An Optimization Approach to Design Monitoring Schemes for CO₂ Leakage Detection

Catherine M. R. Yonkofski^{a*}, Jason A. Gastelum^{a*}, Ellen A. Porter^a, Luke R. Rodriguez^a, Diana H. Bacon^a, Christopher F. Brown^a

^a. Pacific Northwest National Laboratory, Richland, WA 99352, USA

* Correspondence to Catherine Ruprecht Yonkofski, catherine.ruprecht@pnnl.gov

Abstract

This paper demonstrates an approach to identify optimal monitoring designs that minimize the time to first detection of CO₂ leakage from a subsurface storage formation. This research is part of the National Risk Assessment Partnership (NRAP), a DOE project tasked with conducting risk and uncertainty analysis in the areas of reservoir performance, natural leakage pathways, wellbore integrity, groundwater protection, monitoring, and systems level modeling. Our approach applies a simulated annealing algorithm that searches the solution space by iteratively mutating potential monitoring designs. A user-defined leakage signature based on change from initial conditions is implemented to infer CO₂ leakage had occurred. An example application was performed to demonstrate the effectiveness and efficiency of this method compared to an exhaustive search of the entire solution space. We accounted for uncertainty in the example by evaluating the performance of potential monitoring designs across a set of simulated leakage realizations. The example application approached the optimal time to leakage detection in a few minutes using a standard workstation, which was several orders of magnitude faster than an exhaustive search of the solution space, thus exhibiting efficiency and effectiveness. A user-friendly tool, DREAM (Designs for Risk Evaluation and Management), is being developed for use on personal computers to make this method accessible to stakeholders, regulators, and researchers.

Keywords: CO₂ storage, carbon capture and sequestration; CO₂ leakage detection; monitoring verification and accounting; MVA; monitoring systems design; risk and uncertainty

1 **1 Introduction**

2 As advances toward industrial scale deployment of geologic CO₂ storage technologies are
3 made, regulatory requirements have been developed to ensure the safety of United States
4 drinking water aquifers (USDWs) (75 FR 77229, 2010). These requirements are discussed as
5 part of a series of technical guidance documents to support owners and operators of Class VI
6 CO₂ injection wells and Underground Injection Control (UIC) Program permitting authorities
7 (U.S. EPA, 2015). Among other requirements, Class VI well owners and operators will be
8 required to complete Testing and Monitoring Plans particular to each storage site (75 FR 77229,
9 2010, U.S. EPA, 2013). These plans must detail methods to detect signs of leakage during the
10 injection and post-injection phases above the confining zone and within overlying USDWs (75
11 FR 77229, 2010, U.S. EPA, 2013). Similarly, the International Energy Agency (IEA) published a
12 CO₂ storage regulatory framework (2010) including monitoring, reporting, and verification
13 requirements calling for monitoring plans that detail the rationale behind monitoring
14 technologies, techniques, parameters, and locations selected.

15 This research is part of the National Risk Assessment Partnership (NRAP), a DOE project
16 tasked with conducting risk and uncertainty analysis in the areas of reservoir performance,
17 natural leakage pathways, wellbore integrity, groundwater protection, monitoring, and systems
18 level modeling. NRAP is working to develop risk assessment tools that could be used for
19 compliance with federal requirements and ultimately lead to safe, permanent geologic CO₂
20 storage. Reduced-order-models (ROMs) are being developed to predict the behavior of more
21 complex hydrological systems as part of this effort (Bacon, 2013; Keating et al., 2013; Dai et al.,
22 2014; Carroll et al., 2014, Hou et al., 2014). The Edwards Aquifer and the High Plains Aquifer

were selected as characteristic shallow USDWs to base hydrologic systems for a suite of these NRAP models, representing unfractured carbonate aquifers and confined alluvium aquifers, respectively (Bacon, 2013; Keating et al., 2013; Carroll et al., 2014, Hou et al., 2014). The ROMs were based on complex reactive transport models that were developed for each site by the respective NRAP research teams. The complex model development is of particular interest to this study, as multiple realizations used to develop ROMs can also be used to develop site specific monitoring strategies.

Models simulated CO₂ and brine leakage into the Edwards and the High Plains aquifer systems to describe potential impacts to groundwater quality, quantify thresholds relating to these impacts, describe the time scale for impact, and discuss the probability of detecting a groundwater plume should leakage occur (Bacon, 2013; Last et al., 2013; Carroll et al. 2014; Hou et al., 2014; Last et al., 2015). Hypothetical scenarios assumed that wellbores were the most likely conduits for brine and CO₂ leakage (Bacon, 2013; Last et al., 2013; Carroll et al. 2014; Hou et al., 2014; Last et al., 2015). One set of simulations varied hydrogeologic parameters, leakage rates, and total dissolved solids (TDS) and trace metal concentrations of the leakage source, while treating dissolved components as nonreactive tracers (Carroll et al., 2014). Subsequent leakage plumes were likely to persist for long periods of time (>200 years) in the absence of remediation (Carroll et al., 2014). In addition, the probability of detecting plumes using existing wells to sample the groundwater chemistry was very low, because the plumes were relatively small in both aquifers (Carroll et al., 2014). These results highlight the importance of early leakage detection and the need to develop methodologies that prevent and/or directly detect leakage prior to reaching USDWs. A challenge raised by Carroll et al. (2014) is

the need to develop spatially diverse monitoring designs capable of efficiently detecting leakage before persistence in USDWs.

Also as part of the NRAP initiative, Hou et al. (2014) developed complex hydrologic models of the Edwards Aquifer which focused on geochemical uncertainties due to CO₂ and brine leakage scenarios through an abandoned well. The models included equilibrium, kinetic mineral, and adsorption reactions related to the carbonate and clay minerals in the aquifer reacting with major ions and trace metals in groundwater, as well as CO₂ and brine leaking from the wellbore (Hou et al., 2014). Geochemical parameters and wellbore ROM parameters were varied to generate exploratory samples (Hou et al., 2014). For each of these input samples, the STOMP-CO₂ (White et al., 2012) simulator was used to predict the length, width, height, and volume of pH<6.5 plume in the aquifer as well as the CO₂ flux to the atmosphere. A random set of realizations generated from this model was used by this study to demonstrate our proposed monitoring design strategy.

The major objective of this paper is to present a systematic approach for identifying robust site specific monitoring designs using existing data generated from physics-based simulations of CO₂ leakage. The monitoring designs generated by this approach are proven near-optimal with respect to the objective of minimizing the time to first detection of CO₂ leakage from a hypothetical storage formation assuming reasonable impact thresholds. Additionally, methods were developed under computational constraints to meet NRAP goals of providing tools that are accessible to all users via personal computers. The monitoring designs presented are meant to serve as an example application of the discussed optimization method and may not

represent realistic optimal monitoring designs, as appropriate configurations will be highly site specific.

2. 1. Background

Simulation-based optimization is a common approach to solving geoscience decision problems. Optimization strategies can be divided into four categories: (1) gradient-based algorithms that rely on derivative information; (2) heuristic algorithms that follow empirical guidelines and incorporate elements of structured randomness; (3) direct search algorithms that utilize deterministic stencil-based procedures; and (4) hybrid methods that combine aspects of the previous three categories (Matott et al., 2011). Typical CO₂ storage related optimization studies seek to optimize CO₂ and/or water injection strategies by maximizing CO₂ storage efficiency, minimizing cost, or maximizing secondary oil and gas production.

Recent examples of each of the four optimization categories applied to subsurface CO₂ injection applications demonstrate various approaches to these problems. An example of a gradient based method was used by Ruprecht (2014), where the Levenberg-Marquardt minimization algorithm was used to minimize cost while maximizing secondary trapping of CO₂ through varied water injection rates. Zhang and Agrawal (2013) took a heuristic approach in developing GATOUGH, a genetic algorithm based optimizer for the TOUGH2 simulator that can be used to optimize CO₂ storage capacity. Cameron and Durlofsky (2012) used the Hooke-Jeeves Direct Search algorithm to minimize the long-term fraction of mobile CO₂ in a storage reservoir by optimizing the locations and operating parameters for injection and brine production wells. Nghiem et al. (2009) also used a direct search method, performing full enumeration of a small parameter space to maximize the rate of CO₂ trapping by varying brine injection well

locations and injection rates. Jahangiri and Zhang (2012) worked to co-optimize oil production and CO₂ sequestration by using an ensemble based hybrid technique. This technique computed the gradient of the objective function through Monte Carlo evaluation over an ensemble of the states corresponding to the uncertainties in the description of the model (Jahangiri and Zhang, 2012). Each of these examples attempts to optimize CO₂ storage performance within the storage formation.

This study considered CO₂ leakage scenarios in overlying aquifers in an attempt to design a monitoring scheme to minimize the time to leakage detection. Previous studies by Sun et al. (2012, 2013) similarly attempt to optimize monitoring networks based on pressure-based leakage detection in wells. The method discussed in this paper is unique in that it allows for simultaneous consideration of an unlimited number of monitoring parameters and determines monitoring locations in all three dimensions. We also provide the ability to require multiple sensors of potentially distinct types to signal leakage to reduce the risk of false positives and increase the robustness of our solution. The more substantive differences come in the approach to the optimization problem and are discussed in Section 4 after the example problem is demonstrated.

One of the foremost challenges inherent to these types of subsurface optimization tasks is accounting for the uncertainty in the geology of the target CO₂ storage site. Our approach addresses uncertainty by evaluating the performance of potential monitoring designs across multiple simulated leakage realizations. Each realization represents a hypothetical potential leakage scenario. By optimizing over the set of potential leakage scenarios, we identify monitoring designs robust to an unknown future. A secondary challenge specific to this study is the sheer number of potential monitoring designs to consider. Monitoring designs will likely

consist of dozens of individual sensors which could be emplaced anywhere in the area of review. Consequently, the number of monitoring design combinations precludes complete enumeration of the solution space. As a result, our approach applies a simulated annealing heuristic method to approximate the optimal solution.

We present an approach for finding monitoring designs that minimize the expected time to first detection ($E[TFD]$) of CO_2 leakage. First, we formally define the objective function and describe the components of the algorithm. Secondly, we demonstrate that this approach identifies near-optimal monitoring designs by presenting a test case comparison of $E[TFD]$ of monitoring designs identified by this approach with $E[TFD]$ for every feasible design in the solution space.

2 Methodology

This section describes the objective function and summarizes our approach for optimizing monitoring design for a CO_2 storage site. Our decision variable is the monitoring design, defined as a collection of sensors of different types deployed at specified locations within an aquifer overlying a storage formation. The “sensor-types” are not explicitly types of sensors (ie. CO_2 soil flux chambers), but relate to the monitoring parameters that various sensor-types would detect. These parameters may include pressure, temperature, gas saturation, dissolved component concentrations, or any other quantity that can be modeled in a physics-based simulation of CO_2 transport and detected by borehole monitoring technologies. Sensor-types included in the monitoring design are based on the parameters considered in the CO_2 leakage realizations provided as input to the optimization application. From this set of calculated quantities, the user must select which sensor-types to include in the monitoring design as notional sensors capable of measuring their target property to arbitrary precision. The performance of a monitoring design is

computed for an individual realization by calculating the time to first detection of CO₂ leakage, using the criteria defined in a specified leakage signature. Finally, the objective function value for a monitoring design is calculated by applying the expectation function to the vector of times to first detection, $E[TFD]$, across the set of realizations.

2.1. Optimization Formulation

The optimal monitoring design is defined as the monitoring design that minimizes the $E[TFD]$ of CO₂ leakage from the storage reservoir into an overlying aquifer. The objective function can be expressed using the convention of Wang et al. (2011) as follows:

$$\text{Minimize}_{\mathbf{x}} Z(\mathbf{x}, \Omega) = E_{\Omega}[TFD(\mathbf{x}, \omega)] \quad (1)$$

where the decision variable, $\mathbf{x}$, specifies the locations and types of sensors that constitute a monitoring design, Ω represents the set of subsurface simulation realizations, ω represents an individual realization from the set Ω , E_{Ω} is the expectation function over the set of subsurface simulation realizations Ω , and TFD is a function that calculates the time to first detection of CO₂ leakage for a given monitoring design $\mathbf{x}$ and realization ω . This function queries the result of a subsurface simulation realization and uses a leakage signature, discussed in detail in Section 2.4, to determine the earliest time at which leakage can be inferred. Given N simulation realizations ω , indexed by i , the objective function can be rewritten as:

$$\sum_{i=1}^N p(\omega_i) TFD(\mathbf{x}, \omega_i) \quad (2)$$

where $p(\omega_i)$ is the probability of realization ω_i . The optimization is bounded by the following constraints:

$$C(\mathbf{x}) \leq \alpha \quad (3)$$

$$W(\mathbf{x}) \leq \beta \quad (4)$$

where $C(\mathbf{x})$ is the cost of a monitoring design, calculated as the sum of the costs of the constituent sensors, and $W(\mathbf{x})$ is the number of wells in a monitoring design, calculated as the count of the unique x-y coordinate pairs for the sensors in the configuration. It is assumed that all monitoring wells for sensor placement are vertical, which is a likely reality of projects trying to avoid excessive costs. This assumption does not relate to the simulated leaky wells, which could allow for results from a leakage pathway of any orientation.

2.2. *Simulated Annealing*

Since complete enumeration of the solution space may not be feasible at a scale easily realized on a personal computer, the optimal solution is approximated using simulated annealing. Matott et al. (2011) found that simulated annealing performs moderately well amongst the algorithms they evaluated in their comparison of optimizers for geoscience applications. Similarly, in their comparison of several algorithms for optimizing well placement in the context of recovery, Bangerth et al. (2006) found that simulated annealing is very efficient in finding near-optimal solutions.

Simulated annealing is an iterative search heuristic analogous to the physical process of annealing. It utilizes a mutation function implemented as follows:

- If the budget allows, first attempt to add an additional sensor into the monitoring design.
- If multiple sensor-types are allowed in the design, randomly select the new sensor-type.
- If the sensor budget is exhausted, attempt to move a randomly selected sensor currently in the design to an adjacent node. This step is randomly skipped 20% of the time.

- If none of the adjacent nodes are available - because they already contain a sensor of that type - or the previous step is randomly skipped, remove this sensor. Add a new sensor of any type within the sensor budget to a node in the set of nodes that could result in an alarm, given the leakage signature.

The algorithm begins by setting configurations $\mathbf{x}_b$, $\mathbf{x}_c$, and $\mathbf{x}_n$ equal to an initial configuration, $\mathbf{x}_i$. In each iteration, the current solution, $\mathbf{x}_c$, is randomly mutated to generate a new solution, $\mathbf{x}_n$. In the context of our study, this means the location of a sensor in configuration $\mathbf{x}_c$ is changed. The mutation function is designed such that $\mathbf{x}_n$ is guaranteed to satisfy the specified cost and well constraints. The objective function values of both solutions are calculated by querying the results of the subsurface simulation realizations and applying the criteria specified in the leakage signature. If the $E[\text{TFD}]$ for $\mathbf{x}_n$ is less than that of $\mathbf{x}_c$, it replaces $\mathbf{x}_c$. The algorithm then checks to see if $\mathbf{x}_c$ is less than the best incumbent solution, $\mathbf{x}_b$. If so, $\mathbf{x}_c$ replaces $\mathbf{x}_b$. The algorithm then proceeds to the next iteration.

If $\mathbf{x}_n$ is greater than the current solution, the algorithm may nevertheless accept the solution as means of avoiding convergence to local optima. To decide whether to accept an inferior solution, the algorithm calculates an acceptance probability, based on the difference in objective function values between the two candidate solutions and the “temperature”, T , of the iteration (a strictly decreasing function on iteration named from original annealing applications) and compares it to a randomly generated number between zero and one. If the random number is less than or equal to the acceptance probability, $\mathbf{x}_c$ is replaced by $\mathbf{x}_n$, otherwise the algorithm proceeds to the next iteration with a decreased T -value. Given this logic, as the algorithm progresses, it becomes decreasingly likely that a worse solution will be accepted as $\mathbf{x}_c$.

The algorithm continues to loop until the specified number of iterations, k , have been executed. The algorithm returns $\mathbf{x}_b$, the best solution the algorithm has evaluated. Pseudo-code describing the steps of the algorithm is presented in Figure 1.

```

Set  $\mathbf{x}_b = \mathbf{x}_c = \mathbf{x}_n = \mathbf{x}_i$ 
Set  $k = 1$ 
Set  $T = 100$ 
While ( $k \leq$  the total number of iterations to be performed)
    If  $Z(\mathbf{x}_n) < Z(\mathbf{x}_c)$  then
        Set  $\mathbf{x}_c = \mathbf{x}_n$ 
        If  $Z(\mathbf{x}_c) \leq Z(\mathbf{x}_b)$ 
            Set  $\mathbf{x}_b = \mathbf{x}_c$ 
        Else
             $T = T * 0.99$ 
            Set  $r = U(0,1)$ 
            If  $r \leq e^{\frac{Z(\mathbf{x}_n) - Z(\mathbf{x}_c)}{T}}$ 
                Set  $\mathbf{x}_c = \mathbf{x}_n$ 
            Set  $\mathbf{x}_n$  to be a mutation of  $\mathbf{x}_c$ 
             $k = k + 1$ 
End while
Return  $\mathbf{x}_b$ 

```

Figure 1. Simulated annealing algorithm pseudo-code

2. 3. Leakage Signature

A leakage signature is the evidence necessary to conclude that leakage has occurred. We use the term *signature* as defined by Baker et al. (2013) to mean the collection of features assembled for the purpose of detecting a target phenomenon of interest. This is a flexible system that allows site specific signatures to be used within the objective function to infer CO₂ leakage.

The leakage signature is defined hierarchically at two levels. The bottom-level of the signature hierarchy specifies the criteria that cause an individual sensor to alarm. The criteria can

be specified as an exceedance threshold, a relative change from background, or an absolute change from background. The changes may be specified as strictly greater than or strictly less than. For example, a notional alarm signature may be defined as: “CO₂ gas saturation greater 0.05” or “a change in CO₂ from background of +0.05”. The top-level of the signature hierarchy specifies the number of alarming sensors in a configuration necessary to infer leakage of CO₂. The criteria may invoke Boolean logic to specify a combination of different sensor-types. For example, a notional inference signature may be defined as: “four alarming CO₂ gas saturation sensors and three alarming pressure sensors”. The ability to compose leakage signatures means that even a trivial signature which consists of only a single alarming sensor is possible.

The leakage signature is used to calculate the TFD for each simulation realization and sensor configuration combination. To calculate the TFD, we iterate over the simulated time steps of the realization, comparing the corresponding values returned from each sensor location in the configuration against the specified alarm signature to determine whether the sensor has alarmed. We then compare the number of alarming sensors in the configuration against the inference signature. If the inference signature is met, we infer that leakage has occurred at the current time step. If leakage is not inferred after evaluating each time step of the current realization, the configuration is assigned an excessively large penalty to the TFD. Assignment of a penalty value is a conventional technique for optimization problems and is necessary in this case because each configuration must resolve to some value in order to be compared to other configurations. While the selection of the penalty amount is arbitrary, it does not prevent the algorithm from identifying optimal monitoring designs because the ranking of the objective values is preserved. It does, however, render the $E[TFD]$ unusable as an estimate of the expected time until first detection of

leakage in cases in which leakage is not detected across all realizations. Pseudo-code describing how the inference model is implemented is presented in Figure 2.

```
For each time step  $t$  in realization  $\omega$ 
  For each sensor  $s$  in configuration  $x$ 
    If simulated value at location of  $s$  exceeds alarm signature of  $s$ 
       $s_{\text{alarm}} = \text{true}$ 
    Next  $s$ 
    If inference signature is satisfied by  $x$ 
      Return  $t$ 
  Next  $t$ 
Return 1,000,000
```

Figure 2. Inference model pseudo-code to calculate time to first detection for a specified configuration.

2. 4. *Best Time to Detection*

Separate from our simulated annealing algorithm, we developed a “Best TTD” algorithm that assumes an unlimited number of sensors and sensor types, placed in every element of the solution space, and finds the best TTD based on the minimum sensor requirements in every scenario. This method provides no information about the optimal configuration, but provides the absolute best TTD given the requirements specified in seconds. Once the user knows the absolute best TTD, they have more insight to interpret the simulated annealing results and make changes to inputs. For example: if the algorithm is unable to find the optimal solution (or near optimal solution), the user should first choose to increase the iterations and then choose to increase the budget. Plotting results of multiple iterative procedures assists in this decision making process, as the user should be able to recognize the pattern of convergence.

2. 5. *Leakage Simulations*

To simulate leakage of CO₂ from the storage reservoir into an overlying aquifer, we used simulation output from Pacific Northwest National Laboratory's (PNNL) Subsurface Transport Over Multiple Phases simulator, STOMP-CO₂-R (White et al., 2012). STOMP is a general-purpose tool for simulating multiphase subsurface flow and transport. The CO₂-R module specifies systems involving reactive transport of CO₂ and brine at reservoir conditions. An existing model of the Edwards Aquifer described in Hou et al. (2014) was used as a notional groundwater aquifer located above a CO₂ injection site. The three-dimensional, homogeneous, unconfined aquifer was assumed to be 150 m below the ground surface and 19 km × 23 km in area. A mass rate point source at the lower boundary of the aquifer introduced CO₂ into the aquifer from a characteristic storage formation. The values of the model input parameters for each realization were sampled from probability distributions in Table 1, resulting in a set of realizations representing different possible leakage scenarios. Each realization simulated 200 years (50 years of CO₂ injection into the storage formation and 150 years of monitoring) and calculated the CO₂ gas saturation and pressure at each node in the domain.

The assumption of this example is that leakage has occurred above the deep-subsurface zones into a simplified homogeneous aquifer; however, the same methodology is used by in overlying formations of any depth and of any hydrogeologic properties. The complexities of the leakage model are entirely user-dependent as this method does no flow calculations intrinsically. These realizations were selected because of the availability of data and known background values of observation parameters. The use of realizations from the Edwards Aquifer model does

not imply that CO₂ injection is planned beneath this site or that shallow aquifers should be targeted to provide E[TFD].

As a proxy for variation in the location of the leakage source term into the aquifer, synthetic realizations were generated by shifting the output data from simulated realizations along the x - and y -axes. By this method, we considered 20 realizations about four leakage locations. Each of the 20 realizations represents a potential leakage scenario with different geochemical parameters and different leakage source term locations. Each realization is assumed to be equally probable; however, it should be noted that knowledge of specific subsurface systems may lead to a weighted distribution of probabilities considered. Our objective was to demonstrate that our approach identifies near-optimal monitoring designs. When identifying monitoring designs for an operational CO₂ storage site -where the intent is to determine the optimal design rather than simply show that the approach is valid- the set of realizations should be sufficiently thorough to cover all the potential leakage outcomes.

Table 1. Sampled geologic properties used to create leakage realizations (Hou et al., 2014)

Parameter	Minimum	Maximum	Distribution
Porosity, %	13	23	normal
Hydraulic conductivity, m/day	0.09	4.6	log normal
Calcite, solid phase, volumetric %	0.1	0.99	normal
Calcite dissolution rate at 25°C, mol/m ² /s	1.5×10^{-7}	1.5×10^{-5}	log normal
Dolomite dissolution rate at 25°C, mol/m ² /s	2.9×10^{-9}	2.9×10^{-7}	log normal
Surface area, cm ² /g	0.01	1	log normal
Calcite equilibrium constant at 25°C, log	1.35	2.35	normal
Dolomite equilibrium constant at 25°C, log	2.01	3.01	normal
Distance from injection well to leaky well, m	10	4300	log uniform
Maximum leak rate, kg/s	0.01	1	log uniform

3 Application and Results

In this section we demonstrate simulated annealing using a test case which consists of 20 subsurface leakage realizations generated using the procedure discussed in the previous section. This example used two leakage signatures to demonstrate the sensor selection logic employed in our algorithm. Thresholds were set based on no-impact thresholds previously established for this system (Last et al., 2013). The first was total dissolved solids (TDS), which had a leakage threshold of 420 mg/L, and signaled in 4,161 locations across all 20 realizations. The second was pH, which had leakage threshold of <6.6 , and triggered in 3,585 locations. Together, these two sets of data formed the solution space for the iterative procedure (Figure 3). We required at least two TDS and two pH sensors to be selected in the final monitoring configurations, with a budget given for five total. Using these constraints, our algorithm selected a fifth sensor of either TDS or pH based on performance of the configurations. Note that the union of nodes meeting the alarm criteria across realizations and time steps, rather than the intersection, was considered because in cases in which the scenarios are not equally probable, nodes that do not meet the alarm criteria across all realizations may still be part of the optimal solution.

Figure 3. Solution space of 20 realizations for TDS and pH leakage signals.

Since the monitoring designs generated during each iteration of the algorithm are based on random mutations of the previous designs, identically configured instances of the algorithm will result in different solutions. To account for this randomness, the algorithm was executed 10 times, each time executing 3,000 iterations. As discussed in Section 2. 1, our simulated annealing algorithm discarded configurations that never detected leakage by assigning them a penalty.

Computationally, the time to evaluate each monitoring design in the entire solution space would be equivalent to the time it takes to evaluate each monitoring design generated by the simulated annealing algorithm. For reference, it took a standard workstation with a 4th Generation Intel processor and 16 GB memory, 4 minutes to complete this process. While the exact speed of each evaluation is related to specific computer capabilities, it is reasonable to compare the number of monitoring designs that would be tested in a full enumeration versus the number tested for this example. The full enumeration would require a number of configuration combinations on the order of 10^{13} , where this example reduced the amount of tested configurations to the order of 10^4 . This comparison demonstrates the efficiency of the simulated annealing algorithm in designing these configurations.

We evaluated the effectiveness of our simulated annealing algorithm by comparing the E[TFD] of monitoring designs identified by the algorithm to the global optimal objective value. The “Best TTD” algorithm determined the soonest a monitoring configuration with these specifications could detect a leak was 10 years, which was found in all 20 realizations assuming and unlimited number of sensors. Results from this example showed the percent of scenarios where leakage is detected increased throughout the iterative procedure (Figure 4), as configurations became more successful (Figure 5). This method identified 4,855 different configurations with average E[TFD] of 10.5 years, which ranged from 10 to 20 years (Figure 5).

These results confirmed the ability of our algorithm to effectively and efficiently identify near-optimal solutions.

Figure 4. Percent of scenarios in which a leak was successfully detected throughout the iterative procedures for each simulated annealing run.

Figure 5. Iterative progression calculating the average estimated time to first detection of each simulated annealing run. Occasional increases in $E[TFD]$ reflect random steps taken by the algorithm.

4 Discussion

The success of this approach in terms of efficiency and effectiveness at finding near-optimal solutions has prompted development of user-friendly software tool capable of employing this optimization method called “DREAM” (Designs for Risk Evaluation and Management).

DREAM is being built to be highly extensible, from the input data through the algorithm to the format of the results. It is compatible with common subsurface flow models (ex. STOMP, NUFT) and handles large numbers of realizations simultaneously by pruning the solution space, based on user input, to a manageable size prior to execution.

The objective function of the DREAM tool is very similar to that used by Sun et al. (2012, 2013). However, the approaches taken to finding a solution are very different. Sun et al. (2013) used the Gurobi solver in Matlab to produce an optimal solution for a given objective function over the full set of decision variables associated with the problem. The one piece of timing information that they gave is that this computes the exact solution in a one-scenario, three-sensor case with 3,525 decision variables within a couple of seconds. When this is scaled up to 10, 20, and 30 scenarios and three sensors, no timing information is given, but the number of decision variables scales almost exactly linearly. Since we do not know the exact solution method being used inside the Gurobi toolbox, we cannot presume to predict the scaling exactly, but we do know that the two most likely approaches (branch-and-bound or cutting-plane) can be solved in polynomial time for a fixed number of decision variables. It is implied that the scaling is poor in the statement, “We see in this example that as the number of scenario [sic] increases, the number of decision variables grows quickly. Therefore, for cases involving many equally likely scenarios a better approach is to identify bounding scenarios, if possible, and then perform optimization” (Sun et al., 2013). In contrast, the decision variables in the simulated annealing approach in DREAM are the potential locations of the sensors, which tend to have a high degree of overlap between the scenarios. Therefore, we find that there is little increase in the number of decision variables as we include more scenarios in the model. The added overhead of evaluating

each scenario for detection is negligible, allowing DREAM to load and run even a hundred scenarios at a time without much decrease in performance.

With regards to the scalability of the two algorithms with respect to an increase in the number of sensors, we cannot speak to the scalability of the algorithm presented by Sun et al. (2013) as no such information is provided. In the simulated annealing approach implemented in DREAM, the inclusion of more sensors has two effects. First, it increases the number of iterations spent during the setup of the scenario placing sensors, and therefore could require a few more iterations to converge properly. Second, it increases the overhead cost associated with checking each scenario for detection and deciding on a mutation action, as there are more sensors involved. In practice we find that neither of these two things has a significant impact on the runtime, and we can increase the number of sensors by an order of magnitude while seeing the runtime increase far less.

The tradeoff for the speed of our approach and the resilience that we observe to increases in the numbers of sensors and scenarios is that we have no guarantee that we have found the global optimal configuration at the end of the simulated annealing run. This we can address via the ability to quickly determine what the optimal value is if unlimited sensors can be placed (via the “Best Time to Detection” algorithm). Then, since the simulated annealing algorithm is computationally cheap to run, the user can vary the number of wells and sensors until they come to a solution that is within an acceptable range of the optimal solution.

It is clear from the thousands of solution configurations found in the discussed example that additional decision support is needed to determine optimal monitoring configurations beyond minimizing the time to first detection. One simple method to reduce the number of

configurations considered is to exclude nearby monitoring locations from the search following the identification of a successful configuration. To assist in further decision making, DREAM is being built to perform cost/benefit analyses of the selected systems. These analyses will print out along with configuration results and attempt to provide a simple valuation of the monitoring designs.

5 Conclusions

This paper describes an approach for rapidly identifying monitoring designs that minimize the expected time to first detection, $E[TFD]$, of CO_2 leakage from a storage formation. The efficiency and effectiveness of the approach is demonstrated using a notional test case to show that near-optimal monitoring designs identified from among myriad sub-optimal designs. We applied a simulated annealing algorithm that searches the solution space by iteratively mutating potential monitoring designs. We took into account uncertainty by evaluating the performance of potential monitoring designs across a set of simulated leakage realizations. A flexible two-tiered signature composed of TDS and pH monitors set at no-impact thresholds was used to infer that CO_2 leakage had occurred.

While the test case used in this paper was simplistic, it enabled us to demonstrate that the approach is feasible. The validity of the identified configurations was demonstrated by comparing the $E[TFD]$ of monitoring designs generated by this approach against the global optimal $E[TFD]$, identified via our “Best TTD” approach. Our average configurations identified leakage within a year of the global optimal. Additionally, the worst-case designs identified by our approach were only 10 years longer than the global optimal.

In order to provide real world utility, variation of the leakage location should be a function of a detailed geologic model that contains likely pathways for CO₂ leakage from the reservoir, through the cap-rock, and into the overlying aquifers. User-friendly software capable of employing this optimization method called “DREAM” is currently under development to provide stakeholders, regulators, and researchers an accessible platform with which to implement this tool.

Acknowledgments

PNNL is operated by Battelle for the U.S. Department of Energy (DOE) under Contract DE-AC05-76RL01830. This research is part of the National Risk Assessment Partnership, an initiative within DOE’s Office of Fossil Energy.

References

- Bacon, D.H., 2013. Reduced Order Model for the Geochemical Impacts of Carbon Dioxide, Brine, and Trace Metal Leakage into an Unconfined, Oxidizing Carbonate Aquifer, Version 2.1. Pacific Northwest National Laboratory. PNNL-22285.
http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-22285.pdf
- Baker, N.A., Barr, J.L., Bonheyo, G.T. Joslyn C.A., Krishnaswami K., Oxley M., Quadrel R.W., Sego L.H., Tardiff, M.F., Wynne A.S., 2013. Research towards a systematic signature discovery process. IEEE International Conference on Intelligence and Security Informatics (ISI), June 4-7, Seattle, Washington, pp. 301-308. Institute of Electrical and Electronics Engineers, Piscataway, NJ.

441 Bangerth, W., Klie, H., Wheeler, M.F., Stoffa, P.L., Sen, M.K., 2006. On optimization
 442 algorithms for the reservoir oil well placement problem. *Computational Geosciences*. 10:
 443 303-319.

444 Cameron, D.A., Durlofsky, L.J., 2012. Optimization of well placement, CO₂ injection rates, and
 445 brine cycling for geological carbon sequestration. *International Journal of Greenhouse*
 446 *Gas Control*. 10: 100-112.

447 Carroll, S. A., Keating, E., Mansoor, K., Dai, Z., Sun, Y., Trainor-Guitton, W., Brown, C.,
 448 Bacon, D., 2014. Key factors for determining groundwater impacts due to leakage from
 449 geologic carbon sequestration reservoirs. *International Journal of Greenhouse Gas*
 450 *Control*. 29: 153-168.

451 Dai, Z., Keating, E., Bacon, D., Viswanathan, H., Stauffer, P., Jordan, A., Pawar, R., 2014.
 452 Probabilistic evaluation of shallow groundwater resources at a hypothetical carbon
 453 sequestration site. *Scientific reports*, 4.

454 Federal Requirements Under the Underground Injection Control (UIC) Program for Carbon
 455 Dioxide (CO₂). 75 FR 77229 (December 10, 2010). [https://federalregister.gov/a/2010-](https://federalregister.gov/a/2010-29954)
 456 [29954](https://federalregister.gov/a/2010-29954)

457 Hou, Z., Bacon, D. H., Engel, D. W., Lin, G. , Fang, Y. , Ren, H. , Fang, Z. 2014. Uncertainty
 458 analyses of CO₂ plume expansion subsequent to wellbore CO₂ leakage into aquifers.
 459 *International Journal of Greenhouse Gas Control*. 27: 69-80.

460 IEA, 2010. Carbon Capture and Storage Model Regulatory Framework. International Energy
 461 Agency, <https://www.iea.org/topics/ccs/publications/>

462 Jahangiri, H. R., Zhang, D., 2012. Ensemble based co-optimization of carbon dioxide
 463 sequestration and enhanced oil recovery. International Journal of Greenhouse Gas
 464 Control. 8: 22-33.

465 Keating, E., Dai, Z., Viswanathan, H., 2013. Generation II Reduced Order Models for
 466 Groundwater Impacts due to CO₂ and Brine leakage. Los Alamos National Laboratory.
 467 LA-UR-13-22306. [http://permalink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-](http://permalink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-UR-13-22306)
 468 [UR-13-22306](http://permalink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-UR-13-22306)

469 Last, G.V., Murray, C.J., Brown, C.F., Jordan, P.D., Sharma, M., 2013. No-Impact Threshold
 470 Values for NRAP's Reduced Order Models. Pacific Northwest National Laboratory.
 471 PNNL-22077. [http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-](http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-22077.pdf)
 472 [22077.pdf](http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-22077.pdf)

473 Last G.V., Murray, C.J., Bott, Y.J., 2015. Derivation of Groundwater Threshold Values for
 474 Analysis of Impacts Predicted at Potential Carbon Sequestration Sites. International
 475 Journal of Greenhouse Gas Control, submitted for review.

476 Matott, L.S., Leung, K., Sim, J., 2011. Application of MATLAB and Python optimizers to two
 477 case studies involving groundwater flow and contaminant transport modeling. Computers
 478 & Geosciences. 37: 1894-1899.

479 Nghiem, L., Yang, C., Shrivastava, V., Kohse, B., Hassam, M., Card, C., 2009. Risk Mitigation
 480 Through the Optimization of Residual Gas and Solubility Trapping for CO₂ Storage in
 481 Saline Aquifers. Energy Procedia. 1: 3015-3022.

482 Ruprecht, C. M., 2014. The Effects of Secondary Trapping Mechanisms on Geologic Storage of
 483 Carbon Dioxide. All Dissertations. Paper 1284.
 484 http://tigerprints.clemson.edu/all_dissertations/1284

485 Sun, A. Y., Nicot, J. P., 2012. Inversion of pressure anomaly data for detecting leakage at
 486 geologic carbon sequestration sites. *Advances in Water Resources*, 44, 20-29.

487 Sun, A. Y., Nicot, J. P., Zhang, X., 2013. Optimal design of pressure-based, leakage detection
 488 monitoring networks for geologic carbon sequestration repositories. *International Journal*
 489 *of Greenhouse Gas Control*, 19, 251-261.

490 U.S. EPA, 2012. http://water.epa.gov/type/groundwater/uic/wells_sequestration.cfm (accessed
 491 May 2015).

492 U.S. EPA, 2013. Geologic Sequestration of Carbon Dioxide: Underground Injection Control
 493 (UIC) Program Class VI Well Testing and Monitoring Guidance. EPA 816-R-13-001,
 494 March 2013. <http://water.epa.gov/type/groundwater/uic/class6/upload/epa816r13001.pdf>

495 U.S. EPA, 2015. <http://water.epa.gov/type/groundwater/uic/class6/gsguidedoc.cfm> (accessed
 496 May 2015).

497 Wang, H., Echeverria-Ciaurri, D.E., Durlofsky, L.J., Cominelli, A., 2011. Optimal Well
 498 Placement under Uncertainty using a Retrospective Optimization Framework. SPE
 499 Reservoir Simulation Symposium, February 21-23, 2011. The Woodlands, Texas. SPE
 500 141950.

501 White, M.D., Watson, D.J., Bacon, D.H., White, S.K., McGrail, B. P., Zhang, Z.F., 2012.
 502 STOMP Subsurface Transport Over Multiple Phases STOMP CO₂ and CO₂e Guide.

503 PNNL-21268. http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-
504 21268.pdf

505 Zhang, Z., Agarwal, R., 2013. Numerical simulation and optimization of CO2 sequestration in
506 saline aquifers for enhanced storage capacity and secured sequestration. *Int. Journal of*
507 *Energy and Environment*. 4(3): 387-398.