Computers & Fluids 100 (2014) 291–307

Contents lists available at ScienceDirect

Computers & Fluids

journal homepage: www.elsevier.com/locate/compfluid

Adjoint-accelerated statistical and deterministic inversion of atmospheric contaminant transport

Devina Sanjaya [a], Krzysztof Fidkowski [a,*], Ian Tobasco [b]

[a] Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109, United States
[b] Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, United States

ARTICLE INFO

Article history:
Received 17 August 2012
Received in revised form 25 February 2014
Accepted 17 May 2014
Available online 29 May 2014

Keywords:
Inverse problems
Bayesian inference
Adjoint method
Markov chain Monte Carlo

ABSTRACT

In this work we present and compare deterministic and statistical algorithms for efficiently solving large-scale contaminant source inversion problems. The underlying equations of contaminant transport are assumed linear but unsteady and defined over complex geometries. The algorithms presented are accelerated through discrete adjoint solutions that are pre-computed efficiently in an offline stage, yielding savings in the time-critical online stage of several orders of magnitude in computational time. In the deterministic case, adjoints accelerate the application of the Hessian matrix, while in the statistical case, adjoints are used to directly evaluate samples. To address deterioration of statistical sampling efficiency for anisotropic posteriors, we present an application of a recently developed ensemble Markov chain Monte Carlo method. Results for two- and three-dimensional problems demonstrate the feasibility of statistical inversion for large-scale problems and show the advantage of statistical results over single-point deterministic results.

1. Introduction

Real-time modeling and inversion of contaminant release events is crucial in applications ranging from environmental safety monitoring to homeland security. Such events include unintentional industrial or transport accidents, as well as intentional biological or chemical attacks in urban environments. A successful response to these scenarios requires rapid and accurate identification of the extent of the contamination, inversion of the data to find the source of the release, and prediction of the subsequent path of contamination for evacuation and countermeasures. In this work we address the problem of inversion of the contaminant source for large-scale calculations under some simplifying assumptions.

The contaminant source inversion problem is inherently complex: the geometry is often intricate, the flow conditions are uncertain, and the available measurements are limited and noisy. Computationally, the evolution of the contaminant spread is characterized by a system of partial differential equations that must be discretized on complex geometries, resulting in millions of unknowns. Moreover, the problem is generally ill-conditioned in the sense that small changes in the outputs can cause large changes in the calculated inputs [1]. This ill-conditioning makes single-point deterministic calculations, ones that seek the "best" possible answer, not robust, where robustness is measured by the level of risk associated with using computation in a broader context such as design or decision making. The lack of robustness is due to the fact that the "best" input may be only one of many different inputs that produce nearly the same outputs, especially when measurement error is taken into account.

A statistical approach to the inverse problem, in which probability distributions instead of single-point estimates are calculated for the inputs, can be more robust. Probability distributions convey much more information than a most likely value, and they allow for informed decisions that make full use of the available data and associated uncertainties. Statistical approaches often characterize the probability distributions through sampling, which requires numerous forward simulations. However, sampling can be expensive, especially when each forward simulation already taxes computational resources. Thus, for large-scale contaminant inversions, sampling-based statistical approaches quickly become prohibitively expensive.

Multiple previous studies have investigated large-scale inverse problems [2,3], including the societally-relevant application of contaminant transport [4–9]. The high computational cost of practical contaminant transport simulations prevents their direct use for inversion calculations during real-time events. Reducing this

* Corresponding author. Tel.: +1 7346157247.
E-mail address: kfid@umich.edu (K. Fidkowski).

http://dx.doi.org/10.1016/j.compfluid.2014.05.021
0045-7930/

cost has motivated research into inversion approaches that employ approximate solution techniques, such as grid coarsening [2,10], reduced-order modeling [9,11], and stochastic expansions [12], but the utility and accuracy of these approaches for generating real-time results in practical, large-scale simulations is yet to be demonstrated.

To some extent, previous studies have also considered applying adjoint solutions, or at least output gradients, to uncertainty quantification algorithms. These include gradient-enhanced response surface construction via least-squares [13] and gradient-enhanced Kriging [14]; although presented in the context of optimization, these ideas extend to propagation of input uncertainties [15]. For statistical inverse problems, the use of output gradients has been proposed to improve the acceptance ratio in Markov chain Monte Carlo sampling of anisotropic posteriors [16]. In recent work, these gradients are computed using finite differences based on the large-scale model [11] and on cheaper coarse-scale models [17].

2. Forward problem

The forward problem consists of determining system output quantities through the solution of governing equations given prescribed inputs. In the present case of contaminant transport, the inputs are parameters that describe the initial condition and the outputs are contaminant concentration measurements at various points in space and time.

2.1. Governing equations

Transport of the scalar contaminant field is assumed to be governed by a convection–diffusion equation,

$$\frac{\partial u}{\partial t} + \nabla \cdot (\vec{V}u) - \nu \nabla^2 u = 0, \tag{1}$$

$$u(\vec{x}, 0) = u_0(\vec{x}), \tag{2}$$

$$y_k = u(\vec{x}_k, t_k), \quad 1 \leqslant k \leqslant n_{\text{out}} \tag{3}$$

where $u(\vec{x}, t)$ is the contaminant concentration, $\vec{V}(\vec{x}, t)$ is the velocity field, ν is the diffusion coefficient, $t = 0$ is the time of release, $u_0(\vec{x})$ is the initial condition, and y_k are scalar outputs. The outputs are calculated at spatial positions $\vec{x}_k$ and times t_k. We allow for multiple sensors to be located at the same position $\vec{x}_k$ and for multiple outputs to be recorded at the same time t_k. In solving the forward problem, (1), $\vec{V}$, ν, and $u_0(\vec{x})$ are assumed known, and of interest are the outputs y_k, which require the primal solution $u(\vec{x}, t)$ for $t > 0$.

2.2. Discretization

We discretize (1) using a high-order discontinuous Galerkin (DG) finite element method. DG is chosen for this problem because it offers several attractive features including straightforward extension to high order interpolation, a compact stencil, and stability for convection-dominated flows. Details of the DG spatial discretization can be found in previous works [11,18], and here we only mention that we use pure upwinding for the inviscid flux and the second form of Bassi and Rebay (BR2) [19] for the viscous flux. Order p polynomials are used for the spatial approximation on each element, and in time we employ second-order backward differencing (BDF2) with BDF1 on the first time step. Once discretized, (1) takes the following form,

$$\mathbf{R}^i(\mathbf{u}^j) \equiv \mathbf{M}\left(\frac{3}{2}\mathbf{u}^{j+1} - 2\mathbf{u}^j + \frac{1}{2}\mathbf{u}^{j-1}\right) + \mathbf{R}_s(\mathbf{u}^{j+1}) = 0, \tag{4}$$

where $\mathbf{u}^j \in \mathbb{R}^N$ is the discrete state vector at time $t_j, 1 \leqslant j \leqslant N_T$, $\mathbf{u}^0$ is the discrete initial condition, $\mathbf{M} \in \mathbb{R}^{N\times N}$ is the block-diagonal mass matrix, $\mathbf{R}^i \in \mathbb{R}^N$ is an unsteady residual vector at time $t_i, 1 \leqslant i \leqslant N_T$, $\mathbf{R}_s \in \mathbb{R}^N$ is the spatial residual, and N is the number of spatial degrees of freedom.

We denote by $\mathbf{u}$, without any superscripts, all the unknowns in the entire time history rolled into one vector, i.e. $\mathbf{u} \in \mathbb{R}^{N\cdot N_t}$. The system in (4) can then be written as

$$\mathbf{A}\mathbf{u} = \mathbf{F}\mathbf{u}^0, \tag{5}$$

where $\mathbf{F}$ contains the dependence of the unsteady residual on the initial condition, only nonzero for the first two unsteady residuals, and $\mathbf{A}$ is a sparse constant matrix that contains the dependence of the unsteady residual on the state. The outputs in (3) can also be expressed in terms of the discrete state vector according to

$$\mathbf{y} = \mathbf{C}\mathbf{u} = \mathbf{C}\mathbf{A}^{-1}\mathbf{F}\mathbf{u}^0, \tag{6}$$

where $\mathbf{y} \in \mathbb{R}^{n_{\text{out}}}$ and where the matrix $\mathbf{C}$ consists of the spatial approximation functions evaluated at the desired sensor locations and time nodes.

2.3. Initial and boundary conditions

We assume a spatially-Gaussian distribution for the contaminant concentration at $t = 0, u_0(\vec{x})$. Thus the initial condition is described by three parameters: the distribution center, the standard deviation, and the amplitude. These parameters are rolled into one vector, $\boldsymbol{\mu} \in \mathbb{R}^{n_{\text{par}}}$, where n_{par} is at most 4 in two dimensions and 5 in three dimensions. This simplified model is relevant for many single-point release scenarios and allows us to compare statistical and deterministic inversion approaches.

No conditions are imposed at outflow boundaries of the domain, while $u = 0$ is imposed at inflow boundaries. Our cases will include geometrical objects simulating buildings, and on their boundaries a zero flux condition is imposed. Finally, we assume a spatially-varying but temporally constant velocity field, $\vec{V}(\vec{x})$, derived from potential flow. Specifically, we solve Laplace's equation for the velocity potential, $\phi(\vec{x})$, approximated with $p = 4$ polynomials in space on the same mesh. We then differentiate the potential element-wise to obtain the velocity field, $\vec{V} = \nabla\phi$ on each element.

3. Inverse problem

Whereas the forward problem concerns calculation of outputs of a system for given inputs, the inverse problem reverses this relationship and seeks unknown inputs from measured outputs. In the present contaminant transport problem, the inputs $\boldsymbol{\mu}$ are taken as parameters that describe the initial contaminant distribution, and the observed outputs $\bar{\mathbf{y}}$ consist of noisy contaminant concentration measurements at a limited number of sensors. In this section we present two approaches for obtaining $\boldsymbol{\mu}$ given $\bar{\mathbf{y}}$: one deterministic and one statistical.

3.1. Deterministic solution

The inverse problem can be formulated as a deterministic optimization problem of minimizing a cost function that incorporates the error between observed and simulated outputs along with a regularization term that is chosen to penalize certain features of the inputs:

$$\boldsymbol{\mu}^* = \arg\min_{\boldsymbol{\mu}} \mathcal{J}(\boldsymbol{\mu}), \quad \mathcal{J}(\boldsymbol{\mu}) = (\mathbf{y}(\boldsymbol{\mu}) - \bar{\mathbf{y}})^T(\mathbf{y}(\boldsymbol{\mu}) - \bar{\mathbf{y}}) + T(\mathbf{u}^0(\boldsymbol{\mu})). \tag{7}$$

The expression $\mathbf{y}(\boldsymbol{\mu})$ is shorthand for the chain of operations expressed in (6). $T(\boldsymbol{\mu})$ is a regularization term that alleviates ill-conditioning of the inverse problem – the use of parametrized initial