PROCEEDINGS, TOUGH Symposium 2012
Lawrence Berkeley National Laboratory, Berkeley, California, September 17-19, 2012

GEOTHERMAL MODEL CALIBRATION USING A GLOBAL MINIMIZATION ALGORITHM BASED ON FINDING SADDLE POINTS AS WELL AS MINIMA OF THE OBJECTIVE FUNCTION

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

Here we present a specific method for global optimization of parameters within a model for multiphase flow in porous media. This model, which involves a search for first-order saddle points and the minima of an objective function obtained from the iTOUGH2 code, has been implemented within the EON software for distributed and cloud computing. While it can be applied to problems with a large number of parameters, the application presented here is a simple illustrative model of the Laugarnes geothermal area in Reykjavík, Iceland, calibrated in two-dimensional parameter space.

INTRODUCTION

The development of reservoir models often involves inverse modeling, which consists of estimating model parameters from measurements of system response made at discrete points in space and time. The difference between model calculations and the data at calibration points is measured by the objective function. This objective function could, for example, be the sum of the squares of the difference between calculated model output and field-measured data. Estimating model parameters is then formulated as an optimization problem, in which the goal is to find the parameter values that minimize the objective function.

Even for small models with only a few parameters, the objective function can have more than one minimum. This is illustrated in Figure 1, which shows a 1D cut through an objective function for the model described below. Within the interval shown, three local minima appear.

The occurrence of multiple minima is more likely in models with a larger number of parameters. The main task then becomes that of finding the global minimum of a function among the several local minima—a very challenging problem. It is also important to know whether additional local minima, which are insignificantly higher (compared with estimated error bars), are present and could represent equally good parameter sets for practical purposes.

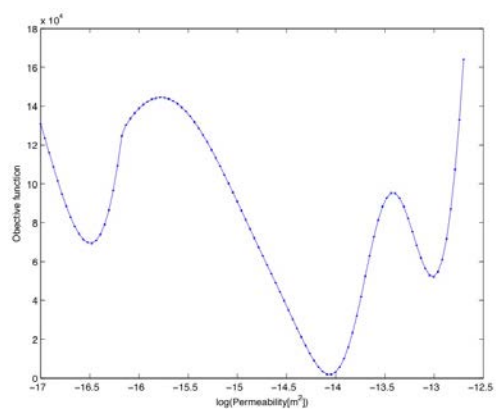

Figure 1. A 1-dimensional cut through the objective function for the model system studied here, illustrating the problem of multiple local minima. The goal is to find the global minimum (vicinity of -14.1) among the local minima.

Global optimization of functions of many variables is often carried out using simulated annealing algorithms that roughly mimic 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 since been cited exten-

sively, and the authors' method applied to many different optimization problems. There, the objective function is taken to represent an "energy" of the system, and a fictitious temperature is introduced. By applying a Monte Carlo algorithm, we can simulate the annealing process by accepting or rejecting changes in the arguments of the objective function. 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 in ultimately reaching lower function values.

A central issue in simulated annealing calculations is the time scale of the "cooling" of the system, from high temperatures to zero temperature. The lower the cooling rate, the more likely the method will find the global minimum. For a given amount of computational effort, a method that can simulate longer time intervals is therefore more likely to reach the global minimum.

One method for long time-scale simulations, known as the adaptive kinetic Monte Carlo (AKMC) algorithm, has recently been developed (Henkelman 2001) in the context of transitions for atomic-scale systems. It can be used for global optimization in a way that is analogous to simulated annealing (Pedersen 2012). The important feature of this approach is the ability to move 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). The kinetic Monte Carlo algorithm is used to select between the determined possible paths through different saddle points. Alternatively, the path lying through the lowest saddle point 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 fewer objective-function evaluations are needed to move from one local minimum to another. Also, the objective function only gets evaluated for parameter values for which its value is relatively small, while the regions with excessively high values for the objective function are avoided.

CONCEPTUAL AND NUMERICAL MODEL

Here, we apply this algorithm to a specific geothermal modeling problem. Laugarnes is a low-temperature field in Iceland, which has been described in some detail by Thorsteinsson and Eliasson (1970). This geothermal area is fed by three aquifers: Aquifer A with water of 110–120°C, Aquifer B with water of 135°C and Aquifer C with water of temperature of 146°C. Tuffs and sediments act as aquicludes between the aquifers. The active reservoir underlies an area of 5 km^2 within the city of Reykjavik and has a base temperature about 145°C (Bodvarsson 1978). Prior to exploitation, the hydrostatic pressure at the surface in the geothermal field was 6–7 bars (Einar Gunnlaugsson et al., 2000) and about 10 L/s of 88°C water issued in free flow from the hot spring. (Thorsteinsson and Eliasson 1970).

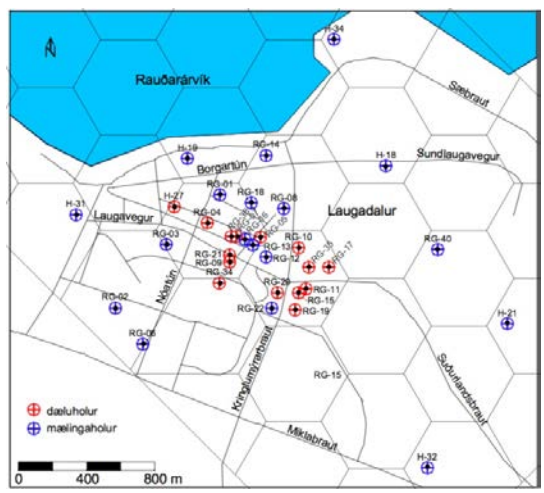

Figure 2. Aerial partial view of Voronoi mesh used to model the area. Red and blue circles production and observation wells respectively.

To build a simplified model for this area, we created a mainly hexagonal Voronoi mesh with 38 volume elements, covering an area of 12 km^2 (Figure 2). The model extends to 2235 m depth in eight layers. There is a single volume element at layer 1 and at layer 8, both of which are inactive and represent the reservoir top and bottom. Layers 3, 5, and 7 represent aquifers A, B and C respectively; layers 2, 4, and 6 represent aquicludes and were assigned lower permeability values (Figure 3).

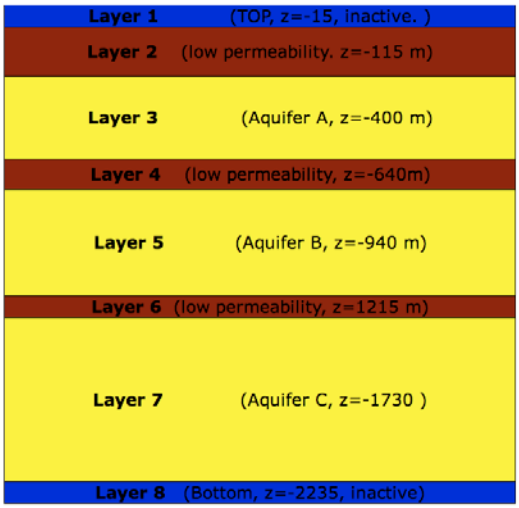

Figure 3. View of layers used in the model. Colors correspond to different material properties.

A 3D representation of the geometry of the model can be seen in Figure 4. The red point at the surface represents the location of an observation well, while the blue point represents the location of a hot spring. Thus, a well on deliverability was defined there acting as a sink.

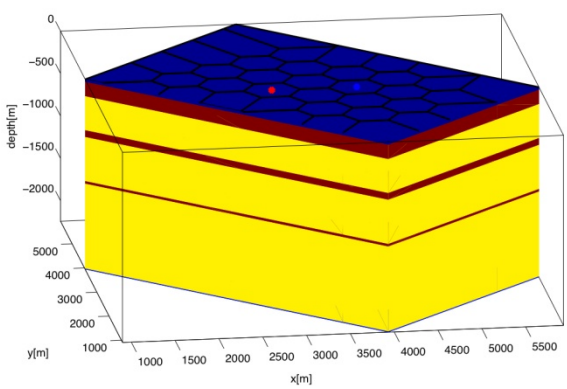

Figure 4. 3D view of the model. Colors correspond to different material properties The blue and red points at the top layer represent the location of hot spring (water sink) and an observation well respectively.

Two types of sources have been included: first, a mass source located at the bottom of the reservoir was positioned in the area where the upflow is thought to be located. Second, heat sources where placed at scattered positions on the bottom (Figure 5).

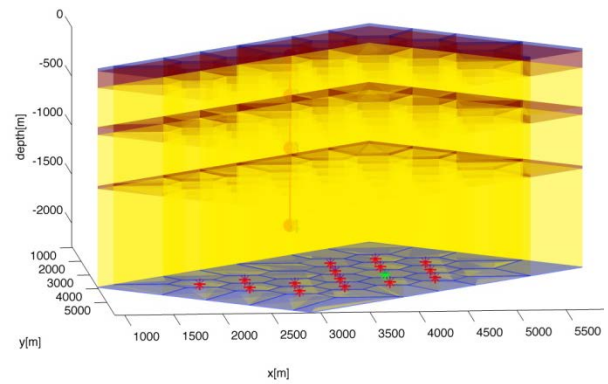

Figure 5. Semi-transparent 3D view of the model. The red starts at the bottom layer represents the heat sources distribution and the green start signifies a water source.

Six calibration points are used (Figure 6). The red line in Figure 6 represents the observation well; it has four calibration points at four different depths. Starting from the top, point 1 represents pressure at the top of reservoir. Points 2, 3 and 4 represent temperatures in aquifer A, B, and C respectively. Points 5 and 6 are in the same location and represent the water flow rate from the hot spring and enthalpy of its water.

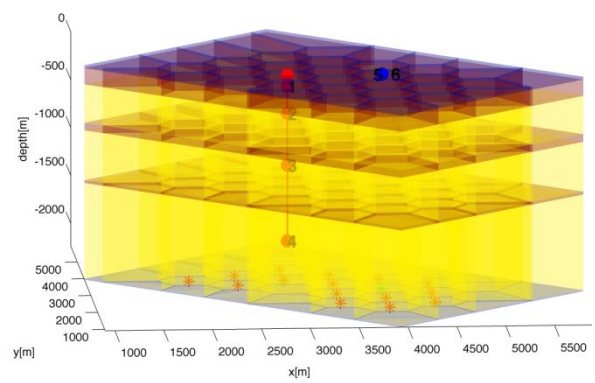

Figure 6. Semi-transparent 3D view of the model. From the top, points from 1 to 4 in red color and 5 to 6 in blue color are calibration points. Points 5 and 6 coincide in the same position.

This model has been constructed to represent a realistic system, but at the same time be simple enough to allow tests and studies of the performance of the optimization algorithm. Focus is mainly on the inverse modeling, so it was decided to generate the data at the calibration points, add some Gaussian noise, and then use this as a representation of the real data for the

natural-state calibration process. Reasonable values were chosen for the fixed parameters in the model, i.e., those that will not change during the optimization, in order to produce a model result close to what already has been reported for this region. Fitting the generated data instead of what is measured has the main advantage that the location of the minimum is known beforehand, which makes the analysis of the optimization algorithm simpler.

OBJECTIVE FUNCTION

The objective function may be close to quadratic or highly nonlinear in nature; it may be continuous, differentiable, and smooth; or discontinuous, not differentiable, and rough. For a nonlinear model, the topography of the objective function away from the minimum becomes intricate, making it difficult for the optimization algorithm to iteratively proceed towards the minimum (Finsterle, 2007).

The two-dimensional optimization problem addressed here originates from the calibration of the simplified model described above. In this case, the objective function is the squared deviation between "observed" and calculated pressure at calibration point 1, temperature at calibration points 2 to 4, and water flow rate and water enthalpy at calibration points 5 and 6 respectively. Figure 7 shows the shape of the objective function in the parameter space defined by the logarithm of mass generation rate (log(q)) and the logarithm of permeability (log(k)) for the aquiclude layers—obtained using the Grid Search Method implemented in iTOUGH2. It should be noticed that within the interval for the permeability between -17.0 and 15.5 the objective function is not smooth. It furthermore shows some point-like discontinuities that probably are caused by numerical instabilities in the forward model. Thus, our focus will mainly be on the smooth regions.

Note that the contour lines in Figure 8 show that the objective function has three minima. The global minimum is known to be located for log(k)=-14.00 and log(q)=1.00, (will be referred as M2), but two local minima also occur, one at log(k)=-16.46 and log(q)=1.10 (M1) and another at log(q)=12.79 and log(q)=1.76 (M3). Between minima M1 and M2, there is a first-order saddle point (SP1), and between M2 and M3, there is another first-order saddle point (SP2).

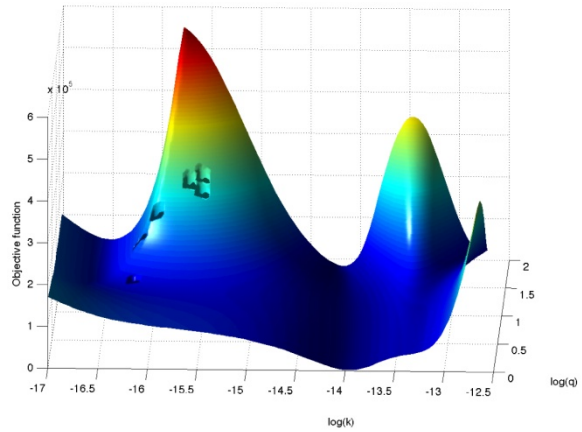

Figure 7. The objective function for the simplified model, as a function of two variables: log(generation rate) on the x-axis and log(permeability) on the y-axis. Between -17.0 and 15.5 it shows some point-like discontinuities that probably are caused by numerical instabilities in the forward model

Both local minima correspond to a significantly higher value of the objective function, but they will attract minimization paths started from nearby regions in parameter space. While it can easily be discounted in this simple two-dimensional problem, this can be a severe problem for models involving many parameters. Furthermore, more complex problems are also likely to have more local minima, several of which can have reasonable values for the parameters.

FINDING THE NEAREST LOCAL MINIMUM

The Levenberg-Marquardt minimization algorithm is found to perform well for most iTOUGH2 applications (Finsterle, 2007). It can be made to converge efficiently by selecting appropriate values for convergence parameters, but as for most other local minimization algorithms, it only strives to converge to the minimum closest to the initial guess of the model parameters. Figure 8 shows the solution paths, for the Levenberg-Marquardt minimization algorithm, starting from four different initial guesses. If the initial guess is in vicinity of the global minimum, the method converges to it. But if the initial guess is not in the proximity of

the global minimum, the method converges to the higher local minima. A region of parameter space leads to convergence to local minima.

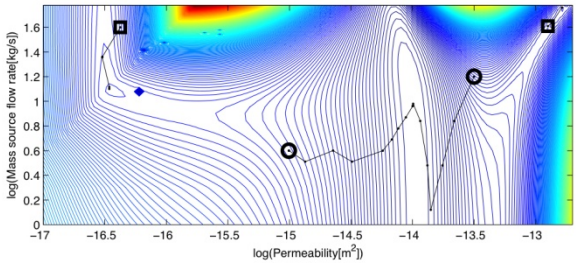

Figure 8: Contour lines of the objective function and solution paths of Levenberg-Marquardt minimization algorithm in the 2D parameter space log (g)-log(k), starting from 4 different points. The squares represent initial guesses that lead to convergence to local minima. The circles represent initial guesses that lead to convergence to the global minimum. 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.

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.

SEARCHING FOR MULTIPLE MINIMA

The task of finding the global minimum of a function with multiple local minima is challenging—the only method that is guaranteed to work is a simulated annealing method requiring an impossibly slow cooling rate and impossibly large computational effort. Therefore, we need a more efficient method for dealing with objective functions with multiple local minima.

The AKMC algorithm can be applied to explore functions with multiple minima. The basic feature of this algorithm is the ability to climb up the objective function surface to focus on regions around first-order saddle points. The algorithm thereby gains the ability to reach new minima adjacent to a known minimum. In a simulated annealing formulation, the new minimum can be accepted or rejected based on the difference in objective-function values 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 on the path to the minimum.

The AKMC algorithm works by the following principle (for a 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. To initialize each search, we impose a small change in the parameter values at the minimum, 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 first-order saddle point. Such searches are continued, using a probabilistic confidence estimate, until additional searches reveal no new low-lying saddle points (Xu 2008).

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 mechanisms 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 both distributed computing and cloud computing (Pedersen 2010), making it possible to use multiple CPUs simultaneously connected by a 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 CPUs, and the saddle points and minima found are reported back to the server, which keeps track of them.

We carried out calculations using this algorithm for the simplified test problem, one starting from an initial point close to the global minimum and

another from an initial point close to a local minimum. In either case, both the global and the local minimum are found (to within a chosen tolerance in the gradient).

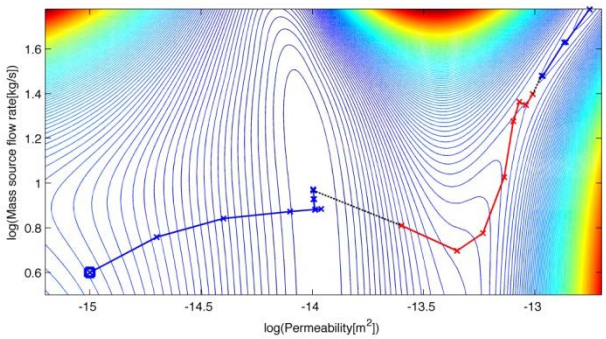

Figure 9. AKMC minimization path from initial guess represented as a blue square. It converges to the global minimum (path in blue), given an initial displacement (black dotted line) it climbs up to converge to the saddle point (path in red), given a displacement from saddle point (black dotted line) it converges also to the local minimum.

Figure 9 shows the minimization path using AKMC. From the initial guess (-15.0, 0.6), represented in the figure as a blue square, the path converges to the global minimum through the blue path. Then, after a small increment in parameter values, it starts a climb up the objective-function surface and converges on a first-order saddle point to a rather loose tolerance. After a displacement along the mode for which the saddle point is a maximum, a minimization converges to the adjacent local minimum.

Similarly, Figure 10 shows the minimization path using a different initial guess (-12.9, 1.61). It converges to a local minimum and, after a displacement, climbs up to converge to a saddle point. Another displacement is given from the saddle point, which then converges to the global minimum.

Both paths go through the vicinity of the first-order saddle point. The tolerance for the 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 objective-function values are avoided, which can be advantageous, since unphysical parameters can lead to ill-defined values of the objective function, as well as a large computational effort.

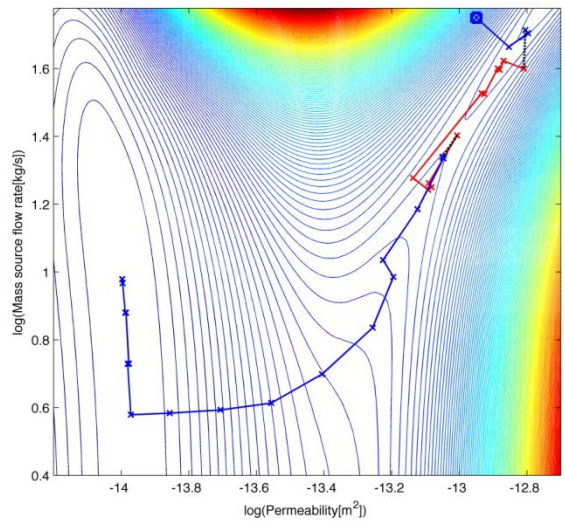

Figure 10. AKMC minimization path from initial guess represented as a blue square. It converges to the local minimum (path in blue), given an initial displacement (black dotted line) it climbs up to converge to the saddle point (path in red), given a displacement from saddle point (black dotted line) it converges also to the global minimum.

DISCUSSION

The problem of finding the global minimum of an objective function that has many local minima is a challenging one—with the only method guaranteed to work being a 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 (Pedersen 2012) can be used to map out the local minima. Not only will this method provide an estimate of the global minimum (as the lowest minimum

found), but it will also give an estimate for the uniqueness of the solution found, and the most important parts of the objective function.

REFERENCES

Bodvarsson, G., and Elliot Zais. A field example of free surface testing. Report, 1978.

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

Gunnlaugsson, E., G. Gislason, G. Ivarsson, and S. P. Kjaran. Low temperature geothermal fields utilized for district heating in reykjavik, iceland. Proceedings World Geothermal Congress 2000 Kyushu - Tohoku, Japan, May 28 - June 10, 2000.

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

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

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

Pedersen, A., L. Pizzagalli and H. Jónsson. Finding mechanism of transitions in complex systems: Formation and migration of dislocation kinks in a silicon crystal, Journal of Physics: Condensed Matter; 21, 084210, 2009.

Pedersen A., and H. Jónsson. Simulations of Hydrogen Diffusion at Grain Boundaries in Aluminum, Acta Material; 57, 4036-4045, 2009.

Pedersen, A., G. Henkelman, J. Schiøtz and H. Jónsson. Long Time Scale Simulation of a Grain Boundary in Copper, New Journal of Physics; 11, 073034, 2009.

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

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

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

Thorsteinsson, T., and J. Eliasson. Geohydrology of the Laugarnes hydrothermal system in Reykjavik, Iceland. U.N. Symposium on the Development and Utilization of Geothermal Resources, Pisa. 1970.

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