PROCEEDINGS, Thirty-Seventh Workshop on Geothermal Reservoir Engineering
Stanford University, Stanford, California, January 30 - February 1, 2012
SGP-TR-194

GLOBAL OPTIMIZATION OF RESERVIOR MODELS BY MAPPING OUT OBJECT FUNCTION MINIMA AND SADDLE POINTS

Manuel Plasencia[1], Andreas Pedersen[1], Andri Arnaldsson[2] and Hannes Jónsson[1,3]

[1]Science Institute of University of Iceland, VR-III
107 Reykjavík , Iceland
[2]Vatnaskil Consulting Engineers, Sudurlandsbraut 50
108 Reykjavík , Iceland
[3]Faculty of Science, University of Iceland VR-III
107 Reykjavík , Iceland
e-mail: mpg2@hi.is

ABSTRACT

The application of a recently developed method for global optimization to determine parameters in a model for multiphase flow in porous media is presented. The calculations are carried out using software for distributed or cloud computing to search for saddle points and minima of an objective function that is obtained from the iTOUGH2 code. The method is illustrated using a simple model calibrated in a two dimensional parameter space, but it can be applied to problems with larger number of parameters and is currently being used to improve a model of the Laugarnes geothermal area in Iceland.

INTRODUCTION

The development of reservoir models often involves inverse modeling, which consists of estimating model parameters from measurements of the system response made at discrete points in space and time. The difference between the model calculation and the data at the calibration points is measured by the objective function. The objective function can, for example, be the sum of squares of the difference between calculated model output and field measured data. The estimation of model parameters is then formulated as an optimization problem where we focus on finding the values of those parameters that minimize the objective function.

Even for small models with only a few parameters, the objective function can have more than one minimum. An example is given below. Multiple minima are more likely in more complex models with larger number of parameters. The main task then becomes that of finding the global minimum of a function, which has several local minima. This is a very challenging problem, illustrated in Figure 1. But, it is also important to know whether additional, local minima which are insignificantly higher (compared with estimated error bars) are also present and could represent equally good parameter sets for practical purposes.

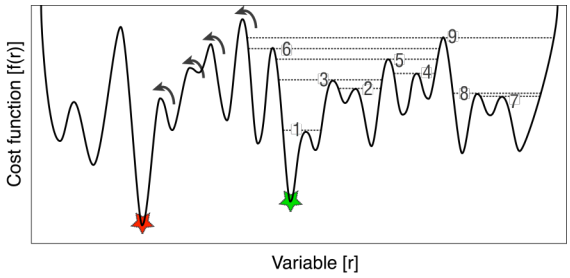

Figure 1: A schematic object function illustrating the problem of multiple local minima. The goal is to find the global minimum (red star) as well as any other local minimum that is only insignificantly higher and represents non-uniqueness in the fit. A calculation starting from the second lowest minimum (green star) using the AKMC algorithm and coarse graining is illustrated. The numbers refer to the sequence of composite states formed.

Global optimization of functions of many variables is often carried out using computer simulated annealing algorithms that mimic roughly the annealing of materials. The 1983 article by Kirkpatrick, Gelatt and Vecchi (Kirkpatrick 1983) illustrated how such an approach could be applied to circuit design. This article has been cited extensively

and the method applied to many different types of objective functions. In every case, the objective function is taken to represent an 'energy' of the system.

A Monte Carlo algorithm based on random numbers is used to simulate an annealing process were changes in the arguments of the objective function are accepted or rejected in accordance with a fictitious 'temperature'. The reason for introducing temperature is to introduce and control the probability of accepting increases in the objective function since they may be an essential intermediate step to ultimately reach lower function values.

A central issue in simulated annealing calculations is the 'time' scale of the 'cooling' of the system from high temperature to zero temperature. The slower the cooling, the more likely the method is to find the global minimum. For a given amount of computational effort, a method that can reach longer time scale is more likely to reach the global minimum.

A method for long time scale simulations based on adaptive kinetic Monte Carlo (AKMC) algorithm has recently been developed (Henkelman 2001). It can be used for global optimization in a way that is analogous to simulated annealing (Pedersen 2012). The basic feature of this approach is an efficient method for moving from one local minimum of the objective function to another via paths that lie close to first order saddle points. The saddle points are found using the minimum mode following method (Henkelman 1999). A kinetic Monte Carlo algorithm is used to select between paths through different saddle points. Alternatively, the path lying through the lowest saddle point and leading to a new local minimum can be chosen, in which case a temperature does not need be defined (Pedersen 2012). The advantage of this algorithm over the original simulated annealing algorithm is that much fewer function evaluations are needed to move from one local minimum to another. Also, the objective function only gets evaluated for parameter values for which the objective function is relatively small, that is parameter values leading to excessively high values of the function are avoided.

A TWO PARAMETER TEST PROBLEM

We illustrate the issues raised and the algorithm proposed here with the two-dimensional optimization problem originate from the inverse modeling of a test sample. This is problem 1 presented by Finsterle as an illustration of the iTOUGH2 code and various optimization methods included there (Finsterle 2007). The task is to determine parameters for a model describing injection of water at constant pressure into a one-dimensional, horizontal column filled with uniform, partially saturated sand. Only two parameters are varied here, the permeability and the porosity. The objective function is the squared deviation between 'observed' and calculated pressure and flow rate at two selected points within the column. Figure 2 shows the shape of the objective function in the parameter space defined by porosity-log(permeability), obtained using the Grid Search Method implemented in iTOUGH2. It is clear from the figure that the objective function has two minima. The global minimun occurs for log(permeability) of -11.7 and porosity of 0.35, but a local minimum occurs at log(permeability) of -12.2 and unphysical large porosity. The two basins surrounding the minima are separated by a ridge along which there is a minimum, corresponding to a first order saddle point on the objective function surface.

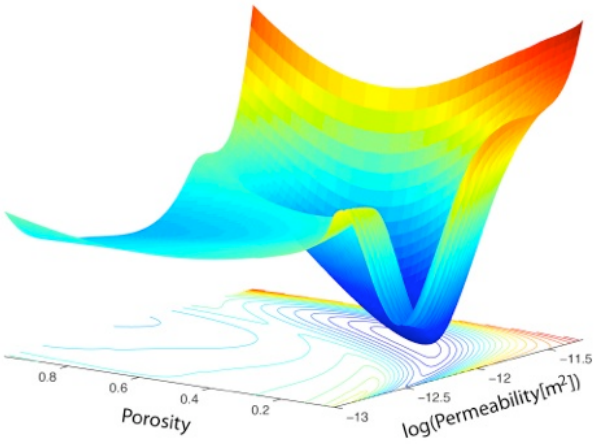

Figure 2: The objective function for the test problem as a function of the two variables: porosity and permeability. (this test problem is taken from the iTOUGH2 sample problem report, LBNL-40042, see Finsterle (2007)).

While the local minimum is not physically reasonable (because the porosity becomes so large) and corresponds to a significantly higher value of the objective function, it will attract minimization paths started within a certain region of parameter space. While it can easily be discounted in this simple two-dimensional problem, this is harder in larger problems involving more parameters. Such problems are also likely to have more local minima, several of which can have reasonable values of the parameters. We use this two parameter example to illustrate the problem, but the goal is to be able to deal with multiple local minima in realistic problems involving several parameters, which likely could arise while

doing a calibration of a model of geothermal reservoir.

FINDING THE NEAREST LOCAL MINIMUM

The Gauss-Newton method is often used for minimizations of objective functions. It is efficient, but only strives to converge onto the local minimum closest to the initial guess. Figure 3 shows contours of the objective function and solution paths, for Gauss-Newton minimization algorithm, obtained using two different initial guesses. We notice that if the initial guess is on the right hand side of the ridge, the method converges to the global minimum. But, if the initial guess is to the left hand side of the ridge, it converges to the local minimum. A large region of parameter space leads to convergence to the local minimum. In this particular case we realize it has no physical meaning but it could be difficult to determine in other situations.

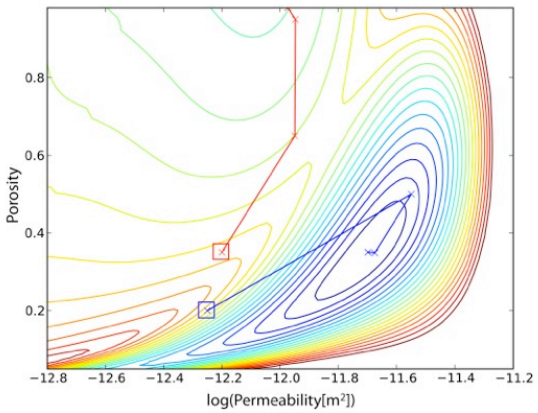

Figure 3: Solution paths of Gauss-Newton minimization algorithm in the 2D parameter space porosity-log(permeability), starting from two different points (squares), one to the right of the ridge and another to the left of the ridge. This illustrates the possibility that a minimization from an initial guess of the parameter values can lead to convergence to a local minimum with a substantially higher value of the objective function than the global minimum.

The Levenberg-Marquardt method can be made to converge more efficiently by selecting appropriate values for convergence parameters, but as the Gauss-Newton it only strives to converge to the minimum closest to the initial guess of the model parameters. Minimization paths generated using this method are shown in Figure 4. The iTOUGH2 code is used to perform the optimization that lead to the solution paths in both Figures 3 and 4. The results for the minimization when using an initial guess to the right of the ridge are similar to the ones given in the iTOUGH2 User's Guide (Finsterle 2007).

This example illustrates the need for exploring the objective function surface beyond just finding the local minimum nearest to the initial guess. While it is easy to envision setting up enough minimization calculations to cover a fine grid of possible initial guesses for all parameters when the number of parameters is small, this will quickly become unmanageable as the number of parameters increases.

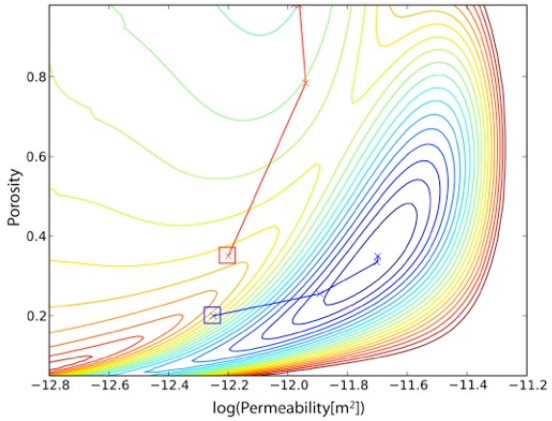

Figure 4: Solution paths of Levenberg-Marquardt minimization algorithm in the 2D parameter space porosity-log(permeability), starting from two different points (squares), one to the right of the ridge and another to the left of the ridge. Similar to Fig. 2, this illustrates the possibility that a minimization from an initial guess of the parameter values can lead to convergence to a local minimum with a substantially higher value of the objective function than the global minimum.

FINDING MULTIPLE MINIMA

The task of finding the global minimum of a function with multiple local minima is a challenging one and the only method that is guaranteed to work is the simulated annealing method with an impossibly slow cooling rate and impossibly large computational effort. A simulated annealing calculation for the two parameter test problem using a finite cooling rate given in the iTOUGH2 User's Guide is shown in Figure 5. While the calculation successfully does find the global minimum, the number of iterations of this kind of approach is large, orders of magnitude larger than for the minimization calculations shown

in Figures 3 and 4. A more efficient method for dealing with objective functions with multiple local minima is clearly needed.

The AKMC algorithm can be used to explore multiple minima of an objective function with computational effort that per minimum is just about two orders of magnitude greater than a single minimization from an initial guess. The basic feature of the algorithm is the ability to climb up the objective function surface to home in on regions around first order saddle points. The algorithm, thereby, identifies new minima that are adjacent to a known minimum. In a simulated annealing formulation the new minimum can be accepted or rejected based in the difference in the values of the objective function and the current value of the temperature (Pedersen 2012). Alternatively, a map of the minima can be generated, with each additional minimum selected based on the height of the first order saddle point in the path to that minimum, as illustrated in Figure 1.

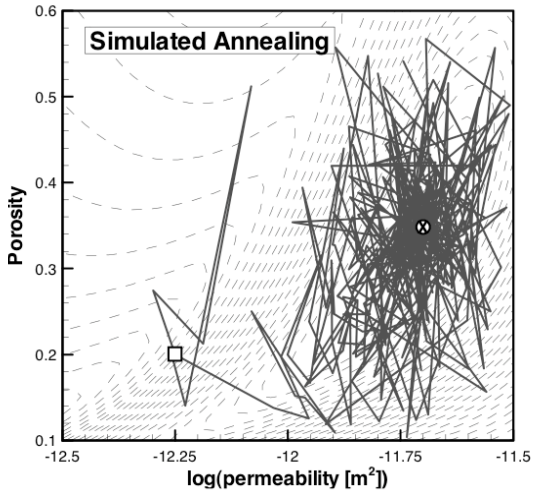

Figure 5: Solution paths of Simulated annealing minimization algorithm in the 2D parameter space porosity-log(permeability), starting from a point to the right of the ridge. The simulation does find the global minimum but requires a very large number of iterations. Picture taken from the iTOUGH2 User's Guide (Finsterle 2007).

The AKMC algorithm is based on the following principle (for more detailed description see Henkelman 2001 and Pedersen 2011). For a given local minimum, several saddle point searches are carried out (on the order of 10 to 100) starting from a random change in the model parameters. A small change in the parameter values at the minimum is generated from a Gaussian random distribution. For each of the perturbed parameter values, the minimum mode following method (Henkelman 1999) is then used to climb up the objective function surface and home in on a saddle point. Such searches are continued until additional searches do not reveal new low-lying saddle points using a probabilistic confidence estimate (Xu 2008).

Two calculations using this algorithm for the two dimensional test problem are shown in Figure 6, one starting from an initial point to the right of the ridge and another from an initial point to the left of the ridge. In either case, both minima are found, the global one and the local one (to within a chosen tolerance in the gradient). Both paths go through the vicinity of the first order saddle point. The tolerance for convergence onto the saddle point can be large since the precise value of the objective function there is not important. The fact that the paths taken from one minimum to another go through the vicinity of saddle points means that parameter regions with very large values of the objective function are avoided, which can be advantageous since unphysical parameters can lead to ill defined values of the objective function and large computational effort.

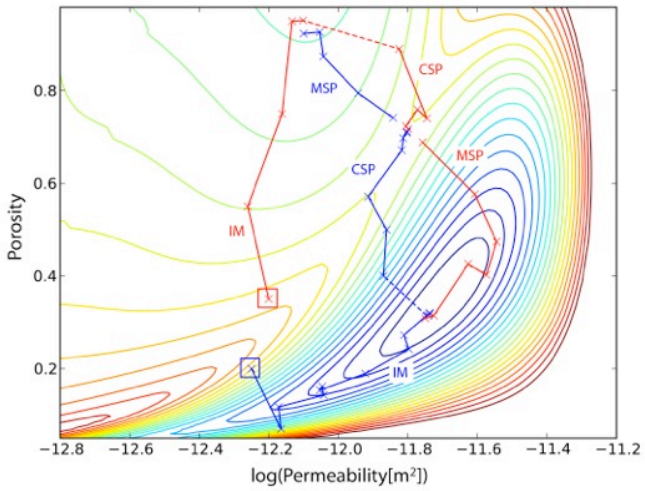

Figure 6: Exploration of the objective function using the AKMC method. Two paths are shown one starting to the right of the ridge (blue) and the other from the left (red). First, a minimization is carried out (IM). Then, an initial increment of parameter values from the minimum is made (dashed line) and the minimum mode following method used to climb up the objective function surface to home in on a saddle point (CSP). Then, an increment of parameter values from the saddle point along the direction of negative curvature away from the initial minimum is made and another minimization carried out (MSP). In either one of the two search paths, both the global and local minima were found.

The most important aspect of the AKMC method is the slow increase in computational effort with the increase in the number of parameters. This method was originally developed to search for transition mechanism and find stable arrangements of atoms in solids. It has been applied successfully to systems with thousands of parameters (atom coordinates in those cases). It has been implemented in software for distributed and cloud computing (Pedersen 2010) making it possible to use multiple CPUs simultaneously connected by simple internet connection. Idle time on computer clusters or personal computers can be used to carry out the calculations. The saddle point searches are farmed out to the various CPU and the saddle points and minima which found are reported back to the server which keeps track of them.

DISCUSSION

The problem of finding the global minimum of an objective function that has many local minima is a challenging one and the only method that is guaranteed to work is the simulated annealing method with an impossibly slow cooling rate requiring infinite computational effort (Kirkpatrick 1983). For object functions that are continuous and differentiable the gradient can be used to navigate on the objective function surface so as to move from one local minimum to another. This assumes the minima can be associated with basins of significant extent and that the surface is not 'rippled'. If these conditions are met, the AKMC method with systematic coarse graining can be used to map out the local minima and not only give an estimate of the global minimum (as the lowest minimum found) but also give an estimate for the uniqueness of the solution found and a few of the essential parts of the objective function.

REFERENCES

S. Finsterle, (2007), "iTOUGH2 User's Guide," *LBNL-40040*, February; "iTOUGH2 Sample Problems," LBNL-40042, February.

G. Henkelman and H. Jónsson, (1999), "A Dimer Method for Finding Saddle Points on High Dimensional Potential Surfaces Using Only First Derivatives," *J. Chem. Phys.*, **111**, 7010.

G. Henkelman and H. Jónsson, (2001), "Long Time Scale Kinetic Monte Carlo Simulations without Lattice Approximation and Predefined Event Table," *J. Chem. Phys.*, **115**, 9657.

S. Kirkpatrick, C.D. Gelatt Jr and M.P. Vecchi, (1983), "Optimization by Simulated Annealing," *Science*, **220**, 671.

A. Pedersen, and H. Jónsson, (2010), "Distributed Implementation of the Adaptive Kinetic Monte Carlo Method," *Mathematics and Computers in Simulation*, **80**, 1487.

A. Pedersen, S.F. Hafstein and H. Jónsson, (2011), "Efficient Sampling of Saddle points with the Minimum Mode Following Method," *SIAM Journal on Computational Science*, **33**, 633.

A. Pedersen, J-C. Berthet and H. Jónsson, (2012), "Simulated Annealing with Coarse Graining and Distributed Computing," *Lecture Notes in Computer Science*, **7134**, 34.

L. Xu and G. Henkelman, (2008), "Adaptive kinetic Monte Carlo for first-principles accelerated dynamics," *J. Chem. Phys.*, **129**, 114104.