Published in final edited form as:

Comput Methods Appl Mech Eng. 2013 January 1; 253: 60–72. doi:10.1016/j.cma.2012.08.023.

Large Scale Parameter Estimation Problems in Frequency-Domain Elastodynamics Using an Error in Constitutive Equation Functional

Biswanath Banerjee^a, Timothy F. Walsh^b, Wilkins Aquino^c, and Marc Bonnet^d

^aSchool of Civil and Environmental Engineering, Cornell University, Ithaca, New York, 14853 USA

^bSandia National Laboratories, P.O. Box 5800, MS 0380, Albuquerque, NM 87185-9999, USA

^cDepartment of Civil and Environmental Engineering, Duke University, Durham, New York, 27708 USA

^dPOems (UMR 7231 CNRS-ENSTA-INRIA), Dept. of Appl. Math., ENSTA, Paris, France

Abstract

This paper presents the formulation and implementation of an Error in Constitutive Equations (ECE) method suitable for large-scale inverse identification of linear elastic material properties in the context of steady-state elastodynamics. In ECE-based methods, the inverse problem is postulated as an optimization problem in which the cost functional measures the discrepancy in the constitutive equations that connect kinematically admissible strains and dynamically admissible stresses. Furthermore, in a more recent modality of this methodology introduced by Feissel and Allix (2007), referred to as the Modified ECE (MECE), the measured data is incorporated into the formulation as a quadratic penalty term. We show that a simple and efficient continuation scheme for the penalty term, suggested by the theory of quadratic penalty methods, can significantly accelerate the convergence of the MECE algorithm. Furthermore, a (block) successive over-relaxation (SOR) technique is introduced, enabling the use of existing parallel finite element codes with minimal modification to solve the coupled system of equations that arises from the optimality conditions in MECE methods. Our numerical results demonstrate that the proposed methodology can successfully reconstruct the spatial distribution of elastic material parameters from partial and noisy measurements in as few as ten iterations in a 2D example and fifty in a 3D example. We show (through numerical experiments) that the proposed continuation scheme can improve the rate of convergence of MECE methods by at least an order of magnitude versus the alternative of using a fixed penalty parameter. Furthermore, the proposed block SOR strategy coupled with existing parallel solvers produces a computationally efficient MECE method that can be used for large scale materials identification problems, as demonstrated on a 3D example involving about 400,000 unknown moduli. Finally, our numerical results suggest that the proposed MECE approach can be significantly faster than the conventional approach of L^2 minimization using quasi-Newton methods.

© 2012 Elsevier B.V. All rights reserved.

*Corresponding author bb435@cornell.edu (Biswanath Banerjee), tfwalsh@sandia.gov (Timothy F. Walsh), wa20@duke (Wilkins Aquino), mbonnet@ensta.fr (Marc Bonnet).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Keywords

error in constitutive equation; inverse problems; parameter estimation; elastodynamics; successive over-relaxation

1. Introduction

The identification of material parameters (*e.g.* elastic moduli) are of paramount importance in science and engineering. For instance, this type of inverse problem arises in connection with applications such as damage detection and structural health monitoring, seismic exploration, and biomechanical imaging, among other applications. Models of complex engineering structures often feature local or heterogeneous parameters that are unknown and, therefore, need to be identified by exploiting experimental information on the mechanical response of the structure. Considerable research effort has been dedicated to develop algorithms for material identification. A review on these methods in the context of elasticity can be found in [1]. Despite current advances, current numerical algorithms are still often challenged by the inherent ill-posedness of inverse problems, especially when, as in this article, the identification of *heterogeneous* material properties is considered.

Parameter identification problems are often solved using nonlinear optimization schemes. The most common form of such approaches involves minimizing the L^2 -norm of the error between computed and measured responses (*e.g.* displacement, strains) [2, 3, 4, 5, 6, 7]. Quasi-Newton methods, which require only the computation of the gradient of the objective function using adjoint methods, are often preferred due to their ease of implementation [3, 8]. Full-Newton methods have also been successfully used for large-scale identification problems [9]. The latter methods converge significantly faster than quasi-Newton methods [8], but are more difficult to implement. In general, gradient-based nonlinear optimization methods have the advantage of allowing the handling of a large number of unknowns and noisy data (using regularization techniques). The main disadvantages of these approaches, when a L^2 -norm misfit functional is used, is their sensitivity to the initial guess due to the non-convexity of the functional and the large number of iterations required to get an acceptable solution when quasi-Newton methods are used.

In recent years, a new paradigm for inverse material identification has emerged which uses the concept of error in constitutive equations (ECE) for defining cost functionals whose physical meaning is stronger than that of usual L^2 functionals and directly related to the material identification problems at hand. The basic premise in the ECE approach is that, given an over-determined set of boundary or internal data (*e.g.* displacements and tractions), a cost functional is defined based on the least residual (measured in terms of an energy norm) in the constitutive equations that connect stresses and strains that are constrained to be dynamically and kinematically admissible, respectively, with respect to the available experimental data. This type of cost functional has the important property of being zero for the exact constitutive equations and strictly positive otherwise. ECE functionals have been initially introduced for error estimation in FEM computations [10] before being also applied to various material parameter identification problems under linear static [11], nonlinear quasistatic [12], time-harmonic [13, 14, 15] or, more recently, transient [16, 17, 18] conditions. ECE functionals were also found to be useful for solving data completion (Cauchy) problems [19]. Similar energy-error functionals have been introduced for the identification of scalar spatially-varying conductivity coefficients in *e.g.* [20, 21], with mathematical and numerical issues also discussed in [22, 23].

In particular, a new ECE-based method, named Modified Error in Constitutive Equations (MECE) approach, was recently proposed for time-domain dynamics problems by [17, 18]. In MECE, fields and equations are separated into reliable and unreliable sets. The reliable set typically includes the equilibrium equations, initial conditions and known boundary conditions, while the unreliable set includes measured data, constitutive properties, and (when applicable) imperfectly known boundary conditions. The identification problem is then posed as the minimization of a MECE functional, with reliable equations enforced strictly (e.g. as constraints, using Lagrange multipliers) and unreliable equations treated through an ECE-type functional and penalty terms. Allix and Feissel [17] showed that MECE is very robust and can produce reliable identification results in problems with high levels of data noise.

To our best knowledge, MECE-based functionals have not yet been applied to largescale dynamical identification problems involving unknown heterogeneous material parameters despite their clear potential advantages for this type of problem. With respect to such large-scale applications, the main drawbacks of existing MECE-based identification formulations are twofold. First, they involve a coupled system of partial differential equations (PDEs) that needs to be solved for each iteration. Second, a penalty parameter, which plays a very important role in accuracy of the inverse problem solution, must be provided *a priori*.

This paper is devoted to the formulation, implementation and validation of the MECE for large-scale heterogeneous material parameter identification problems in the context of frequency domain elastodynamics. Our main contributions in this work are threefold. First, we provide a simple continuation approach for evolving the penalty parameter that appears in MECE, eliminating the need to estimate this parameter *a priori*; second, we show a Successive OverRelaxation (SOR) scheme that can be suitable for solving large scale problems with MECE; and third, we demonstrate the feasibility of using the MECE method in two-dimensional and three-dimensional time-harmonic elasticity imaging problems with numerical experiments involving up to 400,000 unknown material parameters.

The article is organized as follows. The inverse problem and the MECE-based identification approach are introduced in Section 2. Proposed refinements of the MECE approach are presented in Sections 3 and 4, leading to the algorithm of Section 5. Section 6 is devoted to 2D numerical examples designed to assess the performance of the main ingredients of the proposed MECE algorithm (evolutionary penalty parameter, block-SOR technique for solving the MECE optimality equations), and to compare them against a L^2 -BFGS minimization approach. Finally, results on a large-scale 3D identification problem are shown in Section 7, before closing the paper with concluding remarks.

2. Background

2.1. Forward problem

In this section, we briefly describe the strong and weak forms of the forward steady-state elastodynamics problem. Let $\bar{\Omega} = \Omega \cup \partial\Omega \subset \mathbb{R}^d$ ($1 \leq d \leq 3$) denote a bounded and connected body. The governing equations for a linear elastic body undergoing time-harmonic motion can be written as

$$\nabla \cdot \sigma + \mathbf{b} = -\rho\omega^2 \mathbf{u} \quad \text{in } \Omega \quad (1a)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{on } \Gamma_u \quad (1b)$$

$$\sigma \cdot n_s = t \quad \text{on } \Gamma_t \quad (1c)$$

$$\sigma = \mathbb{C} : \epsilon \quad (1d)$$

$$\epsilon[u] = \frac{1}{2} (\nabla u + \nabla u^T) \quad (1e)$$

where u is the displacement field, ω is the angular frequency, ρ is the mass density, b is the body force density, ϵ denotes the linearized strain tensor, σ denotes the stress tensor, t is the applied traction, u_0 is the essential boundary condition, and n_s is the outward unit normal. Γ_t and Γ_u , respectively, denote the traction (natural) and displacement (essential) parts of the boundary. Also, $\Gamma_t \cap \Gamma_u = \emptyset$, $\Gamma_t \cup \Gamma_u = \Omega$, $\Gamma_t \cap \Gamma_u = \emptyset$ and $\mathbb{C}$ is the fourth-order linear constitutive tensor.

Let a and b be second-order tensor fields. In the sequel, we will use the following notation for their inner product.

$$(a, b) := \int_{\Omega} a : \bar{b} \, d\Omega = \int_{\Omega} a_{ij} \bar{b}_{ij} d\Omega \quad (2)$$

where indicial notation (with implicit summation over repeated indices) is used and the overbar denotes complex conjugation. The inner product of vector or scalar fields follow the same notation. When the inner product is defined over a boundary we will use the notation

$$(a, b)_{\Gamma} := \int_{\Gamma} a : \bar{b} \, d\Gamma \quad (3)$$

The weak formulation of (1a) consists in finding $u \in \mathcal{U}$ such that

$$\mathcal{A}(u, w) = \mathcal{F}(w) \quad \forall w \in \mathcal{W} \quad (4)$$

where

$$\mathcal{A}(u, w) := (\mathbb{C} : \epsilon[u], \epsilon[w]) - \rho \omega^2 (u, w) \quad (5)$$

$$\mathcal{F}(w) := (b, w) + (t, w)_{\Gamma_t} \quad (6)$$

and the function spaces $\mathcal{U}$ and $\mathcal{W}$ are defined as

$$\mathcal{U} = \{u | u \in H^1(\Omega), \quad u = u_0 \quad \text{on } \Gamma_u\} \quad (7a)$$

$$\mathcal{W} = \{w | w \in H^1(\Omega), \quad w = 0 \quad \text{on } \Gamma_u\} \quad (7b)$$

In addition, let the space $\mathcal{S}$ of dynamically admissible stresses be defined as

$$\mathcal{S} := \{\sigma | \sigma \in H_{\text{div}}(\Omega), \nabla \cdot \sigma + b = -\rho \omega^2 u \quad \text{in } \Omega, \sigma \cdot n_s = t \quad \Gamma_t\} \quad (8)$$

2.2. Inverse problem and L^2 minimization

The inverse problem associated with the steady-state elastodynamics shown above consists in estimating the spatial distribution of the constitutive tensor $\mathbb{C} \in \mathcal{C}$ for given measured displacements $u^m(\bar{x})$, $\bar{x} \in \Omega_m \subseteq \bar{\Omega}$ obtained at one or more frequencies, where $\mathcal{C}$ is the space of fourth-order tensor fields that are positive definite, symmetric and bounded.

A common approach for reconstructing material parameters is to solve either a constrained or an unconstrained optimization problem by minimizing the mean squared error between the measured and computed responses [3, 9]. We will refer to this method from here on as L^2 minimization. This standard procedure is now briefly described for later reference (Sec. 6.4). Without loss of generality, we consider the case where measurements are obtained for a single frequency and isotropic constitutive models are as given in Eq. (29).

The error functional in L^2 minimization is given as

$$J(u, \mathbb{C}) = \frac{1}{2} \|u - u^m\|_{L^2(\Omega_m)}^2 + \mathcal{R}(\mathbb{C}) \quad (9)$$

where

$$\|u - u^m\|_{L^2(\Omega_m)}^2 = \int_{\Omega_m} |u - u^m|^2 d\Omega_m \quad (10)$$

and $\mathcal{R}$ is a non-negative regularization functional. An approximate solution to the inverse problem is then obtained through the following PDE-constrained nonlinear optimization problem

$$\{u^*, \mathbb{C}^*\} = \underset{u, \mathbb{C}}{\operatorname{argmin}} J(u, \mathbb{C}) \quad \text{subject to Eqs. (1a-e)} \quad (11)$$

which is typically solved using Newton or quasi-Newton methods.

2.3. Error in constitutive equation approach

ECE approaches are based on cost functionals that measure, in terms of an energy norm, the constitutive equation residual between a given displacement field and a given stress field. For linearly elastic materials, the ECE cost functional is defined as

$$U(u, \sigma; \mathbb{C}) := \frac{1}{2} \int_{\Omega} (\sigma - \mathbb{C} : \epsilon[u]) : \mathbb{C}^{-1} : (\sigma - \mathbb{C} : \epsilon[u]) d\Omega \quad (12)$$

It is straightforward to see that $U(u, \sigma; \mathbb{C})$ has the important property of being zero for the exact constitutive equation and strictly positive otherwise (*i.e.* $U = 0 \forall \mathbb{C} \in \mathcal{C}$, $U = 0 \Leftrightarrow \sigma = \mathbb{C} : \epsilon$).

The error in constitutive equation $\varepsilon(\mathbb{C})$ for the given set of measured displacements is then defined through the partial minimization of $U(u, \sigma; \mathbb{C})$ with the material properties $\mathbb{C}$ kept fixed while u and σ fulfill all admissibility and experimental constraints, *i.e.*

$$\varepsilon(\mathbb{C}) := \min_{(u, \sigma) \in \mathcal{U} \times \mathcal{F}} U(u, \sigma; \mathbb{C}) \quad \text{subject to } u = u^m(\bar{x}), \bar{x} \in \Omega_m \subseteq \bar{\Omega} \quad (13)$$

In particular, $\varepsilon(\mathbb{C}) = 0$ in the absence of any measurements (the above minimization problem being then equivalent to the well-posed forward problem (1a-e)), and also for measurements that are consistent with the assumed material property $\mathbb{C}$. Conversely, if the assumed value of $\mathbb{C}$ is inconsistent with the measurements, $\varepsilon(\mathbb{C}) > 0$ and the correct material properties will

be found by minimizing $\varepsilon(\mathbb{C})$. This is the essence of the ECE approach to the inverse problem at hand.

2.4. Modified error in constitutive equation approach

In practice, however, measurement noise is to be expected, in which case the exact verification of experimental values as enforced in (13) is often not desirable. To address this concern, a modified error in constitutive equation (MECE) functional [17, 16] is defined by treating the discrepancy between measured and computed displacements as a penalty term, and is accordingly given by

$$\Lambda(u, \sigma; \mathbb{C}) = U(u, \sigma; \mathbb{C}) + \frac{k}{2} \|u - u^m\|_{L^2(\Omega_m)}^2 \quad (14)$$

where k is a penalization parameter. Then, the inverse problem in the scope of MECE is cast as the optimization problem

$$(u^*, \sigma^*, \mathbb{C}^*) = \arg \min_{(u, \sigma, \mathbb{C}) \in \mathcal{U} \times \mathcal{S} \times \mathcal{C}} \Lambda(u, \sigma; \mathbb{C}), \quad (15)$$

whose solution $(u^*, \sigma^*, \mathbb{C}^*)$ achieves a compromise between (a) satisfying the basic (*i.e.* equilibrium, compatibility and constitutive) mechanical equations and (b) matching the measured displacements. Indeed, the limiting values of the solution $(u^*, \sigma^*, \mathbb{C}^*)$ of (15) as $k \rightarrow 0$ and $k \rightarrow \infty$ are, respectively, the solutions to the L^2 minimization (11) and the “pure ECE” minimization (13). This balance was for example shown to be one of the strongest points of the (transient, spatially 1D) MECE method of Feissel and Allix [17], while a “dual” interpretation of MECE as a penalty approach for the L^2 minimization was proposed in [22]. Furthermore, while our approach is based on the minimization problem in (15), involving a single MECE functional Λ , the MECE method of [17] combines the minimization of (a time-domain version of) Λ with respect to (u, σ) and that of U with respect to $\mathbb{C}$.

2.4.1. Alternating directions—It is important to highlight that in the optimization problem (14) we are searching simultaneously for both the mechanical fields, σ and u , and the elastic parameters in $\mathbb{C}$. A natural approach for solving problem (14), followed in this work, consists in an alternating directions strategy whereby the transition from a current iterate $(u, \sigma, \mathbb{C})^n$ to the next iterate $(u, \sigma, \mathbb{C})^{n+1}$ is based on two successive and complementary partial minimizations of $\Lambda(u, \sigma, \mathbb{C})$. The first partial minimization consists in updating the mechanical fields given the material parameters:

$$(u^{n+1}, \sigma^{n+1}) := \arg \min_{(u, \sigma) \in \mathcal{U} \times \mathcal{S}} \Lambda(u, \sigma; \mathbb{C}^n) \quad (16)$$

Then, the second partial minimization consists in finding updated material parameters given the mechanical fields obtained from (16):

$$\mathbb{C}^{n+1} := \arg \min_{\mathbb{C} \in \mathcal{C}} \Lambda(u^{n+1}, \sigma^{n+1}; \mathbb{C}) \quad (17)$$

This alternating direction strategy serves two purposes: (i) it greatly simplifies the actual implementation of the method as it is straightforward to minimize the MECE functional with respect to the material parameters for fixed mechanical fields, and vice-versa, (ii) it provides greater insight into the MECE approach. It is worth mentioning that the alternating directions approach has been successfully used in other inverse problems to take advantage

of the structure of functionals similar to (14), for instance in the work of Wang *et al.* [24] on image reconstruction.

We now proceed to describe the solution strategy for subproblems (16) and (17), which exploits the optimality conditions for each subproblem using a Lagrange multiplier approach. Accordingly, define a Lagrangian functional $\mathcal{L}: \mathcal{U} \times \mathcal{W} \times \mathcal{S} \times \mathcal{C} \rightarrow \mathbb{R}$ by

$$\mathcal{L}(\mathbf{u}, \mathbf{w}, \sigma; \mathbb{C}) = \Lambda(\mathbf{u}, \sigma; \mathbb{C}) - \text{Re}(\mathcal{B}(\sigma, \mathbf{w}) - \mathcal{F}(\mathbf{w})) \quad (18)$$

where

$$\mathcal{B}(\sigma, \mathbf{w}) := (\sigma, \epsilon[\mathbf{w}]) - \rho \omega^2(\mathbf{u}, \mathbf{w}) \quad (19)$$

Notice that the test function $\mathbf{w} \in \mathcal{W}$ plays the role of a Lagrange multiplier in Eq. (18). The first-order optimality conditions for subproblem (16) are obtained by taking partial derivatives of (18) with respect to $\mathbf{u}$, σ and $\mathbf{w}$ (treating them as independent variables in the derivation), and setting these to zero. First, the partial derivative $\mathcal{L}'_{\sigma}$ of $\mathcal{L}$ with respect to the stress field can be defined in terms of a linear functional $\delta\sigma \mapsto \langle \mathcal{L}'_{\sigma}, \delta\sigma \rangle$ yielding the variation $\delta\mathcal{L}$ induced by any given stress variation $\delta\sigma \in H_{\text{div}}(\Omega)$ through

$$\begin{aligned} \langle \mathcal{L}'_{\sigma}, \delta\sigma \rangle &= \text{Re} \left((\delta\sigma, \mathbb{C}^{-1}:\sigma) - (\delta\sigma, \epsilon[\mathbf{u}]) - (\delta\sigma, \epsilon[\mathbf{w}]) \right) \\ &= \text{Re} \left((\delta\sigma, \mathbb{C}^{-1}:\sigma - \epsilon[\mathbf{u}] - \epsilon[\mathbf{w}]) \right) \quad \forall \delta\sigma \in H_{\text{div}}(\Omega) \end{aligned} \quad (20)$$

Setting $\langle \mathcal{L}'_{\sigma}, \delta\sigma \rangle$ to zero, we get

$$\sigma = \mathbb{C}:\epsilon[\mathbf{u} + \mathbf{w}] \quad (21)$$

Proceeding along similar lines, the partial derivative $\mathcal{L}'_{\mathbf{u}}$ of $\mathcal{L}$ with respect to the displacement field is defined through

$$\langle \mathcal{L}'_{\mathbf{u}}, \delta\mathbf{u} \rangle = \text{Re} \left((\mathbb{C}:\epsilon[\mathbf{u}] - \sigma, \epsilon[\delta\mathbf{u}]) + \kappa(\mathbf{u} - \mathbf{u}^m, \delta\mathbf{u})_{\Omega_m} + \omega^2(\rho\delta\mathbf{u}, \mathbf{w}) \right) \quad \forall \delta\mathbf{u} \in \mathcal{W} \quad (22)$$

Substituting Eq. (21) into Eq. (22), setting the result to zero, and simplifying, we get

$$(\mathbb{C}:\epsilon[\mathbf{w}], \epsilon[\delta\mathbf{u}]) - \omega^2(\rho\delta\mathbf{u}, \mathbf{w}) = \kappa(\mathbf{u} - \mathbf{u}^m, \delta\mathbf{u})_{\Omega_m} \quad \forall \delta\mathbf{u} \in \mathcal{W} \quad (23)$$

We next take the partial derivative $\mathcal{L}'_{\mathbf{w}}$ of $\mathcal{L}$ with respect to the Lagrange multiplier $\mathbf{w}$ to obtain

$$\langle \mathcal{L}'_{\mathbf{w}}, \delta\mathbf{w} \rangle = \text{Re} \left((\sigma, \epsilon[\delta\mathbf{w}]) - \omega^2(\rho\mathbf{u}, \delta\mathbf{w}) - (\mathbf{b}, \delta\mathbf{w}) - (\mathbf{t}, \delta\mathbf{w})_{\Gamma_t} \right) \quad (24)$$

Substituting Eq.(21) into Eq. (24), setting the result to zero, and simplifying, we get

$$(\mathbb{C}:\epsilon[\mathbf{u} + \mathbf{w}], \epsilon[\delta\mathbf{w}]) - \omega^2(\rho\mathbf{u}, \delta\mathbf{w}) = (\mathbf{b}, \delta\mathbf{w}) + (\mathbf{t}, \delta\mathbf{w})_{\Gamma_t} \quad \forall \delta\mathbf{w} \in \mathcal{W} \quad (25)$$

Equations (23) and (25) constitute a pair of coupled variational problems. Together with (21), they constitute the first-order optimality conditions for subproblem (16). In practice, the admissible mechanical fields that are consistent with the current estimate $\mathbb{C}$ of the

material parameters, *i.e.* that solve (16), are thus obtained by solving for $(u, w) \in \mathcal{U} \times \mathcal{W}$ in the coupled variational equations

$$\begin{aligned} \mathcal{A}(u, \delta w) + (\mathbb{C} : \epsilon[w], \epsilon[\delta w]) &= \mathcal{F}(\delta w) \quad \forall \delta w \in \mathcal{W} \\ -\kappa(u, \delta u)_{\Omega_m} + \mathcal{A}(w, \delta u) &= -\kappa(u^m, \delta u)_{\Omega_m} \quad \forall \delta u \in \mathcal{W}, \end{aligned} \quad (26)$$

and then obtain the stress σ from (21).

It is important to highlight that for given elastic properties $\mathbb{C}$, the coupled system (26) has the following properties: (i) the displacement field u approximates the measurement field u^m in a least squares sense; (ii) the Lagrange multiplier w and the displacement field u map to a stress field that satisfies (weakly) the conservation of linear momentum equation and natural boundary conditions; (iii) if $u \rightarrow u^m$ on Ω_m , then $w \rightarrow 0$ and (26) reduces to the forward problem (4); (iv) the formulation includes the case of overdetermined boundary conditions where $\Omega_m \cap \Gamma_t = \emptyset$. Of course, the latter would not be allowed in a classical variational formulation as it would lead to an ill-posed forward problem.

2.4.2. Material update—The still-unexploited first-order optimality equation $\mathcal{L}'_{\mathbb{C}} = 0$ yields the optimality condition for subproblem (17). Since (14) and (18) imply that the explicit dependence of $\mathcal{L}$ on $\mathbb{C}$ occurs only through $\Lambda(u, \sigma, \mathbb{C})$, and hence through $U(u, \sigma, \mathbb{C})$, one obtains

$$\langle U'_{\mathbb{C}}, \delta \mathbb{C} \rangle = 0 \quad \forall \delta \mathbb{C} \in \mathcal{C}, \quad (27)$$

i.e.

$$\left(\delta \mathbb{C}, \epsilon[u] \otimes \epsilon[\bar{u}] - (\mathbb{C}^{-1} : \sigma) \otimes (\mathbb{C}^{-1} : \bar{\sigma}) \right) = 0 \quad \forall \delta \mathbb{C} \in \mathcal{C} \quad (28)$$

(with $(\cdot, \cdot)$ denoting the inner product of type (2) for fourth-order tensor fields). On using the solution u, σ of subproblem (16), the above equation leads to a pointwise updating rule for $\mathbb{C}$. The latter is now shown to become completely explicit for the case of isotropic linear elastic materials, for which $\mathbb{C}$ can be represented as

$$\mathbb{C} = \left(B - \frac{2}{3}G \right) (\mathbf{I} \otimes \mathbf{I}) + 2G\mathbb{I} \quad (29)$$

In the above equation, B and G represent spatially dependent bulk and shear moduli, respectively. $\mathbf{I}$ and $\mathbb{I}$ represent the second and fourth order identity tensor, respectively. Considering perturbations of $\mathbb{C}$ of the form (29) with B, G replaced with $\delta B, \delta G$, the optimality condition (28) reduces to

$$(\epsilon_d, 2\delta G \epsilon_d) + (e_u, \delta B e_u) - \left(\sigma_d, \frac{\delta G}{2G^2} \sigma_d \right) - \left(p, \frac{\delta B}{B^2} p \right) = 0 \quad \forall \delta B, \delta G \in L^\infty(\Omega) \quad (30)$$

where ϵ_d and σ_d represent the deviatoric strain and stress, e_u is the volumetric strain and

$p = \frac{1}{3}\sigma$ is the pressure. Eq. (30) leads to the explicit pointwise updating formulae

$$B^* = \frac{(p, p)^{1/2}}{(e_u, e_u)^{1/2}} = B \frac{(e_u + e_w, e_u + e_w)^{1/2}}{(e_u, e_u)^{1/2}}, \quad (31a)$$

$$G^* = \frac{(\sigma_d, \sigma_d)^{1/2}}{2(\epsilon_d[u], \epsilon_d[u])^{1/2}} = G \frac{(\epsilon_d[u+w], \epsilon_d[u+w])^{1/2}}{(\epsilon_d[u], \epsilon_d[u])^{1/2}} \quad (31b)$$

where all occurrences of $(\cdot, \cdot)$ denote *pointwise* inner products (*i.e.* the implicit integration over a domain is dropped). It is important to bear in mind that σ depends on both the displacement field u and the Lagrange multiplier field w (see (21)), while the volumetric and deviatoric strains are function of u only.

Constitutive updating formulae (31) may alternatively be understood as yielding averaged updates over some region $D \subseteq \Omega$, e.g. over (patches of) elements, with inner products $(\cdot, \cdot)$ now carrying an implicit integration over D . The latter results simply from using piecewise-constant variations of the form $\delta B = \mathcal{X}(D)\delta b$ (where $\mathcal{X}(D)$ is the characteristic function of D and δb is a scalar), and similarly for δG , in (30).

2.5. Discretization of the coupled problems

The coupled problem (26) is discretized in this work using the finite element method. Using standard Voigt notation, the displacement fields and test functions are expressed as

$$\begin{aligned} u^h &= [N] \{u\}, & \delta u^h &= [N] \{\delta u\}, & \epsilon[u^h] &= [B] \{u\}, \\ w^h &= [N] \{w\}, & \delta w^h &= [N] \{\delta w\}, & \epsilon[w^h] &= [B] \{w\}, \end{aligned}$$

where $[N]$ and $[B]$ represents matrices of finite element shape function and their derivatives with respect to physical spatial coordinates, respectively. Substituting the above approximations into Eq. (26), the discrete coupled system of equations is obtained as

$$\begin{bmatrix} [K] - \omega^2 [M] & [K] \\ -\kappa [D] & [K] - \omega^2 [M] \end{bmatrix} \begin{Bmatrix} \{u\} \\ \{w\} \end{Bmatrix} = \begin{Bmatrix} \{P\} \\ -\kappa \{R\} \end{Bmatrix} \quad (32)$$

The matrices in the above block system of equations are defined as

$$[K] = \sum_{\text{elements}} \int_{\Omega^e} [B]^T [C] [B] d\Omega \quad (33a)$$

$$[M] = \sum_{\text{elements}} \int_{\Omega^e} \rho [N]^T [N] d\Omega \quad (33b)$$

$$[D] = \sum_{\text{elements}} \int_{\Omega_m^e} [N]^T [N] d\Omega_m \quad (33c)$$

$$\{P\} = \sum_{\text{elements}} \int_{\Gamma_t^e} [N]^T t d\Gamma + \sum_{\text{elements}} \int_{\Omega^e} [N]^T b d\Omega \quad (33d)$$

$$\{R\} = \sum_{\text{elements}} \int_{\Omega^e} [N]^T u^m d\Omega \quad (33e)$$

where $[C]$ denotes the constitutive matrix. In equations (33)a-e, the data is assumed to consist of a continuous field measured over Ω_m . For data collected sparsely in space,

measurement locations can be made to coincide with finite element nodes. In that case, introducing a diagonal Boolean matrix $[Q]$ with entries being non-zero only for the global degrees of freedoms (dofs) where measurements were taken, $[D]$ is then replaced by $[Q]$ in (32) while $\{R\}$ becomes a sparse vector containing the nodal measured displacements. Note that the matrix $[D]$ (or $[Q]$) is nonnegative as it is associated with the discretization of a least-squares misfit term. When Dirichlet boundary conditions are involved in problem (26) (i.e. if $\Gamma_u \neq \emptyset$), the lines and columns corresponding to prescribed DOFs are removed from system (32), and any nonzero prescribed displacement values appear through equivalent nodal forces in the right-hand side of (32). System (32) is uniquely solvable except in very specific situations, as stated next.

Proposition 1—System (32) admits a unique solution $(\{u\}, \{w\})$ whenever the intersection of the kernels of matrices $[K] - \omega^2 [M]$ and $[D]$ is reduced to the null vector.

If ω is not an eigenfrequency (i.e. a generalized eigenvalue of $[K] - \omega^2 [M]$), then the kernel of $[K] - \omega^2 [M]$ is trivial and Proposition 1 implies unique solvability of (32). If ω is an eigenfrequency, Proposition 1 still implies unique solvability of (32) provided any nonzero $\{v\}$ verifying $[K] - \omega^2 [M] \{v\} = \{0\}$ is such that $[D]\{v\} \neq \{0\}$ (i.e. the chosen measurement configuration is such that any eigenmode is observable). The proof of Proposition 1 is given in the Appendix. This proof in particular makes clear that unique solvability does not depend on whether the kinematic constraints permit rigid-body motions (i.e. whether $[K]$ is singular).

3. Normalization and a continuation scheme for the penalty term

In this section, we address a very important issue in the MECE approach: choosing or determining the penalty coefficient k that appears in the MECE functional (14). As shown in [17], the penalty coefficient plays a fundamental role in the quality of the reconstruction of material parameters, especially in cases where high levels of noise are present in the data. Our goal is to devise a simple strategy for estimating this coefficient adaptively. To this end, we first propose the following form for the penalty coefficient.

$$\kappa = \alpha \frac{U_s(\mathbb{C}_0)}{\|u^m\|_{L^2(\Omega_m)}^2} \quad (34)$$

where $U_s(\mathbb{C}_0)$ is the strain energy in the system for the initial guess $\mathbb{C}_0$ of elastic moduli and $\alpha \in \mathbb{R}^+$ is a non-dimensional penalty coefficient. The scaling (34) of k yields consistent units for the ECE and least-squares components of the MECE functional (14), with both terms having units of energy.

We now derive a selection rule for the penalty parameter α . To this end, since the measured displacements are enforced as a constraint using a quadratic penalty term in (15) (while the governing equations are enforced as constraints using Lagrange multipliers), we turn to the theory of quadratic penalty problems [8]. For any increasing sequence of penalty parameters $\{\alpha_i \in \mathbb{R}^+\}$ with $\alpha_i \rightarrow \infty$ as $i \rightarrow \infty$, it is shown (see e.g. Theorem 17.1 in [8]) that the corresponding sequence of minimizers $\{\sigma_i, u_i, \mathbb{C}_i\}$ for problem (15) converges to a minimizer of the original constrained problem (13). A natural approach to the inverse problem at hand is thus to solve problem (15) with an increasing sequence of penalty parameters $\{\alpha_i\}$, giving increasing weight to the measurement constraint (which is satisfied exactly only in the limit $\alpha_i \rightarrow \infty$). Wang *et al.* [24] showed excellent improvement in convergence speed by applying a monotonically increasing sequence of penalty parameters and an alternating directions minimization, as proposed herein, to image reconstruction problems.

Inspired by the quadratic penalty framework and the results of [24], we thus propose the following simple form for constructing a sequence of penalty parameters:

$$\alpha_{i+1} = \text{Min} \left(10^\beta \alpha_i, \alpha_{\max} \right) \quad (35)$$

where $\beta > 0$ and $\alpha_{\max}$ is a preselected positive number. We note that $\beta = 0$ reduces to the case of a constant α . For the numerical examples studied in this work, the proposed sequence has been found to significantly accelerate the convergence of the minimization (15).

There are several reasons that justify the choice of a progressively increasing penalty term in MECE. First, for small values of α , it is easier (*i.e.* faster convergence) to find a corresponding minimizer for the optimization problem in Eq. (15) than for large values of this coefficient. Second, each minimizer $\{\sigma_i^*, u_i^*, C_i^*\}$ can be used as an improved initial guess to the optimization problem corresponding to the next penalty parameter α_{i+1} in the sequence. In our proposed MECE approach, we execute just one iteration of the minimization scheme for each choice of α , which renders an algorithm that is not more costly than one that uses a fixed value of α . The construction (35) of the sequence α_i is to a large extent arbitrary. More efficient forms to define the sequence may well exist. For instance, by making $\alpha_{\max}$ dependent on the data noise level using optimal regularization arguments such as the Morozov principle or the L -curve. This important topic is left to future investigation.

As opposed to what the quadratic penalty theory suggests, the penalty parameter sequence (35) is prevented to reach arbitrarily large values. Indeed, since the data u^m is expected to be corrupted by measurement noise, exact satisfaction of the measurement constraint is not desirable because of the ill-posed nature of the inverse problem. Achieving a balance between minimizing the ECE and the measurement misfit, implicitly defined by the choice of $\alpha_{\max}$, is a more sensible goal. Like in regularization methods often used for inverse problems [25, 26], the value of $\alpha_{\max}$ should depend on the quality of the measurements (*i.e.* on the level of noise). Typically, $\alpha_{\max}$ has to be selected such that

$$\|u - u^m\|_{L^2(\Omega) \leq \delta} \quad (36)$$

where δ is a number that reflects the quality of the data. We found in practice that a wide range of values of $\alpha_{\max}$ for different levels of noise produce good results. Further work is necessary to identify the exact relationship between $\alpha_{\max}$ and the level of noise in the data.

4. Successive over-relaxation (SOR) technique for the solution of the coupled problem

One of the main computational challenges that arises in the MECE method is obtaining a solution for the coupled system (32). This can be achieved by either employing direct or iterative solvers on the whole system or using a block strategy. Applying direct linear solvers to (32) for large scale problems may be prohibitive. On the other hand, designing efficient and effective pre-conditioners for the entire coupled system is not a trivial task if iterative solvers were to be used. Hence, our goal is to devise an iterative solution strategy that exploits the block structure of the system efficiently and takes advantage of existing massively parallel linear solvers. To this end, we employ a simple successive over-relaxation (SOR) strategy at the block level.

Consider (32) as a block system. The $i + 1$ iterate in the block SOR algorithm [27] is obtained by solving

$$\begin{bmatrix} [K] - \omega^2 [M] & 0 \\ -\eta\kappa [D] & [K] - \omega^2 [M] \end{bmatrix} \begin{Bmatrix} \{u\}^{i+1} \\ \{w\}^{i+1} \end{Bmatrix} = \begin{bmatrix} \widehat{\eta}([K] - \omega^2 [M]) & \eta [K] \\ 0 & \widehat{\eta}([K] - \omega^2 [M]) \end{bmatrix} \begin{Bmatrix} \{u\}^i \\ \{w\}^i \end{Bmatrix} + \begin{Bmatrix} \eta \{P\} \\ -\eta\kappa \{R\} \end{Bmatrix} \quad (37)$$

where $0 < \eta < 2$ and $\widehat{\eta} = 1 - \eta$. Defining

$$\{\widehat{u}\} = \{u\}^{i+1} - \widehat{\eta}\{u\}^i, \quad \{\widehat{w}\} = \{w\}^{i+1} - \widehat{\eta}\{w\}^i, \quad (38)$$

substituting into Eq. (37), and rearranging, we obtain the one-way coupled systems of equations

$$([K] - \omega^2 [M])\{\widehat{u}\} = \eta(\{P\} + [K]\{w\}^i) \quad (39a)$$

$$([K] - \omega^2 [M])\{\widehat{w}\} = \eta\kappa([D]\{u\}^{i+1} - \{R\}) \quad (39b)$$

At each SOR iteration, $\{\widehat{\eta}\}$ is computed first from (39a), and then $\{u\}^{i+1}$ is obtained using the left equation in (38). Subsequently, $\{\widehat{w}\}$ is computed from (39b), and $\{w\}^{i+1}$ is obtained using the right equation in (38).

The system shown in (39) has significant computational advantages. First, equations (39) are standard forced-vibration problems to which linear solvers available in FE codes can be applied directly. Second, both equations in (39) have the same governing matrix, but different right hand sides. This fact allows for very efficient implementation of the SOR iterations. We used the FETI-H algorithm [28, 29] for solving the system (39). This method is based on a non-overlapping domain decomposition, and uses special preconditioners based on plane waves to improve convergence. FETI-H can be used very efficiently in problems with multiple right-hand sides and with a constant system matrix (see for instance [30]) such as the ones studied herein. Note, however, that equations (39) become singular at resonant frequencies even though the original system (32) is in general still uniquely solvable in that case.

5. MECE algorithm

In summary, each iteration in the MECE approach consists in solving the coupled system given in Eq. (32) followed by updating the material parameters using the formulae given in Eq. (31). These steps are executed until a maximum number of iterations is reached or a predefined error tolerance is satisfied. The flow of an implementation of the MECE approach is as follows.

- Set initial values for shear and bulk moduli, fix lower and/or upper bounds and normalization factors.
- Loop until a convergence criterion is met
 1. Obtain a solution of the coupled problem using Eq. (32)
 2. Compute stresses and strains
 3. Update moduli using Eq. (31)

4. Update penalization parameter using Eq. (35)

6. 2D Numerical Examples and Results

This section presents a set of 2D numerical examples that were designed to study the performance of the MECE approach described herein for reconstructing spatially varying bulk and shear moduli in the presence of incomplete and noisy data. A 3D example using a massively parallel iterative solver for (32) will then be presented in Sec. 7. The examples presented herein are motivated from inverse elasticity problems related to biomedical applications [4, 31, 32]. In these problems, displacement (or velocity) data is usually available on the boundary and interior of the body. However, the data may be missing in some directions, as in the case of ultrasound imaging [4], or be fully multidirectional, as in the case of magnetic resonance imaging (MRI) [32].

6.1. Problem description

For the 2D numerical experiments, we considered a $10 \text{ cm} \times 10 \text{ cm}$ square under plane strain conditions and loaded with a uniform pressure at the top. The bottom of the square was free in the X-direction as shown in Figure 1. The domain contained two circular inclusions in a soft matrix (Figures 2a and 2b), whose elastic properties are given in Table 1. Given that the coordinate origin was located at the bottom left corner of the domain, the center of the bottom inclusion had coordinates (2.5 cm, 2.5 cm), while the center of the top inclusion had coordinates (7.5 cm, 7.5 cm). The diameter of the inclusions was 2.5 cm. The mass density of the inclusions and matrix was taken as 1000 kg/m^3 and a single frequency of 50 Hz was used in all cases.

The known “experimental” data used in the numerical examples was generated artificially using a structured finite element mesh consisting of 100×100 four-node element. The y -component of the displacements at all nodes in the mesh was used as known information. The data was then polluted by adding artificial Gaussian white noise. That is, for a degree of freedom i , the synthetic measurement was obtained as

$$u_i^m = u_i^{\text{ref}} (1 + \delta r_i) \quad (40)$$

where u_i^{ref} is a noise-free displacement obtained using the true material properties, δ represents the noise level (with $\delta = 0.01, 0.05$ or 0.1 used in the examples to follow), and r_i denotes a zero-mean and unit-variance Gaussian random variable. The abovedescribed synthetic data was finally interpolated on a coarser, and uniform, 61×61 mesh of bilinear elements which was used for all 2D reconstructions.

In the examples shown in Secs. 6.2 and 6.3, the MECE algorithm was stopped when a maximum number of iterations was reached or when the condition

$$\frac{\|u - u^m\|_{L^2(\Omega_m)}}{\|u^m\|_{L^2(\Omega_m)}} \leq c\delta \quad (41)$$

was met. Without loss of generality, we used $c = 1$ in this work and the maximum number of iterations allowed for the MECE algorithm was set to 1000.

The results of Secs. 6.2 and 6.3 were obtained using a direct solver for the coupled system (32), while the SOR technique was used in Sec. 6.4. Furthermore, in Secs. 6.2 and 6.3, the parameters for the penalty sequence (35) were chosen as $\alpha_0 = 1$, $\alpha_{\max} = 10^9$, and $\beta = 0.25$. From the numerical examples studied in this work, we found that the algorithm was not

sensitive to different choices of these parameters. The 2D numerical examples used an initial guess of constant shear and bulk moduli both with values of 1 MPa.

6.2. Reconstruction of shear and bulk moduli

We consider the spatial reconstruction of the true distributions of G and B defined by Figures 2a, 2b and Table 1. The reconstructed profiles obtained using MECE with the proposed iterative penalization scheme are shown in Figures 2c and 2d for a noise level $\delta = 0.05$ (for visualization purposes, the nodal averages of G and B are plotted). The MECE algorithm was clearly able to reconstruct the location of the inclusions as well as their sharp boundaries.

Figure 3 shows the spatial variation of the reconstructed shear and bulk moduli along the section $A - A'$ indicated in Fig. 1, for different noise levels. These plots demonstrate that the MECE-based algorithm was able to identify accurately the boundaries of the inclusions and correct magnitudes of the moduli with, as expected, a degradation of accuracy as the level of noise increased. While the sharp discontinuities are well captured in the reconstruction, it is important to notice that the error in the reconstructed bulk modulus was consistently higher than that of the shear modulus. The higher error in the bulk modulus may be due to the loading condition and the unidirectionality of the measurements used in the simulated experiments. For instance, the fact that the synthetic data is limited to the vertical component of the displacement field may reduce its sensitivity to volume changes.

There are many applications (*e.g.* those involving near-incompressible materials) in which G is the only material parameter to be identified. As the number of unknowns is reduced for the same level of information, it is expected that the accuracy of the results would increase as compared to the case where both B and G are treated as unknowns. The cross-sectional plot of Figure 4 corroborates this statement for the $\delta = 0$ and $\delta = 0.01$ noise levels. However, it can also be observed that the increased accuracy from having fewer design variables does not hold for higher levels of noise.

6.3. Convergence behavior with sequence of penalty coefficients

Now we compare the convergence behavior of the MECE algorithm when the proposed continuation scheme (35) was used versus using a constant penalty parameter. Figure 5 shows a plot of the value of the cost function $\Lambda(u, \sigma, \mathbb{C})$ defined by (14) and using $k = 1$ versus number of iterations. Notice that this choice of k yields a cost function that represents an equally weighted sum of the ECE and L_2 errors. The plots shown in this figure correspond to different values of α (held constant during the iterations) and α selected using the continuation scheme (35) for a noise level $\delta = 0.05$. For a constant α , the number of iterations required for reaching a certain level of $\Lambda(u, \sigma, \mathbb{C})$ is seen to decrease as the value of α increases. However, the quality of the reconstruction is poorer with higher values of α as can be observed in the cross-sectional plot of Figure 6, showing that with higher values of constant α the reconstruction becomes highly oscillatory. This behavior is due to the fact that when larger, constant penalty parameters are chosen, the noise in the measurements is enforced with increased weight early in the minimization process.

It is interesting to observe that the reconstructed solutions for both the continuation scheme and $\alpha = 1.0$ are similar in accuracy and smoothness. However, the MECE algorithm displays a significantly faster convergence when the continuation scheme is used. Notice that the number of iterations required to reach the lowest value of the cost function is more than ten times higher for $\alpha = 1.0$ than that required by the continuation scheme. These results demonstrate the significant computational advantage of using a continuation scheme with MECE. However, it is also important to recognize that there may be constant values of α

that produce rates of convergence comparable to those displayed by the continuation scheme. However, finding such a constant α can be very difficult in practice.

The sensitivity of the convergence rate of MECE to the choice of the exponent β in the penalty sequence (35) has also been studied. Figure 7 shows plots of the value of the cost function $\Lambda(u, \sigma, \mathbb{C})$ versus number of iterations for five different values of β . For all nonzero values of β , the stopping criterion was reached within 5-10 iterations. In contrast, when $\beta = 0$ (i.e. for keeping the constant value $\alpha = 1$), the algorithm was unable to meet this criterion after 1000 iterations. Cross-sectional plots of the corresponding reconstructed moduli (Figure 8) show that adequate reconstructions were obtained for all choices of β . This numerical experiment highlights the robustness of the proposed continuation scheme over a wide range of values β .

6.4. A comparison between the MECE method and an L^2 /L-BFGS minimization approach

In this section, we compare the performance of the proposed MECE scheme with that of a standard L^2 minimization approach that employs a quasi-Newton method (see Section 2.2). To this end, we used a limited-memory (L-BFGS) algorithm for bound-constrained minimization [33]. The domain used for this comparison had dimensions 1 m $\times$ 1 m and was loaded with a uniform pressure on the top surface and restrained in the vertical direction at the bottom surface. The top pressure was taken as unity and the excitation frequency was 0.025 Hz. The domain contained a concentric circular inclusion with a radius of 0.15 m. The target material properties were $E_{\text{matrix}} = 1 \text{ Pa}$, $\nu_{\text{matrix}} = 0.25$ for the matrix and $E_{\text{inc}} = 10 \text{ Pa}$, $\nu_{\text{inc}} = 0.25$ for the inclusion, the corresponding shear and bulk moduli thus being $G_{\text{matrix}} = 0.4 \text{ Pa}$, $G_{\text{inc}} = 4 \text{ Pa}$ and $B_{\text{matrix}} = 0.67 \text{ Pa}$, $B_{\text{inc}} = 6.7 \text{ Pa}$.

The given target moduli were used to generate the data used for the inversion. Displacements at all nodes in the vertical and horizontal directions were taken as measured data. For this example, the same mesh was used for generating the data and obtaining an inverse problem solution. This is acceptable in this case as we are comparing the performance of the two algorithms under the same conditions. Furthermore, using the same mesh for the inversion and the data generation allows for a straightforward computation of the reconstruction errors. The latter were computed for the shear and bulk moduli, according to

$$e_G := \left(\frac{\sum_e (G_e^h - G_e)^2}{\sum_e G_e^2} \right)^{1/2}, \quad e_B := \left(\frac{\sum_e (B_e^h - B_e)^2}{\sum_e B_e^2} \right)^{1/2} \quad (42)$$

where G_e and B_e represent the values of the true shear and bulk moduli, respectively, at the centroid of element e , while G_e^h and B_e^h represent the corresponding values obtained from the inversion.

The parameters used for the MECE algorithm in this example were $\alpha_0 = 0.1$, $\beta = 0.25$, and $\alpha_{\text{max}} = 5$. The SOR algorithm was used in this example with $\nu = 0.2$, a maximum number of 5 SOR iterations per MECE iteration, and a SOR tolerance of 10^{-8} . The initial guess was taken as a uniform material distribution with $E = 4 \text{ Pa}$ and $\nu = 0.25$ for both the MECE and L-BFGS algorithms.

The relative errors e_G and e_B were computed at each iteration of the MECE and L-BFGS algorithms and are shown in Figure 9 in logarithmic scale plots. The plots show that for a given level of relative error in bulk or shear modulus, the number of iterations for MECE is at least one order of magnitude smaller than the corresponding number for L-BFGS. Similar results that reinforce the improved convergence rates of error in constitutive equation

methods over quasi-Newton methods with L^2 minimization have also been reported for the reconstruction of viscoelastic moduli in steady-state dynamics [34]. Figure 10 shows the distribution in shear modulus after 500 iterations of the L-BFGS algorithm and just 100 iterations of the MECE algorithm. This figure shows significantly better accuracy of the reconstruction achieved with MECE than that obtained with L^2 minimization.

Some remarks are in order about the relative computational cost of an iteration of each algorithm. Each MECE iteration requires solving a coupled block-system of equations (32) once, while for L^2 minimization each cost functional and gradient evaluation entails solving one forward problem and one adjoint problem. We limit our discussion herein to the case where the block-SOR algorithm is used for MECE. First, it is important to notice that in one iteration of MECE, the left-hand side (*i.e.* system matrix) of (39) does not change during the inner SOR iterations. In the case of L^2 -BFGS, the adjoint problem has the same system matrix as the last forward problem solved. However, the line search component of the algorithm may require additional evaluations of the forward problem. For the discussion presented herein, we will conservatively assume that an iteration of the L^2 -BFGS algorithm entails solving just one forward and one adjoint problem with the same system matrix.

The fact that the system matrix does not change, if conveniently exploited, makes the computational cost of one MECE iteration similar to that of one L^2 evaluation. For instance, assuming that a direct approach (*e.g.* Gauss elimination) is used to solve (39a,b), only one factorization at the beginning of the SOR iterations is required, the rest of the computational cost being spent in operations of lower computational complexity such as matrix-vector multiplications, or back and forward substitutions. L^2 -BFGS also requires only one factorization per evaluation. For large FE models, one MECE iteration and one L^2 -BFGS evaluation are thus expected to have similar computational costs (with L^2 -BFGS evaluations needing less lower-complexity operations and thus having a slight edge over MECE iterations) provided that only one factorization is done. Given the vast difference in convergence rate between MECE and L-BFGS displayed in the previous plots, the MECE algorithm is expected to be more computationally efficient than L^2 minimization with L-BFGS even if one evaluation of the latter algorithm is somewhat faster than one MECE iteration.

To our best knowledge, there has not been any rigorous mathematical analysis of the MECE approach presented herein that elucidates the properties of the method and explains the numerical observations presented in this work. However, there has been work on similar methods that indicate that ECE functionals are convex or have “improved” convexity over their L^2 counterpart (see *e.g.* Gockenbach and Khan [23], and Bonnet and Constantinescu [1]). During the course of this work, the MECE method was found to be much less sensitive to initial guesses than the L^2 minimization approach; an observation in agreement with results reported by Gockenbach and Khan [23], Bonnet and Constantinescu [1], and Aguilo [34].

7. 3D Example Using Parallel Algorithm

In this section, we consider a three-dimensional example in which the proposed SOR algorithm and the FETI-H parallel linear solver are used to obtain solutions of the coupled block system (32) at each MECE iteration. The domain of the problem consists of a sti cylinder embedded in a soft cubic matrix as shown in Figure 11. The “true” Young’s modulus and Poisson’s ratio of the matrix were taken as $E_{\text{matrix}} = 1 \text{ Pa}$ and $\nu_{\text{matrix}} = 0.3$, respectively. The corresponding values for the inclusion were $E_{\text{inc}} = 10 \text{ Pa}$ and $\nu_{\text{inc}} = 0.3$. The dimensions of the cube were $(1 \text{ m} \times 1 \text{ m} \times 1 \text{ m})$, while the cylinder, located at the center of the cube, had a diameter of 0.5 m and a length of 0.5 m . The lower face of the block (at y

$= -0.5 \text{ m}$) was restrained in the Y -direction and free in the other two directions. The top face ($y = 0.5$) was given a pressure load of magnitude 1.0 Pa , and the remaining four side faces were given pressure loads of 0.5 Pa . The pressure loads had no difference in phase. This type of loading was used to excite both shear and bulk response. Two different frequencies, 0.025 Hz and 0.050 Hz , were used to generate the data for the inversion. These frequencies were selected to be below the first resonant frequency (0.11 Hz) of the excited structure in order to ensure the well-conditioning of the block system.

The measured data was interpolated onto a coarser mesh that had no representation of the cylinder and contained about 200,000 linear tetrahedral elements. Furthermore, noise was added to the interpolated data with $\delta = 0.01$ (see Eq. (40)). For this example, it was assumed that the three components of displacements were measured (available) at each node location of the coarse mesh. Finally, the inverse problem was solved using this interpolated data and the coarse mesh. The total number of design variables (*i.e.* unknown moduli) for this problem was around 400,000.

At each MECE iteration, the proposed block-SOR algorithm was used to approximate the solution of system (32) with FETI-H used as the linear solver for the systems (39a,b) at each SOR iteration. The SOR coefficient η was set at 0.2 and the maximum number of SOR iterations allowed in one MECE iteration was set at 7. In this example, the SOR tolerance was set to 10^{-8} . The continuation scheme parameters were taken as $\alpha_{\max} = 1.0$, $\alpha_0 = 0.05$, and $\alpha = 0.25$. The value of $\alpha_{\max}$ for this example was selected from numerical experiments and was decided based on the stability of the SOR algorithm.

We determined from numerical experiments that $\eta = 0.2$ produced adequate rates of convergence for the SOR algorithm and selected $\alpha_{\max}$ according to this value of η . We performed 50 MECE iterations and reached a value of $\Lambda(u, \sigma, \mathbb{C})$ (with $k = 2$) of the order of 10^{-4} . This example was carried out on a Linux workstation with 8 processors and 24 GB of RAM.

Figure 12 shows a clip plane view of the bulk modulus through the center of the domain. This plot shows that the matrix and inclusion are clearly distinguished. Furthermore, it can be observed that the reconstructed solution is close to the target distribution of the bulk modulus. The target values for the matrix and the inclusion were $B_{\text{matrix}} = 0.83 \text{ Pa}$ and $B_{\text{inc}} = 8.3 \text{ Pa}$, respectively. Figure 13 shows a threshold plot of the bulk modulus throughout the entire volume, overlaid on top of a perspective view of the original position of the cylinder as shown in Figure 11 to highlight the accuracy of the reconstruction. Again, the stiff region is clearly illuminated and the values are in the expected range. The spurious higher moduli seen at the top and sides of the cube (where the loading is applied) are likely due to the insensitivity of the measured displacements to the bulk modulus in these parts of the domain.

Figure 14 shows a clip plane view of the shear modulus through the center of the domain. As in the case of the bulk modulus, the matrix and inclusion are clearly distinguished in the plot and the reconstructed solution is in an expected range. Recall that the target shear moduli for the matrix and inclusion were given as $G_{\text{matrix}} = 0.38 \text{ Pa}$ and $G_{\text{inc}} = 3.8 \text{ Pa}$, respectively. Figure 15 shows a threshold plot of the shear modulus throughout the whole volume. In this case, there is little aberration in the shear moduli near the surface where the loading is applied. The bounds in the threshold plot in Figure 15 show that the shear modulus values obtained from the inverse problem are indeed close to the expected values, both in the matrix and in the inclusion.

A few remarks are in order. First, the reconstructed moduli are overestimated in some areas of the cylinder and underestimated in others. This trend in the results is expected and can be

attributed to the different sources of error in the simulated example (*i.e.* data, spatial approximation, optimization, etc). Second, the run time for this 3D example was just under six hours, which is indeed a successful outcome given the size of the inverse problem. The performance of our implementation of the parallel MECE algorithm can in fact be further improved as the blocks in the coupled system were factorized in the FETI-H solver during each SOR iteration, rather than just once at the beginning of the SOR iterations, during which the left hand side remains unchanged. We are in the process of improving our implementation to this end. Finally, the algorithm was able to identify a solution with only 50 MECE iterations and 7 SOR iterations per MECE iteration, which are indeed low iteration counts given the large dimension of the parameter search space.

Finally, we would like to comment on the effect of frequency on the performance of the proposed MECE algorithm. The coupled system of equations (32) is still in general solvable at resonant frequencies (see Proposition 1). We investigated the performance of the algorithm when the excitation frequency is close to a resonant frequency of the system. Adequate reconstructions were obtained in 2D using a direct solver (consistently with the established solvability of (32)), with results similar to those shown for the example in Section 6.2. However, we did not perform such exploration with the 3D version of the algorithm based on the iterative block solver, as the latter fails at resonant frequencies (the diagonal blocks then being singular) and presumably requires preconditioning at near-resonant frequencies. A more detailed numerical and analytical study of this issue will be pursued in a future investigation.

8. Concluding Remarks

We proposed a methodology based on the modified error in constitutive equations (MECE) approach that is suitable for large scale inverse identification of elastic parameters in frequency-domain dynamics. To this end, we proposed a continuation scheme for the penalty term that appears in MECE. This continuation scheme leads to faster convergence and improved accuracy in the reconstructed solution as compared to the case when a constant penalty coefficient is used. Furthermore, we put forward a simple block-SOR algorithm that allows for a straight forward implementation of MECE using existing massively parallel solvers. Our results demonstrate that the MECE algorithm can produce moduli reconstructions in as few as 10 iterations in 2D problems and as few as 50 iterations for the 3D examples explored in this paper. This is a significant improvement over the convergence speed usually observed in conventional least squares methods implemented using quasi-Newton schemes where hundreds or even thousands of iterations are commonly needed to obtain adequate results. The latter was demonstrated through a numerical example in the paper.

Acknowledgments

The work of Biswanath Banerjee and Wilkins Aquino was partially supported by the United States National Institute for Biomedical Imaging and Bioengineering, Grant# 5R01EB002640-13. The work of Wilkins Aquino was also supported by the CSAR Program at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy (DE-AC04-94AL85000).

Appendix: A. Proof of Proposition 1

Let $0 = \omega_1 \dots \omega_N$ denote the N eigenfrequencies (counting possible multiplicities), and set $[X] := [X_1, \dots, X_N]$, where X_j is a $[M]$ -orthonormal eigenmode for $\omega = \omega_j$. Set $\{u\} = [X]\{U\}$ (*i.e.* expanding $\{u\}$ on the basis of eigenmodes) and likewise $\{w\} = [X]\{U\}$, in (32),

and left-multiply the resulting equations by $\{X_j^T, X_j^T\}$ for all j . The corresponding projected structural matrices are such that

$$\tilde{K} := [X]^T [K] [X] = \text{Diag}(\omega_j^2), \quad \tilde{Z} := [X]^T ([K] - \omega^2 [M]) [X] = \text{Diag}(\omega_j^2 - \omega^2)$$

Moreover, letting r denote the multiplicity of the eigenvalue $\omega = 0$ (i.e. the number of linearly independent rigid-body motions permitted by $[K]$), the generalized DOFs are partitioned according to $\{U\} = \{U_0, U_1\}$, $\{W\} = \{W_0, W_1\}$ (with $U_0, W_0 \in \mathbb{R}^r$ gathering the generalized DOFs associated to rigid-body modes). By virtue of the $[M]$ -orthonormality of the X_j , the induced partitioning on $[K]$ and $[M]$ is such that $\tilde{K}_{00} = [0]$ and $\tilde{K}_{01} = \tilde{M}_{01} = [0]$; moreover, both diagonal matrices $\tilde{Z}_{00}$ and $\tilde{K}_{11}$ are invertible.

Upon performing all the above-described operations, the system (32) yields the following block-partitioned governing matrix equation for the generalized DOFs:

$$\begin{bmatrix} \tilde{Z}_{00} & 0 & 0 & 0 \\ 0 & \tilde{Z}_{11} & 0 & \tilde{K}_{11} \\ -\kappa \tilde{D}_{00} & -\kappa \tilde{D}_{01} & \tilde{Z}_{00} & 0 \\ -\kappa \tilde{D}_{10} & -\kappa \tilde{D}_{11} & 0 & \tilde{Z}_{11} \end{bmatrix} \begin{bmatrix} U_0 \\ U_1 \\ W_0 \\ W_1 \end{bmatrix} = \begin{bmatrix} \tilde{p}_0 \\ \tilde{p}_1 \\ -\kappa \tilde{r}_0 \\ -\kappa \tilde{r}_1 \end{bmatrix} \quad (\text{A.1})$$

where $\tilde{p} := [X]^T \{p\}$, $\tilde{r} := [X]^T \{r\}$, $\tilde{D} := [X]^T [D] [X]$, and $\tilde{D}_{00}$ etc. result from DOF partitioning of $\tilde{D}$. Using the known invertibility of $\tilde{Z}_{00}$ and $\tilde{K}_{11}$, the first three lines of (A.1) can be solved for U_0, W_0, W_1 , yielding

$$U_0 = \tilde{Z}_{00}^{-1} \tilde{p}_0, \quad W_0 = \kappa \tilde{Z}_{00}^{-1} (\tilde{D}_{01} U_1 + \tilde{D}_{00} \tilde{Z}_{00}^{-1} \tilde{p}_0 - \tilde{r}_0), \quad W_1 = \tilde{K}_{11}^{-1} (\tilde{p}_1 - \tilde{Z}_{11} U_1)$$

Substituting the above results into the last equation of (A.1) thus yields the following matrix equation for the remaining unknown U_1 :

$$R U_1 = \tilde{Z}_{11} \tilde{K}_{11}^{-1} \tilde{p}_1 - \kappa (\tilde{D}_{10} \tilde{Z}_{00}^{-1} \tilde{p}_0 - \tilde{r}_1) \quad \text{with} \quad R := \kappa \tilde{D}_{11} + \tilde{Z}_{11} \tilde{K}_{11}^{-1} \tilde{Z}_{11} \quad (\text{A.2})$$

Since both $k[D]$ and $[K]$ are real symmetric non-negative matrices, so is the resolvent matrix R . Therefore, R is invertible if and only if it is positive definite, i.e. noninvertibility of R occurs only if there exists a nonzero vector V such that $V^T R V = 0$. Since $k V^T [D] V = 0$ and $V^T \tilde{Z}_{11} \tilde{K}_{11}^{-1} \tilde{Z}_{11} V \geq 0$ for any V , R is singular only if $k V^T [D] V = 0$ and $V^T \tilde{Z}_{11} \tilde{K}_{11}^{-1} \tilde{Z}_{11} V = 0$ simultaneously hold for some nonzero V . That is, R is singular only if V belongs

simultaneously to the kernels of $[D]$ and $\tilde{Z}_{11}$ since $\tilde{K}_{11}^{-1}$ is invertible. The kernel of $\tilde{Z}_{11}$ is nontrivial if ω coincides with one of the nonzero eigenvalues ω_j . This kernel is then the corresponding modal subspace. Proposition 1 follows and the proof is complete. The case where $r = 0$, i.e. $[K]$ is invertible, is included as a special, simpler, case where DOF partitioning is not necessary and the proof proceeds otherwise similarly.

References

- [1]. Bonnet M, Constantinescu A. Inverse problems in elasticity. *Inverse Problems*. 2005; 21:R1–R50.
- [2]. Doyley MM, Meaney PM, Bamber JC. Evaluation of iterative reconstruction method for quantitative elastography. *Physics in Medicine and Biology*. 2000; 45:1521–1540. [PubMed: 10870708]

- [3]. Oberai AA, Gokhale NH, Feijoo GR. Solution of inverse problems in elasticity imaging using the adjoint method. *Inverse Problems*. 2003; 19:297–313.
- [4]. Oberai AA, Gokhale NH, Dooley MM, Bamber JC. Evaluation of the adjoint equation based algorithm for elasticity imaging. *Physics in Medicine and Biology*. 2004; 49:2955–2974. [PubMed: 15285258]
- [5]. Dooley MM, Srinivasan S, Dimidenko NSE, Ophir J. Enhancing the performance of model-based elastography by incorporating additional a priori information in the modulus image reconstruction process. *Physics in Medicine and Biology*. 2006; 51:95–112. [PubMed: 16357433]
- [6]. Brigham JC, Aquino W. Inverse viscoelastic material characterization using POD reduced-order modeling in acoustic-structure interaction. *Comp. Meth. Appl. Mech. Eng.* 2009; 198:893–903.
- [7]. Aguilo MA, Brigham JC, Aquino W, Fatemi M. Inverse Reconstruction of Spatially-Varying Elastic Properties in Solids using Vibroacoustics. *IEEE Transactions on Medical Imaging*. 2010; 4:1012–1021. [PubMed: 20335092]
- [8]. Nocedal, J.; Wright, SJ. *Numerical Optimization*. SIAM, Springer; 2000.
- [9]. Epanomeritakis I, Akcelik V, Ghattas O, Bielak J. A Newton-CG method for large-scale three-dimensional elastic full-waveform seismic inversion. *Inverse Problems*. 2008; 24:034015, 26.
- [10]. Ladevèze P, Leguillon D. Error estimate procedure in the finite element method and applications. *SIAM J. Numer. Anal.* 1983; 20:485–509.
- [11]. Geymonat G, Pagano S. Identification of mechanical properties by displacement field measurement: a variational approach. *Meccanica*. 2003; 38:535–545.
- [12]. Latourte F, Chrysochoos A, Pagano S, Wattrisse B. Elastoplastic behavior identification for heterogeneous loadings and materials. *Exp. Mech.* 2008; 48:435–449.
- [13]. Ladevèze P, Nedjar D, Reynier M. Updating of finite element models using vibration tests. *AIAA Journal*. 1994; 32:1485–1491.
- [14]. Ladevèze P, Chouaki A. Application of a *posteriori* error estimation for structural model updating. *Inverse Problems*. 1999; 15:49–58.
- [15]. Barthe D, Deraemaeker A, Ladevèze P, Le Loch S. Validation and updating of industrial models based on the constitutive relation error. *AIAA JOURNAL*. 2004; 42:1427–1434.
- [16]. Allix O, Feissel P, Nguyen HM. Identification strategy in the presence of corrupted measurements. *Engineering Computations*. 2005; 22:487–504.
- [17]. Feissel P, Allix O. Modified constitutive relation error identification strategy for transient dynamics with corrupted data : the elastic case. *Comp. Meth. Appl. Mech. Eng.* 2007; 196:1968–1983.
- [18]. Nguyen H-M, Allix O, Feissel P. A robust identification strategy for rate-dependent models in dynamics. *Inverse Problems*. 2008; 24:065006.
- [19]. Andrieux S, Ben Abda A. Solving Cauchy problems by minimizing an energy-like functional. *Inverse Problems*. 2006; 22:115–133.
- [20]. Kohn R, McKenney A. Numerical implementation of a variational method for electrical impedance tomography. *Inverse Problems*. 1990; 6:389–414.
- [21]. Knowles I. Parameter identification for elliptic problems. *J. Comp. Appl. Math.* 2001; 131:175–194.
- [22]. Chavent G, Kunisch K, Roberts JE. Primal-dual formulations for parameter estimation problems. *Computational and Applied Mathematics*. 1999; 18:173–229.
- [23]. Gockenbach MS, Khan AA. An abstract framework for elliptic inverse problems. Part 1: An output least-squares approach. *Mathematics and Mechanics of Solids*. 2007; 12:259–276.
- [24]. Wang Y, Yang J, Yin YZW. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences*. 2008; 3:248–272.
- [25]. Tikhonov, AN.; Goncharski, AV.; Stepanov, VV.; Yagoda, AG. *Numerical methods for the solution of ill-posed problems*. Kluwer Academic Publishers; 1995.
- [26]. Hansen, PC. *Rank-deficient and discrete ill-posed problems*. SIAM; Philadelphia, USA: 1998.
- [27]. Saad, Y. *Iterative Methods for Sparse Linear Systems*. SIAM; 2004.

- [28]. Farhat C, Macedo A, Lesoinne M, Roux F-X, Magoules F, de la Bourdonnaye A. Two-level domain decomposition methods with Lagrange multipliers for the fast solution of acoustic scattering problems. *Comp. Meth. Appl. Mech. Eng.* 2000; 184:213.
- [29]. Farhat, C.; Macedo, A.; Tezaur, R. FETI-H: A scalable domain decomposition method for high frequency exterior Helmholtz problems. In: Lai, Choi-Hong; Bjorstad, Petter; Cross, Mark; Widlund, Olof, editors. *Eleventh International Conference on Domain Decomposition Methods*; 1999. p. 231-241.
- [30]. Avery P, Farhat C, Reese G. Fast frequency sweep computations using a multi-point Pad+—based reconstruction method and an efficient iterative solver. *Int. J. Num. Meth. Eng.* 2007; 69:2848–2875.
- [31]. Liew HL, Pinsky PM. Recovery of shear modulus in elastography using an adjoint method with B-spline representation. *Finite Element Analysis and Design.* 2005; 41:778–799.
- [32]. Muthupillai R, Lomas D, Rossman P, Greenleaf J, Manduca A, Ehman R. Magnetic resonance elastography by direct visualization of propagating acoustic strain waves. *Science.* 1995; 269:1854–1857. [PubMed: 7569924]
- [33]. Zhu C, Byrd RH, Nocedal J. L-BFGS-B: algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization. *ACM Transaction on Mathematical Software.* 1997; 23:550–560.
- [34]. Aguilo, MA. Ph.D. thesis. Cornell University; USA: 2010. Inverse strategies for characterization of material properties.

- We develop an ECE strategy suitable for large scale parameter identification
- We propose a continuation scheme for accelerating convergence
- We formulate a block-SOR scheme for large scale problems
- The proposed MECE scheme showed faster convergence than a conventional L2/BFGS algorithm

Figure 1.
Schematic of the 2D example problem.

Figure 2.
2D example: reconstructions with data noise level $\delta = 0.05$. Units: Pa.

Figure 3.
2D example: cross-sectional plot at different noise levels. Units: Pa.

Figure 4.
2D example: cross-sectional plot at different noise levels when only G is reconstructed.
Units: Pa.

Figure 5. 2D example: effect of the proposed continuation scheme versus a constant penalization parameter on the convergence of the MECE algorithm (noise level $\delta = 0.05$).

(a) shear modulus

(b) bulk modulus

Figure 6.
2D example: cross-sectional plot corresponding to Figure 5. Units: Pa.

Figure 7.

2D example: effect of β for a noise level $\delta = 0.01$. The case $\beta = 0.0$ represents the constant penalty parameter case.

Figure 8.
2D example: cross-sectional plots corresponding to various choices of β . Units: Pa.

Figure 9. Comparison between MECE and L^2 /BFGS: relative errors e_G , e_B versus iterations

Figure 10.
Comparison between MECE and L^2 /BFGS: solutions. Units: Pa.

Figure 11.
Geometry of the three-dimensional model of a cylinder surrounded by an elastic matrix.

Figure 12.

Clip plane contour plot of the bulk modulus corresponding to the 3D example. Units: Pa

Figure 13.
Threshold plot of the bulk modulus corresponding to the 3D example. Units: Pa

Figure 14.

A clip plane contour plot of the shear modulus from the solution of the inverse problem on a cylinder in an elastic matrix. Units: Pa

Figure 15.

A threshold plot of the shear modulus from the solution of the inverse problem on a cylinder in an elastic matrix. Units: Pa

Table 1

Target Material Properties for 2D Example

Material Properties	Background	Inclusion 1	Inclusion 2
Shear Modulus G (Pa)	1×10^6	2×10^6	4×10^6
Bulk Modulus B (Pa)	2×10^6	3×10^6	5×10^6