

Optimal Experiment Design for Timelapse Tomography : Customizing Crosswell Micro-arrays For Monitoring Applications

Jonathan B. Ajo-Franklin
Earth Resources Laboratory
Dept. of Earth, Atmospheric, and Planetary Sciences
Massachusetts Institute of Technology
Cambridge, MA 02142

May 7, 2007

Abstract

Geophysical monitoring techniques offer the only approach capable of assessing both the spatial and temporal dynamics of subsurface fluid processes. Historically, monitoring datasets have consisted of surveys sequentially collected using acquisition geometries and sensor platforms similar to static measurements. Unfortunately, a host of logistical constraints hamper the repeatability of such surveys, particularly difficulties replicating the source/receiver geometry. Increasingly, permanent sensor arrays in boreholes and on the ocean floor are being deployed to improve the repeatability and increase the temporal sampling of monitoring surveys. Because permanent arrays require a large up-front capital investment and are difficult (or impossible) to re-configure once installed, a premium is placed on selecting a geometry capable of imaging the desired target at minimum cost.

We present a simple approach to optimizing downhole sensor arrays for monitoring experiments making use of differential seismic traveltimes. In our case, we use a design quality metric based on the accuracy of tomographic reconstructions for a suite of imaging targets. By not requiring an explicit SVD of the forward operator, evaluation of this objective function scales to problems with a large number of unknowns. We also restrict the design problem by recasting the array geometry into a low dimensional form more suitable for optimization. A side effect of using these restrictive parameterizations for experiment geometry is a well-behaved objective function more amenable to local search techniques. To demonstrate the efficacy of our algorithm, we consider a series of possible designs optimization problems for a next-generation permanent tomographic monitoring system. We test two search algorithms on the design problem, the Nelder-Mead downhill simplex method and the Multilevel Coordinate Search algorithm.

The complete design algorithm is tested for three crosswell acquisition scenarios relevant to continuous seismic monitoring, a 2 parameter array length optimization, a 4 parameter length/offset optimization, and a comparison of optimal multi-source designs. In the last case, we also examine trade-offs between source sparsity and the quality of tomographic reconstructions. Preliminary results suggest that high-quality differential images can be generated using only a small number of optimally positioned sources, an observation with immediate relevance to several field projects still in the development phase.

1 Introduction

Understanding the combined spatial and temporal characteristics of dynamic subsurface processes is a challenge common to many contemporary problems in the geosciences. Examples include the growth of bacterial communities during enhanced remediation, the dissolution of methane hydrates during stimulated production, and the movement of CO₂ during sequestration activities. These are but three of numerous applications which require knowledge of spatio-temporal variability to improve our fundamental understanding of relevant physical processes.

Geophysical monitoring techniques offer the only approach capable of assessing the dynamics of these systems beyond the limited possibilities afforded by direct *in situ* observations. Historically, monitoring datasets have consisted of surveys sequentially collected using acquisition geometries and sensor platforms similar to static measurements e.g. repeated reflection seismic surveys. Unfortunately, a host of logistical constraints hamper the repeatability of such surveys, particularly difficulties in exactly replicating the source/receiver geometry and sufficiently matching the seismic processing flow [Ross and Altan, 1997].

Increasingly, permanent sensor arrays in boreholes [Blanco *et al.*, 2006; Daley *et al.*, 2007] and on the ocean floor [Smit *et al.*, 2006; Thompson *et al.*, 2006] are being deployed to improve the repeatability of monitoring surveys by eliminating the need to re-position receivers. By allowing continuous monitoring, permanently deployed sensor arrays also have the potential to improve the *temporal* resolution of geophysical measurements by an order of magnitude (or more); this increase in temporal sampling will in turn allow scientists to probe transient subsurface processes with shorter time scales, invisible to the currently available 4D seismic monitoring methods that typically sample on the scale of months or even years.

Despite clear data quality advantages, permanent arrays require both substantial up-front capital investments and long-term maintenance and thus necessitate careful analysis to justify their installation. Since permanent arrays are often difficult (or impossible) to re-configure once in-place, their design assumes an even more important role than in the case of more traditional surveys where mistakes can be potentially be corrected during later deployments. For permanent arrays to be a cost-effective monitoring option, emphasis must be placed selecting the sparsest (and hence least expensive) possible source/receiver geometry capable of producing images of a given quality. Smit *et al.* [2006] cogently presents this philosophy within the context of ocean bottom cable (OBC) installations. While the cost of a given geometry is relatively simple to evaluate, the “quality” of the resulting experiment is more difficult to quantify since it depends not only on our ability to resolve individual geologic features but also on the monitoring questions we are attempting to answer. Binary detection of change within a broad region should be an easier task than reconstructing an accurate subsurface image. Designing a good survey is most difficult in cases for which we lack experience e.g. irregular geometries, unusual spatial constraints, or a strongly heterogeneous background, all of which can render typical “rules of thumb” invalid.

Given the central importance of survey geometry in the context of monitoring, we will consider algorithms for the selection of an optimal (or at least improved) configuration from within the space of possible designs. This task falls within the class of problems often referred to as *optimal experiment design* (OED), a field which has proponents in both geophysics [Maurer and Boerner, 1998; Curtis, 1999a,b; Stummer *et al.*, 2002; van den Berg *et al.*, 2003; Curtis *et al.*, 2004; Routh *et al.*, 2005] and the broader scientific community [Box and Lucas, 1959; John and Draper, 1975; Walter and Pronzato, 1987; Muzic *et al.*, 1996]. When searching for an optimal experiment several considerations rapidly become apparent, mainly what “optimal” means in the context of geophysical measurements. Other crucial issues include the range of experimental designs to be considered, the technique used to solve the optimization problem, and the fashion in which design constraints are implemented.

In our approach, we use the l_2 difference between a suite of reference images and the corresponding tomographic reconstructions as a quality measure for a given geometry. This suite of reference images can be selected to mimic either traditional resolution metrics (e.g. checkerboards in our case) or directly linked to the spatial characteristics of the process being monitored. This simple quality measure avoids computation of the SVD of the forward operator and as a result can be evaluated for large problems. We also adopt a reduced parameterization for describing experimental geometry; instead of allowing sources and receivers to occupy arbitrary locations, we search for an optimal set of “meta-parameters” such as array width,

orientation, or center location. This approach, while lacking the flexibility of the most general descriptions of experiment geometry, results in a more manageable search problem. For simple cases, the global minimum of the resulting objective function can be determined through use of a variety of search techniques. Reduced parameterizations have the additional benefit of being physically realizable. Long cables with arbitrary sensor spacings are generally difficult to manufacture; by limiting the choice of geometries we implicitly reduce the cost of custom fabrication.

In this preliminary investigation, we use OED to improve the design of active seismic monitoring systems within a crosswell geometry. We first consider a 2 parameter array design problem and explore the characteristics of the quality metric. We compare the convergence and performance of two optimization methods, the Nelder-Mead downhill simplex and multilevel coordinate search strategies for selecting array geometries. For this simple example, both techniques converge to the globally optimal design. We also consider a more complicated 4 parameter design problem which optimizes a combination of array length and vertical offset. We conclude by examining trade-offs between source sparsity and the quality of tomographic reconstructions. Preliminary results suggest that high-quality differential images can be generated using only a small number of optimally positioned sources, an observation with immediate relevance to several field projects still in the development phase.

2 Principles & Previous Research

As mentioned previously, optimal designed algorithms can be roughly decomposed into three components, the metric used to discern an “optimal” experiment, the space of experiments considered as possible solutions, and the search technique used to find the best design. Another important component which we will only briefly cover is the implementation of secondary constraints, either direct limitations on the geometry or auxiliary parameters such as survey cost, acquisition time etc. In this section we will outline a framework for expressing OED search problems and their associated components.

2.1 An OED Notation

Implicit in the entire design process is the inversion step itself i.e. the methodology by which measurements are transformed into estimates of earth properties. Most OED algorithms are formally posed within the framework of linear inverse problems, sometimes with ad-hoc extensions to the non-linear case. We will use the traditional notation for such problems [Menke, 1984],

$$\mathbf{G}\mathbf{m} = \mathbf{d}, \tag{1}$$

where \mathbf{G} , referred to as the kernel, maps a model (\mathbf{m}) to a dataset (\mathbf{d}). \mathbf{G} encapsulates both the operative physics and the geometry of the experiment. Linear inversion attempts to undo the action of \mathbf{G} thereby reconstructing an earth model given a dataset. Since \mathbf{G}^{-1} often does not exist; we will use \mathbf{G}^{-g} to denote the generalized inverse of the kernel. Many OED techniques are posed in terms of the singular value decomposition (SVD) of \mathbf{G} ,

$$\mathbf{G} = \mathbf{U}\mathbf{S}\mathbf{V}^T \tag{2}$$

where \mathbf{U} and \mathbf{V} are the data and model space eigenvector matrices respectively. \mathbf{S} is a diagonal matrix whose elements are the singular values of \mathbf{G} .

As mentioned previously, \mathbf{G} contains an implicit dependence on the experiment geometry; since we are searching across these possible geometries this relationship must be written explicitly. Let the geometry of our permanent source/receiver array be encapsulated by a vector ξ . In the most general case where we are free to locate each of d sources and e receivers at any point in space we express the experiment geometry as,

$$\xi = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_d : \mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_e] \tag{3}$$

where \mathbf{s} and \mathbf{r} are source and receiver locations in the appropriate number of spatial dimensions. In 3D, the number of scalar design parameters in the experiment becomes $(d + e) \times 3$. The optimality metric, Ω , maps ξ to a scalar measure of experiment quality. The optimal design geometry, ξ_{optim} can then be written as,

$$\xi_{optim} = \min(\Omega(\xi, \{\alpha_1, \dots, \alpha_t\})), \quad (4)$$

where the set of α variables are auxiliary parameters which are required to evaluate Ω but are not part of the search problem. An example of α would be a background velocity model required to calculate \mathbf{G} . In many cases, the metric Ω can be decomposed into two parts, a raw measure of experiment quality (Ω_u) and a set of f soft secondary constraints (Ω_{c_i}), possibly including information pertaining to spatial limitations of the survey, field costs, or acquisition time. Momentarily dropping the arguments to Ω we can write this decomposition as,

$$\Omega = \Omega_u + \sum_{i=1}^f \beta_i \Omega_{c_i}, \quad (5)$$

where β_i is a weighting factor which scales the strength of the i th soft constraint terms. Equation 5 is in no way encompassing since the constraint term can also be implemented in a “hard” form as part of the search algorithm itself. Previous authors have also decomposed Ω_u into a sum of “global” vs. “local” image attribute terms [Curtis, 1999b] to allow use of focusing in the optimization process; our approach does not require such a distinction. This notation for writing Ω is not well-suited to expressing greedy design algorithms for which the dimensions of ξ vary dynamically as part of the search process.

2.2 Previous Research

Existing OED algorithms suffer from a combination of problems, mainly inappropriate choices for Ω , the experiment quality metric, and/or poor parameterizations for the geometry search space, typically involving the explicit use of ξ . A bad choice of Ω or ξ can in turn result in an unnecessarily large search space with non-ideal properties, particularly the existence of extraneous local minimum.

2.2.1 Experiment Quality Metrics

The quality metrics proposed by previous authors roughly fall into two camps, those who define quality in terms of the eigen-properties of a given design [Barth and Wunsch, 1990; Curtis and Snieder, 1997; Curtis, 1999a,b] and more applied practitioners who optimize heuristic attributes such as fold or ray coverage [Liner et al., 1999; Vermeer, 2003; Galbraith, 2004]. Design algorithms based on SVD analysis of the linear inverse problem were pioneered in a geophysical context by Maurer [Maurer and Boerner, 1998] and Curtis [Curtis and Snieder, 1997; Curtis, 1999a,b] who adapted previous work by Barth [Barth and Wunsch, 1990] and earlier authors [Box and Lucas, 1959; John and Draper, 1975]. All of these techniques seek to maximize the amount of information contained within a given geophysical survey, typically by reducing the linear problem’s null space. Curtis [1999b] provides a catalog of five experiment quality metrics based only on the eigenspectrum of \mathbf{G} . Most of these metrics involve sums or products of eigenvalues, possibly truncated so as to consider only values above or below a given threshold. The most celebrated of these metrics is the so-called D criterion [John and Draper, 1975] which corresponds to the determinant of $\mathbf{G}^T \mathbf{G}$ or,

$$\Omega_u = \prod_{i=1}^N \lambda_i \quad (6)$$

where λ_i corresponds to the i th eigenvalue. Another possible metric is the area beneath the eigenspectrum normalized to the largest eigenvalue advocated by Curtis and Snieder [1997],

$$\Omega_u = \sum_{i=1}^N \frac{\lambda_i}{\lambda_1} \quad (7)$$

Equation 7 measures the “flatness” of the operator’s eigenspectrum. Both methods tend to boost small eigenvalues but have different sensitivities to entries near zero or the specified truncation threshold. For an exhaustive discussion of metrics of this type we refer the reader to *Curtis* [1999a].

Despite their apparent elegance, global SVD based techniques suffer from several clear weaknesses. Methods which rely solely on examination of the eigenspectrum provide no avenue for incorporating zones of interest since all spatial information is encoded in the right eigenvectors (\mathbf{V}); the resulting designs often optimize the experiment to recover properties within unimportant portions of model space. This is particularly problematic in the case of monitoring experiments where we have prior knowledge of regions where change might occur. *Curtis* [1999b] attempts to remedy this problem by incorporating a second “focused” component of Ω by examining projections of model regions onto the right singular vectors. In practice, the global eigenspectrum methods sometimes suffer from a curious problem, mainly that many experiment geometries with “optimal” (or at least improved) design metrics seem to generate poor images.

From a computational perspective, metrics which require the complete SVD of the forward operator are impractical for anything but small problems considering the $O(n^3)$ time and memory complexity of the factorization. In many large scale seismic inversion problems, the sparse forward operator is too large to explicitly store, much less factorize using dense algorithms. For some of the global metrics (such as equation 7) there are short cuts which reduce this effort to slightly more manageable proportions but in general they are not easy to adapt to problems with large model dimensions.

An attractive set of related quality metrics are those which make direct use of the model resolution matrix, \mathbf{R} , which for linear problems maps the true model (\mathbf{m}_{true}) to the estimated model (\mathbf{m}_{est}) i.e. $\mathbf{m}_{est} = \mathbf{R} \mathbf{m}_{true}$. The diagonal elements of \mathbf{R} are scalar measures of how well-resolved each parameter in the inverse problem is. In the case where $\mathbf{R} = \mathbf{I}$, the true model is perfectly reconstructed. \mathbf{R} can be expressed in terms of the right eigenvectors $\mathbf{V}^T \mathbf{V}$ or in terms of the generalized inverse of the kernel $\mathbf{G}^{-g} \mathbf{G}$. Since neither \mathbf{V} nor \mathbf{G}^{-g} can be reasonably computed for large problems, alternative approaches are needed. One possibility is to exploit iterative Lanczos methods to sequentially recover elements from $diag(\mathbf{R})$ as is done by *Cox* [2004] for the related optimal parameterization problem. *Routh et al.* [2005] suggests a computationally feasible approach which sequentially reconstructs the k th row of \mathbf{R} by solving for point-spread functions (PSF). *Routh et al.* [2005] then uses the compactness of the PSF to estimate the suitability of experiment designs yielding a quality metric of the form,

$$\Omega_u(\xi) = \|\mathbf{Q}_k(\mathbf{p}_k(\xi) - \Delta_k)\|_2^2, \quad (8)$$

where \mathbf{p}_k , the point-spread response, is the solution of a regularized linear inverse problem of the form

$$\begin{bmatrix} \mathbf{G}(\xi) \\ \lambda \mathbf{W}_m \end{bmatrix} \mathbf{p}_k = \begin{bmatrix} \mathbf{G}(\xi) \Delta_k \\ 0 \end{bmatrix} \quad (9)$$

In equation 9, \mathbf{W}_m is a model-space regularization operator (\mathbf{I} in the case of damping) and Δ_k is an impulse function in the model domain; the product $\mathbf{G} \Delta_k$ is the synthetic data which would be generated by such an impulse. Equation 8 also includes a weighting operator \mathbf{Q}_k which penalizes elements of \mathbf{p}_k away from the k th cell according to a specified distance function (see *Routh et al.* [2005] for an explicit formulation). In words, the PSF metric measures how closely (in an l_2 sense) a given geometry can reconstruct a delta feature at some location in the model. This provides a tempting alternative to SVD techniques since it requires only an LSQ inversion routine and can take full advantage of sparse solvers. Because regularization can be easily incorporated into the PSF calculation, the effect of prior smoothness constraints on the optimal geometry can also be evaluated. A final more conceptual benefit is that metric 8 truly operates on the inverted model; the benefits of improving Ω are easily visible in the resulting images.

The limitation of the PSF metric as written in equation 8 is that the resolution is only evaluated at single location k ; to evaluate such a metric over a larger zone of interest would require a full inversion for each cell within this region. As we will discuss in the following sections, we adopt a metric similar to 8 except that we replace Δ_k with a single or small suite of models with broader (although less precise) features.

A second class of techniques uses easily computed metrics which provide a more heuristic guide to experiment quality e.g. CMP fold, azimuthal coverage, ray density, or more general measures of operator

sensitivity. In contrast to SVD related methods, these heuristic metrics can be computed for very large problems e.g. 3D seismic surveys for which explicit generation of the modeling operator is not feasible. One example of this type of approach is the work of *Liner et al.* [1999] who use a mix of soft and hard constraints to search for survey geometries which satisfy acquisition criterion including CMP fold and maximum offset. The resulting quality metric is very inexpensive to evaluate; millions of survey geometries can be tested for suitability in a relatively short period of time. Unfortunately, such metrics have several well-known problems, the most crucial is that they can be fulfilled even in cases where the resulting surveys have severe deficiencies from an inversion perspective. *Curtis and Snieder* [1997] documents this type of problem in the context of optimal mesh search, a process intimately related to OED; they show that ray hit-count is a poor metric for mesh adaptation in traveltime tomography.

2.2.2 Parameterizing Survey Geometry

Almost as important as the selection of a quality metric is the set of geometries considered as valid surveys since this determines the dimensions of the required search space. As mentioned previously, the most flexible possible representation is the vector ξ which allows arbitrary positions for all d sources and e receivers which results in a problem of dimensional $(d + e) \times 3$. In some cases such as when the sensors are constrained to a surface [*Barth and Wunsch*, 1990] or a line [*Curtis*, 1999a], each source and receiver has less than 3 degrees of freedom thereby slightly reducing the search space.

However, even for mid-sized seismic surveys $(d + e)$ inevitably becomes quite large, particularly in the context of global search techniques. The resulting optimization problem exhibits several unattractive attributes, primarily the existence of numerous local minima in the objective function. Additionally, the optimal geometries are often difficult to realize in practice; most acquisition techniques rely on some degree of regularity and the custom fabrication of take-out cables or streamers with arbitrary sensor spacings can be prohibitively expensive.

3 Combining Parsimonious Geometry Representations & Image-Based Quality Metrics

Our approach to the survey design problems combines a parsimonious parameterization of survey geometry with an image-based survey quality metric similar to the PSF technique advocated by *Routh et al.* [2005]. The resulting optimization problem is solved using one of two direct search method, either the local Nelder-Mead downhill simplex algorithm [*Nelder and Mead*, 1965] or the global Multilevel Coordinate Search algorithm developed by *Huyer and Neumaier* [1999]. This blend of a reduced search space, an intuitive model domain quality metric, and robust search algorithms results in a method which both scales to mid-sized problems and significantly improves the spatial resolution of the resulting experiments.

As mentioned previously, an alternative to searching across all possible geometries is to introduce a limited subset of feasible configurations described by a second set of parameters. In this case, we have an operator Φ which maps a short vector \mathbf{q} to a full geometry i.e. $\Phi(\mathbf{q}) = \xi$. The elements of \mathbf{q} are typically a small number of geometry descriptors e.g. array length, array center location, line spacing etc. While the mapping function Φ is designed to produce reasonable geometries, some constraints must still be applied to the elements of \mathbf{q} necessitating a penalty function of the form, $\Omega_c(\mathbf{q})$

The resulting problem minimizes the combined Ω with respect to \mathbf{q} ,

$$\xi_{optim} = \min(\Omega_u(\Phi(\mathbf{q})) + \sum_{i=1}^f \beta_i \Omega_{c_i}(\mathbf{q}_i)) \quad (10)$$

While the global minimum of Ω with respect to ξ is in not guaranteed to be located in the lower dimensional space covered by Φ , the geometry produced by optimizing $\Phi(\mathbf{q})$ will typically be better than an arbitrary configuration and will hopefully be easy to physically construct if the parameterization \mathbf{q} is chosen correctly. Since \mathbf{q} is of a small dimension, searching for a quasi-optimal Ω is an achievable task. As will

be shown in later sections, the resulting objective function can be unimodal for design problems with a low number of dimensions making it amenable to local search techniques. This approach to designing surveys by optimizing a limited set of meta-parameters is intuitive to those familiar with real seismic acquisition; *Liner et al.* [1999], *Morrice et al.* [2001], *Vermeer* [2003], and others have posed their survey design algorithms in terms of similar parameters such as line offset and shot density instead of arbitrary locations for each geophone or source.

Up to this point we have not mentioned our formulation for Ω_u . In our case, Ω_u will be a measure of a given geometry’s ability to reconstruct a known test model, \mathbf{m}_{true} , i.e.

$$\Omega_u(\Phi(\mathbf{q})) = \|\mathbf{m}_{true} - \mathbf{m}_{est}\|_2^2 \quad (11)$$

The estimated model, \mathbf{m}_{est} , is calculated by solving a regularized least-squares problem of the form,

$$\begin{bmatrix} \mathbf{G}(\mathbf{m}_{base}, \Phi(\mathbf{q})) \\ \lambda_x \mathbf{D}_x \\ \lambda_z \mathbf{D}_z \end{bmatrix} \mathbf{m}_{est} = \begin{bmatrix} \mathbf{d} \\ 0 \\ 0 \end{bmatrix}. \quad (12)$$

\mathbf{G} , as written in equation 12, depends on a prior model estimate, \mathbf{m}_{base} around which we linearize the problem. \mathbf{m}_{base} and \mathbf{m}_{true} are not the same; \mathbf{m}_{base} is a reference model used for calculating \mathbf{G} while \mathbf{m}_{true} is the true model used for generating linear synthetic data and evaluating the quality metric. \mathbf{D}_x and \mathbf{D}_z are 1st order derivative operators in the x and z directions while λ_x and λ_z are the associated regularization parameters. Since the a background model is known, the \mathbf{d} in equation 12 is the evaluation of a linear forward problem on the current geometry with the addition of noise i.e. $\mathbf{d} = \mathbf{G}(\mathbf{m}_{base}, \Phi(\mathbf{q}))\mathbf{m}_{true} + \mathbf{d}_{noise}$. The regularization matrices included in 12 are not by any means restricted to derivatives; the practitioner should feel free to substitute their favorite operator.

An important component of the inversion quality metric shown in equation 11 is the choice of test model, \mathbf{m}_{true} . As mentioned previously, the PSF approach uses a model with a delta spike at a single location. While point spread functions are an excellent local measure of resolution, it requires a large number of independent evaluations to assess resolving power over an extended model region. An alternative is to use a single or small group of models which provide information on resolution for a larger region which includes the imaging target. One such commonly used metric in the global tomography community is the checkerboard test where a resolution is evaluated by looking at the smearing of a regular array of features. Checkerboard tests cannot recover the same resolution information as a series of PSF evaluations since the smeared components of different features may overlap. In our quality metric, \mathbf{m}_{true} is a the product of a checkerboard and a localized window function which is only positive in the target region as is shown in figure 1.

For monitoring applications, this type of localization is desirable since the perturbed region is often confined to a single geological unit. This is particularly true in cases where the process being monitored is a fluid injection (e.g. CO₂, bio-stimulant fluid, steam) and geophysical measurements are being used to delineate the shape and extent of “fronts” in the flooding process. While checkerboard tests are by no means the ideal test of survey resolution [*Leveque et al.*, 1993], they are relatively inexpensive to perform, a key attribute when used in the inner portion of an objective function evaluated as many as 10⁶ times. In the optimal case, image metrics like equation 11 would be evaluated for a suite of models related to the process being monitored; if the imaging target is a zone of methane hydrate dissolution, the results of initial thermal transfer and flow modeling could be used to generate a suite of \mathbf{m}_{true} realizations for use in experiment design.

From a computational perspective, the quality metric shown in equation 11 requires two main operations, the generation of the forward operator given the background model and the solution of a regularized linear inverse problem to generate the \mathbf{m}_{est} . Since both of these components are available as part of standard imaging packages, this process often requires minimal code modification on the part of the user. Evaluation of equation 11 can also be performed without the explicit construction of any of the operators making it useful in cases where only the action of \mathbf{G} and \mathbf{G}^T on vectors can be performed in core.

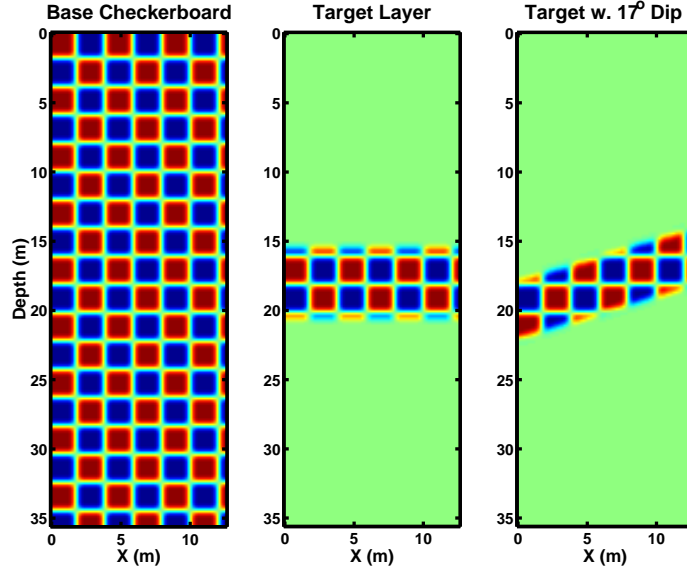


Figure 1: A test checkerboard and with various target screens : The left panel shows a global checkerboard with square 2 m cells while the middle and right panels depict targeted windows of this test model confined to either a horizontal or dipping layers.

3.1 Design Optimization : Two Algorithms

To this point, we have described a simple system for exploring experimental designs which combines an l_2 image reconstruction metric and a parsimonious description of experiment geometry. The resulting objective functions are well-behaved for several common seismic array design parameters although this observation almost certainly does not generalize to arbitrary parameterizations (or geophysical problems). Since the selection of an algorithm for minimizing the RHS of equation 10 strongly depends on the character of the objective function, we will restrict ourselves to the following observations which can be inferred from our later examples,

1. **Low Dimension**

Using a parsimonious geometry description guarantees a low-dimensional objective function by design.

2. **Expensive to Evaluate**

While our l_2 image metric can scale efficiently to large problems, it is still expensive to evaluate, requiring at the very least a forward solve and a full LSQ inversion, although not an $O(n^3)$ operation like the SVD.

3. **No Analytical Form For Gradient**

In our case, no analytic form exists for the gradient of the objective function with respect to the design parameters.

4. **Multimodal?**

For all but the simplest choice of parametrizations the design objective function seems like it should have multiple local minima; however, for the problems we have been able to explore visually (2 and 3 variables) the objective function surface is often dominated by a single broad minimum.

5. **Relatively Smooth**

Our tomographic monitoring problem, particularly when appropriately regularized, does not have exhibit sudden jumps in image quality due to small perturbations in experiment geometry.

6. Broad Plains and the Occasional Valley

While some possible designs are, in fact, very bad (!), large regions of parameter space often exhibit similar l_2 image quality metrics; the objective function surface for these regions is correspondingly flat. In some cases, valleys exist, indicating trade-offs between multiple design parameters.

Previous optimal experiment design studies have made extensive use of stochastic global search methods, with genetic algorithms (GA) [Goldberg, 1989] and simulated annealing (SA) [Ingber, 1989, 1996] being the most popular techniques. In one of the earlier experiment design investigations Barth and Wunsch [1990] used SA to search for optimal oceanographic monitoring networks while Curtis [1999a] used GA for the design of combined crosswell/VSP arrays. While both of these methods have been successfully applied to a large range of optimization problems, they share a disadvantage common to many global search approaches; mainly that they converge slowly, sometimes too slowly to be of practical interest, and as a result require a large number of objective function evaluations. In fact, GA-based approaches lack even exponential upper bounds for the number of evaluations required. Both methods also rely on extensive tuning on the part of the user for effective performance on specific problems.

We examine two alternative derivative-free approaches to minimizing the design metric, the downhill simplex method of Nelder and Mead (NM) [Nelder and Mead, 1965] and the Multilevel Coordinate Search (MCS) algorithm [Huyer and Neumaier, 1999]. Both methods are deterministic in contrast to the stochastic optimization algorithms often used in the design literature. The NM method is intrinsically local while MCS provides attempts to balance local and global search goals with guaranteed (but very slow) convergence even in the case of multi-modal objective functions.

The NM downhill simplex method is a popular [Press et al., 1992] local search algorithm which in one previous case [Muzic et al., 1996] has been successfully applied to OED. The NM approach does not compute numerical derivatives using finite-differences; instead, the objective function is evaluated at a set $n + 1$ points (referred to as a *simplex*) where n is the dimension of the optimization problem. Depending on the value of the objective function at these points, the simplex is modified through a fixed set of heuristic operations including “flipping” the location of individual simplex points and expanding or contracting the entire set of points. We refer the reader to Press for both a complete description of the algorithm and a straightforward implementation. In practice, the algorithm reliably creeps to a local minimum of the objective function. Although the convergence of NM is relatively slow in comparison to more sophisticated quasi-Newton methods which explicitly include derivative information, NM is reliable even in cases where the objective function is noisy. As with any local method, NM converges to the local minimum closest to the starting location making it susceptible to multi-modal objective functions.

In contrast to NM, the MCS algorithm, descended from the DIRECT method proposed by Jones et al. [1993], is a deterministic global search algorithm which attempts to balance large scale exploration of the parameter space with local refinement in regions with promising objective function values. MCS is guaranteed to converge to a global minimum assuming the objective function is continuous in the vicinity of the minimum; unfortunately the convergence bound is weak in the sense that an exponential number of steps might be required. Instead of starting with a parameter guess as does NM, the MCS algorithm is initialized with a hyper-box bounding all of the search variables thus allowing implicit incorporation of hard bounds on design parameters. This box is then subjected to different sub-division and sampling operations; the method of division cycles between global and local refinement steps. The original comparison of Huyer and Neumaier [1999] suggests that MCS outperforms most deterministic global search algorithms for the suite of test functions proposed by Dixon and Szego [1978].

4 Application to the Seismic Monitoring Problem

Our discussion of OED so far has been general with no assumptions regarding the operation which \mathbf{G} performs, the model parametrization represented by \mathbf{m} , or the type of data stored as \mathbf{d} . We will now apply

our formulation to the concrete example of differential seismic traveltime tomography, a technique being used as part of several integrated seismic monitoring projects. All examples which will be shown use a crosswell geometry to match field deployments in which we are currently participating in.

In the case of differential inversion we consider two datasets, \mathbf{d}_1 and \mathbf{d}_2 , acquired at different times, but with the same geometry. Instead of inverting the two datasets independently, differential tomography inverts $\Delta\mathbf{d}$, the difference between measurements made at time 1 and time 2, for $\Delta\mathbf{m}$, the temporal perturbation in model parameters. In our case, $\Delta\mathbf{d}$ is the relative traveltime difference between two different surveys for the same source/receiver pair while $\Delta\mathbf{m}$ is a perturbation in slowness. \mathbf{G} is the ray-theoretic forward modeling operator. The model is defined on a rectilinear mesh of constant slowness cells.

We treat the differential tomography problem as a linear perturbation; we assume that a background velocity model, \mathbf{m}_{base} , is available and calculate curved ray-paths in that initial model. Recovery of $\Delta\mathbf{m}$ is then considered as a purely linear inverse problem. This simplification seems reasonable in the case of seismic monitoring since the property changes induced by human perturbations are often considerably smaller than background variations. Unlike the static imaging problem where the earth model is initially unknown, most monitoring surveys have access to background property estimates from reconnaissance/exploration surveys or well logs. Our OED algorithm is designed to find surveys which optimally recover $\Delta\mathbf{m}_{true}$ given \mathbf{m}_{base} and constraints on the locations of elements in ξ . In this approach, the experiment is not designed to optimally image the background model but only to resolve a localized perturbation. As discussed in the previous section, $\Delta\mathbf{m}_{true}$ will in our case be a depth or layer limited checkerboard model.

Evaluation of Ω for a given \mathbf{m}_{base} , \mathbf{q} , and mapping function can be written in pseudo-code as,

1. Map the reduced parameter set, \mathbf{q} , to the full geometry ξ using Φ
2. Generate \mathbf{G} for the geometry ξ in \mathbf{m}_{base} , a non-linear forward modeling step
3. Calculate a perturbed synthetic dataset using the \mathbf{G} from the background model, $\Delta\mathbf{d} = \mathbf{G} \Delta\mathbf{m}_{true}$
4. Invert $\Delta\mathbf{d}$ for $\Delta\mathbf{m}_{est}$ using a regularized linear solver.
5. Evaluate the image quality objective function, $\Omega_u = \|\Delta m_{est} - \Delta m_{true}\|_2^2$
6. Evaluate the geometry penalty function Ω_c on ξ
7. Compute the combined experiment quality metric, $\Omega = \Omega_u + \Omega_c$

We calculate \mathbf{G} , the forward modeling operator, in a 2.5 dimensional geometry designed to accommodate heterogeneous 2D structure and well trajectories with out-of-plane deviations. The solution to the eikonal equation is calculated using a finite-difference eikonal solver [Vidale, 1988; Sethian and Popovici, 1999]; we took advantage of components in the FAST package developed by Zelt [Hole and Zelt, 1995]. Frechet derivatives are extracted using reverse ray-tracing [Aldridge and Oldenburg, 1993] in the shot-domain. The regularized linear inversion in step 4 is performed using the LSQR algorithm developed by Paige and Saunders [1982] and solves equation 12. The geometry penalty function mentioned in step 6 is simply a barrier term which evaluates to zero within the interior of the modeling domain and smoothly increases to a large value as source and receiver approach the edges. Ω is optimized over the reduced vector \mathbf{q} using either the NM or MCS algorithm as discussed previously.

4.1 Designing an Optimal Asymmetric Crosswell Array

A complete test of our optimal design algorithm was performed in support of a possible extension to the Frio crosswell seismic monitoring installation. The Frio demonstration project [Hovorka et al., 2006] is an on-going multi-institution effort to improve understanding of the *in situ* dynamics of super-critical CO₂ injection within a saline aquifer located in East Texas. Motivated by previous seismic monitoring experiments targeting CO₂ floods [Lazaratos and Marion, 1997], Lawrence Berkeley National Laboratory (LBNL) carried

out imaging experiments during both the first and second phases of the Frio project. The first experiment consisted of a pair of traditional time-lapse crosswell surveys where both sources and receivers were deployed via wireline [Daley *et al.*, 2005]. While the wireline surveys successfully imaged the expanding CO₂ plume as shown by Hovorka *et al.* [2006] and Ajo-Franklin *et al.* [2007], the existence of only a single repeat survey made evaluation of flow dynamics difficult.

The second LBNL imaging experiment evaluated the combination of a single permanently installed down-hole piezoelectric source and a fixed array of 24 hydrophones with both devices deployed on production tubing [Daley *et al.*, 2007]. This experiment clearly demonstrated the high level of repeatability possible with a permanent monitoring system. More importantly, the acquisition of datasets over time intervals as short as 15 minutes allowed real-time monitoring of the seismic signal during the injection process. However, the use of only a single source prevented the Frio II seismic experiment from effectively imaging CO₂ movement although qualitative information on plume geometry was obtained.

Future monitoring experiments will move beyond the single-source geometry used in the Frio II equipment trial and combine continuous monitoring with a multi-source geometry more suitable for tomographic reconstruction. Considering the high cost of each additional source transducer ($\approx \$ 50 \times 10^3$) in comparison to hydrophone pods, an asymmetric source/receiver distribution seems almost inevitable. This combination of high capital cost and minimal previous experience with similar acquisition systems makes this particular design problem a perfect candidate for OED techniques.

4.1.1 A 2 Parameter Design Problem

For our first example, we will consider a simple crosswell array optimization problem where the geometry is fully described by 2 parameters, the vertical extents of the source and receiver arrays. In this case \mathbf{q} , the reduced parameterization is simply $[P_1, P_2]$. Figure 2 shows the intended geometry. Both the source and receiver arrays are centered around a single 5 m target layer with an alternating pattern of 1.5 m wide perturbations. All figures are plotted in terms of array 1/2 widths i.e. the distance from the array center to the top of the array. In order to provide insight into the design challenges at the Frio site, we selected an asymmetric geometry distribution with 8 sources and 80 receivers. Each inversion calculated as part of evaluation of the quality metric was performed on a 75 x 210 sample mesh yielding an underdetermined 640 x 15750 (data x model) problem.

Considering the low dimensionality of the search problem, we first directly evaluated the experiment quality metric on a regular 2D grid to gain some insight into the optimization problem and the corresponding survey geometries. The existence of a calculated reference surface also allows us to examine the convergence properties of both our global and local search schemes. Figure 3 shows the image quality metric for array 1/2 width values between 2 and 16 meters. As can be seen, the objective function surface is relatively well behaved with a localized trough-like minimum corresponding to source array 1/2 widths near 6 meters. The global minimum occurs on the domain boundary at close to the maximum receiver array width.

Figure 4 attempts to capture the linkage between the experiment quality surface and the actual inversion results. While the top panel shows the quality surface, the bottom row shows the true reference checkerboard model [A] and tomographic images corresponding to 3 geometries sampled by the grid search. Panel [B] shows a narrow aperture geometry where the sources and receiver arrays are localized near the zone of interest. Panel [C] shows the geometry and tomogram associated with the global minimum; this “optimal” solution couples a source array of moderate length to a wide receiver array. Panel [D] depicts the result corresponding to a wide angular aperture in both the source and receiver wells. As can be clearly seen, the optimal geometry successfully resolves checkerboard cells near both wells with only a small amount of smearing visible above and below the target zone. In contrast, the narrow aperture geometry tends to smear features along the primary survey diagonals imparting an “X” shaped overprint on the tomogram. The maximum aperture geometry cannot resolve features near the source well due to sparse vertical spacing. This result is significant since it suggests that asymmetry in the number of sources and receivers can be partially compensated for by appropriately adjusting array lengths in cases where the target zone is spatially limited.

Figure 5 compares the convergence of the NM (left) and MCS (right) optimization algorithms when

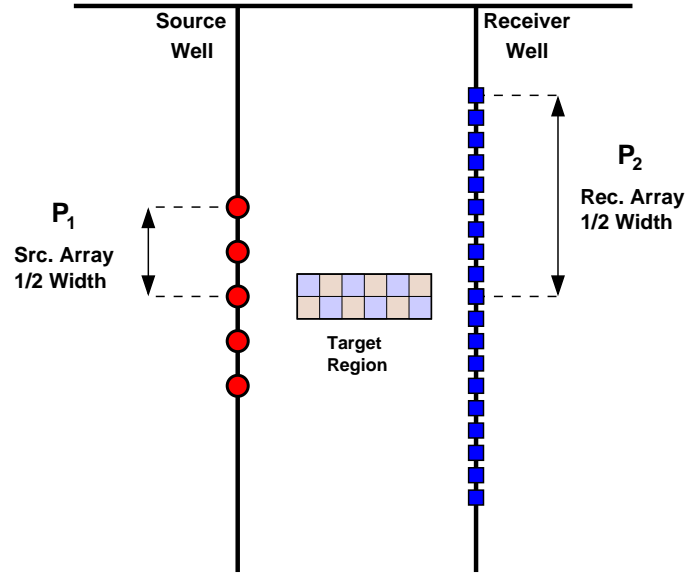


Figure 2: Design variables for a 2 parameter crosswell geometry : The search problem is posed in terms of choosing optimal source (P_1) and receiver (P_2) array dimensions. The target is a depth limited checker-board intended to replicate the reservoir unit targeted for monitoring.

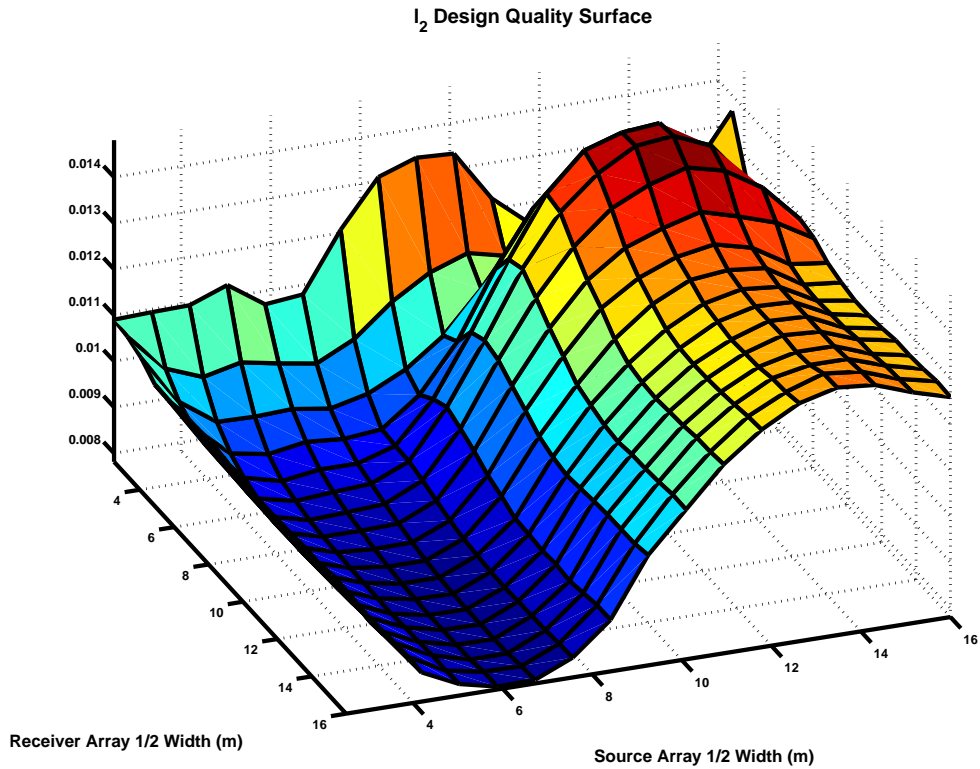


Figure 3: Experiment quality metric for the 2 parameter problem evaluated on a regular grid : Note the trough-shaped minimum corresponding to a narrow range of source array 1/2 widths near 6 m.

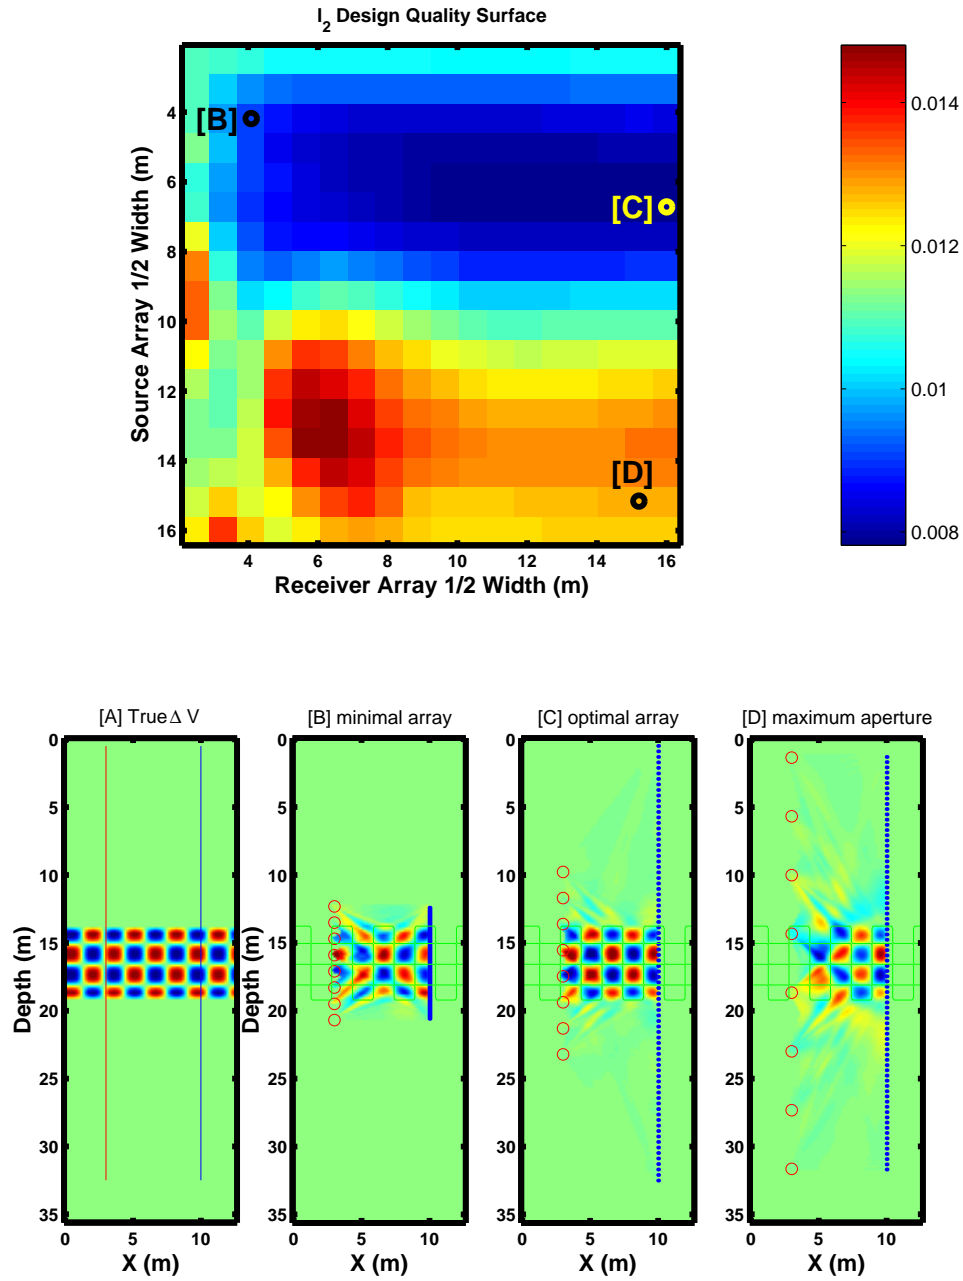


Figure 4: The l_2 objective function surface and 3 designs for the (P_1, P_2) search problem : The top panel depicts the combined objective function sampled on a regular mesh with 3 particular solutions labeled. The bottom panels show the true perturbation model (A), and reconstructed images corresponding to a minimal array (B), a “optimal” array (C), and a wide aperture array (D).

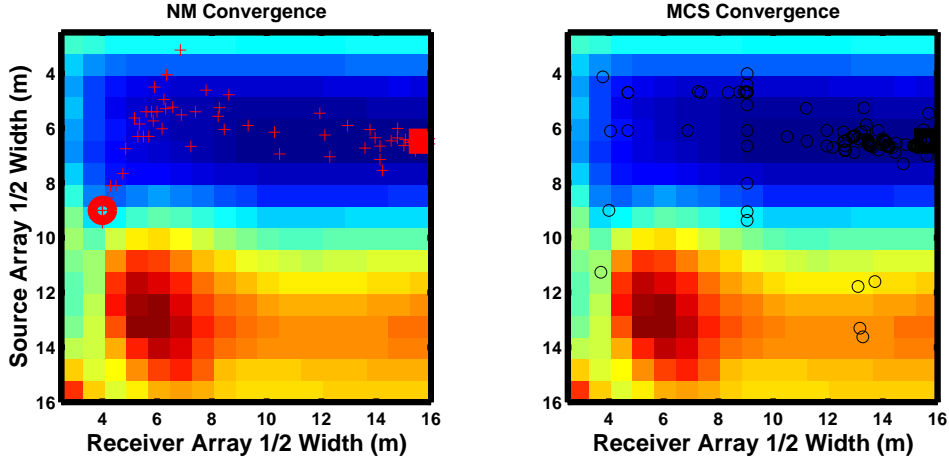


Figure 5: A comparison of the convergence of the Nelder-Mead downhill simplex (NM, shown in the left panel) and the global Multilevel Coordinate Search (MCS, shown in the right panel) algorithms. The bold squares indicate the global optimum which both the NM and MCS methods converge to. The bold red circle on the left panel indicates the NM starting model.

applied to the same test case. In each panel, the background color map is derived from the grid evaluation while the superimposed symbols indicate where each algorithm evaluated the objective function. Both algorithms converge to the correct global minimum as estimated by the original grid search. Given this particular start location (shown as a red circle), the NM solver converged to within tolerance after 130 evaluations of the objective function. The MCS solver was initialized with only the bounds on the domain but required 203 evaluations to achieve the same level of accuracy.

Figure 5 also gives insight into the evaluation patterns of the two algorithms. NM, as with most local search methods, “feels” out the trough in the objective function and does not sample locations outside of this zone of influence. In cases where multiple local minima were present, this behavior would guarantee that NM would be sensitive to start location. In contrast, the MCS solver samples the objective function at several locations outside the trough in zones with significantly worse quality metrics. MCS attempts to balance global and local search performance by cycling through different types of sub-sampling patterns; in this case, after some initial linear sampling steps, the algorithm focused on the region containing the global minimum.

Further testing of the NM algorithm on this particular optimization problem suggests that convergence is not particularly sensitive to start location. Figure 6 shows the convergence behavior of 3 NM runs with different initial solutions. All 3 runs converged to the global minimum and required less than 120 objective function evaluations. Although we suspect that this level of reliability is limited to very simple design problems, the performance of NM in this context suggests that it may be a useful optimization algorithm if given a “hot start” near the global minimum provided by a global search method.

4.1.2 A 4 Parameter Design Problem

While the first test problem provided insight into the simplest class of crosswell array designs, we decided to investigate a higher dimensional problem with more utility to the Frio project. The primary injection unit at Frio is located within relatively close proximity to a salt dome; upward movement of the salt has resulted in faulting near the flank and has imparted an upward dip of about 17° to the target layer. We expected that this significant dip in the target would necessitate vertical re-positioning of the arrays in addition to adjustments in array width. Figure 7 shows this slightly more complicated geometry; \mathbf{q} is now a vector of length 4 including the center depths of the source (P_1) and receiver (P_2) arrays in addition to the source

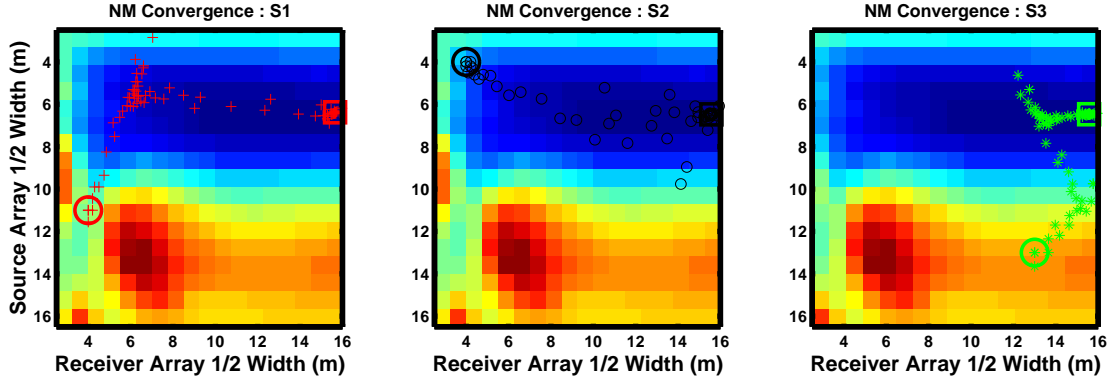


Figure 6: A comparison of the convergence of the NM downhill simplex method for three different starting models : Each panel depicts a single optimization run with the bold circles indicating starting models and the bold squares indicating the final model after convergence.

and receiver array 1/2 widths (P_3, P_4).

Preliminary numerical tests indicated that NM, when applied to this design problem, showed significant sensitivity to starting location; based on this observation we focused on application of the MCS algorithm. Panel [A] of figure 8 shows the test checkerboard model with a dip of 17° to match the injection unit at Frio. Panel [B] depicts the tomographic results for an initial “best guess” geometry where the source and receiver arrays effectively encompassed the target region. Panel [C] shows the optimal geometry as determined by MCS. In this case MCS converged after 333 objective function evaluations. Since only the 4D bounds on \mathbf{q} are required for MCS, panel [B] is included purely for the purpose of illustration.

The optimal geometry shown in figure 8 exhibits several interesting attributes. As in the previous example, the optimal design has improved spatial resolution near the source well and reduces artifacts above and below the imaging target. Like the 2 parameter case, the best design has asymmetric source and receiver array widths. Additionally, the source array is shifted upwards so as to shoot across the dip of the target layer.

4.2 The Next-Generation Permanent Monitoring Array?

While the crosswell arrays shown in the first two examples were small (8 sources x 80 receivers), future pilot deployments are likely to be even sparser due to cost constraints. The number of source levels in particular will likely dominate project budgets. With this reality in mind, we solved the 2 parameter OED problem for source arrays with 2, 4, and 6 levels and a 48 level receiver array. Figure 9 shows the test model (panel [A]) besides the optimal geometries for this set of source array sizes. All three geometries provide some degree of resolution near the receiver well; however, increasing the number of sources provides the coverage and aperture needed to effectively image features near the source well and in the central portion of the imaging target. All three designs also exhibited the asymmetric array lengths present in previous OED tests. Surprisingly, a minimal array with only 6 sources was able to effectively image features across the entire target unit suggesting that similar arrays could provide useful real-time tomographic imagery when permanently installed in the well-bore.

5 Conclusions & Future Developments

Through the examples shown in this paper, we have conclusively demonstrated the benefit of OED algorithms when applied to the design of crosswell seismic acquisition geometries. Our approach combines a simple l_2 image-based experiment quality metric, a parsimonious representation of survey geometry, and a robust

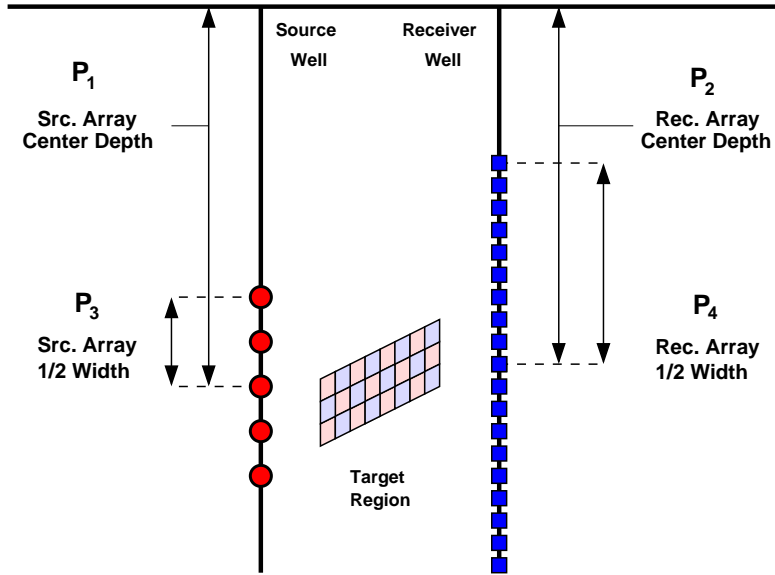


Figure 7: Design variables for a 4 parameter crosswell geometry : The search problem is posed in terms of choosing optimal source and receiver array depths (P_1 , P_2) and 1/2 widths (P_3 , P_4). The target is a depth limited checker-board intended to replicate the reservoir unit targetted for monitoring, in this case possibly tilted to replicate geologic dip.

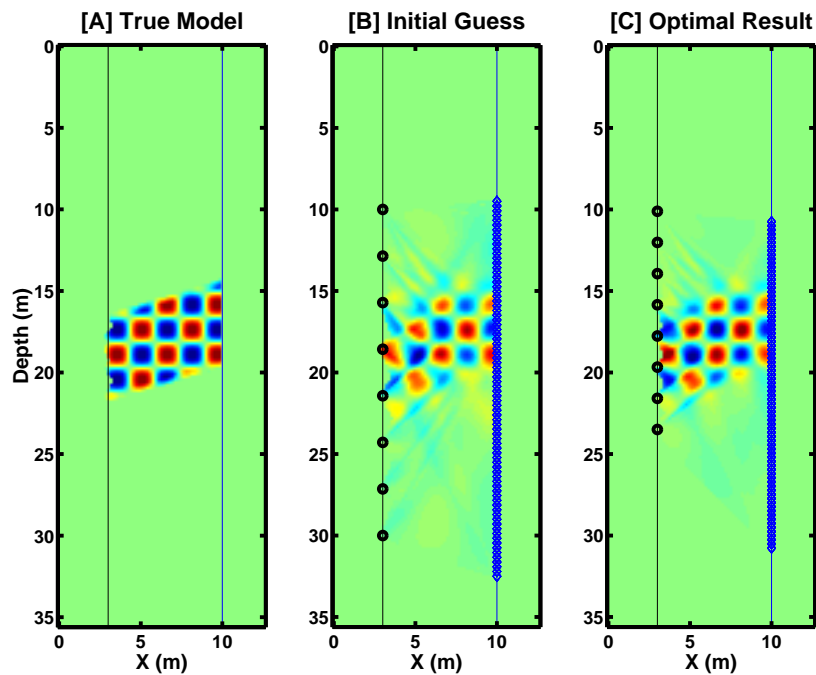


Figure 8: Optimal crosswell design for a 4 parameter system : Shown are the true model perturbation [A], followed by the inversion results for an initial guess at \mathbf{q} [B], and the final optimized experiment [C].

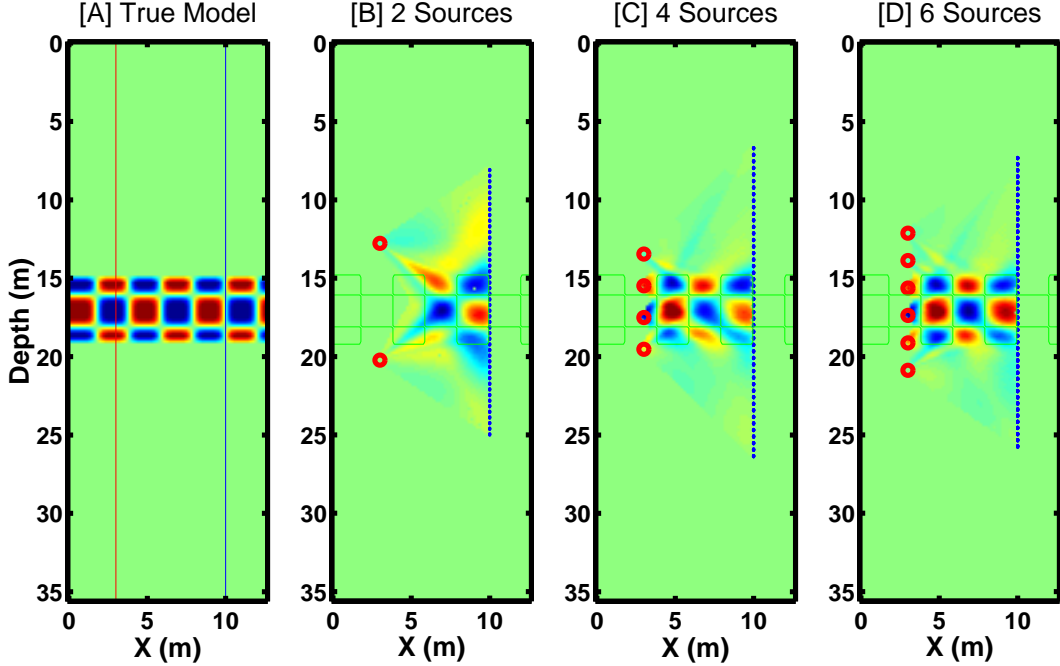


Figure 9: Optimal Designs for 3 sparse arrays : Panel [A] depicts the true checkerboard model while panels [B],[C], and [D] show the optimal acquisition geometries for 2, 4, and 6 sources respectively.

direct search method. These OED components were applied to differential seismic traveltime tomography, a developing imaging technique with application to *in situ* monitoring of flow processes. The resulting methodology performs well for both 2 and 4 parameter optimal experiment design problems and scales to models for which evaluation of SVD metrics would be prohibitively expensive. Our investigation of optimal crosswell monitoring arrays for geometries with significantly fewer sources than receivers has yielded the following observations,

1. Tomographic image quality can be improved by incorporating asymmetry into the length of source and receiver arrays.
2. In addition to asymmetric cable lengths, the resolution of dipping features can potentially benefit from a vertical offset between the two arrays.
3. Useful tomographic imagery can be derived from sparse asymmetric crosswell data, thus opening the door to continuous monitoring of flow processes from within the reservoir.

The examples shown in this paper have been limited to ray-theoretic traveltime tomography. The added benefit of using later arrivals, particularly reflections, in crosswell imaging has been well-established [Harris *et al.*, 1995; Lazaratos *et al.*, 1995]; the next step in our investigation will be to incorporate differential wavefield imaging methods into the same OED framework. Arrays tailored to wavefield imaging will likely have different characteristics than those designed purely for traveltime tomography due to their different spatial sensitivity patterns. Traditional wavefield processing also requires attention to spatial aliasing, such constraints could easily be added to the formulation as soft penalty terms.

Our examination of the Nelder-Mead and Multilevel Coordinate Search algorithms suggests that NM is useful in low-dimensional design problems where the objective function is close to unimodal. For more complicated problems, a global optimization algorithm is almost certainly required but NM may have some

utility after a “hot start” places it near the global minimum. For the 4 dimensional design problem, MCS reliably converged to high quality designs provided only with weak bounds on the design variables. Not having compared MCS to other global optimization techniques, we cannot come to any conclusions concerning its superiority to more commonly used SA or GA methods.

While we are strong advocates of parsimonious descriptions of experiment geometry, exactly how many parameters should be used is a topic open for discussion. Additional degrees of freedom could include linear, quadratic, or higher order variations in source/receiver spacing. While providing this extra flexibility would likely provide some improvement to image quality, the costs of custom fabrication might outweigh these benefits. Ideally, all budgetary components of the survey would be rolled into a single objective function which would provide a trade-off curve between total acquisition cost and image quality. An even more ambitious approach would pose the experiment quality metric in a probabilistic hypothesis testing framework, thus allowing us to formally quantify the cost of accurately answering questions about the subsurface. Our route, while considerably less ambitious, provides useful design parameters which can be calculated in finite time with available computational resources. We hope to utilize this OED approach in the design of several pilot-scale seismic monitoring experiments which will be deployed in the near future.

6 Acknowledgments

We would like to thank the Founding Members Consortium of the Earth Resources Laboratory for their generous support and guidance. The first author benefited from a broad range of discussions with Prof. Jerry Harris, Darrel Coles, and Burke Minsley. Tom Daley provided detailed information on the Frio continuous seismic monitoring project while Dr. Richard Coates and Dr. Chung Chang provided an analogous description of the similar Mallik 2 deployment. We would also like to thank Prof. Colin Zelt for making the source code for the FAST eikonal solver publicly available.

References

- Ajo-Franklin, J., B. Minsley, and T. Daley (2007), Applying compactness constraints to differential seismic traveltimes tomography, *Geophysics*, (*in press*).
- Aldridge, D., and D. Oldenburg (1993), Two-dimensional tomographic inversion with finite-difference traveltimes, *Journal of Seismic Exploration*, *2*, 257–274.
- Barth, N., and C. Wunsch (1990), Oceanographic experiment design by simulated annealing, *J. Phys. Ocean.*, *20*, 1249–1263.
- Blanco, J., S. Knudsen, and F. Bostick (2006), Time-lapse VSP field test for gas reservoir monitoring using permanent fiber optic seismic system, in *Proceedings Of The 93rd Ann. Mtg*, pp. 3447–3451, Soc. of Expl. Geophysics.
- Box, G., and H. Lucas (1959), Design of experiments in non-linear situations, *Biometrika*, *47*, 77–90.
- Cox, B. E. (2004), Tomographic inversion of focusing operators, Ph.D. thesis, Delft University of Technology.
- Curtis, A. (1999a), Optimal design of focused experiments and surveys, *Geophysical Journal International*, *139*, 205–215.
- Curtis, A. (1999b), Optimal experiment design : cross-borehole tomographic examples, *Geophysical Journal International*, *136*, 637–650.
- Curtis, A., and R. Snieder (1997), Reconditioning inverse problems using the genetic algorithm and revised parametrization, *Geophysics*, *62*(4), 524–1532.
- Curtis, A., A. Micheline, D. Leslie, and A. Lomax (2004), A deterministic algorithm for experimental design applied to tomographic and microseismic monitoring surveys, *Geophysical Journal International*, *157*, 595–606.
- Daley, T., L. Myer, and E. Majer (2005), Acquisition of time-lapse, 6-component, P- and S-wave, crosswell seismic survey with an orbital vibrator and of time-lapse VSP for CO₂ injection monitoring, in *75th Ann. Internat. Mtg*, Soc.Of Expl. Geophysocs.
- Daley, T., R. Solbau, J. Ajo-Franklin, and S. Benson (2007), Continuous active-source seismic monitoring of CO₂ injection into a brine aquifer, *Geophysics* (*in review*).
- Dixon, L., and G. Szego (1978), The global optimization problem an introduction, in *Towards Global Optimization 2*, edited by L. Dixon and G. Szego, pp. 1–15, North-Holland, Amsterdam.
- Galbraith, M. (2004), A new methodology for 3d survey design, *The Leading Edge*, pp. 1017–1023.
- Goldberg, D. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison Wesley Longman, Inc.
- Harris, J., R. Nolen-Hoeksema, R. Langan, . M. V. Schaak, S. Lazaratos, and J. R. III (1995), High-resolution crosswell imaging of a West Texas carbonate reservoir: Part 1- Project summary and interpretation, *Geophysics*, *60*, 667–681.
- Hole, J., and B. Zelt (1995), 3-D finite-difference reflection traveltimes, *Geophysical Journal International*, *121*, 427–434.
- Hovorka, S., et al. (2006), Measuring permanence of CO₂ storage in saline formations : the Frio experiment, *Environmental Geosciences*, *13*(2), 1–17.
- Huyer, W., and A. Neumaier (1999), Global optimization by multilevel coordinate search, *Journal of Global Optimization*, *14*, 331–355.

- Ingber, L. (1989), Very fast simulated re-annealing, *Mathematical and Computer Modelling*, *12*, 967–973.
- Ingber, L. (1996), Adaptive simulated annealing (asa) Lessons learned, *Control and Cybernetics*, *25*, 33–54.
- John, R., and N. Draper (1975), D-optimality for regression designs a review, *Technometrics*, *17*, 15–23.
- Jones, D., C. Perttunen, and B. Stuckman (1993), Lipschitzian optimization without the lipschitz constant, *Journal of Optimization Theory and Applications*, *79*, 157–181.
- Lazaratos, S., and B. Marion (1997), Crosswell seismic imaging of reservoir changes caused by CO₂ injection, *The Leading Edge*, *16*, 1300–1306.
- Lazaratos, S., J. Harris, J. R. III, and M. V. Schaack (1995), High-resolution crosswell imaging of a West Texas carbonate reservoir: Part 4 - Reflection imaging, *Geophysics*, *60*, 702–711.
- Leveque, J.-J., L. Rivera, and G. Wittlinger (1993), On the use of the checker-board test to assess the resolution of tomographic inversions, *Geophysical Journal International*, *115*(1), 313,318.
- Liner, C., W. Underwood, and R. Gobeli (1999), 3-d seismic survey design as an optimization problem, *The Leading Edge*, pp. 1054–1060.
- Maurer, H., and D. Boerner (1998), Optimized and robust experimental design : a non-linear application to EM sounding, *Geophysical Journal International*, *132*(2), 458–468.
- Menke, W. (1984), *Geophysical Data Analysis : Discrete Inverse Theory*, Academic Press.
- Morrice, D., A. Kenyon, and C. Beckett (2001), Optimizing operations in 3-d land seismic surveys, *Geophysics*, *66*(6), 1818–1826.
- Muzic, R., A. Nelson, G. Saidel, and F. Miraldi (1996), Optimal experiment design for pet quantification of receptor concentration, *IEEE Transactions on Medical Imaging*, *15*(1).
- Nelder, J., and R. Mead (1965), A simplex method for function minimization, *The Computer Journal*, *7*, 308–313.
- Paige, C., and M. Saunders (1982), An algorithm for sparse linear equations and sparse least squares, *ACM Transactions in Mathematical Software*, *8*, 43–71.
- Press, W., B. Flannery, S. Teukolsky, and W. Vetterling (1992), *Numerical Recipes in C : The Art of Scientific Computing*, Cambridge University Press.
- Ross, C., and S. Altan (1997), Time-lapse seismic monitoring: Some shortcomings in nonuniform processing, *The Leading Edge*, *16*, 931–937.
- Routh, P., G. Oldenborger, and D. Oldenburg (2005), Optimal survey design using the point spread function measure of resolution, in *Proceedings Of The 75th Ann. Mtg, Soc.Of Expl. Geophysics*.
- Sethian, J., and A. M. Popovici (1999), 3-D travelttime computation using the fast marching method, *Geophysics*, *64*(2), 516–523.
- Smit, F., M. Ligtenag, P. Wills, and R. Calvert (2006), Towards affordable permanent seismic reservoir monitoring using the sparse OBC concept, *The Leading Edge*, pp. 454–459.
- Stummer, P., H. Maurer, H. H. and A. G. Green (2002), Optimization of dc resistivity data acquisition : real-time experimental design and a new multielectrode system, *IEEE Transactions on Geoscience and Remote Sensing*, *40*(12), 2727–2735.

- Thompson, M., L. Amundsen, P. Karstadand, J. Langhammer, . H. Nakstad, and M. Eriksrud (2006), Field trial of fibre-optic multi-component sensor system for application in ocean bottom seismic, in *Proceedings Of The 93rd Ann. Mtg*, pp. 1148–1152, Soc. of Expl. Geophysics.
- van den Berg, J., A. Curtis, and J. Trampert (2003), Optimal nonlinear Bayesian experimental design : an application to amplitude versus offset experiments, *Geophysical Journal International*, 155, 411–421.
- Vermeer, G. (2003), 3d seismic survey optimization, *The Leading Edge*, pp. 934–941.
- Vidale, J. (1988), Finite-difference calculation of travel times, *Bull. Seis. Soc. Of America*, 78, 2062–2076.
- Walter, E., and L. Pronzato (1987), Optimal experiment design for nonlinear models subject to large prior uncertainties, *American Journal of Physiology*, 253, 530–534.