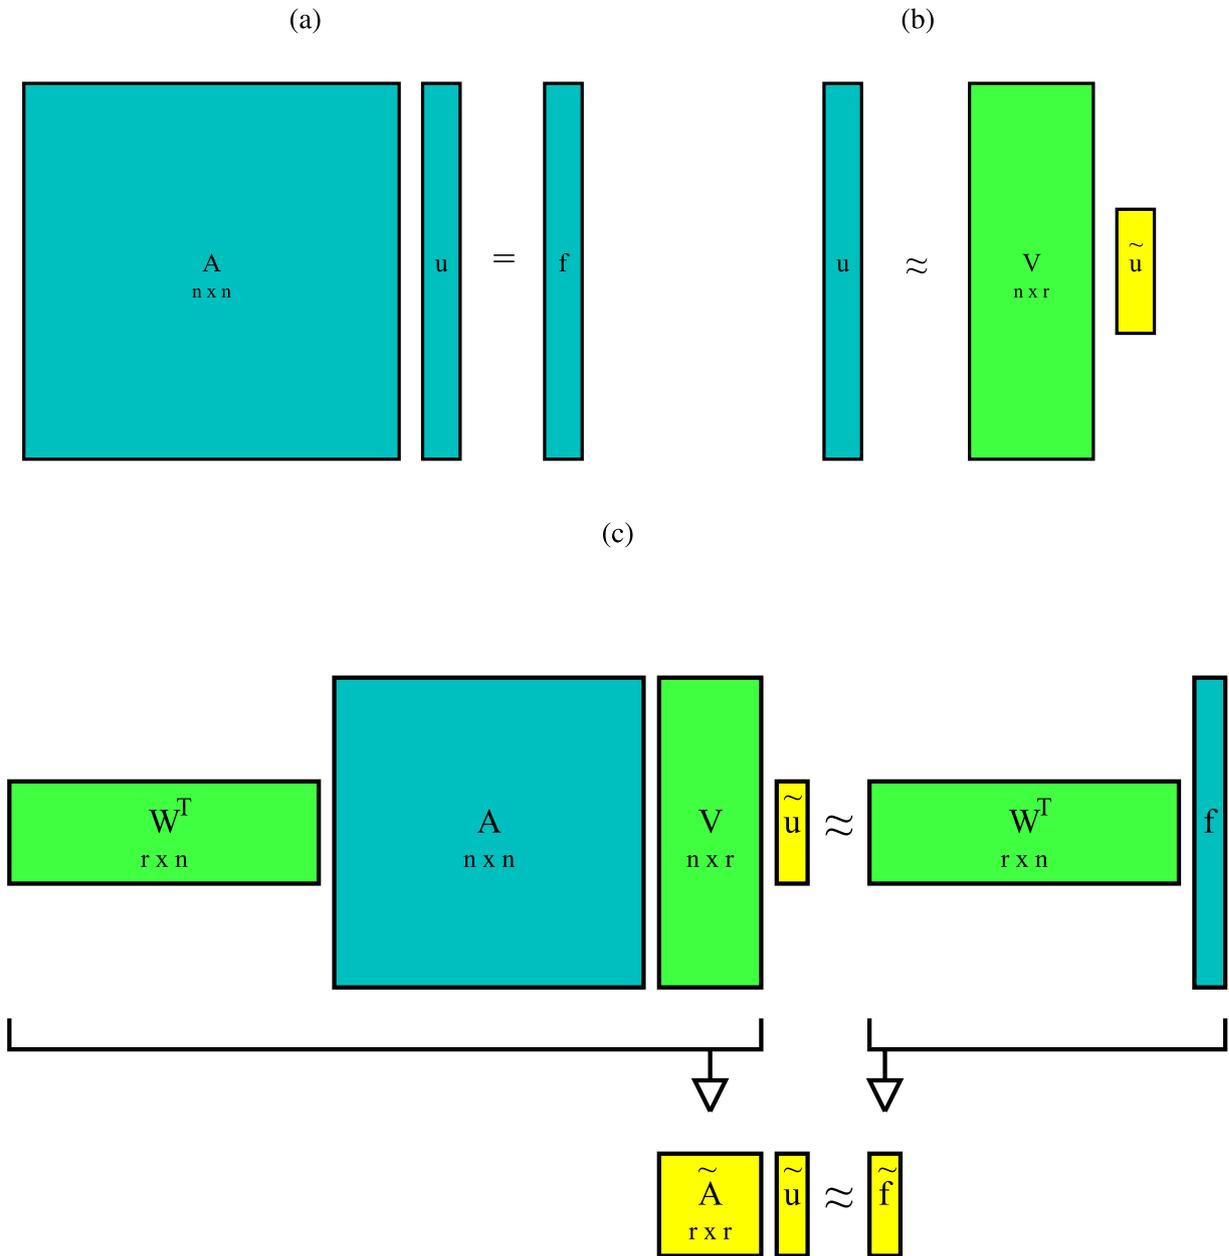


Figure 1. The process of MOR using a linear projection method is shown schematically in this figure. In (a) we have a standard linear system of dimension n with a solution vector \mathbf{u} . In (b) we can use a set of r basis vectors \mathbf{V} with $r \ll n$ to create an approximation of the solution. In (c), by inserting this approximation into the full dimension system, and multiplying from the left by either the same ($\mathbf{W} = \mathbf{V}$) or another set of basis vectors, we can condense the system of equations down to an approximation of the same form with dimension r . In some cases, using $\mathbf{W} = \mathbf{V}$ is sufficient which is called Galerkin projection. For some physical problems, using distinct bases for \mathbf{V} and \mathbf{W} , which is called Petrov–Galerkin projection, provides additional efficacy and stability.

One drawback of the POD method is that it requires solving eq. (1) for a rather large amount of time or frequency and/or parameter values and storing the respective solutions. A common alternative strategy is the so-called ‘Greedy’ algorithm (Veroy *et al.* 2003). The Greedy algorithm iteratively builds an orthonormal basis by selecting a frequency and parameter value from a finite dimensional subset of all admissible parameters, whose corresponding solution maximizes the error between the solution and its approximation with the already build basis vectors. Orthonormality of the basis is maintained via a Gram–Schmidt process. By considering only an estimator for the error instead of the error itself, it is in fact only

necessary to have a number of solutions of the order of the size of the reduced model; a huge improvement.

In principle, any method which allows us to find a reduced orthogonal basis appropriate for our system is possible, provided it is computationally cheap and effective at approximating \mathbf{s} and provides a stable reduced space in which to solve the system of equations through time. Rather than constructing \mathbf{V} from particular solutions of eq. (1), as normally done in MOR, we will use a eigenfunctions of the wave equation and thus be able to make the link with NM theory commonly used in seismology (e.g. Gilbert 1970; Park & Gilbert 1986; Woodhouse 1988; Al-Attar *et al.* 2012)

2.2 Normal modes

For completeness, we present a short summary of NM theory, but refer to Dahlen & Tromp (1998) for more details. In the NM community, the seismic wave equation is often written in the form

$$\left(\mathcal{V}(t) + \mathcal{W} \frac{\partial}{\partial t} + \rho \frac{\partial^2}{\partial t^2}\right) \mathbf{s}(t) = \mathbf{f}(t) \quad (5)$$

with $\mathcal{V}(t)$ an operator representing the material properties of the Earth with the dependence on time due to inelastic effects, \mathcal{W} representing Coriolis forces due to the Earth's rotation, ρ the equilibrium density, \mathbf{s} the displacement and lastly the external forcing \mathbf{f} . Here dependencies on spatial location are omitted for brevity.

The standard approach in NM seismology, is to use the spherically symmetric, non-rotating elastic isotropic (SNREI) approximation, which simplifies eq. (5) to

$$\left(\mathcal{V}_{\text{SNREI}} + \rho_{\text{SNREI}} \frac{\partial^2}{\partial t^2}\right) \mathbf{s}(t) = \mathbf{f}(t). \quad (6)$$

Now the material properties $\mathcal{V}_{\text{SNREI}}$ are independent of time and setting the external forcing to zero and taking the Fourier transform we can obtain a general eigenproblem of the form

$$(\mathcal{V}_{\text{SNREI}} - \omega^2 \rho_{\text{SNREI}}) \mathcal{S} = 0, \quad (7)$$

With \mathcal{S} the Fourier transformed displacement. This eigenproblem can be solved for pairs ω_k and \mathcal{S}_k over a frequency range of interest (Pekeris & Jarosch 1958; Woodhouse 1988).

In eq. (7), the eigenvectors \mathcal{S}_k are orthogonal and normalized so that

$$\int_V \mathcal{S}_j(\mathbf{x})^* \cdot \rho_{\text{SNREI}}(\mathbf{x}) \cdot \mathcal{S}_k(\mathbf{x}) d\mathbf{x} = \delta_{j,k}, \quad (8)$$

where $\delta_{j,k}$ is the Kronecker delta and * denotes complex conjugation. In an early result from Gilbert (1970), a synthetic seismogram can be computed in an SNREI earth model in a relatively straightforward and efficient manner using

$$\mathbf{s}_{\text{SNREI}}(\mathbf{x}, t) = \sum_k \frac{1}{\omega_k^2} \mathcal{F}_k(1 - \cos(\omega_k t)) \mathcal{S}_k, \quad (9)$$

where \mathcal{F}_k represents the earthquake forcing expressed in terms of eigenmode coefficients. This style of synthetic seismogram calculation has come to be known as mode summation.

Speaking more generally, since the eigenvectors are orthogonal, they form a basis with which we can express the displacement field within a 3-D Earth, that is

$$\mathbf{s}(\mathbf{x}, t) \approx \sum_k u_k(t) \mathcal{S}_k(\mathbf{x}), \quad (10)$$

where we have used approximately equal as strictly speaking we are referring to a subset of the eigensolutions. This can be combined with eq. (5) to give

$$\sum_k (u_k \mathcal{V} \mathcal{S}_k + \dot{u}_k \mathcal{W} \mathcal{S}_k + \rho \ddot{u}_k \mathcal{S}_k) = \mathbf{f}. \quad (11)$$

For any eigenfunction \mathcal{S}_j and integration over Earth's volume, we obtain

$$\sum_k \left(u_k \int_V \mathcal{S}_j^* \mathcal{V} \mathcal{S}_k d\mathbf{x} + \dot{u}_k \int_V \mathcal{S}_j^* \mathcal{W} \mathcal{S}_k d\mathbf{x} + \ddot{u}_k \int_V \mathcal{S}_j^* \rho \mathcal{S}_k d\mathbf{x} \right) = \int_V \mathcal{S}_j^* \mathbf{f} d\mathbf{x}. \quad (12)$$

Owing to the fact that a subset of eigensolutions are used, this can in turn be expressed in matrix form as

$$\mathbf{V}\mathbf{u} + \mathbf{W}\dot{\mathbf{u}} + \mathbf{P}\ddot{\mathbf{u}} = \mathbf{f}', \quad (13)$$

where the size of the matrices \mathbf{V} , \mathbf{W} and \mathbf{P} depends on the number of used eigenmodes. This equation is referred to as the 'full coupling' approach and can be used for low frequency studies of the Earth for density (Akbarashrafi *et al.* 2017), attenuation and anisotropy (Beghein & Trampert 2003; Deuss *et al.* 2010).

In the context of the previous section on MOR, the construction of eq. (13) represents a Galerkin projection of the full system of equations of the Earth in response to external forcing onto the reduced basis of the NMs of an SNREI earth model (Park & Gilbert 1986). Thus the method of using the NMs of an approximate and simpler earth model as a basis set and the subsequent Galerkin projection of heterogeneous and higher physical fidelity earth models can be seen as an early form of MOR for the calculation of global synthetic seismograms.

3 APPLICATION TO LOCAL SEISMOLOGY

3.1 Time-domain full waveform simulation

For the purposes of this work, we limit the discussion to time domain solutions using the spectral element method (Patera 1984). However, we stress that the presentation here equally applies to similar schemes such as finite differences, the finite element method, the discontinuous Galerkin method or, more generally, to schemes for which a spatial discretization results in a set of equations that can be arranged into matrix form.

The spectral element approach has found broad appeal in seismology since its introduction to wave propagation problems (Seriani *et al.* 1992; Komatitsch *et al.* 1999) and has been applied for regional 3-D (Komatitsch & Tromp 1999; Fichtner *et al.* 2009b; Peter *et al.* 2011) and global studies (Komatitsch & Tromp 2002a, b; Nissen-Meyer *et al.* 2014). The advantages of the spectral element method for wave propagation problems, referring specifically to Gauss–Lobatto–Legendre (GLL) quadrilateral or hexahedral elements, are a diagonal mass matrix, good convergence properties, ease of applying free surface boundary conditions, a computationally efficient tensorized basis and maintenance of high accuracy in distorted elements (Fichtner 2011; Oliveira & Seriani 2011).

For a time domain solution of the wave equation, the spectral element method discretizes the space into quadrilateral (2-D) or hexahedral (3-D) cells and within each of these cells approximates the continuous displacement field with GLL polynomials. The solution to the weak of the wave equation is approximated using numerical quadrature of the same basis.

For the boundary conditions of the regional simulations considered in this manuscript, we use a stress free condition on the top of the model and absorbing boundaries on all other boundaries. The stress free condition is enforced through the weak form of the equations of motion. There are many strategies for implementing absorbing boundaries but we choose to use Stacey absorbing boundaries (Clayton & Engquist 1977; Stacey 1988) which are sufficient for body wave applications considered here. This absorbing boundary condition applies tractions at the boundary in response to normal and tangential velocities of particle motion at the boundary and works well for absorbing normal incidence waves but is less applicable for surface wave studies.

Although not explicitly done in codes such as SpecFEM2D, the set of equations to be solved through time can be assembled in to a matrix equation of the form

$$\left(\mathbf{M} \frac{\partial^2}{\partial t^2} + \mathbf{C} \frac{\partial}{\partial t} + \mathbf{K} \right) \mathbf{s} = \mathbf{f}(t), \quad (14)$$

where the matrix \mathbf{M} is diagonal and dependent on density. The matrix \mathbf{C} is sparse diagonal and encodes the absorbing boundaries where for the particular absorbing boundaries we use, boundary velocities are used to generate tractions at the boundary to absorb outgoing wave energy (note that \mathbf{C} may be block diagonal if absorbing boundaries are not axis aligned). Lastly \mathbf{K} is a symmetric sparse matrix dependent on the elastic properties. The displacement vector \mathbf{s} written in this form contains the displacement in all Cartesian directions of discrete nodes, that is $\mathbf{s}^T = [s_x, 1s_y, 1s_z, 1 \dots]$, with the ordering of nodes generally optimized for performance.

Typically, the solution of the equations is numerically integrated through time using an explicit time marching scheme such as Newmark–Beta (Newmark 1959) or a Runge–Kutta variant (Berland *et al.* 2006). Alternatively, frequency domain techniques can be used by taking the Fourier transform of the equations and solving the system of equations for a set of desired frequencies (Pratt 1999).

While many decades of effort has focused on the application of NMs to whole Earth problems, there has been no published application of free body oscillations to regional studies that we are aware of. We speculate that the primary reasons for this are: firstly that there appears to be the view that a large number of eigenmodes are required for the calculation of synthetic seismograms (Cummins *et al.* 1997). Secondly, within the global seismology community, NMs are equated predominantly with spherical harmonics expansions and therefore global problems. Lastly, the calculation of eigendecompositions of large systems of equations as are typically seen in regional systems has historically been significantly more expensive than explicit time domain solutions or even frequency domain solutions. We will demonstrate that eigenmode solutions to regional problems are possible, computationally efficient and thus useful for applications in regional seismology.

3.2 A simple illustrative example

We begin with a small toy example to demonstrate the basic principles in a sufficient small model that can be rapidly solved with modest computational resources. This example is a 2-D homogeneous isotropic P–SV system with a domain of 1 km² with 200 m² 4th order cells and has only 882 degrees of freedom. The bottom, left -and right -handboundaries are absorbing with the top being a free surface. The values of the homogeneous elastic parameters are 2600 kg m⁻³ for density, 5.8 km s⁻¹ for P -wave velocity and 3.2 km s⁻¹ for S -wave velocity. The mesh and the location of a source (red star) and receiver (green triangle) are shown in Fig. 2

For a spectral element solution, Seriani & Priolo (1994) demonstrated that approximately 4.5 points per wavelength were sufficient to model wave propagation with minimal numerical dispersion. We use the SpecFEM2D code to compute a seismogram which suggests using 5 points per wavelength. Given this numerical constraint and the S -wave velocity, we can calculate that the model as configured will be accurate for waveforms up to approximately 12.5 Hz based upon the mean distance between points.

In order the compute the eigendecomposition for this small system, we can use standard LAPACK (Anderson *et al.* 1999) numerical

routines by reformulating the spectral element matrix equation from quadratic form into a linearized companion form (Hammarling *et al.* 2013; Hawkins 2018). This results in a doubling of the size of the matrices and the system of equations becomes

$$\underbrace{\begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}}_{-\mathbf{A}} \frac{\partial}{\partial t} \underbrace{\begin{bmatrix} \dot{\mathbf{s}} \\ \mathbf{s} \end{bmatrix}}_{\mathbf{B}} + \underbrace{\begin{bmatrix} \mathbf{C} & \mathbf{K} \\ -\mathbf{I} & \mathbf{0} \end{bmatrix}}_{\mathbf{B}} \begin{bmatrix} \dot{\mathbf{s}} \\ \mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{f}(t) \\ \mathbf{0} \end{bmatrix} \quad (15)$$

which can be seen as the velocity-displacement form of the wave equation. Similarly the stress–displacement form is another alternative that could equally be used. The equation is Fourier transformed and the forcing term set to zero leaving a general eigenproblem of the form

$$\lambda \mathbf{A} \tilde{\mathbf{y}} = \mathbf{B} \tilde{\mathbf{y}}, \quad (16)$$

with matrices \mathbf{A} and \mathbf{B} as shown in eq. (15), $\lambda = i\omega$ and $\tilde{\mathbf{y}}$ the Fourier transform of vector $\mathbf{y}^T = [\dot{\mathbf{s}} \mathbf{s}]^T$.

An additional implementation detail is that the matrices generally need to be rescaled for numerical accuracy owing to the large discrepancy in magnitudes between density and the elastic moduli resulting in a poorly balanced eigen equation (see Hammarling *et al.* (2013) and Hawkins (2018) for details of the scaling procedure). An alternatively strategy for the conditioning of the system of equations is to recast the equations in non-dimensionalized form, as is used in Mineos (Masters *et al.* 2011) and SpecFEM3D_Globe (Komatitsch & Tromp 2002a, b). It should also be noted that these two strategies are compatible and can be used jointly.

The eigenvalues of this simple model are plotted in Fig. 3 where for convenience we plot only the positive imaginary eigenvalues; the negative eigenvalues are simply the complex conjugate of the positive ones. We have scaled the eigenvalues by dividing by 2π meaning that the imaginary (vertical) axis can be read directly as frequency in Hertz.

A key aspect to note is that even with a spectral element mesh relatively optimized for a desired frequency range of interest, there are significant numbers of eigenvalues with imaginary components greater than the maximum frequency desired. Above we argued that our seismograms will only be accurate up to 12.5 Hz based on the mean distance between points. Quantitatively, considering only eigensolutions with zero or positive imaginary components, there are 896 total eigensolutions, of which 28 are real solutions, 90 are solutions with imaginary components corresponding to less than 12.5 Hz leaving 806 that contribute to frequencies outside the range considered by the mesh in a time domain solution. In principle, we should thus be able to compute accurate seismograms with only 118 eigensolutions. We will reexamine this point when we consider approximate seismograms through eigensolution truncation.

3.3 Modal summation

To compute a seismogram from the eigendecomposition of the system of equations, one approach is to follow a similar procedure as outlined by Gilbert (1970), with some small modifications. In the general eigenproblem in eq. (16), the matrix \mathbf{B} is non-symmetric ensuring that there will be distinct left and right eigenvectors. If we label these vectors \mathcal{L}_k and \mathcal{R}_k respectively with k indexing a distinct eigenpair, they can be normalized so that

$$\mathcal{L}_k^* \mathbf{A} \mathcal{R}_j = \delta_{k,j}, \quad (17)$$

where $*$ is the conjugate transpose.

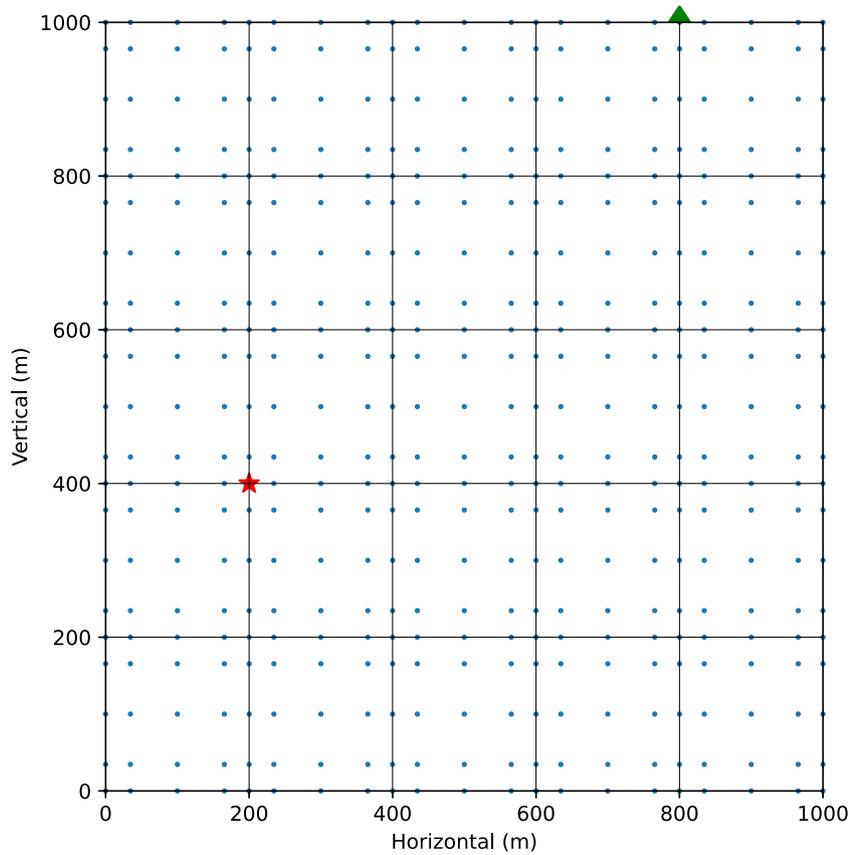


Figure 2. The simple homogeneous mesh used for demonstration purposes. The mesh is a five by five cell using 4th order spectral elements. The Gauss–Lobatto–Legendre (GLL) nodes are shown with small points. A source and receiver are also shown with a red star and green triangle respectively.

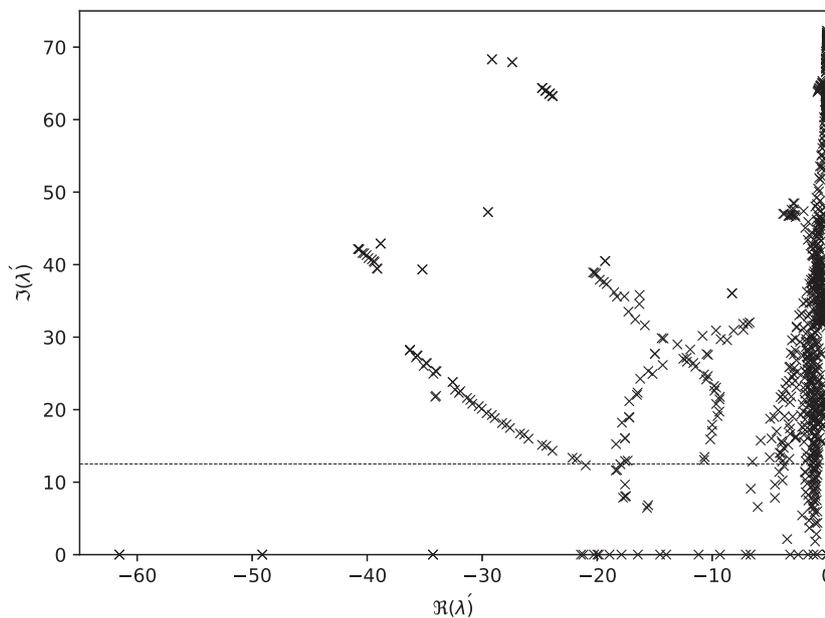


Figure 3. The eigenvalues computed from a small spectral element model with absorbing boundaries. The imaginary (vertical) axis is labelled in frequency for convenience. The horizontal dashed line represents the maximum frequency based on the requirement of a minimum of 5 points per wavelength.

Using similar arguments to Gilbert (1970), we can express the state vector \mathbf{y} as a sum over the right eigenvectors

$$\mathbf{y} = \sum_k y_k \mathcal{R}_k, \tag{18}$$

where y_k is a scalar weight of the k th right eigenvector. Strictly speaking, for equality the summation needs to be over all right eigenvectors. A key difference between this formulation and that of Gilbert (1970) is that the vectors \mathcal{R}_k here are linearly independent rather than orthogonal. The combination of

left and right eigenvectors form a biorthogonal basis so that for any \mathbf{y}

$$y_k = \mathcal{L}_k^* \mathbf{A} \mathbf{y}, \quad (19)$$

which follows from eq. (17).

It is useful here to separate the forcing term in eq. (15) into a constant vector in space and a source–time function, for example

$$\mathbf{f}(t) = \mathbf{x}_s f(t). \quad (20)$$

Then the wavefield for a given forcing can be written as

$$\mathbf{y}(t) = \sum_k \left[\underbrace{(\mathcal{L}_k^* \cdot \mathbf{x}_s)}_{\mathcal{F}_k} \underbrace{\mathcal{R}_k}_{S_k} \underbrace{\exp(\lambda_k t)}_{1 - \cos(\omega_k t)} \right] * f(t). \quad (21)$$

This can be seen as a further generalization of eq. (9) with approximate correspondences indicated with under braces. The other difference here is the inclusion of a convolution operation with the general source time function $f(t)$ which is absent in the derivation by Gilbert (1970) as they used a fixed Heaviside moment tensor solution.

To compute a single component seismogram at a particular receiver we can use a vector \mathbf{x}_r to represent the sensitivity of the seismogram and write

$$s(t) = \left[\sum_k (\mathcal{L}_k^* \cdot \mathbf{x}_s) (\mathbf{x}_r^T \cdot \mathcal{R}_k) \exp(\lambda_k t) \right] * f(t). \quad (22)$$

In a spectral element mesh, the vector \mathbf{x}_r will be constructed as an interpolation using the GLL polynomials within a single cell.

In Fig. 4, we show the full modal seismogram (using all 896 eigensolutions) calculated with eq. (22) for our simple P–SV system using all eigensolutions. As can be seen from the plots, the computed seismograms are visually identical to a direct time integration of eq. (14) with mean squared errors of 0.025 and 0.057 mm² for the horizontal and vertical components, respectively.

3.4 Direct diagonal time integration

An alternative method for simulating wave propagation, especially in combination with a Rayleigh–Ritz approach, is to use the properties of the eigendecomposition to diagonalize the system of equations in companion form eq. (15). Rewriting the state vector expansion (eq. 18) in matrix notation we obtain

$$\mathbf{y} = \mathbf{R} \mathbf{z}, \quad (23)$$

where the coefficients y_k are used to form the vector \mathbf{z} and the right eigenvectors form the columns of matrix \mathbf{R} . The vector \mathbf{y} represents the velocity–displacement state $[\dot{\mathbf{s}} \ \mathbf{s}]^T$ as before. If we restrict \mathbf{R} to a subset of eigenvectors, \mathbf{z} represent a projection of the state vector onto a smaller dimensional space. We can substitute this into eq. (15) to obtain

$$\mathbf{A} \mathbf{R} \dot{\mathbf{z}} = \mathbf{B} \mathbf{R} \mathbf{z} - \mathbf{x}_s f(t). \quad (24)$$

Multiplying this by the left eigenvectors we obtain

$$\dot{\mathbf{z}} = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} \mathbf{z} - \mathbf{L}^* \mathbf{x}_s f(t). \quad (25)$$

That is, the mass matrix is reduced to the identity matrix owing to the normalization in eq. (17) and and the biorthogonality of the eigenvectors. The matrix on the right-hand side, \mathbf{B} , is reduced to a diagonal matrix with the eigenvalues along the diagonal. The solution \mathbf{z} represents an approximation of the velocity and displacement of the domain and the displacement can be evaluated, for example for a synthetic seismogram at a point, using a receiver vector \mathbf{x}_r , that is,

$$\mathbf{u}(\mathbf{x}_r, t) = \mathbf{x}_r \mathbf{y}(t) \approx \mathbf{x}_r \mathbf{R} \mathbf{z}(t). \quad (26)$$

It can be shown that numerically integrating this system using all the eigenvalues gives the same result as the full modal summation shown in Fig. 4 within machine precision. We will illustrate the usefulness of this type of solution in subsequent examples.

3.5 Eigensolution truncation

In Fig. 4 it was demonstrated that the modal summation produced faithful seismograms that compared well to a direct time domain solution. In theory, the seismograms should be exactly the same but there will be small differences between them due to limited precision floating point calculations, with contributions both from the calculation of the eigensolutions and the different ways the synthetic seismograms are computed.

Once the eigensolution is computed, using the full eigensolution in the calculation of seismograms is significantly cheaper than the time domain solution, by around an order of magnitude for this simple problem. However, as shown in Fig. 3, there are many eigenvalues with frequencies much higher than that which the mesh has been optimized for. Additionally, we can see some large negative real values in eigenvalues, indicating that these eigensolutions decay rapidly in time and are as such less likely to impact the seismogram of interest.

We can construct a seismogram using modal summation with only those eigensolutions within the desired frequency band of interest, that is seismograms using 118 eigensolutions, consisting of 3 rigid body modes, 25 eigensolutions with real eigenvalues, and 90 eigensolutions with imaginary components corresponding to frequencies less than 12.5 Hz. These seismograms are shown in Fig. 5, where a key difference between this truncated solutions and the full eigensolution is that there are larger errors at the beginning of the seismogram that decay to near zero at 0.3 s. It has been suggested that a larger number of eigensolutions are required above a given threshold to more accurately reconstruct seismic signals using modal summation (Akbarashrafi *et al.* 2017). However, here a different mechanism is the cause for these errors.

3.6 Green’s function oscillations

From the NM summation approach, we can stably compute the point to point Green’s function for both the full eigensolution and the truncated eigensolution. This is shown in Fig. 6 with the full eigensolution Green’s function in red and the truncated to 12.5 Hz Green’s function in black.

It is evident that the truncated curve is in essence a low-pass filtered version of the full solution. A problem with the truncated Green’s function is the fact that it is non-zero at zero time as a consequence of the truncation. A similar phenomena has been observed in using NM methods for the calculation of sensitivity kernels (Capdeville 2005). The action of convolving the Green’s function with a Ricker wavelet source has the effect of low-pass filtering

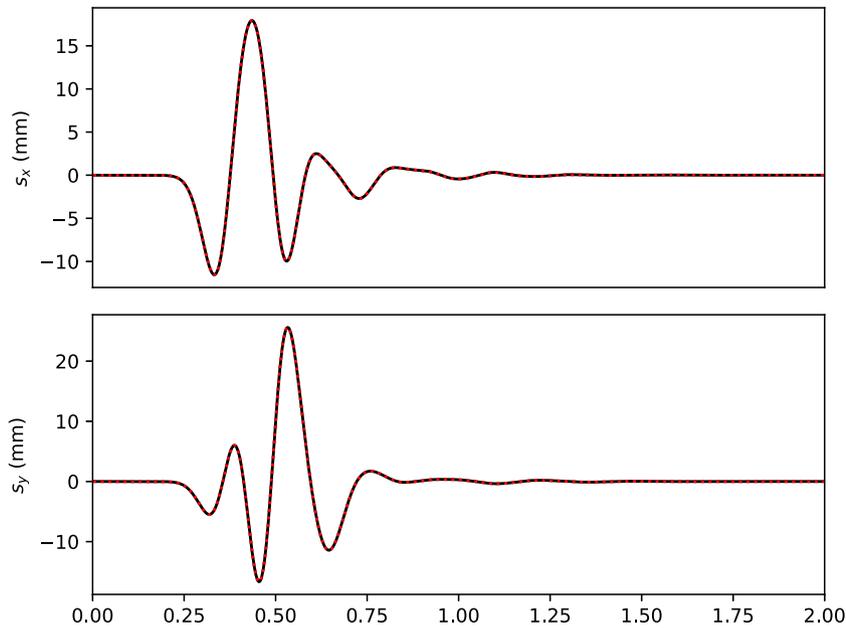


Figure 4. The seismograms for the horizontal (top) and vertical (bottom) components calculated using mode summation is shown in black. The same seismograms calculated using a time domain numerical scheme is shown in red dotted lines for comparison. In this case the rigid body modes (see Appendix A1) are replaced with analytical solutions to improve the static offset.

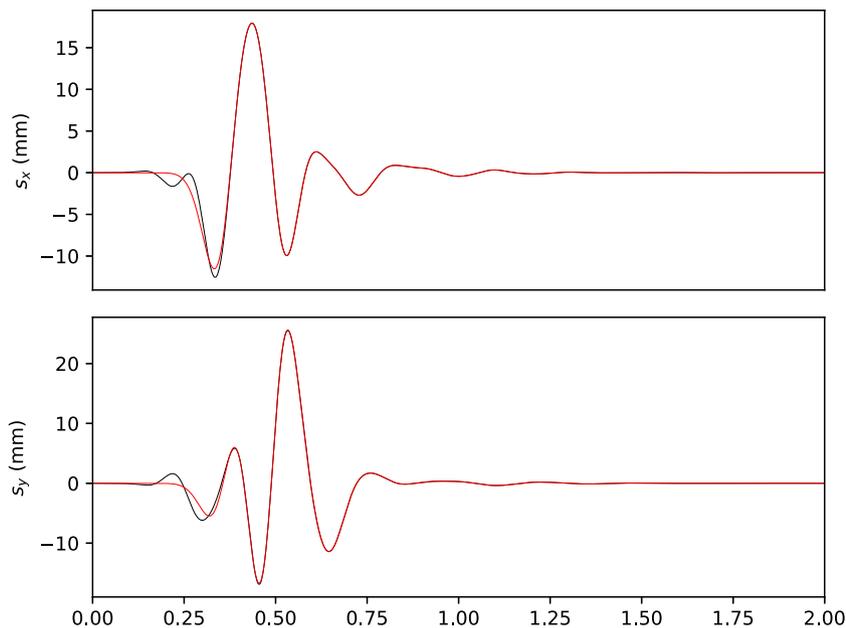


Figure 5. With reference to Fig. 3, this plot shows the reconstruction of the simple seismogram using only eigensolutions with a corresponding frequency less than 12.5 Hz.

the Green's functions to produce seismograms without large oscillations. The step at zero-time in the truncated version acts as a Heaviside function during convolution with the source time function, and is at the origin of the large errors at the beginning of the seismogram.

Using the red curve as a guide in Fig. 6, we can observe that the approximately first 0.125 s of the Green's function should be zero. This represents the time for the direct propagation of seismic energy from the source to the receiver. If we use a simple cosine tapering up to 0.125 s of the truncated Green's functions before convolving with

the Ricker source time function, the computed seismograms are now as shown in Fig. 7. In comparing the seismograms in Figs 5 and 7 we can see that there is a dramatic improvement achieved by tapering the early parts of the truncated Green's function. In practice, this problem only exists in the case where source–receivers are close in terms of the frequency range of interest. In this small example, the source and receiver are less than two wavelengths apart. More generally, when sources and receivers are further apart, the issues caused by the truncation of Green's functions can be safely ignored using appropriate windowing of seismograms.

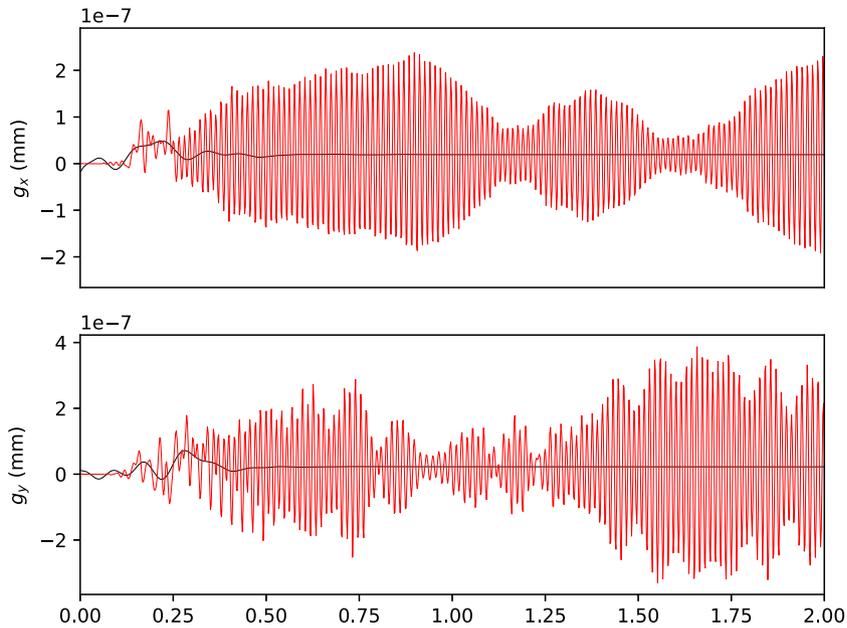


Figure 6. The Green's functions for a vertical impulse constructed from the full eigensolution (red) and the eigensolutions with frequencies less than 12.5 Hz.

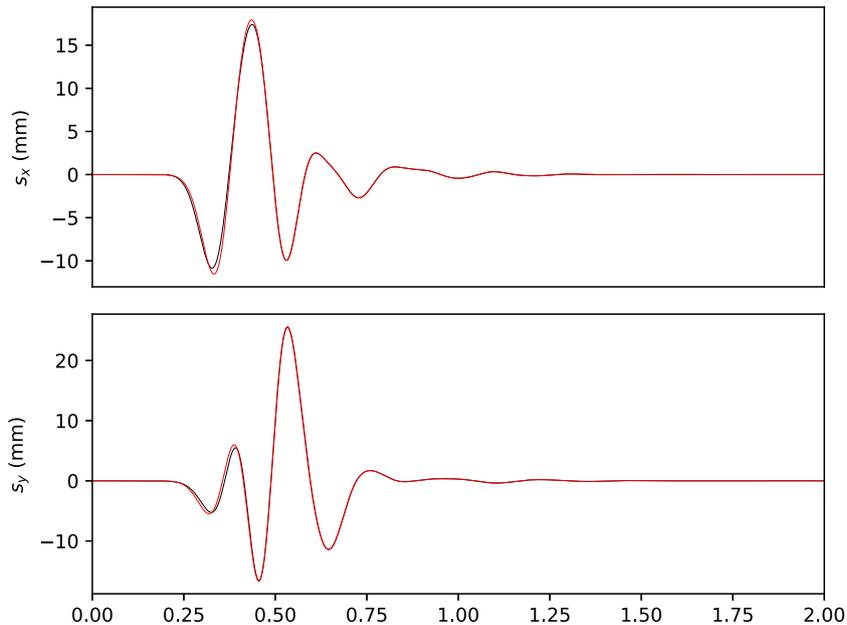


Figure 7. This plot shows the reconstruction of the simple seismogram using only eigensolutions with a corresponding frequency less than 12.5 Hz with tapering used to eliminate the wiggles in the beginning of the Green's functions.

3.7 Summary

The point of this small scale example is to show that applying free body oscillation techniques to a regional model may in fact be useful in various contexts.

Even a simple and reasonably optimal spectral element mesh has a large number of eigensolutions that contain frequencies well outside the range of interest. Using a reduced subset of eigensolutions in conjunction with tapering results in accurate seismograms with a dramatic reduction in memory and computational requirements. Although not applicable in this small example, the tapering is only required when the duration of the source is sufficiently large to

interfere with the first arriving pulse of the Green's function. In typical seismology problems, the source duration is small compared to the travel time between a source and a receiver. In a scenario with longer traveltimes, this would cause a small error in the seismogram around the source time, but would not be relevant in many applications as important parts of the synthetic seismograms, for example, the first arrival, would be unaffected. In the case where near source accuracy is important, then there are additional strategies to treat this artefact (Capdeville 2005).

A key problem with the use of eigensolutions is that the method used for this small scale problem, that is, linearization of the quadratic eigenproblem and then solving for all eigensolutions, does

not scale well. However some recent advances in eigensolvers have in part made this more tractable allowing application of these ideas to larger and more complex problems.

4 LARGER HETEROGENEOUS MODELS

A key problem with extending eigen approaches to full waveform modelling is the solution the general eigenproblem. The approach in the previous section becomes computationally prohibitive at only a modest number of degrees of freedom. For larger eigenproblems, there are many different approaches well-suited to linear symmetric eigenproblems. A recent advance are methods that directly solve polynomial eigenproblems which are particularly useful to our formulation.

4.1 The FEAST non-linear eigensolver

A relatively recent advance in eigensolvers is the use of a contour integration technique in eigenvalue space to efficiently compute a subset of eigenvalues and their eigenvectors. The original description of this class of eigensolvers was introduced by Polizzi (2009) for general symmetric eigenproblems. It has since been extended to non-Hermitian matrices (Kestyn *et al.* 2016) and importantly for our application, polynomial and non-linear eigenproblems (Gavin *et al.* 2018). This is of particular interest here in that we have a quadratic eigenvalue problem.

There are three key features of the FEAST algorithm that dramatically improve the efficiency of computing eigensolutions. First, the algorithm does not need to linearize a quadratic eigenproblem and therefore increase the size of the problem by two. The non-linear FEAST algorithm solves a polynomial problem without increasing the size of the problem and can additionally take advantage of the fact that the matrices \mathbf{M} , \mathbf{C} and \mathbf{K} are symmetric (i.e. an arbitrarily weighted summation of these matrices is symmetric whereas in companion form the matrices are no longer symmetric). The process of linearizing the general matrix equation into companion form both increases the size of the eigenproblem and destroys symmetry resulting in a far more computationally challenging problem. Secondly, the contour integral formulation means that the FEAST algorithm is simple to parallelize in that we can compute discrete regions of complex eigenspace in parallel in a trivial way. Lastly, the FEAST algorithm can be parallelized at different levels from threaded parallelism for small scale problems that can fit entirely in memory, up to very large systems that can then be solved with parallel iterative solvers operating across multiple nodes of a cluster supercomputer (Gavin & Polizzi 2018).

4.2 Complex model

At many scales the Earth exhibits potentially large variations in composition and in turn seismic velocities including sharp jumps in material properties. To perform an experiment on a larger and more realistic structure we constructed a quadrilateral mesh using a velocity model of the Groningen gas field in the Netherlands. The elastic properties of the model and significant interfaces were obtained from Romijn (2017). Some small scale features were removed for the purposes of making the meshing of the domain more tractable, and then a quadrilateral mesh was generated using the GMSH software (Geuzaine & Remacle 2009).

The resulting mesh has approximately 9000 cells, and 300 000 degrees of freedom. The mesh is shown in Fig. 8 and has been

designed to satisfy five points per wavelength for frequencies up to 2.8 Hz. In the figure, shading is used to indicate the shear wave velocity and shows the location of sharp discontinuities in the material properties in the model. Of particular note is the very slow shear wave velocity near the surface that requires a large number of small cells to accurately simulate wave propagation.

4.3 Eigensolutions

We computed a range of eigenvalues for this system using the FEAST polynomial eigensolver routines. We used 2.8 Hz as a threshold for the imaginary axis. The set of eigenvalues was computed on four nodes of a cluster super computer (each node consists of two Xeon Cascade Lake processors with 24 cores each and 192 Gb of memory). The computational walltime for this calculation of the eigensolutions of this system was 11 min 27 s of wall time, and the corresponding measured CPU time was 18 hr and 45 min.

The rigid body modes were computed from analytical results and the quadratic eigensolutions were combined into companion form eigenvectors as shown in Appendix A1. The plot of eigenvalues obtained is shown in Fig. 9.

4.4 Seismograms

In Fig. 10, we present the results of computing the seismograms using the 2744 eigensolutions obtained in the previous section. The waveform simulations have been computed using the direct diagonal time integration scheme from Section 3.4 and no tapering has been applied. As a result of this, a clear error is seen around one second due to Green's function truncation. In Fig. 10, we have included the result across a variety of receivers at the surface to show the approximation across a broader range of the domain, and with the exception of the truncation errors around the one second mark, the reconstruction is very accurate. The ground truth is obtained from a direct time integration of eq. (14).

4.5 Computational savings

For the seismograms above we have 2744 eigensolutions, either real or complex with positive imaginary components. Strictly speaking, the seismograms are calculated including the conjugate eigensolutions as well, but we can incorporate the effect of the conjugate eigensolutions with a simple rescaling.

To quantify the performance increase in computing a seismogram from a theoretical point, we first present an estimate of the number of floating point operations required to solve for the time derivative for both the time domain solution and the modal summation solution. The reason for using this as a measure of performance improvement is that it is a concrete way of comparing theoretical performance independent of implementation optimizations and the numerical time integration method used.

In the example presented, the time domain solution system has 297 378 degrees of freedom and the stiffness matrix has 21 209 476 non-zero matrix entries in total. The absorbing boundary matrix is diagonal with 1346 non-zero entries. From this we can calculate that there are approximately 43 Mflops required to solve for the gradient (dominated by the stiffness matrix multiply).

Solving eq. (25) for a single seismogram using 2744 eigensolutions, since the \mathbf{B} matrix is now diagonal, we only need to perform effectively 3 complex vector operations in this case requiring 8232 complex floating point operations. Using a factor of 4, this equates

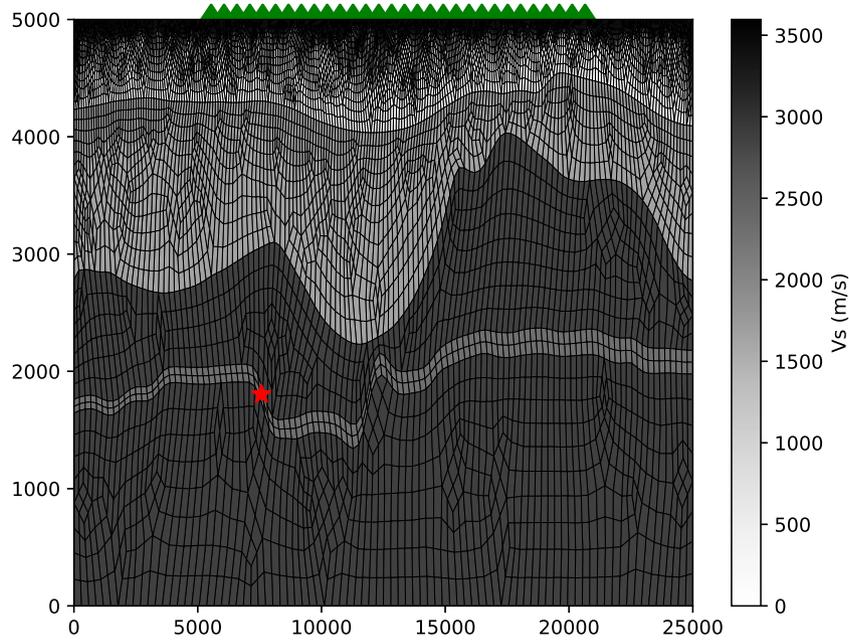


Figure 8. The quadrilateral mesh for a 2-D slice through the Groningen region in the Netherlands. The velocity model was derived from Romijn (2017). An earthquake location is indicated with a red star, and a series of receivers are indicated with green triangles on the surface. The mesh is shaded with the shear wave velocity to highlight the sharp interfaces in the model and the very low shear wave speed near the surface.

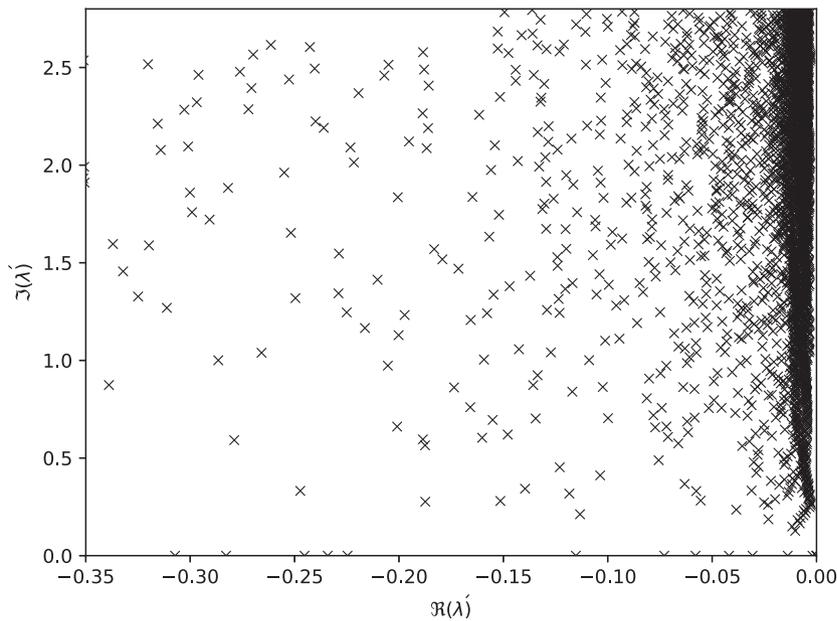


Figure 9. The computed eigenvalues for the 2-D mesh in Fig. 8 up to a frequency of 2.80 Hz. The eigenvalues are scaled by two pi so that the imaginary axis can be read as frequency.

to approximately 33 kflops. For this example, the number of floating point operations is reduced by more than three orders of magnitude. In both these case, this represents the computational complexity of advancing a simulation of the domain to the next time step and is independent of the number of sources or receivers.

For a direct timing comparison, we ran the SpecFEM2D time domain code versus the direct diagonal time integration code serially using the same configuration, that is, time step size, number of iterations, numerical integration scheme, and numbers of sources and receivers. For the time domain code for this test we obtained

a wall time of 2460 s and for the direct diagonal time integration, based on modal decomposition, an average of 9.69 s, or around a 250 times speed-up. For this test, 30 seismograms were computed. When computing the theoretical speed-up based on floating point operations as in the previous paragraph, we obtain a expected 230 times reduction in computational complexity which compares well with that obtained through direct measurement. This reinforces a well-known result that a direct diagonal time integration, based on modes, can be significantly faster for computing seismograms, but it has an additional feature that can be used to improve this further.

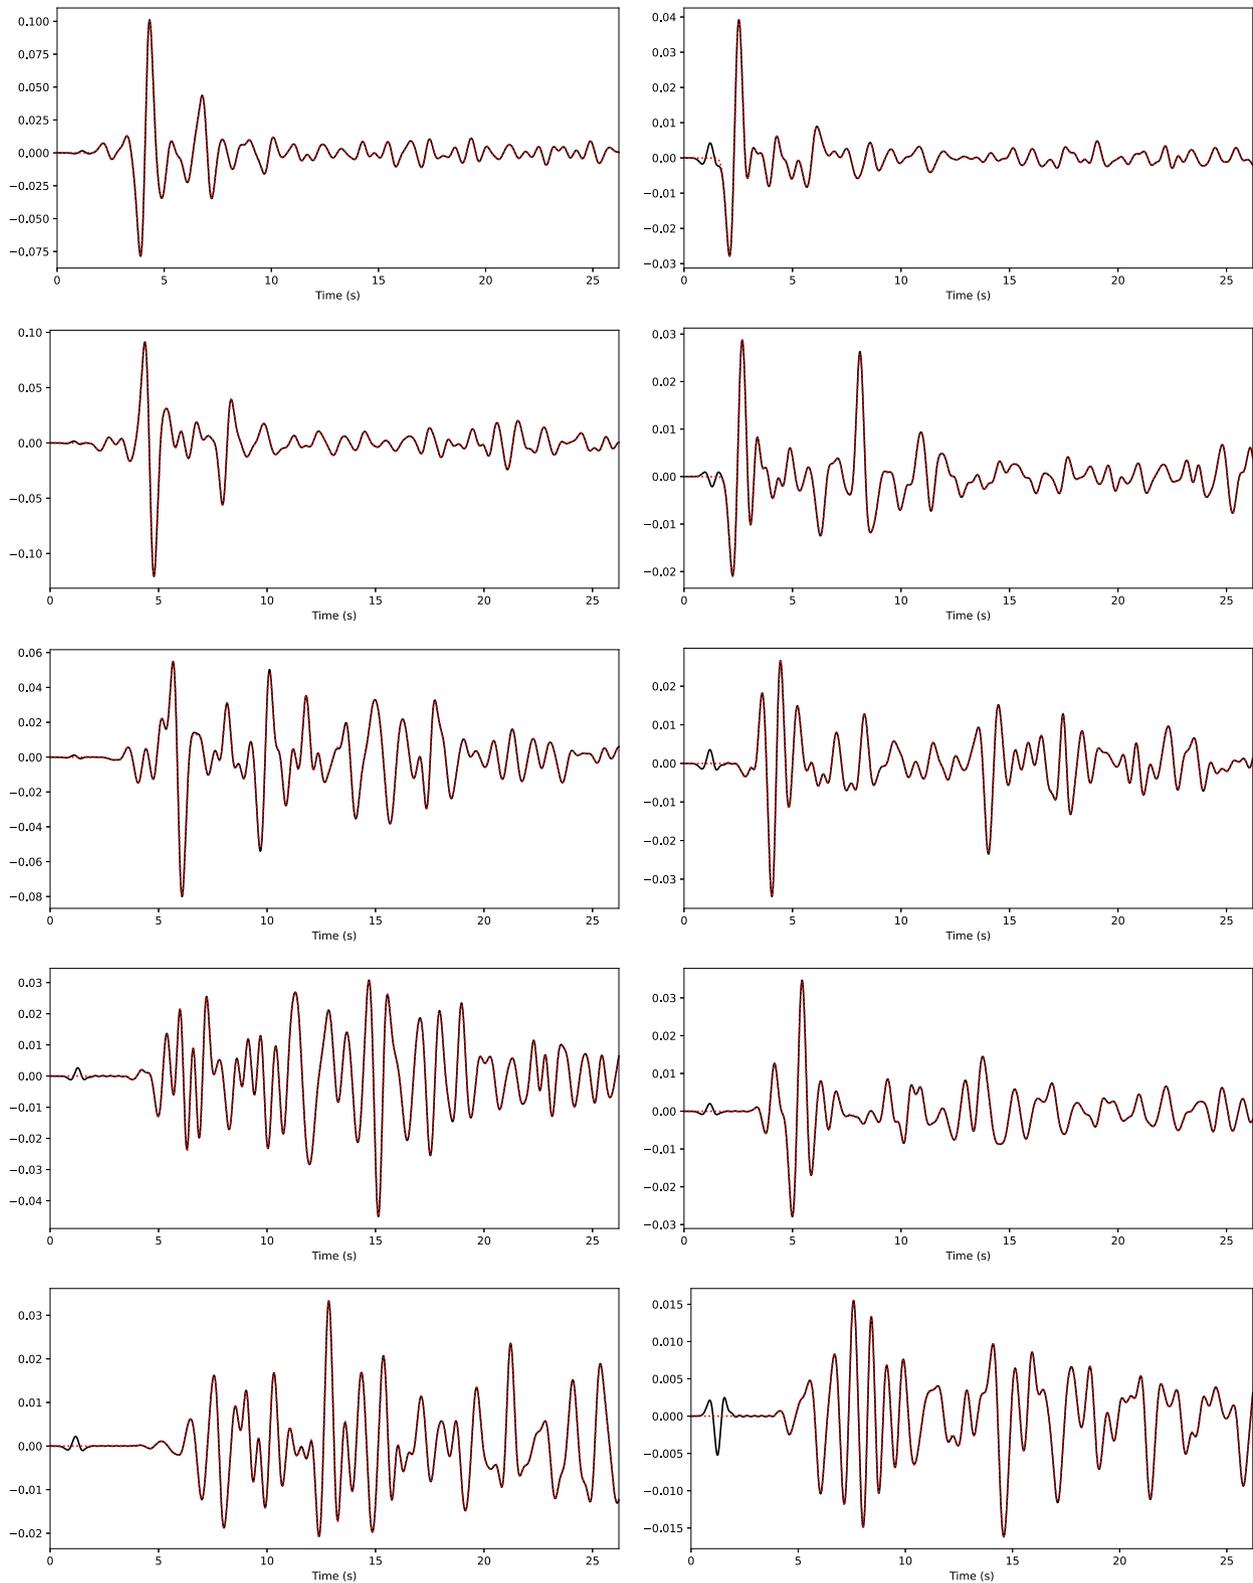


Figure 10. A selection of horizontal (left-hand column) and vertical seismicograms computed using a direct time domain solution (red) and the same seismicogram approximated using a truncated modal summation solution (black).

4.6 Courant–Friedrichs–Lewy for the truncated system

With an explicit time integration scheme, we are limited by the Courant–Friedrichs–Lewy (CFL) condition (Courant *et al.*

1928). This condition effectively specifies a maximum time step size based on the velocity and node spacing in the spatial mesh. Another way of thinking about the CFL limit is

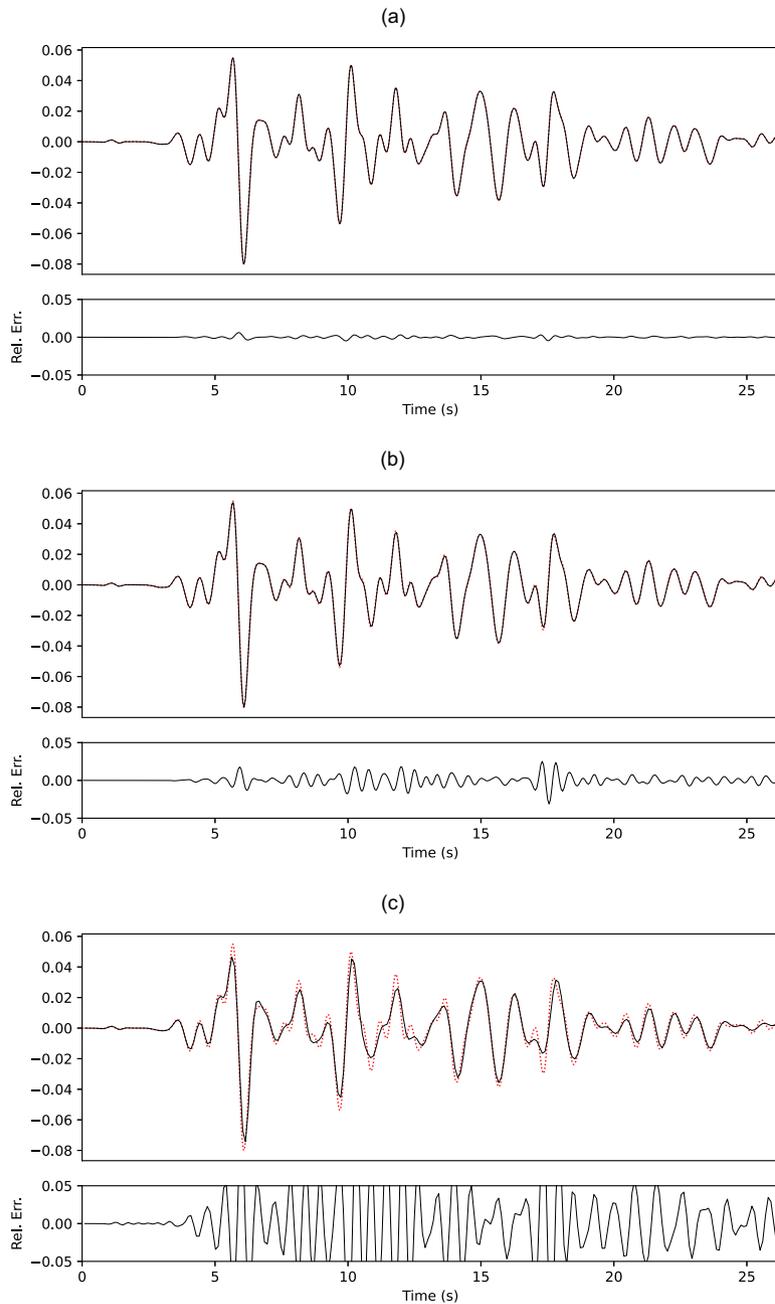


Figure 11. Three representative plots comparing simulations with increasing time step size relative to a reference solution. In (a) the time step is 32 times the size, 64 in (b) and 128 in (c). The reference seismogram is plotted with a red dotted line and the simulation with a larger step size in black. Even for dramatically larger step sizes, the seismogram is still faithfully reconstructed.

that it represents the time step required to maintain stable numerical simulations of the oscillations of the highest frequency eigenmode.

In the previous examples on the comparison of wall-clock times, we have used the same time step size. We tested the accuracy of the direct diagonal time integration, based on a reduced set of modes, as the time step size is progressively doubled and we show some representative plots in Fig. 11. In general there is very little error up to a 64 times increase in time step size, but for a 256 times increase the simulation becomes unstable. For each doubling of the time step size, we are approximately halving the computational time. In Fig. 12, we demonstrate this and with a combination of

both the direct diagonal time integration technique and its ability to use time steps well beyond the original time domain CFL limit for explicit numerical integration in SpecFEM2D, it is not unreasonable to achieve accurate seismograms with a 4 orders of magnitude reduction in computational time.

Another feature of the direct diagonal integration technique which we have not addressed here is that since the matrix on the right-hand side is diagonalized, implicit numerical integration schemes are possible without requiring a matrix solve operation at each step. This could further improve on the accuracy of larger step sizes and provide further incremental performance improvements.

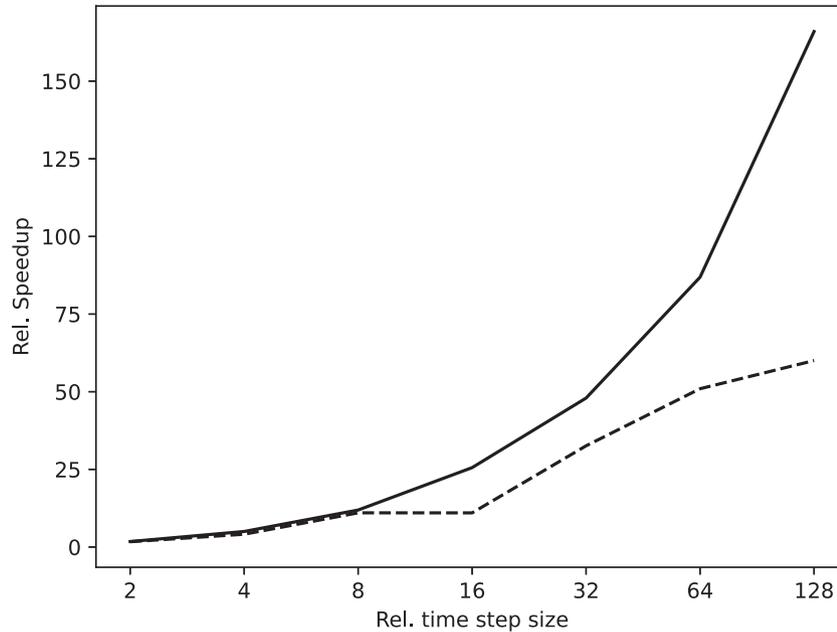


Figure 12. As the time step size is increasingly doubled, we naturally achieve a similar performance improvement. This plot demonstrates this scaling with the solid line indicating the speed up of the seismogram computation and the dashed line indicating the overall speed up (wall time) which includes initialization and I/O.

4.7 Simulating small model perturbations

In many applications where we simulate seismic wave propagation, we often wish to use the difference between simulated waveforms and those observed to update our model of the Earth, for example, in seismic full waveform tomography. In the case where we perturb a model by a small amount, for example, as a result of an iteration in a seismic tomography optimization problem, it would be possible to recompute the new eigensolutions using the FEAST library. However, directly recomputing eigenfunctions from scratch at each model update would be prohibitive. Alternatively, since we have shown that we can very accurately reproduce a waveform using a truncated eigensolution, a question arises as to whether we can equally compute seismograms for nearby models, that is, models with small perturbations. To test this we have applied a range of perturbations to the model in Fig. 8.

The procedure we can follow for computing the seismograms for small model perturbations is again inspired from NM theory and uses an analog to the ‘full coupling’ approach. We assume that we have small perturbations $\delta\mathbf{M}$ and $\delta\mathbf{K}$ and that we have not perturbed the absorbing boundary (i.e. the matrix \mathbf{C} is unchanged). Then we can use the projection technique to obtain

$$\mathbf{L}^* \begin{bmatrix} -\mathbf{M} - \delta\mathbf{M} & 0 \\ 0 & -\mathbf{I} \end{bmatrix} \mathbf{R}\dot{\mathbf{z}} = \mathbf{L}^* \begin{bmatrix} \mathbf{C} & \mathbf{K} + \delta\mathbf{K} \\ -\mathbf{I} & 0 \end{bmatrix} \mathbf{R}\mathbf{z} - \mathbf{L}^* \mathbf{x}_s f(t), \quad (27)$$

which potentially results in a non-diagonal complex matrix on the left-hand side that we would need to invert during the simulation. In the case of no perturbations to the density of the model, the equation remains still quite efficient to integrate by stepping through time, but on the right-hand side the addition of the $\delta\mathbf{K}$ term results in a dense matrix instead of having a diagonal matrix reducing computational efficiency.

4.8 The Rayleigh–Ritz method

Taking again inspiration from global seismology, where we invoke the Rayleigh–Ritz method by expanding the eigensolutions of a 3-D earth model as a linear combination of the eigensolutions of the 1-D earth model (eq. 10), to arrive at a smaller size eigensystem (eq. 13) for the expansion coefficients.

We invoke a similar argument to arrive at eq. (27), which is now not diagonal, but has reduced dimension, and is thus easy to rediagonalize. This is similar to what is referred to as full mode coupling in global seismology (eq. 13). So starting from a perturbed equation

$$\mathbf{A}'\dot{\mathbf{z}} = \mathbf{B}'\mathbf{z} - \mathbf{L}^* \mathbf{x}_s f(t), \quad (28)$$

we can solve this system as another general eigenproblem in the frequency domain

$$\lambda \mathbf{A}'\tilde{\mathbf{z}} = \mathbf{B}'\tilde{\mathbf{z}}, \quad (29)$$

where $\lambda = i\omega$ and $\tilde{\mathbf{z}}$ is the Fourier transform of \mathbf{z} . This procedure comes at the additional cost of another, albeit modest, eigensolution, but the result is the ability to very rapidly compute potentially multiple seismograms. For perturbations to density, this procedure is generally always preferable as it reduces the left-hand side matrix in eq. (28) to the identity matrix making explicit numerical integration far more efficient. If only the elastic parameters are perturbed, the computational saving depends on the number of seismograms to be computed. For a single seismogram, it is perhaps better to simply proceed without using the rediagonalization step and numerically integrate using the dense matrix on the right-hand side. However, at some number of seismograms, the cost of recomputing the smaller eigenproblem using a Rayleigh–Ritz approach will be amortized.

4.9 Results

To verify that accuracy is preserved when we use a reference set of eigensolutions to compute seismograms in perturbed models, we

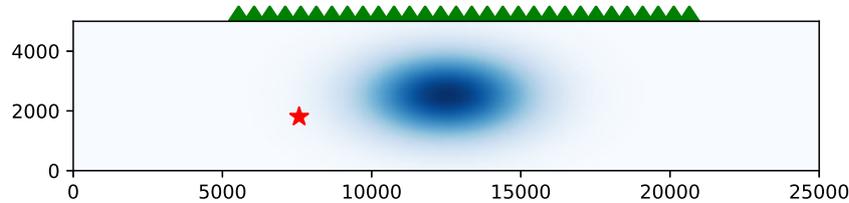


Figure 13. The Gaussian perturbation location used for the seismograms with varying degrees of perturbations to density, S -wave velocity and P -wave velocity.

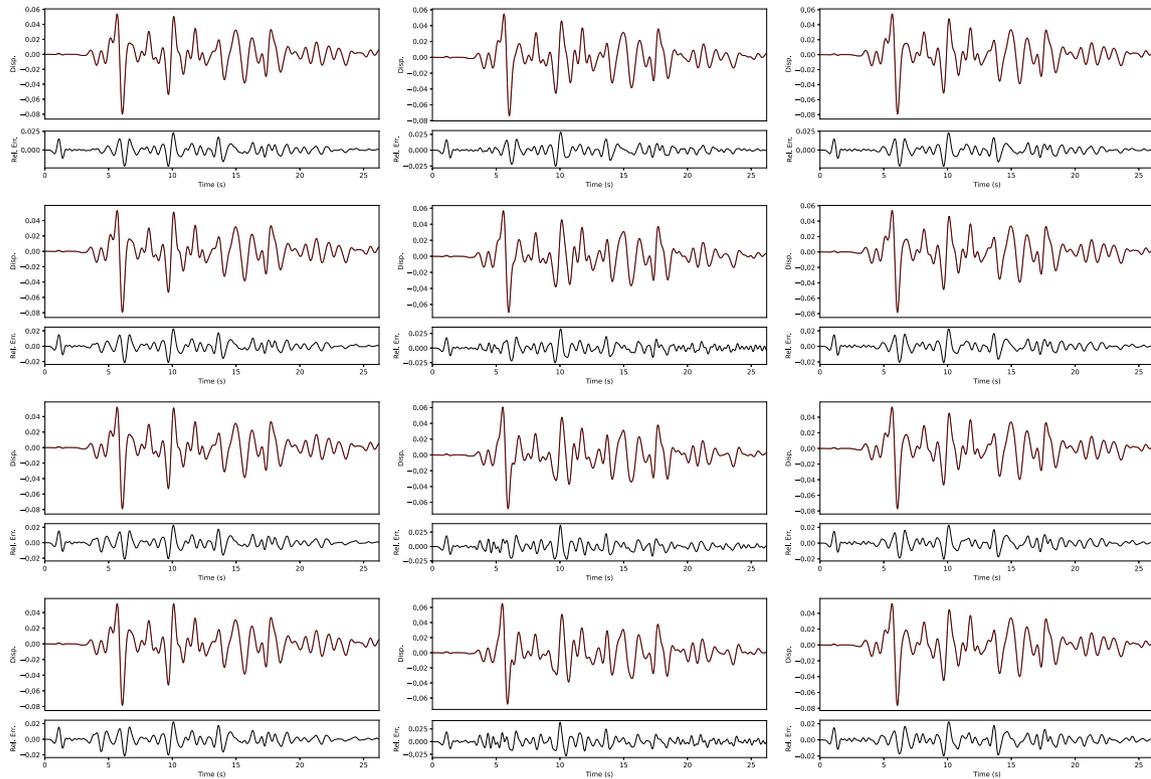


Figure 14. Horizontal displacement seismograms for a single station with increasing perturbations to density (left-hand column), S -wave speed (middle column) and P -wave speeds (right-hand column). The perturbation magnitudes for each row of results is 10, 20, 30 and 40 per cent, respectively. In each plot, the eigensolution is shown with a black line and the equivalent time domain solution is shown with a dotted red line. Below each plot we show the relative error between the two.

used the model construction with the amplitude of the Gaussian perturbation shown in Fig. 13 varying between 1 and 40 per cent for density, P -wave velocity and S -wave velocity. The results of these independent solutions are shown for the horizontal seismograms in Fig. 14 and the vertical seismograms in Fig. 15. We show the seismograms only for the central seismometer and the rows in the figures represent 10, 20, 30 and 40 per cent perturbations and the columns represent density, S -wave and P -wave perturbations. For each seismogram we show the equivalent time-domain solution as a dotted red line and the Rayleigh–Ritz direct diagonal time integration solution in black. The relative error is also shown below each plot. We observe from the range of the plots that the error due to the approximate eigensolution is small. The key error is the ringing artefact prior to the first arrival waves that can be ignored in any application using appropriate data windowing. In general the maximum relative error is approximately 2 per cent but, importantly, information such as phase are correctly reproduced throughout the seismograms.

This example demonstrates that for applications such as seismic tomography or full waveform monitoring where we have a reasonable starting model, we may only need to compute the eigendecomposition once. For nearby models we can subsequently compute wave propagation using reprojection and optionally a Rayleigh–Ritz procedure. In this example, the difference is dramatic: for the full eigensolution we need of the order of 10 min of computation on 4 nodes of a cluster (192 cores), whereas for reprojection and optionally rediagonalizing using a Rayleigh–Ritz procedure takes of the order of 10 s on a single core.

5 CONCLUSION

We have demonstrated that with the FEAST library we can compute a sufficiently large set of eigensolutions for reasonable sized models efficiently that then allows us to compute accurate seismograms up to 4 orders of magnitude faster than an equivalent time domain

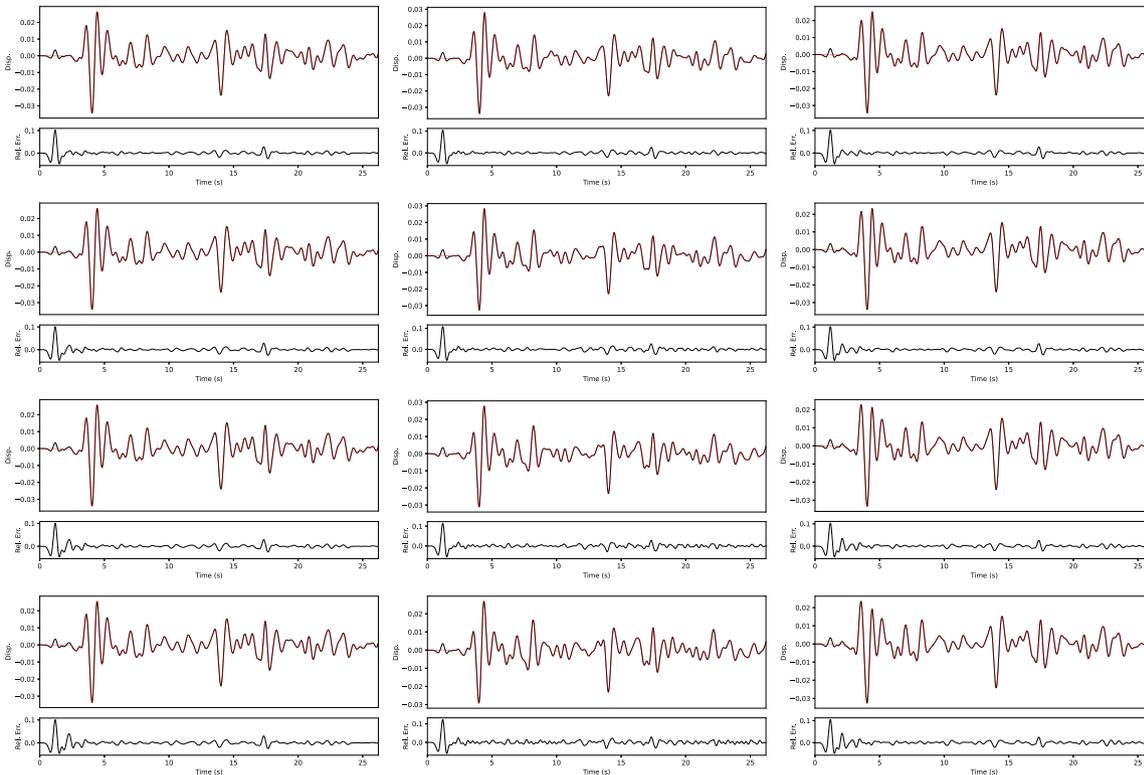


Figure 15. Vertical displacement seismograms for a single station with increasing perturbations to density (left-hand column), S -wave speed (middle column) and P -wave speeds (right-hand column). The perturbation magnitudes for each row of results is 10, 20, 30 and 40 per cent, respectively. In each plot, the eigensolution is shown with a black line and the equivalent time domain solution is shown with a dotted red line. Below each plot we show the relative error between the two.

code. The cost for computing the eigendecomposition is still significant meaning that for large-scale wave propagation or for a small number of solutions, that a time domain simulation is still preferable. However, we also showed that by computing a reference set of eigensolutions for nearby models, we can use a Rayleigh–Ritz technique to compute seismograms of nearby models, again with high efficiency and with limited loss of accuracy, even for relatively high perturbations to the model.

6 APPLICATIONS TO INDUCED STRESS MONITORING

As an illustration of a full-waveform monitoring application, we construct a synthetic experiment based on recovering a simulated induced stress in a gas reservoir. A recent formulation of induced stress has been derived and implemented using the spectral element method (Tromp & Trampert 2018). From the point of view of the matrix equation (eq. 14), the addition of an induced stress field effectively applies a perturbation to the stiffness matrix \mathbf{K} ,

$$\left(\mathbf{M} \frac{\partial^2}{\partial t^2} + \mathbf{C} \frac{\partial}{\partial t} + \mathbf{K} + \mathbf{T}^0 \right) \mathbf{s} = \mathbf{f}(t). \quad (30)$$

This is similar (with $\mathbf{T}^0 = \delta\mathbf{K}$) to the examples from the previous section. However, the induced stress \mathbf{T}^0 is not an arbitrary perturbation but instead must satisfy equilibrium stress conditions; specifically the stress field must be divergence free and must satisfy boundary conditions (e.g. free surface conditions).

We could follow the same procedure as in Section 4, that is, for each proposed model of induced stress, form the full sized matrix

\mathbf{T}^0 and project the new stiffness matrix using selected eigensolutions (eqs 27 and 28). However, this would prove prohibitive in an MCMC setting for a potential real-time or near real-time monitoring application as each new model needs to be first constructed at the full model dimension and then subsequently projected onto the reduced subspace using the set of eigensolutions. Here, we suggest an approach that can improve performance during inversion by pre-computing the projection of a predefined basis used to encode a proposed induced stress field model.

The benefits of this are twofold: First, the computational demands of matrix building and reprojection are removed from the online component of seismogram calculation and secondly, by judicious tuning of an induced stress basis, the convergence of any inverse problem improves owing to a reduction of dimensionality. The potential drawback of this approach is that such a reduced basis will not be able to capture the full range of induced stress models that may be possible, and this is a choice in the inversion akin to selecting a grid coarseness for a traditional tomographic inversion.

6.1 Parametrizing induced stress tensors in spectral element meshes

The parametrization of induced stress fields is non-trivial as an induced stress field in equilibrium must be divergence free throughout the domain and satisfy boundary conditions. In the context of the whole Earth, Al-Attar & Woodhouse (2010) presented an approach to constructing a basis for equilibrium stress fields using generalized spherical harmonics. For our regional 2-D study, we follow a

similar line, but use spectral elements directly to parametrize an induced stress field tensor that varies across a 2-D domain.

As a demonstration, we will use SH-wave propagation in a 2-D setting, which reduces the induced stress tensor to being composed in terms of four parameters. The symmetric induced stress tensor at a point in space can be written as

$$\mathbf{T}^0 = -p^0 \mathbf{I} + \begin{bmatrix} \tau_{xx} & \tau_{xz} \\ \tau_{xz} & \tau_{zz} \end{bmatrix}. \quad (31)$$

The components of the induced stress tensor such as p^0 , τ_{xx} , τ_{xz} can be discretized over a spectral element mesh. Taking this discretization, we can solve for a basis set by constructing a large system of equations which specifies continuity and boundary conditions at each point of the spectral element mesh. Notionally, we construct a large vector with all the discrete values of the stress components and we can then write a simple matrix equation

$$\mathbf{D} \begin{bmatrix} p_1^0 \\ \vdots \\ p_n^0 \\ \tau_{x,x,1} \\ \vdots \end{bmatrix} = 0, \quad (32)$$

where \mathbf{D} is the matrix which represents the spectral element discretized restrictions, that is the stress is divergence free, $\nabla \cdot \mathbf{T}^0 = 0$ with the boundary condition $[\mathbf{n} \cdot \mathbf{T}^0]_{\pm}$.

Eq. (32) can be solved numerically using a complete singular value decomposition. From this we obtain a set of vectors \mathbf{t}_i representing solutions satisfying $\mathbf{D} \cdot \mathbf{t}_i = 0$ (i.e. the null space of \mathbf{D}) and that are mutually orthogonal. Since each solution vector \mathbf{t}_i is an equilibrium stress field, any linear combination of this set of solutions forms an equilibrium stress field, and hence, this set of orthogonal vectors can then be used as a basis for representing the induced stress in our 2-D problem.

The construction of the induced stress basis in this manner does not need to be computed on the same mesh as the seismic wave simulation. In fact, a coarser mesh can be used and the resulting basis can be interpolated to the finer simulation mesh actually used for waveform calculations. The interpolation properties of spectral elements ensures that numerical imprecision that may cause violations of the induced stress conditions will be minimal (under the assumption that the embedded mesh is of a coarser resolution than the simulation mesh). For the purposes of this study, we have used a coarse mesh to define the induced stress field which is subsequently interpolated onto the simulation mesh using the interpolation property of the GLL polynomials used by the spectral element method. The resolution of the the embedded mesh from which the induced stress basis was obtained was selected by trial an error to obtain a reasonable synthetic stress model.

This procedure then allows an induced stress model to be written as

$$\mathbf{T}^0 = \sum_i \alpha_i \mathbf{T}_i, \quad (33)$$

where the α_i are scalar weights and \mathbf{T}_i are the induced stress matrices constructed using the null space vectors, \mathbf{t}_i that are solutions to eq. (32). The procedure for calculating each \mathbf{T}_i is: first take the i th null space vector solution from eq. (32), which defines the induced stress in the embedded mesh. Next, interpolate the embedded mesh to the global mesh used for simulation to create induced fields of pressure and deviatoric stresses. Finally, assemble the global matrix \mathbf{T}_i for the induced stress field using the same global assembly procedure

for the full spectral element mesh, that is the same procedure as used for the stiffness matrix. Inserting this back into eq. (30) and performing the projection we have

$$\mathbf{L}^* \begin{bmatrix} -\mathbf{M} & 0 \\ 0 & -\mathbf{I} \end{bmatrix} \mathbf{R} \mathbf{z} = \mathbf{L}^* \begin{bmatrix} \mathbf{C} & \mathbf{K} + \sum_i \alpha_i \mathbf{T}_i \\ -\mathbf{I} & 0 \end{bmatrix} \mathbf{R} \mathbf{z} - \mathbf{L}^* \mathbf{x}_s f(t), \quad (34)$$

and by using precomputed projections of the individual \mathbf{T}_i matrices, that is

$$\tilde{\mathbf{T}}_i = \mathbf{L}^* \begin{bmatrix} 0 & \mathbf{T}_i \\ 0 & 0 \end{bmatrix} \mathbf{R} \quad (35)$$

we can simulate the effect of induced stress using a projected equation of the form

$$\dot{\mathbf{z}} = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_r \end{bmatrix} \mathbf{z} + \sum \alpha_i \tilde{\mathbf{T}}_i \mathbf{z} - \mathbf{L}^* \mathbf{x}_s f(t) \quad (36)$$

which can be very efficiently updated for changes to the model parameters α_i as the matrices involved are significantly smaller.

6.2 Prior

In solving a non-linear inverse problem for the set of model parameters α_i outlined in the previous section, in almost all cases, some a priori constraint is required to stabilize the inverse problem. Ideally, in an inversion we would aim to constrain model parameters in a physically meaningful way, yet the α_i model parameters that we wish to solve for have little physical meaning.

In a general non-linear inverse problem (Tarantola & Valette 1982), we often assume Gaussian prior information and seek to minimize

$$(\mathbf{g}(\mathbf{m}) - \mathbf{d})^T \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{m}) - \mathbf{d}) + (\mathbf{m} - \mathbf{m}_0)^T \mathbf{C}_m^{-1} (\mathbf{m} - \mathbf{m}_0), \quad (37)$$

where \mathbf{g} is our non-linear forward operator, \mathbf{m} the model vector and \mathbf{d} the data vector. \mathbf{C}_d and \mathbf{C}_m are the data and model prior covariance matrices, respectively. The second term is a choice in formulating the inverse problem and in this case penalizes deviations from some reference or mean model \mathbf{m}_0 .

In the general case where, for example, the model parameters are associated directly with the spectral element mesh, a spatial prior consisting of some mean and covariance can be intuitively derived. What is needed is to project physical prior constraints onto the model parameters α_i . Given *a priori* information on the spatial variation of induced stress over the embedded mesh, a basis set $\mathbf{S} = [\mathbf{t}_1 \mathbf{t}_2 \dots \mathbf{t}_n]$ and a vector of model parameters $\boldsymbol{\alpha}$, we can write

$$(\mathbf{m} - \mathbf{m}_0)^T \mathbf{C}_m^{-1} (\mathbf{m} - \mathbf{m}_0) \approx (\boldsymbol{\alpha} - \boldsymbol{\alpha}_0)^T \mathbf{S}^T \mathbf{C}_m^{-1} \mathbf{S} (\boldsymbol{\alpha} - \boldsymbol{\alpha}_0), \quad (38)$$

where the approximation comes from the fact that \mathbf{S} is non-square and hence $\mathbf{S} \boldsymbol{\alpha}_0 \approx \mathbf{m}_0$. Since \mathbf{S} is non-square we need to solve for $\boldsymbol{\alpha}_0$ in order to specify the prior. From the point of view of an induced stress reference model, we have taken the approach that the induced stress will be pressure dominant with the deviatoric components having zero mean. As such, we have chosen to compute $\boldsymbol{\alpha}_0$ using a linear least squares method that minimizes the deviatoric components. The rationale for this choice stems from the observation that at the near surface, changes in induced pressure have a greater influence on the wave propagation. Other choices are available, but there is an inherent compromise in this approach whereby the reduced basis restricts the space of models available and as such the application of prior information.

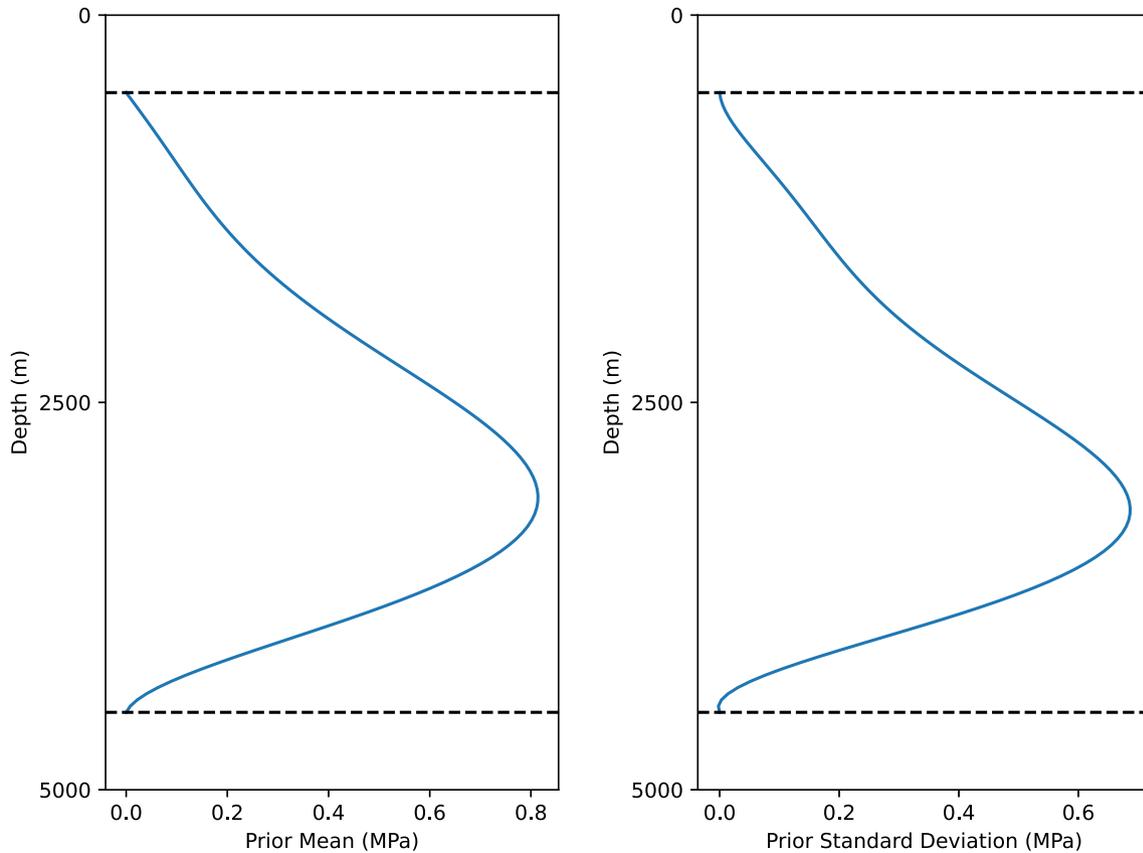


Figure 16. The spatial prior mean and standard deviations used in the synthetic inversion. We assume an independent Gaussian distribution of the pressure field with a mean and standard deviation solely dependent on depth. For the deviatoric components of the induced stress, we use a zero mean prior with a standard deviation of 0.2 MPa.

For the synthetic problem, we first construct a spatial prior in the embedded mesh that varies with depth for the pressure field as shown in Fig. 16. This profile is used to define the mean and standard deviation of an independent Gaussian prior for the induced pressure as a function of depth in the 2-D region. For the deviatoric components, the prior is also an independent Gaussian but with a zero mean. This spatial prior is then translated to the scalar coefficients of the basis vectors, α , using eq. (38).

6.3 Synthetic target

To construct a synthetic test problem, we utilize the same mesh from Fig. 8 and construct a synthetic gas extraction scenario where gas is extracted from the thin layer at around 3 km depth (i.e. the Rotliegend layer which is a porous gas bearing sandstone layer in the Groningen region). Rather than computing a physically based induced pressure from modelling methods (Postma & Jansen 2018), we opt for constructing a pressure field using a peak pressure drop of 2.5 MPa at the base of an imagined well with exponential decay horizontally along the layer with an additional vertical smearing applied through a Gaussian filter. The schematic of the true synthetic model is shown in Fig. 17, with the extraction well shown with a red line.

Synthetic seismograms were then created for the stations indicated with green triangles in Fig. 17 from an earthquake located at the red star which was simulated with a vertical force and a Ricker wavelet source time function. Independent Gaussian noise

was added to the seismograms with a noise level chosen to be approximately 2 per cent of the peak amplitude of the seismogram at the central station. A simple independent Gaussian likelihood for the measure of fit between model predicted seismograms and our ground truth model and manually selected windows were used. Although the inversion strategy here is simple, the main point is to demonstrate the rapid calculation of synthetic seismograms using a reduced free body oscillation basis approach that could be used more broadly in more complex inversion schemes and other forward modelling applications.

For this inversion, we ran 48 independent Markov chains which were started from random models sampled from the prior outlined in the previous section. For computing the synthetic seismograms from model parameters, the procedure is as follows: we first sum the induced stress contributions as in eq. (36), then use the Rayleigh–Ritz procedure and compute the eigendecomposition of the reduced system, before finally calculating the seismograms. In this case, even though there is only a single source, the seismograms are quite long in time and recomputing an eigensolution provided the best performance. Each forward calculation took a fraction over 2 s with this approach and this includes summing the induced stress bases, recomputing the eigensolution, and subsequent time integration.

The MCMC simulation was run for in total 70 000 iterations with the first 20 000 iterations removed as burnin iterations. The trajectory of the negative log-likelihood is shown in Fig. 18 collectively for the independent chains. The convergence of the Markov chains is rapid and consistent, in the Figure we show two zoomed regions of the negative log likelihood showing that convergence is likely

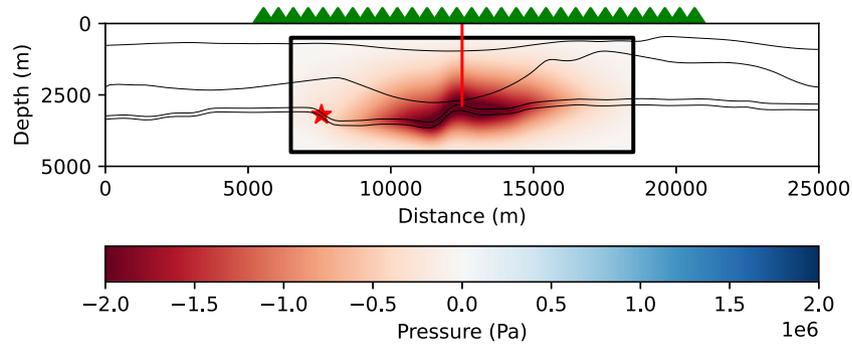


Figure 17. Schematic of the induced stress pressure field used for the calculation of synthetic data. The extraction well shown in red creates an induced stress (pressure drop) that radiates from the base of the well.

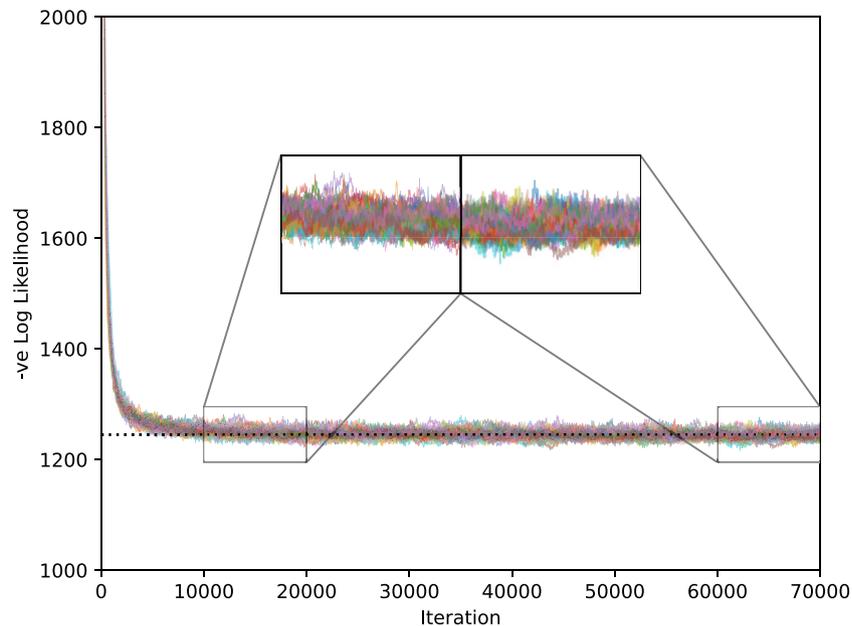


Figure 18. The trajectory of the negative log likelihood as a function of iteration number. Each independent Markov chain is shown with a different colour for each. The convergence is rapid and we show that the chains have essentially converged after approximately 20 000 iterations with the two zoomed regions.

reached sometime between 10 000 and 20 000 iterations with each iteration taking a fraction over 2 s. In addition, this is a serial calculation and requires very little memory, meaning that many such solutions could be computed in parallel with ease and the format of the calculation is suitable for GPU computing.

In Fig. 19, we show the initial model generated seismogram in (a) and the seismograms generated from the mean model. The perturbations to the wavefield from the induced stresses are small and this in general explains the rapid convergence. While generally speaking the effects of induced stress are small, in a more realistic problem there would be modelling errors, that is, uncertainties caused by imperfect knowledge of elastic moduli and densities, that would make such an inversion more complex.

6.4 Results

The tomographic results of the MCMC inversion are shown in Fig. 20 with the posterior mean shown in (b) and standard deviation shown in (c). The recovery of the central region of the induced stress field is reasonable however the broader and lower intensity tails of the induced stress field are poorly recovered. The standard deviation

plot in (c) shows a marked increase in uncertainty below the second interface. While there is a contribution to this sudden increase in uncertainty due to the geometry of the source and receivers, a primary cause of this jump in uncertainty is the change in elastic parameters across the second interface. The increase in elastic moduli across this interface result in a relative change in the influence of induced stress on the wavefield, that is, as the elastic moduli increase in magnitude, a 1 MPa induced pressure for example will have less of an effect on observed wavefields. The resulting standard deviation demonstrates this quite clearly.

In the construction and inversion of the induced stress field we have focused on the induced pressure. The rationale for doing so is in the application area, the change in induced pressure is both expected to dominate the induced stress field, and that the propagation of seismic energy is affected more by changes in pressure than changes in deviatoric components. However, the inversion obtains estimates of all components of the induced stress tensor. In Fig. 21, we show the recovered mean models of two of the deviatoric stress components, τ_{xx} and τ_{xz} , against the true model. In this case, the recovery of the deviatoric components appears quite good across the range of the embedded mesh.

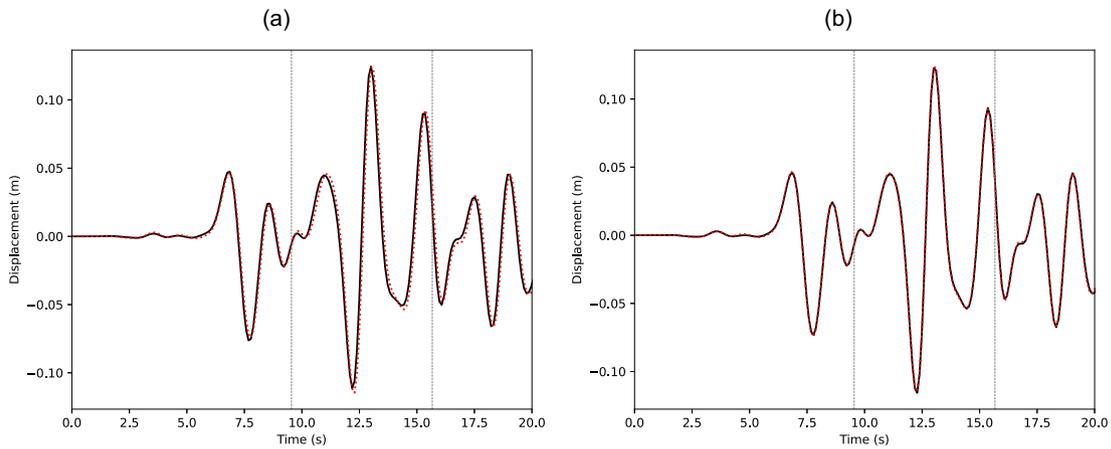
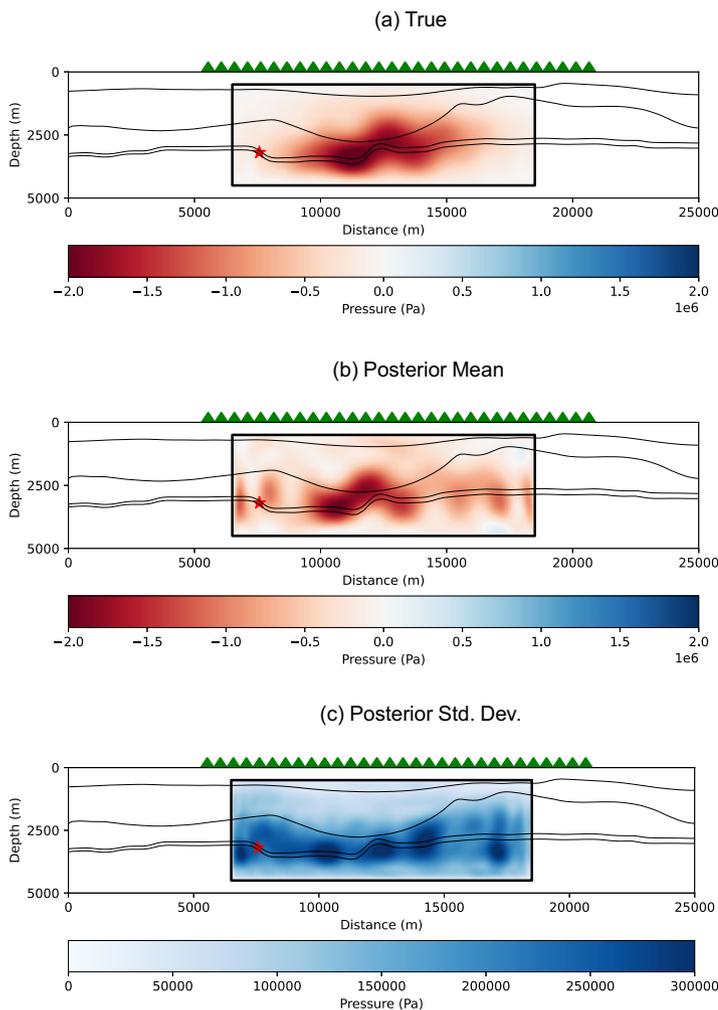


Figure 19. An example of the seismogram fits for a random starting model sampled from the prior in (a) and the converged mean model in (b). The perturbations due to stress are small and cause only small changes in the observed seismograms, however the approximate method used to compute the seismograms is able to be effectively used in a difficult inverse problem.



(d) The posterior mean (a) and standard deviation pressure (b) from the MCMC inversion of synthetic data generated from the true model (a).

Figure 20. The posterior mean (a) and standard deviation pressure (b) from the MCMC inversion of synthetic data generated from the true model (a).

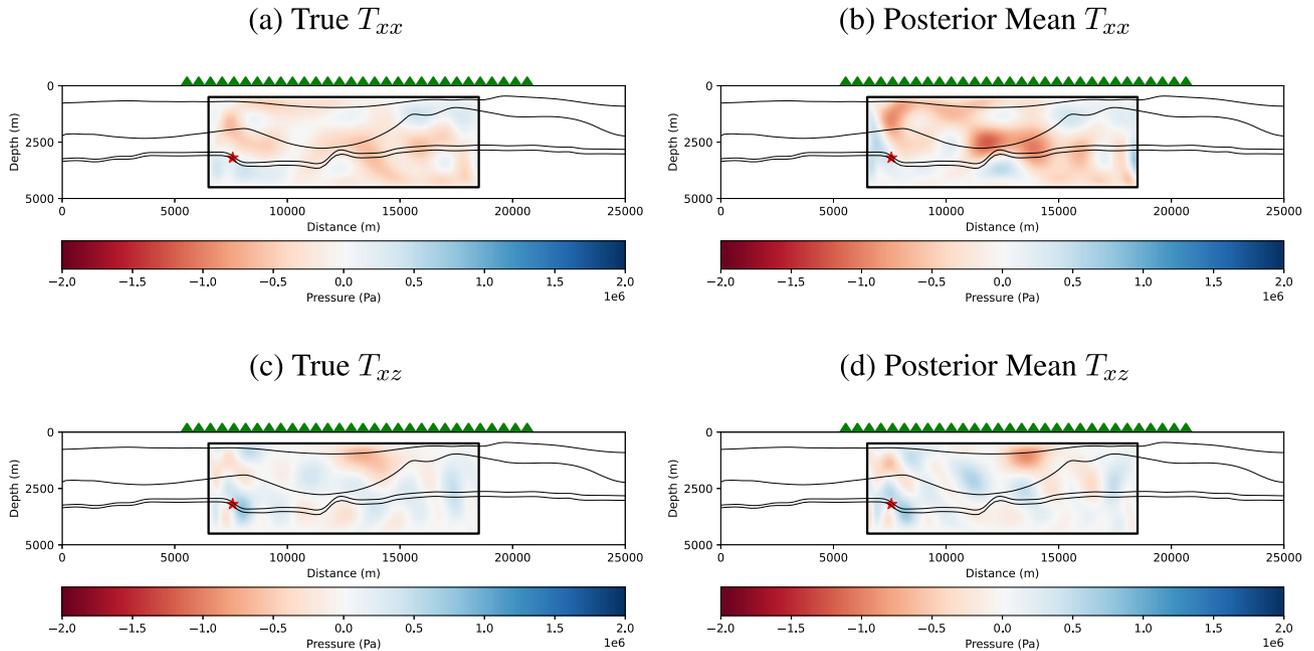


Figure 21. Comparison plots between the true model deviatoric induced stresses and posterior mean models.

6.5 Summary

In this section, we have demonstrated the use of a direct integration based on reduced eigenmodes coupled with a reduced basis encoding equilibrium stress fields to solve a synthetic inverse problem with a simple MCMC approach. The speed at which this inverse problem can be run, and potential ways it can be further improved upon, show that this method has promise in applications of full-waveform monitoring. We have chosen induced stress monitoring as a target and generated a model basis through the use of the null space operator to obtain a space of valid stress models. However we could equally apply this procedure to other interesting sparse parametrizations (Hawkins & Sambridge 2015) and other monitoring applications such as carbon sequestration, volcanic magma chamber events, and near surface effects (Fokker *et al.* 2021).

The synthetic problem as constructed assumes perfect knowledge of the underlying true model. In a realistic setting, there would need to be an additional accounting of errors in the underlying model and some way of partitioning observed changes in the seismograms due to systematic model errors versus actual stress changes. However the main point of this example was to demonstrate the use of a direct integration based on reduced eigenmodes in an inverse problem and how it applies to full waveform monitoring applications.

While the changes in the wavefield owing to induced stresses are small, we have demonstrated that a reduced basis and eigen-projection method are still of a high enough fidelity to be used for their inversion. The main point is that the computational efficiency is dramatically improved in terms of speed and memory requirements, while some degree of accuracy is sacrificed. However, the subtle changes in the seismic waves are still able to be relatively accurately inverted for reasonable tomographic approximations of the true input induced stress model.

This example showed the inversion of an induced stress field using one earthquake source. In passive monitoring applications it would be more appropriate to use passive ambient seismic noise. We have repeated this experiment with a source at the surface as an approximation of the sensitivities of ambient noise data to the same

induced stress and geometry and the results are shown in Fig. A2. In this second example the recovery is not as good owing to poorer sensitivity of the ambient noise wavefield to the induced stress field at depth.

7 DISCUSSION AND CONCLUSIONS

In this manuscript we have demonstrated the link between traditional NM theory and MOR techniques: full coupling in NM seismology can be thought of as an early MOR method where we perform a Galerkin projection onto a set of reduced basis functions, which in NM theory are simply eigenmodes of a simple model of the Earth. We also show that free body oscillations can also be used in the context of regional seismic studies to generate a Petrov–Galerkin MOR scheme. While the computational expense in terms of solving for eigensolutions is now more tractable owing to recent advances in eigensolvers, namely FEAST, it is still not cost competitive compared to time domain solvers more generally. However, as demonstrated with our synthetic example, certain more niche applications involving near real-time applications such as monitoring or generating many realizations, using an eigen-based approach may nonetheless still be feasible and in some cases preferable where the initial cost of eigendecomposition can be amortized over many subsequent waveform calculations.

More generally, the field of MOR research is very active with many recent publications focused on the wave equation in different contexts but broadly applicable to seismology (Glas *et al.* 2020; Bigoni & Hesthaven 2020; Henning *et al.* 2022). In this study we have used eigenvectors as a means to project the system of equations onto a lower dimensional space, however any choice of vectors can be made. The potential advantage of MOR techniques over the free body oscillation approach used here is twofold: first, solving for suitable vectors can be more efficient meaning that high-dimensional problems can be more tractable, and secondly, MOR techniques are generally able to find spaces of lower dimension than free body oscillation techniques for equivalent error levels. This suggests that traditional global NM seismology might also

benefit from recent advances from the MOR community, and that hybrid approaches utilizing MOR and eigensolutions could be used to extend the range of applicability of global NM studies.

At the regional scale, the approach demonstrated here invites comparison to more common frequency domain approaches for exploration scale problems (Pratt 1999). The solution of eigenmodes requires the solution of a discrete number of complex frequencies which is in essence the same as solving the problem in the frequency domain. There are two key differences however: firstly, and this is a subtle point, for the eigensolution we compute solutions for the homogeneous equation (source free) meaning that the solutions can be used for an arbitrary source. This is in some way true for the frequency domain approach as typically for a given frequency the left-hand side is decomposed into an LU factorization which then makes solving for any source relatively inexpensive compared to the initial factorization. A more compelling second point is that we can reuse eigensolutions from a reference model in modelling wave propagation with nearby models using a Rayleigh–Ritz approach. We have demonstrated this with perturbations to elastic moduli and density and incorporated this ability in an inversion of induced stress in a synthetic example.

In the inverse problem presented, we used a very simple MCMC approach. Alternative approaches to the use of eigensolutions for the solution of inverse problems exist (Grote & Nahum 2019) and the use of eigensolutions in computing forward models does not preclude the use of adjoint state techniques in a seismology context (Capdeville 2005; Hawkins 2018; Hawkins & Sambridge 2019). A subject of future work would be the extension of methods presented here for computing relevant gradients using MOR and NM techniques. Given both the speed and memory efficiency demonstrated in the induced stress example, in that we can trivially compute many realizations with different model parametrizations (in approximately two seconds on a single core), another benefit of the approaches outlined here is to be able to generate ensembles of training data that can be used in rapid hazard determination using prior sampling (Käufel *et al.* 2016) or similarly for the generation of large amounts of data that can be used in machine learning. Alternatively, the use of a reference set of eigensolutions could be used in data-driven applications to directly estimate true model parameters through an autoregressive approach (Masters *et al.* 2000).

A final point is that this manuscript has established through some theory and examples, the use of free body oscillation approaches to regional studies and linked this to MOR techniques. The problems demonstrated herein are small compared to those typically tackled with time domain approaches. Scaling eigensolvers or indeed MOR techniques to one or two orders of magnitude larger than presented here is not a trivial problem. To demonstrate this, we have computed the eigensolutions for a range of 3-D homogeneous models of increasing size and the computational and memory requirements are shown in Appendix C. We can reliably compute the eigensolutions for a 3-D model with a size of 1 million degrees of freedom. Typically, large 3-D simulations will be of the order of 100 million degrees of freedom or more. For such large problems, rather than treating the problem as a whole, we believe that domain decomposition could be used as has already been demonstrated in a MOR context (Babuška & Lipton 2011; Buhr *et al.* 2020; Ma *et al.* 2022; Schleuß & Smetana 2022; Smetana & Patera 2016), and in the context of coupling spectral element solutions to NM approaches (Capdeville *et al.* 2003). Another motivation for the domain decomposition approach is that in many seismic wave simulations we require coupling between fluid and solid regions. In this manuscript, we have demonstrated that free body oscillations can

be applied in 3-D elastic, 2-D P–SV and 2-D SH simulations and physically speaking SH simulations are analogous to wave propagation in fluids, so in principle fluid–solid simulations could be treated globally. However, a more scalable approach, both from a performance and flexibility point of view, would be to pursue domain decomposition strategies that also allow different physical models in the different subdomains.

SOFTWARE AVAILABILITY

The python scripts for running the small example are available from <https://github.com/rhyshawkins/MORNormalModeExample>.

SpecFEM2D is available at <https://github.com/geodynamics/specfem2d>

The FEAST eigensolver code is available from its authors at <https://www.feast-solver.org/>

ACKNOWLEDGMENTS

We would like to thank Prof Yann Capdeville, Prof Jeroen Tromp and the editor Prof Carl Tape for their thoughtful and constructive comments on this paper.

This work is part of research programme DeepNL, financed by the Dutch Research Council (NWO) under project number DeepNL.2018.033.

This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government.

DATA AVAILABILITY

The synthetic data generated for this study are available from the authors upon reasonable request.

REFERENCES

- Afanasiev, M., Boehm, C., van Driel, M., Krischer, L., Rietmann, M., May, D.A., Knepley, M.G. & Fichtner, A., 2019. Modular and flexible spectral-element waveform modelling in two and three dimensions, *J. geophys. Int.*, **216**(3), 1675–1692.
- Akbarashrafi, F., Al-Attar, D., Deuss, A., Trampert, J. & Valentine, A.P., 2017. Exact free oscillation spectra, splitting functions and the resolvability of Earth’s density structure, *J. geophys. Int.*, **213**(1), 58–76.
- Al-Attar, D. & Woodhouse, J.H., 2010. On the parametrization of equilibrium stress fields in the Earth, *J. geophys. Int.*, **181**(1), 567–576.
- Al-Attar, D., Woodhouse, J.H. & Deuss, A., 2012. Calculation of normal mode spectra in laterally heterogeneous earth models using an iterative direct solution method, *J. geophys. Int.*, **189**(2), 1038–1046.
- Anderson, E. *et al.*, 1999. *LAPACK Users’ Guide*, 3rd edn, Society for Industrial and Applied Mathematics.
- Babuška, I. & Lipton, R., 2011. Optimal local approximation spaces for generalized finite element methods with application to multiscale problems, *Multiscale Model. Simul.*, **9**(1), 373–406.
- Beghein, C. & Trampert, J., 2003. Robust normal mode constraints on inner-core anisotropy from model space search, *Science*, **299**(5606), 552–555.
- Benner, P., Ohlberger, M., Cohen, A. & Willcox, K., 2016. *Model Reduction and Approximation: Theory and Algorithms*, SIAM.
- Benner, P., Grivet-Talocia, S., Quarteroni, A., Rozza, G., Schilders, W. & Silveira, L.M., 2020a. *Model Order Reduction, Vol. 3: Applications*, De Gruyter.
- Benner, P., Grivet-Talocia, S., Quarteroni, A., Rozza, G., Schilders, W. & Silveira, L.M., 2020b. *Model Order Reduction, Vol. 2: Snapshot-Based Methods and Algorithms*, De Gruyter.

- Berkooz, G., Holmes, P. & Lumley, J.L., 1993. The proper orthogonal decomposition in the analysis of turbulent flows, *Annu. Rev. Fluid Mech.*, **25**(1), 539–575.
- Berland, J., Bogey, C. & Bailly, C., 2006. Low-dissipation and low-dispersion fourth-order Runge–Kutta algorithm, *Comput. Fluids*, **35**(10), 1459–1463.
- Bigoni, C. & Hesthaven, J.S., 2020. Simulation-based anomaly detection and damage localization: an application to structural health monitoring, *Comput. Methods Appl. Mech. Eng.*, **363**, 112896, doi:10.1016/j.cma.2020.112896.
- Buhr, A., Iapichino, L., Ohlberger, M., Rave, S., Schindler, F. & Smetana, K., 2020. Localized Model Reduction for Parameterized Problems, in *Model Order Reduction. Vol. 2 Snapshot-Based Methods and Algorithms*, pp. 245–306, eds Benner, P., Grivet-Talocia, S., Quarteroni, A., Rozza, G., Schilders, W.H.A. & Sileira, L.M., De Gruyter.
- Capdeville, Y., 2005. An efficient Born normal mode method to compute sensitivity kernels and synthetic seismograms in the Earth, *J. geophys. Int.*, **163**(2), 639–646.
- Capdeville, Y., Chaljub, E. & Montagner, J.P., 2003. Coupling the spectral element method with a modal solution for elastic wave propagation in global earth models, *J. geophys. Int.*, **152**(1), 34–67.
- Clayton, R. & Engquist, B., 1977. Absorbing boundary conditions for acoustic and elastic wave equations, *Bull. seism. Soc. Am.*, **67**(6), 1529–1540.
- Courant, R., Friedrichs, K. & Lewy, H., 1928. Über die partiellen differenzgleichungen der mathematischen physik, *Math. Ann.*, **100**, 32–74.
- Cummins, P.R., Takeuchi, N. & Geller, R.J., 1997. Computation of complete synthetic seismograms for laterally heterogeneous models using the direct solution method, *J. geophys. Int.*, **130**(1), 1–16.
- Dahlen, F.A. & Tromp, J., 1998. *Theoretical Global Seismology*, Princeton Univ. Press.
- Dahlen, F.A., Hung, S.-H. & Nolet, G., 2000. Fréchet kernels for finite-frequency travel times - I. theory, *Geophys. J. Int.*, **141**, 157–174.
- Deuss, A., Irving, J.C. & Woodhouse, J.H., 2010. Regional variation of inner core anisotropy from seismic normal mode observations, *Science*, **328**(5981), 1018–1020.
- Fichtner, A., 2011. *Full Seismic Waveform Modelling and Inversion*, Springer.
- Fichtner, A., Igel, H., Bunge, H.P. & Kennett, B. L.N., 2009a. Simulation and inversion of seismic wave propagation on continental scales based on a spectral-element method, *J. Numer. Anal., Indust. Appl. Math.*, **4**(1–2), 11–22.
- Fichtner, A., Kennett, B. L.N., Igel, H. & Bunge, H.-P., 2009b. Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods, *J. geophys. Int.*, **179**(3), 1703–1725.
- Fokker, E., Ruigrok, E., Hawkins, R. & Trampert, J., 2021. Physics-based relationship for pore pressure and vertical stress monitoring using seismic velocity variations, *Remote Sens.*, **13**(14), 2684, doi:10.3390/rs13142684.
- Gavin, B. & Polizzi, E., 2018. Krylov eigenvalue strategy using the feast algorithm with inexact system solves, *Numer. Linear Algebra Appl.*, **25**(5), e2188, doi:10.1002/nla.2188.
- Gavin, B., Miedlar, A. & Polizza, E., 2018. Feast eigensolver for nonlinear eigenvalue problems, *J. Comput. Sci.*, **27**, 107–117.
- Geuzaine, C. & Remacle, J.-F., 2009. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities, *Int. J. Numer. Methods Eng.*, **79**(11), 1309–1311.
- Gilbert, F., 1970. Excitation of the normal modes of the Earth by earthquake sources, *Geophys. J. R. astr. Soc.*, **22**, 223–226.
- Glas, S., Patera, A.T. & Urban, K., 2020. A reduced basis method for the wave equation, *Int. J. Comput. Fluid Dyn.*, **34**(2), 139–146.
- Grote, M.J. & Nahum, U., 2019. Adaptive eigenspace for multi-parameter inverse scattering problems, *Comput. Math. Appl.*, **77**(12), 3264–3280.
- Hammarling, S., Munro, C.J. & Tisseur, F., 2013. An algorithm for the complete solution of quadratic eigenvalue problems, *ACM Trans. Math. Softw.*, **39**(3), 1–19, doi:10.1145/2450153.2450156.
- Hawkins, R., 2018. A spectral element method for surface wave dispersion and adjoints, *J. geophys. Int.*, **215**(1), 267–302.
- Hawkins, R. & Sambridge, M., 2015. Geophysical imaging using trans-dimensional trees, *J. geophys. Int.*, **203**(2), 972–1000.
- Hawkins, R. & Sambridge, M., 2019. An adjoint technique for estimation of interstation phase and group dispersion from ambient noise cross correlations, *Bull. seism. Soc. Am.*, **109**(5), 1716–1728.
- Henning, J., Palitta, D., Simoncini, V. & Urban, K., 2022. An ultraweak space-time variational formulation for the wave equation: analysis and efficient numerical solution, *ESAIM Math. Model. Numer. Anal.*, **56**(4), 1173–1198.
- Hesthaven, J.S., Rozza, G. & Stamm, B., 2016. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, SpringerBriefs in Mathematics, Springer International Publishing.
- Hung, S.-H., Dahlen, F.A. & Nolet, G., 2001. Wavefront healing: a banana-doughnut perspective, *Geophys. J. Int.*, **146**, 289–312.
- Iollo, A. & Lombardi, D., 2014. Advection modes by optimal mass transfer, *Phys. Rev. E*, **89**(2), 022923, doi:10.1103/PhysRevE.89.022923.
- Käuffl, P., Valentine, A.P., de Wit, R.W. & Trampert, J., 2016. Solving probabilistic inverse problems rapidly with prior samples, *J. geophys. Int.*, **205**(3), 1710–1728.
- Kestyn, J., Polizzi, E. & Tang, P.P.T., 2016. Feast eigensolver for non-Hermitian problems, *SIAM J. Sci. Comput.*, **38**(5), S772–S799.
- Komatitsch, D. & Tromp, J., 1999. Introduction to the spectral element method for three-dimensional seismic wave propagation, *Geophys. J. Int.*, **139**, 806–822.
- Komatitsch, D. & Tromp, J., 2002a. Spectral-element simulations of global seismic wave propagation - I. Validation, *Geophys. J. Int.*, **149**, 390–412.
- Komatitsch, D. & Tromp, J., 2002b. Spectral-element simulations of global seismic wave propagation - II. Three-dimensional models, oceans, rotation and self-gravitation, *Geophys. J. Int.*, **149**, 390–412.
- Komatitsch, D., Vilotte, J.-P., Vai, R., Castillo-Covarrubias, J.M. & Sánchez-Sesma, F.J., 1999. The spectral element method for elastic wave equations - applications to 2-D and 3-D seismic problems, *Int. J. Numer. Methods Eng.*, **45**, 1139–1164.
- Komatitsch, D., Ritsema, J. & Tromp, J., 2002. The spectral-element method, beowulf computing, and global seismology, *Science*, **298**(5599), 1737–1742.
- Kunisch, K. & Volkwein, S., 2001. Galerkin proper orthogonal decomposition methods for parabolic problems, *Numer. Math.*, **90**, 117–148.
- Levander, A., 1988. Fourth-order finite-difference P-SV seismograms, *Geophysics*, **53**(11), 1425–1436.
- Ma, C., Scheichl, R. & Dodwell, T., 2022. Novel design and analysis of generalized finite element methods based on locally optimal spectral approximations, *SIAM J. Numer. Anal.*, **60**(1), 244–273.
- Masters, G., Laske, G. & Gilbert, F., 2000. Autoregressive estimation of the splitting matrix of free-oscillation multiplets, *J. geophys. Int.*, **141**, 25–42.
- Masters, G., Woodhouse, J.H. & Freeman, G., 2011. *Mineos v1.0.2, [software] Computational Infrastructure for Geodynamics*.
- Nair, N.J. & Balajewicz, M., 2019. Transported snapshot model order reduction approach for parametric, steady-state fluid flows containing parameter-dependent shocks, *Int. J. Numer. Methods Eng.*, **117**(12), 1234–1262.
- Newmark, N.M., 1959. A method of computation for structural dynamics, *J. Eng. Mech. Div.*, **85**(3), 67–94.
- Nissen-Meyer, T., van Driel, M., Stähler, S.C., Hosseini, K., Hempel, S., Auer, L., Colombi, A. & Fournier, A., 2014. Axisem: broadband 3-D seismic wavefields in axisymmetric media, *Solid Earth*, **5**(1), 425–445.
- Ohlberger, M. & Rave, S., 2013. Nonlinear reduced basis approximation of parameterized evolution equations via the method of freezing, *Comput. Rend. Math.*, **351**(23–24), 901–906.
- Oliveira, S.P. & Seriani, G., 2011. Effect of element distortion on the numerical dispersion of spectral element methods, *Commun. Comput. Phys.*, **9**(4), 937–958.
- Park, J. & Gilbert, F., 1986. Coupled free oscillations of an aspherical, dissipative, rotating earth: Galerkin theory, *J. geophys. Res.*, **91**(B7), 7241–7260.
- Patera, A.T., 1984. A spectral element method for fluid dynamics: laminar flow in a channel expansion, *J. Comput. Phys.*, **54**(3), 468–488.
- Pekeris, C. & Jarosch, H., 1958. The free oscillations of the Earth, in *Contributions in Geophysics in Honor of Beno Butenberg*, Pergamon Press.

- Peter, D. et al., 2011. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes: SPEC3D version 2.0 'sesame', *Geophys. J. Int.*, **186**(2), 721–739.
- Polizzi, E., 2009. Density-matrix-based algorithm for solving eigenvalue problems, *Phys. Rev. B*, **79**, 115112, doi:10.1103/PhysRevB.79.115112.
- Postma, T. & Jansen, J.D., 2018. The small effect of poroelastic pressure transients on triggering of production-induced earthquakes in the Groningen natural gas field, *J. geophys. Res.*, **123**(1), 401–417.
- Pratt, R.G., 1999. Seismic waveform inversion in the frequency domain, part 1: theory and verification in a physical scale model, *Geophysics*, **64**(3), 888–901.
- Quarteroni, A., Manzoni, A. & Negri, F., 2016. *Reduced Basis Methods for Partial Differential Equations: An Introduction*, Springer.
- Reiss, J., Schulze, P., Sesterhenn, J. & Mehrmann, V., 2018. The shifted proper orthogonal decomposition: a mode decomposition for multiple transport phenomena, *SIAM J. Sci. Comput.*, **40**(3), A1322–A1344.
- Romijn, R., 2017. Groningen velocity model 2017, Tech. rep., Nederlandse Aardolie Maatschappij, Assen, The Netherlands.
- Roza, G., Huynh, D. & Patera, A., 2008. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics, *Arch. Comput. Meth. Eng.*, **15**(3), 229–275.
- Schleuß, J. & Smetana, K., 2022. Optimal local approximation spaces for parabolic problems, *Multiscale Model. Simul.*, **20**(1), 551–582.
- Seriani, G. & Priolo, E., 1994. A spectral-element method for acoustic wavesimulation in heterogeneous media, *Finite Elem. Anal. Des.*, **16**, 337–348.
- Seriani, G., Priolo, E., Carcione, J. & Padovani, E., 1992. High-order spectral element method for elastic wave modeling, in *SEG Technical Program Expanded Abstracts*, pp. 1285–1288, Society of Exploration Geophysicists.
- Sirovich, L., 1987. Turbulence and the dynamics of coherent structures. I. Coherent structures, *Quart. Appl. Math.*, **45**(3), 561–571.
- Smetana, K. & Patera, A.T., 2016. Optimal local approximation spaces for component-based static condensation procedures, *SIAM J. Sci. Comput.*, **38**(5), A3318–A3356.
- Stacey, R., 1988. Improved transparent boundary formulations for the elastic-wave equation, *Bull. seism. Soc. Am.*, **78**(6), 2089–2097.
- Taddei, T., 2020. A registration method for model order reduction: data compression and geometry reduction, *SIAM J. Sci. Comput.*, **42**(2), A997–A1027.
- Tarantola, A., 1984. Inversion of seismic reflection data in the acoustic approximation, *Geophysics*, **49**(8), 1259–1266.
- Tarantola, A. & Valette, B., 1982. Generalized nonlinear inverse problems solved using the least squares criterion, *Rev. Geophys.*, **20**(2), 219–232.
- Tharastarson, S., van Driel, M., Krischer, L., Boehm, C., Afanasiev, M., van Herwaarden, D.-P. & Fichtner, A., 2020. Accelerating numerical wave propagation by wavefield adapted meshes. Part II: full-waveform inversion, *J. geophys. Int.*, **221**(3), 1591–1604.
- Tromp, J. & Trampert, J., 2018. Effects of induced stress on seismic forward modelling and inversion, *J. geophys. Int.*, **213**(2), 851–867.
- Tromp, J., Komatitsch, D. & Liu, Q., 2008. Spectral-element and adjoint methods in seismology, *Commun. Comput. Phys.*, **3**(1), 1–32.
- van Driel, M., Boehm, C., Krischer, L. & Afanasiev, M., 2020. Accelerating numerical wave propagation using wavefield adapted meshes. Part I: forward and adjoint modelling, *J. geophys. Int.*, **221**(3), 1580–1590.
- Veroy, K., Prud'homme, C., Rovas, D.V. & Patera, A.T., 2003. A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations, in *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference*, Vol. **3847**.
- Virieux, J., 1986. P-SV wave propagation in heterogeneous media: velocity-stress finite-difference method, *Geophysics*, **51**(4), 889–901.
- Virieux, J. & Operto, S., 2009. An overview of full-waveform inversion in exploration geophysics, *Geophysics*, **74**(6), WCC1–WCC26.
- Virieux, J., Calandra, H. & Plessix, R.-E., 2011. A review of the spectral, pseudo-spectral, finite-difference and finite-element modelling techniques for geophysical imaging, *Geophys. Prospect.*, **59**(5), 794–813.

- Welper, G., 2017. Interpolation of functions with parameter dependent jumps by transformed snapshots, *SIAM J. Sci. Comput.*, **39**(4), A1225–A1250.
- Welper, G., 2020. Transformed snapshot interpolation with high resolution transforms, *SIAM J. Sci. Comput.*, **42**(4), A2037–A2061.
- Woodhouse, J., 1988. The calculation of the eigenfrequencies and eigenfunctions of the free oscillations of the Earth and the Sun, in *Seismological Algorithms*, Academic Press.

APPENDIX A: CONSTRUCTING SOLUTIONS FOR THE COMPANION FORM DIRECTLY FROM QUADRATIC SOLUTIONS

An advantage of the FEAST algorithm is that solutions of the quadratic eigenvalue problem can be computed directly, rather than having to form companion matrices. However, for the projection of the eigensystem, it is preferable to use a companion form as the resulting projected system of equation is diagonal whereas direct projection of the quadratic matrix equation results in dense matrices and hence poorer performance owing to the larger number of floating point operations.

Given a quadratic system of the form

$$(\lambda_i^2 \mathbf{M} + \lambda_i \mathbf{C} + \mathbf{K}) \phi_i = 0 \quad (\text{A1})$$

with matrices \mathbf{M} , \mathbf{C} and \mathbf{K} symmetric, and λ_i and ϕ_i a right eigenvalue and eigenvector, respectively. We also know due to the matrices being symmetric that $\phi_i^* (\bar{\lambda}_i^2 \mathbf{M} + \bar{\lambda}_i \mathbf{C} + \mathbf{K}) = 0$, where the bar represents the complex conjugate. Given the use of the companion form

$$\lambda_i \begin{bmatrix} -\mathbf{M} & 0 \\ 0 & -\mathbf{I} \end{bmatrix} \mathcal{R}_i = \begin{bmatrix} \mathbf{C} & \mathbf{K} \\ -\mathbf{I} & 0 \end{bmatrix} \mathcal{R}_i \quad (\text{A2})$$

we can construct the right eigenvectors of the companion form directly from the quadratic eigensolutions λ_i and ϕ_i using

$$\mathcal{R}_i = \begin{bmatrix} \lambda_i \phi_i \\ \phi_i \end{bmatrix} \quad (\text{A3})$$

and similarly for the left eigenvectors,

$$\mathcal{L}_i = \begin{bmatrix} \bar{\lambda}_i \bar{\phi} \\ -(\bar{\phi}^* \mathbf{K})^T \end{bmatrix} \quad (\text{A4})$$

and normalize these vectors such that

$$\mathcal{L}_i^* \begin{bmatrix} -\mathbf{M} & 0 \\ 0 & -\mathbf{I} \end{bmatrix} \mathcal{R}_j = \delta_{i,j}. \quad (\text{A5})$$

A1 Rigid body modes

The preceding construction works for non-zero eigensolutions. In a system without Dirichlet boundary conditions, that is, with no fixed nodes, there will be a set of eigensolutions where $\lambda = 0$. These are generally called 'rigid body modes' as these solutions represent translations and rotations of the system. If not correctly treated in a truncated set of eigensolutions, these can cause large errors in synthetic seismograms.

By inspection of eq. (A1), these solutions satisfy

$$\mathbf{K} \phi_i^r = 0, \quad (\text{A6})$$

where the superscript r indicates a rigid body solution. In a 2-D P-SV system, there will be 3 rigid body modes (2 orthogonal translations and a rotation) whereas in a 3-D P-SV system there will be 6 (3 orthogonal translations and 3 orthogonal rotations). Given some

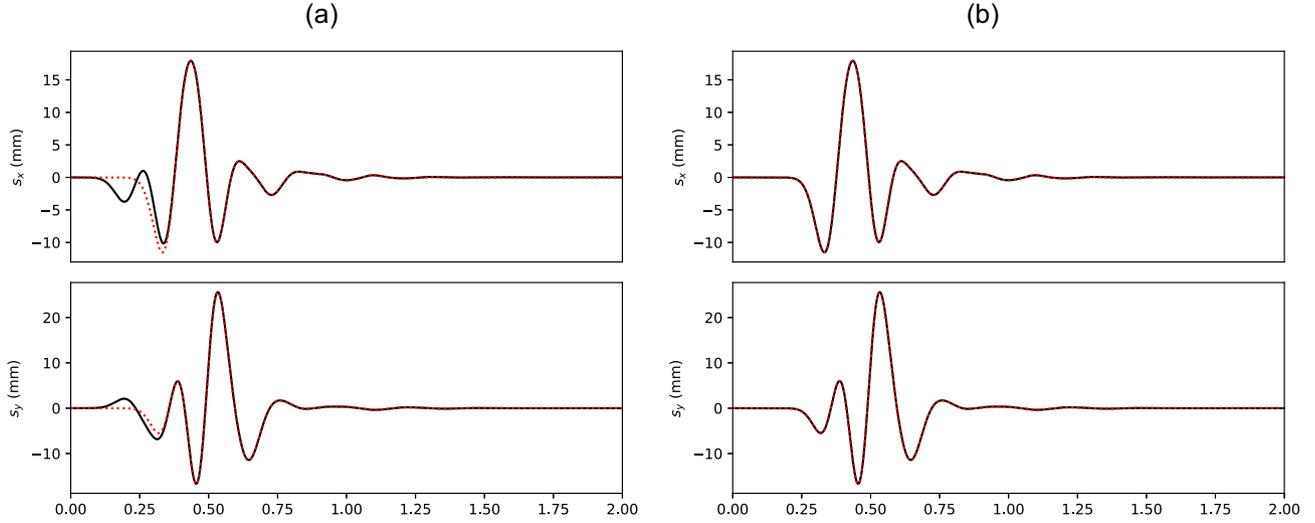


Figure A1. The seismograms for the horizontal (top panels) and vertical (bottom panels) components calculated using mode summation is shown in black. The same seismograms calculated using a time domain numerical scheme is shown in red dotted lines for comparison. In (a) the rigid body modes are not included whereas in (b) the rigid body modes eigenvectors for the companion form are calculated using the method in Section A1.

number of rigid body modes it is evident that if for example, Θ is a linear combination of these modes, that $K\Theta = 0$ as well. We can use this fact to construct valid left and right eigenvectors of the companion form directly from the rigid body modes of the quadratic equation.

To first determine the appropriate linear weighting to use, we form the matrix

$$\hat{\mathbf{C}} = [\phi_1^r \ \phi_2^r \ \dots \ \phi_n^r]^T \mathbf{C} [\phi_1^r \ \phi_2^r \ \dots \ \phi_n^r] \quad (\text{A7})$$

and compute its singular value decomposition to obtain unitary matrices \mathbf{U} and \mathbf{V} such that $\hat{\mathbf{C}} = \mathbf{U}\mathbf{s}\mathbf{V}^T$ where \mathbf{s} is the diagonal matrix of singular values. Assuming matrix $\hat{\mathbf{C}}$ is full rank, the columns of matrices \mathbf{U} and \mathbf{V} can be used to construct left and right eigenvectors of the companion form. If V_i represents the i th column of matrix \mathbf{V} , the i th right eigenvector of the companion form is given by

$$\mathcal{R}_i = \begin{bmatrix} 0 \\ V_i^T [\phi_1^r \ \phi_2^r \ \dots \ \phi_n^r] \end{bmatrix}, \quad (\text{A8})$$

and similarly for the left eigenvector

$$\mathcal{L}_i = \begin{bmatrix} U_i^T [\phi_1^r \ \phi_2^r \ \dots \ \phi_n^r] \\ \mathbf{C} (U_i^T [\phi_1^r \ \phi_2^r \ \dots \ \phi_n^r]) \end{bmatrix}. \quad (\text{A9})$$

These rigid body vectors can be normalized in the same way as in eq. (A5) and it can be trivially demonstrated that these rigid body vectors are mutually bi-orthogonal to the non-rigid body eigenvectors giving a stable and effective way to construct companion form eigenvectors from solutions to the quadratic eigenequation. The one caveat is that the matrix \mathbf{C} must have sufficient entries to ensure that the matrix $\hat{\mathbf{C}}$ has full rank. This will generally be the case in regional simulations as generally multiple boundaries are absorbing and hence many cell faces will contribute to the construction of the \mathbf{C} matrix.

The effect of including properly calculated rigid body modes in either modal summation or for projected matrices is shown in Fig. A1 where in (a) the rigid body modes are not used whereas in (b) they are. The seismograms in the case where the rigid body modes are used are a visually indistinguishable from a time domain solution.

APPENDIX B: INDUCED STRESS REFRACTION

While not directly relevant to the main topic of this manuscript, we also repeated the simple MCMC treatment using a refraction type experiment where the source was located near the surface. The sensitivities of this type of system in some way mimic the sensitivities (particularly to depth) of passive monitoring techniques at the surface. The results of the inversion are shown in Fig. A2 and demonstrate that the input true model is not recovered. In addition, the uncertainties are sufficiently large in the region of interest that suggests that passive experiments would in practice not be able to recover sufficient information about the subsurface stress fields without further advances in modelling or inference techniques. In part this is due to a reasonably strong interface layer directly above the main area of induced stress which in effect reflects a large proportion of energy back towards the receivers, meaning that differences caused by the induced stress field are small relative to this reflected energy.

APPENDIX C: SCALING OF FEAST EIGENSOLVER FOR A 3-D HOMOGENEOUS MODEL

To demonstrate the feasibility of the methods presented here for larger models more applicable to real world applications, we have computed a series of eigensolutions for homogeneous models of increasing size. We do this by creating a regular mesh that has a fixed number of cells in the vertical or depth direction while varying the number of cells laterally. The cells are fixed at 1 km in size with density of 2600 kgm^{-3} , P -wave velocity of 6 kms^{-1} and shear wave velocity of 3.5 km^{-1} . The top surface has a free stress condition with all other surfaces being absorbing boundaries using a Stacey absorbing condition.

C1 Computational cost

For the scaling tests we have performed a weak scaling test where as we increase the problem size, we similarly increase the number

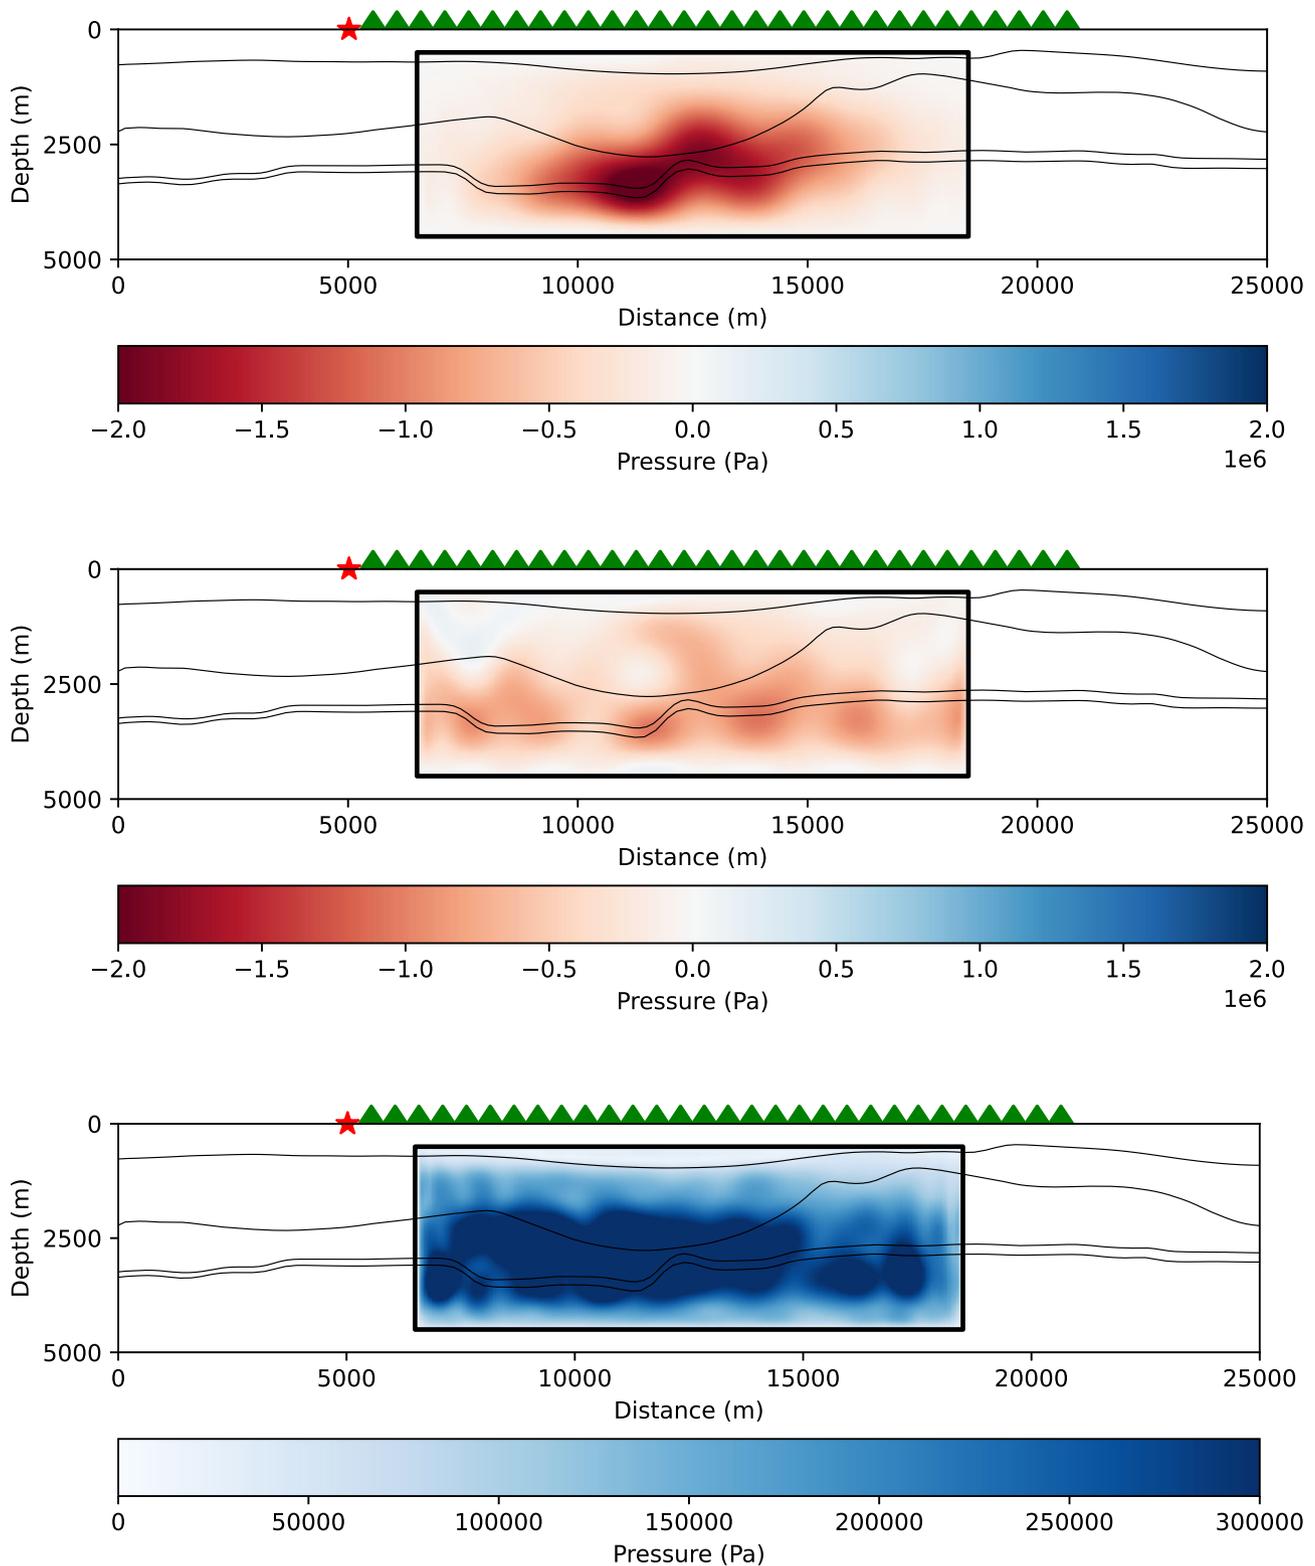


Figure A2. The posterior mean and standard deviation pressure from the MCMC inversion of synthetic data generated from the true model (a).

of parallel processors used to compute the eigensolutions. While strictly speaking we have not kept the work per processor constant while increasing the model size, these scaling tests nonetheless give a good appreciation of the effort required as system size increases.

The total CPU time is computed directly through operating system calls and amortized over all processors used in the parallel solution. The results are plotted in Fig. A3 in a Log–Log plot of total CPU time versus degrees of freedom. The best fitting line to the scaling data has a gradient of approximately 1.6 indicating that

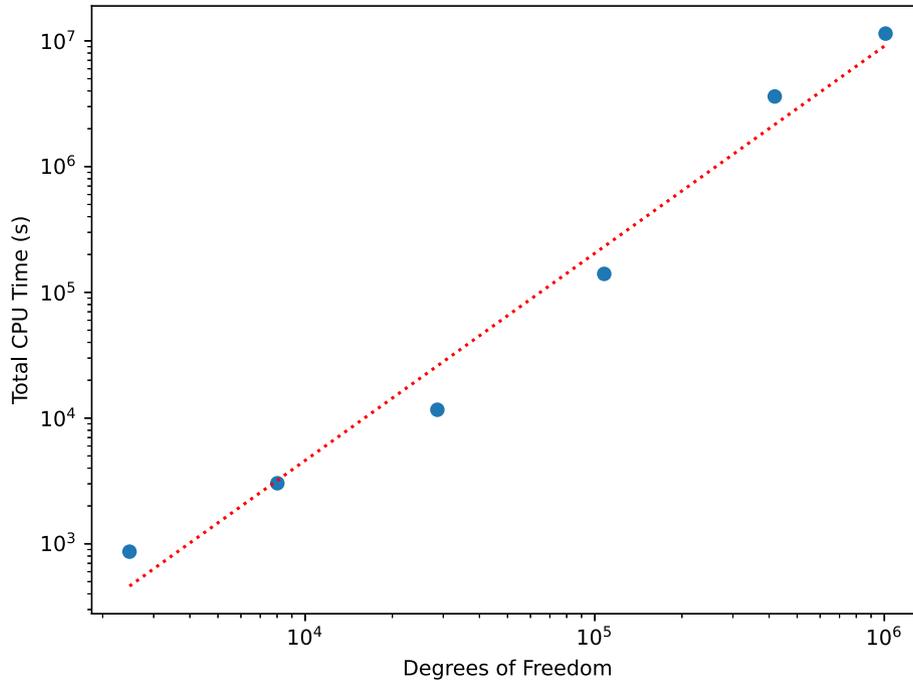


Figure A3. Log plot of the total CPU time used versus the model size expressed as degrees of freedom (DOF). Also shown with a red-dotted line is the best fitting linear fit in log space of the scaling curve which has a slope of approximately 1.6 indicating approximately quadratic scaling of CPU time versus degrees of freedom.

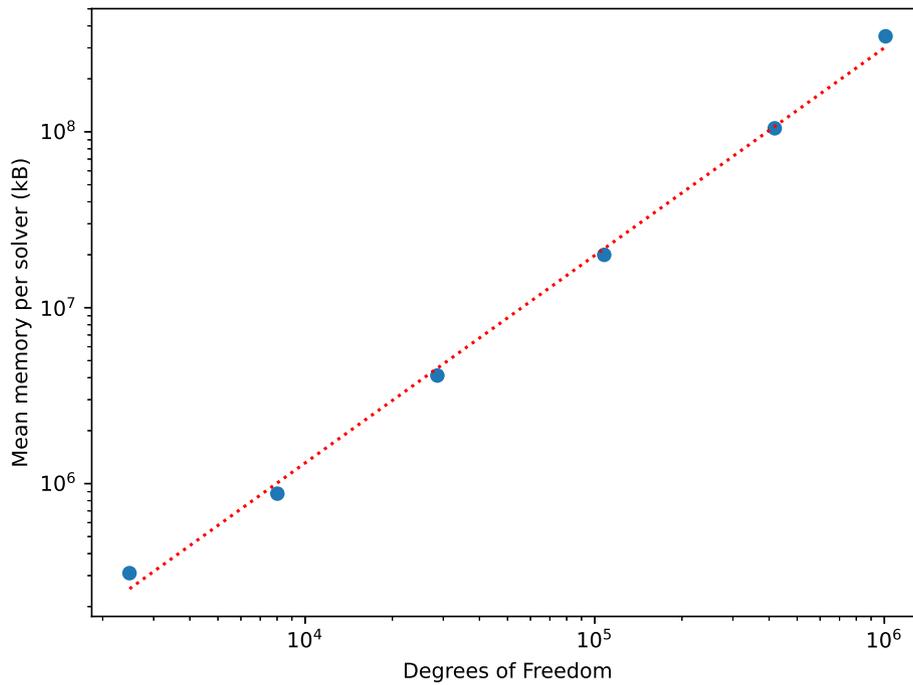


Figure A4. Log plot of the total Memory use versus the model size expressed as degrees of freedom (DOF). Also shown with a red-dotted line is the best fitting linear fit in log space of the scaling curve which has a slope of approximately 1.2 indicating approximately linear scaling of memory usage versus degrees of freedom.

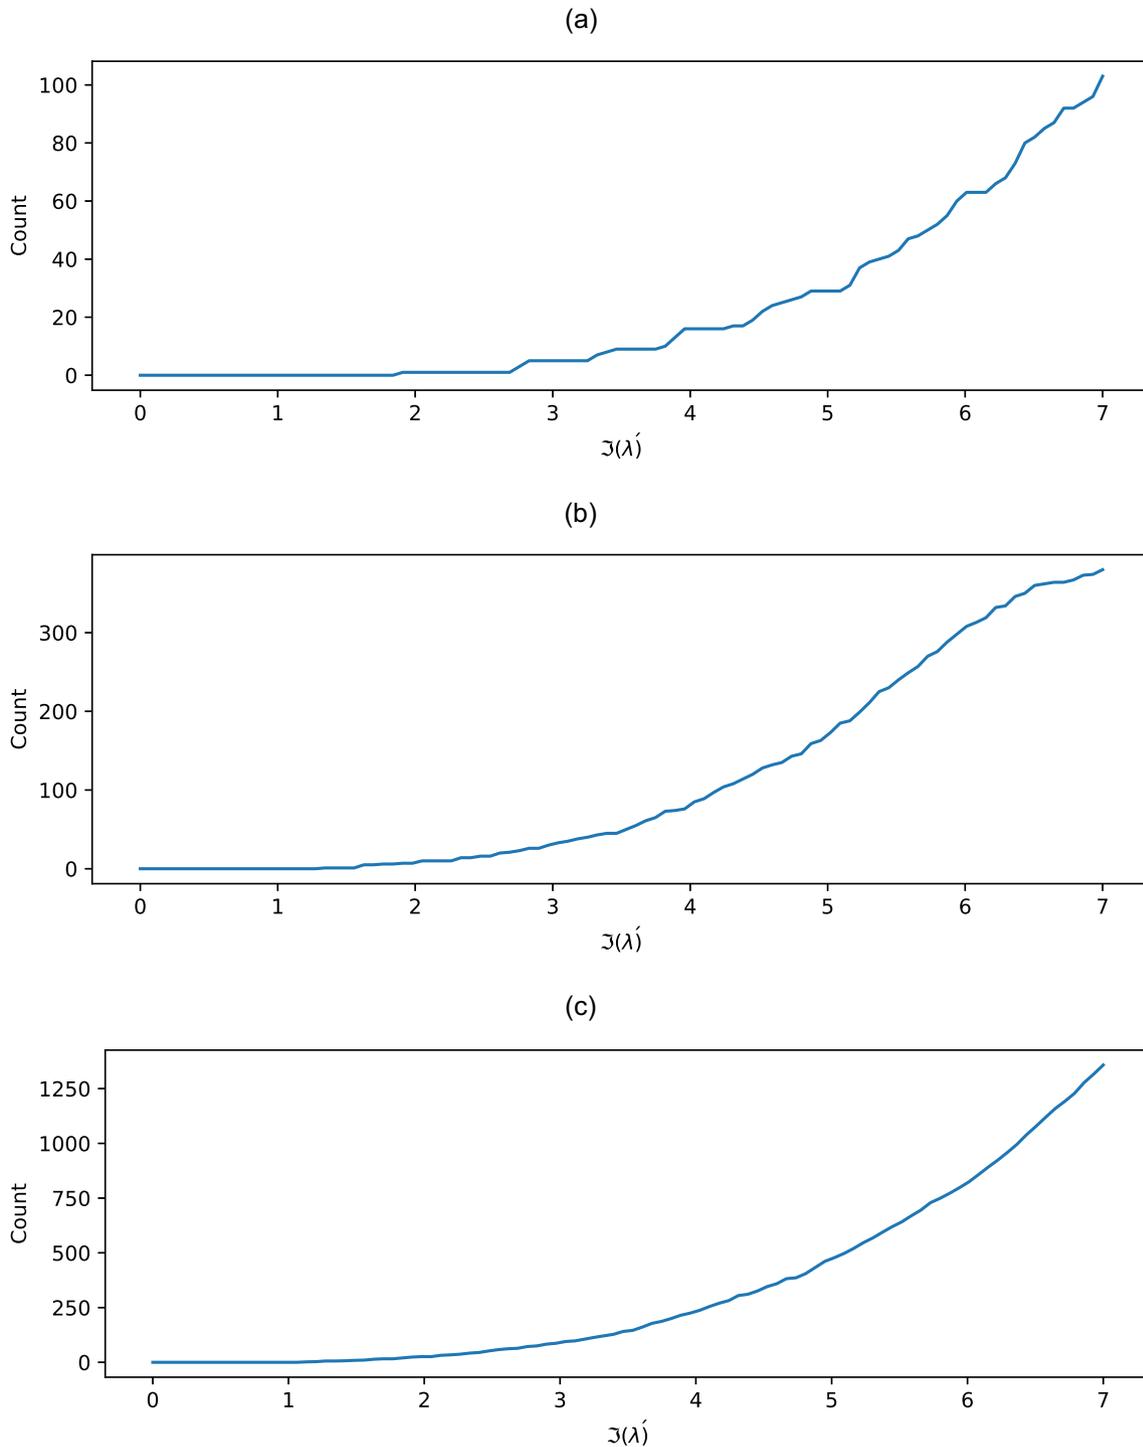


Figure A5. Plot of the cumulative number of eigensolutions as a function of frequency for the 100kdof (a), 400kdof (b) and 1Mdof (c) models. The growth in the number of eigensolutions as a function of frequency is well described by a cubic function.

the CPU timescales approximately quadratically with degrees of freedom.

C2 Memory

Another important consideration is the memory use as a function of model size. From the scaling tests, we also have diagnostics on the total memory use for each of the parallel solvers. We run multiple FEAST solvers in parallel and each reports its own memory usage and we subsequently average these over all solvers to give a mean memory usage per solver.

The mean memory usage is shown in Fig. A4. The best-fitting line to the memory scaling data has a gradient of approximately 1.1 indicating that memory scales nearly linearly with degrees of freedom.

C3 Number of eigensolutions

As a byproduct, the results of the scaling tests also give us a set of eigensolutions from which we can see the growth in the number of eigenvalues as a function of frequency. In Fig. A5 we show

the growth of the number of eigensolutions as a function of the imaginary component (or angular frequency) for the three largest models. In each of these plots, the growth is well described by a cubic polynomial which is consistent with results obtained by Al-Attar *et al.* (2012) for NMs of the Earth.

C4 Summary

In this scaling study, we have shown that we can reliably compute the eigensolutions for systems up to 1 million degrees of freedom.

The scaling of CPU time and Memory show that the limiting factor is likely to be CPU time as this scales with higher degree (approximately quadratically).

Calculating eigensolutions for much larger models is feasible, but the cost would likely be uncompetitive compared to more common time domain solutions using explicit numerical time integration. For this reason, and to permit coupling of different physical domains, our view is that the way forward is to use domain decomposition techniques and couple parts of larger models similar to strategies used before with SEM and NM approximations (Capdeville *et al.* 2003).