

Home Search Collections Journals About Contact us My IOPscience

A penalty method for PDE-constrained optimization in inverse problems

This content has been downloaded from IOPscience. Please scroll down to see the full text. 2016 Inverse Problems 32 015007 (http://iopscience.iop.org/0266-5611/32/1/015007)

View the table of contents for this issue, or go to the journal homepage for more

Download details:

IP Address: 142.103.160.110 This content was downloaded on 24/12/2015 at 13:31

Please note that terms and conditions apply.

Inverse Problems 32 (2016) 015007 (26pp)

doi:10.1088/0266-5611/32/1/015007

A penalty method for PDE-constrained optimization in inverse problems

T van Leeuwen¹ and F J Herrmann²

¹ Mathematical Institute, Utrecht University, Utrecht, The Netherlands ² Dept. of Earth, Ocean and Atmospheric Sciences, University of British Columbia, Vancouver (BC), Canada

E-mail: T.vanLeeuwen@uu.nl

Received 10 April 2015, revised 23 September 2015 Accepted for publication 13 November 2015 Published 16 December 2015



Abstract

Many inverse and parameter estimation problems can be written as PDE-constrained optimization problems. The goal is to infer the parameters, typically coefficients of the PDE, from partial measurements of the solutions of the PDE for several right-hand sides. Such PDE-constrained problems can be solved by finding a stationary point of the Lagrangian, which entails simultaneously updating the parameters and the (adjoint) state variables. For large-scale problems, such an *all-at-once* approach is not feasible as it requires storing all the state variables. In this case one usually resorts to a *reduced* approach where the constraints are explicitly eliminated (at each iteration) by solving the PDEs. These two approaches, and variations thereof, are the main workhorses for solving PDE-constrained optimization problems arising from inverse problems. In this paper, we present an alternative method that aims to combine the advantages of both approaches. Our method is based on a quadratic penalty formulation of the constrained optimization problem. By eliminating the state variable, we develop an efficient algorithm that has roughly the same computational complexity as the conventional reduced approach while exploiting a larger search space. Numerical results show that this method indeed reduces some of the nonlinearity of the problem and is less sensitive to the initial iterate.

Keywords: inverse problems, PDE-constrained optimization, penalty method

(Some figures may appear in colour only in the online journal)

1. Introduction

In inverse problems, the goal is to infer physical parameters (e.g., density, soundspeed or conductivity) from indirect observations. When the underlying model is described by a partial

0266-5611/16/015007+26\$33.00 © 2016 IOP Publishing Ltd Printed in the UK

differential equation (PDE) (e.g., the wave-equation or Maxwell's equations), the observed data are typically partial measurements of the solutions of the PDE for multiple right-hand sides. The parameters typically appear as coefficients in the PDE. These problems arise in many applications such as geophysics [1–4], medical imaging [5, 6] and non-destructive testing.

For linear PDEs, the inverse problem can be formulated (after discretization) as a constrained optimization problem of the form

$$\min_{\mathbf{m},\mathbf{u}} \frac{1}{2} \| P\mathbf{u} - \mathbf{d} \|_2^2 \quad \text{s.t.} \quad A(\mathbf{m})\mathbf{u} = \mathbf{q},$$
(1)

where $\mathbf{m} \in \mathbb{R}^M$ represents the (gridded) parameter of interest, $A(\mathbf{m}) \in \mathbb{C}^{N \times N}$ and $\mathbf{q} \in \mathbb{C}^N$ represent the discretized PDE and source term, $\mathbf{u} \in \mathbb{C}^N$ is the state variable and $\mathbf{d} \in \mathbb{C}^L$ are the observed data. The measurement process is modeled by sampling the state with $P \in \mathbb{R}^{L \times N}$. Throughout the paper we assume that *A* stems from a well-posed PDE and hence is invertible. We further use T to denote the (complex-conjugate) transpose.

Typically, measurements are made from multiple, say *K*, independent experiments, in which case $\mathbf{u} \in \mathbb{C}^{KN}$ is a block vector containing the state variables for all the experiments. Likewise, $\mathbf{q} \in \mathbb{C}^{KN}$ and $\mathbf{d} \in \mathbb{C}^{KL}$ are block vectors containing the right-hand sides and observations for all experiments. The matrices $A(\mathbf{m})$ and P will be block-diagonal matrices in this case. Typical sizes for M, N, K, L for seismic inverse problems are listed in table 1.

In practice, one usually includes a regularization term in the formulation (1) to mitigate the ill-posedness of the problem. To simplify the discussion, however, we ignore such terms with the understanding that appropriate regularization terms can be added when required.

1.1. All-at-once and reduced methods

In applications arising from inverse problems, the constrained problem (1) is typically solved using the method of Lagrange multipliers [7, 8] or sequential quadratic programming (SQP) [9, 10]. This entails optimizing over both the parameters, states and the Lagrange multipliers (or adjoint-state variables) simultaneously. While such *all-at-once* approaches are often very attractive from an optimization point-of-view, they are typically not feasible for large-scale problems since we cannot afford to store the state variables for all *K* experiments simultaneously. Instead, the so-called *reduced* approach is based on a (block) elimination of the constraints to formulate an unconstrained optimization problem over the parameters:

$$\min_{\mathbf{m}} \frac{1}{2} \left\| PA(\mathbf{m})^{-1} \mathbf{q} - \mathbf{d} \right\|_{2}^{2}.$$
 (2)

While this eliminates the need to store the full state variables for all *K* experiments, evaluation of the objective and its gradient requires PDE-solves. Moreover, by eliminating the constraints we have dramatically reduced the search-space, thus arguably making it more difficult to find an appropriate minimizer. Note also that the dependency of the objective on **m** is now through $A(\mathbf{m})^{-1}\mathbf{q}$ in stead of through $A(\mathbf{m})\mathbf{u}$. For linear PDEs, the latter can often be made to depend linearly on **m** while the dependency of $A(\mathbf{m})^{-1}\mathbf{q}$ on **m** is typically nonlinear.

1.2. Motivation

The main motivation for this work is the observation that the stated inverse problem would be much easier to solve if we had a *complete* measurement of the state (i.e., *P* is invertible). In this case, we could reconstruct the state from the data as $\mathbf{u} = P^{-1}\mathbf{d}$ and subsequently recover the parameter by solving

Table 1. Typical size of seismic inverse problem in terms K: of the number of experiments, L: the number of measurements per experiment, N: the number of discretization points and M: the number of parameters.

	Small 2D	Large 2D	Industrial 3D
K	10^{2}	10 ³	10 ⁶
L	10^{2}	10^{3}	10^{6}
М	10 ⁶	10 ⁹	10^{12} 10^{9}
Ν	10^{3}	10 ⁶	10 ⁹

$$\min_{\mathbf{m}} \frac{1}{2} \|A(\mathbf{m})\mathbf{u} - \mathbf{q}\|_2^2, \tag{3}$$

which for many linear PDEs would lead to a linear least-squares problem. This approach is known in the literature as the *equation-error approach* [11, 12]. When we do not have complete measurements, this method does not apply directly since we cannot invert P. We can, however, aim to recover the state by solving the following (inconsistent) overdetermined system

$$\begin{pmatrix} A(\mathbf{m}) \\ P \end{pmatrix} \mathbf{u} \approx \begin{pmatrix} \mathbf{q} \\ \mathbf{d} \end{pmatrix},$$
 (4)

which combines the physics and the data. We can subsequently use the obtained estimate of the state to estimate **m** by solving (3). These steps can be repeated in an alternating fashion as needed. In a previous paper [13], we proposed this methodology for seismic inversion and coined it wavefield reconstruction inversion (WRI). We showed—via numerical experiments —that this approach can mitigate some of the (notorious) nonlinearity of the seismic inverse problem. In the current paper we seek to analyze this approach in more detail and broaden the scope of its application to inverse problems involving PDEs.

1.3. Contributions and outline

We give the above sketched method a sound theoretical basis by showing that it can be derived from a *penalty* formulation of the constrained problem:

$$\min_{\mathbf{m},\mathbf{u}} \frac{1}{2} \|P\mathbf{u} - \mathbf{d}\|_2^2 + \frac{\lambda}{2} \|A(\mathbf{m})\mathbf{u} - \mathbf{q}\|_2^2,$$
(5)

the solution of which (theoretically) satisfies the optimality conditions of the constrained problem (1) as $\lambda \to \infty$. Such reformulations of the constrained problem are well-known but are, to our best knowledge, not being applied to inverse problems.

The main contribution of this paper is the development of an efficient algorithm based on the penalty formulation for inverse problems and the insight that we can approximate the solution of the constrained problem (1) up to arbitrary finite precision with a finite λ . In (illposed) inverse problems, we often do not require very high precision because we can only hope to resolve certain components of **m** anyway. We can understand this by realizing that not all constraints are equally important; those that constrain the null-space components of **m** need not be enforced as strictly. Thus, the parameter λ need only be large enough to enforce the dominant constraints. The numerical experiments suggest that a single fixed value of λ is typically sufficient.

Our approach is based on the elimination of the state variable, \mathbf{u} , from (5). This reduces the dimensionality of the optimization problem in a similar fashion as the *reduced* approach

(2) does for the constrained formulation (1) by solving *K* systems of equations. This elimination leads to a cost function $\phi_{\lambda}(\mathbf{m})$ whose gradient and Hessian can be readily computed. The main difference is that the state \mathbf{u} in this case is not defined by solving the PDE, but instead is solved from an overdetermined system that involves both the PDE *and* the data (4). Due to the special block-structure of the problems under consideration, this elimination can be done efficiently, leading to a tractable algorithm. Contrary to the conventional *reduced* approach, the resulting algorithm does *not* enforce the constraints at each iteration and arguably leads to a less nonlinear problem in \mathbf{m} . It is outside the scope of the current paper to give a rigorous proof of this statement, but we present some numerical evidence to support this conjecture.

The outline of the paper is as follows. First, we give a brief overview of the constrained and penalty formulations in sections 2 and 3. The main theoretical results are presented in section 4 while a detailed description of the proposed algorithm is given in section 5. Here, we also compare the penalty approach to both the all-at-once and the reduced approaches in terms of algorithmic complexity. Numerical examples on a 1D dc-resistivity and 2D ultrasound and seismic tomography problems are given in section 6. Possible extensions and open problems are discussed in sections 7 and 8 gives the conclusions.

2. All-at-once and reduced methods

A popular approach to solving constrained problems of the form (1) is based on the corresponding Lagrangian:

$$\mathcal{L}(\mathbf{m}, \mathbf{u}, \mathbf{v}) = \frac{1}{2} ||P\mathbf{u} - \mathbf{d}||_2^2 + \mathbf{v}^T (A(\mathbf{m})\mathbf{u} - \mathbf{q}),$$
(6)

where $\mathbf{v} \in \mathbb{C}^{KN}$ is the Lagrange multiplier or adjoint-state variable [7, 14]. A necessary condition for a solution (\mathbf{m}^* , \mathbf{u}^* , \mathbf{v}^*) of the constrained problem (1) is that it is a stationary point of the Lagrangian, i.e. $\nabla \mathcal{L}(\mathbf{m}^*, \mathbf{u}^*, \mathbf{v}^*) = 0$. The gradient and Hessian of the Lagrangian are given by

$$\nabla \mathcal{L}(\mathbf{m}, \mathbf{u}, \mathbf{v}) = \begin{pmatrix} \mathcal{L}_{\mathbf{m}} \\ \mathcal{L}_{\mathbf{u}} \\ \mathcal{L}_{\mathbf{v}} \end{pmatrix} = \begin{pmatrix} G(\mathbf{m}, \mathbf{u})^T \mathbf{v} \\ A(\mathbf{m})^T \mathbf{v} + P^T(P\mathbf{u} - \mathbf{d}) \\ A(\mathbf{m})\mathbf{u} - \mathbf{q} \end{pmatrix},$$
(7)

and

$$\nabla^{2} \mathcal{L}(\mathbf{m}, \mathbf{u}, \mathbf{v}) = \begin{pmatrix} R(\mathbf{m}, \mathbf{u}, \mathbf{v}) & K(\mathbf{m}, \mathbf{v})^{T} & G(\mathbf{m}, \mathbf{u})^{T} \\ K(\mathbf{m}, \mathbf{v}) & P^{T} P & A(\mathbf{m})^{T} \\ G(\mathbf{m}, \mathbf{u}) & A(\mathbf{m}) & 0 \end{pmatrix},$$
(8)

where

$$G(\mathbf{m}, \mathbf{u}) = \frac{\partial A(\mathbf{m})\mathbf{u}}{\partial \mathbf{m}}, \quad K(\mathbf{m}, \mathbf{v}) = \frac{\partial A(\mathbf{m})^T \mathbf{v}}{\partial \mathbf{m}},$$
$$R(\mathbf{m}, \mathbf{u}, \mathbf{v}) = \frac{\partial G(\mathbf{m}, \mathbf{u})^T \mathbf{v}}{\partial \mathbf{m}}.$$

These Jacobian matrices are typically sparse when A is sparse and can be computed analytically.

2.1. All-at-once approach

So-called all-at-once approaches find such a stationary point by applying a Newton-like method to the Lagrangian [7, 15]. A basic algorithm for finding a stationary point of the Lagrangian up to a given tolerance ϵ is given in algorithm 1. If the algorithm converges it returns iterates ($\mathbf{m}^*, \mathbf{u}^*, \mathbf{v}^*$) such that $\|\nabla \mathcal{L}(\mathbf{m}^*, \mathbf{u}^*, \mathbf{v}^*)\|_2 \leq \epsilon$. Many variants of algorithm 1 exist and may include preconditioning, inexact solves of the

Algorithm 1. Basic Newton algorithm for finding a stationary point of the Lagrangian via the all-at-once method

Require: initial guess \mathbf{m}^0 , \mathbf{u}^0 , \mathbf{v}^0 , tolerance ϵ k = 0while $\|\nabla \mathcal{L}(\mathbf{m}^k, \mathbf{u}^k, \mathbf{v}^k)\|_2 \ge \epsilon$ do $\left(\delta \mathbf{m}^k\right)$

 $\begin{pmatrix} \delta \mathbf{m}^{k} \\ \delta \mathbf{v}^{k} \\ \delta \mathbf{v}^{k} \end{pmatrix} = -(\nabla^{2} \mathcal{L}(\mathbf{m}^{k}, \mathbf{u}^{k}, \mathbf{v}^{k}))^{-1} \nabla \mathcal{L}(\mathbf{m}^{k}, \mathbf{u}^{k}, \mathbf{v}^{k})$ determine steplength $\alpha^{k} \in (0, 1]$ $\mathbf{m}^{k+1} = \mathbf{m}^{k} + \alpha^{k} \delta \mathbf{m}^{k}$ $\mathbf{u}^{k+1} = \mathbf{u}^{k} + \alpha^{k} \delta \mathbf{u}^{k}$ $\mathbf{v}^{k+1} = \mathbf{v}^{k} + \alpha^{k} \delta \mathbf{v}^{k}$ k = k + 1d while

end while

KKT system $(\nabla^2 \mathcal{L})^{-1} \nabla \mathcal{L}$ and a linesearch to ensure global convergence [8, 15, 16]. For an extensive overview we refer to [17].

An advantage of such an all-at-once approach is that it eliminates the need to solve the PDEs explicitly; the constraints are only (approximately) satisfied upon convergence. However, such an approach is not feasible for the applications we have in mind because it involves simultaneously updating (and hence storing) all the variables.

2.2. Reduced approach

In inverse problems one usually considers a *reduced* formulation that is obtained by eliminating the constraints from (1). This results in an unconstrained optimization problem:

$$\min_{\mathbf{m}} \left\{ \phi(\mathbf{m}) = \frac{1}{2} \left| \left| P \mathbf{u}_{\text{red}}(\mathbf{m}) - \mathbf{d} \right| \right|_{2}^{2} \right\},\tag{9}$$

where $\mathbf{u}_{red}(\mathbf{m}) = A(\mathbf{m})^{-1}\mathbf{q}$. The resulting optimization problem has a much smaller dimension and can be solved using black-box nonlinear optimization methods. In contrast to the *all-at-once* method, the constraints are satisfied at each iteration.

The gradient and the Hessian of ϕ are given by

$$\nabla \phi(\mathbf{m}) = G(\mathbf{m}, \mathbf{u}_{\text{red}})^T \mathbf{v}_{\text{red}},\tag{10}$$

$$\nabla^2 \phi(\mathbf{m}) = G(\mathbf{m}, \mathbf{u}_{\text{red}})^T A(\mathbf{m})^{-T} P^T P A(\mathbf{m})^{-1} G(\mathbf{m}, \mathbf{u}_{\text{red}})$$
$$- K(\mathbf{m}, \mathbf{v}_{\text{red}})^T A(\mathbf{m})^{-1} G(\mathbf{m}, \mathbf{u}_{\text{red}})$$
$$- G(\mathbf{m}, \mathbf{u}_{\text{red}})^T A(\mathbf{m})^{-T} K(\mathbf{m}, \mathbf{v}_{\text{red}})$$
$$+ R(\mathbf{m}, \mathbf{u}_{\text{red}}, \mathbf{v}_{\text{red}}), \tag{11}$$

where $\mathbf{v}_{red} = A(\mathbf{m})^{-T}P^{T}(\mathbf{d} - P\mathbf{u}_{red})$.

۲

A basic (Gauss–Newton (GN)) algorithm for minimizing ϕ (**m**) is given in algorithm 2. Note that this corresponds to a block-elimination of the KKT system and the iterates automatically satisfy $\mathcal{L}_{\mathbf{u}}(\mathbf{m}^k, \mathbf{u}_{red}^k, \mathbf{v}_{red}^k) = \mathcal{L}_{\mathbf{v}}(\mathbf{m}^k, \mathbf{u}_{red}^k, \mathbf{v}_{red}^k) = 0$. If the algorithm terminates successfully, the final iterates ($\mathbf{m}^*, \mathbf{u}_{red}^*, \mathbf{v}_{red}^*$) additionally satisfy $\|\mathcal{L}_{\mathbf{m}}(\mathbf{m}^*, \mathbf{u}_{red}^*, \mathbf{v}_{red}^*)\|_2 \leq \epsilon$, so that $\|\nabla \mathcal{L}(\mathbf{m}^*, \mathbf{u}_{red}^*, \mathbf{v}_{red}^*)\|_2 \leq \epsilon$.

Algorithm 2. Basic GN algorithm for find a stationary point of the Lagrangian via the reduced method

Require: initial guess \mathbf{m}^0 , tolerance ϵ

k = 0 $\mathbf{u}_{red}^{0} = A(\mathbf{m}^{0})^{-1}\mathbf{q}$ $\mathbf{v}_{red}^{0} = A(\mathbf{m}^{0})^{-T}P^{T}(\mathbf{d} - P\mathbf{u}_{red}^{0})$ while $\|\mathcal{L}_{\mathbf{m}}(\mathbf{m}^{k}, \mathbf{u}_{red}^{k}, \mathbf{v}_{red}^{k})\|_{2} \ge \epsilon$ do $\mathbf{g}_{red}^{k} = G(\mathbf{m}^{k}, \mathbf{u}_{red}^{k})^{T}\mathbf{v}_{red}^{red}$ $H_{red}^{k} = G(\mathbf{m}^{k}, \mathbf{u}_{red}^{k})^{T}A(\mathbf{m}^{k})^{-T}P^{T}PA(\mathbf{m}^{k})^{-1}G(\mathbf{m}^{k}, \mathbf{u}_{red}^{k})$ determine steplength $\alpha^{k} \in (0, 1]$ $\mathbf{m}^{k+1} = \mathbf{m}^{k} - \alpha^{k}(H_{red}^{k})^{-1}\mathbf{g}_{red}^{k}$ $\mathbf{u}_{red}^{k+1} = A(\mathbf{m}^{k+1})^{-1}\mathbf{q}$ $\mathbf{v}_{red}^{k+1} = A(\mathbf{m}^{k+1})^{-T}P^{T}(\mathbf{d} - P\mathbf{u}_{red}^{k+1})$ k = k + 1end while

A disadvantage of this approach is that it requires the solution of the PDEs at each update, making it computationally expensive. It also strictly enforces the constraint at each iteration, possibly leading to a very nonlinear problem in \mathbf{m} .

3. Penalty and augmented Lagrangian methods

It is impossible to do justice to the wealth of research that has been done on penalty and augmented Lagrangian methods in one section. Instead, we give a brief overview highlighting the main characteristics of a few basic approaches and their limitations when applied to inverse problems.

A constrained optimization problem of the form (1) can be recast as an unconstrained problem by introducing a non-negative penalty function $\pi : \mathbb{C}^N \to \mathbb{R}_{\geq 0}$ and a penalty parameter $\lambda > 0$ as follows

$$\min_{\mathbf{m},\mathbf{u}} \frac{1}{2} ||P\mathbf{u} - \mathbf{d}||_2^2 + \lambda \pi (\mathbf{A}(\mathbf{m})\mathbf{u} - \mathbf{q}).$$
(12)

The idea is that any departure from the constraint is penalized so that the solution of this unconstrained problem will coincide with that of the constrained problem when λ is large enough.

A quadratic penalty function $\pi(\cdot) = \frac{1}{2} ||\cdot||_2^2$ leads to a differentiable unconstrained optimization problem (12) whose minimizer coincides with the solution of the constrained optimization problem (1) when $\lambda \to \infty$ [14, theorem 17.1]. Practical algorithms rely on repeatedly solving the unconstrained problem for increasing values of λ .

A common concern with this approach is that the Hessian may become increasingly illconditioned as $\lambda \to \infty$ when there are fewer constraints than variables. For PDE-constrained optimization in inverse problems, there are enough constraints ($A(\mathbf{m})$ is invertible) to prevent this. We discuss this limiting case in more detail in section 5.

For certain non-smooth penalty functions, such as $\pi(\cdot) = || \cdot ||_1$, the minimizer of (12) is a solution of the constrained problem for $any \lambda \ge \lambda^*$ for some λ^* [14, theorem 17.3]. In practice, a continuation strategy is used to find a suitable value for λ . An advantage of this approach is that λ does not become arbitrarily large, thus avoiding the ill-conditioning problems mentioned above. A disadvantage is that the resulting unconstrained problem is no longer differentiable. With large-scale applications in mind, we therefore do not consider exact penalty methods any further in this paper.

Another approach that avoids having to increase λ to infinity is the *augmented* Lagrangian approach (see [14]). In this approach, a quadratic penalty $\lambda \|\mathbf{A}(\mathbf{m})\mathbf{u} - \mathbf{q}\|_2^2$ is added to the Lagrangian (6). A standard approach to solve the constrained problem based on the augmented Lagrangian is the *alternating direction method of multipliers*. In its most basic form it relies on minimizing the augmented Lagrangian w.r.t. (\mathbf{m}, \mathbf{u}) and subsequently updating the multiplier \mathbf{v} and the penalty parameter λ [18, 19]. This would require us to store the multipliers, which is not feasible for the problems we have in mind.

In the next two sections, we discuss a computationally efficient algorithm for solving the constrained optimization problem (1) based on a quadratic penalty formulation. This formulation is attractive because it leads to a differentiable, unconstrained, optimization problem. Moreover, the optimization in \mathbf{u} has a closed-form solution which can be computed efficiently, making it an ideal candidate for the type of problems we have in mind.

4. A reduced penalty method

Using a quadratic penalty function, the constrained problem (1) is reformulated as

$$\min_{\mathbf{m},\mathbf{u}} \left\{ \mathcal{P}(\mathbf{m},\mathbf{u}) = \frac{1}{2} ||P\mathbf{u} - \mathbf{d}||_2^2 + \frac{1}{2}\lambda ||A(\mathbf{m})\mathbf{u} - \mathbf{q}||_2^2 \right\}.$$
 (13)

The gradient and Hessian of \mathcal{P} are given by

$$\nabla \mathcal{P} = \begin{pmatrix} \mathcal{P}_{\mathbf{m}} \\ \mathcal{P}_{\mathbf{u}} \end{pmatrix} = \begin{pmatrix} \lambda G(\mathbf{m}, \mathbf{u})^{T} (A(\mathbf{m})\mathbf{u} - \mathbf{q}) \\ P^{T} (P\mathbf{u} - \mathbf{d}) + \lambda A(\mathbf{m})^{T} (A(\mathbf{m})\mathbf{u} - \mathbf{q}) \end{pmatrix},$$
(14)

and

$$\nabla^{2} \mathcal{P} = \begin{pmatrix} \mathcal{P}_{\mathbf{m},\mathbf{m}} & \mathcal{P}_{\mathbf{m},\mathbf{u}} \\ \mathcal{P}_{\mathbf{u},\mathbf{m}} & \mathcal{P}_{\mathbf{u},\mathbf{u}} \end{pmatrix},\tag{15}$$

where

$$\mathcal{P}_{\mathbf{m},\mathbf{m}} = \lambda \Big(G(\mathbf{m}, \mathbf{u})^T G(\mathbf{m}, \mathbf{u}) + R(\mathbf{m}, \mathbf{u}, A(\mathbf{m})\mathbf{u} - \mathbf{q}) \Big), \tag{16}$$

$$\mathcal{P}_{\mathbf{u},\mathbf{u}} = P^T P + \lambda A(\mathbf{m})^T A(\mathbf{m}), \tag{17}$$

$$\mathcal{P}_{\mathbf{m},\mathbf{u}} = \lambda \Big(K(\mathbf{m}, A(\mathbf{m})\mathbf{u} - \mathbf{q}) + A(\mathbf{m})^T G(\mathbf{m}, \mathbf{u}) \Big).$$
(18)

Of course, optimization in the full (\mathbf{m}, \mathbf{u}) -space is not feasible for large-scale problems, so we eliminate \mathbf{u} by introducing

$$\mathbf{u}_{\lambda}(\mathbf{m}) = \operatorname{argmin}_{\mathbf{u}} \mathcal{P}(\mathbf{m}, \mathbf{u}), \tag{19}$$

and define a reduced objective:

$$\phi_{\lambda}(\mathbf{m}) = \mathcal{P}(\mathbf{m}, \mathbf{u}_{\lambda}(\mathbf{m})).$$
⁽²⁰⁾

The optimization problem for the state (19) has a closed-form solution:

$$\mathbf{u}_{\lambda} = \left(A(\mathbf{m})^{T} A(\mathbf{m}) + \lambda^{-1} P^{T} P\right)^{-1} \left(A(\mathbf{m})^{T} \mathbf{q} + \lambda^{-1} P^{T} \mathbf{d}\right).$$

The modified system $A^TA + \lambda^{-1}P^TP$ is a low-rank modification of the original PDE and incorporates the measurements in the PDE solve. This is the main difference with the conventional reduced approach (see algorithm 2); the estimate of the state is not only based on the physics and the current model, but also on the data.

Following [20, theorem 1], it is readily verified that the gradient and Hessian of ϕ_{λ} are given by

$$\nabla \phi_{\lambda}(\mathbf{m}) = \mathcal{P}_{\mathbf{m}}(\mathbf{m}, \mathbf{u}_{\lambda}), \tag{21}$$

$$\nabla^2 \phi_{\lambda}(\mathbf{m}) = \mathcal{P}_{\mathbf{m},\mathbf{m}}(\mathbf{m},\mathbf{u}_{\lambda}) - \mathcal{P}_{\mathbf{m},\mathbf{u}}(\mathbf{m},\mathbf{u}_{\lambda}) \big(\mathcal{P}_{\mathbf{u},\mathbf{u}}(\mathbf{m},\mathbf{u}_{\lambda}) \big)^{-1} \mathcal{P}_{\mathbf{u},\mathbf{m}}(\mathbf{m},\mathbf{u}_{\lambda}).$$
(22)

Note that $\nabla^2 \phi_{\lambda}$ is the Schur complement of $\nabla^2 \mathcal{P}$.

A basic GN algorithm for minimizing ϕ_{λ} is shown in algorithm 3. Note that the computation of the adjoint-state \mathbf{v}_{λ} does *not* require an additional PDE-solve in this algorithm. Instead, the forward and adjoint solve are done simultaneously via the normal equations.

Algorithm 3. Basic GN algorithm for find a stationary point of the Lagrangian via the penalty method

Require: initial guess \mathbf{m}^0 , penalty parameter λ , tolerance ϵ k = 0 $\mathbf{u}_{\lambda}^0 = (A(\mathbf{m}^0)^T A(\mathbf{m}^0) + \lambda^{-1} P^T P)^{-1} (A(\mathbf{m}^0)^T \mathbf{q} + \lambda^{-1} P^T \mathbf{d})$ $\mathbf{v}_{\lambda}^0 = \lambda (A(\mathbf{m}^0) \mathbf{u}_{\lambda}^0 - \mathbf{q})$ **while** $\|\mathcal{L}_{\mathbf{m}}(\mathbf{m}^k, \mathbf{u}_{\lambda}^k, \mathbf{v}_{\lambda}^k)\|_2 \ge \epsilon$ do $\mathbf{g}_{\lambda}^k = G(\mathbf{m}^k, \mathbf{u}_{\lambda}^k)^T \mathbf{v}_{\lambda}^k$ $H_{\lambda}^k = \lambda G^T (I - A(A^T A + \lambda^{-1} P^T P)^{-1} A^T) G$ determine steplength $\alpha^k \in (0, 1]$ $\mathbf{m}^{k+1} = \mathbf{m}^k - \alpha^k (H_{\lambda}^k)^{-1} \mathbf{g}_{\lambda}^k$ $\mathbf{u}_{\lambda}^{k+1} = (A(\mathbf{m}^{k+1})^T A(\mathbf{m}^{k+1}) + \lambda^{-1} P^T P)^{-1} (A(\mathbf{m}^{k+1})^T \mathbf{q} + \lambda^{-1} P^T \mathbf{d})$ $\mathbf{v}_{\lambda}^{k+1} = \lambda (A(\mathbf{m}^{k+1}) \mathbf{u}_{\lambda}^{k+1} - \mathbf{q})$ k = k + 1**end while**

Next, we show how the states, \mathbf{u}_{λ}^{k} and \mathbf{v}_{λ}^{k} , generated by this algorithm relate to the states generated by the reduced approach and subsequently that if the algorithm successfully terminates the iterates ($\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*}$) satisfy

$$\|
abla \mathcal{L}ig(\mathbf{m}^*,\,\mathbf{u}^*_{\!\lambda},\,\mathbf{v}^*_{\!\lambda}ig)\|_2\leqslant \epsilon+\mathcal{O}ig(\,\lambda^{\!-1}ig).$$

Lemma 4.1. For a fixed **m**, the states \mathbf{u}_{λ} and \mathbf{v}_{λ} used in the reduced penalty approach (algorithm 3) are related to the states \mathbf{u}_{red} and \mathbf{v}_{red} used in the reduced approach (algorithm

2) as follows

$$\mathbf{u}_{\lambda} = \mathbf{u}_{\rm red} + \mathcal{O}(\lambda^{-1}), \tag{23}$$

$$\mathbf{v}_{\lambda} = \mathbf{v}_{\text{red}} + \mathcal{O}(\lambda^{-1}). \tag{24}$$

Proof. The state variables used in the penalty approach are given by

$$\mathbf{u}_{\lambda} = \left(A^{T}A + \lambda^{-1}P^{T}P\right)^{-1}\left(A^{T}\mathbf{q} + \lambda^{-1}P^{T}\mathbf{d}\right),$$

and

$$\mathbf{v}_{\lambda} = \lambda \big(A \mathbf{u}_{\lambda} - \mathbf{q} \big).$$

The former can be re-written as

$$\mathbf{u}_{\lambda} = A^{-1} \left(I + \lambda^{-1} A^{-T} P^{T} P A^{-1} \right)^{-1} \left(\mathbf{q} + \lambda^{-1} A^{-T} P^{T} \mathbf{d} \right)$$

For $\lambda > ||PA^{-1}||_2^2$ we may expand the inverse as $(I + \lambda^{-1}B)^{-1} \approx I - \lambda^{-1}B + \lambda^{-2}B^2 + \dots$, and find that

$$\mathbf{u}_{\lambda} = A^{-1}\mathbf{q} + \lambda^{-1}A^{-1}A^{-T}P^{T}(\mathbf{d} - PA^{-1}\mathbf{q}) - \lambda^{-2}A^{-1}(PA^{-1})^{T}(PA^{-1})A^{-T}P^{T}(\mathbf{d} - PA^{-1}\mathbf{q}) + \mathcal{O}(\lambda^{-3}) = \mathbf{u}_{red} + \lambda^{-1}A^{-1}(I - \lambda^{-1}(PA^{-1})^{T}(PA^{-1}))\mathbf{v}_{red} + \mathcal{O}(\lambda^{-3}).$$
(25)

We immediately find

$$\mathbf{v}_{\lambda} = \mathbf{v}_{\text{red}} - \lambda^{-1} (PA^{-1})^{T} (PA^{-1}) \mathbf{v}_{\text{red}} + \mathcal{O}(\lambda^{-2}).$$
⁽²⁶⁾

Remark. lemma 4.1 suggests a natural scaling for the penalty parameter; $\lambda > ||PA^{-1}||_2^2$ can be considered large, while $\lambda < ||PA^{-1}||_2^2$ can be considered small.

Theorem 4.2. At each iteration of algorithm 3, the iterates satisfy $\|\mathcal{L}_{\mathbf{u}}(\mathbf{m}^{k}, \mathbf{u}_{\lambda}^{k}, \mathbf{v}_{\lambda}^{k})\|_{2} = 0$ and $\|\mathcal{L}_{\mathbf{v}}(\mathbf{m}^{k}, \mathbf{u}_{\lambda}^{k}, \mathbf{v}_{\lambda}^{k})\|_{2} = \mathcal{O}(\lambda^{-1})$. Moreover, if algorithm 3 terminates successfully at \mathbf{m}^{*} for which $\|\mathcal{L}_{\mathbf{m}}(\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*})\|_{2} \leq \epsilon$, we have $\|\nabla \mathcal{L}(\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*})\|_{2} \leq \epsilon + \mathcal{O}(\lambda^{-1})$.

Proof. Using the definitions of \mathbf{u}_{λ} and \mathbf{v}_{λ} we find for any \mathbf{m}

$$\mathcal{L}_{\mathbf{u}}(\mathbf{m}, \mathbf{u}_{\lambda}, \mathbf{v}_{\lambda}) = A(\mathbf{m})^{T} \mathbf{v}_{\lambda} + P^{T} (P \mathbf{u}_{\lambda} - \mathbf{d})$$
$$= \lambda A^{T} (A \mathbf{u}_{\lambda} - \mathbf{q}) + P^{T} (P \mathbf{u}_{\lambda} - \mathbf{d}) = 0.$$
(27)

Using the approximations for \mathbf{u}_{λ} and \mathbf{v}_{λ} for $\lambda > \|PA^{-1}\|_{2}^{2}$ presented in lemma 4.1, we find

$$\mathcal{L}_{\mathbf{v}}(\mathbf{m}, \mathbf{u}_{\lambda}, \mathbf{v}_{\lambda}) = A(\mathbf{m})\mathbf{u}_{\lambda} - \mathbf{q}$$

= $\lambda^{-1}\mathbf{v}_{red} + \mathcal{O}(\lambda^{-2}).$ (28)

Thus we find

$$\|\mathcal{L}_{\mathbf{v}}(\mathbf{m}^*, \mathbf{u}^*_{\lambda}, \mathbf{v}^*_{\lambda})\|_2 \leqslant \mathcal{O}(\lambda^{-1}).$$
⁽²⁹⁾

At a point \mathbf{m}^* for which $\|\mathcal{L}_{\mathbf{m}}(\mathbf{m}^*, \mathbf{u}^*_{\lambda}, \mathbf{v}^*_{\lambda})\|_2 \leq \epsilon$ we immediately find that

. .

$$\begin{aligned} \|\nabla \mathcal{L} \left(\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*}\right)\|_{2}^{2} \\ &= \|\mathcal{L}_{\mathbf{m}} \left(\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*}\right)\|_{2}^{2} + \|\mathcal{L}_{\mathbf{u}} \left(\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*}\right)\|_{2}^{2} + \|\mathcal{L}_{\mathbf{v}} \left(\mathbf{m}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*}\right)\|_{2}^{2} \\ &\leqslant \epsilon^{2} + \mathcal{O} \left(\lambda^{-2}\right), \end{aligned}$$
(30)

and hence that

$$\|\nabla \mathcal{L}(\mathbf{m}^*, \mathbf{u}^*_{\lambda}, \mathbf{v}^*_{\lambda})\|_2 \leqslant \epsilon + \mathcal{O}(\lambda^{-1}).$$
(31)

This means we can use algorithm 3 to find a stationary point of the Lagrangian within finite accuracy (in terms of $\|\nabla \mathcal{L}\|_2$) with a finite λ . In order to reach a given tolerance, we need $\lambda^{-1} \|\mathcal{L}_{\mathbf{v}}\|_2$ to be small compared to $\|\mathcal{L}_{\mathbf{m}}\|_2$. This condition can easily be verified for a given **m** and thus serve as the basis for a selection criterion for λ .

In an ill-posed inverse problem, we can only hope to reconstruct a few components of \mathbf{m}^* , roughly corresponding the dominant eigenvectors of the reduced Hessian. Thus, the tolerance ϵ can be quite large, even for an acceptable reconstruction. Driving down the norm of the gradient any further would only refine the reconstruction in the eigenmodes corresponding to increasingly small eigenvalues and would hardly affect the final reconstruction and datafit.

Next, we derive an expression for the distance of the final iterate of algorithm 3, \mathbf{m}_{λ}^* , to a stationary point of the Lagrangian, \mathbf{m}^* in terms of the chosen tolerance, ϵ , and the dataerror $\eta = \|\mathbf{d} - PA(\mathbf{m}_{\lambda}^*)^{-1}\mathbf{q}\|_2$.

Theorem 4.3. Given a stationary point of the Lagrangian $(\mathbf{m}^*, \mathbf{u}^*, \mathbf{v}^*)$ and the final iterate of algorithm $3 \mathbf{m}^*_{\lambda}$ such that $\|\mathcal{L}_{\mathbf{m}}(\mathbf{m}^*_{\lambda}, \mathbf{u}^*_{\lambda}, \mathbf{v}^*_{\lambda})\|_2 \leq \epsilon$, we have $\|\mathbf{m}^*_{\lambda} - \mathbf{m}^*\|_2 \leq \kappa(H)_{(\tilde{\epsilon} + \tilde{\lambda}^{-1}\tilde{\eta})}$ where $H = J^T J$ is the reduced GN Hessian with condition number $\kappa(H)$, $\tilde{\epsilon} = \epsilon/\|H\|_2$ is the scaled tolerance, $\tilde{\eta} = \|\mathbf{d} - PA(\mathbf{m}^*_{\lambda})^{-1}\mathbf{q}\|_2/\|J\|_2$ is the scaled data-error and $\tilde{\lambda} = \lambda/\|PA^{-1}\|_2^2$ is the scaled penalty parameter.

Proof. We expand the gradient of the Lagrangian at the stationary point as

$$\begin{pmatrix} \mathcal{L}_{\mathbf{m}} \left(\mathbf{m}_{\lambda}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*} \right) \\ \mathcal{L}_{\mathbf{u}} \left(\mathbf{m}_{\lambda}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*} \right) \\ \mathcal{L}_{\mathbf{v}} \left(\mathbf{m}_{\lambda}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*} \right) \end{pmatrix} \approx \begin{pmatrix} R \left(\mathbf{m}^{*}, \mathbf{u}^{*}, \mathbf{v}^{*} \right) & K \left(\mathbf{m}^{*}, \mathbf{v}^{*} \right)^{T} & G \left(\mathbf{m}^{*}, \mathbf{u}^{*} \right)^{T} \\ K \left(\mathbf{m}^{*}, \mathbf{v}^{*} \right) & P^{T}P & A \left(\mathbf{m}^{*} \right)^{T} \\ G \left(\mathbf{m}^{*}, \mathbf{u}^{*} \right) & A \left(\mathbf{m}^{*} \right) & 0 \end{pmatrix} \begin{pmatrix} \mathbf{m}_{\lambda}^{*} - \mathbf{m}^{*} \\ \mathbf{u}_{\lambda}^{*} - \mathbf{u}^{*} \\ \mathbf{v}_{\lambda}^{*} - \mathbf{v}^{*} \end{pmatrix}.$$
(32)

Using the expressions for the gradient of the Lagrangian at $(\mathbf{m}_{\lambda}^*, \mathbf{u}_{\lambda}^*, \mathbf{v}_{\lambda}^*)$ obtained in theorem 4.2 we can obtain an expression for $\mathbf{m}_{\lambda}^* - \mathbf{m}^*$ etc by solving

$$\begin{pmatrix} R & K^{T} & G^{T} \\ K & P^{T}P & A^{T} \\ G & A & 0 \end{pmatrix} \begin{pmatrix} \mathbf{m}_{\lambda}^{*} - \mathbf{m}^{*} \\ \mathbf{u}_{\lambda}^{*} - \mathbf{u}^{*} \\ \mathbf{v}_{\lambda}^{*} - \mathbf{v}^{*} \end{pmatrix} = \begin{pmatrix} \mathcal{L}_{\mathbf{m}} \left(\mathbf{m}_{\lambda}^{*}, \mathbf{u}_{\lambda}^{*}, \mathbf{v}_{\lambda}^{*} \right) \\ 0 \\ \lambda^{-1} \mathbf{v}^{*} + \mathcal{O} (\lambda^{-2}) \end{pmatrix}.$$
(33)

Eliminating the bottom two rows we find

$$H(\mathbf{m}_{\lambda}^{*}-\mathbf{m}^{*}) = \mathcal{L}_{\mathbf{m}}(\mathbf{m}_{\lambda}^{*},\mathbf{u}_{\lambda}^{*},\mathbf{v}_{\lambda}^{*}) + \lambda^{-1}F\mathbf{v}^{*} + \mathcal{O}(\lambda^{-2}), \qquad (34)$$

where

$$H = R - K^{T} A^{-1} G - G^{T} A^{-T} K + G^{T} A^{-T} P^{T} P A^{-1} G$$

is the reduced Hessian and

$$F = G^T A^{-T} P^T P A^{-1} - K^T A^{-1}.$$

Expressing v^* as

$$\mathbf{v} = A^{-T} P^{T} \Big(\mathbf{d} - P A \Big(\mathbf{m}_{\lambda}^{*} \Big)^{-1} \mathbf{q} \Big),$$

and ignoring the second order terms in H and F, i.e

$$H \approx G^T A^{-T} P^T P A^{-1} G,$$

and

$$F \approx G^T A^{-T} P^T P A^{-1},$$

we find

$$\|\mathbf{m}_{\lambda}^{*}-\mathbf{m}^{*}\|_{2} \leq \left|\left|H^{-1}\right|\right|_{2}\left(\epsilon+\widetilde{\lambda}^{-1}\|G^{T}A^{-T}P^{T}\|_{2}\eta\right),$$

where $\eta = \|\mathbf{d} - PA(\mathbf{m}_{\lambda}^*)^{-1}\mathbf{q}\|_2$ denotes the data-error. Introducing the scaled tolerance and data-error we express this as

$$\|\mathbf{m}_{\lambda}^{*}-\mathbf{m}^{*}\|_{2} \leq \kappa(H) \Big(\tilde{\epsilon} + \tilde{\lambda}^{-1} \tilde{\eta} \Big).$$

As noted before, we can only hope to reconstruct a few components of \mathbf{m}^* , given by the dominant eigenvectors of H. Thus, even for an acceptable reconstruction, the error $\|\mathbf{m}_{\lambda}^* - \mathbf{m}^*\|_2$ can be relatively large. On the other hand, the data-error for the corresponding \mathbf{m}_{λ}^* can be quite small, even if $\|\mathbf{m}_{\lambda}^* - \mathbf{m}^*\|_2$ is large. We expect the penalty method to yield an acceptable reconstruction when $\hat{\lambda}^{-1}\tilde{\eta}$ is small compared to $\tilde{\epsilon}$.

5. Algorithm

In this section, we discuss some practicalities of the implementation of algorithm 3. We slightly elaborate the notation to explicitly reveal the multi-experiment structure of the problem. In this case, the data are acquired in a series of *K* independent experiments and $\mathbf{d} = [\mathbf{d}_1; ...; \mathbf{d}_K]$ is a block-vector. We partition the states and sources in a similar manner. Since the experiments are independent, the system matrix *A* is block-diagonal matrix with *K* blocks $A_i(\mathbf{m})$ of size $N \times N$. Similarly, the matrix *P* consists of blocks P_i . Recall that we collect *L* independent measurements for each experiment, so the matrices $P_i \in \mathbb{R}^{N \times L}$ have full rank.

5.1. Solving the augmented PDE

Due to the block structure of the problem, the linear systems can be solved independently. We can obtain the state \mathbf{u}_i by solving the following inconsistent overdetermined system

$$\begin{pmatrix} A_i(\mathbf{m})\\ \lambda^{-1/2}P_i \end{pmatrix} \mathbf{u}_i \approx \begin{pmatrix} \mathbf{q}_i\\ \lambda^{-1/2}\mathbf{d}_i \end{pmatrix},\tag{35}$$

in a least-squares sense. Assuming that all the blocks A_i and P_i are identical, we will drop the subscript i for the remainder of this subsection. Next, we will discuss various approaches to solving the overdetermined system (35).

Factorization: If both A and P are sparse, we can efficiently solve the system via a QR factorization or via a Cholesky factorization of the corresponding Normal equations. In many applications, $P^T P$ is a (nearly) diagonal matrix and thus the augmented system $A^T A + \lambda^{-1} P^T P$ has a similar sparsity pattern as the original system. Thus, the fill-in will not be worse than when factorizing the original system.

Iterative methods: While we can make use of factorization techniques for small-scale applications, industry-scale applications will typically require (preconditioned) iterative methods. Obviously, we can apply any preconditioned iterative method that is suitable for solving least-squares problems, such as LSQR, LSMR or CGLS [21–23]. Another promising candidate is a generic accelerated row-projected method described by [24, 25] which proved useful for solving PDEs and can be easily extended to deal with overdetermined systems [26].

To get an idea of how such iterative methods will perform we explore some of the properties of the augmented system. The augmented system $A^{T}A + \lambda^{-1}P^{T}P$ is a rank L modification of the original system $A^{T}A$. It follows from [27, theorem 8.1.8] that the eigenvalues are related as

$$\mu_n (A^T A + \lambda^{-1} P^T P) = \mu_n (A^T A) + a_n \lambda^{-1}, n = 1, 2, ..., N,$$
(36)

where $\mu_1(B) > \mu_2(B) > ... > \mu_N(B)$ denote the eigenvalues of B and the coefficients a_n satisfy $\sum_{n=1}^{N} a_n = L$. This means that at worst, 1 eigenvalue is shifted by $L\lambda^{-1}$ while at best, all the eigenvalues are shifted by $LN^{-1} \lambda^{-1}$. For the condition numbers $\kappa(B) = \mu_1(B)/\mu_N(b)$ we find

$$C_N^{-1}\kappa(A^T A) \leqslant \kappa(A^T A + \lambda^{-1} P^T P) \leqslant C_1\kappa(A^T A),$$
(37)

where $C_i = \left(1 + \frac{L}{\lambda \mu_i(A^T A)}\right)$. To illustrate this, we show a few examples for a 1D (time-harmonic) parabolic PDE $(\iota \omega - \partial_x^2)u = 0$ and the 1D Helmholtz equation $(\omega^2 + \partial_x^2)u = 0$, both with Neumann boundary conditions. Both are discretized using first-order finite-differences on $x \in [0, 1]$ with N = 51 points for $\omega = 10\pi$. The sampling matrix P consists of L rows of the identity matrix (regularly sampled). The ratio of the condition numbers of $A^{T}A$ and $A^{T}A + \lambda^{-1}P^{T}P$ for the parabolic and Helmholtz equation are shown in tables 2 and 3. For these examples, the condition number of the augmented system is actually lower than that of the original system. The eigenvalues are shown in figures 1 and 2. These show that the actual eigenvalues distributions do not change significantly. We expect that iterative methods will perform similarly on the augmented system as they would on the original system. How to effectively precondition the augmented system given a good preconditioner for the original system is a different matter which is outside the scope of this paper.

Direct methods: When the matrix has additional structure we might actually prefer a direct method over an iterative method. An example is explicit time-stepping, where the

Table 2. Ratio of the condition numbers of $A^TA + \lambda P_L^TP_L$ and A^TA for various λ and L, where A is a finite-difference discretization of $\iota(10\pi) - \partial_x^2$ on $x \in [0, 1]$ and P_L is a restricted identity matrix of rank L.

	$\lambda = 0.1$	$\lambda = 1$	$\lambda = 10$
L = 1 $L = 10$ $L = 20$	9.83e-01	9.84e-01	9.89e-01
	9.68e-02	5.07e-01	9.11e-01
	9.19e-02	5.01e-01	9.09e-01

system matrix A exhibits a lower-triangular block-structure. In this case the action of A^{-1} can be computed efficiently via forward substitution, requiring storage of only a few time-slices of the state. The adjoint system A^{-T} can be solved by backward substitution, however, the full time-history of the state variable is needed to compute the gradient [28]. For the penalty method, the augmented system $A^{T}A + \lambda^{-1}P^{T}P$ will have a banded structure and the system can be solved using a block version of the Thomas algorithm, which again would require storage of the full time-history. So even in this setting it seems possible to apply the penalty method at roughly the same per-iteration complexity as the reduced method.

5.2. Gradient and Hessian computation

Given these solutions \mathbf{u}_k of (35), the gradient, \mathbf{g}_{λ} and GN Hessian H_{λ} of ϕ_{λ} are given by (see equations (21) and (22))

$$\mathbf{g}_{\lambda} = \lambda \sum_{k=1}^{K} G_{k}^{T} (A_{k} \mathbf{u}_{k} - \mathbf{q}_{k}),$$
(38)

$$H_{\lambda} = \lambda \sum_{k=1}^{K} G_{k}^{T} \Big(I - A_{k} \Big(A_{k}^{T} A_{k} + \lambda^{-1} P_{k}^{T} P_{k} \Big)^{-1} A_{k}^{T} \Big) G_{k},$$
(39)

where $G_k = G(\mathbf{m}, \mathbf{u}_k)$. We can compute the inverse of $(A_k^T A_k + \lambda^{-1} P_k^T P_k)$ in the same way as used when solving for the states. In practice, we would solve for one state at a time and aggregate the results on the fly. Moreover, the GN Hessian admits a natural sparse approximation $H_\lambda \approx \lambda \sum_{k=1}^{K} G_k^T G_k$ which has proved to work well in practice [29].

5.3. Choosing λ

An important aspect of the proposed method is the choice of the penalty parameter, λ . Theorem 4.2 essentially states that we can expect to find a stationary point of the Lagrangian within finite accuracy with finite λ as long as $\lambda^{-1} || \mathcal{L}_{\mathbf{v}} ||_2$ is small compared to $|| \mathcal{L}_{\mathbf{m}} ||_2$. We can easily keep track of this quantity—at no significant additional computational cost—during the iterations and increase λ as needed. Initialization can be done by directly enforcing $\lambda^{-1} || \mathcal{L}_{\mathbf{v}} ||_2$ to be some fraction of $|| \mathcal{L}_{\mathbf{m}} ||_2$ at the initial iterate. A natural scaling for λ is suggested by lemma 4.1; $\lambda = \tilde{\lambda} || PA^{-1} ||_2^2$, where $\tilde{\lambda} > 1$ is considered large and $\tilde{\lambda} < 1$ is considered small.

5.4. Complexity estimates

A summary of the leading order computational costs *per iteration* of the penalty, reduced and all-at-once approaches is given in table 4. The storage and computation required for the reduced and penalty methods are of the same order in terms of K, N and M. The PDE-solves in both the penalty and reduced approaches can be done independently and in parallel. This

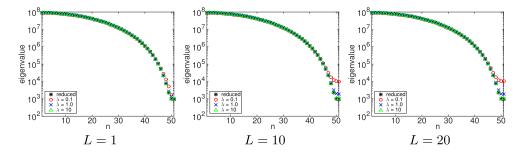


Figure 1. Eigenvalues of the augmented system, $A^T A + \lambda P_L^T P_L$, for various λ and L, where A is a finite-difference discretization of $\iota(10\pi) - \partial_x^2$ on $x \in [0, 1]$ and P_L is a restricted identity matrix of rank L. For comparison, the eigenvalues of the original system $A^T A$ are also shown.

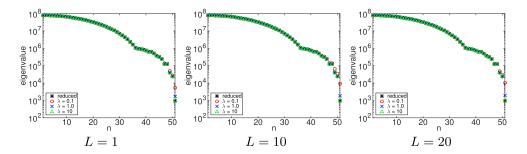


Figure 2. Eigenvalues of the augmented system, $A^TA + \lambda P_L^T P_L$, for various λ and L, where A is a finite-difference discretization of $(10\pi)^2 m + \partial_x^2$ on $x \in [0, 1]$ and P_L is a restricted identity matrix of rank L. For comparison, the eigenvalues of the original system A^TA are also shown.

Table 3. Ratio of the condition numbers of $A^TA + \lambda P_L^TP_L$ and A^TA for various λ and L, where A is a finite-difference discretization of $(10\pi)^2 m + \partial_x^2$ on $x \in [0, 1]$ and P_L is a restricted identity matrix of rank L.

	$\lambda = 0.1$	$\lambda = 1$	$\lambda = 10$
L = 1 $L = 10$ $L = 20$	1.80e-01	5.44e-01	9.17e-01
	1.03e-01	5.06e-01	9.10e-01
	9.10e-02	5.00e-01	9.09e-01

makes these approaches more attractive for large-scale problems from a computational pointof-view.

We have argued in section 5.1 that it is plausible that the augmented system can be solved as efficiently as the original PDE. However, it is not clear how the penalty and reduced methods will compare in the required number of iterations, though we expect that for small λ the optimization problem is less nonlinear and hence easier to solve. To reach a given tolerance with the penalty method, however, we need a continuation strategy in λ , adding to the cost of the penalty method. In the next section we compare the actual computational costs on a few test-cases.

Table 4. Leading order computational and storage costs per iteration of different methods; K denotes the number of experiments and N denotes the number of gridpoints and M denotes the number of parameters.

	# PDE's	Storage	Gauss-Newton update
Penalty	K	N + M	Solve matrix-free linear system in <i>M</i> unknowns, requires <i>K</i> (overdetermined) PDE solves per mat-vec
Reduced	2 <i>K</i>	2N + M	Solve matrix-free linear system in <i>M</i> unknowns, requires 2 <i>K</i> PDE solves per mat-vec
All-at-once	0	2KN + M	Solve sparse symmetric, possibly indefinite system in $(2KN + M) \times (2KN + M)$ unknowns

6. Case studies

The following experiments are done in Matlab, using direct factorization to solve the PDEs (with Matlab slash). We consider both a GN and a quasi-Newton (QN) variant of the algorithms and use a weak Wolfe linesearch to determine the steplength. In the GN method the Hessian is inverted using conjugate gradients (pcg) up to a relative tolerance of δ . The matrix-vector products are computed on the fly. For the QN method we use the L-BFGS inverse Hessian with a history size of 5 [14]. We measure the cost of the inversion by counting the number of PDE solves as outlined in table 4. In all experiments, we set λ relative to the largest eigenvalue of $A^{-T}P^TPA^{-1}$ at the initial iterate. This scaling is justified by lemma 4.1. We compare the results for various fixed values of λ and additionally show the results for an ad hoc continuation strategy; performing a few iterations for increasing values of λ using the result for each λ as initial guess for the next.

To avoid the inverse crime, we compute the data for the ground truth model on a finer grid than used for the inversion.

In these experiments, we illustrate that the penalty method:

- converges to a stationary point of the Lagrangian within the predicted tolerance of $\mathcal{O}(\lambda^{-1})$;
- can give practically the same or better results as the reduced method at a lower computational cost;
- is not overly sensitive to noise;
- is less sensitive to the initial guess than the conventional approach.

The Matlab code used to perform the experiments is available from https://github.com/ tleeuwen/Penalty-Method.

6.1. 1D DC resistivity

We consider the PDE

$$\partial_t u(t, x) = \partial_x \big(m(x) \partial_x \big) u(t, x), \tag{40}$$

on the domain $x \in [0, 1]$ with Neumann boundary conditions. A finite-difference discretization in the temporal Fourier domain gives

$$A(\mathbf{m}) = \iota \omega \operatorname{diag}(\mathbf{w}) + D^T \operatorname{diag}(\mathbf{m})D, \tag{41}$$

where ω is the angular frequency, $\mathbf{w} = \left[\frac{1}{2}; 1, ..., 1; \frac{1}{2}\right]$, **m** represents the medium parameter in the cell-centers and *D* is the $N - 1 \times N$ finite-difference matrix

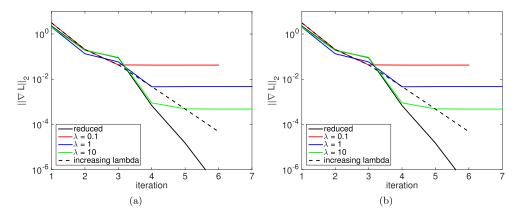


Figure 3. Convergence history and reconstructions for 1D resistivity problem. Even though the penalty method does not converge to same tolerance as the reduced method in terms of the gradient of the Lagrangian, the resulting parameter estimates are almost the same.

$$D = \frac{1}{h} \begin{pmatrix} -1 & 1 & & \\ & -1 & 1 & \\ & & \ddots & \ddots \\ & & & -1 & 1 \end{pmatrix},$$

with h = 1/(N - 1). The Jacobian is given by

$$G(\mathbf{m}, \mathbf{u}) = D^T \operatorname{diag}(D\mathbf{u}).$$

The ground-truth model is $m(x) = 1 + e^{-10(x-1/2)^2}$ and we locate two sources and receivers on either end of the domain. The data are generated on a grid with N = 201 points and we have K = L = 2.

For the inversion we use N = 101 points. We use a GN method with $\epsilon = 10^{-9}$, $\delta = 10^{-3}$ and include a regularization term $\frac{\alpha}{2} \|D\mathbf{m}\|_2^2$ with $\alpha = 10^{-6}$. The initial parameters are $\mathbf{m}^0 = \mathbf{1}$.

The results are shown in figure 3. The convergence plot, figure 3(a), shows the predicted behavior of the penalty method; the norm of the gradient of the Laplacian stalls at $\mathcal{O}(\lambda^{-1})$. The convergence of the continuation strategy shows that it is possible to reach the desired tolerance by gradually increasing λ (using $\lambda = \{0.1, 1, 10, 100\}$ with a few iterations each). The resulting parameter estimates are very similar as can be seen in figure 3(b). The actual costs of the inversion are listed in table 5. The computational cost for the various approaches are of the same order of magnitude, except for $\lambda = 10$, where more than twice as many iterations are required.

6.2. 2D Acoustic tomography

Consider the 2D scalar wave-equation

$$m(x)\partial_t^2 u(t,x) = \nabla^2 u(t,x), \tag{42}$$

on $x \in \Omega \subseteq \mathbb{R}^2$ with radiation boundary conditions $\sqrt{m}(x)\partial_t u(t, x) - n(x) \cdot \nabla u(t, x) = 0$ on $\partial \Omega$ where n(x) is the outward normal vector.

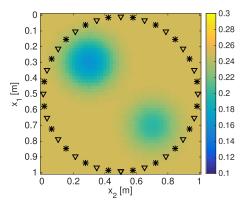


Figure 4. Ground truth model (s^2/km^2) and locations of the sources (*) and receivers (∇)

Table 5. Costs of the 1D dc resistivity inversion.

	Reduced	$\lambda = 0.1$	$\lambda = 1$	$\lambda = 10$	Increasing λ
Iterations	7	5	6	7	6
PDE solves	496	222	223	280	292

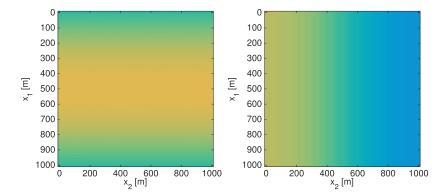


Figure 5. Perturbations \mathbf{v}_1 and \mathbf{v}_2 used to plot the misfit $\phi(\mathbf{m}^* + a\mathbf{1}\mathbf{v}_1 + a_2\mathbf{v}_2)$ and $\phi_{\lambda}(\mathbf{m}^* + a_1\mathbf{v}_1 + a_2\mathbf{v}_2)$ in figure 6.

Discretization in the temporal Fourier domain leads to a scalar Helmholtz equation

$$A(\mathbf{m}) = \operatorname{diag}(\mathbf{s}) - D^T D, \tag{43}$$

where $D = [I_2 \otimes D_1; D_2 \otimes I_1]$ with D_i the $(N_i - 1) \times N_i$ finite-difference matrix, I_i the $N_i \times N_i$ identity matrix and $s_i = \omega^2 m_i$ in the interior and $s_i = \omega^2 m_i/2 + \iota \omega \sqrt{m_i}/h$ on the boundary. The Jacobian is given by

$$G(\mathbf{m}, \mathbf{u}) = \operatorname{diag}(\mathbf{s}')\operatorname{diag}(\mathbf{u}), \tag{44}$$

where $s'_i = \omega^2$ in the interior and $s'_i = (\omega^2 + \iota\omega/\sqrt{m_i})/2$ on the boundary.

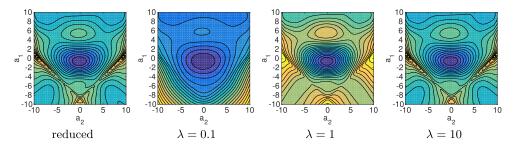


Figure 6. Misfit in the direction of the perturbations shown in figure 5. For small λ , the reduced penalty objective ϕ_{λ} is less nonlinear than the reduced objective ϕ .

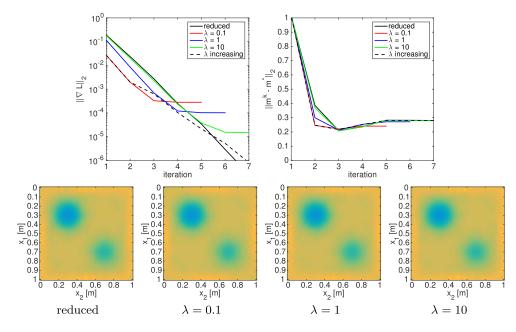


Figure 7. Convergence history, reconstruction error and reconstructions for data without noise using a GN method. Even though the penalty method does not converge to same tolerance as the reduced method in terms of the gradient of the Lagrangian, the resulting parameter estimates are almost the same.

The observation matrix P samples the solution at the receiver locations using 2D linear interpolation while the point sources are defined using adjoint 2D linear interpolation.

6.2.1. Ultrasound tomography. The domain $\Omega = [0, 1] \times [0, 1]$ m is discretized using $N_1 \times N_2$ points. The ground-truth \mathbf{m}^* as well as the source and receiver locations are shown in figure 4. We use a single frequency of 5 kHz (i.e., $\omega = 10^4 \pi$). The data for the ground-truth model are generated using $N_1 = N_2 = 101$ while the following experiments are done with $N_1 = N_2 = 51$.

Nonlinearity: First, we investigate the sensitivity of the misfit functions ϕ and ϕ_{λ} by plotting ϕ (**m**^{*} + a_1 **v**₁ + a_2 **v**₂) and ϕ_{λ} (**m**^{*} + a_1 **v**₁ + a_2 **v**₂) as a function of (a_1 , a_2). We take **v**₁, **v**₂ to be slowly oscillatory modes as shown in figure 5. The misfit as a function of (a_1 , a_2)

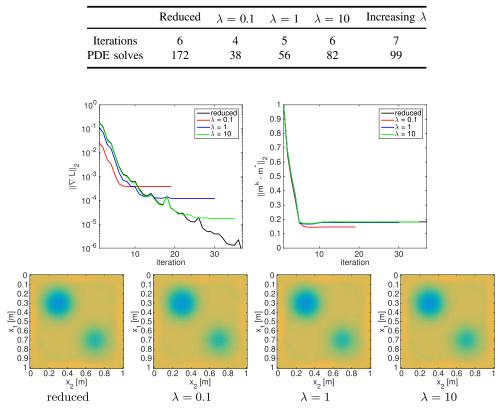
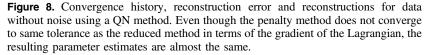


Table 6. Costs of the 2D ultrasound inversion with a GN method.



is shown in figure 6. We see a radically different behavior for the reduced and penalty methods.

The first exhibits strong nonlinearity and some spurious stationary points while for small λ the penalty misfit is much better behaved. For larger values λ the penalty misfit starts to behave more like the reduced misfit as expected.

Inversion: For the inversion, we include a regularization term $\frac{\alpha}{2} \|D\mathbf{m}\|_2^2$ with $\alpha = 2$ and compare the GN method ($\epsilon = 10^{-6}$, $\delta = 10^{-1}$) to the QN method ($\epsilon = 10^{-6}$). The initial parameter \mathbf{m}_0 is constant at $\frac{1}{4}s^2/m^2$.

The results for the GN method are shown in figure 7. The convergence history, figure 7 (top, left), shows the predicted behavior of the penalty method; the norm of the gradient of the Lagrangian stalls at $\mathcal{O}(\lambda^{-1})$ when using the penalty method. The convergence history of the continuation strategy shows that is possible to reach the desired tolerance by gradually increasing λ (using $\lambda = \{0.1, 1, 10, 100, 1000\}$ with a few iterations each). Figure 7 (top, right) shows that the methods perform similarly in terms of reconstruction error. The resulting parameter estimates are very similar as can be seen in figure 7 (bottom). The actual costs of the inversion are listed in table 6. The penalty method converges to the same error $||\mathbf{m}^k - \mathbf{m}^*||$ in less iterations and uses less PDE solves. Note that all methods start overfitting after a few

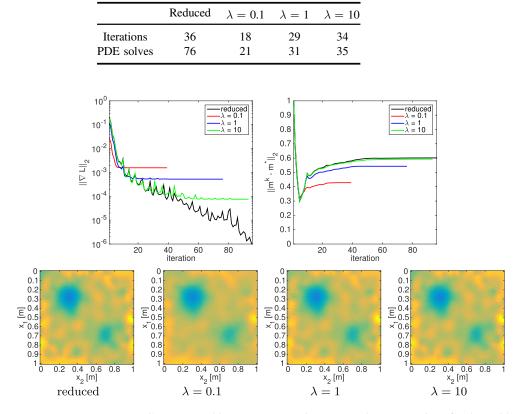


Table 7. Costs of the 2D ultrasound inversion with a QN method.

Figure 9. Convergence history, reconstruction error and reconstructions for data with 10% Gaussian noise using a QN method. Even though the penalty method does not converge to same tolerance as the reduced method in terms of the gradient of the Lagrangian, the resulting parameter estimates are almost the same. In fact, for small λ , result is even a little better.

iterations. This can be countered by including more appropriate regularization or stopping the iterations early. The point here is to show that the penalty method gives similar results as the reduced method.

The results for the QN method are shown in figure 8. The convergence history shows the same behavior as the previous experiment. The costs of the inversion, shown in table 7, are slightly less than those of the GN method. As with the GN method, the penalty method converges in less iterations and uses less PDE solves per iterations.

Sensitivity to noise: Results for the QN method on data with 10% Gaussian noise are shown in figure 9. Figure 10 shows the results on data with 20% Gaussian noise. These results show that the penalty approach is not overly sensitive to noise and gives very similar—even slightly better—results compared to the reduced approach.

6.2.2. Seismic tomography. Here, the domain $\Omega = [0, 5] \times [0, 20]$ km is discretized using $N_1 \times N_2$ points. The ground-truth \mathbf{m}^* as well as the source and receiver locations are shown in figure 11. We use a frequency of 2 Hz (i.e., $\omega = 4\pi$). The data for the ground-truth are

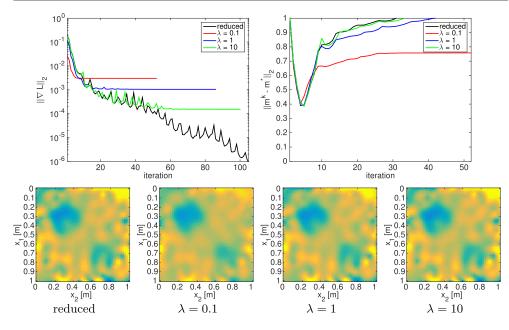


Figure 10. Convergence history, reconstruction error and reconstructions for data with 20% Gaussian noise using a QN method. Even though the penalty method does not converge to same tolerance as the reduced method in terms of the gradient of the Lagrangian, the resulting parameter estimates are almost the same. In fact, for small λ , result is even a little better.

generated using $N_1 = 101$, $N_2 = 401$ while the following experiments are done with $N_1 = 51$, $N_2 = 201$.

Sensitivity to the initial guess: For the inversion, we include a regularization term $\frac{\alpha}{2} \|D\mathbf{m}\|_2^2$ with $\alpha = 5$ and use the QN method ($\epsilon = 10^{-6}$). We will use two different initial guesses, I and II, depicted in figure 11, and see whether the methods converge to the same final iterate. Initial iterate I is much closer to the ground truth than the initial iterate II. This can also be observed when looking at the data produced by these iterates. The first initial iterate produces data that differs only slightly from the observed data and inversion is considered to be easy. The second initial iterate produces data that is shifted significantly with respect to the observed data and inversion is considered to be difficult. It should be noted, however, that a significant source of the error in the initial model is the region near the surface (z = 0). In practical applications, such large errors in this region of the model might not occur.

The results for initial guess I (figure 11, middle) are shown in figure 12. We see that both the reduced and penalty methods converge to roughly the same final iterate and are able to fit the data equally well. Starting from initial guess II (figure 11, bottom), however, we see in figure 13 that the reduced and penalty methods converge to different final iterates. For small λ , however, the penalty method converges to roughly the same iterate as when starting from a better initial guess. Looking at the data-fit, we observe that the penalty method for small λ is still able to fit the data perfectly while the reduced method is not. In seismology, this phenomenon is called *cycle-skipping*.

Figure 14 shows the convergence of the methods in terms of the data misfit $||P\mathbf{u} - \mathbf{d}||_2$ and the distance to the constraint $||A(\mathbf{m})\mathbf{u} - \mathbf{q}||_2$. We observe that, when starting from initial guess I, both the penalty and reduced methods converge to approximately the same point. For

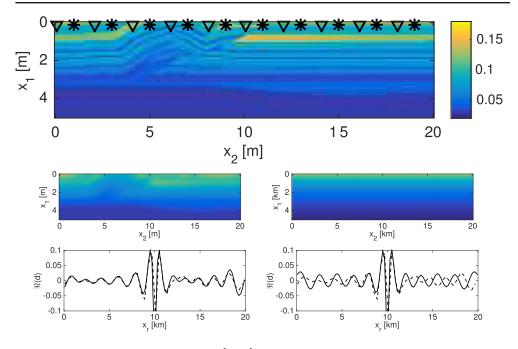


Figure 11. Ground truth (s^2/km^2) (top) with locations of the sources (*) and receivers (∇) and initial iterates I (middle, left) and II (bottom, right). The bottom row shows the data for a source in the center for the ground truth (dashed line) as well as the data for the two initial iterates. The first initial iterate produces data that differs only slightly from the observed data and inversion is considered to be easy. The second initial iterate produces data that is shifted significantly with respect to the observed data and inversion is considered to be difficult.

initial guess II the penalty method for $\lambda = 0.1$ and $\lambda = 1$ needs a few more iterations, but still converges to the same point as for initial guess I. For $\lambda = 10$ and the reduced method, however, the iterations stall at a relatively high data misfit.

These experiments suggests that the penalty method indeed mitigates some of the nonlinearity of the problem, allowing the optimization to converge to the same final iterate, even when the initial guess is further away from the ground truth.

7. Discussion

This paper lays out the basics of an efficient implementation of the penalty method for PDEconstrained optimization problems arising in inverse problems. While the initial results are promising, some aspects of the proposed method warrant further investigation.

Even though the theoretical results suggest that the penalty approach can find a stationary point of the Lagrangian with finite precision with a finite λ , it is not clear how to choose a suitable value for λ a priori. Our analysis and results suggest that choosing λ to be a small fraction of $||PA^{-1}||_2^2$ at the initial iterate yields good results. A continuation strategy for λ is needed if we want to guarantee finding a stationary point of the Lagrangian with preset tolerance. A natural way to do this seems to be detecting when the penalty method stalls and subsequently reducing λ . The numerical results suggest that such an approach is viable, but further study is needed in order to develop a robust continuation strategy.

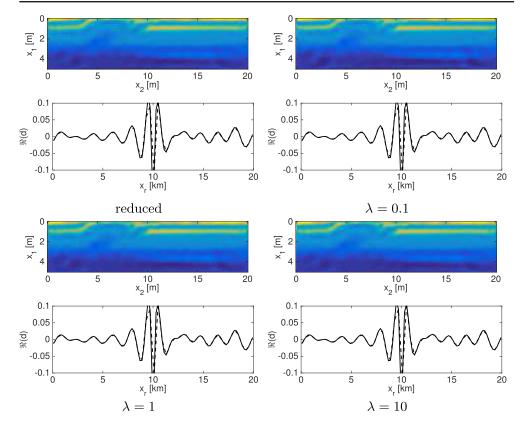


Figure 12. QN reconstructions after 50 iterations and corresponding data for a source in the center, starting from the initial iterate I. Both the penalty and reduced methods converge to the same final iterate when starting from this initial guess and are able to fit the data equally well.

The penalty formulation essentially relaxes the constraints and therefore allows for errors in the physics as well as the data. As a result, the penalty formulation leads to reduced sensitivity of the final reconstruction to the initial guess. Further investigation is needed to characterize this robustness.

Finally, the Hessian of the penalty objective exhibits additional structure that could potentially be exploited. In particular, the penalty-method GN Hessian is full rank and allows for a natural sparse approximation $H_{\lambda} \approx \lambda G^T G$ (see equation (18)). The reduced GN Hessian, on the other hand, has rank of at most *ML* and does not permit such a natural sparse approximation.

8. Conclusions

We have presented a penalty method for PDE-constrained optimization with linear PDEs with applications to inverse problems. The method is based on a quadratic penalty formulation of the constrained problem. This reformulation results in a an unconstrained optimization problem in both the parameters and the state variables. To avoid having to store and update the state variables as part of the optimization, we explicitly eliminate the state variables by

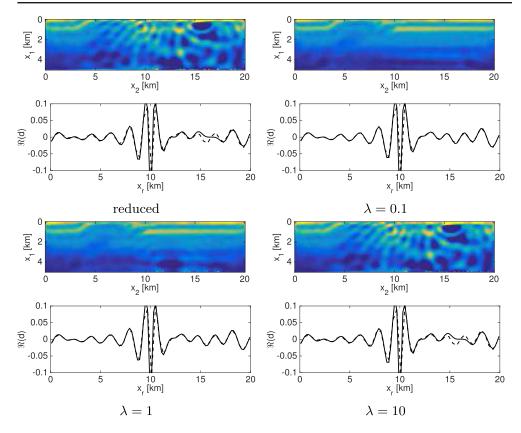


Figure 13. QN reconstructions after 50 iterations and corresponding data for a source in the center, starting from the initial iterate II. For small λ , the penalty method converges to the same final iterate as when starting from initial guess II, showing stability against changes in the initial guess. The reduced method converges to a completely different model, suggesting that the optimization method is stuck in a local minimum. This is confirmed when looking at the data-fit.

solving an overdetermined linear system. The proposed method combines features from both the *all-at-once* approach, in which the states and parameters are updated simultaneously, and the conventional *reduced* approach, in which the PDE-constraints are eliminated explicitly. While having a similar computational complexity as the conventional reduced approach, the penalty approach explores a larger search space by not satisfying the PDE-constraints exactly.

We show that we can (theoretically) find a stationary point of the Lagrangian of the constrained problem within a given tolerance as long as the penalty parameter, λ , is chosen large enough. While theoretically we need $\lambda \uparrow \infty$, we can suffice with solving the problem for a finite λ to reach the stationary point within finite precision.

The main algorithmic difference with the conventional reduced approach is the way the states are eliminated from the problem. Instead of solving the PDEs, we formulate an overdetermined system of equations that consists of the discretized PDE and the measurements. We discuss the properties of this augmented system and show with a few numerical examples that both the structure of the system as well as the eigenvalues are not altered dramatically as compared the original PDE. Thus, it is plausible that the augmented system can be solved as efficiently using the same approach as is used for the original PDE.

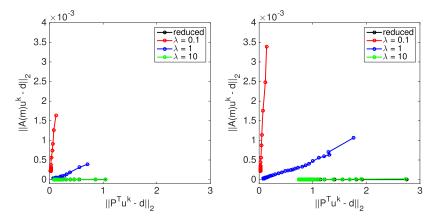


Figure 14. Convergence history in terms of the data-fit and distance to the constraint, starting from initial iterate I (left) and starting from initial iterate II (right). These plots show that, for small λ , the penalty method is able to reduce both the data and PDE misfit to the same level when starting from either initial guess. The reduced method, however, cannot reduce the data misfit to the same level when starting from initial guess II, suggesting that it got stuck in a local minimum.

The numerical examples show that very good results can be obtained by using even a single, relatively small, value of λ . An ad hoc continuation strategy further shows that it is viable to gradually increase λ in order to reach the desired tolerance.

The numerical examples further show that when using the penalty formulation, the optimization problem may actually be less nonlinear and that in some cases a better parameter reconstruction is obtained as compared to the conventional reduced approach. In particular, the results show that the penalty method is not overly sensitive to noise and less sensitive to the initial model than the conventional reduced approach.

Thus, the proposed approach is a viable alternative to the conventional reduced approach for solving inverse problems with PDE-constraints.

Acknowledgments

This work was in part financially supported by the Natural Sciences and Engineering Research Council of Canada via the Collaborative Research and Development Grant DNOISEII (375142-08) and the Netherlands Organisation of Scientific Research (NWO) as part of research programme 613.009.032. This research was carried out as part of the SIN-BAD II project which is supported by the following organizations: BG Group, BGP, CGG, Chevron, ConocoPhillips, DownUnder GeoSolutions, Hess, Petrobras, PGS, Schlumberger, Sub Salt Solutions and Woodside.

References

- Tarantola A and Valette A 1982 Generalized nonlinear inverse problems solved using the least squares criterion *Rev. Geophys. Space Phys.* 20 129–232
- [2] Haber E, Ascher U M and Oldenburg D W 2004 Inversion of 3D electromagnetic data in frequency and time domain using an inexact all-at-once approach *Geophysics* 69 1216
- [3] Epanomeritakis I, Akçelik V, Ghattas O and Bielak J 2008 A Newton-CG method for large-scale three-dimensional elastic full-waveform seismic inversion *Inverse Problems* 24 034015

- [4] van Leeuwen T and Herrmann F J 2014 3D frequency-domain seismic inversion with controlled sloppiness SIAM J. Sci. Comput. 36 S192–217
- [5] Abdoulaev G S, Ren K and Hielscher A H 2005 Optical tomography as a PDE-constrained optimization problem *Inverse Problems* 21 1507–30
- [6] Wang K, Matthews T, Anis F, Li C, Duric N and Anastasio M A 2015 Breast ultrasound computed tomography using waveform inversion with source encoding *Proc. SPIE* 94190C
- [7] Haber E, Ascher U M and Oldenburg D 2000 On optimization techniques for solving nonlinear inverse problems *Inverse Problems* 16 1263–80
- [8] Grote M J, Huber J, Kourounis D and Schenk O 2014 Inexact interior-point method for PDEconstrained nonlinear optimization SIAM J. Sci. Comput. 36 A1251–76
- [9] Dennis J E, Heinkenschloss M and Vicente L N 1998 Trust-region interior-point SQP algorithms for a class of nonlinear programming problems SIAM J. Control Optim. 36 1750–94
- [10] Heinkenschloss M and Ridzal D 2008 An inexact trust-region SQP method with applications to PDE-constrained optimization *Numer. Math. Adv. Appl.* 613–20
- [11] Richter R G 1981 Numerical identification of a spatially varying diffusion coefficient Math. Comput. 36 375–86
- [12] Banerjee B, Walsh T F, Aquino W and Bonnet M 2013 Large scale parameter estimation problems in frequency-domain elastodynamics using an error in constitutive equation functional *Comput. Methods Appl. Mech. Eng.* 253 60–72
- [13] van Leeuwen T and Herrmann F J 2013 Mitigating local minima in full-waveform inversion by expanding the search space *Geophys. J. Int.* 195 661–7
- [14] Nocedal J and Wright S J 2006 Numerical optimization Springer Series in Operations Research (Berlin: Springer)
- [15] Haber E and Ascher U M 2001 Preconditioned all-at-once methods for large, sparse parameter estimation problems *Inverse Problems* 17 1847–64
- [16] Biros G and Ghattas O 2003 Inexactness issues in the Lagrange–Newton–Krylov–Schur method for pde-constrained optimization *Large-Scale PDE-Constrained Optimization* (Berlin: Springer) pp 93–114
- [17] Herzog R 2010 Lectures notes algorithms and preconditioning in pde-constrained optimization Technical Report Chemnitz University of Technology
- [18] Eckstein J 2012 Augmented lagrangian and alternating direction methods for convex optimization: a tutorial and some illustrative computational results *Technical Report* RRR 32-2012 Rutgers University
- [19] Curtis F E, Jiang H and Robinson D P 2014 An adaptive augmented Lagrangian method for largescale constrained optimization *Math. Progam.* 152 201–45
- [20] Aravkin A Y and van Leeuwen T 2012 Estimating nuisance parameters in inverse problems Inverse Problems 28 115016
- [21] Paige C C and Saunders M A 1982 LSQR: an algorithm for sparse linear equations and sparse least squares ACM Trans. Math. Softw. 8 43–71
- [22] Fong D C-L and Saunders M A 2011 LSMR: an iterative algorithm for sparse least-squares problems SIAM J. Sci. Comput. 33 2950–71
- [23] Bru R, Marín J, Mas J and Tma M 2014 Preconditioned iterative methods for solving linear least squares problems SIAM J. Sci. Comput. 36 A2002–22
- [24] Björck Å and Elfving T 1979 Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations BIT 19 145–63
- [25] Gordon R and Gordon D 2013 Robust and highly scalable parallel solution of the Helmholtz equation with large wave numbers J. Comput. Appl. Math. 237 182–96
- [26] Censor Y, Eggermont P P B and Gordon D 1983 Strong underrelaxation in Kaczmarz's method for inconsistent systems *Numer. Math.* 41 83–92
- [27] Golub G H and van Loan C F 1996 Matrix Computations 3rd edn (Baltimore: The Johns Hopkins University Press)
- [28] Rothauge K, Haber E and Ascher U 2015 Numerical computation of the gradient and the action of the Hessian for time-dependent PDE-constrained optimization problems arXiv:1509.03801
- [29] Esser E, Guasch L, van Leeuwen T, Aravkin A Y and Herrmann F J 2015 Automatic salt delineation wavefield reconstruction inversion with convex constraints., SEG technical program expanded abstracts Soc. Explor. Geophys. 1337–43