

Inversion of Airborne Contaminants in a Regional Model

Volkan Akcelik¹, George Biros², Andrei Draganescu⁴, Omar Ghattas³, Judith Hill⁴,
and Bart van Bloemen Waanders⁴

¹ Stanford Linear Accelerator Center
volkan@slac.stanford.edu

² Department of Mechanical Engineering and Applied Mechanics
University of Pennsylvania
biros@seas.upenn.edu

³ Institute for Computational Engineering and Sciences
The University of Texas at Austin
omar@ices.utexas.edu

⁴ Optimization and Uncertainty Estimation Department
Sandia National Laboratories {aidraga, jhill, bartv}@sandia.gov

Abstract. We are interested in a DDDAS problem of localization of airborne contaminant releases in regional atmospheric transport models from sparse observations. Given measurements of the contaminant over an observation window at a small number of points in space, and a velocity field as predicted for example by a mesoscopic weather model, we seek an estimate of the state of the contaminant at the beginning of the observation interval that minimizes the least squares misfit between measured and predicted contaminant field, subject to the convection-diffusion equation for the contaminant. Once the “initial” conditions are estimated by solution of the inverse problem, we issue predictions of the evolution of the contaminant, the observation window is advanced in time, and the process repeated to issue a new prediction, in the style of 4D-Var. We design an appropriate numerical strategy that exploits the spectral structure of the inverse operator, and leads to efficient and accurate resolution of the inverse problem. Numerical experiments verify that high resolution inversion can be carried out rapidly for a well-resolved terrain model of the greater Los Angeles area.

1 Introduction

We are interested in a DDDAS problem of localization of airborne contaminant releases in regional atmospheric transport models from sparse observations. Given measurements of the contaminant over an observation window at a small number of points in space, and a velocity field as predicted for example by a mesoscopic weather model, we seek an estimate of the state of the contaminant at the beginning of the observation interval that minimizes the least squares misfit between measured and predicted contaminant field, subject to the convection-diffusion equation for the contaminant. Once the “initial” conditions are estimated by solution of the inverse problem, we issue predictions of the evolution of the contaminant, the observation window is advanced in time, and the process repeated to issue a new prediction, in the style of 4D-Var.

The main difficulty of this inverse problem is that the inverse operator is formally dense and of grid dimension, and requires as many forward/adjoint solves to create as the number of grid points. Thus, the inverse problem is intractable for highly-resolved problems, even with the most powerful supercomputers available today. To address this issue, we invoke a conjugate gradient method that is tailored to the spectral structure of the inverse operator, in particular that the operator behaves like a Fredholm integral equation of the second kind. Thus, convergence can be obtained in a mesh-independent number of iterations, typically very small. Numerical experiments suggest that high resolution inversion can be carried out rapidly for a sufficiently well-resolved terrain model of the greater Los Angeles area.

In Section 2 we formulate the inverse problem as an output least squares optimization problem with a convection-diffusion PDE constraint. First order optimality conditions produce a coupled system of partial differential-algebraic equations, which includes the initial value convection-diffusion PDE, the terminal-value adjoint convection-diffusion PDE, and an algebraic equation for the initial concentration. This system is an ill-posed boundary value problem in 4D space-time; block-elimination to a 3D system in just the initial condition variables creates the aforementioned dense inverse operator (a Schur complement) that is of dimension of the grid points. A conjugate gradient method then makes effective use of its spectral structure.

Section 3 provides a prototype inversion scenario: localization of the release of a contaminant in the Los Angeles harbor from short-term measurements of its transport by onshore winds, followed by longer-term prediction of the transport of the contaminant throughout the Greater LA Basin. We conduct numerical experiments that demonstrate the capability and effectiveness of the inversion.

2 Formulation and optimality conditions

Mathematically, transport of the contaminant is described by the convection-diffusion equation. We seek to reconstruct the initial concentration from measurements of the concentration over a short time horizon, taken at a small number of sensor locations throughout the domain. In this section we give details on the mathematical formulation of the inverse problem, its discretization, and the strategy for its numerical solution.

Given observations of the concentration $\{u_j^*\}_{j=1}^{N_s}$ at N_s locations $\{\mathbf{x}_j\}_{j=1}^{N_s}$ inside a domain Ω , we wish to estimate the initial concentration $u_0(\mathbf{x})$ that leads to the closest reproduction of the observed concentrations. The inverse problem is formulated as a constrained, regularized, least squares optimization problem:

$$\begin{aligned} \min_{u, u_0} \mathcal{J}(u, u_0) &\stackrel{\text{def}}{=} \frac{1}{2} \sum_{j=1}^{N_s} \int_0^T \int_{\Omega} (u - u^*)^2 \delta(\mathbf{x} - \mathbf{x}_j) \, d\mathbf{x} \, dt + \frac{\beta}{2} \int_{\Omega} u_0^2 \, d\mathbf{x}, \\ \text{subject to} \quad &\frac{\partial u}{\partial t} - \nu \Delta u + \mathbf{v} \cdot \nabla u = 0, \quad \text{in } \Omega \times (0, T), \\ &\nu \nabla u \cdot \mathbf{n} = 0, \quad \text{on } \Gamma \times (0, T), \\ &u = u_0, \quad \text{in } \Omega \times \{t = 0\}. \end{aligned} \tag{1}$$

The first term in the objective functional \mathcal{J} represents a least-squares misfit of predicted concentrations $u(\mathbf{x}_j)$ with observed concentrations $u^*(\mathbf{x}_j)$, where the delta function localizes the u and u^* fields to the sensor locations. The second term in \mathcal{J} , scaled by the constant $\beta/2$, is a regularization term that introduces well-posedness in the inverse problem. In the absence of the regularization term, the inverse problem would be ill-posed even if measurements are available at all points in space and time: small perturbations in measurements, e.g. due to sensor faultiness or data transmission, would produce exponentially large errors in the recovered solution. It is also the case that we cannot hope to recover components of the initial concentration that are much more oscillatory than dictated by the spacing of the sensors. Therefore, oscillatory components of u_0 lie in the null space of the inverse operator and, in the absence of regularization, will appear as arbitrary noise in the reconstructed initial concentration field.

The constraints in the optimization problem (1) are the contaminant transport convection-diffusion equation, boundary condition, and initial condition, where $\mathbf{v}(\mathbf{x}, t)$ is the velocity field and ν is the diffusion coefficient. The transport of the pollutant is driven by the initial conditions, the diffusion, and the velocity field. In practice, the velocity field would be provided by a regional numerical weather prediction model. For our present purposes, however, we are interested in assessing the viability of our inversion method, and thus for simplicity, we employ a steady laminar incompressible Navier-Stokes solver to generate wind velocity fields over a terrain of interest.

The inverse problem then is to determine the initial concentration field $u_0(\mathbf{x})$, and the resulting space-time evolution of the concentration $u(\mathbf{x}, t)$, by solving the optimization problem (1). First-order necessary conditions for optimality, the *KKT conditions*, may be derived by introducing a Lagrangian functional:

$$\begin{aligned} \mathcal{L}(u, u_0, p) \stackrel{\text{def}}{=} & \frac{1}{2} \sum_{j=1}^{N_s} \int_0^T \int_{\Omega} (u - u^*)^2 \delta(\mathbf{x} - \mathbf{x}_j) \, d\mathbf{x} \, dt + \frac{\beta}{2} \int_{\Omega} u_0^2 \, d\mathbf{x} \\ & + \int_0^T \int_{\Omega} \left(p \frac{\partial u}{\partial t} + \nu \nabla u \cdot \nabla p + p \mathbf{v} \cdot \nabla u \right) \, d\mathbf{x} \, dt + \int_{\Omega} p(u - u_0) \, d\mathbf{x}, \end{aligned} \quad (2)$$

in which the *adjoint concentration* $p(\mathbf{x}, t)$ is used to enforce the convection-diffusion equation and initial condition. Requiring stationarity of the Lagrangian \mathcal{L} with respect to p , u , and u_0 , respectively yields the KKT conditions, which consist of:

The forward convection-diffusion problem

$$\begin{aligned} \frac{\partial u}{\partial t} - \nu \Delta u + \mathbf{v} \cdot \nabla u &= 0 \quad \text{in } \Omega \times (0, T), \\ \nu \nabla u \cdot \mathbf{n} &= 0 \quad \text{on } \Gamma \times (0, T), \\ u &= u_0 \quad \text{in } \Omega \times \{t = 0\}. \end{aligned} \quad (3)$$

The adjoint convection-diffusion problem

$$\begin{aligned} -\frac{\partial p}{\partial t} - \nu \Delta p - \nabla \cdot (p\mathbf{v}) &= -\sum_{j=1}^{N_s} (u - u^*) \delta(\mathbf{x} - \mathbf{x}_j), \quad \text{in } \Omega \times (0, T) \\ (\nu \nabla p + p\mathbf{v}) \cdot \mathbf{n} &= 0, \quad \text{on } \Gamma \times (0, T), \\ p &= 0 \quad \text{in } \Omega \times \{t = T\}. \end{aligned} \quad (4)$$

The initial concentration equation

$$\beta u_0 - p|_{t=0} = 0 \quad \text{in } \Omega. \quad (5)$$

The equations in (3) are the original forward convection-diffusion transport problem for the contaminant field. The adjoint convection-diffusion problem (4) resembles the forward problem, but with some essential differences. First, it is a terminal value problem, the adjoint p is specified at the final time $t = T$, rather than an initial value problem. Second, the convection is directed backward along the streamlines. Finally, it is driven by a source term given by the negative of the misfit between the predicted and measured concentrations at sensor locations. The initial concentration equation (5) is, in the case of L^2 regularization, an algebraic equation. Together, (3), (4), and (5) furnish a coupled system of linear PDEs for (u, p, u_0) . The principal difficulty in solving this system is that, though the forward and adjoint transport problems are parabolic-hyperbolic problems, the KKT optimality system is a *coupled boundary value problem in 4D space-time*.

To simplify discussion of solution approaches, we introduce the operators A, T, B and R . Here, A denotes the forward transport operator and A^{-1} its inverse; T extends a spatial field at initial time into space-time; B is an observation operator that localizes space-time to points at which sensors are placed; R is the regularization operator (in the present case the identity); A^* is the adjoint transport operator and A^{-*} its inverse; and T^* restricts a space-time field to a spatial field at $t = 0$. With these definitions, we can write the KKT conditions in operator form:

$$\begin{bmatrix} B & 0 & A^* \\ 0 & \beta R & -T^* \\ A & -T & 0 \end{bmatrix} \begin{bmatrix} u \\ u_0 \\ p \end{bmatrix} = \begin{bmatrix} Bu^* \\ 0 \\ 0 \end{bmatrix} \quad (6)$$

Special-purpose Krylov solvers and parallel preconditioners can be very effective at solving discretized versions of optimality systems such as (6) for optimization problems constrained by *steady-state* PDEs [1, 2]. Here, however, the 4D space-time nature of (6) presents prohibitive memory requirements for large scale problems. Solution of (6) in its *full-space* form is essentially intractable for such problems using present computing resources. Instead, we pursue a *reduced-space* method that amounts to a block elimination combined with a matrix-free Schur complement solver.

Eliminating the concentration u and forward transport equation (third row of (6)) and adjoint concentration p and adjoint transport equation (first row of (6)) from the KKT optimality system, we obtain the Schur complement system for the initial concentration u_0 :

$$Hu_0 = g, \quad (7)$$

where the *reduced Hessian* (or inverse) operator H is defined by

$$H \stackrel{\text{def}}{=} T^* A^{-*} B A^{-1} T + \beta R, \quad (8)$$

and the *reduced gradient* g is defined by

$$g \stackrel{\text{def}}{=} T^* A^{-*} B u^*.$$

It is immediately clear that H is a symmetric and boundedly invertible operator. Thus, the optimization problem has a unique solution for non-vanishing β .

Notice that H is a non-local operator and while explicit formation of H and its singular value decomposition constitutes an attractive and popular approach for small-scale inverse problems [3], alternative approaches are essential for large-scale problems, and in particular those for which rapid response is mandated. Therefore, we opt to solve the discretized form of (7) using the conjugate gradient (CG) method. H is never formed explicitly; instead, we compute $w = H v$, the action of the reduced Hessian on a given spatial field v , in matrix-free fashion as follows. (i) Set $u_0 = v$ and solve the forward transport equation (3) to obtain the concentration evolution u . (ii) Compute the misfit between the measurements u^* and the predicted concentrations u at the sensor locations. (iii) Use this misfit as a source to solve the adjoint transport equation (4) backward in time to obtain $p|_{t=0}$, the adjoint at $t = 0$. (iv) Set $w = \beta v - p|_{t=0}$. Therefore, each application of the reduced Hessian requires solving two transport equation, one forward in time and one backward. Besides u_0 , we need to store the entire time history of u (if we have measurements over the entire time interval $(0, T)$) but *only* at the sensor locations. In contrast with full-space methods, we avoid storing the forward and adjoint concentration time histories. Thus, the memory requirements for the inverse problem (1) are similar to those of the forward problem.

The overall computational cost of the (unpreconditioned) CG method is the work per iteration — dominated by the two transport equations solutions — multiplied by the number of CG iterations. The latter depends on the condition number of the reduced Hessian. One can show that, for fixed β , H is a compact perturbation of the identity, and its discrete version, denoted by H_h (h represents the discretization parameter), has a mesh-independent condition number [4]. Furthermore, its spectrum has a small number of clusters; it collapses exponentially onto β . Therefore, if we use a Krylov method, such as CG, to solve (7), the number of iterations for a specific relative residual reduction will also be mesh-independent. This is an optimal number of iterations, and we have verified numerically [5]. Thus, mesh-independent convergence comes for free, with the constant mildly deteriorating with $\beta \rightarrow 0$. The constant also depends on the Peclet number, the length of the time horizon, the complexity of the velocity field, and the topography.

3 Implementation and Numerical Results for Los Angeles region

We demonstrate our inversion framework on a hypothetical atmospheric contamination event in the Greater Los Angeles Basin (GLAB) region. Using real topographical data

and synthesized velocity fields, we conduct numerical experiments in which sparse observations are extracted from forward simulations and subsequently used in the inverse problem. Our implementation builds on PETSc [6] to manage parallel data structures, interface with linear solvers and domain decomposition preconditioners, and utilize a range of software services.

In our actual implementation, we first discretize the optimization problem (1), and then write optimality conditions (as opposed to the writing the infinite-dimensional optimality conditions (6) and then discretizing; in the present case the two are not identical [7]). We employ Streamline Upwind Petrov-Galerkin (SUPG) finite elements [8] in space and a Crank-Nicolson discretization in time. For problems with high Peclet number, stabilized methods such as SUPG are more accurate than standard Galerkin on coarse meshes. We use a logically-rectangular topography-conforming isoparametric hexahedral finite element mesh on which piecewise-trilinear basis functions are defined. Since the Crank-Nicolson method is implicit, we “invert” the time-stepping operator using a restarted GMRES method, accelerated by an additive Schwarz domain decomposition preconditioner, both from the PETSc library.

Contaminant transport is modeled over a $360 \text{ km} \times 120 \text{ km} \times 5 \text{ km}$ region in the GLAB. Land surface elevations are obtained at 1 km spacing from the USGS Land Processes Distributed Active Archive Center (GTOPO30 digital elevation model).⁵ The three-dimensional mesh is created from the surface elevations by inserting equally spaced grid points vertically from the surface grid to the top of the domain at 5 km.

To simulate a contamination event, an initial contaminant plume with a Gaussian concentration given by $20 \exp(-0.04|\mathbf{x} - \mathbf{x}_c|)$ is centered at $\mathbf{x}_c = (120 \text{ km}, 60 \text{ km}, 0 \text{ km})$. The plume is transported over the GLAB region by solving the forward convection-diffusion equation with specified velocity field over a time horizon of 120 minutes. Sensor measurements are taken every 3 minutes to develop a time history from which to invert. For this contaminant, the diffusion coefficient is taken as $\nu = 0.05$. The regularization parameter is fixed at $\beta = 0.01$. The forward and inverse problems are solved on a mesh with $361 \times 121 \times 21$ grid points, representing 917,301 concentration unknowns at each time step. The time step is the same as the sensor recording rate, i.e. 3 minutes, for a total of 40 time steps. Therefore, there are about 74×10^6 total space-time variables in the KKT optimality system (6).

A velocity field \mathbf{v} for the convection-diffusion equation is synthesized by solving the steady-state incompressible Navier-Stokes equations, $\rho(\mathbf{v} \cdot \nabla)\mathbf{v} + \nabla p - \mu\Delta\mathbf{v} = \mathbf{0}$, $\nabla \cdot \mathbf{v} = 0$, where p is the fluid pressure and (ρ, μ) are its density and viscosity. To simulate an onshore wind, an inflow Dirichlet boundary condition with $v_x = v_{\max} (z/(5.0 - z_{\text{surface}}))^{0.1}$ and zero for the other components is applied to the $x = 0$ plane, where v_{\max} is specified as 30 km/hr. Traction-free boundary conditions are applied to the outflow plane at $x = 360 \text{ km}$. Traction-free tangential and no-slip normal boundary conditions are applied to the remaining portions of the boundary. An SUPG-stabilized finite element method is employed to solve the Navier-Stokes equations using linear tetrahedral elements derived by subdividing the convection-diffusion hexahedral mesh.

⁵ <http://edcdaac.usgs.gov/gtopo30/dem.img.asp>

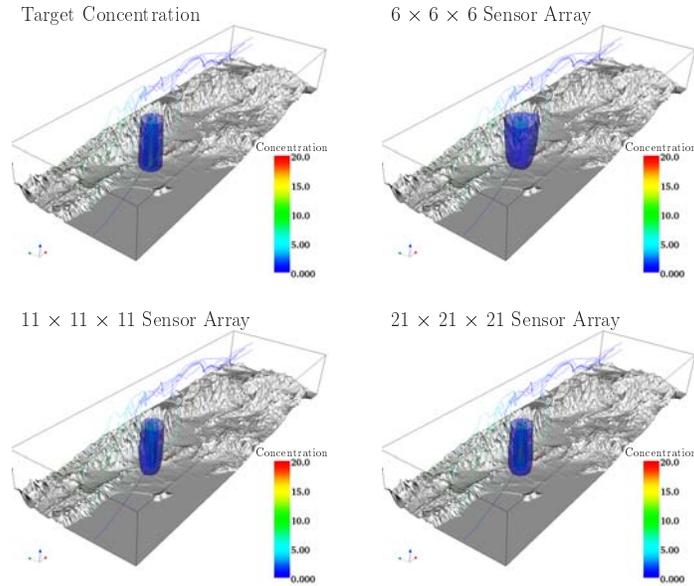


Fig. 1. Sensitivity of the inversion result to the sensor array density. The target initial concentration is shown in the upper-left corner, and inversion results using successively-finer sensor arrays are shown in the subsequent images. As the number of sensors in each direction increases, the quality of the reconstruction of the initial concentration plume improves. Inversion using the $21 \times 21 \times 21$ sensor array takes 2.5 hours on 64 processors of the Alphaserver EV68 system at the Pittsburgh Supercomputing Center.

We sample the concentrations from the forward transport simulations on a uniformly-spaced array of sensors, and use them as synthetic observations to drive the inverse problem. Figure 1 depicts inversion results for different sensor array densities, along with the actual initial concentration (labeled *Target* in the figure). One of the critical issues is to determine the number of sensors required to resolve the initial concentration. Recall that due to the non-vanishing regularization parameter and the fixed mesh size, we cannot expect to recover the initial concentration exactly.

What is of ultimate interest is how successful the reconstructed initial field is in predicting the actual transport of the contaminant. Figure 2 compares the actual evolution (left) and predicted evolution (right) of the contaminant plume in time. It is evident from this figure that although the reconstructed concentration does not match the actual concentration exactly at $t = 0$, the difference between the two diminishes over time, due to the dissipative nature of the forward convection-diffusion problem.

Acknowledgements

Supported in part by the Computer Science Research Institute at Sandia National Laboratories, the U.S. Department of Energy under the SciDAC Terascale Optimal PDE

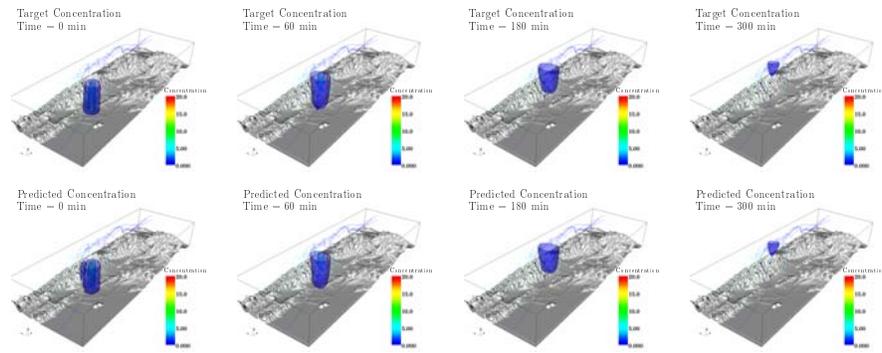


Fig. 2. Illustration of the predictive capabilities of our inversion algorithm, using an $11 \times 11 \times 11$ sensor array. Forward transport of the actual initial concentration is compared with forward transport of the reconstructed initial concentration plume. The trajectories are close to each other, and the comparison improves with time.

Simulations (TOPS) Center and grant DE-FG02-04ER25646, and the National Science Foundation under grants ACI-0121667, EAR-0326449, CCF-0427985, and CNS-0540372. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000. Computing resources at the Pittsburgh Supercomputing Center provided under NSF TeraGrid award MCA04N026P.

References

1. Biros, G., Ghattas, O.: Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization. Part I: The Krylov-Schur solver. *SIAM Journal on Scientific Computing* **27**(2) (2005) 687–713
2. Biros, G., Ghattas, O.: Parallel Lagrange-Newton-Krylov-Schur methods for PDE-constrained optimization. Part II: The Lagrange Newton solver, and its application to optimal control of steady viscous flows. *SIAM Journal on Scientific Computing* **27**(2) (2005) 714–739
3. Hansen, P.C.: Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion. SIAM (1997)
4. Drăgănescu, A.: Two investigations in numerical analysis: Monotonicity preserving finite element methods, and multigrid methods for inverse parabolic problems. PhD thesis, University of Chicago (2004)
5. Drăgănescu, A., Dupont, T.F.: Optimal order multi-level preconditioners for regularized ill-posed problems (2005) To appear.
6. Balay, S., Buschelman, K., Gropp, W.D., Kaushik, D., McInnes, L.C., Smith, B.F.: PETSc home page. <http://www.mcs.anl.gov/petsc> (2001)
7. Abraham, F., Behr, M., Heinkenschloss, M.: The effect of stabilization in finite element methods for the optimal boundary control of the Oseen equations. *Finite Elements in Analysis and Design* **41**(3) (2004) 229–251
8. Brooks, A.N., Hughes, T.J.R.: Streamline upwind/Petrov–Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier–Stokes equations. *Computer Methods in Applied Mechanics and Engineering* **32** (1982) 199–259