FINAL SCIENTIFIC/TECHNICAL REPORT

FOR AWARD DE-FE0009260

"AN ADVANCED JOINT INVERSION SYSTEM FOR CO_2 STORAGE MODELING WITH LARGE DATE SETS FOR CHARACTERIZATION AND REAL-TIME MONITORING – ENHANCING STORAGE PERFORMANCE AND REDUCING FAILURE RISKS UNDER UNCERTAINTIES"

APRIL 30, 2016

SUBMITTED BY

Stanford University
Bldg. Y2E2
473 Via Ortega
Stanford, CA 94305

PRINCIPAL INVESTIGATOR

Peter K. Kitanidis
Professor of Civil and Environmental Engineering

Stanford University
Bldg. Y2E2, Rm 147
473 Via Ortega
Stanford, CA 94305
Phone: (650) 723-8321
Email: peterk@stanford.edu

Project/Grant Period: October 1, 2012 to January 31, 2016

SUBMITTED TO

U.S. Department of Energy
National Energy Technology Laboratory
Karen Kluger– Federal Project Manager
412-386-6667, Karen.kluger@netl.doe.gov

DISCLAIMER

ABSTRACT

As large-scale, commercial storage projects become operational, the problem of utilizing information from diverse sources becomes more critically important. In this project, we developed, tested, and applied an advanced joint data inversion system for CO_2 storage modeling with large data sets for use in site characterization and real-time monitoring. Emphasis was on the development of advanced and efficient computational algorithms for joint inversion of hydro-geophysical data, coupled with state-of-the-art forward process simulations.

The developed system consists of (1) inversion tools using characterization data, such as 3D seismic survey (amplitude images), borehole log and core data, as well as hydraulic, tracer and thermal tests before CO_2 injection, (2) joint inversion tools for updating the geologic model with the distribution of rock properties, thus reducing uncertainty, using hydro-geophysical monitoring data, and (3) highly efficient algorithms for directly solving the dense or sparse linear algebra systems derived from the joint inversion. The system combines methods from stochastic analysis, fast linear algebra, and high performance computing. The developed joint inversion tools have been tested through synthetic CO_2 storage examples.

TABLE OF CONTENTS

1 EXECUTIVE SUMMARY 6

2 INTRODUCTION 10

2.1 CONTEXT: REAL TIME MONITORING OF CARBON DIOXIDE STORAGE 10

2.2 PROJECT OBJECTIVES 10

2.3 REPORT ORGANIZATION 12

3 BAYESIAN JOINT STOCHASTIC INVERSION: METHOD DEVELOPMENT AND VALIDATION 13

3.1 EFFICIENT JOINT INVERSION FOR OFF-LINE INVERSION 13

3.1.1 QUASI-LINEAR GEOSTATISTICAL APPROACH WITH THE FAST MULTIPOLE METHOD 13

3.1.2 PRINCIPAL COMPONENT GEOSTATISTICAL APPROACH 17

3.2 EFFICIENT JOINT INVERSION METHODS FOR DYNAMIC MONITORING 21

3.2.1 HIERARCHICAL KALMAN FILTER (HIKF) 21

3.2.2 SPECTRAL KALMAN FILTER (SPECKF) 25

3.2.3 COMPRESSED STATE KALMAN FILTER (CSKF) 28

3.2.4 SMOOTHING BASED COMPRESSED STATE KALMAN FILTER (SCSKF) 30

3.3 EFFICIENT SOLVERS FOR LARGE DENSE LINEAR SYSTEMS 32

3.3.1 BLACK BOX FAST MULTIPOLE METHOD (BBFMM) 32

3.3.2 FAST DIRECT LINEAR SOLVER AND DETERMINANT COMPUTATION FOR DENSE LINEAR SYSTEMS 33

3.3.3 BLOCK BASIS FACTORIZATION METHOD (BBF) 34

3.4 MATRIX VECTOR MULTIPLICATION WITH GPUS 36

4 METHOD TESTING FOR 3D SYNTHETIC CCS EXAMPLE 38

4.1 SYNTHETIC EXAMPLE DESCRIPTION 38

4.2 SYNTHETIC DATA GENERATION 39

4.3 SYNTHETIC INVERSION OF PUMPING DATA 40

4.4 SYNTHETIC JOINT HYDROTHERMAL TRACER TEST INVERSION 42

5 APPLICATION OF METHODOLOGY TO FIELD TEST SITES 44

5.1 FRIO-I PILOT TEST 44

5.1.1 PROBLEM DESCRIPTION 44

5.1.2 MODEL DEVELOPMENT & INVERSION 47

5.2 IN SALAH 49

5.2.1 PROBLEM DESCRIPTION 49

5.2.2 MODEL DEVELOPMENT & INVERSION 50

6 CODE DEVELOPMENT AND VALIDATION 52

6.1 CODE DEVELOPMENT 52

6.2 VALIDATION OF DEVELOPED ALGORITHMS 52

7 SUMMARY & CONCLUSIONS 53

8 REFERENCES 56

1 EXECUTIVE SUMMARY

A major factor that impedes the understanding and widespread acceptance of carbon dioxide storage in the subsurface for climate change mitigation is the difficulty in monitoring the spatial distribution and movement of injected CO_2, in order to prevent leakages to groundwater systems or to the ground surface. Sparsity of observations and the geologic complexity of deep subsurface formations introduce significant uncertainties in methods that attempt to reliably monitor injected carbon dioxide. Existing tools that use conventional methods for general-purpose model calibration are only limited to systems with few unknowns and observations and do not provide uncertainty quantification. Data assimilation techniques can be used to make progress in this important area; however, existing or "textbook" methods are not computationally capable to deal with realistic large-scale cases and multiple types of characterization and monitoring data.

In this project, we have developed fast and reliable methods for real-time monitoring of CO_2 geologic storage sites under incomplete information, with considerations of uncertainty and risk. The development of advanced sensor and computational technologies is changing the technological state of the art in real-time monitoring, allowing the collection of more frequent and more diverse measurements. Our work has focused on the development of computational tools that can utilize such complex monitoring datasets to provide accurate real time monitoring of CO_2 operations. Such tools can be used to (a) improve our understanding of the natural heterogeneity of storage systems and (b) to improve predictions of CO_2 spatial distribution.

This project has developed innovative data assimilation methods and tools that are based on sound fundamentals and are adapted to the physical and computational challenges encountered in CO_2 storage monitoring. We have developed, implemented, and tested methods that can estimate many unknowns, with quantification of uncertainty, with speeds that are orders of magnitude higher than traditional data processing methods, such that they are applicable for real large-scale projects. These methods utilize fast linear algebra tools, which are revolutionizing the way we do mathematical modeling. Such tools take advantage of special features (or "structure") in the mathematical methods, focus on computing only what is needed, perform computations only at an accuracy that is sufficient for the application at hand, and take advantage of available computational environments.

In particular, the tools developed comprise two algorithms for fast offline inversion and four algorithms for fast dynamic inversion. The common characteristics of these algorithms are that:

- They are fast, i.e., able to process large data sets using modest computer resources. This is achieved through the use of numerical techniques that improve computational efficiency by orders of magnitude. The computational **cost** of all methods developed **scales linearly** with the number of unknowns, as opposed to the quadratic scaling of conventional methods.
- They provide risk-based estimates to inform decision-making. This is achieved using sophisticated statistical techniques that provide **uncertainty quantification** measures for each estimated quantity.
- Use forward models in a **black-box fashion** so that the data assimilation code is generalizable for any data assimilation problem as long as a forward model exists. No modification of forward models is necessary as in adjoint state methods.

The following table lists the algorithms developed and the particular attributes and applicability of each one. The conventional methods are included for comparison:

Method	Cost (# simulations)	Attributes
Offline inversion		
Conventional	n (with adjoint methods)	
Fast linear GA (Geostatistical Approach)	none	Any linear problem
	The method combines the Quasi-Linear Geostatistical approach with the hierarchical matrices technique to reduce the cost of matrix-matrix multiplications including the large covariance matrix.	
PCGA (Principal Component Geostatistical Approach)	K*niter	Smooth problems
	The method combines the Quasi-Linear Geostatistical approach with a matrix factorization technique that compresses the error covariance matrix based on its eigenspectrum, thereby reducing the effective number of the unknown quantities.	
Dynamic inversion / Kalman Filter		
Conventional	n (with adjoint methods)	
HiKF (Hierarchical Kalman Filter)	none	Fast data acquisition, linear models
	The method modifies the conventional Kalman Filter by assuming a random walk forward model, thereby reducing the Jacobian calculations and accelerates matrix-matrix	

	multiplications using the hierarchical matrices approach.	
SpecKF (Spectral Kalman Filter)	n	Fast data acquisition, non linear models
	The method modifies the conventional Extended Kalman Filter by employing an approximation to the forward model that allows efficient cross-covariance updates.	
CSKF (Compressed State Kalman Filter)	K	Smooth problems
	The method modifies the conventional Extended Kalman Filter by using a low rank approximation of the covariance matrix based on its eigenspectrum and fixed eigenbases.	
sCSKF (Smoothing-based Compressed State Kalman Filter)	2K	Smooth problems, parameter estimation
	The method modifies the conventional Extended Kalman Filter by using a low rank approximation of the covariance matrix based on its eigenspectrum and a fixed basis. For improved parameter estimation and reducing linearization errors, the method employs a one-step ahead smoothing.	

n: number of measurements
m: number of unknowns (typically corresponds to gridded field of unknowns)
K: effective rank of covariance matrix
niter: number of iterations (typically 4-5)

Select tools were applied to synthetic monitoring field data and with two specific CO_2 injection sites in mind, including the Frio, Texas pilot and the commercial site at InSalah, Algeria, where the number of unknowns is close to 10^5. Application to these synthetic datasets illustrates the importance of computational efficiency for data assimilation applications with more than a few thousands and demonstrates the computational scalability achieved by our novel algorithms. Conventional algorithms would have prohibitive costs that would necessitate significantly reducing the resolution of the forward models and the reliability of the data assimilation method. In particular, application of our algorithm sCSKF to synthetic three-dimensional monitoring data for characterization of a site similar to the Frio-I site in Texas, demonstrated that:

- Even for strongly heterogeneous fields with significant anisotropy, large-scale formations of high and low conductivity can be identified based on a small number of wells, as long as vertically distributed measurements are available.
- Including thermal and chemical tracer data improves estimation accuracy. The relative worth of each dataset can be assessed using the Kalman Gain.

- Uncertainty estimates are extremely valuable as to where more measurements should be obtained to improve estimation and reliability.

Overall, the tools and methods developed in this project enhance both the actual reliability and the perceived safety of CCS facilities and are well suited to the policy and management challenges in CO_2 capture and storage (CCS).

2 INTRODUCTION

2.1 CONTEXT: REAL TIME MONITORING OF CARBON DIOXIDE STORAGE

While the deep subsurface can store much of the CO_2 produced by coal-powered power plants, the injection into the subsurface of millions of tons of carbon dioxide per year involves many technological challenges. A major difficulty is related to the heterogeneity of geologic formations and the *challenges in characterizing the sites and in monitoring the progress of injection and storage*. Not only is it difficult to identify the ideal sites, but also it is hard to manage a site, for example to regulate injection rates in order to optimize performance and minimize the risk of CO_2 leakage. The consideration of possible large-scale leakage, however unlikely such a leakage may be considered by some experts, is expected to be an impediment to the regulatory acceptance of many sites. The public demands assurances that the storage facility is properly monitored and an early warning system is in place. For the operators of the storage facility, slow long-term leakage is a significant financial risk if it must be recompensed by purchasing CO_2 emission permits. *Reliable monitoring both reduces actual risks and allays fears associated with perceived risks*. Existing methods for monitoring are based on either non-computational, experience-based evaluation of monitoring data, or rely on history matching of monitoring data and calibration of a small number of unknowns. There is need for a systematic, statistically sound, framework for utilization of monitoring data within numerical simulations, that can provide reliable estimation of unknown quantities along with uncertainty and risk analysis, in real-time. This need has been the primary motivation for the work done in this project.

2.2 PROJECT OBJECTIVES

The primary objective of the proposed research is to develop, test, and apply an advanced joint inversion tool system for CO_2 storage modeling with large sets of different types of characterization and monitoring data. The system will enable existing efficient numerical simulators for forward modeling of CO_2/brine two-phase flow and transport to be utilized *efficiently* in the inversion of field scale characterization and monitoring data. The system will also provide a stochastic and statistical framework for (real-time) monitoring, modeling, and inversion, which can work consistently with uncertainty quantification, risk assessment and mitigation, and optimal control and operation. To this end, the developed methods have the following characteristics:

- They are fast, i.e., able to process large data sets in real time using modest computer resources. This is achieved through the development of innovative numerical techniques that improve computational efficiency by orders of magnitude.
- They provide risk-based estimates to inform decision-making. This is achieved using sophisticated statistical techniques, that provide estimated quantities along with their uncertainties, calculated from the sensitivity between the data and the unknowns.

The advanced joint inversion tool system developed in this project enhances the predictive capability of models (*e.g.*, regarding CO_2 plumes) and improves storage performance through better understanding of storage systems (*e.g.*, spatially varying hydrogeologic properties and potential pathways for leakage), with dynamic reduction in uncertainty. It is an ultimate goal that the joint inversion tool system will be the central piece for a decision-making system for optimal control of CO_2 injection and storage by linking forward simulation, dynamic monitoring and inversion, uncertainty quantification, and risk assessment under a consistent framework. This project achieved the following specific objectives:

- A package was developed that includes efficient algorithms and fast solvers of linear algebraic system arising in inverse optimization problems in the joint inversion methodology with a very large number (N) of unknowns, such as Fast Direct Solvers with a computational complexity of O(N log N).
- The algorithms developed can be used for stochastic inversion for optimizing spatially varying rock properties using characterization and monitoring data; to demonstrate the updates of rock properties with incremental addition of field testing and monitoring data and the reduction in the uncertainties of the rock properties; and to demonstrate the efficiency of the joint inversion system with the developed fast linear solvers on practical data sets.
- The developed efficient joint inversion tools were tested with synthetic datasets for simple test cases for which they can be compared to conventional, exact but much more computationally expensive methods, as well as for large-scale synthetic examples for characterization with single-phase field tests.
- The developed algorithms are discussed in the context of the two real CO_2 storage sites with large characterization/ monitoring data.

The system consists of (1) statistical and computational inversion tools for providing prior distribution of rock properties using characterization data, such as 3D seismic survey (amplitude images), borehole log and core data, as well as hydraulic, tracer and thermal tests before CO_2 injection, (2) statistical and computational joint inversion tools for estimating the distribution of rock properties and CO_2 with uncertainty quantification using monitoring data for CO_2 injection and storage, and (3) highly efficient algorithms for directly solving the dense or sparse linear algebra systems derived from the joint

inversion and for performing matrix matrix multiplications involved in inversion algorithms. The developed joint inversion tool system is tested with simulated data using the TOUGH2-MP/EOS1 and TOUGH2-MP/ECO2N simulations for synthetic single phase and CO_2 storage examples respectively. The applicability and efficiency of developed methods is also discussed in the context of real datasets from two field sites: the Frio I Pilot Test and the In Salah Storage Project.

2.3 REPORT ORGANIZATION

This report is organized to be consistent with the Statement of Project Objectives for this project. An overview of the developed methods and applications is given and the reader is directed to journal articles for more technical details. This is followed by Summary and Conclusions section.

3 BAYESIAN JOINT STOCHASTIC INVERSION: METHOD DEVELOPMENT AND VALIDATION

The main feature of the methods presented under this section is that they are able to handle the excessive computational and storage costs of large-scale stochastic (*i.e.*, with quantification of uncertainty) inversions that prohibit the use of traditional inversion methods. Each method takes advantage of different properties or characteristics of a category of problems in order to simplify or approximate the computations. Our newly developed algorithms achieve this objective by capitalizing on the structure and the associated properties of the large covariance and sensitivity matrices involved in the inversion process, and by using linear algebra techniques that perform fast matrix-vector and matrix-matrix multiplications. An important feature of these approximate methods is that they are constructed with the requirement that the introduced approximation error is small and controllable.

3.1 EFFICIENT JOINT INVERSION FOR OFF-LINE INVERSION

3.1.1 Quasi-Linear Geostatistical Approach with the Fast Multipole Method

The first algorithm developed utilizes the fast multipole method (FMM) (Fong and Darve, 2009), in order to speed up a crucial part of linear and nonlinear inversion methods and is presented in detail in Ambikasaran et al. (2013). The bottleneck of stochastic linear inversion using the quasi linear geostatistical approach (QLGA) is the multiplication of the prior covariance matrix Q with the transpose of the measurement matrix H, *i.e.*, the Jacobian matrix of observations with respect to the unknowns. When the number of unknowns to be estimated (m) is large compared to the number of measurements (n), directly performing this multiplication can have tremendous computational and storage cost, as these costs scale with m^2. The proposed algorithm has a respective cost that scales with m, so that as the problem grows the cost can still be manageable. A major advantage of the proposed algorithm compared to other fast approaches is that it is not constrained to regular grids, as fast Fourier transform (FFT) type methods do, and is not constrained to smooth covariance functions, as low rank methods do. The method is general and applicable to irregular grids and a wide range of covariance matrices and can be easily applied for both two- and three-dimensional (2D and 3D) problems.

The basic premise of the proposed method can be summarized as follows:

A wide range of covariance matrices used in stochastic inversion have a structure called hierarchical low-rank, falling into the category of hierarchical (H^2) matrices. This is due to the fact that the interactions between far away clusters of points (far-field interaction) can be efficiently represented by a smooth function, which enables the use of low-rank approximations for each cluster. The fast multipole method (FMM) is an algorithm that uses the properties of H^2 matrices to compute dense matrix-vector products in a computationally efficient way and with sharp a priori error bounds. The accuracy of the method depends on the number of Chebyshev nodes used (*i.e.*, depth of the hierarchy used) as well as the effective rank of each low rank cluster (Fong and Darve, 2009).

Combining the QLGA with efficient ways to compute matrix matrix multiplications and with harnessing the sparsity of some of the matrices involved results in a method that overall scales linearly with the number of unknowns, which is the characteristic of this algorithm.

a. Validation

The method proposed is an approximation of the full QLGA for stochastic inversion. The approximation allows for the computational speed up, but also introduces an error, that is however small and controllable. To validate the method, we used a synthetic dataset of crosswell tomography created to resemble data acquired during monitoring at the Frio II test site in Texas. The data were created based on a collection of time-series seismic travel times recorded from a seismic source in one borehole and seismic receivers in the second borehole (Figure 1).

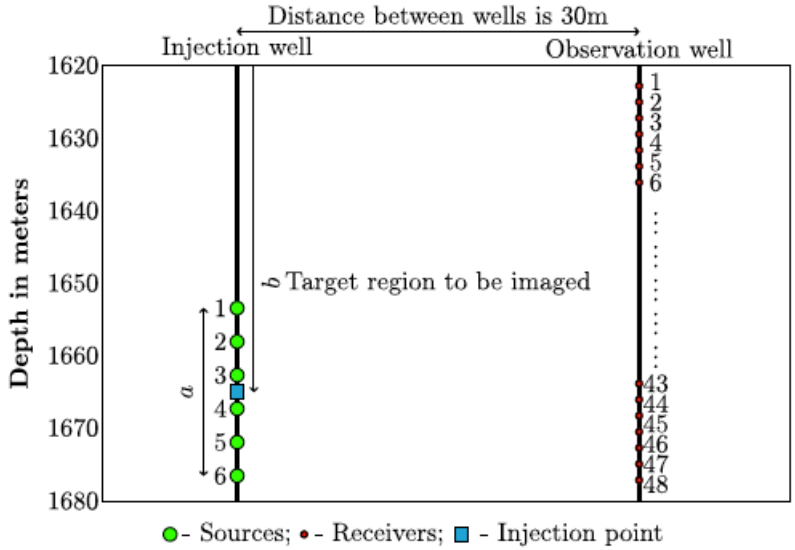

Figure 1: Schematic of Frio-II seismic monitoring experiment showing the sources and receivers in the two boreholes, and the point of CO_2 injection.

From the baseline measurements made, a delay (or "slowness") in travel time is produced over a period of time due to the presence of CO_2. This delay suggests the ray path transverses through material of lower seismic velocity, in this case, the CO_2 plume. To create the synthetic data, TOUGH2/ECO2N was used to produce a CO_2 saturation image, which was converted to an image of the "true" CO_2 induced slowness based on a petrophysical relationship. The objective of our algorithm is to use noisy travel time data obtained for the synthetic "true" CO_2 induced slowness image to reconstruct the CO_2 induced slowness.

Select results are shown in Figure 2, which shows the final estimates of our proposed method compared to the true slowness:

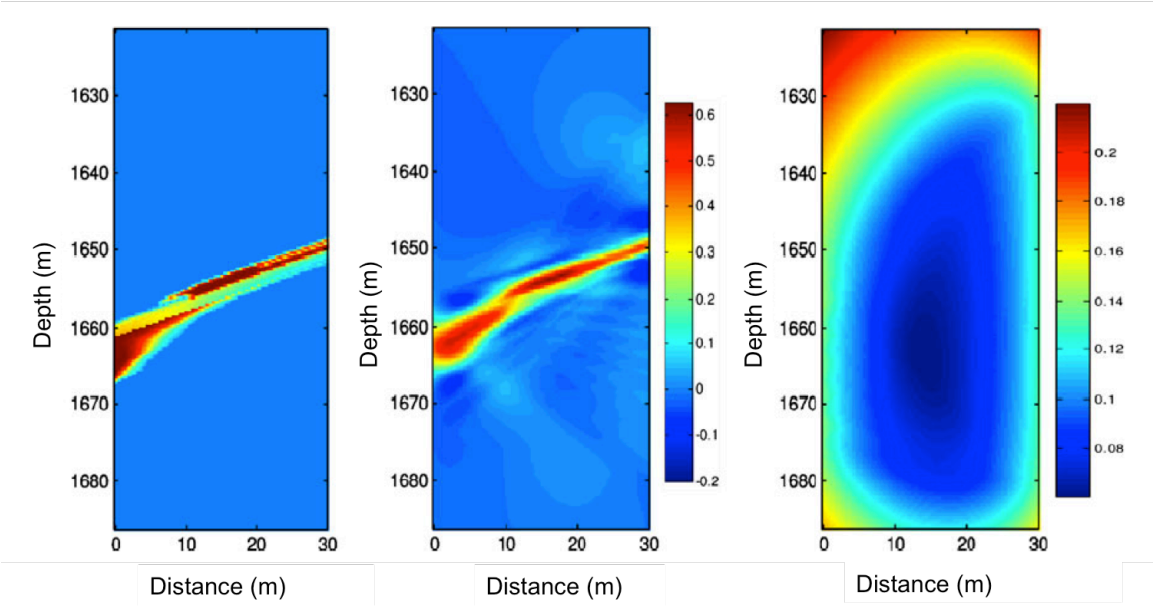

Figure 2: (left) true slowness simulated with TOUGH2/ECO2N (middle) reconstructed slowness using the proposed algorithm and (right) uncertainty in the estimated solution.

In terms of our methods accuracy, the only source of error in the computations is the approximation of the QH^T matrix. To investigate this, we computed the error as a function of the Chebyshev nodes:

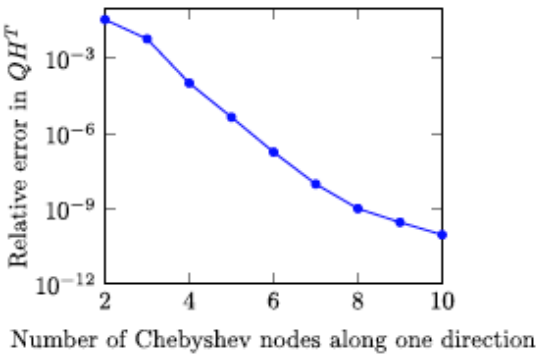

Figure 3: Relative error in QH^T versus the number of Chebyshev nodes along one direction for m=51249-

As shown in Figure 3, using as few as five Chebyshev nodes results in a negligible relative error in QH^T, and as a consequence, in the final estimates (Figure 2). An important attribute of the method is that it provides uncertainty estimates (right panel of Figure 2). Quantification of uncertainty is crucial in the context of inverse problems as it allows the assessment of the solution provided by inversion for a given set of data.

b. Computational efficiency

The approach speeds up the geostatistical method of inversion dramatically. The same crosswell tomography problem was solved for different discretization schemes so as to vary the number of unknowns and the results are shown in Figure 4:

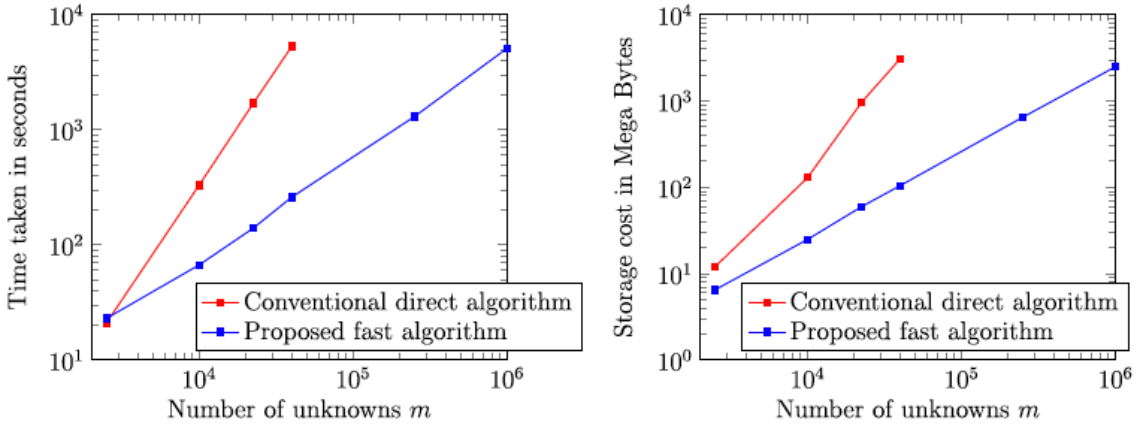

Figure 4: (left) Comparison of the time taken by the fast QLGA versus the conventional direct algorithm (right) same comparison for the storage cost.

In conclusion, we have developed a fast inversion algorithm for large-scale linear inversion that can be used for problems with large number of unknowns and relatively few observations. The computational speed up achieved by the proposed algorithms allows estimation of the unknowns and of the associated uncertainty, with a negligible

and controllable error. The computational speed achieved enables the use of our algorithm for use with large-scale systems with fine discretization, as well as for purposes of optimizing data collection. For example, and as discussed in more detail in Ambikasaran et al. (2013), the ability to run an inversion of 250,000 unknowns in less than 20 minutes allowed the optimization of the tomography sources and receivers and consequently of the capture zone. This attribute has a great impact for real field applications where collection of data is expensive and time and labor intensive.

3.1.2 Principal Component Geostatistical Approach

The second algorithm we developed to improve the computational efficiency of offline large scale inversions, is based on state compression methods. We proposed a new state compression method that can be used to significantly reduce both the computational time and the storage space requirements when solving large inverse problems. The method can be used to solve static inverse problems and is called Principal Component Geostatistical Approach (PCGA) (Kitanidis and Lee, 2014, Lee and Kitanidis, 2014). PCGA utilizes a matrix-free approach to calculate the required derivative information, overcoming the need for computationally expensive computation of a full Jacobian matrix and of performing matrix multiplications. In addition, PCGA exploits the properties of covariance matrices with low effective rank to reduce the size of the problem by orders of magnitude. The method is ideal for problems with smooth solutions and noisy measurements, such as hydraulic tomography and geophysical monitoring methods.

a. Validation

The PCGA is developed and detailed in Kitanidis and Lee, 2014. The premise of the method is that certain smooth covariance matrices can be well approximated by much smaller matrices obtained by eigenvalue decomposition. In particular, covariance matrices with an eigenspectrum that drops rapidly are most amenable to this approximation, as good accuracy can be obtained with very few principal components (N) compared to the number of unknowns ($N<<m$). This is shown for a synthetic one-dimensional flow with variable and unknown conductivity example with 100 unknowns. With just 20 principal components, the PCGA closely resembles the solution given by the full GA, as shown in Figure 5. The efficiency of the method is even greater for cases with noisy observations.

b. Evaluation

The PCGA was implemented for a large-scale two dimensional problem of hydraulic tomography, and was used to illustrate the efficiency and accuracy of the method as well as the advantages of joint inversion. The method is based on low rank approximation of the covariance matrix, which is performed in a computationally effective manner using the randomized eigenvalue decomposition method. This is combined with a matrix-free approach, which avoids direct evaluation of Jacobian matrices. The matrix-free approach also has the advantage that uses the forward simulation as a black box, and no modifications are necessary as in adjoint-state methods. The accuracy and efficiency of PCGA is determined by the number of principal components N used, which in turn determines the number of forward model evaluations. When the number of principal components is small, which is the case in many geoscience applications, the PCGA is scalable for very high dimensional problems.

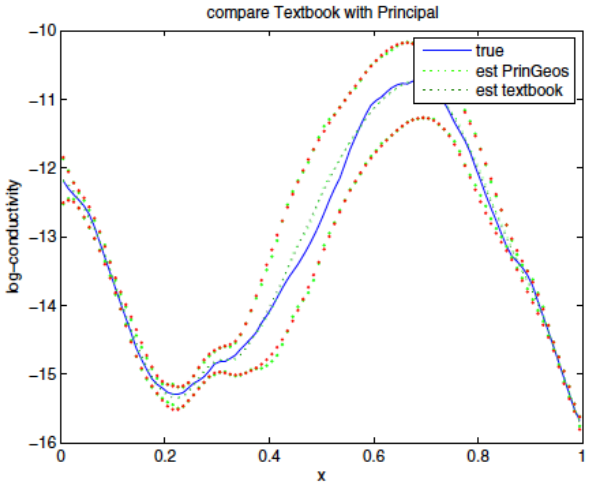

Figure 5: Comparison between the true conductivity, the estimates obtained by the textbook geostatistical approach (est textbook) and the estimates obtained by the proposed method PCGA (estPrinGeos).

The computational efficiency of the method was demonstrated using synthetic benchmark applications showing the applicability of the algorithm for inversion problems with more than one million of unknowns (Lee and Kitanidis, 2014). Two synthetic applications are considered for characterizing a two-dimensional heterogeneous domain: one where only pressure data are used, and one where pressure and tracer data are used. One of the two heterogeneous domains considered as well as the well locations used for monitoring are shown in Figure 6:

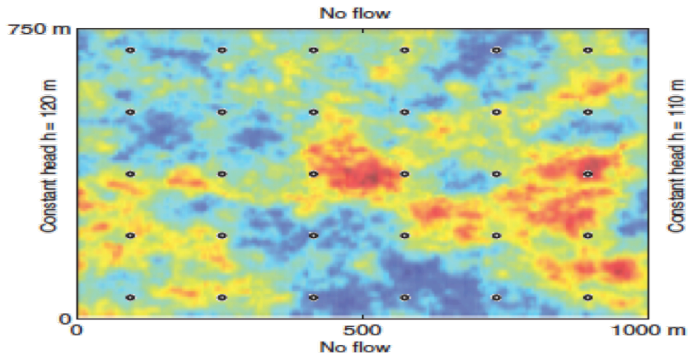

Figure 6: Domain and boundary conditions for synthetic applications of PCGA. Background color shows one of the two heterogeneous permeability fields (referred to as Case 2 below and in the paper) used to evaluate the method (Lee and Kitanidis, 2014).

Figure 7, shows the estimation for the case where only pressure data were used. It can be seen that the two heterogeneous conductivity fields are reconstructed, even with only 36 and 48 principal components respectively. The point-wise accuracy of estimation is related to the location and information carried in each measurement, as also reflected in the estimated uncertainty.

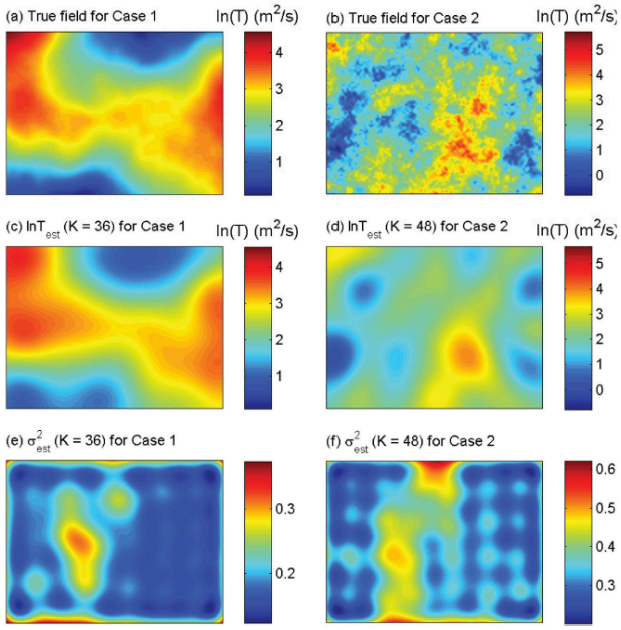

Figure 7: True log transmissivity fields (a) and best estimates (b) and uncertainties (c) using the PCGA.

The PCGA was also used to illustrate the benefits of joint inversion. For the same domain as above, the inversion was conducted using both head data, and tracer travel time data. The latter dataset is considered more informative as solute transport is more sensitive to conductivity variations compared to single-phase flow. Figures 8 and 9 show the

improvement in the data-fit and in the estimated conductivities respectively, when tracer data were included in the inversion.

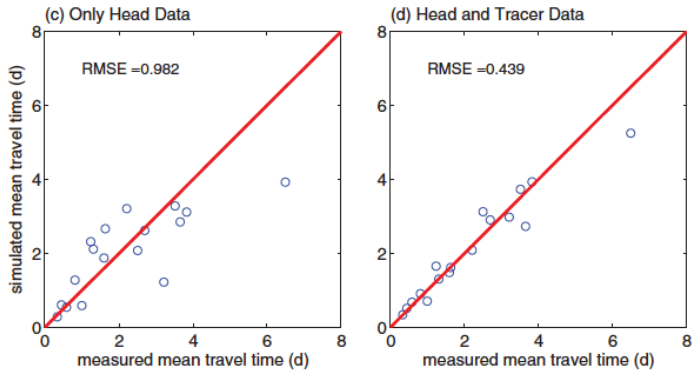

Figure 8: Measurement data fitting: simulated versus measured mean travel time for only head data (left) and from joint inversion of head and tracer data (right

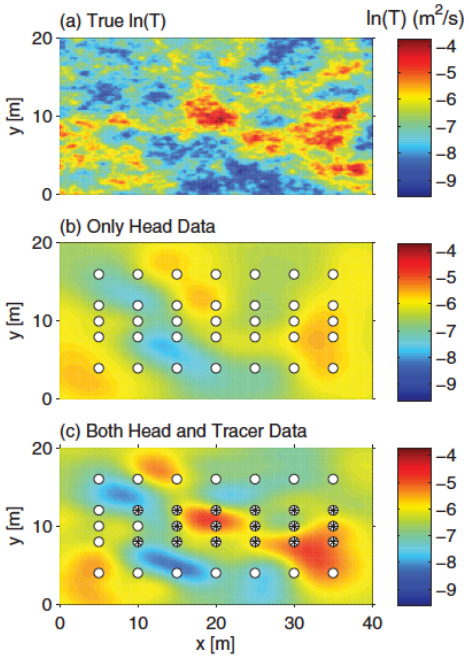

Figure 9: True log transmissivity field and best estimate for head data inversion (b) and head and mean travel time data inversion (c). The monitoring locations are indicated by circle for head and asterisk for tracer data.

In conclusion, the PCGA is a method that makes joint inversion for systems with millions of unknowns possible. This is particularly relevant for large-scale applications where it is important to characterize fine scale heterogeneity of the subsurface. The method takes advantage of fast linear algebra techniques to accelerate computations, and its accuracy is controllable and depends on the eigenspectrum of the error covariance of the unknown field. PCGA is a general-purpose inversion method and can be used with any simulation model.

3.2 EFFICIENT JOINT INVERSION METHODS FOR DYNAMIC MONITORING

Online or real-time estimation methods, in which observations are assimilated as they are obtained in time, are required for dynamic monitoring of a process. Such is the case for monitoring CO_2 during and after injection. Such data assimilation can be performed with various methods, the most well established being the Kalman Filter (KF). In simple terms, the Kalman Filter is a Bayesian approach (or minimum variance unbiased estimation method) for correcting the predictions of an uncertain dynamic model using observations that may be affected by error. Being a stochastic inversion approach, the Kalman Filter models the unknowns and the data as random variables and gives optimal estimates of the mean and standard deviation in the linear Gaussian case, *i.e.*, when the errors are zero mean Gaussian random variables, and the forward model and measurement operators are linear. When these assumptions do not hold, the Extended Kalman Filter (EKF) can be used, which applies the KF equations after linearization. Notwithstanding the linearization errors, the EKF can provide reasonable accuracy in many non-Gaussian, non-linear cases encountered in practice. However, the computational cost of applying the textbook version of the KF and the EKF is prohibitive for large-scale applications due to the large covariance matrices involved, the cost of computing Jacobian (*i.e.*, derivative) matrices, and matrix products involving the covariance matrix. In this section, we present four methods that have been developed to approximate the KF and EKF with a much smaller computational cost. Error analysis for all methods indicates that the proposed methods have a small and controllable error and have advantages over other, ensemble based, Kalman Filter methods.

3.2.1 HIERARCHICAL KALMAN FILTER (HIKF)

The first method is the Hierarchical Kalman Filter (HiKF) which is a fast implementation of the Kalman Filter that utilizes the FMM method and hierarchical matrices to improve storage and computational time requirements.

The version of HiKF that is presented in Li et al., (2014) is customized to the random-walk dynamical model, which is tailored to a class of data assimilation problems in which measurements are collected quasi-continuously. Using the random walk dynamical model, Jacobian-matrix computations are significantly simplified, such that the major bottleneck becomes the computation of the product QH^T, where Q is the error covariance matrix, and H is the measurement operator. As discussed in section 3.1.1, this bottleneck can be tackled by the FMM approach, which harnesses the structure of the covariance, which possesses the properties of H2 matrices. In addition, the method adopts a dynamic

update of cross-covariances (mxn) instead of the much larger covariance matrices. These features of the new method result in an overall computational cost that scales linearly with the number of unknowns and gives high accuracy as it does not involve any low-rank approximations.

The new method was validated and evaluated using the application of time-lapse cross-well seismic tomography discussed in Section 3.1.1, where measurements are collected periodically and the observations are used to track a CO_2 plume in the subsurface (Figure 10). Measurements of travel time between a network of sources and receivers (Figure 1) are used to estimate the CO_2 induced changes in slowness (i.e. delay in travel time) in a two dimensional domain due to the seismic velocity decrease in the reservoir from the CO_2 plume. This case is based on the model developed for the Frio-II brine pilot CO_2 injection experiment (Daley et al., 2007). The true slowness data associated with the presence of CO_2 were generated based on saturations from TOUGH2-ECO2N and using a petrophysical relationship between slowness and saturation.

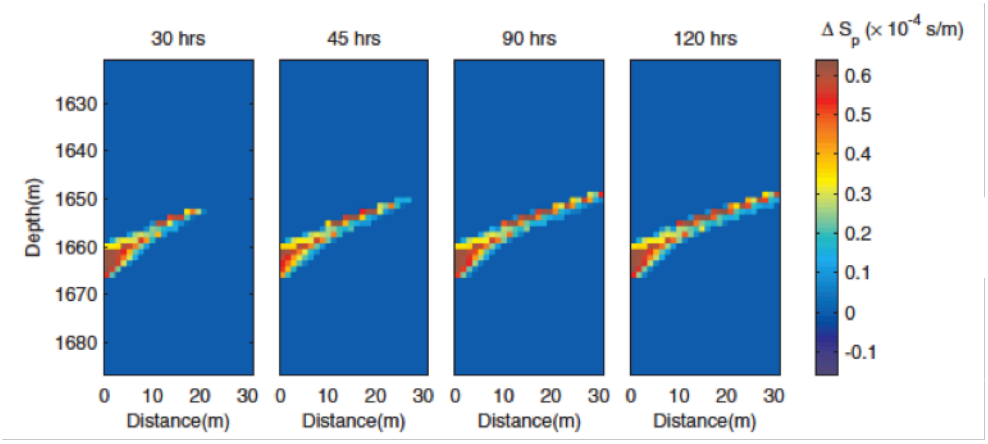

Figure 10: True CO_2 induced changes in compressional wave slowness between two wells used for validation of HiFK generated.

a. Validation

The HiKF is an approximation of the KF, with the approximation error introduced by the FMM method for the calculation of QH^T. First, for validation purposes, the HiKF was compared to the KF. As can be seen in Figure 11, the results from KF and HiKF are indistinguishable. This high accuracy is due to the fact that the FMM considers the full spectrum of the covariance, unlike low-rank methods. This is especially important for CO_2 applications, where the image to be estimated may have sharp interfaces and high frequency components that require a covariance with a slowly decreasing spectrum.

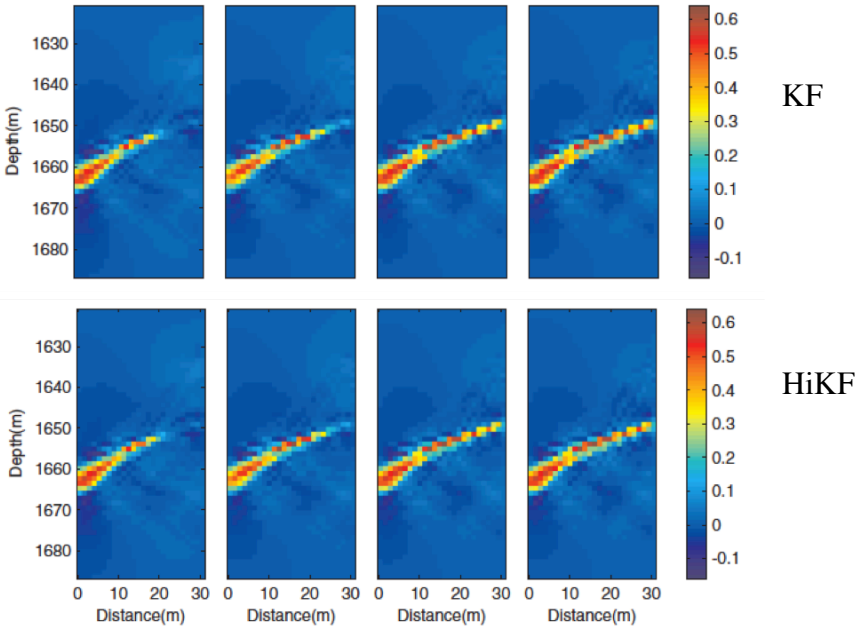

Figure 11: Estimated CO_2 induced changes in compressional wave slowness between two wells for KF (top) and HiKF (bottom).

b. Accuracy

The HiKF was able to reconstruct the true image of the CO_2 plume with high accuracy using the travel time measurements (Figure 12). The predictions of slowness become increasingly accurate over time as more data are assimilated. Such high accuracy is expected when data are available at frequent time intervals such that the assumption of the random walk model is applicable.

The HiKF results were also compared to results obtained by the Ensemble Kalman Filter, which is a popular low rank method used for data assimilation. As mentioned above, the test case discussed here, as is typical of CO_2 injection scenarios, has sharp interfaces that a low rank filter is not expected to capture well. As shown in Figure 13, the EnKF has a higher estimation error relative to both the true solution and the KF solution for the same computational cost, as expected. More specifically, the EnKF solution is smeared and has numerous random features that are not present in the true solution. These shortcomings are also prevalent in the uncertainty estimated by EnKF, which is significantly noisier compared to HiKF. In contrast the HiKF gives accurate uncertainty estimation, practically equivalent to that given by KF, with lower uncertainty in the region with high seismic ray coverage.

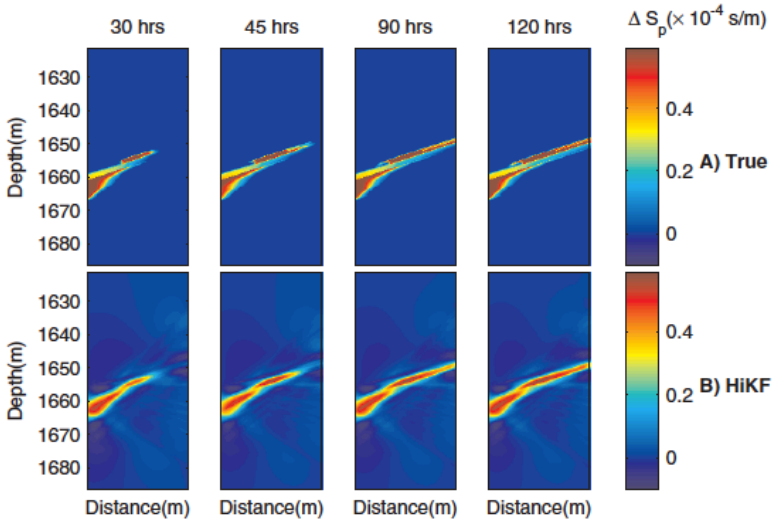

Figure 12: True and HiKF estimates for the compressional wave slowness.

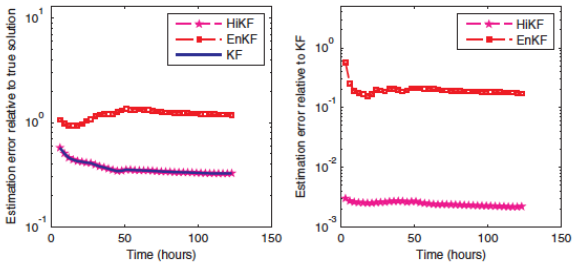

Figure 13: Estimation error relative to the true solution (left) and relative to the KF solution (right).

c. Efficiency

The HiKF is significantly faster than the textbook version of Kalman filter and has much lower storage costs, as it scales linearly with the number of unknowns, as opposed to the quadratic KF scaling. In the particular application discussed here, for a number of unknowns m=50778, the computational time is reduced from approximately 4 hours to less than 2 minutes. The improvement that our method provides in storage and computational time requirements increases as the inversion problem becomes larger, which manifests the advantage of HiKF for large-scale applications, and as shown in Figure 14.

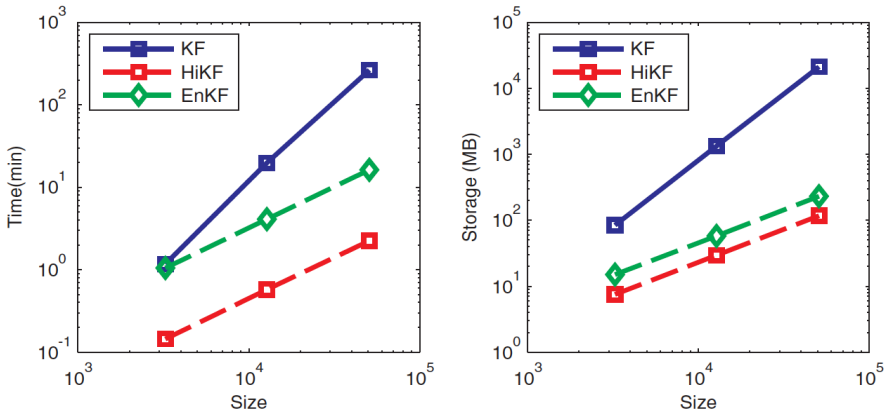

Figure 14: Computational time and storage cost of each data assimilation method plotted as a function of the number of unknowns. The proposed HiKF algorithm outperforms the full Kalman Filter (KF) and the Ensemble Kalman Filter (EnKF) in both storage and computational time requirements (Li et al., 2014)

Overall, the Hierarchical Kalman filter (HiKF) provides high accuracy and significant computational savings and is ideally suited for cases with large number of unknowns, relatively smaller number of measurements at a given point in time, and when measurements are collected frequently such that a random walk model can be used to approximate the forward model operator. Details about the experiments and results can be found in Li et al. (2014).

3.2.2 Spectral Kalman Filter (SPECKF)

The second dynamic inversion method developed is the Spectral Kalman Filter (SpecKF) and is described in detail in Ghorbanidehno et al. (2015). The SpecKF is an extension of the HiKF for more general cases where the transition matrix is not the identity matrix (i.e. random walk model), such as forward models in hydrogeologic applications (*e.g.*, TOUGH2). The method uses a forward model approximation and is better suited for cases where the time interval between observations is reasonably small. The algorithm is designed such that it can be used for any non-linear problem for which a forward model is available. The forward model is used as a black box to compute the Jacobian of the state transition equation using a finite differences approach. The computational speed-up of the SpecKF is achieved by updating cross-covariance matrices instead of the larger covariance matrices. The benefit can be considerable, especially in large systems, because the computational complexity of the SpecKF scales with the number of measurements, as opposed to the effective rank of the covariance matrix in low-rank KFs or the number of ensemble members in ensemble methods. The filter sacrifices a little accuracy in the operations involving the forward model to achieve a large improvement

in implementation speed. A great advantage of the SpecKF is that it does not suffer from errors from low rank approximation of the state covariance matrix.

a. Validation

We have validated the SpecKF by comparing it to the full KF. To be able to run KF with reasonable computational times for this comparison we used a simple linear diffusion problem.

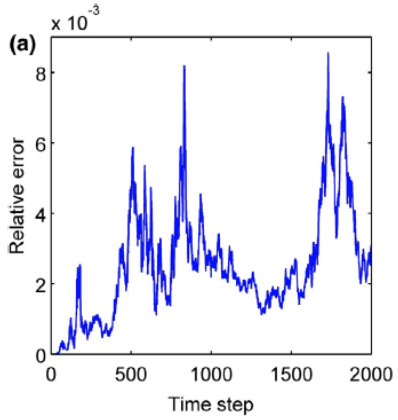

Figure 15: Average relative error between the SpecKF and the KF estimate with time.

Agreement with the full Kalman Filter verifies that the algorithm and implementation are correct. A very small relative error of the order of 10^{-3} confirms that the performance of the SpecKF is satisfactory (Figure 15). Our error analysis demonstrated that the SpecKF also provides accurate estimates of the uncertainty at the measurement points. The SpecKF algorithm does not provide the entire spatial distribution of the uncertainty to save on computational cost, but it is possible to obtain uncertainty estimates at a specified location, at a cost of one additional forward simulation per location.

b. Accuracy and efficiency

We have implemented and tested our method for a non-linear forward model where CO_2 is injected in a homogeneous two-dimensional domain. The filter is initiated with erroneous initial conditions representing a case where the initial CO_2 saturation is unknown, and is using 25 noisy measurements of saturation and 9 pressures collected every 15 days to update the field of pressure and CO_2 distribution on a 45x45 computational grid (2x2025=4050 unknowns).

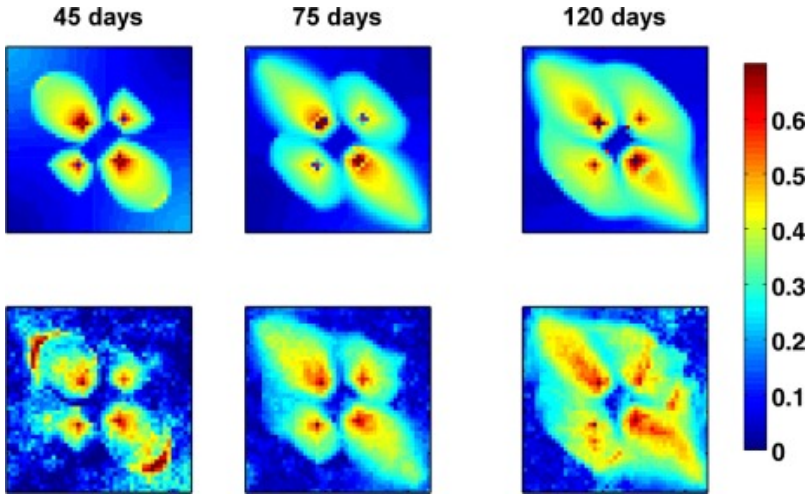

Figure 16: Comparison of true saturation with the SpecKF and the EnKF estimations for the case with small measurement error. First row: the SpecKF estimation for saturation; second row: the EnKF estimation for saturation with the same computational cost.

We compare the results of SpecKF to the EnKF, which is the most commonly applied fast Kalman Filter. We tested the algorithm under a variety of scenarios and found that the SpecKF provides same or better accuracy than the EnKF. More specifically, we demonstrate that SpecKF has superior performance than EnKF, which is known to fail under two specific scenarios: firstly, for cases with low measurement error, the EnKF is known to diverge, producing spurious features that are inconsistent with the physical model. This is shown in Figure 16, which demonstrates that SpecKF is more robust than EnKF for low measurement error and for the same computational cost.

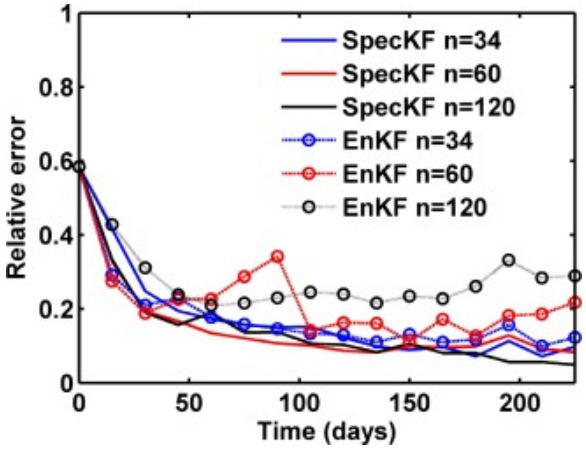

Figure 17: Comparison of the relative error of EnKF and SpecKF saturation estimation with respect to the true solution for three different number of measurements

Secondly, the EnKF is known to require a large number of realizations as the number of measurements increases to achieve the same accuracy, thus increasing its computational cost. The SpecKF computational cost scales exactly linearly with the number of

measurements, having a predictable increase in computational cost as the number of measurements increases, unlike EnKF where the number of realizations required is unknown. This is shown in Figure 17, where the estimation error is shown for the two algorithms for the same computational cost, and varying the number of measurements used. As shown, the EnKF error increases with the number of measurements, while the SpecKF error remains relatively constant.

Overall, the SpecKF is advantageous for non-linear estimation in systems with sharp interfaces since it considers the full covariance spectrum with no low-rank approximations and since its accuracy depends on controllable, physically meaningful parameters. Its computational cost scales linearly with the number of measurements and unknowns, thus being applicable for large-scale systems where the traditional KF is computationally prohibitive to apply. Compared to EnKF, SpecKF provides more reliable uncertainty quantification.

3.2.3 Compressed State Kalman Filter (CSKF)

The third dynamic inversion method developed is the Compressed State Kalman Filter (CSKF) (Kitanidis, (2014); Li et al., (2015)). The CSKF uses N preselected orthogonal bases to compute an accurate rank-N approximation of the covariance that is close to the optimal spectral approximation given by SVD. The CSKF uses an efficient matrix-free approach that propagates uncertainties using $N + 1$ forward model evaluations, where N is the effective rank. As the CSKF is a low-rank method, it is ideally suited for smooth-state problems where the covariance eigenspectrum of the state covariance decays rapidly. In such cases, a great computational speed up can be achieved.

a. Validation

We first validate the CSKF for a linear state-space model, as in this case it is possible to evaluate the full KF and use the KF estimates (*i.e.*, the optimal estimate) to validate the generalized CSKF algorithm. We show that as the rank N increases, the CSKF converges to the KF (Figure 18). The CSKF accurately reproduced the Kalman gain and state estimates of KF with a small number of bases and gave high quality uncertainty quantification.

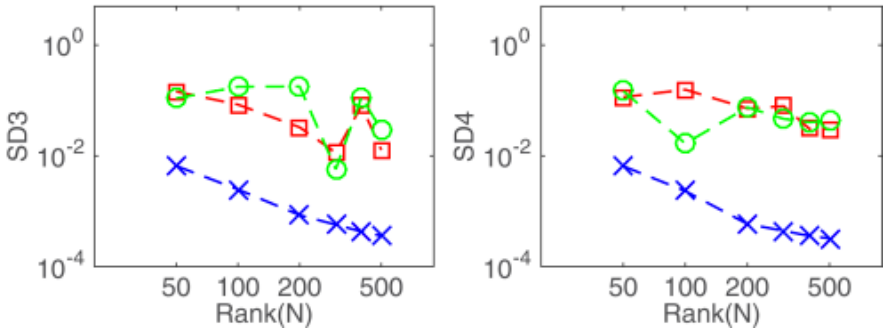

Figure 18: Convergence analysis for validation of CSKF (left) SD3: log-log plot of errors in the Frobenius norm of the posterior covariance and (right) SD4: log-log plot of errors in the Frobenius norm of Kalman gain *K*. All metrics measure errors relative to KF.

b. Accuracy and computational efficiency

The CSKF was further evaluated for a CO_2 monitoring case, which is a more complex, nonlinear problem governed by multiphase physics. CO_2 saturation in a two-dimensional homogeneous domain is estimated from measurements of pressure, flow rate, and saturation, given erroneous initial conditions, i.e. assuming that no CO_2 is present initially in the domain. The forward simulator TOUGH2-ECO2N is used for the forward simulations. The CSKF was shown to perform equally well to the EnKF in terms of the estimated mean, identifying the main features of the saturation distribution (Figure 19). However, large discrepancies were observed between the two methods in the computed posterior variance (Figure 20). In particular, the CSKF provided robust Kalman gain and uncertainty estimates that did not degrade as more data were assimilated, as opposed to drastic variance reductions predicted by the EnKF.

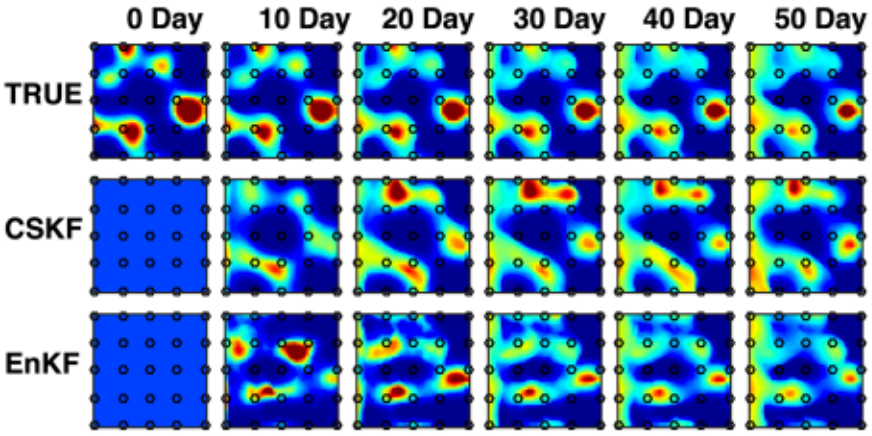

Figure 19: Solution given by CSKF and EnKF with $N = 50$: true CO_2 saturation and its posterior mean given by each method. The sampling locations of saturation are marked by circles.

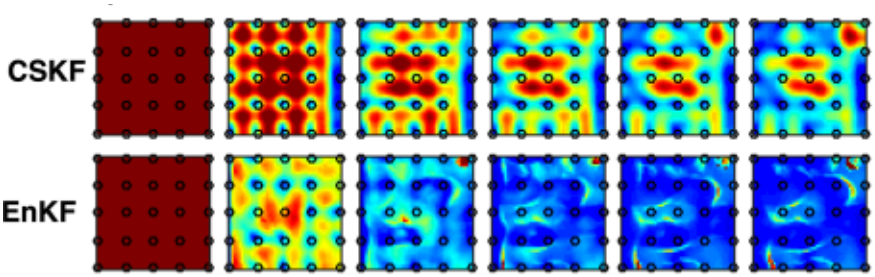

Figure 20: Posterior standard deviation given by CSKF and EnKF, and EnKF+LOC with $N = 50$ The sampling locations of saturation are marked by circles.

The CSKF provides a reliable and robust method to control the trade-off between the accuracy and the computational cost of low-rank Kalman filtering for nonlinear estimation problems and is a practical alternative to conventional computationally intensive inversion techniques and ensemble-based compression techniques.

3.2.4 Smoothing based Compressed State Kalman Filter (sCSKF)

The fourth dynamic inversion method developed is a variant of the CSKF, and is termed the smoothing-based Compressed State Kalman Filter (sCSKF). The sCSKF was created to handle cases in which strong non-linearities cause significant linearization errors during the estimation, which are often manifested with state updates that are not consistent with each other and the physics of the problem. One such case is when the parameters of the forward model are unknown and are being estimated as part of the data assimilation problem. In the CCS context, this would be a case where the permeability of the injection formation is unknown and is being estimated from the data, together with the CO_2 saturation and pressure. In this case, the uncertainty in the permeability may cause CO_2 saturation estimations that are not consistent with multiphase physics.

To address these problems, the sCSKF algorithm includes the addition of a smoothing step for each data assimilation step that improves the initial guess around which the update is linearized, thus reducing the linearization error. This algorithm has a higher computational cost than CSKF, but is better suited for joint parameter and state estimation, and is still significantly more computationally efficient than the full Kalman Filter.

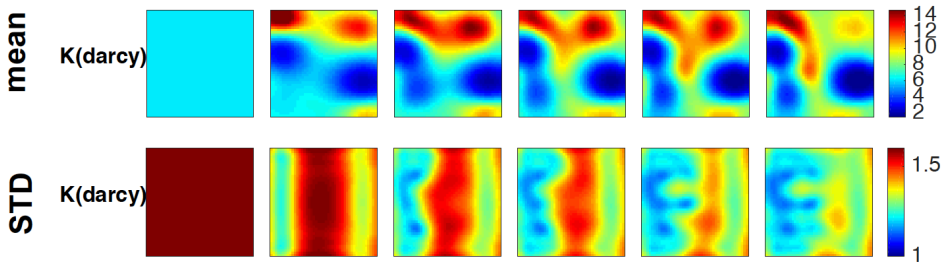

Figure 21: Estimated permeability and associated uncertainty at different times as more data are assimilated (left to right).

Figure 21 shows the results for a two dimensional application similar to the one shown in Figure 19, where in addition to CO_2 saturation and pressure, the underlying heterogeneous permeability field is also estimated. It can be seen that as more data are being assimilated (left to right) more features of the heterogeneous permeability distribution are being identified, and the uncertainty associated with these features is reduced.

3.3 EFFICIENT SOLVERS FOR LARGE DENSE LINEAR SYSTEMS

3.3.1 BLACK BOX FAST MULTIPOLE METHOD (BBFMM)

The Fast Multipole Method (FMM) (Fong and Darve, 2009) is a numerical technique to calculate matrix-vector products, of matrices of a special structure, or equivalently sums of the form

$$f(x_i) = \sum_{j=1}^{N} K(x_i, y_j)\, \sigma_j$$

where $i \in \{1,2, \ldots M\}$, in $O(M + N)$ operations as opposed to $O(MN)$ with a controllable error ϵ, where M and N are the dimensions of the matrix K.

The FMM is applicable for a class of matrices called hierarchical matrices, which are data-sparse approximations of dense matrices arising in applications like boundary integral equations or stochastic analysis. An example of a hierarchical matrix arising out of a two-dimensional application is shown in Figure 22. Of the class of all hierarchical matrices, H2 -matrices are matrices whose structure can be harnessed by the FMM in order to efficiently calculate matrix vector products.

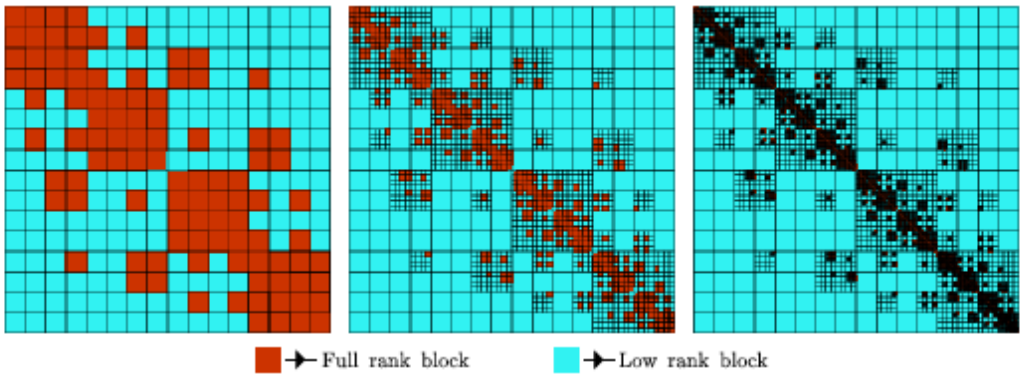

Figure 22: Hierarchical matrices arising out of a two dimensional problem at different levels in the tree

The Black Box Fast Multipole Method (BBFMM) is a tool developed in this project that has extended the kernels (covariance matrices) for which the FMM is applicable. The approximation scheme used in the BBFMM relies on Chebyshev interpolation to construct low-rank approximations for well-separated clusters. In addition the use of Singular Value Decomposition ensures that the computational cost is minimal. In particular the rank is optimally chosen for a given error. Some significant advantages of

the BBFMM are that it is very convenient for complex kernels for which analytical expansions might be difficult to obtain; it is easy to implement and have known error bounds. Its computational efficiency can be significant and can be increased depending on the accuracy required, while the pre-computation cost is very small.

The BBFMM forms the backbone of our fast offline inversion algorithm presented in Section 2, and of the fast Kalman Filters developed in this project, HiKF and the SpecKF, where they are used to perform matrix-matrix and matrix-vector multiplications with dense covariance matrices in an efficient way.

The BBFMM code has been developed for 2D and 3D systems and is available to the public at:

https://github.com/sivaramambikasaran/BBFMM2D

https://github.com/ruoxi-wang/BBFMM3D

https://github.com/judithyueli/mexBBFMM2D

https://github.com/ruoxi-wang/mexBBFMM3D

3.3.2 Fast Direct Linear Solver and Determinant Computation for Dense Linear Systems

In addition to the fast numerical algorithms specifically developed for Kalman filters and stochastic inversing, we have also developed a series of general fast algorithms for linear algebra, in particular for solving linear systems. A series of algorithms and codes were developed to address a range of problems:

1. Dense matrices. Extending ideas from the fast multipole method, we have designed algorithms to solve Ax = b where A is a dense matrix defined by an RBF (radial basis function) kernel such as a Gaussian or exponential kernel. This solver has cost O(n) where n is the size of the matrix. The details of this work can be found at Coulier et al., 2015.

2. Sparse matrices: multifrontal methods. Multifrontal techniques form the basis of efficient direct linear solvers for sparse matrices. In 3D, however, these solvers have a cost of $O(n^2)$ and become prohibitive. We have been able to reduce this cost to $O(n^{4/3})$ by accelerating the calculation with the dense blocks using the hierarchical matrix format. The details of this work can be found at Aminfar et al., 2016.

3. Sparse matrices: multilevel methods. The hierarchical matrix format was used to extend and generalize algebraic multigrid methods (AMG). This has led to a new class of O(n) preconditioners that are more robust and black-box than conventional AMG methods, which typically require optimizations specific to the problem at hand. The details of this work can be found at Pouransari et al., 2015.

3.3.3 Block Basis Factorization Method (BBF)

The block basis factorization (BBF) algorithm was developed to speed up matrix-vector multiplications where the kernel matrix is generated from high dimensional data. The BBF overcomes the complexity issue in Fast Multipole Method where there is an exponential dependence on the dimension; it also generalizes vanilla low-rank matrix approximations to the cases where low-rank approximation fails to be efficient. The latter characteristic is most relevant to the applications of interest to this problem.

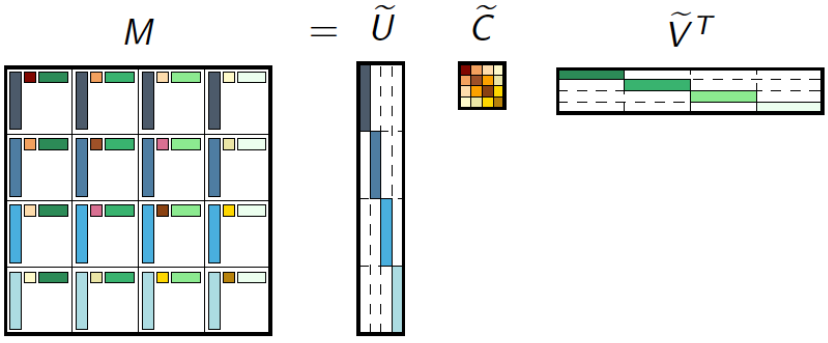

Figure 23: Schematic showing the Block Basis Factorization of a dense matrix M.

The BBF conceptualization, shown in Figure 23, gives a rank-rk approximation with O(nr + (rk)^2) storage. This should be contrasted with a traditional low-rank scheme that gives a rank-r approximation using O(nr) memory. Since the memory cost is a close approximation of the running time for a matrix-vector multiplication, BBF therefore offers an efficient scheme that is linear in both memory and application time.

For the low dimensional cases of interest to this project, the BBF structure provides a generalization of the low-rank scheme, especially for the cases where the low-rank property breaks and traditional low-rank methods fail. It also provides an alternative to the FMM in cases where a large number of Chebyshev nodes is needed for the desired accuracy.

BBF method is highly memory efficient, since the basis U and V are block-diagonal instead of dense. In terms of applicability, it is more accurate when h is small: it considers local interactions instead of only global ones (which low-rank methods do);

We have tested the efficiency of the BBF by comparing it to various low rank schemes using different sampling methods (Wang et al., 2015). Figure 24 compares the accuracy of the BBF structure vs. the low-rank structure; the results show that BBF approximates the kernel matrix better than the low-rank structure, regardless of the sampling method used in the algorithm.

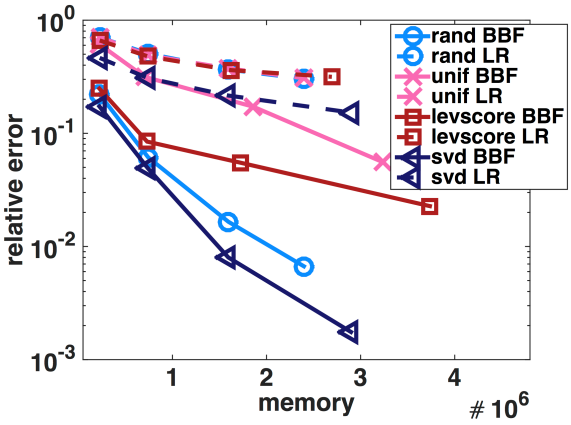

Figure 24: Kernel approximation error versus memory for BBF and low rank structure with different sampling methods. Different symbols represent different sampling methods.

For further testing, we benchmarked BBF against other state-of-art kernel approximation methods on some real datasets (Figure 22). BBF exhibits a significantly higher accuracy than the competing methods. We are comparing against the naive Nystrom (Nys), k-means Nystrom (kNys), leverage score Nystrom (lsNys), the Memory Efficient Kernel Approximation method (MEKA), and the Random Kitchen Sinks (RKS). It is worth noting that due to its flexible structure, it outperforms the exact SVD (cubic time).

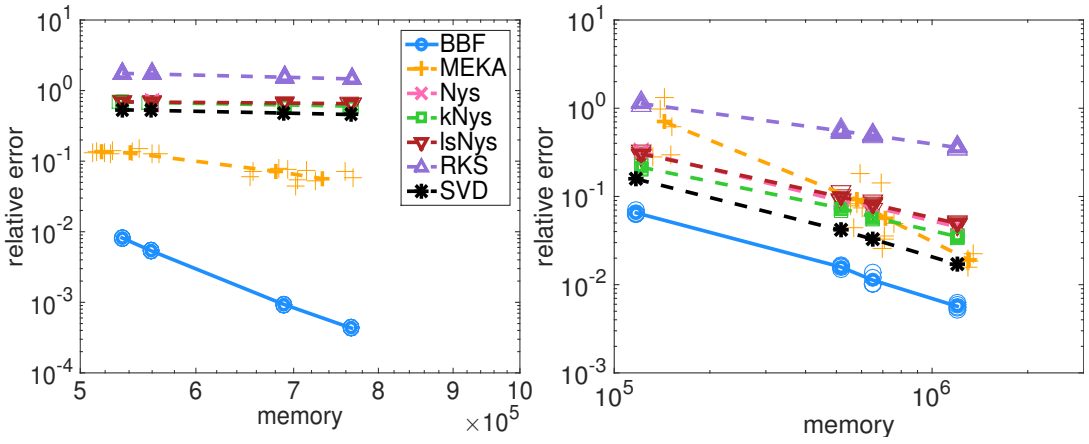

Figure 25: Kernel reconstruction error versus memory footprint (on a loglog scale) for BBF and different state-of-art kernel approximation methods on real datasets.

3.4 MATRIX VECTOR MULTIPLICATION WITH GPUS

This work focuses on harnessing GPUs for performing matrix vector multiplications efficiently (Wong et al., 2015). This type of calculation is critical when solving linear systems arising from PDEs, in particular in reservoir simulations (see TOUGH2, etc).

Recently, graphics processing units (GPUs) have been increasingly leveraged in a variety of scientific computing applications. However, architectural differences between CPUs and GPUs necessitate the development of algorithms that take advantage of GPU hardware. Inversion algorithms and finite-element analysis typically involve sparse matrix vector (SPMV) multiplication operations that can be accelerated by parallelization using GPUs. We developed a new SPMV algorithm and several variations for unstructured finite element meshes on GPUs. The effective bandwidth of current GPU algorithms and the newly proposed algorithms were measured and analyzed for 15 sparse matrices of varying sizes and varying sparsity structures. The effects of optimization and differences between the new GPU algorithm and its variants were then subsequently studied. Lastly, both new and current SPMV GPU algorithms were utilized in the GPU CG solver in GPU finite element simulations.

To validate the newly developed algorithm, we compare our results against parallel PETSc finite element implementation results. The effective bandwidth tests indicated that the new algorithms compare favorably with current algorithms for a wide variety of sparse matrices and can yield very notable benefits. GPU finite element simulation results demonstrate the benefit of using GPUs for finite element analysis and also show that the proposed algorithms can yield speedup factors up to 12-fold for real finite element applications. More details on these results can be found in Wong et al., 2015.

Figure 26 shows the speed-up of various methods compared to the single core CPU. CUSP, CUSPARSE, MGPU are existing GPU libraries. K1 and K2 are two variants of our new GPU algorithm. Jacobi and Block Jacobi are multicore solvers in PETSc (benchmarked with 1 to 4 cores).

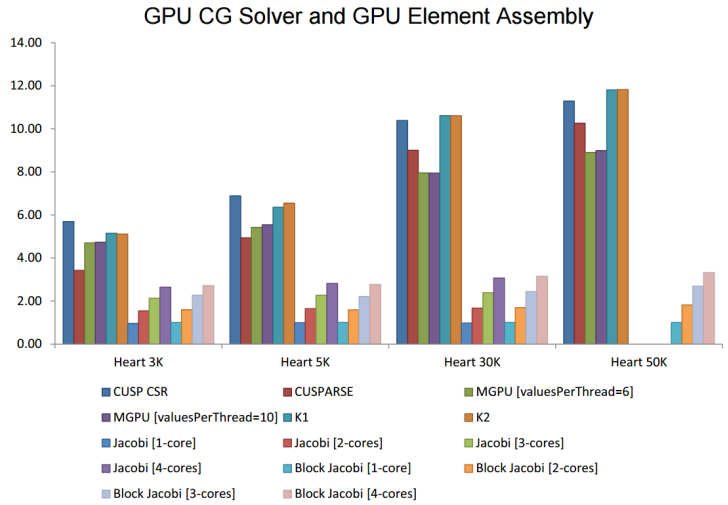

Figure 26: The resulting factor of the increase in speed for the GPU CG solver (with a standard single core CPU-based element assembly routine) and the GPU finite element method.

4 METHOD TESTING FOR 3D SYNTHETIC CCS EXAMPLE

4.1 SYNTHETIC EXAMPLE DESCRIPTION

A synthetic (*i.e.*, simulated on the computer) field scale case was used to demonstrate the applicability of the developed methods for large-scale problems. The synthetic example was generated so that it resembles the Frio-I CO_2 storage pilot experiment in terms of its size and for a domain with strong heterogeneity and high connectivity of large-scale features, and small-scale heterogeneity within the large-scale features. To generate the true permeability field, for the large-scale features the program TProGs (Carle and Fogg, 1997) was used, which is a transition probability based hydrofacies generator, while small-scale heterogeneity was generated using an exponential covariance model. The code for generating realizations requires only a small set of basic input parameters, like a training image for the first type of heterogeneity, and a correlation length and variance for the second type. This code can easily produce a variety of "true" fields. An example of one such "true" field generated, and the one used as true in the synthetic inversions, is shown in Figure 27:

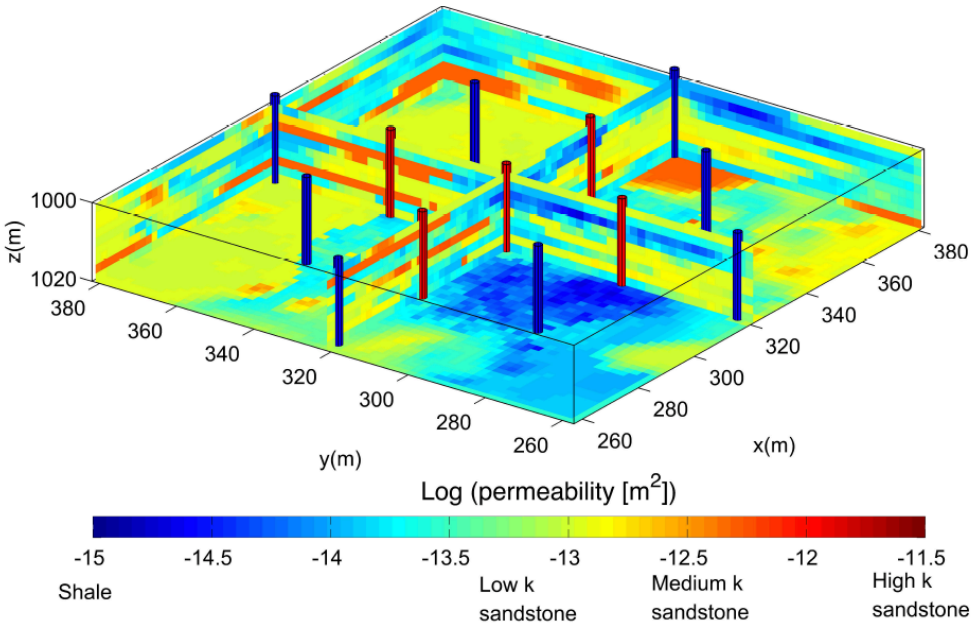

Figure 27: One example of a "true" field generated with TPROGS for large scale features connectivity, combined with small scale heterogeneity.

The domain for the synthetic experiments was chosen to be similar to Test Site one (Frio C Sand formation), in terms of the injection well spacing (30 m) and the thickness of the heterogeneous sand formation (20 m). The entire domain being considered is 640 m x

640 m x 20 m (width x length x thickness). Within the 640 m x 640 m horizontal extent of the domain, we focus on a 128 m x 128 m rectangular domain of heterogeneity. The discretization is chosen to be fine enough to allow a reasonably accurate representation of the multiphase physics (Figure 28). This inversion problem results in a total of *24030* unknown permeability values for each grid block that the inversion algorithm will estimate, providing a high-resolution characterization.

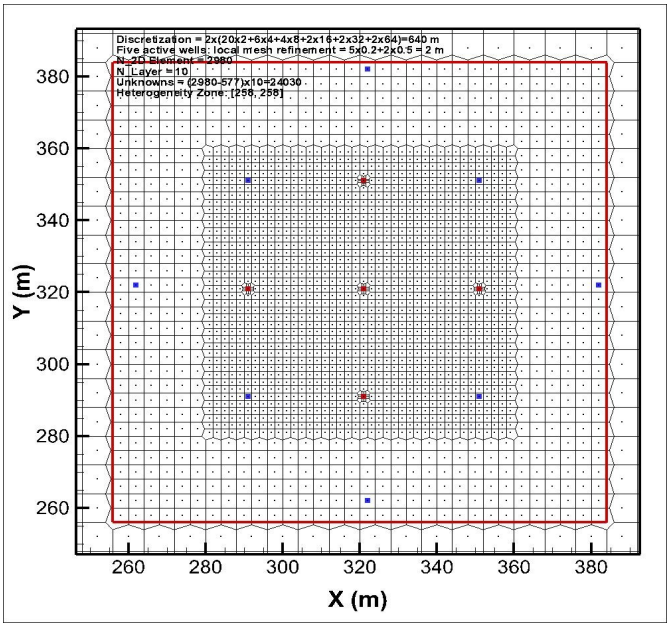

Figure 28: Plan view of the domain and discretization of synthetic geologic carbon storage example. The domain within the red rectangle is considered heterogeneous (the "true" field), and is padded to a large extent by a homogeneous domain to avoid excessive boundary condition effects.

4.2 SYNTHETIC DATA GENERATION

The synthetic experiments performed using the "true" fields to generate the synthetic monitoring dataset include a hydraulic tomography experiment and a hydrothermal tracer test experiment. A monitoring network consisting of eight monitoring wells has been chosen so that it covers sufficiently the domain of unknown permeabilities.

The first phase of inversion includes data from synthetic pumping tests (hydraulic tomography). The pumping tests were designed as a sequence of five 1-day pumping tests, each performed at a different pumping well and followed by 2 days of rest (no pumping), for a total of 15 days. Pressure data collected at all wells during each pumping test can be assimilated in the inversion algorithm to infer the unknown heterogeneous permeability field. In terms of spatial density of monitoring data, a maximum of thirteen wells (active and monitoring wells) can be used for pressure measurements; each well has 10 vertical ports. In total, a maximum of 130 transient pressure measurements (simulated

and contaminated with noise) can be used in the stochastic inversion at each data assimilation step.

Following the pumping tests, a sequence of four dipole hydraulic-thermal tests were conducted. The design of the tests is as follows: a 10-day injection in the central active well with simultaneous pumping at one of the other four active wells are followed by 20 days of pressure temperature decay. Measurements of pressure are taken at all wells, and temperature profiles at the five active wells are used for the inversion; the eight monitoring wells are not used because they have no signal of thermal perturbations.

For the simulation of the pumping tests and thermal tests, we used TOUGH2-MP combined with the single-phase flow module EOS1 (isothermal or non-isothermal flow of water with one or two components). For the sequence of dipole hydraulic-thermal test, non-isothermal flow of single-component water is used (Figure 29).

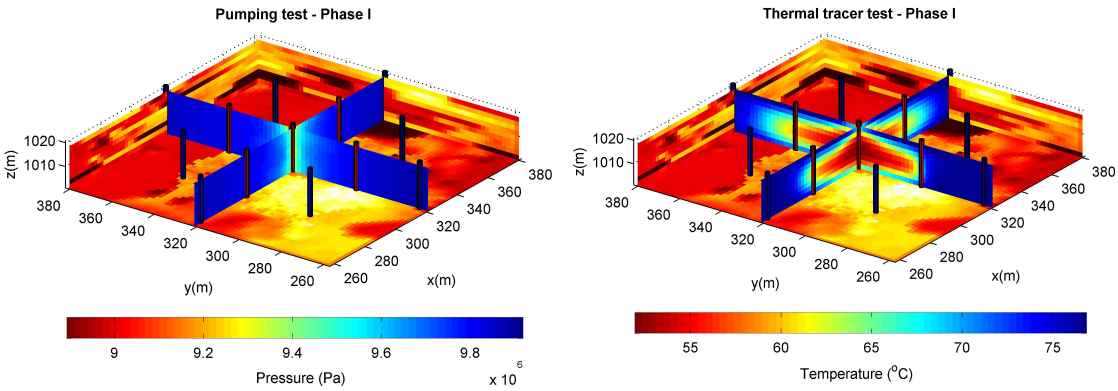

Figure 29: Representative results from forward simulations of synthetic experiments (left) Pumping test showing pressure distribution and (right) Hydrothermal test showing temperature distribution.

4.3 SYNTHETIC INVERSION OF PUMPING DATA

Since we are interested in parameter estimation, for the inversion of the hydro-tracer-thermal and CO_2 injection synthetic datasets we used the smoothing based Compressed State Kalman Smoother (sCSKF). The results showed that sCSKF was able to reconstruct the major features of the permeability field (Figure 30) and provide uncertainty estimates that are informative (Figure 31).

We explored the sensitivity of the results to the number of bases selected and found that 25 bases sufficed for reasonable accuracy and exceptional computational efficiency (each inversion step lasted only 2 hours).

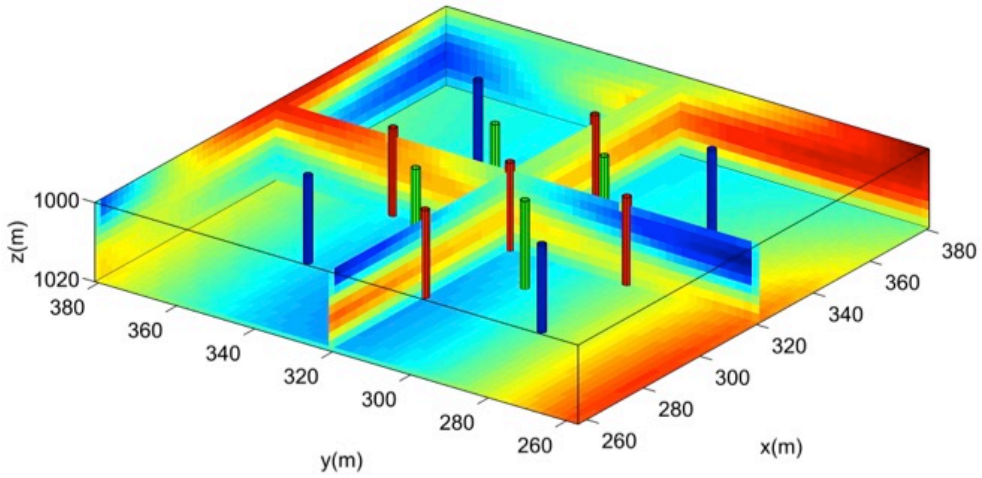

Figure 30: Estimated mean log permeability for synthetic pumping tests.

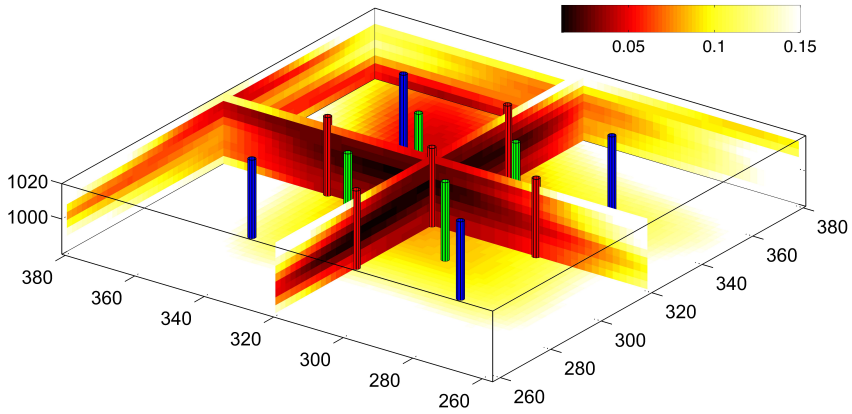

Figure 31: Standard deviation for estimated log permeabilities for synthetic pumping tests.

4.4 SYNTHETIC JOINT HYDROTHERMAL TRACER TEST INVERSION

For the inversion of the hydro-tracer-thermal and CO_2 injection synthetic datasets, measurements of both pressure and temperature were used to estimate the same permeability field using the sCSKF (Figure 32). The results showed a marked reduction in uncertainty of estimation procured by the additional measurements and increased sensitivity of heat transport to the permeability variations in the domain (Figure 33). The added value of temperature was evident in the results, where more detailed features of the permeability were identified compared to the pumping tests.

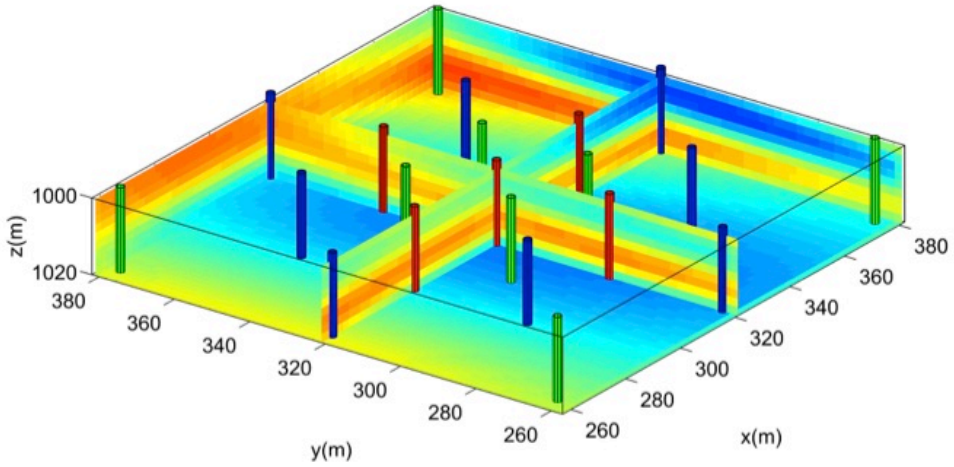

Figure 32: Estimated mean log permeability for synthetic hydrothermal tracer tests

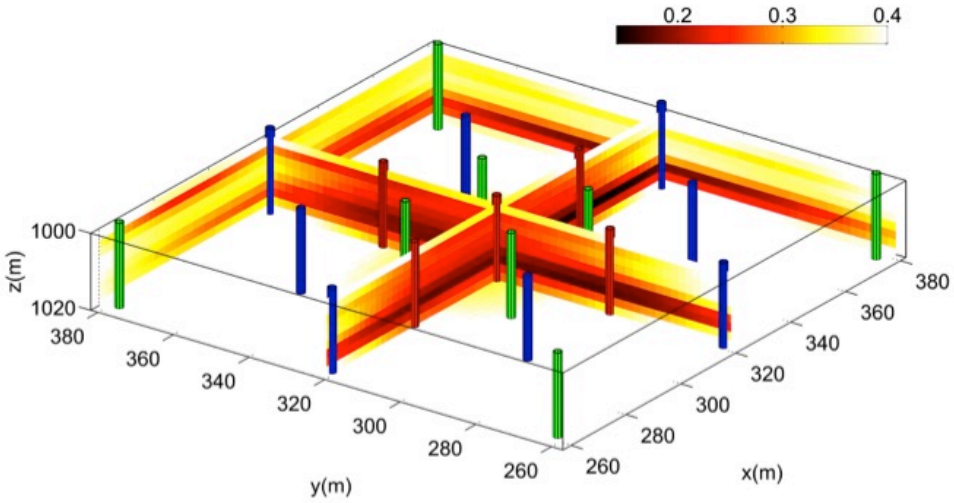

Figure 33: Standard deviation associated with estimated mean log permeability for synthetic hydrothermal tracer test.

In this inversion, the same compression rate of 25 bases was again possible, due to the smoothness of the temperature distribution, such that the inversion was highly computationally efficient with only 25x2 forward simulations being required for assimilation of each dataset.

Application of the sCSKF to the 3D synthetic datasets illustrated that the computational scalability achieved by our algorithms is imperative for data assimilation at a reasonably fine resolution where the number of unknowns exceeds a few thousands. Conventional algorithms would have prohibitive costs that would necessitate significantly reducing the resolution of the forward models and the reliability of the data assimilation method. In particular, application of our algorithm sCSKF to synthetic three-dimensional monitoring data for characterization of a site similar to the Frio I site in Texas, demonstrated that:

- Even for strongly heterogeneous fields with significant anisotropy, large-scale formations of high and low conductivity can be identified even based on a small number of wells, as long as vertically distributed measurements are available.
- Including thermal and chemical tracer data improves estimation accuracy. The relative worth of each dataset can be assessed using the Kalman Gain.
- Uncertainty estimates are extremely valuable as to where more measurements should be obtained to improve estimation and reliability.

5 APPLICATION OF METHODOLOGY TO FIELD TEST SITES

5.1 FRIO-I PILOT TEST

5.1.1 PROBLEM DESCRIPTION

The Frio I pilot test of CO_2 injection was conducted in October 2004 by injecting ~1600 metric tonnes of supercritical CO_2 into the Frio C sand at a depth of ~1500 m. During and after the 10-day CO_2 injection test, the pressure and temperature were continuously monitored at two bottomhole gauges at the injection well and the monitoring well, as well as at their wellheads. Figure 34 shows the regional caprock of Anahuac Shale and the Frio Formation with Frio A, B, and C sands, the injection well (IW) and the monitoring well (MW), as well as the two bottomhole pressure/temperature gauges.

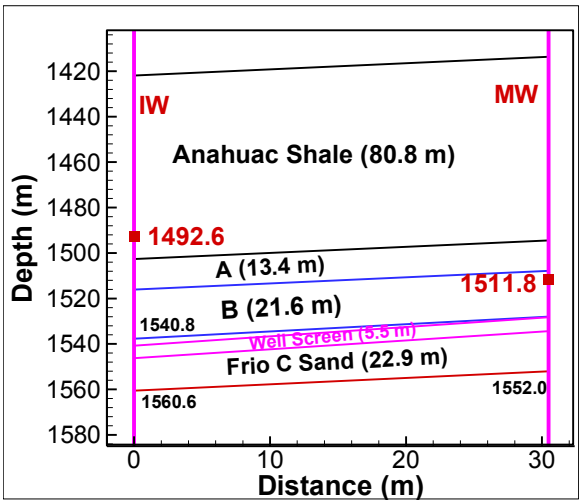

Figure 34: Thickness of the caprock (Anahuac Shale), the storage formation (Frio Sand), and the injection unit (Frio C Sand) in the vicinity of the injection well (IW) and the monitoring well (MW) at the Frio I test site. Also shown are the depth of the pressure/temperature gauges in both wells and formation/unit tops and bottoms.

The vertical profiles of pressure, temperature and CO_2 saturation were obtained three times at the injection well and seven times at the monitoring well. Different tracers were released with injected CO_2 and their concentrations were monitored using the U-tube system installed.

Before the CO_2 injection test, a pumping test was conducted for one day by pumping formation water from the monitoring well, and then a pumping and re-injection dipole test was conducted for 15 days. The pressure and temperature at the two bottomhole

gauges and wellheads were continuously monitored. Significant temperature perturbations were also observed in the temperature profiles obtained before the CO_2 injection test started. A conservative tracer was released with injected water under the quasi-steady-state flow condition induced by the dipole test.

During the experiment, significant temperature perturbations along the wellbores and in the storage formation were observed. Such large perturbations can be seen in the vertical temperature profiles along both wells after the dipole test was completed (Figure 35). More specifically, Figure 35 shows a significant temperature spike at the injection point (far left panel), which at closer look (middle and right panels) was also changing in time, likely in response to the movement of the CO_2 plume. The temperature time series and vertical temperature profiles provide invaluable data for characterizing the Frio formation.

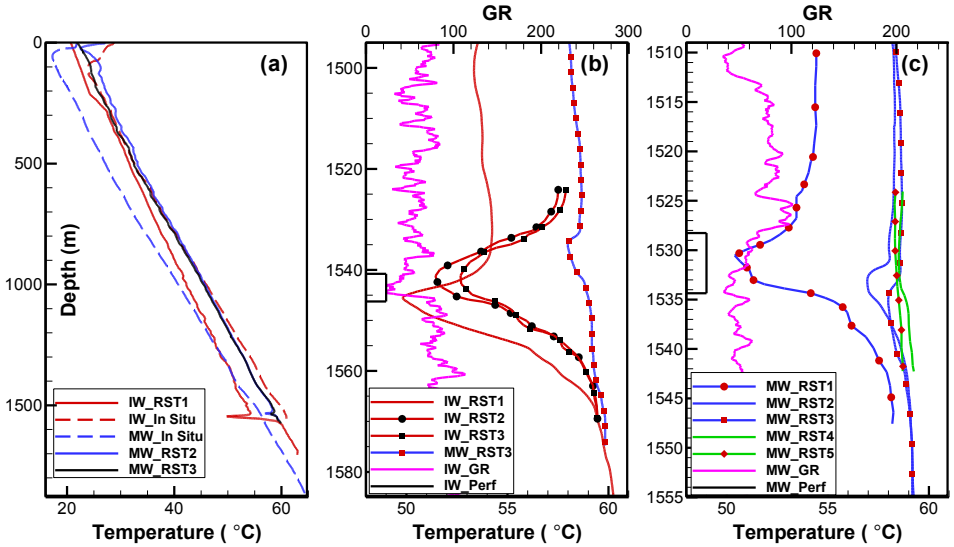

Figure 35: Comparison of vertical profiles of temperature observed at different times (a) at the injection well (IW) and the monitoring well (MW) before and during CO_2 injection, (b) Close up at the injection interval at the injection well after CO_2 injection, and (c) close-up at the monitoring well before, during and after CO_2 injection. Also shown are the Gamma Ray logs (in pink), RST logs, and well perforation intervals (in black rectangles). Note different scale between graphs.

For the CO_2 injection test, Figure 36 shows the vertical profiles of CO_2 saturation for six Reservoir Saturation Tool (RST) logs. We can see that the CO_2 plume is confined in a 4 m thick Frio C sand and the CO_2 saturations as high as 0.88 are observed. Figure 37 shows the breakthrough curves of three tracers released with injected CO_2 at three different times. The CO_2 arrival is quick, 50 hours after CO_2 injection started. Note that injected CO_2 may not have a sufficient temperature signal at the monitoring well because all the monitored temperature perturbations stems from the dipole test conducted before CO_2 injection and decays with time.

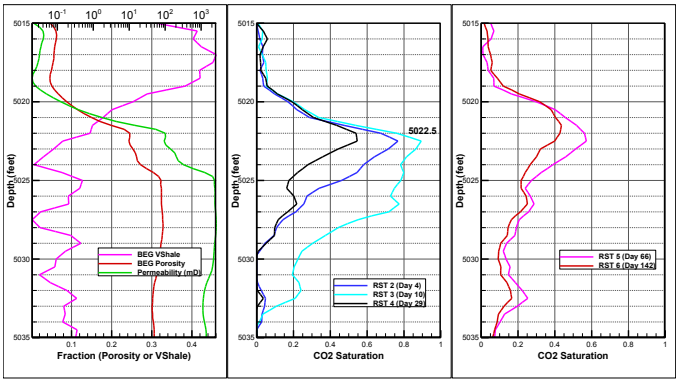

Figure 36: Vertical profiles of CO_2 saturation in the Frio C sand in the monitoring wells estimated using the pulsed neutron logging on October 8 (RST2), October 14 (RST3), November 2 (RST4), December 9, 2004 (RST5), and February 23, 2005 (RST6), as compared to the volumetric shale, porosity and permeability logs.

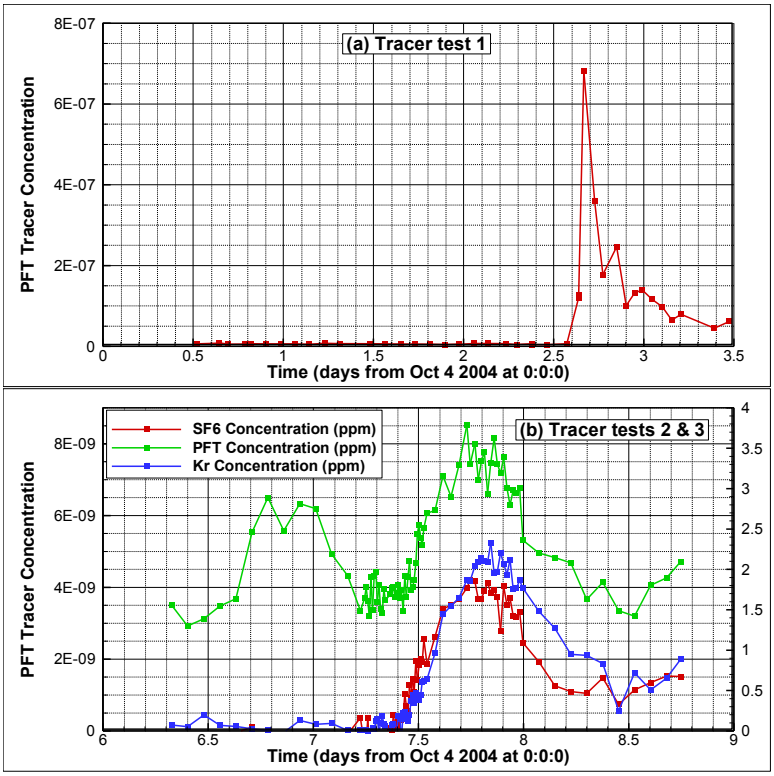

Figure 37: Time-dependent concentrations of PFT in tracer tests 1, 2 and 3, and gas tracer concentrations (SF6 and Kr) in tracer test 3.

The above datasets are available for use in the inversion in order to identify main pathways that CO_2 followed during the Frio I experiment. The objective is to confirm findings from geophysical data that preferential pathways were present that led to a non-uniform distribution of the CO_2 (Figure 38).

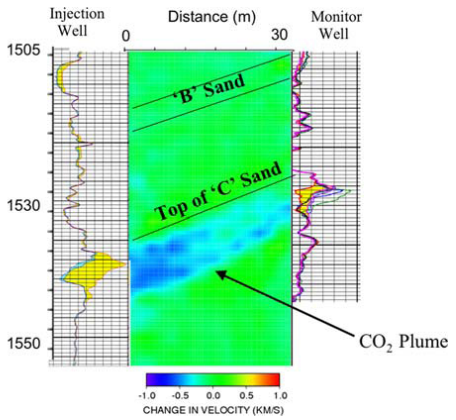

Figure 38: P-wave tomograph with RST logs at wells.

5.1.2 MODEL DEVELOPMENT & INVERSION

Figure 39 shows the geological model in a transformed coordinate system, with Transformed X to the northeastern direction. Two bounding faults form an impervious boundary of the model domain, along with three minor faults acting as flow barriers. Figure 40 shows the plan view of the generated 3D mesh with zoom in the vicinity of the two wells. The 3D mesh is 1000 m by 1000 m in the horizontal direction, and 80 m in thickness including Frio A, B, and C sands. We also included the two wells of 1500 m in length from the ground surface to the top of the model domain for the reservoir. 150 1D elements were used to represent each well. The thermal gradient is 0.0252 °C/m. The heat exchange between the wellbores and their surrounding formations was simulated using a semi-analytical solution implemented in TOUGH2. In total, we have 2072 2D elements and 45584 3D elements, in addition to the 300 well elements.

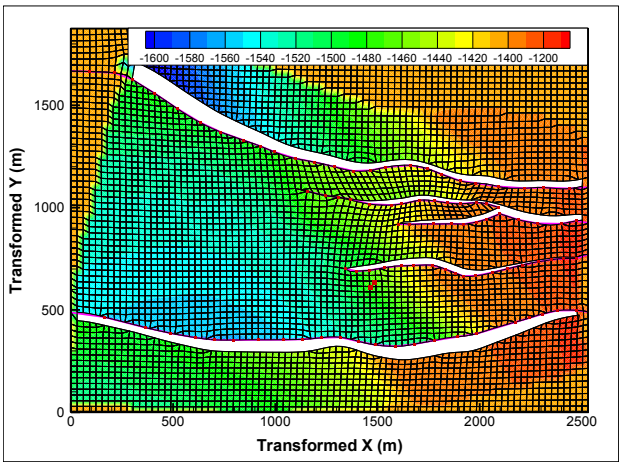

Figure 39: The elevation (in flood contour) of the top of Frio C sand in the transformed coordinate system, the injection and monitoring wells, and the two bounding faults used to define model boundary, and three internal minor faults.

The inversion of this data is performed in two phases: pressure gauge data are used in a calibration mode to obtain wellbore information, while vertical temperature profiles and tracer concentrations are used for characterizing the main preferential pathways in the Frio formation. For this first modeling phase a reduced model is used, shown in Figure 40. The second modeling phase focuses on CO_2 migration prediction for which a second TOUGH2 model is built which includes the complex geology of the Frio-I site, including local faults that constrain the CO_2 plume, as shown in Figure 39. The results are tested for their consistency with available geophysical data that indicate the presence of two major preferential pathways in the region of interest, likely due to high permeability sand channels, between the injection and monitoring well.

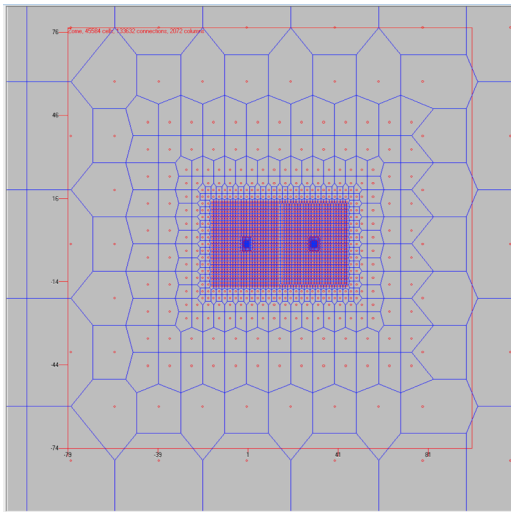

Figure 40: Plan view of the generated 3D mesh, with zoom in the vicinity of the injection well and monitoring well for locally refined mesh.

Work on the inversion of the real Frio dataset is ongoing. A few considerations regarding the inversion of the Frio-II pilot experiment include the following:

a) Inversion and monitoring through data assimilation requires the development of a forward model for the domain and experiments of interest. The bottleneck of the inversion is often the development of such a model.
b) The use of monitoring data often requires judgment calls from an expert. For example, the Frio-II monitoring data for wellhead pressure provide more information regarding the wellbore, rather than the underlying formation. For this purpose, such information can be used in a pre-processing calibration step to obtain a small number of model parameters, in order to simplify the forward modeling and the inversion scheme for the spatially distributed unknowns (e.g. permeability) and make better use of computational resources. For the same purpose, preprocessing of data collected continuously in time may be necessary to reduce the cost of data assimilation.

c) Availability of monitoring data directly related to features of the underlying formation is crucial for heterogeneity characterization. In the Frio case, the vertical temperature profiles provide the best available information for characterization of heterogeneity, while tracer tests are useful for constraining the overall average permeability, but not for delineating fine scale features of the formation.

5.2 In Salah

5.2.1 Problem description

The application of our methods to the In Salah pilot test aims to understand the relation between CO_2 injection and resulting ground surface deformation, measured and monitored by Interferometric Synthetic Aperture Radar (InSAR) data. The InSalah CO_2 project in Algeria, one of the first large scale commercial carbon sequestration projects, injected 3 million tons of CO_2 in three horizontal wells, which resulted in a measureable surface displacement of approximately 5 mm/year (Vasco et al., 2010).

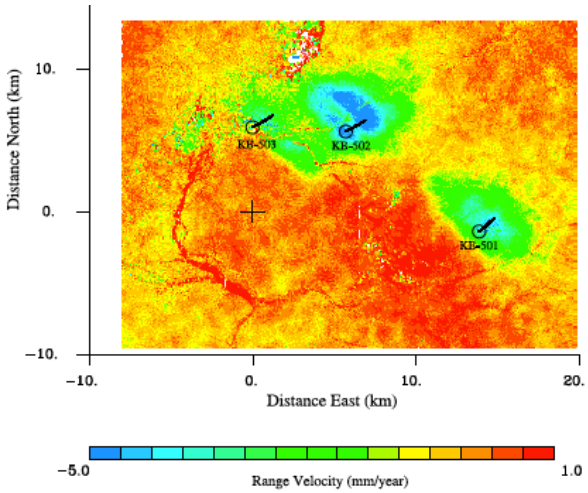

Figure 41: InSAR surface uplifting velocity measurements in mm/year (Vasco et al., 2010)

The focus of our inversion is to identify and characterize the source of the deformation, and specifically to answer the question of, whether the volume increase in the reservoir or a fracture opening or a combination of those disturbances resulted in the observed surface uplifting. This has been a major point of discussion in the literature, in particular, related to the double lobed appearance of the uplifting above well KB 502 (Figure 41).

The advantages of our analysis compared to previous inversions done with the InSAR dataset are the following: Firstly, satellite data collected at 41 different times are used

simultaneously in the same inversion in our dynamic Kalman Filter framework, in order to capture the dynamic expansion of the CO_2 plume in the reservoir, possible CO_2 leakage in the nearby fault, and the expansion or creation of the fault due to the injection as this changed in time. Secondly, we evaluate the above in a statistical framework considering uncertainty in the forward model as well as in the rock hydrological and geomechanical properties and obtain uncertainty estimates associated with the inverted quantities and their variation in space. This allows us to evaluate different scenarios that may have caused the double lobe formation above KB 502 and establish which scenario is more likely, based on its uncertainty.

5.2.2 MODEL DEVELOPMENT & INVERSION

For the InSalah inversion, we employ a combination of the HiKF and SpecKF, in order to facilitate the special nature of the InSAR dataset. The forward model for the evolution of pressure and fracture changes is assumed to be the random walk model, as in HiKF, instead of a hydrogeomechanical forward model. Despite the simplified forward model, this setup of the Kalman Filter allows the consideration of continuity in the physics between the 41 InSAR images, without requiring a costly computational forward model to simulate the pressure and geomechanics of this complex system. The observation model that relates the state (volume and fracture change) to the observations (deformation) is given by a semi- analytical model that is used in a black box fashion to calculate sensitivities, in the same manner done in SpecKF. Combining these two filters reduces greatly the computational cost, and allows characterization at a fine resolution with quantification of uncertainty.

The observation model used for the InSalah field site is a semi-analytical model that simulates the surface deformation resulting from CO_2 injection at the InSalah site. More specifically, the forward model used calculates the co- and post-seismic deformation that is caused by four possible dislocation sources in a viscoelastic-gravitational homogeneous half-space (Wang et al., 2006). The model comprises two calculation steps: One that calculates the Green functions corresponding to a given half space, and one where the dislocations are calculated for the given dislocation source; the latter step is the least computationally expensive and is the only one that needs to be repeatedly run in an inversion. For this reason, the forward model is well suited for evaluating our inversion algorithms.

The observation model development follows previous work by Vasco et al., (2010), where the same InSAR dataset was used in a deterministic inversion to obtain the changes in reservoir volume and fracture aperture that resulted in the observed surface deformation. The model assumes an improved layered earth model that is derived from well log observations, and that has been previously found to provide more physically

consistent results (Vasco et al., 2010). Furthermore, a composite model of both the reservoir and the damage zone (fault) is considered here, as previous work has shown that neither the reservoir nor the damage zone alone can account for the double lobed pattern observed.

Work on the inversion of the InSalah InSAR dataset is ongoing. The InSAR dataset provides an opportunity to tailor our inversion algorithms for spatially distributed datasets, i.e. datasets provided as images like the InSAR data. In such cases, although the number of data points available is significant, it is not necessary that they all provide useful information for the inversion. In addition, the more data are considered the higher the computational cost in terms of sensitivities that need to be calculated and forward model simulations. To reduce the computational cost we examine different ways of compressing the datasets, while maximizing the information that can be obtained from them.

This inversion provides better understanding the behavior of the fault during CO_2 injection at the In Salah site, which is essential in order to conduct forward simulations of the CO_2 injection at the site.

6 CODE DEVELOPMENT AND VALIDATION

6.1 CODE DEVELOPMENT

All code developed for stochastic joint inversion has been developed in Matlab ®. Code for fast linear algebra methods is developed in C++. Fast linear algebra codes area available to the public at:

https://github.com/sivaramambikasaran/BBFMM2D

https://github.com/ruoxi-wang/BBFMM3D

https://github.com/judithyueli/mexBBFMM2D

https://github.com/ruoxi-wang/mexBBFMM3D

6.2 VALIDATION OF DEVELOPED ALGORITHMS

Validation of all developed stochastic inversion algorithms has been conducted and is described in the respective section.

7 SUMMARY & CONCLUSIONS

This project developed a suite of computational methods for accurate and computationally efficient data utilization for characterization and monitoring of large-scale subsurface systems. A summary of these computational tools along with attributes and applicability is listed in the following:

Method	Cost (# simulations)	Attributes
Offline inversion		
Conventional	n (with adjoint methods)	
Fast linear GA (Geostatistical Approach)	none	Any linear problem
	The method combines the Quasi-Linear Geostatistical approach with the hierarchical matrices technique to reduce the cost of matrix-matrix multiplications including the large covariance matrix.	
PCGA (Principal Component Geostatistical Approach)	K*niter	Smooth problems
	The method combines the Quasi-Linear Geostatistical approach with a matrix factorization technique that compresses the error covariance matrix based on its eigenspectrum, thereby reducing the effective number of the unknown quantities.	
Dynamic inversion / Kalman Filter		
Conventional	n (with adjoint methods)	
HiKF (Hierarchical Kalman Filter)	none	Fast data acquisition, linear models
	The method modifies the conventional Kalman Filter by assuming a random walk forward model, thereby reducing the Jacobian calculations and accelerates matrix-matrix multiplications using the hierarchical matrices approach.	
SpecKF (Spectral Kalman Filter)	n	Fast data acquisition, non linear models
	The method modifies the conventional Extended Kalman Filter by employing an approximation to the forward model that allows efficient cross-covariance updates.	
CSKF	K	Smooth problems

(Compressed State Kalman Filter)		
	The method modifies the conventional Extended Kalman Filter by using a low rank approximation of the covariance matrix based on its eigenspectrum and fixed eigenbases.	
sCSKF (Smoothing-based Compressed State Kalman Filter)	2K	Smooth problems, parameter estimation
	The method modifies the conventional Extended Kalman Filter by using a low rank approximation of the covariance matrix based on its eigenspectrum and a fixed basis. For improved parameter estimation and reducing linearization errors, the method employs a one-step ahead smoothing.	

n: number of measurements
m: number of unknowns (typically corresponds to gridded field of unknowns)
K: effective rank of covariance matrix
niter: number of iterations (typically 4-5)

These algorithms were built, tested and demonstrated in the context of Carbon Storage projects in deep subsurface geological formations. Such applications involve diverse datasets and large numbers of unknowns that render the use of traditional inversion methods infeasible. The project addressed the need for new algorithms that use advanced mathematical ideas to reduce the computational cost and provide accurate estimates with reliable quantification of uncertainty at a reasonable cost.

The fast inversion algorithms developed in the project combine Bayesian methods, for stochastic inversion, with fast linear algebra techniques that accelerate matrix-matrix multiplications, the latter being the bottleneck of traditional inversion. Two fast linear algebra techniques that we harness are the fast multipole method, and efficient matrix decomposition techniques, such as SVD and eigenvalue decomposition. With these mathematical tools, alternatives to traditional inversion techniques were created, each of which is best suited for different CCS applications. Parameters that determine which of the developed algorithms is most suitable include the frequency of data acquisition, the type of data used, and the heterogeneity of the system. The major advantages of the algorithms developed are:

- They are fast, i.e., able to process large data sets using modest computer resources.
- They provide reliable uncertainty quantification to inform decision-making.
- Use forward models in a black-box fashion so that the methods are generalizable to any data type for which a forward model exists.

We demonstrated the efficiency of our methods using a large realistic three-dimensional synthetic scenario with datasets similar to those collected at real field sites, and were able to demonstrate that in reasonable time frames, good quality estimates are possible even with limited and diverse datasets, a typical scenario in real CCS projects.

The contributions of this work to the CCS state of knowledge are multifold. With these novel algorithms for inversion that go beyond traditional deterministic, low-resolution inversion, it is possible to better characterize potential CCS field sites, thus being able to better evaluate carbon capacities. It also becomes possible to reliably monitor ongoing operations and prevent undesirable incidents like CO_2 leakage or migration to sensitive areas. With efficient estimation algorithms like those developed in this project, optimization of CCS operations and more informed decision-making will also become possible. Furthermore, these methods can be used to perform data worth analysis as to better design data collection and thereby minimize the overall cost of CCS projects. Combination of these methods to automatic control systems is the logical next step of this work, which will further enhance our ability to implement CCS as a possible climate change mitigation measure.

8 REFERENCES

Agullo, E., Bramas, B., Coulaud, O., Darve, E.F., Messner, M. and Takahashi, T., 2014, Task-based FMM for heterogeneous architectures. *Concurrency Computat.: Pract. Exper.*; doi: 10.1002/cpe.3723.

Ambikasaran, S., Li, J. Y., Kitanidis, P.K., and Darve, E. F., 2013, Large-scale stochastic linear inversion using hierarchical matrices, *Journal of Computational Geosciences*. 17:913–927, doi 10.1007/s10596-013-9364-0

Ambikasaran, S., and Darve, E.F., 2013, An O(N log N) fast direct solver for hierarchical semi-separable matrices, *Journal of Scientific Computing*. 57:477–501, doi 10.1007/s10915-013-9714-z

Aminfar, A., and Darve, E.F., 2014, A Fast Sparse Solver for Finite-Element Matrices, *Parallel Computing Special Issue on PMAA14*. http://arxiv.org/abs/1410.2697

Aminfar, Amirhossein, Sivaram Ambikasaran, and Eric Darve, 2016, A fast block low-rank dense solver with applications to finite-element matrices. *Journal of Computational Physics* 304: 170-188. http://arxiv.org/abs/1403.5337

Carle, F. S. and Fogg, G.E., 1997, Modeling Spatial Variability with one and multidimensional Continuous-Lag Markov Chains, *Math. Geol.,* 29(7)

Coulier, Pieter, Hadi Pouransari, and Eric Darve, 2015, The inverse fast multipole method: using a fast approximate direct solver as a preconditioner for dense linear systems. arXiv preprint arXiv:1508.01835 (2015).

Daley, Thomas M and Ajo-Franklin, Jonathan B and Doughty, Christine, 2011, Constraining the reservoir model of an injected CO2 plume with crosswell CASSM at the Frio-II brine pilot, *International Journal of Greenhouse Gas Control*, 5(4), 1022-1030

Fong, William, and Eric Darve, 2009, The black-box fast multipole method. *Journal of Computational Physics* 228, no. 23: 8712-8725

Ghorbanidehno, H., Kokkinaki, A., Li, Y. J., Darve, E. F. and Kitanidis, P. K., 2015, Real time data assimilation for large-scale systems: the Spectral Kalman Filter. *Advances in Water Resources*, doi: 10.1016/j.advwatres.2015.07.017

Kitanidis, P. K., 2014, Compressed State Kalman Filter for large systems, *Advances in Water Resources*, 76, 120-126, doi: 10.1016/j.advwatres.2014.12.010

Kitanidis, P. K., and Lee, J., 2014, Principal Component Geostatistical Approach for Large-Dimensional Inverse Problems, *Water Resour. Res.*, 50. doi:10.1002/2013WR014630.

Lee, J. and Kitanidis, P.K., 2014, Large-scale hydraulic tomography and joing inversion of head and tracer data using the Principal Component Geostatistical Approach (PCGA), *Water Resour. Res.*, 50. doi:10.1002/2014WR015483.

Li, J. Y., Ambikasaran, S., Darve, E. F. and Kitanidis, P. K., 2014, A Kalman filter powered by H^2 matrices for quasi-continuous data assimilation problems, *Water Resour. Res.*, 50, doi:10.1002/2013WR014607

Li, J. Y., A. Kokkinaki, H. Ghorbanidehno, E. F. Darve, and P. K. Kitanidis, 2015, The compressed state Kalman filter for nonlinear state estimation: Application to large-scale reservoir monitoring, *Water Resour. Res*., 51, 9942–9963, doi:10.1002/2015WR017203.

Liu, X., Zhou, Q., Kitanidis, P.K., and Birkholzer, J. T., 2014, Fast iterative implementation of large-scale nonlinear geostatistical inverse modeling, *Water Resour. Res*., 50, 198–207, doi: 10.1002/2012WR013241.

Pouransari, Hadi, Pieter Coulier, and Eric Darve, 2015, Fast hierarchical solvers for sparse matrices. arXiv preprint arXiv:1510.07363.

Vasco, D. W., A. Rucci, A. Ferretti, F. Novali, R. C. Bissell, P. S. Ringrose, A. S. Mathieson, and I. W. Wright, 2010, Satellite-based measurements of surface deformation reveal fluid flow associated with the geological storage of carbon dioxide, *Geophys. Res. Lett.*, 37, L03303, doi:10.1029/2009GL041544.

Wong, J., Kuhl, E., and Darve, E.F, 2015, A new sparse matrix vector multiplication graphics processing unit algorithm designed for finite element problems. *Intern. Journal for Numerical Methods in Engineering* 102.12 p. 1784-1814, doi: 10.1002/nme.4865

Wang, R., Li, Y., Mahoney, M. W. and E. F. Darve, 2015, Structured Block Basis Factorization for Scalable Kernel Matrix Evaluation. *arXiv preprint arXiv:1505.00398*