



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Final Report: Improved Site Characterization And Storage Prediction Through Stochastic Inversion Of Time-Lapse Geophysical And Geochemical Data

A. Ramirez, W. McNab, Y. Hao, D. White, J.
Johnson

April 18, 2011

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Final Report: Improved Site Characterization And Storage Prediction Through Stochastic Inversion Of Time- Lapse Geophysical And Geochemical Data

A. Ramirez¹, W. McNab¹, Y. Hao¹, D. White², J. Johnson³

1-Lawrence Livermore National Laboratory

2-Geological Survey of Canada

3-Schlumberger Doll Research

LLNL-TR-480694

Abstract

During the last months of this project, our project activities have concentrated on four areas: 1) performing a stochastic inversion of pattern 16 seismic data to deduce reservoir bulk/shear moduli and density; the need for this inversion was not anticipated in the original scope of work, 2) performing a stochastic inversion of pattern 16 seismic data to deduce reservoir porosity and permeability, 3) complete the software needed to perform geochemical inversions and 4) use the software to perform stochastic inversion of aqueous chemistry data to deduce mineral volume fractions. This report builds on work described in progress reports previously submitted (Ramirez et al., 2009, 2010, 2011 -- reports fulfilled the requirements of deliverables D1 – D4) and fulfills deliverable D5: Field-based single-pattern simulations work product.

The main challenge with our stochastic inversion approach is its large computational expense, even for single reservoir patterns. We dedicated a significant level of effort to improve computational efficiency but inversions involving multiple patterns were still intractable by project's end. As a result, we were unable to fulfill Deliverable D6: Field-based multi-pattern simulations work product.

This work has been performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344

Table of Contents

Abstract	1
Introduction	2
Chapter 1: Seismic Results, Pattern 16, Phase 1A	4
1.1 Introduction	4
1.2 Seismic moduli and density inversion	7
1.3 Porosity/permeability inversions	10
1.4 Comparison of Predicted and Observed Waveforms	18
1.5 Computational Expense.....	26
1.6 Summary and conclusions	27
Chapter 2: Development and Application of a Markov Chain Scheme for Inverting the Distribution of Mineralogy and Reaction Rates Along a Flow Path in the Weyburn Reservoir	27
2.1 Introduction	27
2.2 Reservoir Geochemistry Model	29
2.3 Geochemistry Inversion Using the Markov Chain Monte Carlo Algorithm: Development and Testing Using Synthetic data	33
2.4 Inversions using Brine Chemistry Data from Pattern 16	42
2.5 Conclusions	52
Summary and Conclusions	54
Implications for Best Practices Manual	55
Acknowledgements	55
References	56
Appendix	57

Introduction

Our stochastic inversion approach uses reactive transport modeling, facies-based geostatistical methods, and a novel stochastic inversion technique to optimize agreement between observed and predicted storage performance. Such optimization is accomplished through stepwise refinement of: 1) the reservoir model—principally its permeability magnitude and heterogeneity—and 2) geochemical parameters—primarily key mineral volume fractions and kinetic data. We suggest that these refinements may lead to significantly improved history matching and forward modeling of CO₂ storage. Our tool uses the Markov Chain Monte Carlo (MCMC) methodology.

Deliverable D1, previously submitted as a report titled “Development of a Stochastic Inversion Tool To Optimize Agreement Between The Observed And Predicted Seismic Response To CO₂ Injection/Migration in the Weyburn-Midale Project” (Ramirez et al., 2009), described the stochastic inversion approach that will identify reservoir models that optimize agreement between the observed and predicted seismic response. The software that implements this approach has been completed, tested, and used to process seismic data from pattern 16. A previously submitted report titled “Model verification: synthetic single pattern simulations using seismic reflection data”, Ramirez et al. 2010, partially fulfilled deliverable D3 by summarizing verification activities that evaluate the performance of the seismic software and its ability to recover reservoir model permeabilities using synthetic seismic reflection data. Ramirez et al. 2011, partially fulfilled deliverable D4 requirements by summarizing model development activities required for application of TProGS software and the use of TProGS within the MCMC tool. We were unable to apply the TProGS tool to the seismic and geochemical data because we had to concentrate project resources on unexpected challenges presented by the MCMC inversions.

The D5 deliverable requires a summary of the results of applying NUFT/MCMC to refine the reservoir model and geochemical parameters by optimizing observation/prediction agreement for the seismic/geochemical response to CO₂ injection/migration within a single pattern of Phase 1A/1B. This document includes a chapter that shows and discusses permeability models produced by seismic inversion that used seismic data from pattern 16 in Phase 1A. It also presents geochemical inversions of mineral volume fractions that use aqueous chemistry data from pattern 16 and presents the results of verification tests using synthetic data. It fulfills the requirements of deliverable D5 and completes the requirements associated with deliverable D3 by verifying the performance software of the geochemistry software and its ability to recover mineral volume fractions.

The main challenge with our stochastic inversion approach is its large computational expense, even for single reservoir patterns. We dedicated a significant level of effort to improve computational efficiency but inversions involving multiple patterns were still intractable by project’s end. As a result, we were unable to fulfill Deliverable D6: Field-based multi-pattern simulations work product.

We had to significantly modify our original plans for geochemical inversions. Originally, we planned to use a 3D reactive transport simulator (NUFT) as the centerpiece of our geochemical inversion approach. The computational expense of this approach proved to be intractable. The 3D simulator requires about 1.75 hours and 5 hours (per model) to simulate 1.3 years and 9 years of CO₂ injection, respectively. We expect that our MCMC approach will require a few thousand iterations to converge (each iteration consists of one 3D reactive transport calculation). The estimated run-time is 145 days assuming that 2000 MCMC iterations are required and that 1.3 years of CO₂ injection are simulated. For 9 years of injection, the estimated run-time is 416 days. Realistically, we believe that the runtime could be reduced by a factor of 2 (72 days, 1.3 years of CO₂ injection), but these run-times would still be too long for practical use. This situation required that we modify our original plans by considering alternative inversion approaches.

A second challenge is that the mineral volume fraction and kinetic rate constants are spatially heterogeneous and can show significant variability on the scale of meters to tens of meters. Ideally, an inversion approach models this level of variability by allowing changes at each node in the calculation grid (our pattern 16 grid has 26896 nodes). However, the amount of aqueous chemistry data is quite limited even at the Weyburn project where an unprecedented water sampling program has collected hundreds of water samples from multiple wells over a period of about 10 years. Near pattern 16 there are several water-sampling wells where multiple aqueous species are measured for each sampling episode. In inversion parlance, this means that the ratio of knowns (measurements) to unknowns (geochemical properties at each node, i.e., one rate constant for each reaction of interest, and one volume fraction for each mineral of interest) is extremely poor. The sparsity of information will produce solutions to the inverse problem that are highly non-unique, i.e., there will be a very large (possibly infinite) number of solutions that fit the data equally well. This means that it will be impossible to uniquely determine the 3D distribution of the mineral volume fractions and rate constants in the reservoir because there is insufficient enough information to do so.

To mitigate the computational expense and data sparseness problems, we developed and tested a 1D inverse approach that uses a faster reactive transport simulator (PHREEQC). PHREEQC can solve the reactive transport problem faster than NUFT because it makes a few simplifying assumptions, namely that there is only one fluid phase (brine) instead of the 3 phases assumed by NUFT (brine, oil, free phase CO₂). The 1D flow-tube approach simplifies the inversion problem into a single flow path that connects the CO₂ injection well to a water-sampling point. The reduction of dimensionality from 3D to 1D and the faster simulator reduce the runtime to about 2 days for 2000 MCMC iterations. We used the flow-tube approach to invert for reservoir mineral volume fractions in pattern 16 using the aqueous chemistry data provided by Maurice Shevalier and others at the University of Calgary. Chapter 2 will describe the 1D approach and inversion results in detail.

Chapter 1: Seismic Results, Pattern 16, Phase 1A

1.1 Introduction

This chapter will describe a stochastic inversion approach of seismic properties. We have performed several permeability stochastic inversions using seismic data from pattern 16, Phase 1A area. Here we present the results of one of those inversions. The permeability inversions shown in this chapter assumed well log and laboratory values of seismic properties.

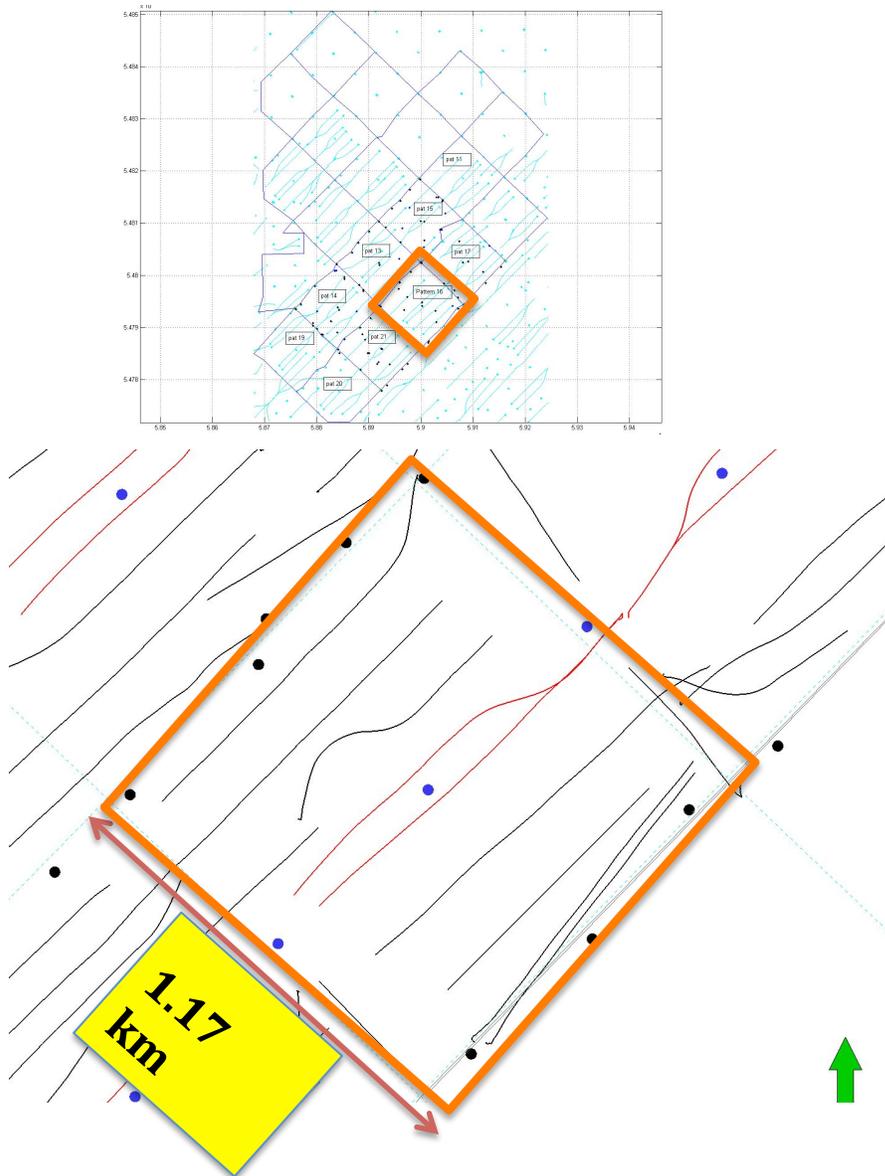


Figure 1.1. The orange square shows the location of pattern 16 where seismic data used by our inversion was collected. This pattern is within the Phase 1A area of the Weyburn-Midale reservoir. Borehole layout provided by Barbara Dietiker (Geological Survey of Canada). The red lines indicate the location of the CO₂ injector, the blue circles represent the WAG injectors and the black lines and circles represent the oil producers.

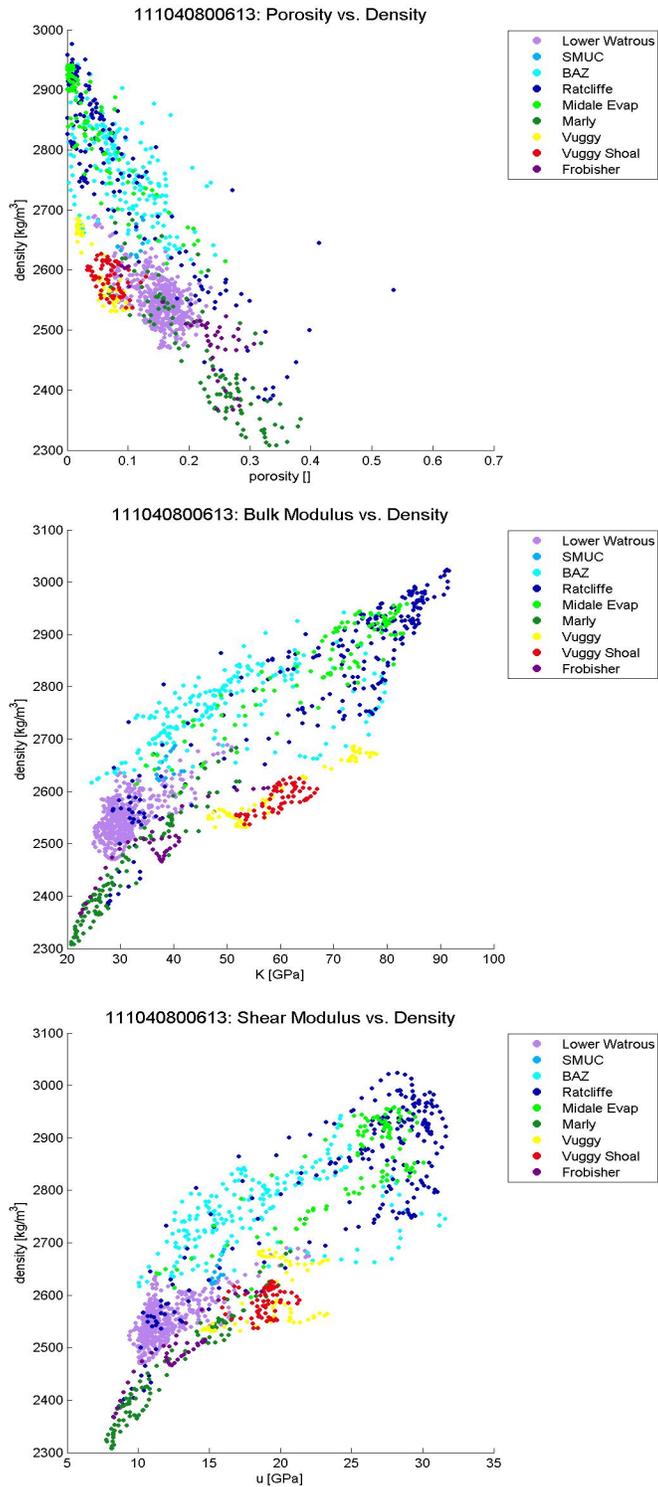


Figure 1.2 shows well-logs used to constrain the moduli inversions. The plots show the correlation between density (vertical axis, kg/m^3) and porosity (top), density and bulk modulus (middle plot, horizontal axis, GPa), and density and shear modulus (bottom plot, horizontal axis, GPa). Barbara Dietiker (Geological Survey of Canada) provided the well log plots.

1.2 Seismic moduli and density inversion

In a previous report (Ramirez et al., 2011) we identified the need to obtain estimates of seismic properties (bulk and shear moduli, density) that honored their spatial variability, and proposed to develop a new inversion approach to obtain them using the pre-CO₂ injection seismic data. The shape of the waveforms is very sensitive to the vertical impedance contrast caused by moduli variability; waveform shape is much less sensitive to impedance changes caused by CO₂ invasion of the pore space. Our original plan had been to use homogeneous values for each layer based on laboratory (Brown, 2002) and well log measurements but this approach proved to be inadequate. The need for this new inversion was discovered late in the project and was not included in the original scope of work.

A stochastic inversion approach that uses empirical petrophysical relationships derived from well logs as constraints was developed. It is a straightforward modification of the inversion approach used for permeability inversion and described in detail by Ramirez et al, 2009. The process starts by proposing realizations of the porosity/permeability fields. These realizations honor the porosity/permeability distribution present in the Cenovus “calibrated” model (see Figure 1.7, Ramirez et al, 2011). Petrophysical relationships based on well log measurements were used to map porosity to density and density to bulk and shear moduli. Figure 1.2 shows cross-plots of density, bulk and shear moduli determined from well logs (provided by Barbara Dietiker, Geological Survey of Canada). Tables 1 – 3 show the equations used to relate density to porosity, bulk modulus and shear modulus. We then used the NUFT flow simulator to predict the pressure distribution within the reservoir during water injection (before CO₂ injection). Predicted waveforms were then calculated (see Ramirez et al., 2009 for details). The predicted waveforms were compared to pre-CO₂ injection waveforms (data collected in 1999) and likelihood values were calculated. We then used the MCMC stochastic inversion technique (Ramirez et al., 2009) to find spatially-variable models of bulk/shear moduli and density that were most consistent with the data.

We were unable to produce acceptable moduli inversions with this method. The likelihoods calculated (proportional to data misfit) turned progressively worse (likelihood becomes increasingly negative) during the inversion (see Figure 1.3), the opposite of the expected behavior. We spent several days looking for software bugs (one minor bug found and fixed) and later attempted a few modifications to the approach; none of the changes improved the outcome. Note that the curves in Figure 1.3 show that the likelihood sometimes slightly improves (becomes more positive) from iteration to iteration but more frequently slightly degrades. After many iterations, this tendency to degrade tends to dominate the search.

We do not understand the root cause of this behavior but can offer one possible explanation. The realizations of moduli and density are chosen completely at random. The number of realizations that are likely to fit the data poorly is vastly larger than the number of realizations that are likely to fit the data well. The acceptance criterion used by

the algorithm (Metropolis-Hasting, MH) that decides whether to accept or reject a new proposal is probabilistic and behaves as follows (see Ramirez et al., 2009 for further details). MH always accepts proposals that improve the likelihood (proportional to misfit) while sometimes accepting proposals that degrade the likelihood particularly if the likelihood for the new proposal is only slightly worse. This behavior insures that the search does not get “stuck” in local maxima. The odds of proposing a realization that slightly degrades the likelihood are larger than the odds of slightly improving the likelihood and it is likely that some of these will be accepted. This tends to point the search away from solutions that improve the likelihood.

We believe that this problem is solvable but the project ended before we were able to develop an algorithm that yields acceptable results. One possible solution is to alter the approach such that the odds of proposing realizations likely to increase the likelihood improve substantially. This can be achieved by adding new information such as a constraint on spatial variability that is based on well log variability. Such an analysis would inform the likely impedance contrast range between each layer pair, thereby providing powerful new constraints to the stochastic inversion. There are stochastic inversion schemes of seismic impedance using waveform and traveltime data reported in the literature (e.g., Quan and Harris, 2008). They use an ensemble Kalman filter (EKF) approach that estimates the impedance spatial variability from well log values. This approach proved successful when using synthetic data. An alternative solution would make use of well-known, deterministic impedance inversion techniques (D. White, personal communication); we would use the inverted impedances directly to compute the predicted waveforms. A third alternative would be a modification of the stochastic approach we used, this time incorporating the spatial variability trends present in the well logs. We would use TProGS (described in Ramirez et al., 2011) to model these trends, and it would produce realizations of moduli/density that exhibit realistic spatial correlation (or variability). A fourth alternative would change the way we use the waveforms: the observed and predicted time-lapse difference waveforms could be compared rather than direct comparison of the observed and predicted post-injection waveforms. This would partly remove the dependence on the pre-injection seismic properties model.

Lithology unit	Equation used to map from porosity to density (kg/m ³)
Ratcliffe	density = 2900.0 + poro*-1818.18 + uniform(-100.0, 100.0)
Evaporite	Porosity = ((density - 2850)/-1818.18) + uniform(-0.06, 0.06)
Marly	Porosity = ((density - 2850)/-1818.18) + uniform(-0.06, 0.06)
Vuggy	Porosity = ((density - 2700)/-1818.18) + uniform(-0.06, 0.06)
Frobisher	Porosity = ((density - 2900)/-1818.18) + uniform(-0.06, 0.06)

Table 1 shows the linear relationships used by the seismic property inversion to remap density realizations to porosity realizations. These relationships are based on well log measurements of density and porosity shown in Figure 1.2. The term “uniform(-0.06,

0.06)” means that random number is sampled from a uniform distribution on the range -0.6 to 0.6.

Lithology unit	Equation used to map from density (kg/m ³) to bulk modulus (Pa)
Ratcliffe	$Bulk_mod = (3.0E-08 * density^{5.3155}) + uniform(-1.5E+10, 1.5E+10)$
Evaporite	$Bulk_mod = (1.0E-09 * density^{5.7370}) + uniform(-1.0E+10, 1.0E+10)$
Marly	$Bulk_mod = (2.0E-08 * density^{5.3453}) + uniform(-6.0E+09, 4.0E+09)$
Vuggy	$Bulk_mod = (2.4E-11 * density^{6.2289}) + uniform(-1.6E+10, 1.4E+10)$
Frobisher	$Bulk_mod = (8.0E-08 * density^{5.1969}) + uniform(-1.1E+10, 9.0E+09)$

Table 2 shows the linear relationships used by the seismic property inversion to remap density realizations to bulk modulus realizations. These relationships are based on well log measurements of density and bulk modulus shown in Figure 1.2. The term “uniform(-6.0E+09, 4.0E+09)” means that random number is sampled from a uniform distribution on the range -6.0E+09 to 4.0E+09. Note that some of the uniform distributions are slightly biased towards negatives values in order to reduce the bulk modulus in layers where fractures are likely to be present.

Lithology unit	Equation used to map from density (kg/m ³) to shear modulus (Pa)
Ratcliffe	$Shear_mod = (5.0E-09 * density^{5.4204}) + uniform(-1.0E+10, 1.0E+10)$
Evaporite	$Shear_mod = (2.0E-06 * density^{4.6575}) + uniform(-5.0E+09, 5.0E+09)$
Marly	$Shear_mod = (1.0E-11 * density^{6.1800}) + uniform(-3.5E+09, 2.5E+09)$
Vuggy	$Shear_mod = (3.3E-02 * density^{3.4213}) + uniform(-6.0E+09, 4.0E+09)$
Frobisher	$Shear_mod = (1.0E-07 * density^{5.0142}) + uniform(-6.0E+09, 4.0E+09)$

Table 3 shows the linear relationships used by the seismic property inversion to remap density realizations to shear modulus realizations. These relationships are based on well log measurements of density and shear modulus shown in Figure 1.2. The term “uniform(-6.0E+09, 4.0E+09)” means that random number is sampled from a uniform distribution on the range -6.0E+09 to 4.0E+09. Note that some of the uniform distributions are slightly biased towards negatives values in order to reduce the shear modulus in layers where fractures are likely to be present.

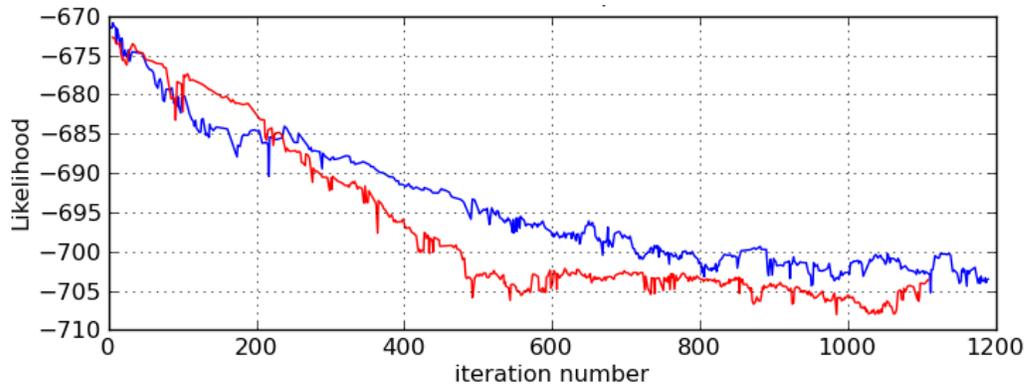


Figure 1.3. The plot shows the unexpected degradation (likelihood becomes increasingly negative) in likelihood observed during the stochastic inversion for bulk/shear moduli and density using seismic data.

1.3 Porosity/permeability inversions

We used pattern 16 seismic data to perform several permeability porosity inversions. We were forced to use bulk and shear moduli values from core and from well logs measurements because the moduli/density inversion described previously did not produce acceptable results. The seismic inversions described below used homogeneous moduli values (one for each layer) that were adjusted by trial and error to reduce the waveform misfit as much as possible (see Ramirez et al, 2010, for details).

The inversion approach we followed consisted of the following steps; for additional details of our MCMC stochastic inversion technique, see Ramirez et al. (2009, 2010). First, we generated random realizations of porosity/permeability. A smoothing filter is applied to these realizations to reduce the gradients in porosity/permeability, thereby adding another constraint to the stochastic inversion. In this document, we call these realizations reservoir models. These models honored geostatistical trends in Cenovus' model (provided by Barbara Dietiker, Geol. Survey of Canada) calibrated against several decades of production data; these trends identified the statistical distributions of porosity and permeability in the reservoir layer, and their cross-correlation. The realizations consist of porosity and permeability fields that vary spatially within each layer and honor the porosity/permeability trends embedded in Cenovus' calibrated model (previously shown in Ramirez et al., 2010, and reproduced in this document's appendix as Figures A.1 and A.2).

Each realization also honored lithology designations and layer boundaries in Cenovus' model. The process that produces one realization is initialized with the porosity/permeability values of the realization that was last accepted by the MCMC process. Then, we randomly select a small subset ($\sim 0.5\%$) of grid nodes whose porosity and permeability values will be changed. A new porosity is randomly assigned to each chosen node by sampling from the distribution in Figure A.1; nodes within the Vuggy layer are populated with porosities that honor the left mode in A.1 while nodes within the

Marly layer are populated with porosities sampled from the right mode (a similar process is used to assign values to nodes within the Ratcliffe, Evaporite and Frobisher layers, each layer honoring a different porosity distribution).

Once porosities are assigned we use the porosity-permeability cross-correlation shown in Figure A.2 to assign the corresponding permeability values. Figure A.2 shows that each porosity value is associated with a range of possible permeabilities. We randomly select a permeability value from this range and assign it to the node being processed. After all chosen nodes have been modified, we apply a smoothing filter to the porosity and permeability fields.

The next step consisted of running the flow simulator. The simulator injected CO₂ into the reservoir at the same rate used in the field. The production rate was set such that the pressure around the producers remained at constant, ambient pressure. The flow simulation predicted various reservoir parameters such as fluid densities, CO₂ saturation and pore pressure for each realization. We then predicted seismic velocities throughout the model using the calculated reservoir parameters and Gassmann's equation. The velocity model was then used to compute seismic reflectivities and zero-offset, 1D seismograms. The algorithm then compared the predicted and observed seismograms, calculated the likelihood function (proportional to data misfit), and used the likelihood value to decide whether to accept or reject the current realization. Our MCMC stochastic inversion technique will find those permeability models that best fit the seismic data and the "prior" constraints (lithology boundaries and geostatistical trends).

The flow simulation assumed that water was injected for 2 years, and that CO₂ injection started after 0.7 years of water injection; the flow simulator modeled reservoir conditions after 1.3 years of CO₂ injection. The seismic data used for the inversion were collected circa December 2001.

The Evaporite layer plays a significant role as an impermeable boundary above the reservoir. The vertical resolution of the flow simulation grid is 4.3 m. The thickness of the Evaporite layer in pattern 16 was sometimes significantly less than this resolution, thereby artificially disappearing from some parts of the grid. In order to preserve the sealing properties above the Marly layer, we set the permeabilities of the Ratcliffe to the same low values associated with the Evaporite. This approach insures that there is a continuous low permeability layer above the reservoir.

Figure 1.4 shows likelihood value as a function of MCMC iteration for two seismic runs. Each plot has red and blue curves representing separate Markov chains that start in different places and independently sample the posterior distribution (solution space). Each MCMC iteration involves one porosity/permeability realization, one flow calculation and one likelihood calculation (i.e., comparison of predicted and observed waveforms).

The step size associated with the top plot in Figure 1.4 was about ten times larger (on average) than the step size for the bottom plot. Both plots show that the likelihood (proportional to misfit) is improving (moving closer to 0). The larger step-size run reaches a stable likelihood value (iteration 5700) that is significantly better than the corresponding small step size value. However, the larger step size also produced an undesirable result: the blue curve stops around iteration 800, indicating that the Markov chain is trapped in a local likelihood maximum. The smaller step-size run shows both chains progressing normally and reaching relatively stable values after iteration 4600, thereby indicating that both chains are producing models that fit the data about the same. This means that both Markov chains are sampling models that fit the data equally well and that any further improvement in the likelihood is likely to be small or negligible. In statistical parlance, this means that the seismic inversion has converged after 4600 iterations and that the Markov chains are now sampling from the posterior distribution, the collection of models that contains the solution to our stochastic inversion problem. Thus, we can use models produced after iteration 5000 to obtain a statistically reliable estimate of the heterogeneous reservoir permeabilities.

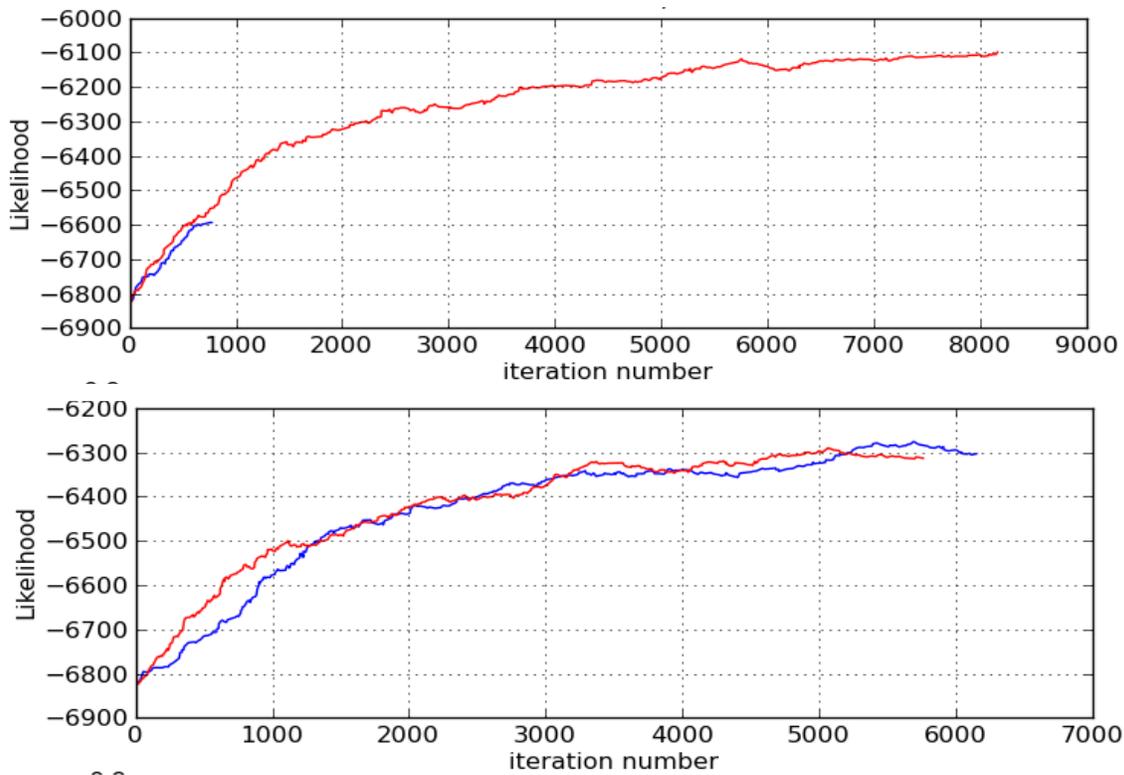


Figure 1.4. The plots show likelihood value as a function of iteration for two permeability inversion runs that used different step sizes. The step size associated with the top plot was about ten times larger (on average) than the step size for the bottom plot. The top plot indicates that a solution with better likelihood (closer to 0) has been found using the larger step size. Likelihood is a measure of the similarity between the predicted and

observed seismic waveforms. Note that a stable value is achieved after about 5700 iterations (top plot) and 5000 iterations (bottom plot).

Ideally, we would have liked to see the larger step-size run behave in a similar fashion but were unable to do so before the project ended. Considering all the factors, we have chosen to analyze the results of the larger step size run because of its better likelihood values and because it reached relatively stable values after iteration 5700.

We now look at permeability models from the posterior distribution. The images in Figure 1.5 show horizontal slices through one of the permeability models that best fit the seismic data (we chose iteration 5751 because it has a likelihood value near maximum). The horizontal slices show the lithologies (1st column), most likely permeability (2nd column) and most likely porosity (3rd column). The corresponding CO₂ saturations and pore pressures are shown in the 4th and 5th columns. The horizontal slices are located 8 m above, 4 m above, 0 m, and 4 m below the location of the CO₂ injector; the injector depth at this location is approximately 1428 m. The permeabilities are plotted using a logarithmic color scale.

Note that the slice located 4m above the injector consists mostly of the Marly unit. In the slice at the injector depth, the upper half of the slice intersects the Vuggy layer and the lower half the Marly layer. The slice located 4m below the injector mainly intersects the Vuggy layer.

The permeabilities of the Marly and Vuggy layers range from about 40 - 45 to 150 – 160 millidarcies. The average permeability of the Vuggy layer is larger than the average Marly permeability, as expected. The Marly and Vuggy porosities range from 6% to 20%. Also, the average porosity of the Marly layer is larger than the average Vuggy porosity.

We next examine the full ensemble of models in the posterior distribution. Figure 1.6 shows the mean (2nd column) and standard deviation (3rd column) of all the permeability models in the posterior distribution, i.e., models with iteration number greater than 5700. Comparing the 2nd column of images in Figures 1.5 and 1.6, we can see that the permeability images are very similar. This suggests that most of the models in the posterior are similar to one another. The variability between models can be evaluated by looking at the standard deviation images, 3rd column of images in Figure 1.6. The standard deviations in the Marly layer range from about 4 to 18 md. The standard deviations in the Vuggy layer range from about 6 to 24 md.

One question we might ask is how much influence does the seismic data exert on the stochastic inversion results. To answer this question, we ran our algorithm without using the seismic data. This run produced realizations that only honored the prior constraints such as the statistical trends in Cenovus calibrated model: histograms of porosity and cross-correlations between porosity and permeability described in Ramirez et al., 2010 and 2011, lithology designations and layer boundaries. Approximately 1100 porosity and permeability models produced by this run were averaged and standard deviations were calculated. The resulting models are presented in Figure 1.7.

A comparison of the average and standard deviations in Figures 1.6 (inversion guided by seismic data) and 1.7 (only used prior constraints, seismic data not used) suggests that the seismic data is influencing the inversion results. For example, consider the images corresponding to the Marly and Vuggy (bottom three rows of images in the two figures). The seismic inversion image shows a wider range of permeabilities and a smaller range of standard deviations than the non-inversion run. This behavior and the likelihood plot in Figure 1.4 suggest that the seismic data is helping guide the inversion by selecting a subset of the realizations that reduces the data misfit. The seismic data determined the locations of the permeability highs and lows within each layer.

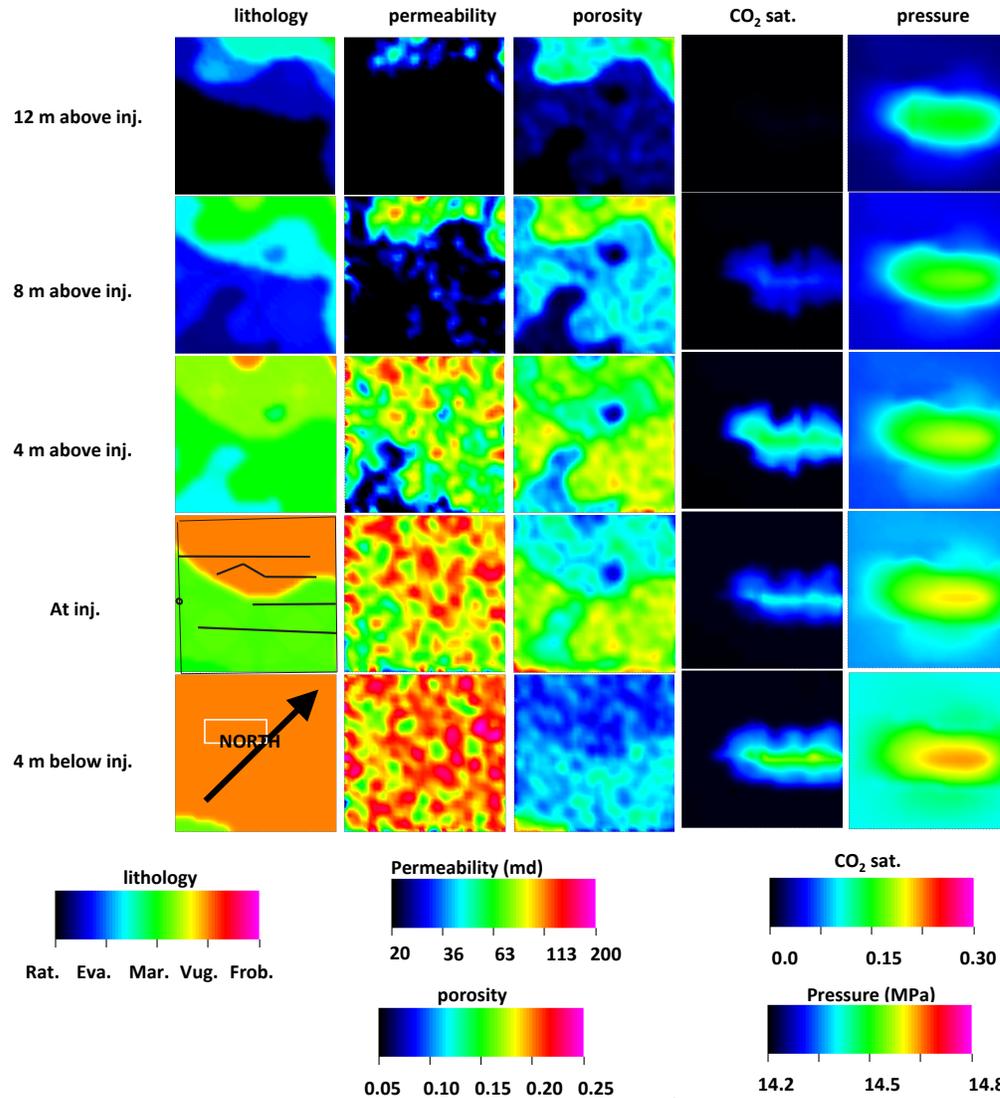


Figure 1.5. The horizontal slices show lithologies (1st column), most likely permeability (2nd column) and most likely porosity (3rd column). CO₂ saturation and pressure are shown in the 4th and 5th columns. The most likely model exhibits the smallest misfit between the observed and calculated seismograms. The permeabilities are shown in millidarcies, on a logarithmic scale.

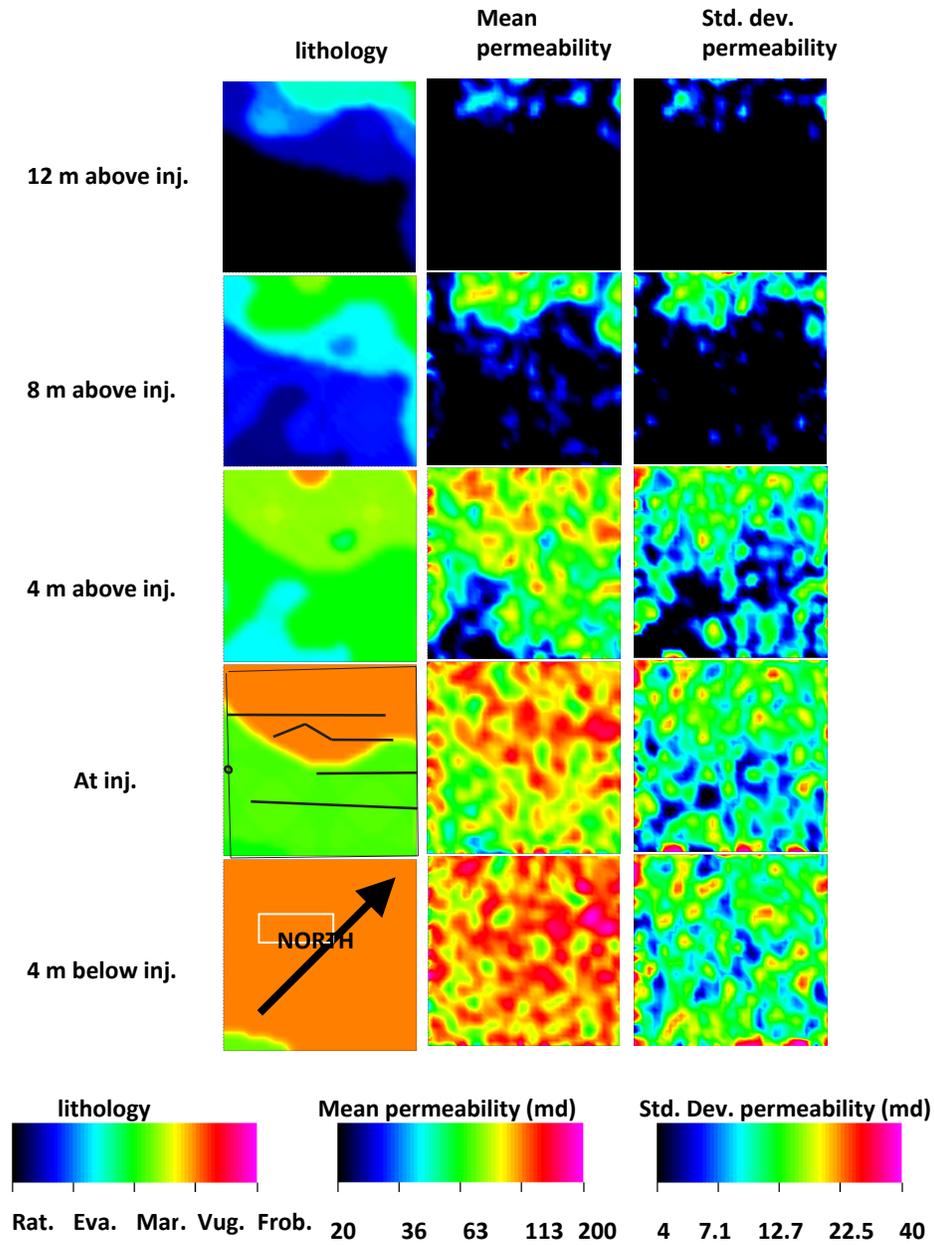


Figure 1.6. The images show the mean and standard deviation of all the permeability models with iteration number greater than 5700. The permeabilities are shown in millidarcies, on a logarithmic scale.

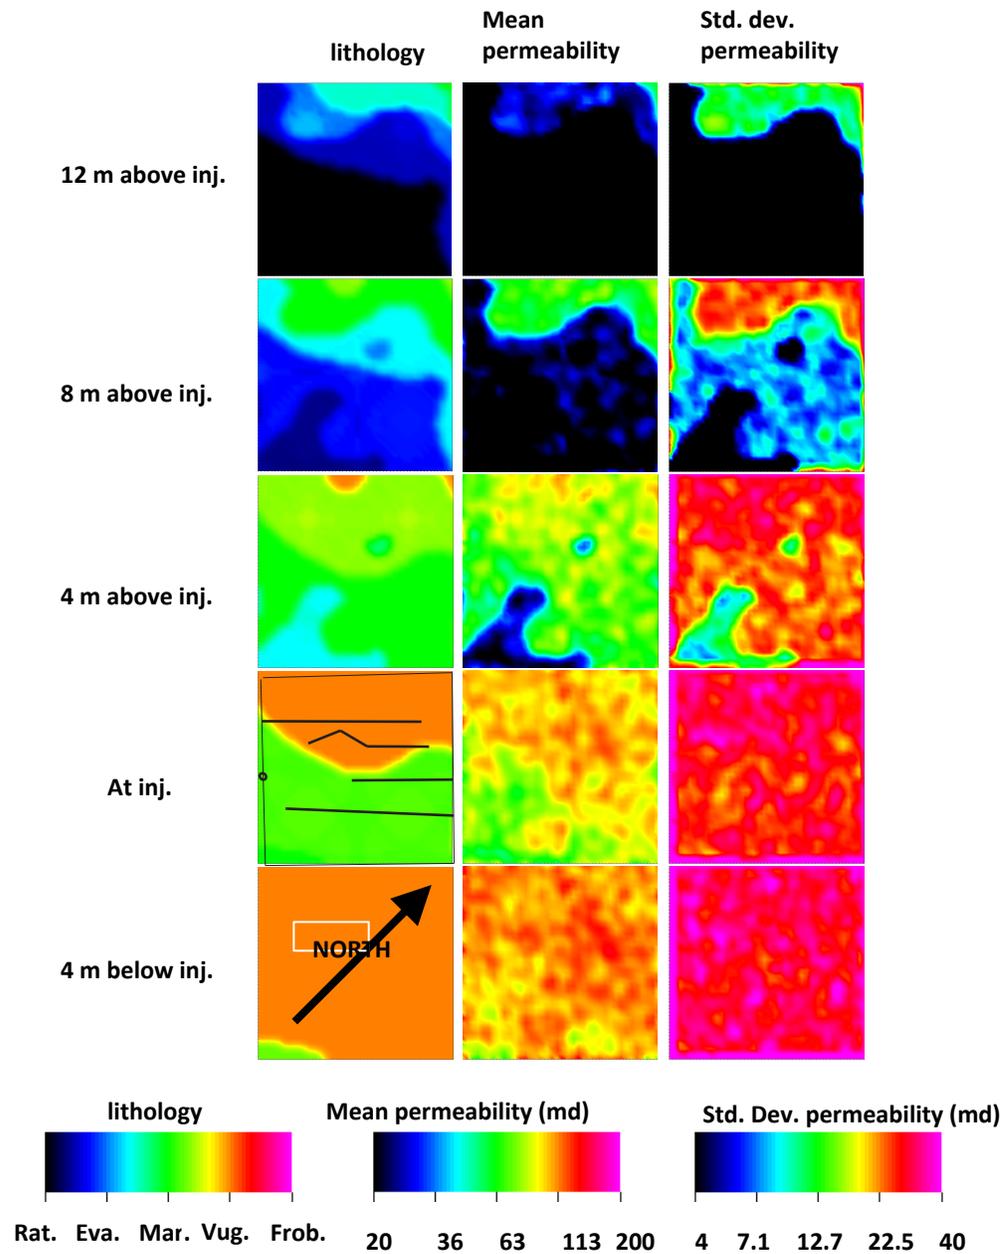


Figure 1.7. The images show the mean and standard deviation of all the permeability models from a run where only the prior information was used (no seismic data used). The permeabilities are shown in millidarcies, on a logarithmic scale.

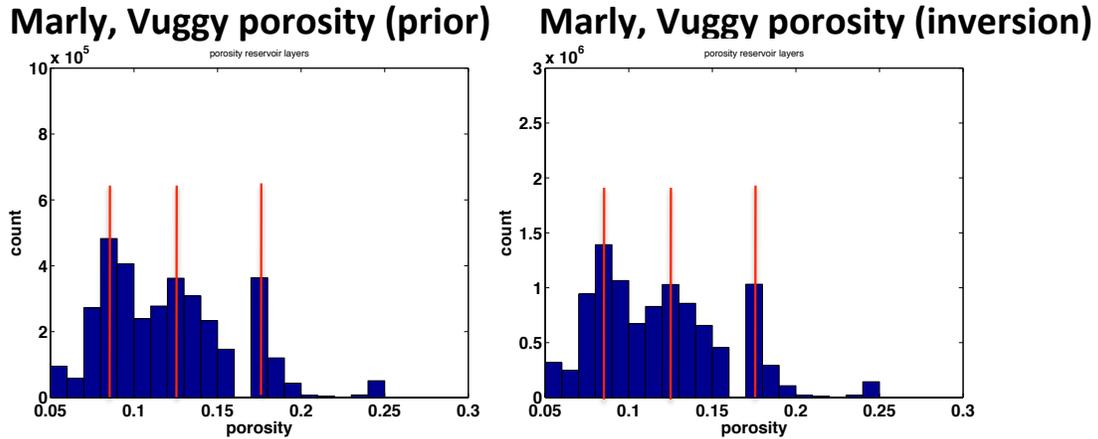


Figure 1.8 shows histograms of porosity using prior information only (left plot, before stochastic inversion) and after stochastic inversion (right plot). The left histogram is used to constrain the stochastic inversion. The two histograms are very similar.

Figure 1.8 and 1.9 provide another view of the inversion results. The figures compare porosity and permeability histograms using only prior information (no seismic data) and after stochastic inversion (using seismic data). The porosity histograms indicate that the inversion does not change the porosity distribution. Similar comments apply to the permeabilities distributions in Figure 1.9. The histograms indicate that the prior constraints determine the modes of the distributions (permeability/porosity values that occur most frequently) and that the seismic data has no discernible effect on them.

1.4 Comparison of Predicted and Observed Waveforms

We will now examine the match (or misfit) predicted and observed waveforms. Table 5 compares the mean misfit (average of the absolute differences between predicted and observed waveforms) for various cases. The table column labeled “Mean misfit using full waveform” shows the mean misfit calculated using all points in the waveforms. The table column labeled “Mean misfit using short waveform” shows the mean misfit calculated using only points that correspond to the Marly, Vuggy and Frobisher units; the points associated with the Ratcliffe unit are not included because these points generate large waveform differences that dominate the misfit calculation and are not related to reservoir properties.

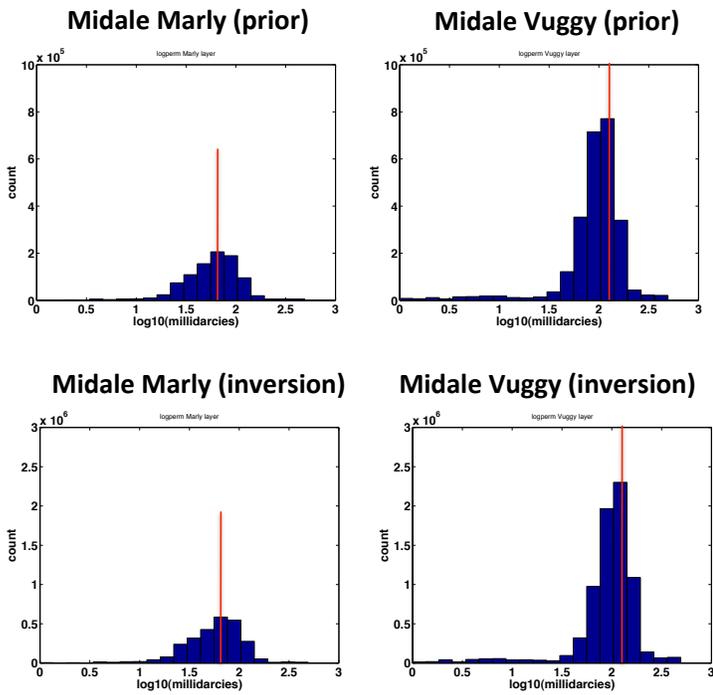


Figure 1.9 shows histograms of reservoir permeability using prior information only (top row plots, before stochastic inversion) and after stochastic inversion (bottom row plots). The top row histograms were used to constrain the stochastic inversion. The prior and inversion histograms are nearly identical.

Case	Mean misfit using full waveform	Mean misfit using short waveform
Pre-inversion, lab measured properties	0.831	0.908
Pre-inversion, Vuggy moduli reduced	0.772	0.574
Inversion calculated likelihoods using full waveforms	0.484	0.756
Inversion calculated likelihoods using short waveforms, iter. 5751	-----	0.533
Inversion calculated likelihoods using short waveforms, use all iter. > 5700	-----	0.511

Table 5 compares average misfits (average of the absolute differences between predicted and observed waveforms) for permeability inversion runs.

The case labeled “Pre-inversion, lab measured properties” shows the mean misfit obtained when laboratory measured bulk/shear moduli values are used and the stochastic inversion does not adjust the permeability/porosity fields. This case produced the largest misfit of all the cases considered, as expected.

In a previous progress report (Ramirez et al., 2011), we indicated that the peak associated with the Vuggy layer was larger in the predictions than in the observations. Following a trial and error approach we reduced the shear and bulk moduli assigned to the Vuggy unit in order to reduce the misfit as much as possible. We also decreased the values of the layer above the Ratcliffe. This had the effect of increasing the Ratcliffe peak and decreasing the Vuggy peak because the predicted waveform is scaled relative to the observations (see Ramirez et al., 2009 for details). The 2nd case “Pre-inversion, Vuggy moduli reduced” shows how the misfit decreased when the Vuggy adjustments were made and the stochastic inversion has not been used to adjust the permeability/porosity fields.

We then ran two stochastic inversions that used the adjusted Vuggy values. After the inversion finished, we calculated the mean misfit for the model that best fit the seismic data. The first inversion used all waveform points to calculate the likelihood values that guided the stochastic search (case “Inversion calculated likelihoods using full waveforms”, in Table 5). Notice that full waveform misfit is substantially larger than the short waveform misfit even though we have normalized for the difference in the number of points. We saw that the largest waveform differences tend to occur along the Ratcliffe waveform section, suggesting that the likelihood calculations were dominated by these differences. This dominance decreased the sensitivity to the reservoir units thereby increasing the short waveform misfit.

Based on these observations, we changed the way the likelihoods were calculated during the inversion, this time excluding the Ratcliffe waveform section. The 4th and 5th cases in Table 5 show that the short waveform misfits substantially decrease when the likelihoods are calculated in this manner. The 4th case, “Inversion calculated likelihoods using short waveforms, iter. 5751”, shows the misfit for one of the models from the run that produced the upper likelihood plot in Figure 1.4. Notice that the likelihood values reach reasonably stable maximum after iteration 5700; the model chosen as the 4th case comes from this region. We also averaged all the models after iteration 5700 and calculated the misfit (last row in table 5).

The 4th and 5th cases both show short waveform misfits that are smaller than the case 2 and case 3 misfits. This confirms that the inversion is reducing the misfit, as previously indicated by the Figure 1.4 plots showing that the likelihood is improving (misfit is decreasing). However, a comparison of the case 2 (pre-inversion) and cases 4 and 5 (post-inversion) misfits indicates that the misfit improvement is small.

We believe that the small improvement is due to the following factors. A) The seismic data is sensitive to seismic impedance that is strongly dependent on bulk/shear moduli and density. Thus, for our application, it is important to have accurate knowledge (in space and magnitude) of bulk/shear moduli and density for all layers of interest. The inversions described here use homogeneous bulk/shear moduli and density values that ignored their spatial variability. We attempted to mitigate this problem by inverting for the spatial distribution of these properties but our attempt was unsuccessful. B) The impedance is much less sensitive to the changes in fluid moduli and pressure associated with CO₂ injection. This means that the permeability/porosity realizations could only change the waveform a small amount and could not compensate for the relatively large waveform differences associated with factor A. C) The location of the layer boundaries from Cenovus’ reservoir model maybe off to some extent, affecting the location of the impedance interfaces and thus, the waveform shape.

Figures 1.10 – 1.15 show examples of the waveform comparisons at 5 locations throughout the reservoir. Each figure compares the waveforms for 3 cases (top to bottom): a) pre-inversion, using laboratory-measured properties, b) pre-inversion after the Vuggy moduli were adjusted by trial and error, and c) after the inversion had adjusted the permeabilities. The diagram on the right shows a map view of pattern 16 and the red circle indicates the waveform locations. Notice the large peak associated with Ratcliffe formation (peak centered around index 10). For the short waveform misfits tabulated in Table 5, we used all points above index 12, thereby ignoring the Ratcliffe. The trough near indexes 13 and 14 corresponds to the top of the Marly and the peak centered near index 15 is associated with the top of the Vuggy unit.

The waveform sections corresponding to the Marly and Vuggy show that the misfit improvement is largest when the adjustments were made to the Vuggy layer and the layer above the Ratcliffe by trial and error; this supports our argument that the waveforms depend strongly on the bulk/shear moduli. The differences between the middle plot and

bottom plot are barely perceptible indicating that the change in pressure and fluid moduli and concomitant permeability distributions only change the waveforms a small amount. We believe that the poor waveform fit probably affected the accuracy of the inverted porosity/permeability fields shown here; we do not know how to estimate the magnitude of this effect at present.

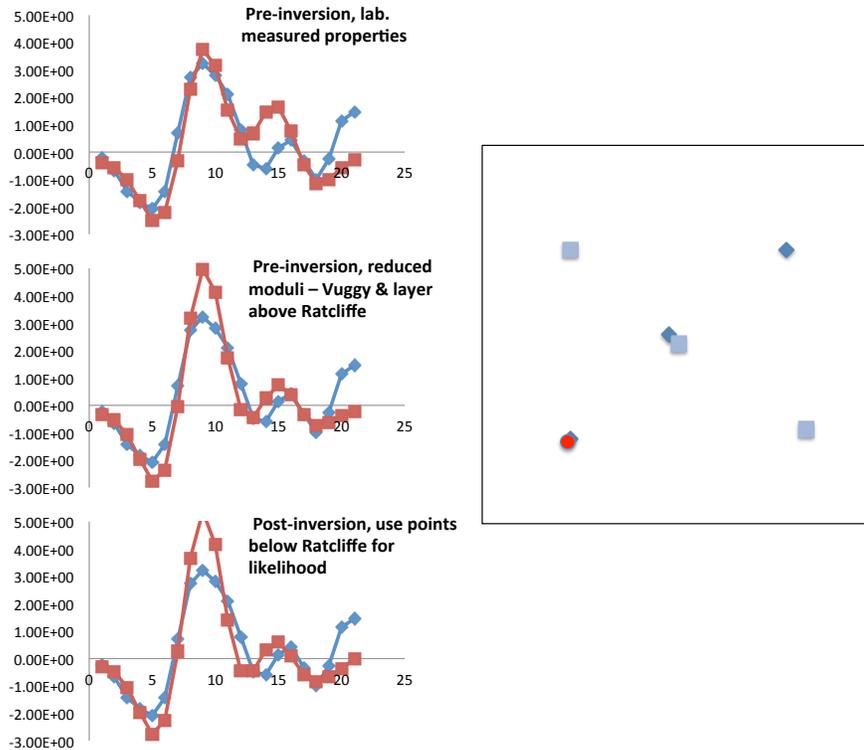


Figure 1.10. The left column of plots compares the observed (blue) and the predicted (red) seismograms for various cases. The right diagram shows the location of the observed seismogram indicated by the red circle. The likelihoods calculated by the stochastic inversion use waveform points below the large Ratcliffe peak (peak centered near index 10 and is assumed to end at index = 12).

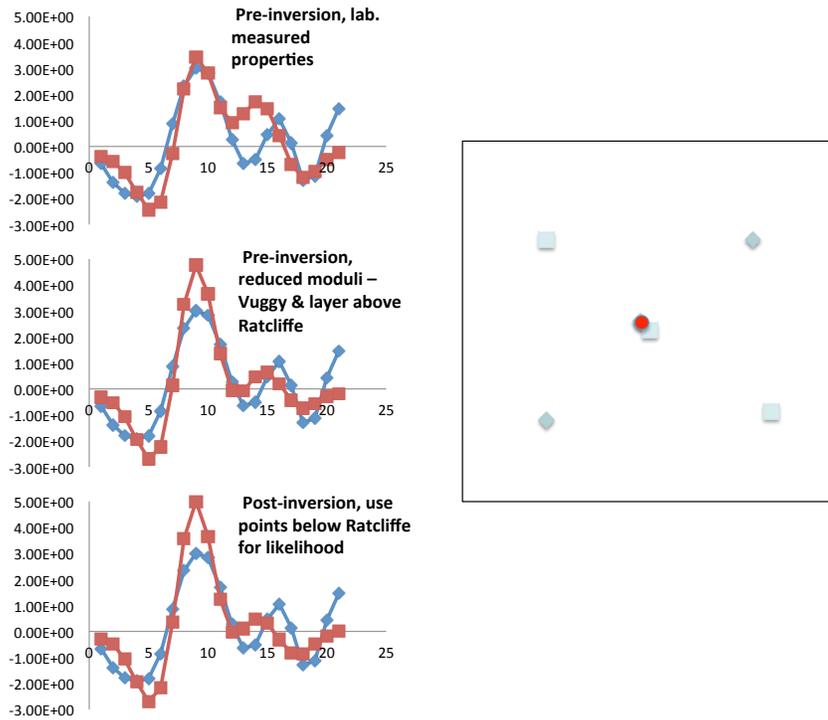


Figure 1.11. The left column of plots compares the observed (blue) and the predicted (red) seismograms for various cases. The right diagram shows the location of the observed seismogram within pattern 16, indicated by the red circle. The likelihoods calculated by the stochastic inversion use waveform points below the large Ratcliffe peak (peak centered near index 10 and is assumed to end at index = 12).

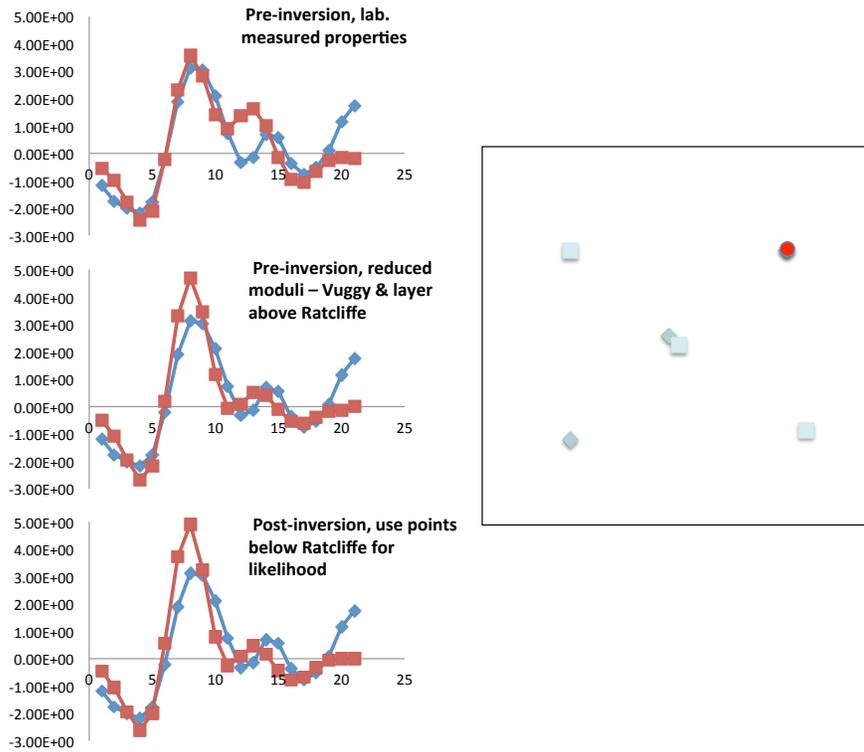


Figure 1.12. The left column of plots compares the observed (blue) and the predicted (red) seismograms for various cases. The right diagram shows the location of the observed seismogram within pattern 16, indicated by the red circle. The likelihoods calculated by the stochastic inversion use waveform points below the large Ratcliffe peak (peak centered near index 10 and is assumed to end at index = 12).

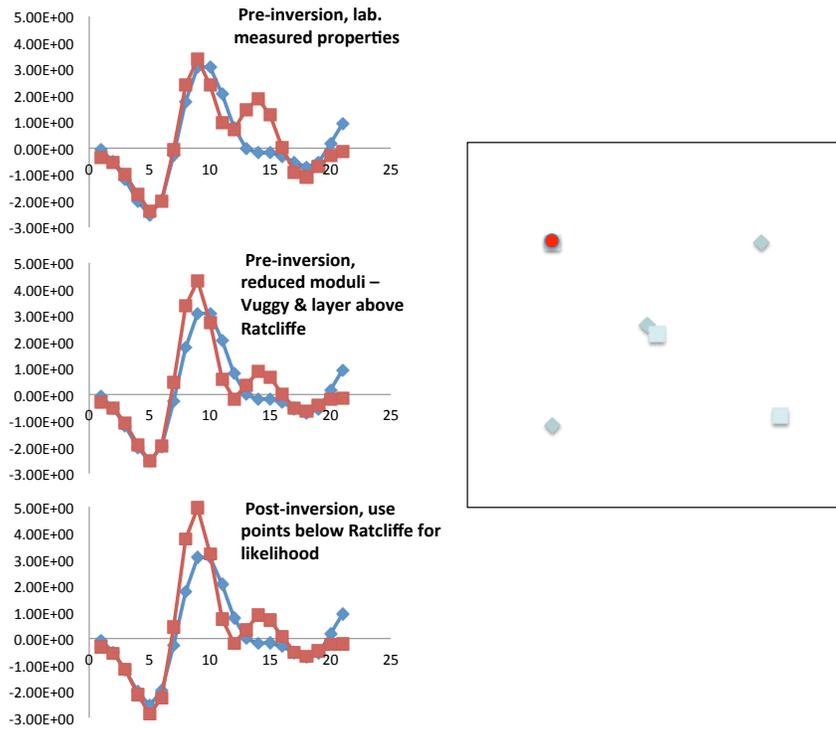
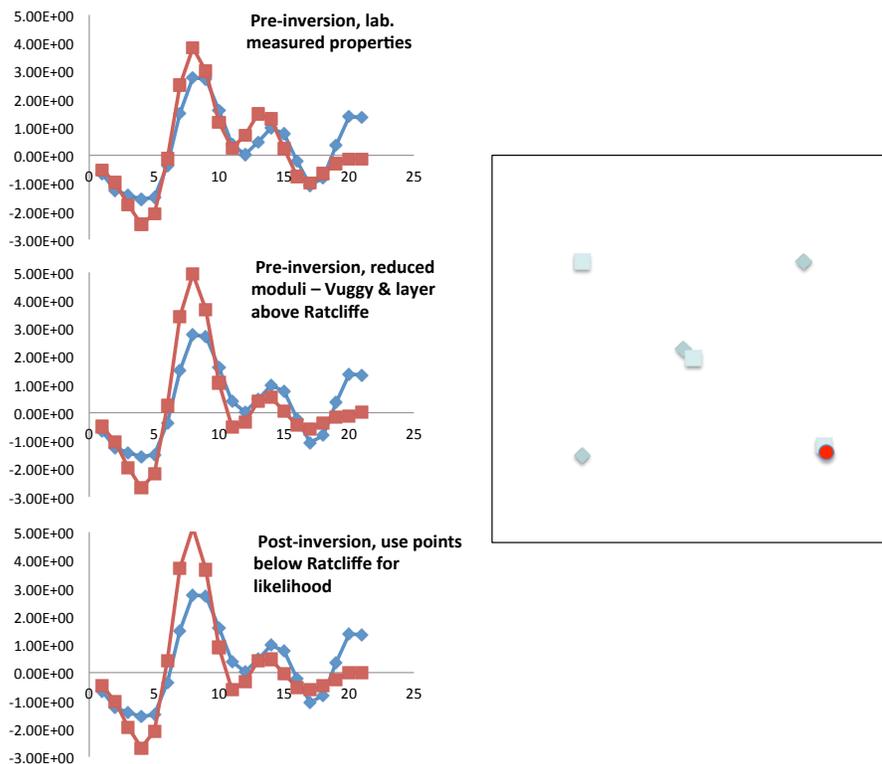


Figure 1.13. The left column of plots compares the observed (blue) and the predicted (red) seismograms for various cases. The right diagram shows the location of the observed seismogram within pattern 16, indicated by the red circle. The likelihoods calculated by the stochastic inversion use waveform points below the large Ratcliffe peak (peak centered near index 10 and is assumed to end at index = 12)..



The left column of plots compares the observed (blue) and the predicted (red) seismograms for various cases. The right diagram shows the location of the observed seismogram within pattern 16, indicated by the red circle. The likelihoods calculated by the stochastic inversion use waveform points below the large Ratcliffe peak (peak centered near index 10 and is assumed to end at index = 12).

1.5 Computational Expense

Perhaps the greatest challenge associated with the use of our stochastic inversion approach is its computational expense. Almost all the expense is in running the flow simulator that predicts reservoir conditions caused by injection/extraction operations.

The inversion described here required about 5700 iterations for the Markov chains to reach convergence. A total of 32,000 porosity/permeability models were evaluated by performing flow simulations for each. About 2500 of the models became part of the posterior distribution (i.e., the solution to the stochastic inverse problem), and the rest were discarded. Each run used 112 processors, running for about 6.8 days (processor time). The wall clock time needed to complete the run was about 9 days. The wall clock time is larger than the processing time because we used a multi-user machine where each run is limited to 16 hours of run time; when a run ends, it has to be re-submitted and wait in the queue for a few hours before execution resumes.

1.6 Summary and conclusions

The seismic inversions suggest that the permeability and porosity are spatially heterogeneous in the Marly and Vuggy layers, as expected. The Marly permeabilities tend to be smaller than the Vuggy's whereas the Marly porosities tend to be larger. Pre and post inversion histograms of permeability and porosity look very similar, thereby suggesting that the prior constraints determine the modes of the distributions (permeability/porosity values that occur most frequently). The seismic data did influence the locations of the permeability highs and lows within the reservoir layers.

A key challenge is that the seismic waveforms are more sensitive to the bulk modulus, shear modulus and density than they are to the permeability, porosity and CO₂. This means that it is necessary to obtain accurate estimates of seismic moduli and density before using seismic data to invert for permeability and porosity. Our original plan was to use homogeneous bulk and shear moduli for each layer (from either well logs or laboratory measurements) but this plan proved inadequate because these properties exhibit substantial heterogeneity. When this was discovered late in the project we developed and tested an inversion scheme that solved for the magnitude and spatial distribution of the moduli; this attempt was unsuccessful by the time the project ended.

We believe that this moduli inversion will work, given additional time and resources to continue development.

The seismic inversions reduced the misfit between the predicted and observed waveforms but the misfit improvements were less than expected. We believe that the poor waveform fit probably affected the accuracy of the inverted porosity/permeability fields. The root cause of this undesirable behavior probably stems from use of homogeneous bulk and shear moduli for each layer (from either well logs or laboratory measurements).

Chapter 2: Development and Application of a Markov Chain Scheme for Inverting the Distribution of Mineralogy and Reaction Rates Along a Flow Path in the Weyburn Reservoir

2.1 Introduction

The geochemistry component of the parameter inversion effort for the Weyburn reservoir assessment involves an attempt to estimate reaction rates associated with reservoir minerals deemed likely to undergo dissolution in response to the addition of CO₂. In addition to calcite and dolomite, these include silicate minerals which are unstable in acidic brine, such as feldspars, as well as anhydrite which may dissolve or precipitate in response to changing concentrations of calcium.

Inversion of “lumped” reaction rates entails both intrinsic, pH-dependent dissolution rates normalized per unit surface area and the associated mineral surface area per unit volume:

$$\frac{dM}{dt} = k A(\Omega - 1) \quad (\text{Eq.-2.1})$$

where M is the mineral mass, k the intrinsic rate constant (adjusted for temperature), A the mineral surface area, and,

$$\Omega = \frac{IAP}{K_{sp}} \quad (\text{Eq.-2.2})$$

where the saturation ratio, Ω , of a particular phase is given by the ion activity, IAP , product of the constituent reactant and product species in the mineral phase's dissolution equation and the solubility product K_{sp} . One approach to the inversion problem is to assume literature-derived values for the intrinsic dissolution rates and invert the mineral specific surface areas (and hence mineral mass or volume, assuming a relationship between volume and surface area) which reproduce observed changes in brine chemistry using a forward reactive transport model. An inversion attempt based on this approach must either (1) yield forward model results that are consistent with both observed data and other, independently-constrained model parameters, or (2) indicate that flawed assumptions exist, including, for example, the values of the intrinsic dissolution rates, and thus warrant modification of the inversion problem definition. This approach quantifies the reactive mineral assemblage members present at the start of the problem. Nonetheless, inversion to even a reduced set of parameters is highly problematic because of a number of key issues and uncertainties:

- In contrast to the seismic data inversion, the geochemical inversion entails multiple input parameters (e.g., lumped reaction rates for each of the proposed mineral assemblage members) and output values (e.g., pH, concentrations of indicators cations such as calcium and magnesium). The degree to which different output values must be weighted in compiling an overall goodness-of-fit or likelihood value, is not known *a priori*.
- In contrast to the high degree of spatial resolution offered by the seismic data for inverting the permeability field in three dimensions, brine geochemical parameters are available from a limited number of wells (e.g., only four wells with adequate data quality are located within Pattern 16). Moreover, samples drawn from a given well represent a spatial average of conditions across some (unknown) portion of the reservoir where the well intersects relatively permeable materials. The implied limit to associating water quality with a specific location within the reservoir is particularly true of horizontal wells.
- Blending of injected water from sources outside of the reservoir obfuscates the definition of ambient background brine composition, rendering initial conditions and boundary conditions inherently ill-defined.

- Geochemical models used to simulate water-rock interactions depend on thermodynamic data associated with putative reactive mineral phases to calculate saturation indices. These thermodynamic data, their temperature corrections, and the stoichiometry of individual phases included in the model are idealizations that may not be accurate.
- To invert reaction rates based on well data, a flow pathway connected the well to a CO₂-injection-induced geochemical perturbation must be inferred or assumed. Inconsistencies between the modeled flow pathway and reality will introduce reactive transport modeling errors associated with residence time, dispersion, and other factors.

The results of the parameter inversion are discussed in the context of these constraints, below.

2.2 Reservoir Geochemistry Model

The Vuggy and Marly units – the permeable constituents of the Weyburn-Midale reservoir – consist primarily of dolomite and calcite, with lesser abundances of aluminosilicate minerals such as feldspars, illite, and kaolinite. To posit a plausible geochemical model of the reservoir with which to inform the geochemical parameter inversion calculations, the brine water quality data set, including baseline data as well as data collected during subsequent monitoring events following the commencement of CO₂ injection, must be reconciled with the inferred mineral abundances and presumed reactivities of key mineral phases.

Monitoring data were collected beginning in August 2000 (Baseline sampling event). The Baseline data set as well as data from subsequent sampling events through Monitoring 11 (September 2004) were speciated using the PHREEQC geochemical model (Parkhurst and Appelo, 1999) to assess the possible impact of CO₂ injection across the reservoir. Initial speciation of 328 water complete water samples collected from multiple wells during these sampling events indicates that the majority of the samples are thermodynamically highly supersaturated with respect to calcite (Figure 2.1), based on laboratory-measured pH samples. It is likely that these samples reflect some off-gassing of CO₂ under ambient atmospheric temperature and pressure during sample collection, recovery, handling, and analysis, resulting in a pH rise and calcite supersaturation which are not indicative of *in situ* conditions (Emberley et al., 2005). This explanation is supported by comparing measured pH values with a set of subsequent of downhole pH measurements collected during some of the early monitoring events, which reveals laboratory-measured pH values that are appreciably higher than those obtained via downhole probe at pH values less than 7.0 (Figure 2.2), presumably those samples characterized by the highest CO₂ partial pressures and hence the most likely to be affected by off-gassing.

To compensate for the loss of CO₂, and hence the presumably skewed pH values across the 328 samples, PHREEQC was used to acidify individual samples. At a fixed alkalinity, this implies the addition of CO₂ to the solution composition. To constraint this

calculation, log saturation indices for calcite were reduced to +0.3 in a majority of the water samples which were initially identified as being highly supersaturated. The pH in a small subset of samples with calcite log saturation indices between +0.3 and +0.5 was not adjusted because of difficulty reconciling total alkalinity and bicarbonate alkalinity. The effect of this adjustment on computed pH versus downhole pH and on carbonate saturation indices is shown on Figures 2.3 and 2.4, respectively. This adjusted water quality data set comprises the brine chemistry data, including concentrations of dissolved CO₂, with which reactive transport model results may be compared and hence assessed, as discussed below.

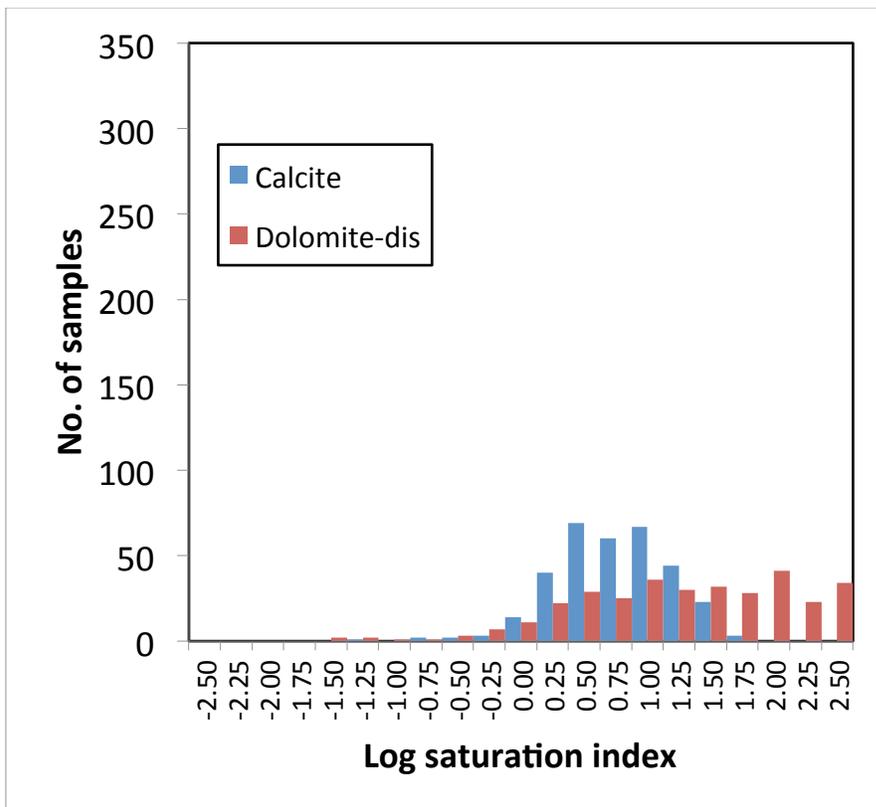


Figure 2.1. Calcite and disordered dolomite saturation indices in brine samples collected during monitoring events.

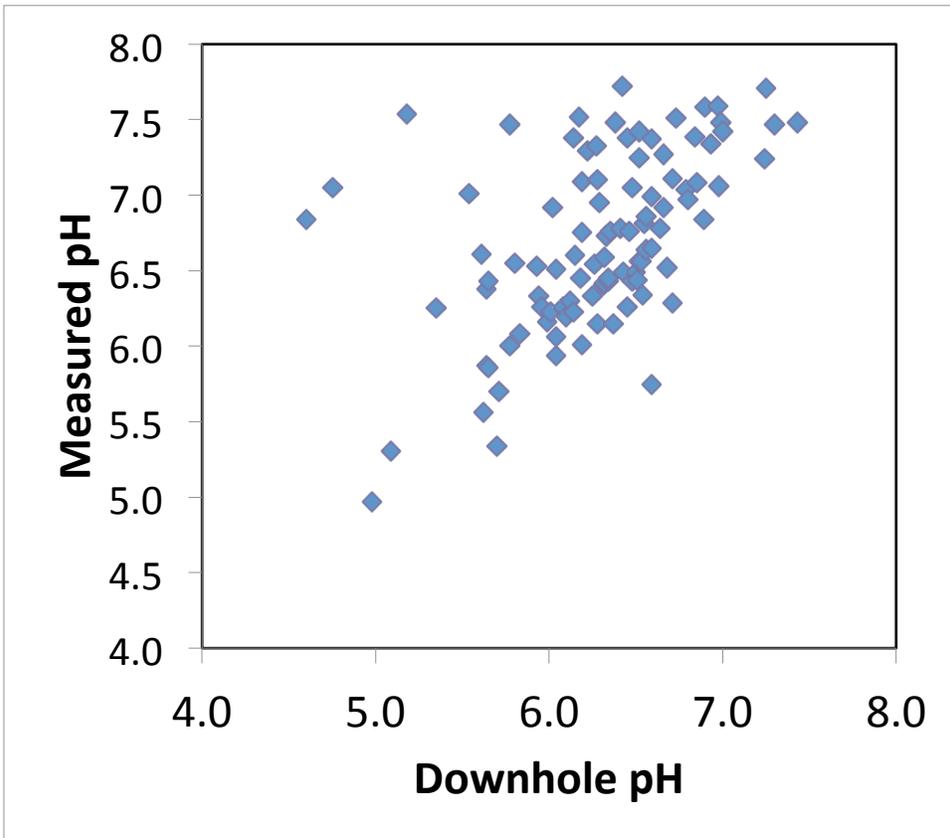


Figure 2.2. Laboratory-measured pH versus downhole pH.

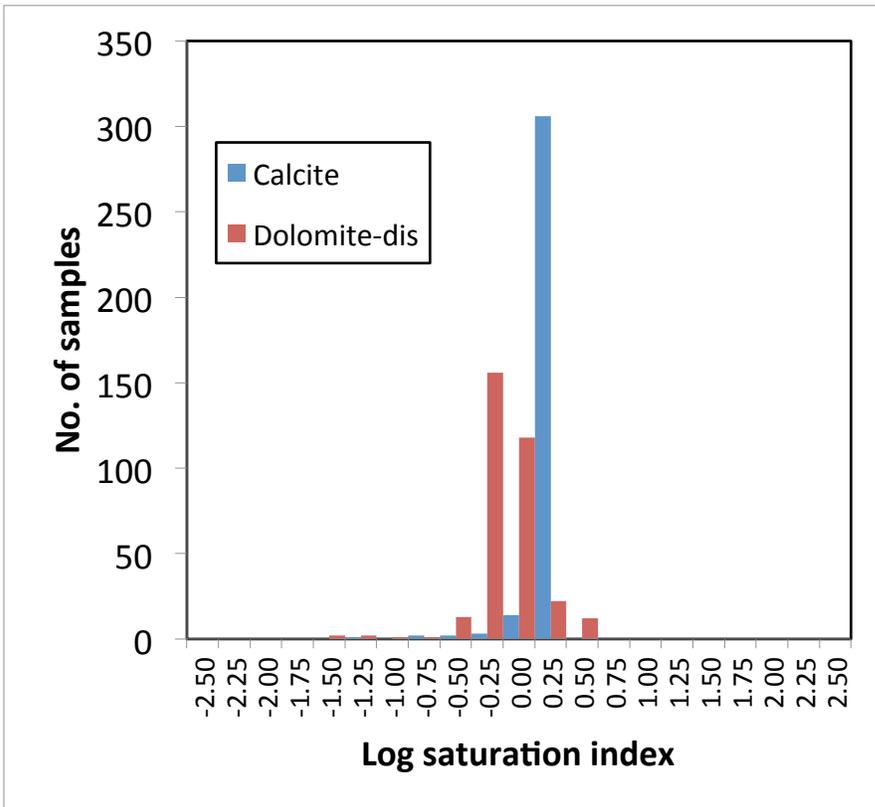


Figure 2.3. Calcite and disordered dolomite saturation indices in brine samples, corrected by simulated pH decrease/CO₂ addition.

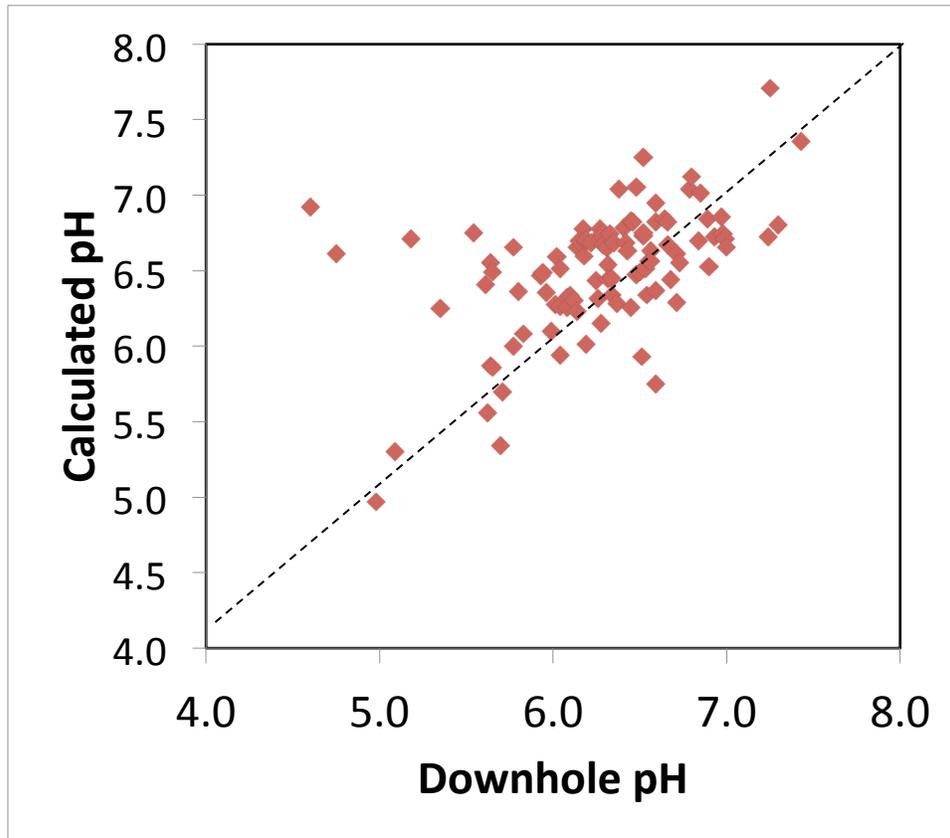


Figure 2.4. Modeled pH (pH-corrected samples) versus downhole pH. Dashed line indicates ideal match.

2.3 Geochemistry Inversion Using the Markov Chain Monte Carlo Algorithm: Development and Testing Using Synthetic data

The Markov Chain Monte Carlo (MCMC) method utilizes a likelihood function to determine whether or not a set of model parameter values proposed by a Markov chain Monte Carlo simulation will be accepted and serve as a basis for a subsequent proposal. The likelihood function for a particular model output (e.g., Ca^{2+} or Mg^{2+} concentration) is defined by:

$$L(\underline{x}) = k \exp\left(-\frac{1}{N} \sum_{i=1}^N \frac{|d(\underline{x})_{pred,i} - d_{0,i}|^n}{\sigma_i^n}\right) \quad (\text{Eq. 2.3})$$

where $L(\underline{x})$ is a measure of the degree of fit between the model predictions, assuming model \underline{x} , and the observed data, k a normalizing constant, N the number of data points, $d(\underline{x})_{pred,i}$ is the predicted data for a given model \underline{x} , $d_{0,i}$ the vector of observed measurements, σ_i the estimated data uncertainty, and $n \geq 1$. Because the geochemical

inversion involves multiple outputs, different likelihood functions must be calculated for each, via Eq.-2.3, and blended as a weighted sum.

To demonstrate the MCMC approach for inverting the initial distributions of mineral mass (and hence, by extension, surface area and lumped reaction rates) in the context of reactive transport, a synthetic test problem was created, based on the Weyburn reservoir mineralogy and brine chemistry, as a means for assessing the capability of the algorithm to find known initial conditions. The test problem was based on the following set of assumptions, as applied to a forward reactive transport model:

1. To reduce the computational burden, a 1-D column 20 m in length is divided uniformly into 20 cells.
2. The uniform permeability of the column is 100 mD, with a porosity of 0.15.
3. The initial brine chemistry in the column is based on the average background brine composition, as described in Section 2.2.
4. Single phase, steady-state flow of CO₂-impacted brine occurs through column. This is accomplished by specifying a fixed fluid pressure gradient across the column of 0.01 bar/m. The influent CO₂-impacted brine composition is based on equilibration of the ambient background brine composition with CO₂ at a partial pressure of 100 bars.
5. For advective-dispersive solute transport, the column dispersivity (characteristic length) is set equal to 1 m.
6. A set of reactants – calcite, dolomite (disordered), anhydrite, and K-feldspar (representing a silicate phase most likely to undergo dissolution in response to a pH decrease) - is distributed long the column, each with fixed intrinsic dissolution rates. Overall reaction rates depend also on mineral surface area (Eq.-2.1), which is assumed to be a function of mineral mass, as well as a mixing efficiency factor that is assumed to depend on fracture porosity, as described below.

Intrinsic mineral dissolution rates are assumed to depend on temperature and pH:

$$k = k_{25}^{nu} \exp\left(\frac{-E_a^{nu}}{R} \left(\frac{1}{T} - \frac{1}{298.15}\right)\right) + k_{25}^H \exp\left(\frac{-E_a^H}{R} \left(\frac{1}{T} - \frac{1}{298.15}\right)\right) \{H+\}^n \quad (\text{Eq.-2.4})$$

where R is the universal gas constant, T the absolute temperature, $\{H+\}$ the activity of the hydrogen ion, n an exponential factor, and k_{25}^{nu} and k_{25}^H the neutral and acid rate constants at 25°C and E_a^{nu} and E_a^H the neutral and acid activation energies, respectively. Using data summarized by Palandri and Kharaka (2004) and an assumed average pH of 5.5 across the column during the progression of the CO₂-rich brine, the resulting intrinsic rate constants for calcite, dolomite, and K-feldspar are 3.1×10^{-6} , 1.2×10^{-6} , and 4.7×10^{-11} mol m⁻² sec⁻¹, respectively. For anhydrite, an intrinsic dissolution rate one order-of-

magnitude faster than the carbonates, or $2 \times 10^{-5} \text{ mol m}^{-2} \text{ sec}^{-1}$, is assumed. Dissolution reactions for both dolomite and K-feldspar are constrained to be irreversible.

For all four minerals, assuming a fixed intrinsic dissolution rate based on an average pH ignores the pH dependence implied by Eq.2-4. This assumption was employed in the construction of the synthetic test problem for the express purpose of computational convenience, as the relatively high intrinsic dissolution rates implied near the injection well (i.e., a low pH/high H^+ activity zone) would require small time steps and thus significantly increased computation time. The effects of modeling Eq.-2.4 directly versus employing fixed approximations is explored in Section 2.4, below.

To convert between mineral volume and specific surface in the reactive transport model, it is assumed that the mineral grain geometry can be approximated by spheres. Based on unpublished experiments involving core samples from the Weyburn reservoir (S. Carroll et al., LLNL), specific surface areas for calcite and dolomite are assumed to be on the order of approximately $0.03 \text{ m}^2/\text{gm}$, implying spheres with effective diameters of $100 \text{ }\mu\text{m}$. For K-feldspar, which is distributed much more sparsely than the other mineral phases, an effective particle diameter of approximately $3 \text{ }\mu\text{m}$ is assumed, based on the analyses of Weyburn lithology samples conducted by Durocher et al. (2005).

Finally, for fluid flow and transport through the one-dimensional column, the rate law implied by Eq.-2.1 only applies for the case of ideal mixing when all of a mineral's mass within a given volume element is exposed to reaction with the influent brine. In the presence of preferential flow conduits, particularly in instances in which flow through fractures predominates, this mixing assumption will not hold. As such, a multiplier, Φ , is applied to the right-hand-side of Eq.-2.1 to account for reduced water-rock contact; i.e., $\Phi \ll 1$. A value of 5×10^{-5} was assigned to Φ based on the *in situ* fracture porosity estimated by Cardona et al. (2002).

Running the inverse problem entailed conducting multiple forward reactive transport simulations for the 1-D column, using the reactive transport modeling capability of PHREEQC with the tacit assumption of single-phase flow. Individual trial simulations, or proposals, were based on a postulated distribution of mineral assemblages placed along the column, spatially correlated as a single Gaussian field using a mineral index number function. This *ad hoc* function, consisting of a series of superimposed Gaussian distribution functions about mean and standard deviation values set between 0 and 1, allows mineral assemblages (e.g., silicate mixture, calcite- or dolomite-rich carbonates, anhydrite) to be "extracted" according to some particular scalar value between 0 and 1 (Figure 2.5). Given the initial masses of each mineral phase in each volume element, surface areas and hence overall reaction rates can then be calculated, dynamically, as the simulation progresses.

An initial mineral assemblage distribution was created using a spatially correlated mineral index function along the column to serve as a synthetic truth data set to test the MCMC inversion approach (Figure 2.6). Subsequent proposals were then constructed using different mineral index number distributions placed along the column. A likelihood function was calculated for four measured outputs which are potentially

indicative of mineral dissolution: pH, Ca^{2+} , Mg^{2+} , and dissolved Si. The likelihood function values for all four parameters, weighted equally, were summed, with the respective values for σ (Eq.-2.3) adjusted to produce a reasonable dynamic range of composite values. Proposals were accepted when the ratio of the composite likelihood value to that of the prior proposal was greater than that of a random number selected from a uniform distribution between 0 and 1; this is the well known Metropolis-Hastings acceptance criteria. After a proposal was accepted, its mineral index distribution served as the basis for a new distribution, implemented by replacing a contiguous section of the column, up to 0.25 of its length, with new (spatially-correlated) mineral index numbers. A python script was used to generate proposals, write PHREEQC input files, run the reactive transport model, and read the output file, and score the proposals using the likelihood function.

The distribution of composite likelihood values, computed using the four measured outputs from each of the 20 volume elements after one pore volume of CO_2 -enriched brine has passed through the column, versus the proposal number (i.e., MCMC iteration number) is shown on Figure 2.7. This case simulates an ideal (and unrealistic) field scenario where there is one sampling well in each of the cells along the 1D flow-path. The behavior of the MCMC algorithm in this example is one characterized by occasional peak values, representing the best fits to the synthetic data set, followed by descents into poorer solutions as the parameter space (i.e., mineral index function values) is explored. Among this set of proposals, the mineral assemblage utilized in the best likelihood case (Figure 2.8) provides a good match to the synthetic truth mineralogy (Figure 2.6) and, by association, the synthetic truth brine chemistry (Figure 2.9).

A second MCMC demonstration involved an attempt to invert to the same synthetic truth data set, but using only the data corresponding to the last volume element in the column, i.e., the effluent port in the 1D column. This simulates the typical field situation where water samples are collected within a flow path connecting the CO_2 injector and the sampling well. This inversion has much less data with which to choose among the proposals, thereby producing results that are inherently less unique, and hence more difficult to assess, because the composition of the brine arriving at the end of the column encompasses all of the time- and space-integrated reactions which have occurred upgradient. The distribution of composite likelihood values for the four measured outputs from the last volume element after one pore volume versus the sequence of proposals is shown on Figure 2.10. In comparison to the 20-volume-element inversion (Figure 2.7), the behavior of the MCMC algorithm is much more erratic for this application. Nonetheless, relatively good fits to the synthetic truth brine chemistry can be distinguished in the set of accepted proposals. For example, averages of inverted mineral index function values along the column taken from all proposals with composite likelihood function scores exceeding 1.5 (Figure 2.11) yields a coarse replica of the synthetic truth mineralogy distribution, with dolomite replacing calcite in the downgradient (right) portion of the column, along with an appearance of anhydrite. This demonstration of some success of the MCMC approach in this example is important for establishing applicability to field data sets with low spatial resolution.

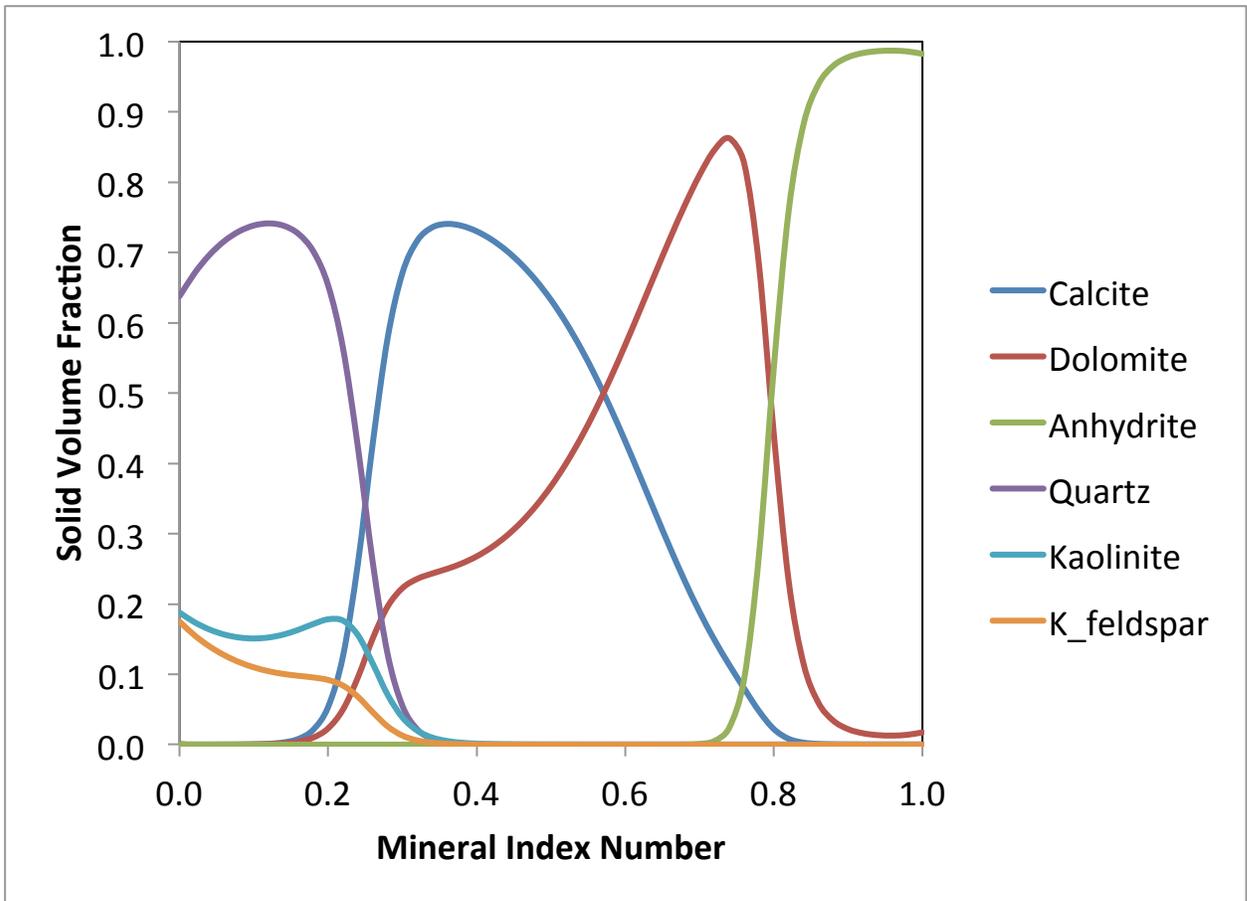


Figure 2.5. Mineral solid volume fractions as a function of the mineral index number, as defined via Eq. (4).

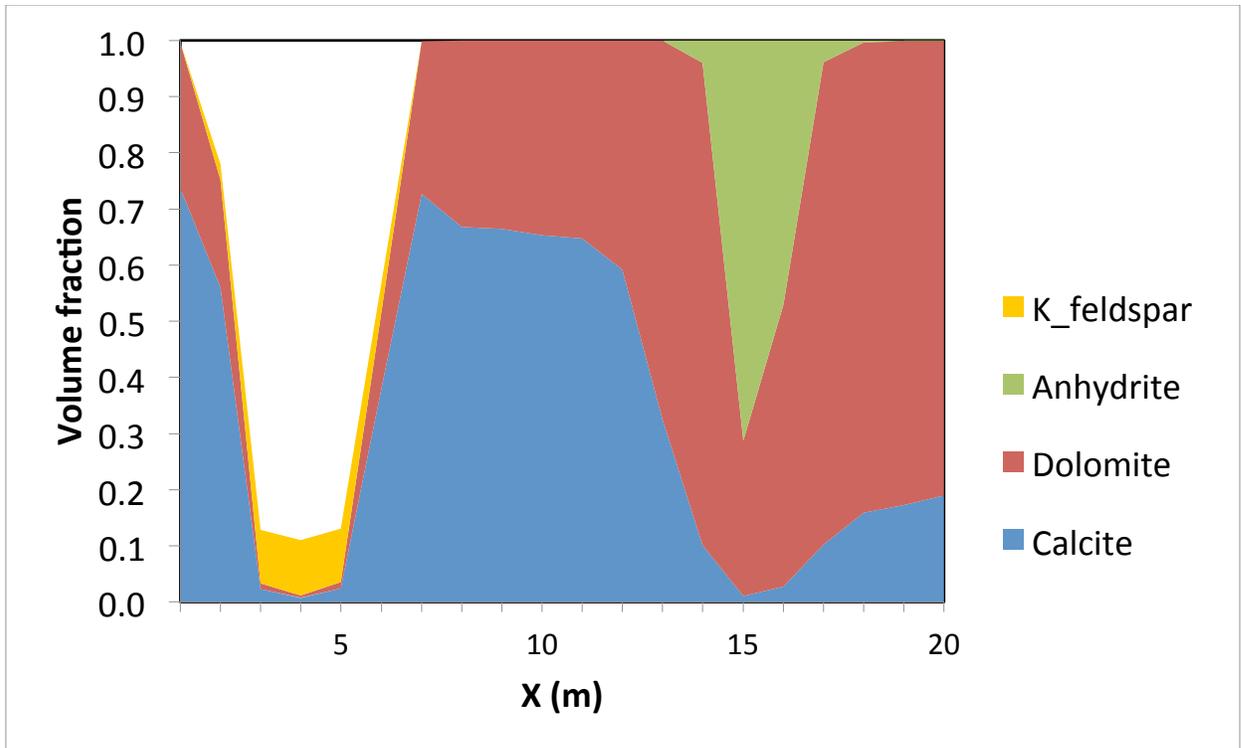


Figure 2.6. “Synthetic truth” initial distribution of reactive mineral phases along the column. At time $t = 0$, CO_2 -enriched brine enters the column at the left end.

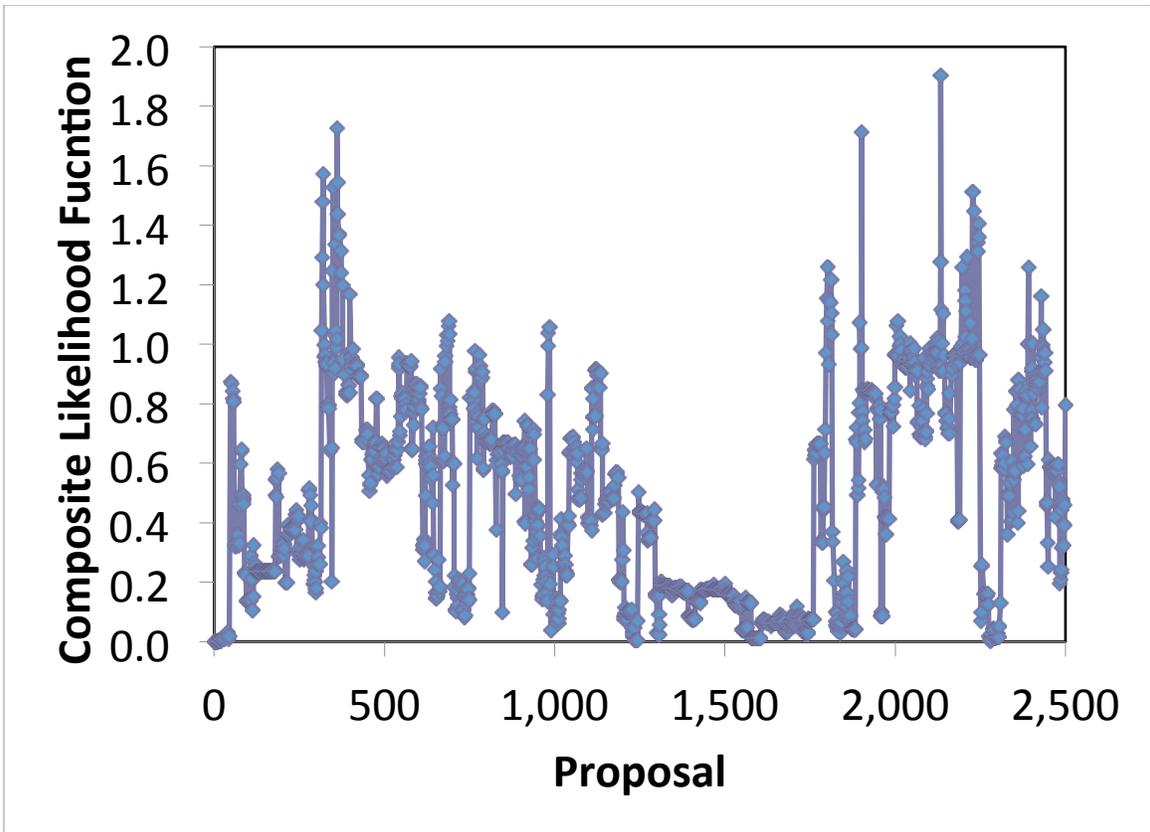


Figure 2.7. Behavior of the composite likelihood function as a function of proposal number when the synthetic data for all 20 cells are used in the inversion. Proposal rejection rate is approximately 36%.

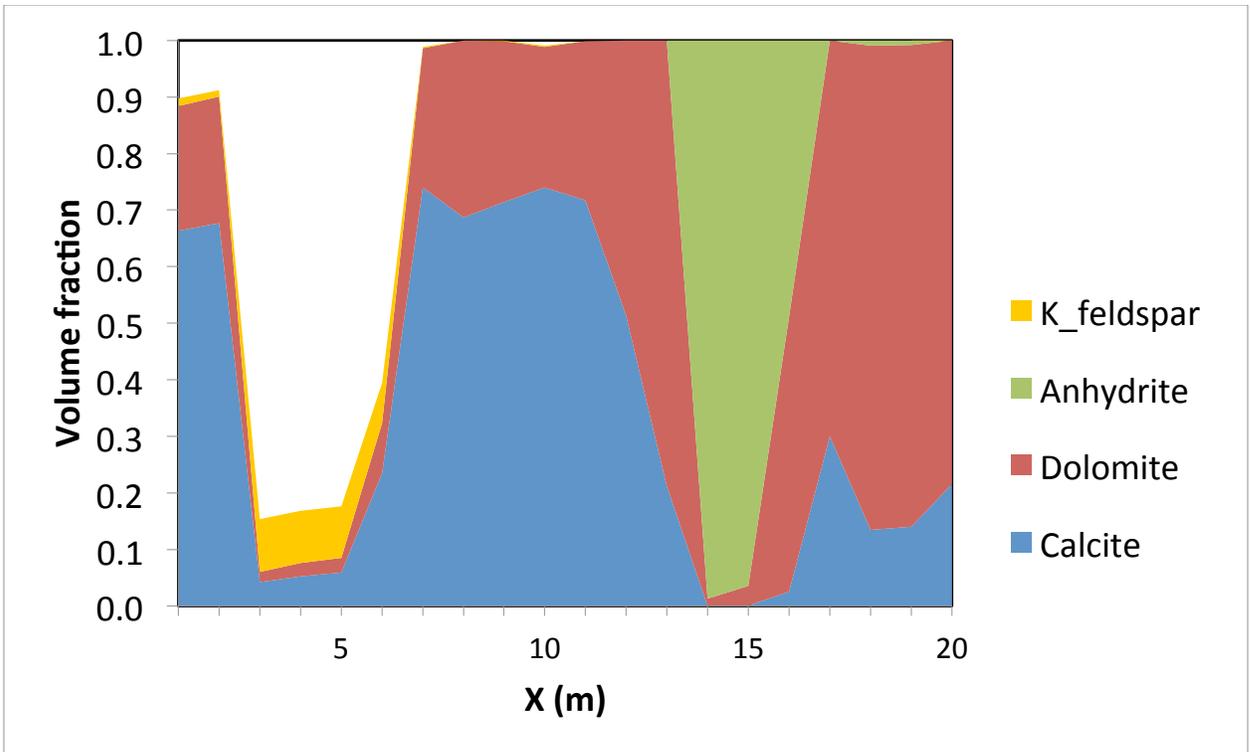


Figure 2.8. Initial mineralogy distribution implied by Proposal #2,134, representing the highest likelihood function score among the 2,500 proposals generated for the full 20-cell-based inversion.

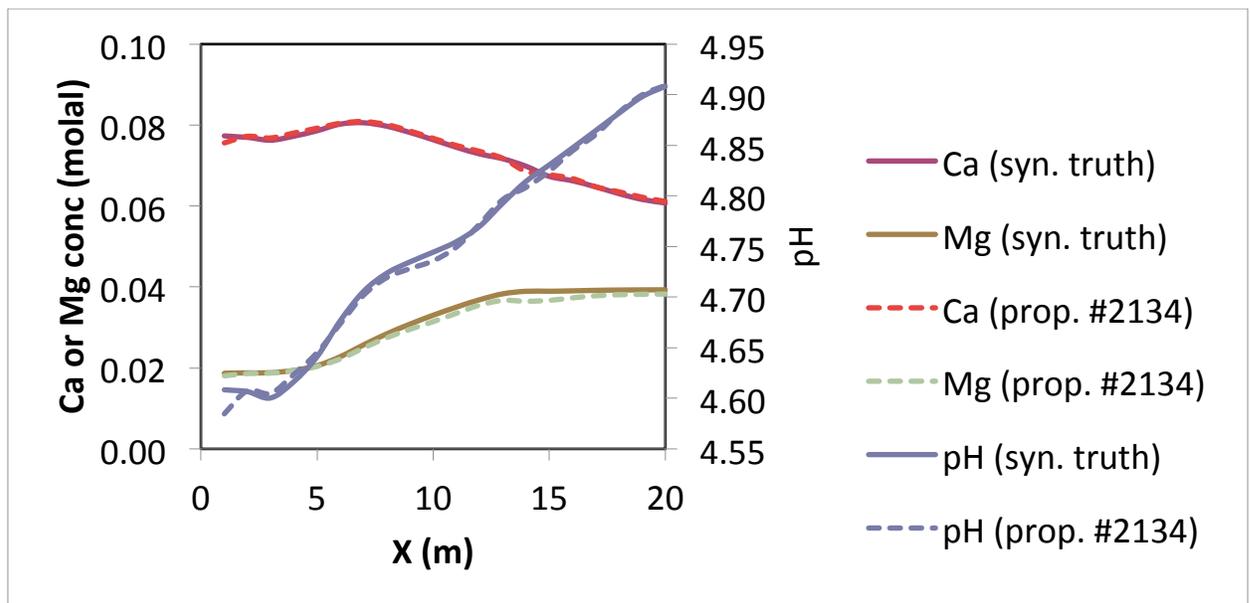


Figure 2.9. Comparison of model output and synthetic truth data for Ca, Mg, and pH along the entire column length (20 cells).

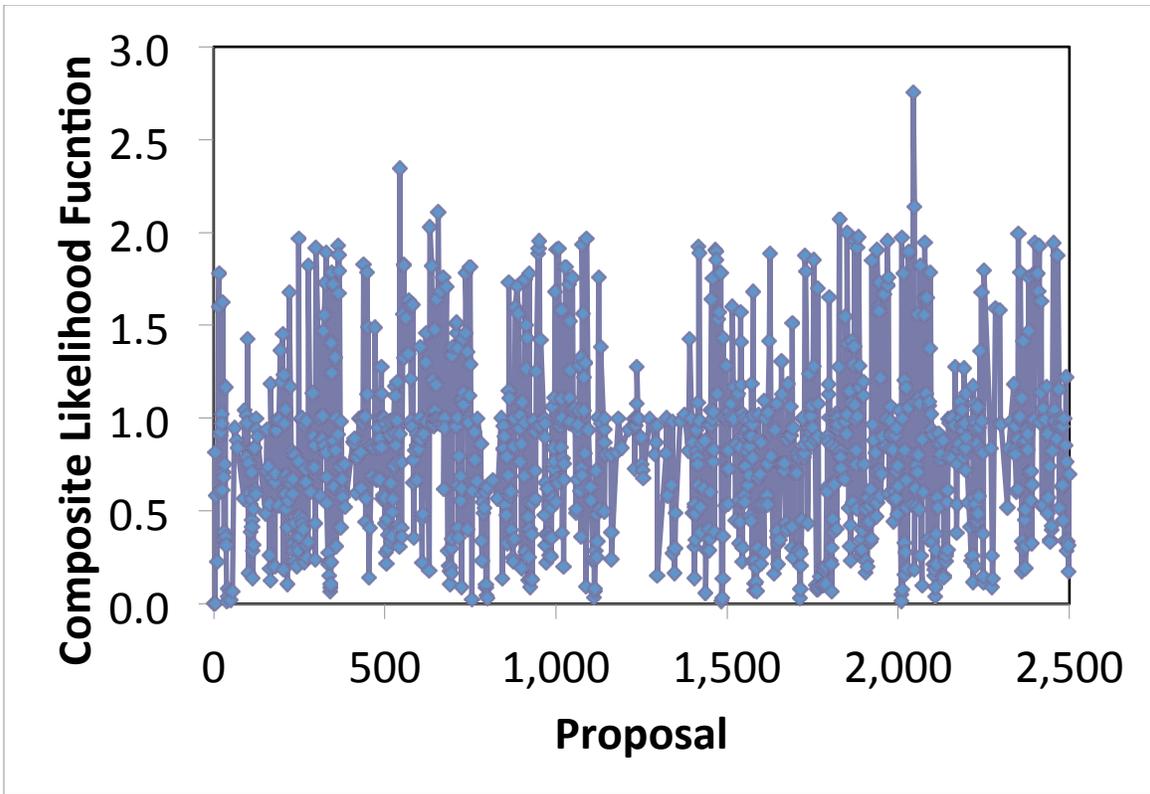


Figure 2.10. Behavior of the composite likelihood function as a function of proposal number when only the synthetic data for the last column cell are used in the inversion. Proposal rejection rate is approximately 54%.

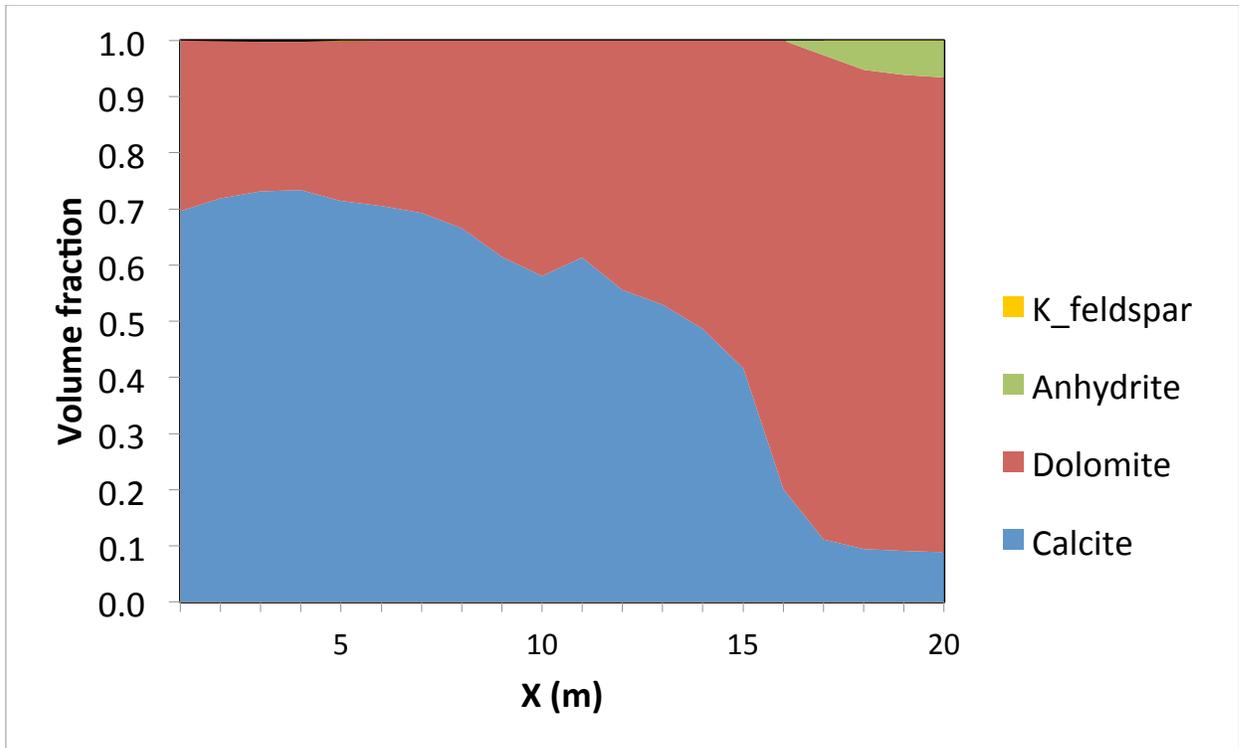


Figure 2.11. Inverted mineralogy distribution implied by the average mineral index values from all proposals with likelihood function scores above 1.5 – or approximately the top 10% of all accepted proposals – generated by inverting from the synthetic data for the last volume element in the column.

2.4 Inversions using Brine Chemistry Data from Pattern 16

Pattern 16 was selected as the target for the MCMC inversion of seismic data to yield putative distributions of permeability. For the geochemistry inversion, complete, corrected brine chemistry data (see Section 2.2) are available from several wells – primarily oil producers – within or near the borders of Patterns 16 (Figure 2.12). Of these, two wells, D08-12-06-14 and 02-12-06-14, exhibit brine chemistry changes during the first 1,000 days of CO₂ injection, including decreased pH and increases in the concentrations of Ca²⁺, Mg²⁺, dissolved Si, and total dissolved CO₂, that are consistent with increased carbonic acid content. Inversion of the D08-12-06-14 data was attempted, using the single-volume-element version of the 1-D test problem described in Section 2.3 as a template, to help constrain possible mineral reaction rates along a possible flow path from the Pattern 16 injector. Three modifications to the test problem model configuration were required:

1. The 1-D flow tube, divided into 40 equal volume elements, was extended to a length of 400 m to represent the distance scale of the shortest path between the CO₂ injector and well D08-12-06-14. This represents a significant change with respect to the original plan of performing a 3-D inversion in order to reduce the

computational expense of the approach, as discussed in the Introduction chapter of this final report.

2. The column permeability (120 mD), porosity (0.1), and pressure gradient (0.015 bar/m) were set to values which permit arrival of a solute front, subject to dispersion, at the well location within a 1,000-day time frame, as suggested by monitoring data.
3. The partial pressure of CO₂ in equilibrium with the brine entering the column at the injection well location was set randomly, per proposal, between 30 and 70 bars, as full CO₂-saturated brine partial pressures are not present in the data set.
4. The dispersivity (characteristic length) of the column varied randomly between 10 m and 40 m, per proposal, since the flow rate and hence the position of the solute front is unknown (unlike the test problem).
5. The intrinsic mineral dissolution rates were modeled using two different approaches, yielding two separate inversions. The first approach entailed employing fixed intrinsic dissolution rates at pH 5.5, as employed for the synthetic problem presented in Section 2.3. The second approach entailed directly solving Eq.-2.4 dynamically (i.e., in space and time) as part of each proposal's forward model.

The forward reactive transport models were all run out to approximately 1,280 days of injection, representing the point at which the solute front migrated 90% of the way through the column.

For the fixed intrinsic dissolution rate inversion, a total of 2,500 proposals were tested; the rejection rate was approximately 32%. The range of mineral index functions associated with the top 5% of accepted proposals at different locations along the column (i.e., proposals with the best likelihoods) are summarized on Figure 2.13, while the inverted distribution of mineralogy from the best likelihood proposal is shown on Figure 2.14. Both plots suggest, as expected, a reactive mineralogy dominated by calcite and dolomite, with a silicate mineral (K-feldspar) also present. The highest scoring mineralogy proposals appear to all exclude anhydrite dissolution or precipitation as a major contributor to the observed brine chemistry.

For the variable intrinsic dissolution rate inversion, 300 proposals were generated as a consequence of appreciably increased computation time. The rejection rate for this inversion was also approximately 32%. The inverted distribution of mineralogy associated with the best likelihood proposal from this inversion is illustrated on Figure 2.16. This distribution is similar in character to that of the fixed intrinsic dissolution rate inversion (i.e., predominantly dolomite dissolution at the upgradient portion of the column, largely replaced by calcite dissolution toward the downgradient end), indicated a relative lack of sensitivity of the inversion to the differences in the mineral dissolution kinetics model with regard to pH dependence. It is also worth noting that both inversions employed different seed values for random number generation and yet proceeded to generally convergence toward the same overall mineralogical trend.

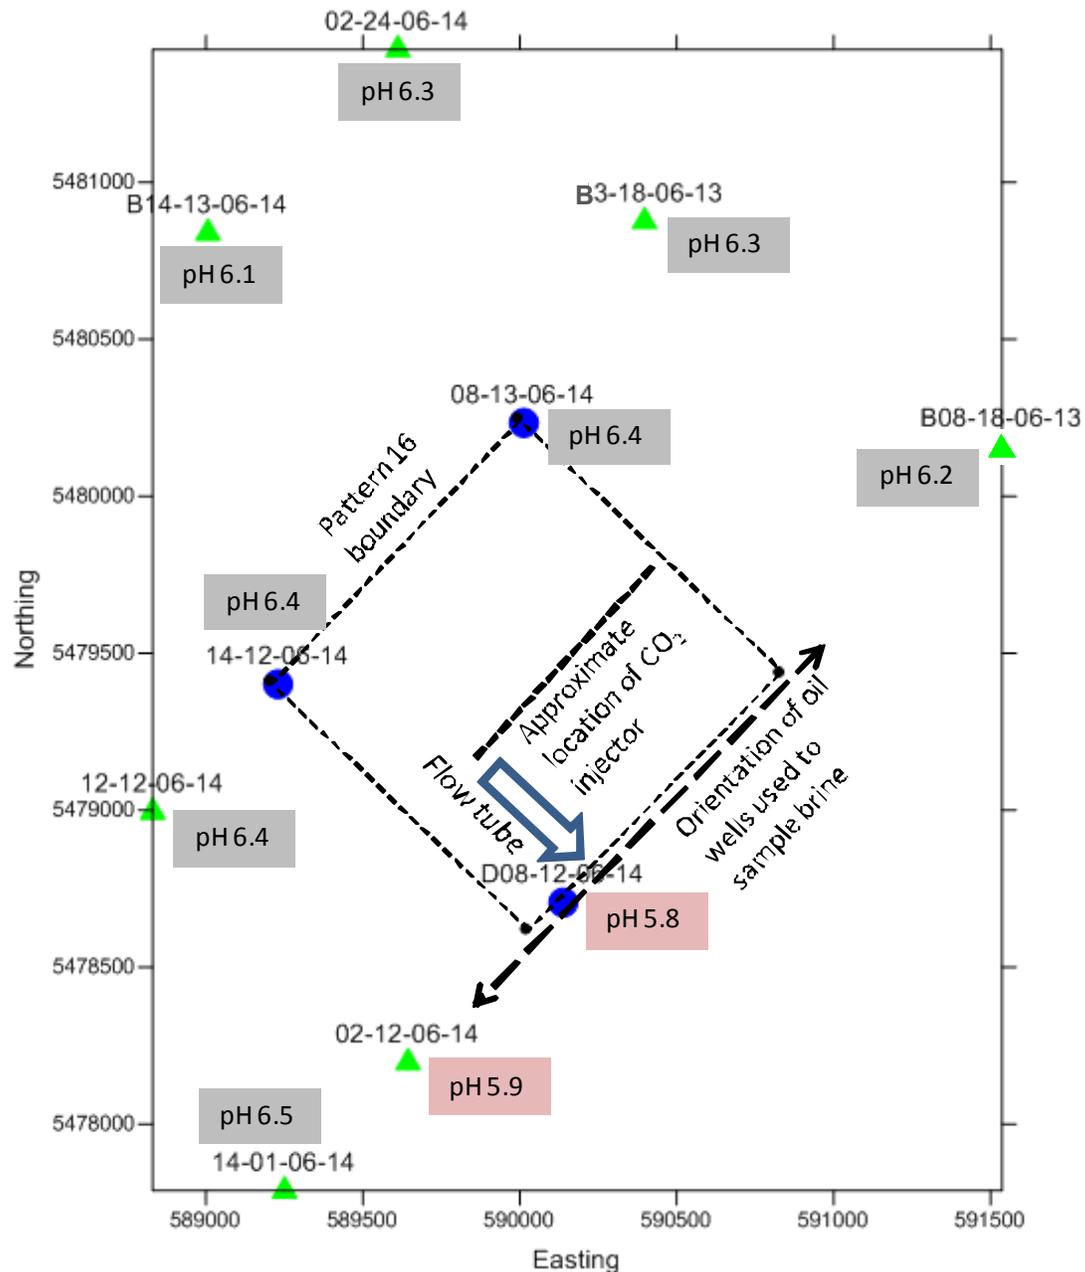


Figure 2.12. Schematic map of the Pattern 16 vicinity, showing approximate well locations within the pattern (blue circles) as well as outside (green triangles); depicted well locations indicate well head locations and do not take into consideration subsurface geometry. Data from well D08-12-06-14 were chosen to inform the inversion because (1) brine chemistry changes consistent with the introduction of CO₂, and (2) proximity to the injector.

Comparisons of best likelihood proposal modeled resultant brine chemistry from both inversions and monitoring data from wells D08-12-06-14 and 02-12-06-14 as a function

of time are shown on Figures 2.16 through 2.20 for pH, Ca^{2+} , Mg^{2+} , dissolved Si, and total dissolved CO_2 , respectively. With the exception of dissolved Si, for which the model is incapable of explaining the apparent concentration increase with time, agreement with the monitoring data is reasonably good. Particularly noteworthy is the inversion results for dissolved CO_2 (Fig. 2.20). Note that total dissolved CO_2 values were not used to calculate the likelihood values that guided the MCMC search. The inverted result agrees well with the trend in the observations thereby suggesting that the inversion is producing internally consistent results.

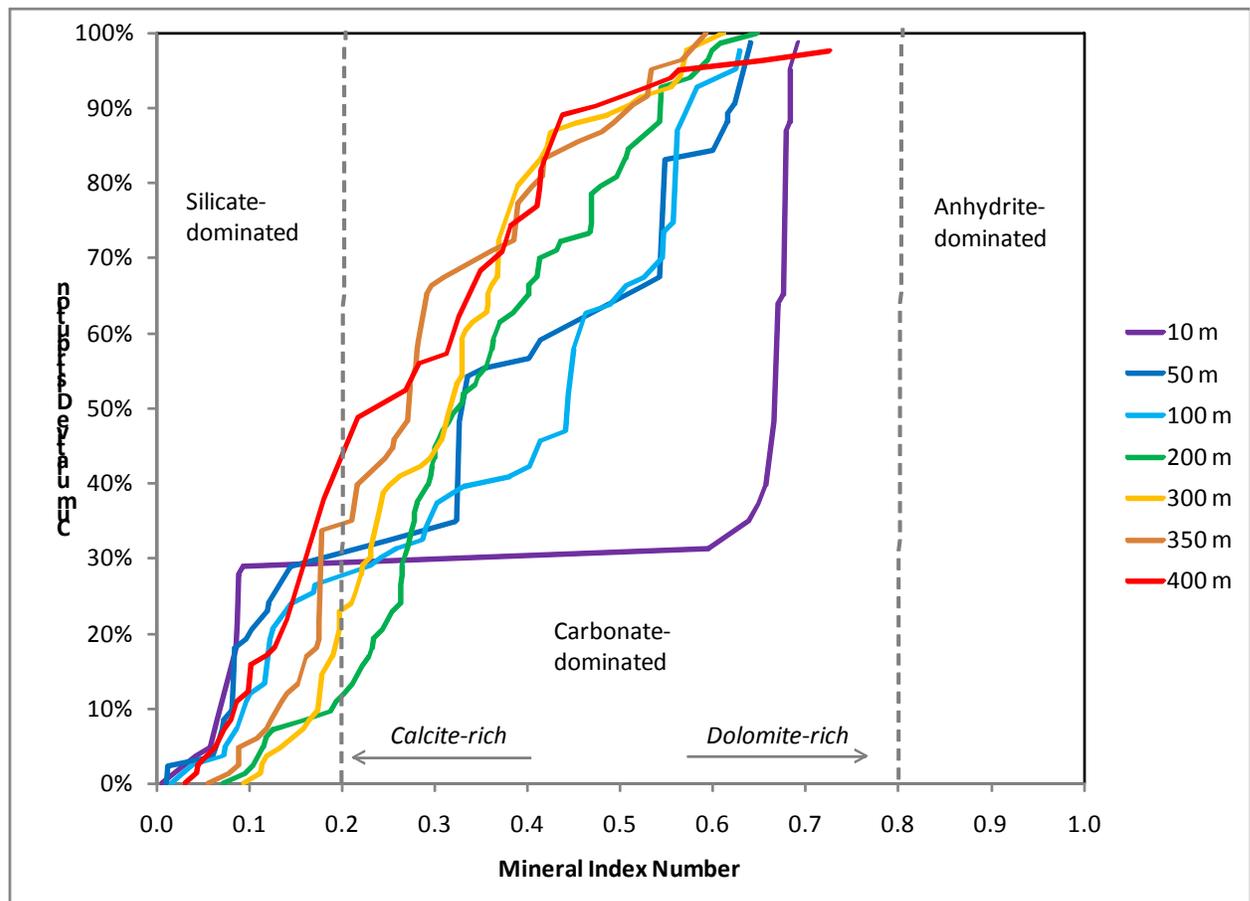


Figure 2.13. Distribution of mineral index numbers for proposals with likelihood function scores above 1.65 (approximately the top 5% of accepted proposals) at different locations along the flow path, inverted from data from Well D08-12-06-14 assuming fixed average intrinsic dissolution rates at pH 5.5. The distributions generally indicate a shift from a reactive dolomite near the influent portion of the flow path to a reactive calcite closer to the effluent end (400 m).

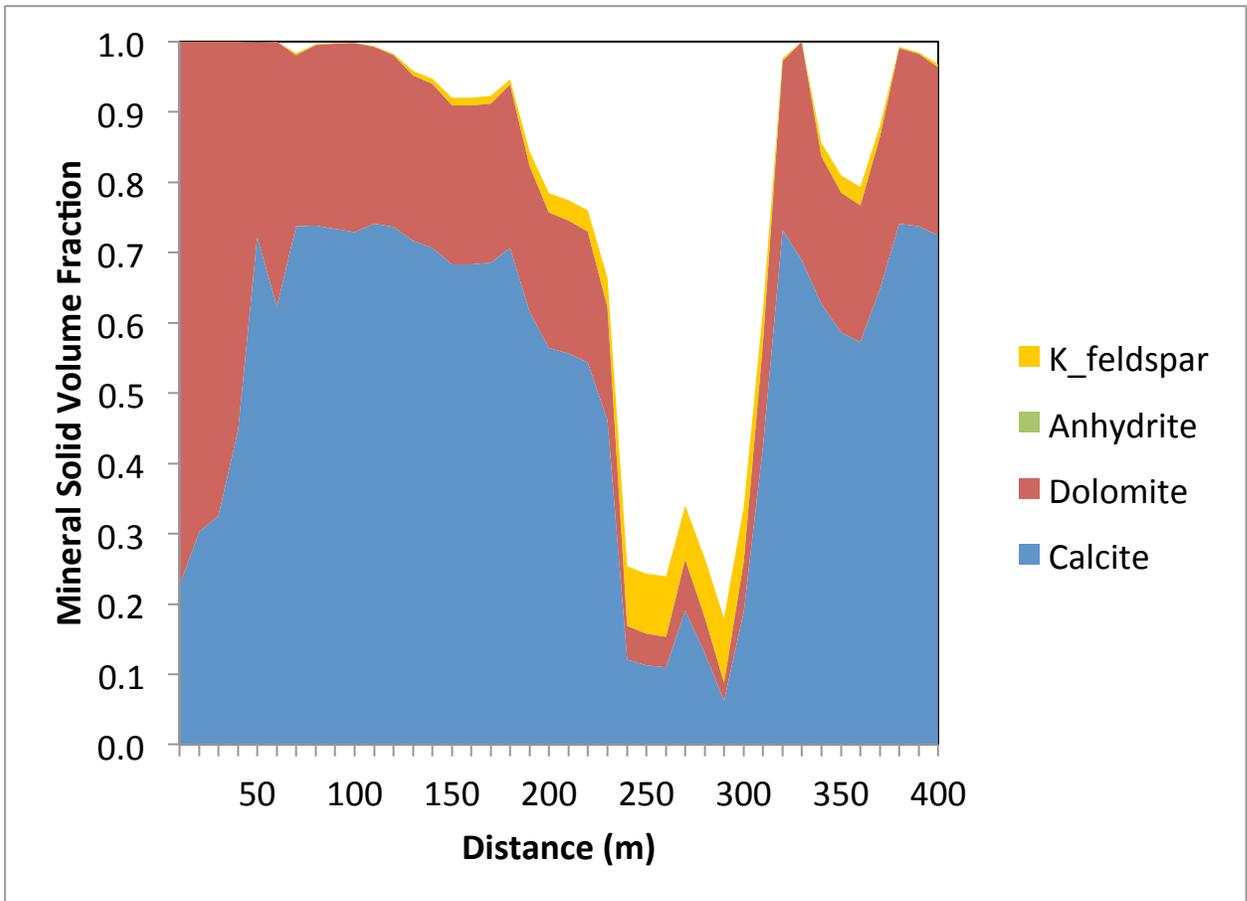


Figure 2.14. Mineralogy distribution along the flow path implied by the best likelihood proposal generated by inverting from the data from Well D08-12-06-14 assuming fixed average intrinsic dissolution rates at pH 5.5. Flow of CO₂-enriched brine is from left to right. Because a silicate-rich mineralogy is not common at $x = 200$ m or $x = 300$ m among the top 5% of the proposals (refer to Figure 2.13), it is likely that the carbonate-poor region predicted to exist between 250 m and 300 m is a low-consequence artifact.

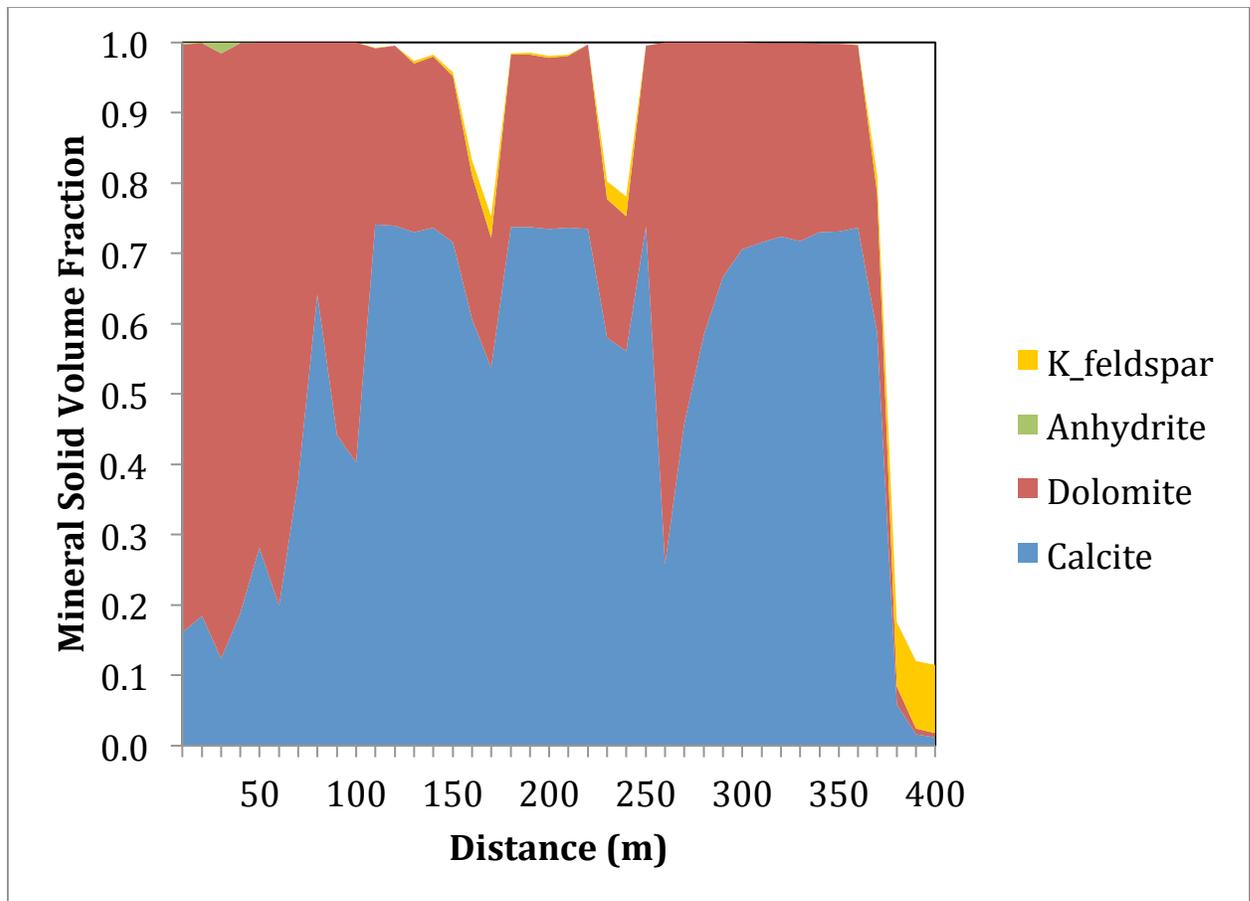


Figure 2.15. Mineralogy distribution along the flow path implied by the best likelihood proposal generated by inverting from the data from Well D08-12-06-14, assuming pH dependence through time and position along column in accordance with Eq.-2.4.

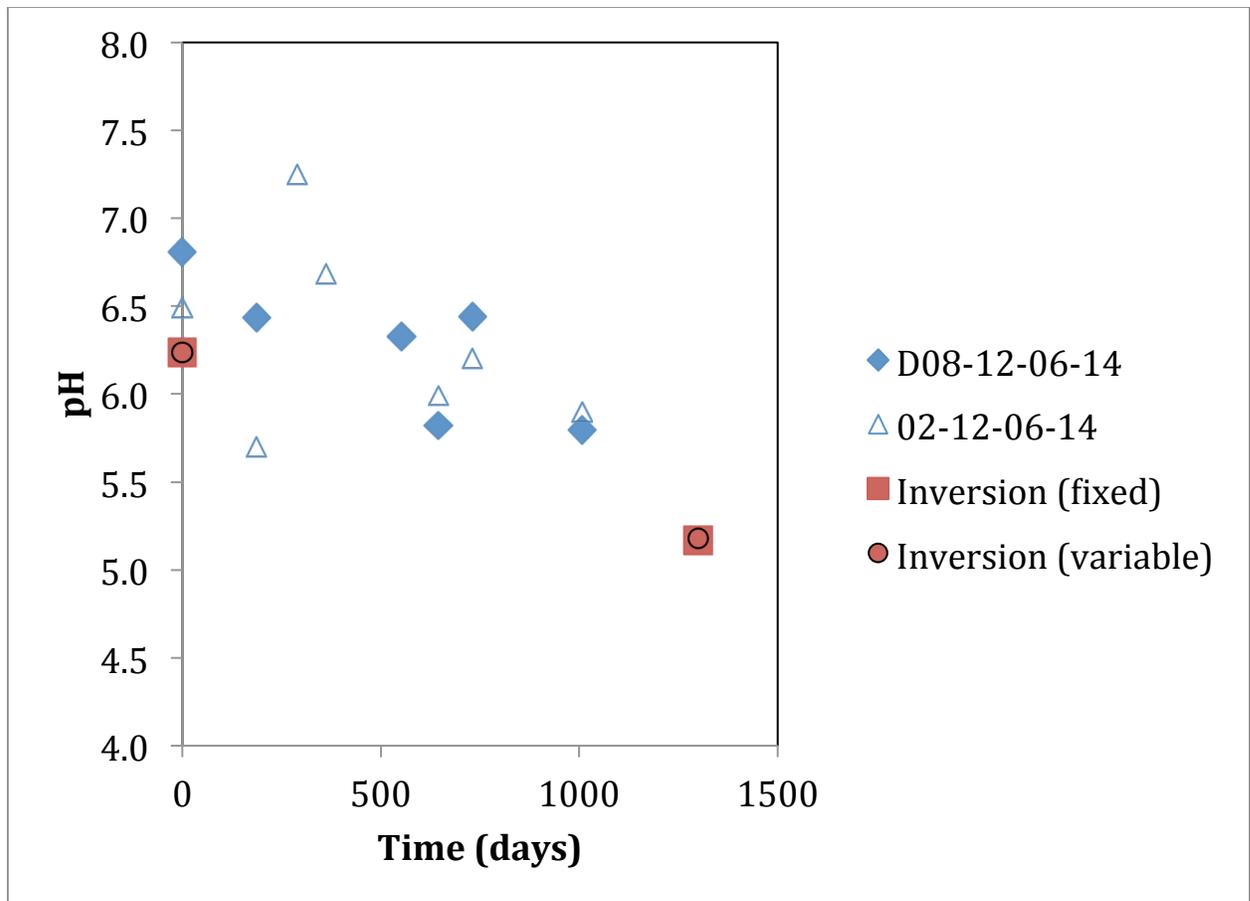


Figure 2.16. Comparison of measured pH, corrected by the procedure discussed in Section 2.2, as a function of time in wells D08-12-06-14 and O2-12-06-14 and modeled pH based on the best likelihood proposal generated by the MCMC inversions using different mineral dissolution kinetic models (i.e., fixed versus variable intrinsic dissolution kinetics). The inversion values at time $t = 0$ represents the initial condition prior to the introduction of CO₂-rich brine into the flow path.

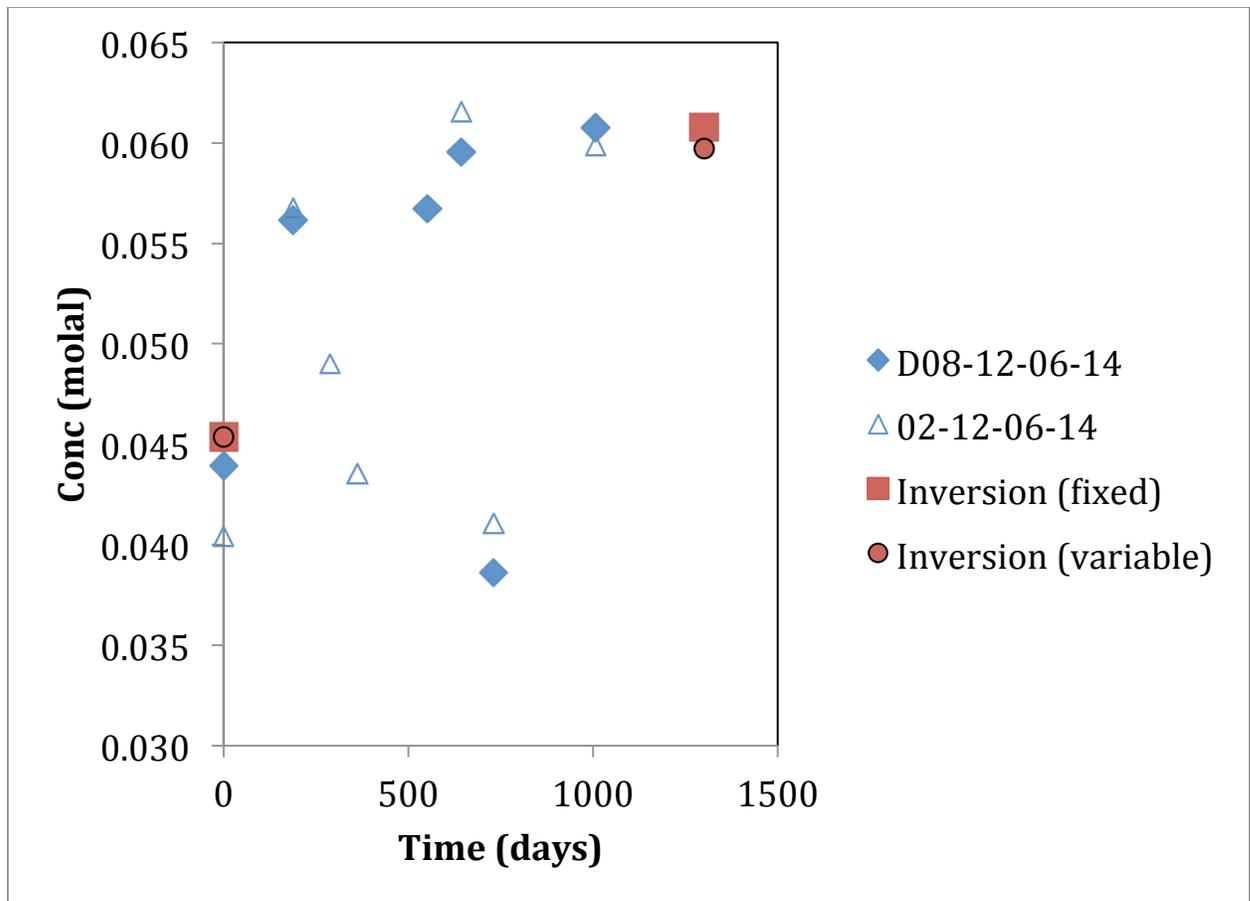


Figure 2.17. Comparison of measured Ca^{2+} as a function of time in wells D08-12-06-14 and O2-12-06-14 and modeled pH based on the best likelihood proposal generated by the MCMC inversions using different mineral dissolution kinetic models (i.e., fixed versus variable intrinsic dissolution kinetics). The inversion values at time $t = 0$ represents the initial condition prior to the introduction of CO_2 -rich brine into the flow path.

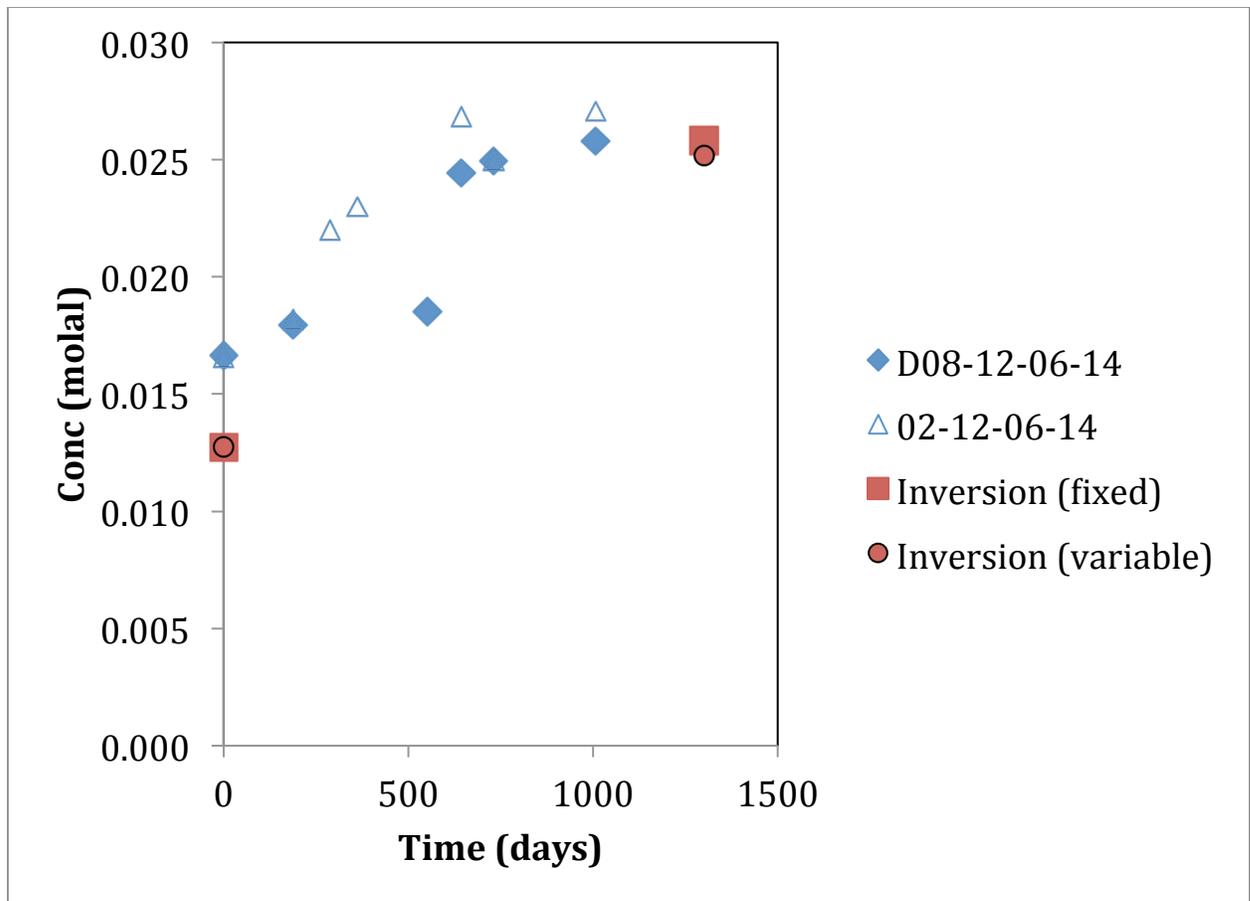


Figure 2.18. Comparison of measured Mg^{2+} as a function of time in wells D08-12-06-14 and O2-12-06-14 and modeled pH based on the best likelihood proposal generated by the MCMC inversions using different mineral dissolution kinetic models (i.e., fixed versus variable intrinsic dissolution kinetics). The inversion values at time $t = 0$ represents the initial condition prior to the introduction of CO_2 -rich brine into the flow path.

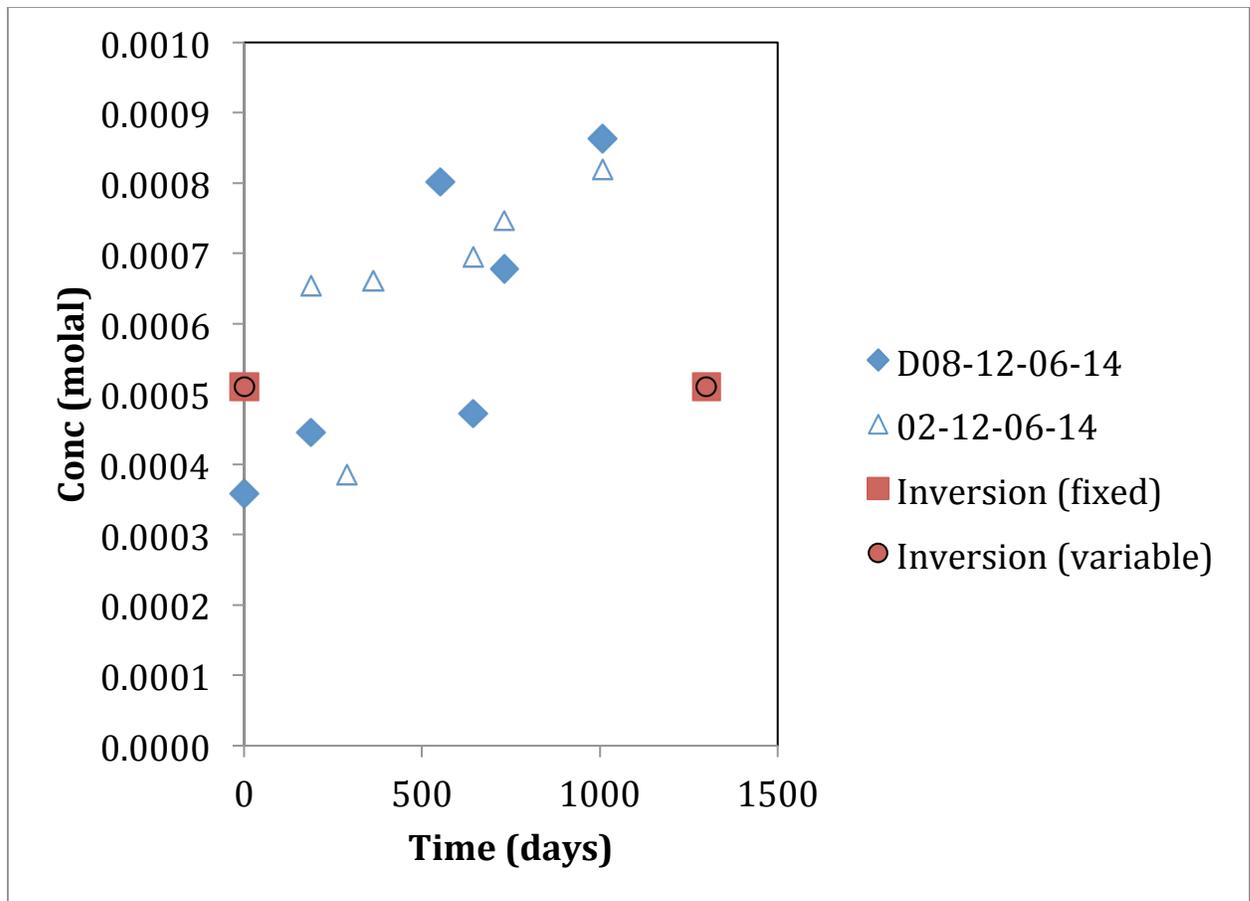


Figure 2.19. Comparison of measured dissolved Si as a function of time in wells D08-12-06-14 and O2-12-06-14 and modeled pH based on the best likelihood proposal generated by the MCMC inversions using different mineral dissolution kinetic models (i.e., fixed versus variable intrinsic dissolution kinetics). The inversion values at time $t = 0$ represents the initial condition prior to the introduction of CO₂-rich brine into the flow path.

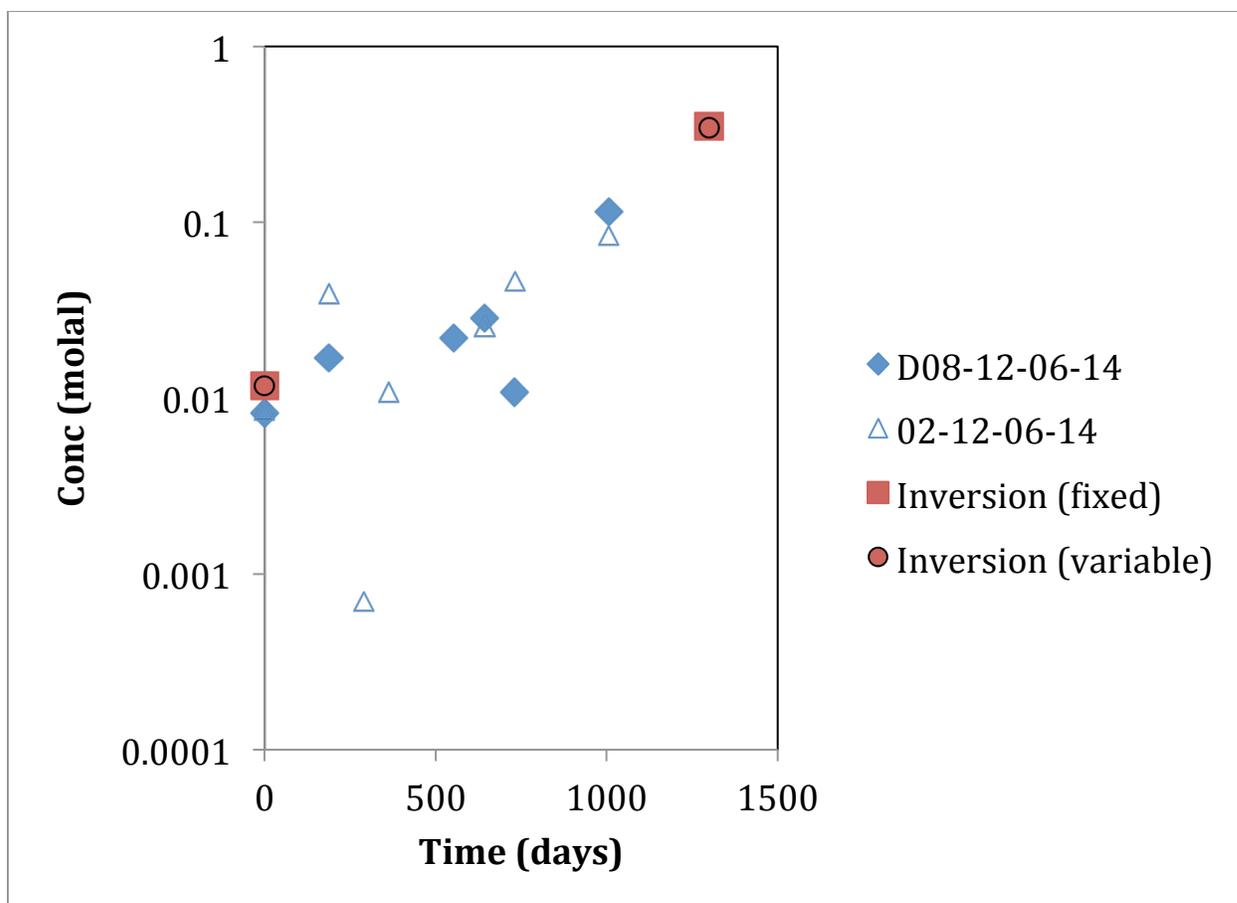


Figure 2.20. Comparison of measured total dissolved CO₂ as a function of time in wells D08-12-06-14 and O2-12-06-14 and modeled pH based on the best likelihood proposal generated by the MCMC inversions using different mineral dissolution kinetic models (i.e., fixed versus variable intrinsic dissolution kinetics). The inversion value at time $t = 0$ represents the initial condition prior to the introduction of CO₂-rich brine into the flow path. Note that total dissolved CO₂ was not included among the metrics used to compute composite likelihood functions used to rank proposals.

2.5 Conclusions

Constrained by the computational burden imposed by a full 3-D model and limited spatial resolution inherent in the monitoring data, the MCMC algorithm was applied to a simplified, 1-D idealization of a reactive transport in the Pattern 16 area to (1) demonstrate the applicability of the methodology, and (2) provide some degree of constraint on mineral dissolution rates. It is important to recognize that the inversion approach as employed herein is highly idealized, not only in assuming 1-D flow but also in assuming that fluid flow is single-phase. While the single-phase model can capture salient features of the water-rock interactions by assuming the influx of a CO₂-enriched brine, it neglects the multiphase physics of the flow system. Nonetheless, inversion of the inherently low-resolution Pattern 16 geochemistry data indicates a reactive flow path between the injection well and well D08-12-06-14 that is characterized by interactions with carbonate minerals, predominantly dolomite-rich near the injection well and

progressing to more calcite-rich further downgradient. This observation is consistent, for example, with the initial interaction of the CO₂-rich brine with the dolomite-rich Marly in contact with the injection well, followed by subsequent interaction with the calcite-rich Vuggy along the flow path (Figure 2.21). With the exception of dissolved Si, the agreement between the modeled and observed geochemical indicators implies a reasonable degree of self-consistency between the intrinsic reaction rates calculated via Eq.-2.4 and the Palandri and Kharaka (2004) data set, the fracture porosity, the mineral specific surfaces, and the assumed cubic relationship between surface area and mineral mass. The inversions do not appear to be sensitive to the pH dependence of the dissolution kinetics implied by Eq.-2.4.

In summary, the MCMC approach for inversion of reactive transport for the geochemical conditions specific to the Weyburn-Midale reservoir appears capable of yielding plausible results that are consistent with data and constrained by independent parameter estimates. Application is tempered, however, by computational constraints, a complex operating history, and – significantly – by limited spatial resolution associated with brine samples. Therefore, this methodology is perhaps most suited as a means for independently estimating lumped dissolution rates at the field scale for comparison with laboratory data.

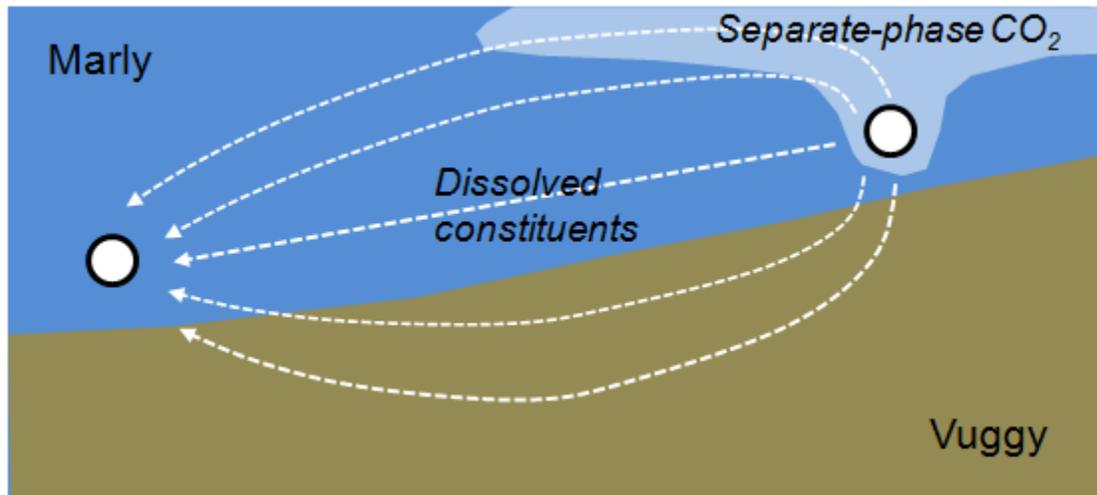


Figure 2.21. A schematic model of fluid flow between the CO₂ injector and the monitoring well (not to scale). While supercritical CO₂ will migrate upwards as a result of buoyancy, brine will also be displaced in response to the applied pressure gradient. Dissolution of dolomite occurs at the point of injection in the Marly, followed by subsequent reaction with calcite where the flow path crosses into the more permeable Vuggy.

Summary and Conclusions

The goals of this project were to develop and use stochastic inversion algorithms that would integrate seismic reflection and aqueous chemistry measurements made during Phase 1A. This final report and previous deliverables describe the inversion approaches we developed and the inversion results obtained. The seismic data was used to reconstruct the reservoir's permeability and porosity in three dimensions. The aqueous chemistry data were used to invert for mineral volume fractions along 1D flow tubes in the reservoir.

A key challenge associated with the stochastic inversion using seismic data and chemistry is its computational expense. We used large amount of project resources to mitigate this problem as much as possible. The seismic runs required about 9 days and used 112 compute cores. The reactive transport runs would have taken hundreds of days if we had continued with our original plan of performing full 3D calculations. Instead, we simplified the problem to reactive transport along 1D tubes and assumed that a single liquid phase was present in the pore space. These simplifications reduced the run times to about 2.0 days.

A second challenge is that the seismic waveforms are more sensitive to the bulk modulus, shear modulus and density than they are to the permeability, porosity and CO₂. This means that it is necessary to obtain accurate estimates of seismic moduli and density before using seismic data to invert for permeability and porosity. Our original plan was to use homogeneous bulk and shear moduli for each layer (from either well logs or laboratory measurements) but this plan proved inadequate because these properties exhibit substantial heterogeneity.

The seismic inversions reduced the misfit between the predicted and observed waveforms but the misfit improvements were less than expected. We believe that the root cause of this behavior stems from use of homogeneous bulk and shear moduli for each layer. We originally believed that homogeneous values would be good enough, but discovered late in the project that our assumption was incorrect and quickly developed a stochastic inversion scheme to solve for bulk modulus, shear modulus and density, using the baseline 1999 reflection survey data (pre- CO₂ injection). By the time the project ended, this new inversion approach was unable to improve the misfit. We believe that this moduli inversion will work, given additional time and resources to continue development.

The seismic inversions described in this report used homogeneous moduli values that were adjusted by trial and error to reduce the waveform misfit as much as possible. The waveform misfits improved after this was done but were still poor. This issue probably affected the accuracy of the inverted porosity/permeability fields shown here; the magnitude of this effect has not been quantified.

The seismic inversions suggest that the permeability and porosity are spatially heterogeneous in the Marly and Vuggy layers, as expected. The Marly permeabilities tend to be smaller than the Vuggy's whereas the Marly porosities tend to be larger. Pre and post inversion histograms of permeability and porosity look very similar, thereby

suggesting that the prior constraints determine the modes of the distributions (permeability/porosity values that occur most frequently) in the inversions. The seismic data did influence the locations of the permeability highs and lows within each layer.

The aqueous chemistry inversions indicate that the reservoir's reactive mineralogy is dominated by calcite and dolomite, as expected. Inversion of the inherently low-resolution geochemistry data indicates a reactive flow path between the injection well and well D08-12-06-14 that is characterized by interactions with carbonate minerals, predominantly dolomite-rich rock near the injection well and progressing to more calcite-rich further down-gradient. This observation is consistent, for example, with the initial interaction of the CO₂-rich brine with the dolomite-rich Marly in contact with the injection well, followed by subsequent interaction with the calcite-rich Vuggy along the flow path. The inversions also suggest that mineral volume fractions are spatially heterogeneous, ranging in value from 0.07 to 0.7 (calcite) and 0.02 to 0.75 (dolomite). The inversions suggest that anhydrite dissolution or precipitation is not a major contributor to the observed brine chemistry. Comparison of the predicted and observed brine chemistries indicates that pH, Ca²⁺, Mg²⁺ and dissolved CO₂ agree reasonably well; dissolved Si does not.

Implications for Best Practices Manual

The seismic waveforms are more sensitive to the bulk modulus, shear modulus and density than they are to permeability and porosity. This means that it is necessary to obtain good estimates of seismic moduli and density before using seismic data to invert for permeability and porosity.

Synthetic data tests (Ramirez et al, 2010) suggest that it may be possible to estimate permeability in reservoir regions that do not contain CO₂ but where significant pressure changes are caused by CO₂ injection. This possibility exists because bulk and shear moduli change with reservoir pressure.

Given currently available reactive transport simulators, it is unfeasible to use MCMC approach to invert for mineral volume fractions and rate constants using the 3D volumes we considered. 1D inversions are currently feasible and can yield plausible results that are consistent with data and constrained by independent parameter estimates. Application is tempered, however, by computational constraints, a complex operating history, and – significantly – by limited spatial resolution associated with brine samples. Therefore, this methodology is perhaps most suited as a means for independently estimating lumped dissolution rates at the field scale for comparison with laboratory data.

Acknowledgements

We are grateful for the funding support provided by the Petroleum Technology Research Center, Saskatchewan, Canada. We are also grateful for the data provided by Barbara Dietiker (Geological Survey of Canada) and Erik Nickel (Saskatchewan Ministry of

Energy and Resources). We also want to acknowledge the contributions of colleague Jeff Wagoner (LLNL) who provided various adaptations of the reservoir model originally developed by Cenovus.

References

- Brown, L. 2002, *Integration of rock physics and reservoir simulation for the interpretation of time-lapse seismic data at Weyburn Field, Saskatchewan*, M. Sc. thesis, Colorado School of Mines, Golden, CO.
- Cardona, R., Ozkan, E., and Batzle, M., 2002. *Pressure-transient experiments and the elastic characterization of fractures*, 2002 SEG Annual Meeting Expanded Abstracts.
- Durocher, K.E., Kotzer, T.G., and Whittaker, S.G., 2005. Physical and chemical characterization of subterranean CO₂ storage sites using synchrotron-based computed microtomography techniques; in Summary of Investigations 2005, Volume 1, Saskatchewan Geological Survey, Sask. Industry Resources, Misc. Rep. 2005-4.1, CD-ROM, Paper A-3, 7p.
- Emberley, S., I. Hutcheon, M. Shevalier, K. Durocher a, B. Mayer a, W.D. Gunter, and E.H. Perkins, 2005. Monitoring of fluid–rock interaction and CO₂ storage through produced fluid sampling at the Weyburn CO₂-injection enhanced oil recovery site, Saskatchewan, Canada, *Applied Geochemistry*, 20, 1131–1157.
- Palandri, J., Kharaka, Y.K., 2004. *A compilation of rate parameters of water–mineral interaction kinetics for application to geochemical modeling*, U.S. Geol. Surv. Open File Report 2004-1068, 64 pp.
- Parkhurst, D.L. and Appelo, C.A.J., 1999. *User's Guide to PHREEQC (version 2)--A Computer program for Speciation, Batch-Reaction, One-Dimensional Transport, and Inverse Geochemical Calculations*, U.S. Geological Survey Water-Resources Investigations Report 99-4259, 312 p.
- Parkhurst, D.L., K. L. Kipp, P. Engesgaard, and S. R. Charlton, 2004. PHAST—*A Program for Simulating Ground-Water Flow, Solute Transport, and Multicomponent Geochemical Reactions*, Techniques and Methods 6–A8, U.S. Geological Survey, 154 p.
- PTRC, 2004, *IEA CHG Weyburn CO₂ Monitoring & Storage Project Summary Report 2000-2004*, M. Wilson and M. Monea, Eds., Petroleum Technical Research Centre, 283 p.
- Quan, Y., and J. Harris, 2008, Stochastic seismic inversion using both waveform and travelttime data and its application to time-lapse monitoring, SEG Expanded Abstracts 27, 1915 (2008); doi:10.1190/1.3059273
- Ramirez, A., Y. Hao, D. White, S. Carle, K. Dyer, X. Yang, W. McNab, W. Foxall, and J. Johnson, 2009, Development of a Stochastic Inversion Tool To Optimize Agreement

Between The Observed And Predicted Seismic Response To CO₂ Injection/Migration in the Weyburn-Midale Project, Lawrence Livermore National Laboratory, LLNL-TR-420989, Livermore CA.

Ramirez, A., K. Dyer, D. White, Y. Hao, and X. Yang, 2010, Model verification: synthetic single pattern simulations using seismic reflection data, Lawrence Livermore National Laboratory, LLNL-TR-441596, Livermore CA.

Ramirez, A., W. McNab, S. Carle, Y. Hao, D. White, J. Johnson, 2011, Progress Report, December 2010: Improved Site Characterization And Storage Prediction Through Stochastic Inversion Of Time-Lapse Geophysical And Geochemical Data, Lawrence Livermore National Laboratory, LLNL-TR-464614, Livermore, CA 94550.

Appendix

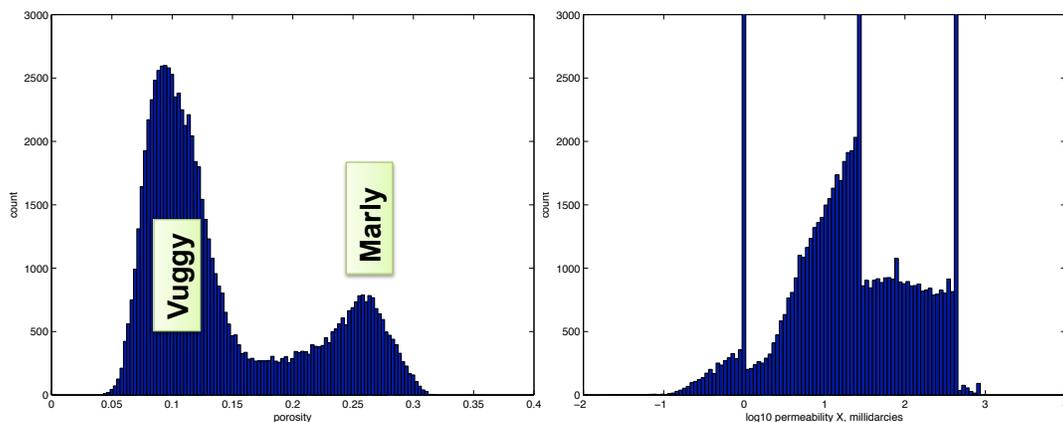


Figure A1. Histograms of porosity (left) and permeability (right) in Cenovus' model, calibrated against production and injection history.

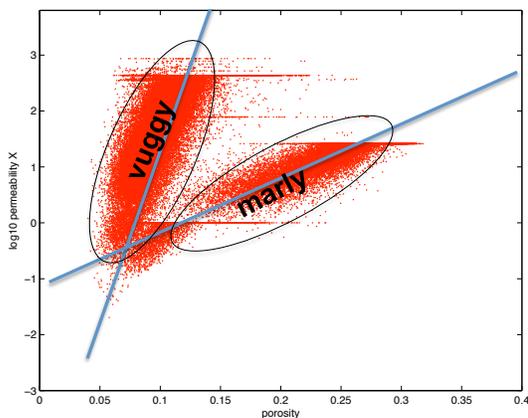


Figure A.2 shows the correlation between porosity and permeability expressed in Cenovus calibrated model.

