

Lawrence Berkeley National Laboratory

Recent Work

Title

Analysis of the Hydraulic Data from the MI Fracture Zone at the Grims el Rock Laboratory, Switzerland

Permalink

<https://escholarship.org/uc/item/14c536dz>

Authors


Davey, A.
Karasaki, K.
Long, J.C.S.
et al.

Publication Date

1989-10-01

**NAGRA - DOE
COOPERATIVE
PROJECT REPORT**

Analysis of the Hydraulic Data from the MI Fracture Zone at the Grimsel Rock Laboratory, Switzerland



*Amy Davey, Kenzi Karasaki,
Jane C. S. Long, Martin
Landsfeld, Antoine Mensch,
and Stephen J. Martel*

October 1989

For Reference

Not to be taken from this copy

Swiss National Cooperative for
the Storage of Nuclear Waste
NAGRA
CH-5401 Baden, Switzerland

Earth Sciences Division
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720, USA

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

Bldg. 50 Library.
Copy 1

LBL-27864

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL-27864
NDC-15

**Analysis of the Hydraulic Data from the MI Fracture Zone
at the Grimsel Rock Laboratory, Switzerland**

Amy Davey, Kenzi Karasaki, Jane C. S. Long, Martin Landsfeld,
Antoine Mensch, and Stephen J. Martel

Earth Sciences Division
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720

October 1989

This work was supported by the Manager, Chicago Operations, Repository Technology Program, Repository Technology and Transportation Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098 and by the Swiss National Cooperative for the Storage of Nuclear Waste (NAGRA).

Preface

This report is one of a series documenting the results of the Nagra-DOE Cooperative (NDC-I) research program in which the cooperating scientists explore the geological, geophysical, hydrological, geochemical, and structural effects anticipated from the use of a rock mass as a geologic repository for nuclear waste. This program was sponsored by the U. S. Department of Energy (DOE) through the Lawrence Berkeley Laboratory (LBL) and the Swiss Nationale Genossenschaft für die Lagerung radioaktiver Abfälle (Nagra) and concluded in September 1989. The principal investigators are Jane C. S. Long, Ernest L. Majer, Karsten Pruess, Kenzi Karasaki, Chalon Carnahan and Chin-Fu Tsang for LBL and Piet Zuidema, Peter Blümling, Peter Hufschmied and Stratis Vomvoris for Nagra. Other participants will appear as authors of the individual reports. Technical reports in this series are listed below.

1. Determination of Fracture Inflow Parameters with a Borehole Fluid Conductivity Logging Method by Chin-Fu Tsang, Peter Hufschmied, and Frank V. Hale (NDC-1, LBL-24752).
2. A Code to Compute Borehole Fluid Conductivity Profiles with Multiple Feed Points by Frank V. Hale and Chin-Fu Tsang (NDC-2, LBL-24928; also NTB 88-21).
3. Numerical Simulation of Alteration of Sodium Bentonite by Diffusion of Ionic Groundwater Components by Janet S. Jacobsen and Chalon L. Carnahan (NDC-3, LBL-24494).
4. P-Wave Imaging of the FRI and BK Zones at the Grimsel Rock Laboratory by Ernest L. Majer, John E. Peterson Jr., Peter Blümling, and Gerd Sattel (NDC-4, LBL-28807).
5. Numerical Modeling of Gas Migration at a Proposed Repository for Low and Intermediate Level Nuclear Wastes at Oberbauenstock, Switzerland by Karsten Pruess (NDC-5, LBL-25413).
6. Analysis of Well Test Data from Selected Intervals in Leuggern Deep Borehole — Verification and Application of PTST Method by Kenzi Karasaki (NDC-6, LBL-27914).
7. Shear Wave Experiments at the U. S. Site at the Grimsel Laboratory by Ernest L. Majer, John E. Peterson Jr., Peter Blümling, and Gerd Sattel (NDC-7 LBL-28808).
8. The Application of Moment Methods to the Analysis of Fluid Electrical Conductivity Logs in Boreholes by Simon Loew, Chin-Fu Tsang, Frank V. Hale, and Peter Hufschmied (NDC-8, LBL-28809).
9. Numerical Simulation of Cesium and Strontium Migration through Sodium Bentonite Altered by Cation Exchange with Groundwater Components by Janet S. Jacobsen and Chalon L. Carnahan (NDC-9, LBL-26395).
10. Theory and Calculation of Water Distribution in Bentonite in a Thermal Field by Chalon L. Carnahan (NDC-10, LBL-26058).
11. Prematurely Terminated Slug Tests by Kenzi Karasaki (NDC-11, LBL-27528).
12. Hydrologic Characterization of Fractured Rocks — An Interdisciplinary Methodology by Jane C. S. Long, Ernest L. Majer, Stephen J. Martel, Kenzi Karasaki, John E. Peterson Jr., Amy Davey, and Kevin Hestir, (NDC-12, LBL-27863).
13. Exploratory Simulations of Multiphase Effects in Gas Injection and Ventilation Tests in an Underground Rock Laboratory by Stefan Finsterle, Erika Schlueter, and Karsten Pruess (NDC-13, LBL-28810).
14. Joint Seismic, Hydrogeological, and Geomechanical Investigations of a Fracture Zone in the Grimsel Rock Laboratory, Switzerland by Ernest L. Majer, Larry R. Myer, John E. Peterson Jr., Kenzi Karasaki, Jane C. S. Long, Stephen J. Martel, Peter Blümling, and Stratis Vomvoris (NDC-14, LBL-27913).
15. Analysis of Hydraulic Data from the MI Fracture Zone at the Grimsel Rock Laboratory, Switzerland by Amy Davey, Kenzi Karasaki, Jane C.S. Long, Martin Landsfeld, Antoine Mensch, and Stephen J. Martel (NDC-15, LBL-27864).
16. Use of Integrated Geologic and Geophysical Information for Characterizing the Structure of Fracture Systems at the US/BK Site, Grimsel Laboratory, Switzerland by Stephen J. Martel and John E. Peterson Jr. (NDC-16, LBL-27912).

Table of Contents

List of Figures	vii
List of Tables	xiii
Abstract	xv
1.0. Introduction	1
2.0. Simulated Annealing	3
2.1. Annealing Theory	7
2.2. A Synthetic Example	11
2.3. Model Validation and Measures of Uncertainty	12
2.3.1. Annealing and Model Uncertainty	12
2.3.2. Prediction Error	16
2.3.3. Sensitivity Analysis	18
3.0. Synthetic Case Studies	21
3.1. The Synthetic Case	21
3.2. Effect of the Geometry of the Template	22
3.3. Effect of the Starting Point	24
4.0. The Migration Experiment Site	27
4.1. Geology of the Southern Part of the Grimsel Rock Laboratory	33
4.2. Structure of the MI Fracture Zone	36
4.3. Hydrogeologic Implications of Fracture Structure at the MI Site	40
5.0. Hydraulic Tests at MI	41
5.1. Hydraulic Tests	41
5.2. Boundary Conditions	43
6.0. The Migration Site Case	49
6.1. Steady State Results	50
6.2. Transient Results	56
7.0. Conclusions and Recommendations	65
7.1. Modelling Approach and Conclusions	65
7.2. Recommendations for Future Study	69

8.0. References	71
Appendix A: Implementation of the Simulated Annealing Algorithm	73
Appendix B: The Cross-Validation Studies	79

List of Figures

	Page	
Figure 2.1.	The synthetic case used to generate well test data for use in annealing. Dots represent points where "well" data were generated, the central being the pumping well. Scale can be considered dimensionless.	13
Figure 2.2.	An example template developed for annealing the synthetic well test data.	13
Figure 2.3.	A configuration resulting from annealing the synthetic well test data.	14
Figure 2.4.	The dimensionless energy versus iteration curve for the synthetic annealing case.	14
Figure 3.1.	Annealing solutions to the synthetic case shown in Figure 2.1 for different templates 1 through 6.	23
Figure 4.1.	Map of Switzerland showing the location of the Grimsel Laboratory.	28
Figure 4.2.	The layout of the Grimsel Rock Laboratory showing the location of the MI fracture zone.	29
Figure 4.3.	Geologic map showing the major structures at the surface above the Grimsel laboratory. The MI zone is part of a group of fractured shear zones (shown by heavy lines) that extend southwest from Lake Raterichsboden. From Nagra Technical Report 87-14.	30
Figure 4.4.	Geologic cross section and generalized map showing major geologic structures in and near the MI site extending to the surface. The line of cross section is shown on Figure 4.3. From Nagra Technical Report 87-14.	31

Figure 4.5.	Geologic setting of the southern part of the Grimsel Rock Laboratory.	32
Figure 4.6.	Logs of fractures in the MI fracture zone in (a) the AU tunnel and (b) the laboratory tunnel. From preliminary draft of Nagra Technical Report 87-14 by Keusen and others.	34
Figure 4.7.	The AU tunnel and MI fracture zone intersected by 8 boreholes. From Nagra Report 1988.	35
Figure 4.8.	Photograph showing braided structure of fractures in a NE-striking fracture zone exposed at the surface above the Grimsel laboratory. Fifteen-cm ruler for scale.	37
Figure 4.9.	Diagram showing the three-dimensional braided fracture structure of the NE-striking fracture zones at Grimsel. The braided pattern appears more drawn out in vertical cross sections than in plain view.	38
Figure 4.10.	Diagram of the MI "fracture" from Frick and others, 1988.	39
Figure 5.1.	Pressure data observed during six constant pressure tests (Solexperts, 1988).	42
Figure 5.2.	Pressure records from the constant flow withdrawal test followed by a buildup test (Solexperts, 1989).	44
Figure 5.3.	The adjusted flow rate in BOMI 87.009.	44
Figure 5.4.	Pressures in the exploratory boreholes configuration before the laboratory tunnel excavation (after Nagra, 1985).	46
Figure 5.5.	Estimated pressure distribution in meters around the tunnels. The rectangle represents the region modeled.	47
Figure 6.1.	The MI template.	51
Figure 6.2a.	The first of five solutions which have similar geometries, case 1.	51
Figure 6.2b.	The second of five solutions which have similar geometries, case 2.	52
Figure 6.2c.	The third of five solutions which have similar geometries, case 3.	52

Figure 6.2d.	The fourth of five solutions which have similar geometries, case 4.	53
Figure 6.2e.	The fifth of five solutions which have similar geometries, case 5.	53
Figure 6.3.	Energy versus iteration curves for the five steady state annealing solutions.	55
Figure 6.4.	The "good" portion of well test curves used to model the system.	58
Figure 6.5.	Template used for transient annealing of MI data.	60
Figure 6.6.	Simulated well test response on the full template before annealing.	61
Figure 6.7.	The annealed configuration at iteration number 15510.	61
Figure 6.8.	Simulated well test response on the annealed configuration.	62
Figure 6.9	The energy versus iteration number for the transient annealing case.	62
Figure A.1.	An annealing solution found starting from a configuration with 60% of the possible elements.	74
Figure A.2.	An annealing solution found starting from a configuration with 60% of the possible elements.	74
Figure A.3.	An annealing solution found starting from a configuration with 70% of the possible elements.	75
Figure A.4.	An annealing solution found starting from a configuration with 70% of the possible elements.	75
Figure A.5.	An annealing solution found starting from a configuration with 80% of the possible elements.	76
Figure A.6.	An annealing solution found starting from a configuration with 80% of the possible elements.	76
Figure A.7.	An annealing solution found starting from a configuration with 90% of the possible elements.	77

Figure A.8.	An annealing solution found starting from a configuration with 90% of the possible elements.	77
Figure A.9.	An annealing solution found starting from a configuration with 100% of the possible elements.	78
Figure A.10.	An annealing solution found starting from a configuration with 100% of the possible elements.	78
Figure B.1.	Cross-Validation Study. Annealing solution case 1, with well 6 left out.	80
Figure B.2.	Cross-Validation Study. Annealing solution case 2, with well 6 left out.	80
Figure B.3.	Cross-Validation Study. Annealing solution case 3, with well 6 left out.	81
Figure B.4.	Cross-Validation Study. Annealing solution case 4, with well 6 left out.	81
Figure B.5.	Cross-Validation Study. Annealing solution case 5, with well 6 left out.	82
Figure B.6.	Cross-Validation Study. Annealing solution case 1, with well 7 left out.	82
Figure B.7.	Cross-Validation Study. Annealing solution case 2, with well 7 left out.	83
Figure B.8.	Cross-Validation Study. Annealing solution case 3, with well 7 left out.	83
Figure B.9.	Cross-Validation Study. Annealing solution case 4, with well 7 left out.	84
Figure B.10.	Cross-Validation Study. Annealing solution case 5, with well 7 left out.	84
Figure B.11.	Cross-Validation Study. Annealing solution case 1, with well 8 left out.	85
Figure B.12.	Cross-Validation Study. Annealing solution case 2, with well 8 left out.	85

Figure B.13.	Cross-Validation Study. Annealing solution case 3, with well 8 left out.	86
Figure B.14.	Cross-Validation Study. Annealing solution case 4, with well 8 left out.	86
Figure B.15.	Cross-Validation Study. Annealing solution case 5, with well 8 left out.	87
Figure B.16.	Cross-Validation Study. Annealing solution case 1, with well 10 left out.	87
Figure B.17.	Cross-Validation Study. Annealing solution case 2, with well 10 left out.	88
Figure B.18.	Cross-Validation Study. Annealing solution case 3, with well 10 left out.	88
Figure B.19.	Cross-Validation Study. Annealing solution case 4, with well 10 left out.	89
Figure B.20.	Cross-Validation Study. Annealing solution case 5, with well 10 left out.	89
Figure B.21.	Cross-Validation Study. Annealing solution case 1, with well 11 left out.	90
Figure B.22.	Cross-Validation Study. Annealing solution case 2, with well 11 left out.	90
Figure B.23.	Cross-Validation Study. Annealing solution case 3, with well 11 left out.	91
Figure B.24.	Cross-Validation Study. Annealing solution case 4, with well 11 left out.	91
Figure B.25.	Cross-Validation Study. Annealing solution case 5, with well 11 left out.	92

List of Tables

	Page
Table 3.1. Annealing performance measures for templates 1 through 6.	24
Table 3.2. Annealing performance measures based on ten independent annealing solutions for each initial density.	26
Table 6.1. The observed steady state head values at each well and the predicted head values found using the median value for five annealing solutions. In each case the steady state head at the indicated well was left out of the energy function.	56
Table B.1. Prediction error in meters. Each well was left out of the energy function.	92

Abstract

One of the major problems in analyzing flow and transport in fractured rock is that the flow may be largely confined to a poorly connected network of fractures. In order to overcome some of this problem, LBL has been developing a new type of fracture hydrology model called an "equivalent discontinuum" model. In this model we represent the discontinuous nature of the problem through flow on a partially filled lattice. A key component in constructing an equivalent discontinuum model from this lattice is removing some of the conductive elements such that the system is partially connected in the same manner as the fracture network. This is done through an statistical inverse technique called "simulated annealing." The fracture network model is "annealed" by continually modifying a base model, or "template" such that the modified systems behave more and more like the observed system. This template is constructed using geological and geophysical data to identify regions which are possibly conducting fluid, and the probable orientations of channels which conduct fluid.

In order to see how the simulated annealing algorithm works, we have developed a series of synthetic "real" cases. In these cases, the "real" system is completely known so that the results of annealing to steady state data can be evaluated absolutely. We have studied the effect of the starting configuration by varying the percent of conducting elements in the initial configuration. Results have shown that the final configurations converge to about the same percentage of conducting elements. Further, we have tried various geometries for the template. It is not surprising that building a template that has conducting elements oriented similarly to the real system appears to give better results.

An example using Nagra field data from the Migration Experiment (MI) at Grimsel Rock Laboratory in Switzerland is also analyzed. The MI provided a unique opportunity to develop a model using geological information and hydrologic field data. We were able to construct a

model, identify possible solutions for the flow geometry given the hydrologic test results, assess the predictive performance of our solutions, and propose improvements for future use. The steady state results show that we can easily match the data and demonstrate the utility of multiple solutions to the inverse problem. We found using cross-validation that multiple solutions can be used to find a better estimator than a single solution and a realistic prediction error can be calculated. The transient response of a system is more sensitive to the distribution of permeability than the steady state response. Theoretically this makes transient annealing more attractive, but in the problems studied here we could not match the results in a low permeability zone. This indicates a model with uniform conductance and capacitance for all channels is not adequate for the MI site transient model, and so a variable permeability model is being developed.

1.0. Introduction

The disposal of nuclear waste is the subject of research internationally, and many countries are considering storage in underground facilities. The need to accurately predict the long term effects of storing waste underground has led to greater effort in developing and validating models of fluid flow and transport in fractured rock.

One of the major problems in analyzing flow and transport in fractured rock is that the flow may be largely confined to a poorly connected network of fractures. In these cases, the equivalent continuum models for flow and transport that were developed for porous media problems may not be reliable. Alternative modeling approaches have been explored for these cases. One approach is discrete fracture flow modeling where every fracture which carries flow is explicitly represented in a stochastic model (Billaux et al., 1989; Geier et al., 1990; Robinson, 1984; Long, 1983; and others). A problem with these models is that they focus on determining the hydrologic behavior from a statistical description of the fracture geometry. This can be very difficult because many fractures do not conduct fluid and because much of the flow may be carried by a few large features that are not well sampled.

In order to overcome some of these problems, LBL has been developing a new type of fracture hydrology model called an "Equivalent Discontinuum" model. In these models we represent the discontinuous nature of the problem through flow on a partially filled lattice. Essentially, we look at flow through the fractured rock as an equivalent percolation problem on a complex lattice.

The lattice we choose is designed to contain a set of conductors which are possibly important and is called a "template". The choice of template depends on a geologic evaluation of the medium. If there is no information about the structure of the rock, the template could be a regular lattice in space.

A key component in constructing an equivalent discontinuum model from this lattice is removing some of the conductive elements such that the system is partially connected in the same manner as the fracture network. This is done through an statistical inverse technique called "simulated annealing". The simulated annealing algorithm makes changes to the lattice and examines whether the change causes the model to behave more or less like insitu tests. Changes are accepted according to a stochastic process which results in a model that can reproduce hydrologic data observed in the field.

This report gives a detailed description of simulated annealing and some synthetic examples which provide insight about the best way to apply this technique. An example application of this technique has also been included. The example is based on data from the Migration Experiment (MI) at Grimsel Rock Laboratory in Switzerland. The MI experiment has been carried out by the Swiss Nationale Genossenschaft fur die Lagerung radioaktiver Abfalle (NAGRA) (Frick, Baertschi, and Hoehn, 1988)

The MI experiment is located in a simple, sub-vertical fracture zone which is intersected by several of the tunnels in the Grimsel Rock Laboratory in the Swiss Alps. From this tunnel, a series of boreholes were drilled into various parts of the fracture zone in order to carry out a series of hydraulic and tracer tests. MI provided a unique opportunity to develop a model using geological information and hydrologic field data. We were able to construct a model, identify possible solutions for the flow geometry given the hydrologic test results, assess the predictive performance of our solutions, and propose improvements for future use. The results of the analysis of the MI data are preliminary but offer encouragement for the utility of this new approach.

2.0. Simulated Annealing

“Simulated Annealing” is an inversion technique which can be used to construct a model which simulates observed behavior. This inversion technique can incorporate geological, geophysical, and hydrological data into one model. This method is particularly useful for fracture networks, since the system behavior is controlled by the geometry of the network. Annealing is used to find an equivalent fracture network model. The fracture network model is “annealed” by continually modifying a base model, or “template” such that the modified systems behave more and more like the observed system. This template is constructed using geological and geophysical data to identify regions which are possibly conducting fluid, and the probable orientations of channels which conduct fluid.

Hydrologic inversion models, such as the conjugate gradient method, or maximum likelihood method (Carrera and Neuman, 1986) were designed to determine the conductivity values in the equivalent continuum or porous medium. Annealing could theoretically be used to do this type of inversion, but would be relatively inefficient in this role. On the other hand, these equivalent continuum techniques work poorly when they are asked to completely turn off the conductivity of a portion of the region. Thus, they are not the technique of choice for poorly connected systems such as fracture systems when we wish to determine how the conductive features are connected.

One approach has been to use fracture mapping to determine the geometry which controls flow. However, at some sites, many of the fractures present do not appear to carry fluid. An interpretation of fracture trace maps from drift walls, fracture logs from boreholes, and single hole packer test data from Fanay-Augères mine in France illustrates the problem (Billaux, et al., 1989). Two drifts were mapped in the mine, one wet (S1) and one dry (S2). For both drifts the fracture geometry seemed to indicate highly connected fracture networks. If all the fractures

inferred from the mapping were hydrologically active the medium would have behaved like an equivalent porous medium. However, crosshole and tracer tests show that this was definitely not the case.

Further evidence of heterogeneous connection in fracture networks comes from a crosshole test program at the Crosshole Site, Stripa Mine, Sweden. Hydraulic tests results could not be explained using single fissure, regularly fissured, or porous medium models (Black, et al., 1987). Investigators at the Stripa Mine later reported that although the rock is ubiquitously fractured; 94% of the hydraulic transmissivity is found in only 4% of the tested rock (Olsson, et al., 1988a). Similar conditions exist at many sites, including the Grimsel Rock Laboratory. These examples demonstrate that the pattern of conductors is often responsible for the first order hydrologic behavior of fracture systems. At these sites, we need to identify the hydraulically active system of connected fractures.

In simulated annealing, we set up a "template" of allowed conducting elements. Then we look at different configurations of these elements by turning some of them off, i.e. making them non-conducting. For each configuration we can compute the behavior of a well test that was also conducted in the field. The "energy" of the configuration is then defined as a function of the difference between the observed and the simulated response. The problem of finding the appropriate model now becomes one of finding configurations which have low values of the energy function. Searching for a low energy configuration is a difficult task because there are many possible configurations.

This search is analogous to the problem of a hiker who has been dropped into a very hilly region, and he or she is expected to find the lowest point. The hiker can tell how far up or down a proposed step will take him, but he can not see farther than this next step. In addition, the hiker has a very short memory, and a bad sense of direction. He can only remember the last position and altitude.

The hiker may begin searching with a simple strategy. He will only take steps which lower his altitude. When a point is reached with no downhill step available, the hiker will stop.

Unfortunately, it is unlikely that this is the lowest point in the region, since there are many hills and valleys. We can say the hiker has found a local minimum, but probably not a global minimum. If the hiker wants a better chance of finding a global minimum, the search strategy will have to be modified.

The hiker may decide that a good search strategy would always allow him to take a proposed downhill step, however he would also sometimes take an uphill step. This would allow him to jump out of local minima, and continue looking for a lower spot. The searcher might choose to take or reject the uphill step randomly. Further, he might decide to base the probability of taking an uphill step on how far up it would take him. A slightly uphill step would be more likely than walking up a cliff.

Simulated annealing uses this type of search strategy to find solutions to similar optimization problems. The algorithm starts from some arbitrarily selected configuration and computes the energy, which is proportional to the difference between observed and measured values. Then an alternative configuration is selected, and the energy for this configuration is computed. If the alternative energy is lower than the energy for the current configuration, the alternative matches the observed data better, and the algorithm will decide to move to the alternative configuration. This is analogous to a downhill step. An "uphill step" to an alternative with a higher energy function will be taken randomly, with a probability which depends on the magnitude of the increase in energy and on a weighting parameter called the temperature. The temperature, T , is decreased as the number of iterations increases to make it more and more unlikely that an unfavorable change will be accepted.

The algorithm and the concept of configurations, energy functions, and a controlling temperature is based on the physical process of metal annealing. Annealing a metal can be the process of slow cooling that allows the metal molecules to form a regular internal pattern. A metal with particles packed in a precise internal pattern, which is repeated in all directions, is called a perfect crystal. A perfect crystal is the minimum energy structure for a metal. At high temperatures, the molecules will move about freely, in and out of different configurations. As the metal

cools, the molecules move about sluggishly, until they fall into a fixed configuration. If the metal cooled too fast, the molecules may fall into a different, higher energy, configuration than the pattern required to form a crystal.

Metropolis and others (1953) used an algorithm to simulate changes in a system of interacting molecules at a fixed temperature T . The simulation was based on a probability distribution for the range of energies called the Boltzmann distribution,

$$P(Q(C)) \propto e^{-Q(C)/bT} \quad (2.1)$$

where b is the Boltzmann constant, and $Q(C)$ is the energy of a configuration of atoms. Thermodynamically, low energy states are more likely, but at any temperature, there is still some chance of being in a high energy state.

Metropolis used an algorithm to simulate the changes a system of molecules could make from configuration to configuration. Starting from some random configuration, the system was assumed to have the option to change from C_1 to C_2 , with probability, P :

$$P = \begin{cases} 1 & \text{if } Q(C_1) - Q(C_2) > 0 \\ e^{-\frac{[Q(C_2) - Q(C_1)]}{bT}} & \text{if } Q(C_1) - Q(C_2) < 0 \end{cases} \quad (2.2)$$

Thus, the system would always move to a new configuration if it was of lower energy, and would sometimes move to a new configuration which was of higher energy. A high temperature would allow the system to jump from low to high energy states often, making it easy for the system to escape from local minima in the energy function. However, high energy states would be almost as likely as low energy states. A low temperature would make low energy states likelier but would increase the average time needed to reach one from an initial high energy state.

The simulated annealing algorithm is a generalized version of the Metropolis algorithm (Kilpatrick, et al., 1983; Tarantola, 1987). The temperature is held fixed for a certain number of configuration changes and then lowered. At first, a high value of T allows the algorithm to jump out of local minima and continue searching for a better region of the function. Later, lowering the

temperature tends to confine the search for a minima, so the algorithm can converge.

Simulated annealing also uses an energy function which defines certain characteristics one wants the final configuration to have. For example, one could use an energy function which defines low energy configurations as those which match an observed response. Then, simulated annealing can be used to search for these configurations. The minimum energy state for molecules in a metal is known, but the minimum energy states for the simulated annealing algorithm are unknown configurations we need to find.

2.1. Annealing Theory

To use simulated annealing on a general problem, one needs a set of possible configurations, a way of randomly changing the configurations, a function one would like to minimize, and an annealing schedule of temperature changes (Press, et al., 1986). Let

Q = an energy function

C = a configuration of elements

M = the finite set of all possible elements, ordered from 1 to M .

We can define the set of all possible configurations using our template or base model; M , the set of all possible pipes or channels. The channels have two possible states: they are either on or off, i.e. conducting or nonconducting. The set of all possible configurations is the set of all combinations of on and off pipes. Let $C = \{C_m, m = 1 \dots M\}$ denote a configuration of on and off pipes, where C_m is a binary random variable associated with each pipe.

We now must decide how to change the system. We choose to try removing or replacing one randomly chosen conductor at a time. Consider some configuration C . We will use some probability function to randomly select a pipe. If the pipe is on, we turn it off, and if the pipe is off, we turn it on. We can define the neighborhood of C to be all configurations one step away from C with one pipe missing or one pipe added. Thus any new configuration, C^i , will vary only slightly from C , that is C^i will be in the neighborhood of C . Let

$\{C^i\}$ = the configuration formed by selecting pipe i , and removing the pipe if it is on, or adding the pipe if it is off.

Let C be the configuration at iteration n and G_C be its neighborhood, $G_C = \bigcup_{i=1}^M C^i$. When we anneal the system, we randomly select a configuration C^i from G_C at each iteration n , and compare the two energy functions $Q(C)$ and $Q(C^i)$.

The energy functions we use are a measure of the difference between the observed and the simulated system response. We consider energy functions of the form:

$$Q = \sum_j (o_j - s_j)^n \quad (2.3)$$

where

$n = 2$ for hydrologic measurements, and

o_j = a vector of observed responses, and

s_j = a vector of simulated responses.

The observed measurements could be hydrologic, geological, or geophysical. Also we can use transient hydrologic measurements. For example:

$$Q = \sum_t \sum_j (h_{oj}(t) - h_{sj}(t))^2$$

where $h_{oj}(t)$ is the observed head response at well j and time t , and $h_{sj}(t)$ is the simulated head response at well j and time t .

The energy function, scaled by the temperature, is used to decide whether the system should make a transition to a new configuration. The temperature is lowered as the algorithm progresses, to make it increasingly unlikely that a transition to a higher energy state will occur.

If we let the algorithm run at a fixed temperature, we are sampling configurations using a Gibb's distribution, a generalization of the Boltzmann distribution:

$$P(C) = \frac{1}{k} e^{-\left[\frac{Q(C)}{T} \right]}$$

Thus the likelihood of occupying a configuration at any iteration is related to the energy of the configuration.

The normalizing constant, k , assures that the sum of the probabilities of all possible configurations is unity. We know that this constant exists, but it is very difficult to evaluate because we must know the energy for every possible configuration to compute k . So, we can not compute the absolute probability of any given configuration because we do not know k . However, we can compute the relative probability of any given configurations. For instance, we could say that a configuration would be twice as probable as another, given our conceptual model.

Further, we know that if the probability function is a Gibbs distribution then this is equivalent to modeling C , the current configuration as a Markov Random Field (Geman and Geman, 1984). A Markov Random Field exists if the probability defined meets two conditions. The first condition is that the probability of selecting any configuration in the system is greater than zero. The second is that the probability of making a transition from C to any other configuration C' given we are at C , depends on C , C' and whether C' is in the neighborhood of C . Past history, such as the configuration we selected before C , does not tell us anything about the probability of moving from C to C' . So, the probability of moving from one configuration to another can change with the iteration, but does not depend on which configurations have been examined in the past. This means we can examine a series of configurations without remembering how we moved from one to the next and we can still compute the relative probability of each configuration.

At each iteration k , given C , G_C , the neighborhood, and T , the temperature, we can find a matrix of transition probabilities. The probability we will move from configuration C to C' , given our current configuration C , ($P(C \rightarrow C'|C)$), is equal to the probability that we select C' to compare with C , ($P(C'|C)$), multiplied by the probability that the system would make the transition to a given C' . That is:

$$P\{C \rightarrow C' | C\} = \begin{cases} 0 & \text{if } C' \notin G_C \\ P(C' | C) \cdot 1 & \text{if } C' \in G_C, C' \neq C \\ & Q(C') - Q(C) \leq 0 \\ P(C' | C) \cdot e^{-\left[\frac{Q(C') - Q(C)}{T}\right]} & \text{if } C' \in G_C, C' \neq C \\ & Q(C') - Q(C) > 0 \end{cases} \quad (2.4)$$

and the probability of not accepting the change to C' is:

$$P(C \not\rightarrow C' | C) = P(C \rightarrow C) \\ = 1 - \sum_{\{C': Q(C') \leq Q(C)\}} P(C' | C) - \sum_{\{C': Q(C') > Q(C)\}} P(C' | C) \cdot e^{-\left[\frac{Q(C') - Q(C)}{T}\right]} \quad (2.5)$$

The temperature schedule is used to lower the temperature, or scaling parameter, as annealing progresses. This means that as the annealing progresses we are less and less likely to keep changes which increase the energy of the system. The length of the temperature schedule controls the number of iterations the algorithm runs. If the number of times we have examined new configurations and decided to accept them is equal to a predetermined number, then we lower the temperature. This new step of the algorithm continues until the same number of changes have been made. Then, we lower the temperature again. This continues until we reach the end of the temperature schedule. When the end of the annealing run is reached, we expect to have reached a low energy state. If this end configuration gives simulated well test results which are within the measurement error of the observed results, we say that the algorithm has converged to an acceptable solution.

There is a theorem which relates the temperature schedule to the convergence properties of annealing. This theorem (Hajek, 1988) shows that annealing done with a temperature schedule which is inversely proportional to the log of the iteration number will converge in probability to a set of minimum energy states. This means that if you follow this temperature schedule, and allow the algorithm to run for a very long time at the last step, the probability that you will be in a certain set of low energy states will approach one. Unfortunately, using a temperature inversely proportional to the log of the iteration number requires sampling a very large number of

configurations. Also, we are not really interested in this form of convergence. We are searching for several fairly different good solutions. Hajek's temperature schedule is over constraining for our purpose.

The temperature schedule we use here is only justified heuristically: it finds low energy solutions. We have followed the suggestion of Press, et al. (1979) and decreased the temperature whenever some fixed number n of changes have been accepted at the current temperature. Each interval of the schedule with constant temperature is called a step. At the end of each iteration, k , the temperature, T_k , is decreased using a geometric series,

$$T_{k+1} = T_k u^k \quad (2.6)$$

where u is a parameter chosen arbitrarily,

$$0 < u < 1 .$$

The initial temperature is chosen such that it is of the same order of magnitude as the energy difference between the first two configurations. This is done in an attempt to scale the energy difference between successive configurations between zero and one. Other choices of temperature schedule are possible and these are currently a topic of research (Dougherty, et al. 1990).

2.2. A Synthetic Example

In order to see how the simulated annealing algorithm works, we have developed a series of synthetic "real" cases. In these cases, the "real" system is completely known so that the results of annealing to steady state data can be evaluated absolutely. These cases are given in Section 3 and an example is included here.

A synthetic case was generated using the fracture network generator FMG, (Long, et al., 1982; Long, 1983). FMG is used to produce random realizations of a population of one-dimensional fractures in a two-dimensional square region called the generation region. A dimensionless network with two fracture sets was generated on a 100×100 grid (Figure 2.1).

On this network, we model a hydraulic interference test by creating a constant flux internal

boundary at a centrally located well. The program TRINET (Karasaki, 1989) is then used to calculate the head response at a series of observations wells. These heads become the "real" data that we try to match with annealing.

A template for annealing was developed using a grid with orientations close to those of the two fracture sets in the synthetic case. Figure 2.2 shows the template. The annealing algorithm found a minimum energy solution which appears by eye to match the flow geometry well. Figure 2.3 shows the minimum energy solution and Figure 2.4 shows the energy versus the iteration number for the annealing run. For this illustration dimensionless energy is defined by

$$E = \sum_i \left[\frac{h_{o_i}}{l} - \frac{h_{s_i}}{l} \right]^2$$

where l is the length scale used in the simulations (e.g. 1 cm or 1 m).


A series of synthetic cases have been generated in order to learn how annealing works best (see Section 3). We have studied the effect of the starting configuration by varying the percent of conducting elements in the initial configuration. Results have shown that the final configurations converge to about the same percentage of conducting elements.

Further, we have tried various geometries for the template. It is not surprising that building a template that has conducting elements oriented similarly to the real system appears to give better results. These studies are ongoing and can also be used to indicate which schemes for choosing configurations to test are best.

2.3. Model Validation and Measures of Uncertainty


2.3.1. Annealing and Model Uncertainty

Once a model has been built we will wish to use the model to make predictions. At this point we will want to know how good the model predictions are. There is a school of thought that maintains we should attempt to "validate" models in order to have confidence in the predictions we make. Model validation has been taken to mean the establishment of the soundness of the models and the legitimacy of specific applications. What actually constitutes validation is not a trivial matter. We can prove that the model is wrong but we can never prove it is right. A given



XBL 905-6392

Figure 2.1. The synthetic case used to generate well test data for use in annealing. Dots represent points where "well" data were generated, the central being the pumping well. Scale can be considered dimensionless.



XBL 905-6393

Figure 2.2. An example template developed for annealing the synthetic well test data.




Figure 2.3. A configuration resulting from annealing the synthetic well test data.




Figure 2.4. The dimensionless energy versus iteration curve for the synthetic annealing case.

model can be "wrong" for two reasons. There may be errors caused by incomplete and unreliable data, and there may be errors in the numerical codes and in basic modeling assumptions i.e., the conceptual model.

For the earth systems we model, the data available is insufficient to completely characterize the system. As an alternative to obtaining one deterministic model with large error based on incomplete and unreliable data, one could view the data as determining a probability distribution on the set of possible models. We may be able to find many models in this distribution which fit the data. The more incomplete and unreliable the data, the larger the possible solution set. We can quantify some of this uncertainty by considering a range of solutions.

There are several ways to choose different solution configurations of conductors. The simplest is to use a series of configurations defined at the end of an annealing process. These configurations are easily available, but they will probably be very similar to each other. Another way to find different configurations is to perform annealing several times, each time starting with a different initial configuration.

Errors caused by modeling assumptions are harder to define. We know that our model is a simplification, and only one of a very large number of possible conceptual models. We have to decide if the model is appropriate for our purposes. We think that the only way to approach this problem is through "peer review" or confidence building. In the peer review, the approach to modeling is scrutinized. What assumptions were made? Do they make sense? What is the evidence supporting these assumptions? What data are used? What is the sensitivity of the calculation to poor assumptions or inaccuracies in the data? Which are the parameters that control the result?

The main function of the model is to make predictions about the behavior of the system, so the model should be mainly judged by its ability to accurately predict the system response. This leads to the use of prediction error as a lump measure of error caused by incomplete data and model assumptions. The estimate of prediction error is made by using the model to make a series of predictions. For each prediction we obtain a prediction error by comparing the calculated

result to the measured result.

We are also concerned with about the robustness of our model: is it insensitive to small deviations from the assumptions? According to Huber (1981) slight deviations from the model assumptions should impair model performance only slightly, and larger deviations should not be catastrophic. Sensitivity analysis is one way to examine the robustness of the model.

The advantages of the quantitative description of uncertainty are many. First, and fundamentally, one is being more scientifically honest with a quantitative description of uncertainty. Second, an objective description of uncertainty gives one new insights into modeling problems. We can try to understand the relationship between our models and the kind of measurements that are taken in the field. This can lead to recommendations about what kind of data should be taken and where it should be taken.

The measures of uncertainty and model assessment we use are the results from sensitivity studies, and prediction error calculated using cross-validation. Both of these estimate the effect of incomplete and unreliable data. Sensitivity analysis measures uncertainty caused by basic assumptions in the model which are wrong, or changed. And the prediction error is a measure which lumps together all sources of error.

2.3.2. Prediction Error

One way science has advanced is through the development of theories or models. A theory or model is useful if it successfully predicts behavior. The pure truth of the theory or model is not always relevant, for instance, there are two parallel theories often used to predict the behavior of light. One theory holds that light is a wave and the other theory holds that light is a particle. Physicists have known for some time that neither theory is strictly true. However, both models are useful, since under different conditions they do predict the behavior of light.

In the same sense, our hydrologic model is not a true representation of the fracture flow system. We are justified in using the model if it can accurately predict behavior. One measure of model goodness is the prediction error. We define the prediction error as the error between some

independent quantity we predict using our model, and the measured value of that quantity.

It is important to recognize that a good prediction of a certain type of response, requires appropriate data to build the model. For example, if one built a model using pressure measurements, then you can expect the model to predict head measurements. However, this same model may not be as useful for predicting flow. If you want to predict flow, you need to use flow measurements to build the hydrologic model. Also, a prediction error estimated for one kind of response should not be used to estimate the prediction error for a different response.

The best possible way to evaluate prediction error is to make a prediction for a known quantity, that hasn't been used to build the model. Unfortunately, one usually needs all the available data to build a good model. One way around this problem is to set aside one data point, construct a model using the rest of the data, predict the value left out, and calculate a prediction error. If we do this for each data point in turn, we have a distribution of prediction errors we may use to estimate the prediction error for a model using all the data. This process is called "cross-validation". Cross-validation may be extended to calculate multiple solutions for each data point we set aside.

For the MI study, we have steady state head values at 8 wells. If we leave out one well at a time and anneal, we can calculate a range of prediction errors for pressure measurements. We can use this range of errors to estimate the prediction error for the full model.

So far the discussion has assumed that there is only one model and therefore only one possible prediction available. However, for an inverse problem one may have a range of good possible models. This can be an embarrassment of riches since making a decision, such as where to put the next well, requires a single prediction. We solve this problem by using a loss function and the prediction error to choose a predictor.

When several annealing solutions are available to predict each measured data point value used in a cross-validation study, we would expect that using a mean or median of the several possible predicted values would give a lower prediction error than using a value from one annealing solution. We might then want to decide if the mean or the median is the "best" choice of a

predictor. Choosing the “best” predictor requires some criteria defining the meaning of “best”. The standard statistical criterion for this type of decision is based on a function of the estimated prediction error, called a loss function. The loss function quantifies the loss we suffer when our predictions are off the true value. Often this loss is not a linear function of the error. One commonly used loss function is the squared error loss function,

$$L(y, \hat{y}) = (y - \hat{y})^2$$

The observed value is y , and the predicted value is \hat{y} . For hydrologic data, we believe the percent difference is a good measure of the loss incurred. Hydrologic data often covers a few orders of magnitude, and it is most important that predictions should be of the same order. By using this loss function, we normalize all the data points so they may be compared. Otherwise, a small prediction error for a large value would count as much as an order of magnitude error for a small value. Our loss function is:

$$L(y, \hat{y}) = \frac{|(y - \hat{y})|}{y}$$

Given a loss function, our criterion for choosing the “best” predictor is, “Choose the estimator which minimizes the sum of the loss functions in a cross-validation study.” For example, in the steady state MI case, we can leave one well out at a time, and then find several annealing solutions. We can then find the mean and median predicted values for each well left out. The loss function for each well left out can be computed for each estimator, and summed over the wells. The predictor, mean or median, with the smallest sum of loss functions is chosen as the “best” estimator. The average of this sum is a measure of the goodness of fit for the composite full model.

2.3.3. Sensitivity Analysis

The reliability of our model depends on certain parameters, such as boundary conditions which may change over time or be inaccurately measured. We have essentially three kinds of parameters: boundary conditions, the value of the conductance assigned to each conductor, and the geometry of the template. Given a range of possible models, we could change the boundary

to boundary conditions can be done by trying different scenarios such as a free surface boundary or adding various constant head boundaries.

As for the conductances, any prediction of flow rate will be directly proportional to the conductance assigned to the elements, so the sensitivity of steady flow rate to conductance is known *a priori* for constant conductance. On the other hand, for steady conditions the head distribution is independent of the conducting value. We might want to consider the sensitivity of the model to the assumption of constant conductance. Our experience has lead us to believe that our model would be improved by allowing conductance to vary.

The geometry of the template could effect the model since the path length between two points will depend on the orientation of the elements. We used a synthetic case to study the effect of this parameter as described in the next section. Other aspects of the template such as the level of discretization are important.

An investigation of the sensitivity of our modeling process to these parameters is ongoing. We expect to find that steady state pressure models are sensitive to boundary conditions, and transient pressure and flow models are sensitive to variable conductance. A preliminary study of the geometry of the template indicates the importance of geologic information, but more study is needed before a firm conclusion can be drawn.

3.0. Synthetic Case Studies

In order to see how the simulated annealing algorithm works, we have developed a series of synthetic cases. For a synthetic case, the system is completely known so that the results of annealing can be evaluated absolutely.

We designed an experiment to determine the effect of the template and the beginning configuration on the end configuration found by annealing with steady state head values. In an analysis of a real site, geological information can be used to set the orientation of the elements in the template. To see if this was important we studied cases where the elements were aligned with the prevailing fracture orientations in our synthetic case and cases where they were not. To see the effect of the starting configurations on the end configurations, we defined initial configurations with different percentages of elements turned on. These were annealed using the same temperature schedule and different seeds. The final annealed geometry, the density, the number of iterations, and the energy of the solutions were compared.

3.1. The Synthetic Case

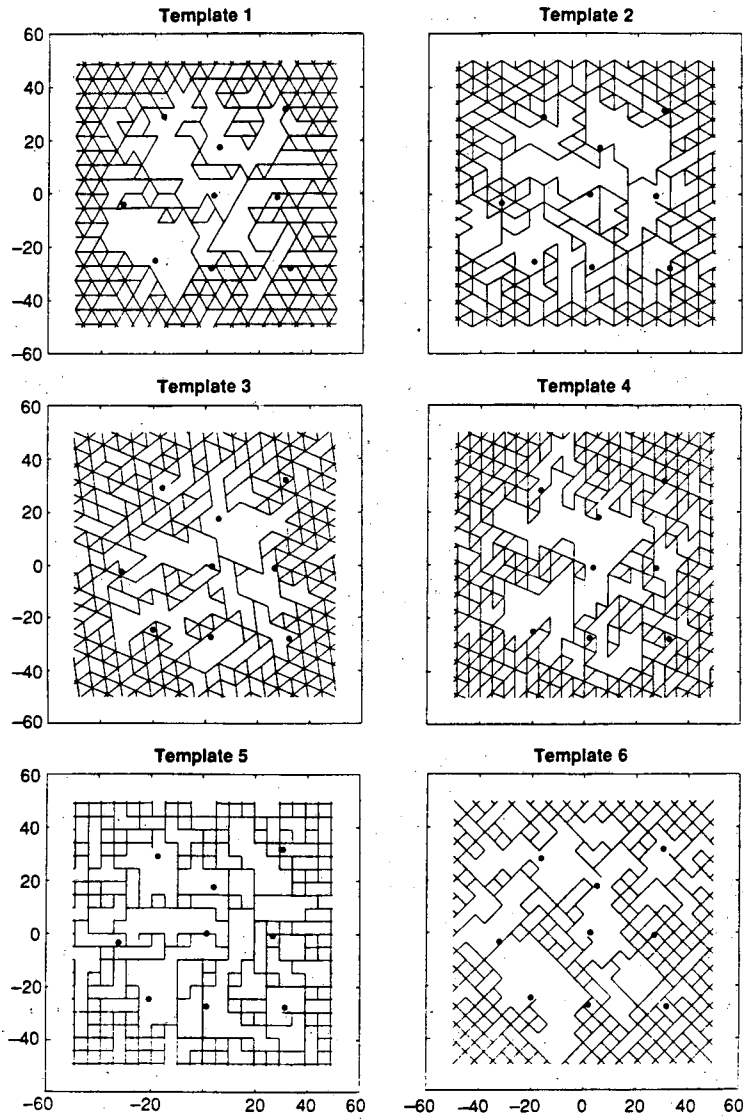
The synthetic case was generated using the fracture network generator FMG, (Long, et al., 1982, Long, 1983). FMG produces random realizations of a population of one-dimensional fractures in a two-dimensional square region called the generation region. A dimensionless network with two fracture sets was generated on a 100×100 grid. Figure 2.1 shows the synthetic case. The first fracture set had a density per unit area of 0.013 fractures, and orientations distributed normally with a mean orientation of 45° measured counter-clockwise from horizontal and a standard deviation of 30° . The length and conductance of the fractures were constant values set at 25 and 0.001. The second fracture set had a density per unit area of 0.025 fractures, and orientations distributed normally with a mean of 90° and a standard deviation of 40° . The length and conductance of the fractures were constant values set at 2 and 0.001 respectively.

3.2. Effect of the Geometry of the Template

Information on the orientation of the fracture channels can often be inferred from geologic investigations. We wanted to see if *a priori* information on the orientation of conductors significantly improved the results. It seems obvious that the flow geometry found by annealing should be more like the "true" flow geometry if the elements had orientations close to those of the real fracture channels. Also, a template with orientations very different from the true channels might slow down convergence of the algorithm, and the energy might tend to stay higher.

We did a preliminary investigation of the effect of template orientation, using 6 different templates and the synthetic well test data case. The templates were 120×120 grids with fracture channels in each set evenly spaced 2 units apart. Both hexagonal and square grids were used. Figure 3.1 shows the annealing solutions found using the 6 templates. The orientations of the fracture set, and the minimum energy and number of iterations until the end of the temperature schedule are given in Table 3.1. The temperature schedule fixes the number of changes made, which is the same for each annealing run. Since the temperature is lowered after every 50 changes, a lower number of iterations indicates that the grid provided a greater chance of moving to a new configuration at each iteration. We would expect that one might find a good solution sooner under this circumstance. The minimum energies are treated as though they are equal if they fall under 0.001, since this is likely to be under the measurement error for this type of data.

In cases 2 through 6 there is no significant difference in the minimum energy or the number of iterations. A visual examination shows that template four, which had channels inclined along the mean directions of the "real" system, gives the best match to the flow geometry. The differences are not dramatic and several of the other templates also give a good match. For example, template 6 looks much like the synthetic case. However, it seems that finding a flow geometry close to that of the real system is harder if the channel orientations are very different from those of the real system as in case 1. Cases 4 and 6 also had a low number of iterations.



XBL 905-6395

Figure 3.1. Annealing solutions to the synthetic case shown in Figure 2.1 for different template geometries (1 through 6).

Table 3.1. Annealing performance measures for templates 1 through 6.

Templates	1	2	3	4	5	6
Inclination set 1	0	30	37.5	45	0	45
Inclination set 2	60	90	97.5	90	90	135
Inclination set 3	120	150	157.5	157.3	-	-
Number of iterations until convergence	1370	1410	1441	1277	1315	1264
Minimum energy	4.0 e^{-03}	1.01 e^{-05}	17.6 e^{-05}	11 e^{-05}	5.7 e^{-05}	5.8 e^{-05}

The energies are all low enough to give negligible mean squared error. The energy is the sum of the squared difference between the “real” head values at the wells and the values of the wells in the annealing solution. The “real” steady state head values range from zero to one. We consider any head difference less than 0.01 to be effectively zero, and all the energies are below 0.005. The average head difference for each well is therefore under 0.01.

Based on this limited sample we believe that the geologic information incorporated into the template for the MI site will improve the solution.

3.3. Effect of the Starting Point

Many different configurations of channels can equally well match the hydrologic data available at a site. We are interested in obtaining a range of flow geometries. However, we expect that the flow geometries should have approximately the same density of channels in order to have the same connectivity.

One might believe that the set of configurations which match the hydrological data can be grouped or categorized. For example, some solutions may tend to have a “hole” in a certain location and others in a different location. These groups may each be associated with a different valley in the energy function. One might be able to reach a certain valley starting from some initial points but not others. One way to find flow geometries in these different valleys might be to start from widely separated configurations and different random seeds. The beginning configurations could also have different percentages of “on” pipes. If we can start from points

with very different percentages, and arrive at different flow geometries with approximately the same density, then this method may give us a good range of solutions.

We have studied the effect of the starting point by running a case with 60% of the pipes turned on in the initial configuration. The end configuration had almost the same density, so we wished to determine if the initial density had a big influence on the density of the solution. We hoped the initial density had no effect, since the density of the solution should depend on the hydrologic behavior and approximate the connectivity of the real system.

A study of the effect of the starting point was designed. We examined configurations with 60, 70, 80, 90, and 100% of the channels in the template initially turned on. Ten different starting arrangements for each density were randomly selected. Each starting configuration was annealed using the same temperature schedule, but starting from a different random seed.

Initially we tried randomly selecting "on" pipes with a uniform distribution. This did not work well, since the spherical distribution we use to randomly change elements during the annealing process tends to examine pipes in the middle of the grid much more often. Pipes on the boundary are not touched during the annealing process. This means that final configurations found starting from high density grids have more pipes along the boundary than final configurations found starting from low density grids. We then decided to remove pipes from the template to find starting configurations using the same spherical distribution. The final configurations found in this manner had slight differences in density across the starting percentage of elements. We believe these can still be attributed to boundary effects. Table 3.2 shows the initial densities and the median density of the ten solutions for each case.

The flow geometries of the solutions for each percentage were compared. Using one starting point we can find a set of solutions almost as diverse as the set found using multiple starting points. However, this case is relatively small, in terms of the number of elements. Multiple starting points may be more advantageous on problems with larger configurations sets.

No significant difference was seen in the number of iterations to convergence or the minimum energy found along the annealing path.

Table 3.2. Annealing performance measures based on ten independent annealing solutions for each initial density.

Initial Density	Median Minimum Energy	Median Number of Iterations	Median Density
60%	6.3 e^{-06}	2943	61
70%	2.2 e^{-06}	2959	60
80%	4.7 e^{-06}	2985	62
90%	7.9 e^{-06}	2945	63
100%	4.65 e^{-06}	2928	62

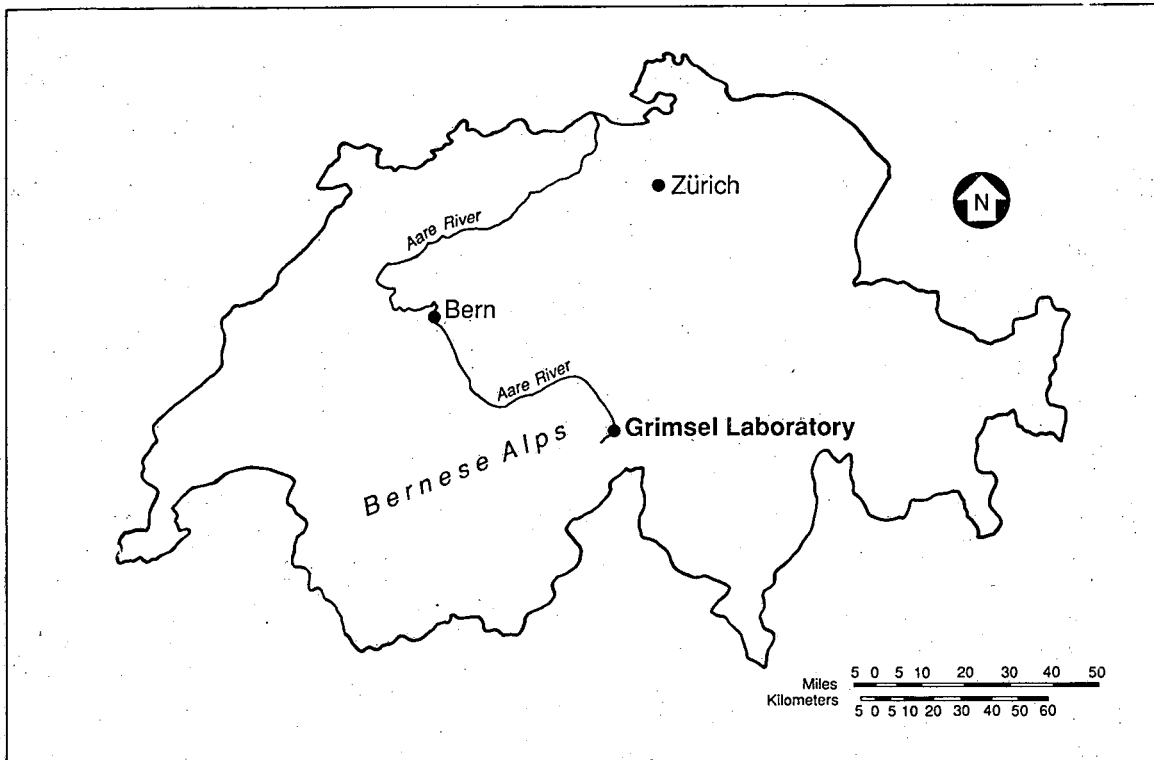
4.0. The Migration Experiment Site

An application of simulated annealing was made to data from the MI experiment at the Grimsel Rock Laboratory in Switzerland. The geology of this site is explained here. This geological data is used to identify and characterize possible hydrologically active structures and assess the solutions found by the inversion. Most of the data discussed was provided by the Swiss National Cooperative for the Storage of Radioactive Waste (Nagra).

Nagra has hosted a variety of experiments in the past few years at its underground Grimsel Rock Laboratory directed towards improving understanding of fracture flow. This facility is located inside a mountain (the Juchlistock) in the Bernese Alps near the headwaters of the Aare River (Figure 4.1). The laboratory is at an elevation of approximately 1730 m, a few hundred meters below the surface. Figure 4.2 shows the drifts and large boreholes of the Grimsel Rock Laboratory. Lawrence Berkeley Laboratory (LBL) has worked with Nagra for five years to help develop site characterization techniques at the Grimsel laboratory.

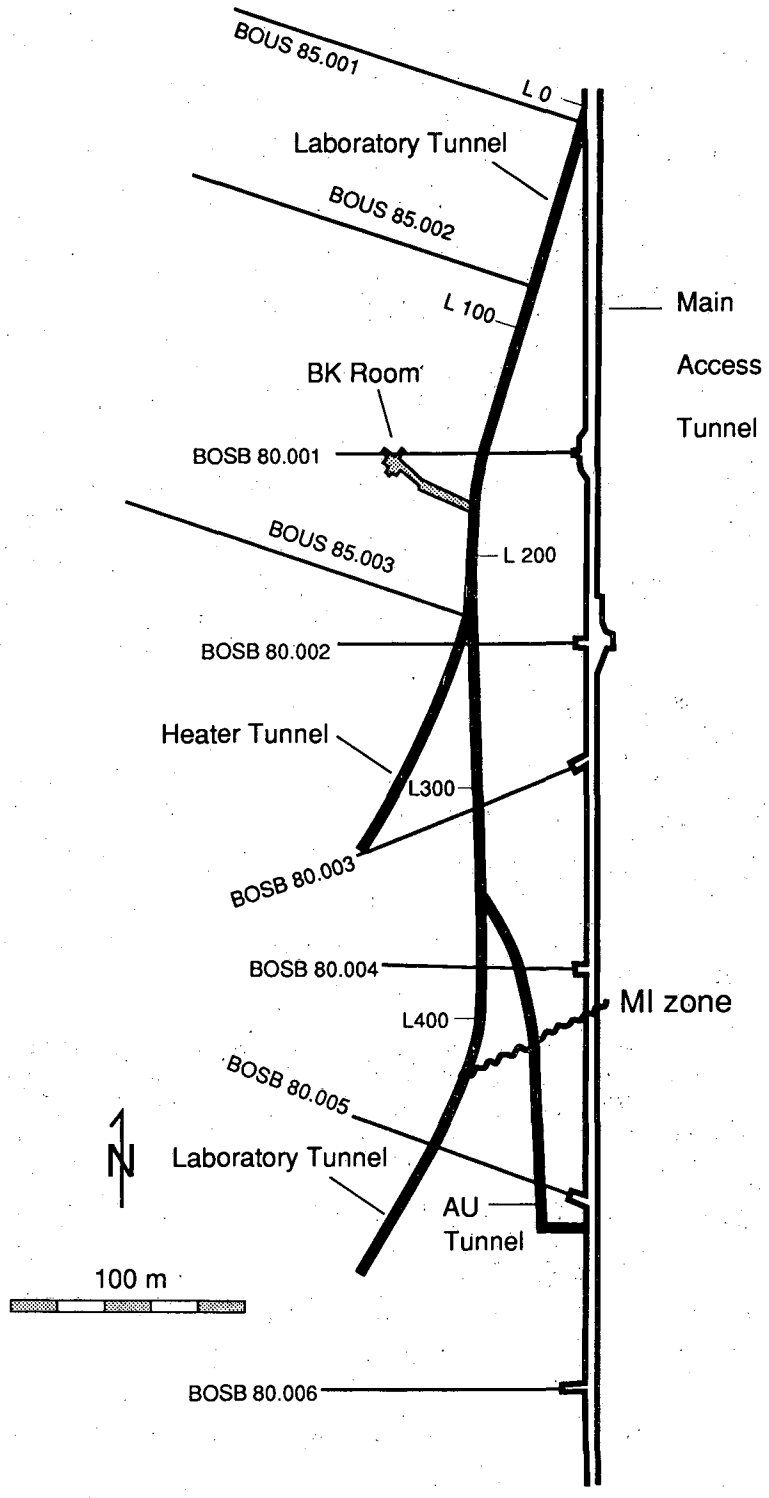
Three main types of steeply-dipping, fracture-bearing structures transect the Juchlistock: S-zones, K-zones, and lamprophyres (Figures 4.3 and 4.4). These form the principal hydrologic features in the subsurface. The S-zones are shear zones that generally dip steeply to the southeast, parallel to the foliation in the host rock. Both the fractures and the grain-scale mineral fabric of the S-zones dip steeply to the southeast. The K-zones are fracture zones that generally strike west or northwest; they cut the host rock fabric at a high angle. The lamprophyres are mafic igneous dikes. These have been metamorphosed and contain abundant biotite. Like the K-zones, the lamprophyres generally strike west or northwest.

In the Grimsel facility, a study of transport phenomena is being conducted at an S-zone in the southern part of the Grimsel laboratory. Figure 4.5 shows the MI zone in the geological setting of the southern part of the Grimsel Rock Laboratory. This zone was chosen for study because its structure is relatively simple. The MI zone intersects the AU tunnel between



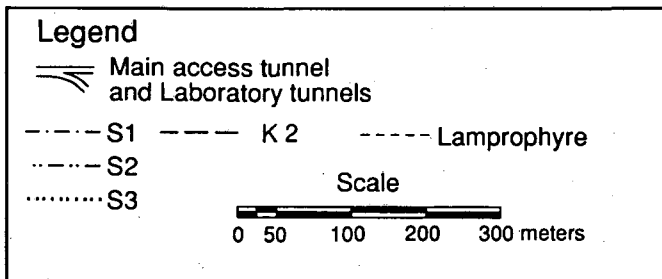

XBL 899-6312

Figure 4.1. Map of Switzerland showing the location of the Grimsel laboratory.



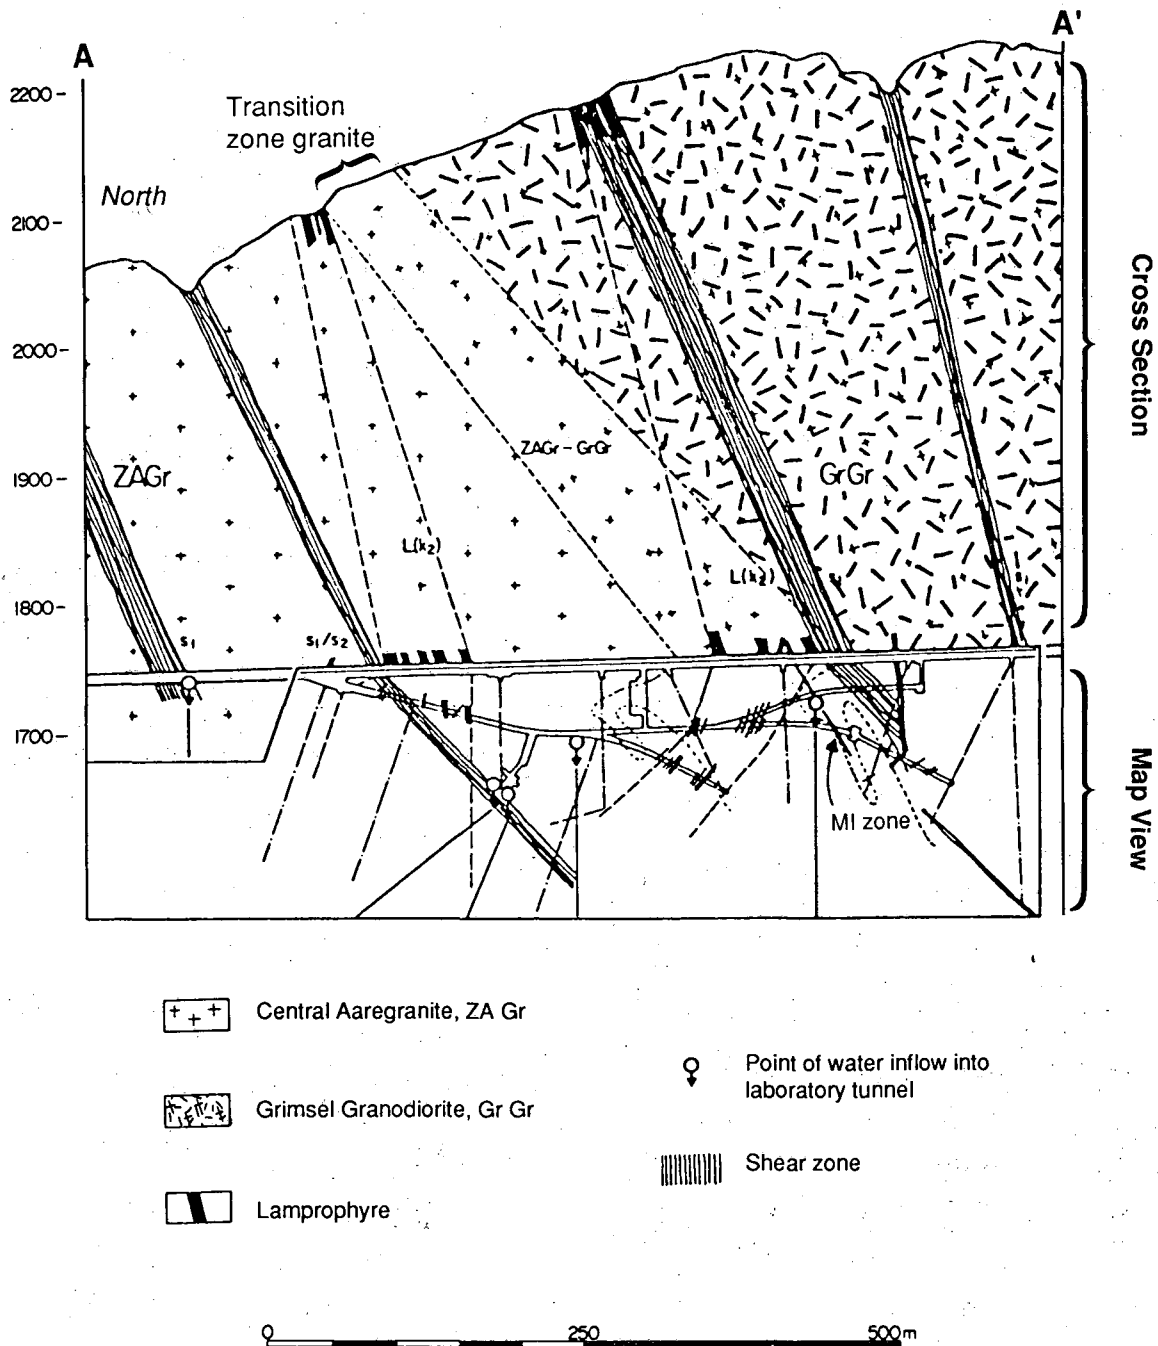
XBL 905-1790

Figure 4.2. The layout of the Grimsel Rock Laboratory showing the location of the MI fracture zone.



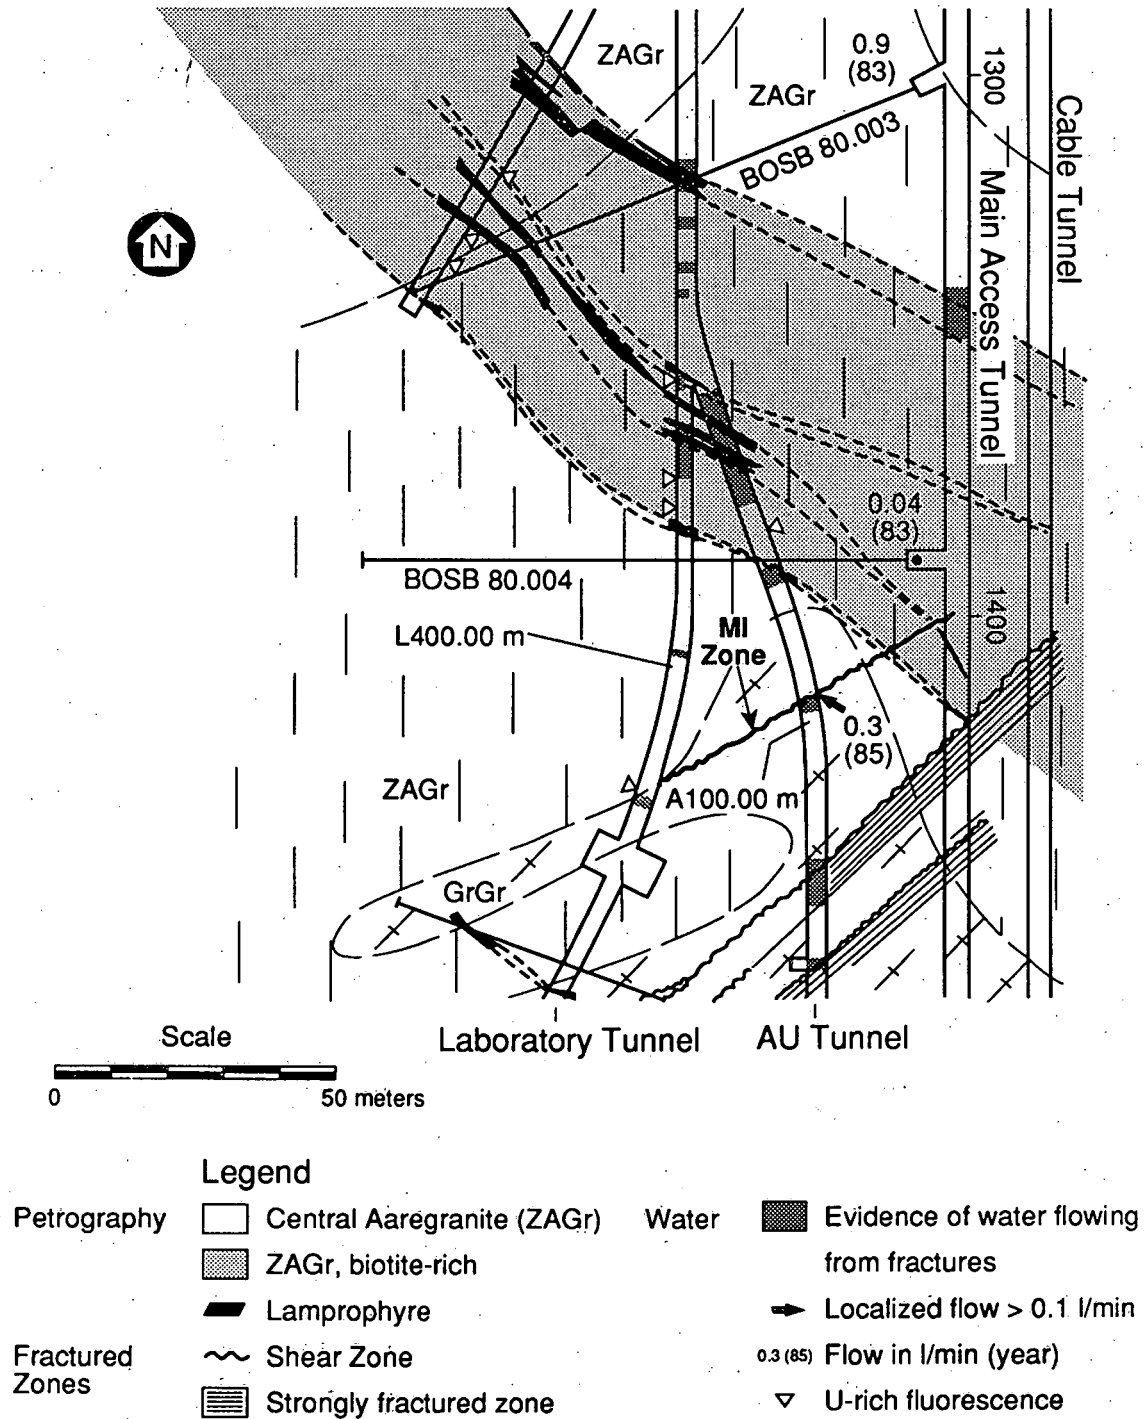
XBL 898-7692

Figure 4.3. Geologic map showing the major structures at the surface above the Grimsel laboratory. The MI zone is part of a group of fractured shear zones (shown by heavy lines) that extend southwest from Lake Räterichsboden. From Nagra Technical Report 87-14.



XBL 8911-4261

Figure 4.4: Geologic cross section and generalized map showing major geologic structures in and near the MI site extending to the surface. The line of cross section is shown on Figure 4.3. From Nagra Technical Report 87-14.



XBL 898-7693

Figure 4.5. Geologic setting of the southern part of the Grimsel Rock Laboratory.

reference marks A95 and A98.5 and the laboratory tunnel between reference marks L419 and L425.5. Figure 4.6 shows the traces of the fractures in the MI zone on the walls of these two tunnels. Eight boreholes have also been drilled through this zone from the AU tunnel for the hydrologic investigations, (Figure 4.7).


4.1. Geology of the Southern Part of the Grimsel Rock Laboratory

The southern part of the Grimsel laboratory (Figure 4.5) intersects three granitic rock types: the dark colored Grimsel Granodiorite, the light-colored Central Aaregranite, and a transition zone of biotite-rich Central Aaregranite between the two main units. The contacts between these units have very irregular shapes. The exposures of the MI shear zone in the AU and laboratory tunnels have been mapped near the edge of a finger of Grimsel Granodiorite that is surrounded by transition zone granite.

The three bodies of granitic rock are strongly foliated, the foliation being defined by aligned grains of biotite and bands of mylonite. The foliation strikes northeast and generally dips about 65° to the southeast. Observations in the laboratory tunnel (P. Bossart, oral communication, 1989) and the AU tunnel indicate that the granitic rocks also contain a linear fabric element. This lineation is defined by elongated feldspar grains that are aligned approximately parallel to the dip of the foliation.

The most prominent shear zones in the southern part of the Grimsel laboratory strike NE, parallel to the foliation in the rock. Three shear zones compose this set; the MI zone is the northernmost. These shear zones appear to be part of a structure that is prominently expressed at the surface (Figures 4.3, 4.4). The second set contains a single zone that strikes nearly east-west.

The other major structures exposed in the southern part of the laboratory are lamprophyres. The lamprophyres dip steeply and generally strike northwest. Many gently-dipping Alpine tension fissures, *Zerrklüfte*, extend from lamprophyres in the southern part of the laboratory. The apertures of these fissures locally exceeds a meter, but more typically are several centimeters. Quartz and chlorite fill these fissures.



XBL 905-1795

Figure 4.6. Logs of fractures in the MI fracture zone in (a) the AU tunnel and (b) the laboratory tunnel. From preliminary draft of Nagra Technical Report 87-14 by Keusen and others.




Figure 4.7. The AU tunnel and MI fracture zone intersected by 8 boreholes. From Nagra Report 1988.

4.2. Structure of the MI Fracture Zone


The MI zone, like other shear zones at Grimsel, contains fractures that are subparallel to the zone. On the east side of the AU tunnel (Figure 4.6a) the zone contains three prominent NE-striking fractures; and they define a zone about 2 m thick. Two of these fractures merge into one at the A96 mark on the west side of the tunnel. In contrast, the MI zone exposure in the laboratory tunnel contains many NE-striking fractures and is about 3 m thick (Figure 4.6b). The zone apparently incorporates many more fractures and thickens as it approaches the laboratory tunnel from the AU tunnel. Measured fracture strikes at the exposures of the MI zone range from N58°E to N80°E; these correspond to dip directions of 148° and 179°, respectively. Measured fracture dips range from 70° to 85° to the southeast. These observations show that the MI zone contains a series of subparallel fractures that strike southeast. Fractures in other S-zones at Grimsel consistently form a braided pattern (Figure 4.8). This braided pattern appears to be more drawn out in the vertical direction than in plan view (Figure 4.9).

Frick and others (1988) have interpreted the MI boreholes as intersecting one nearly planar main fracture within the MI zone. A best-fit plane regressed to the eight borehole-fracture intersections strikes N54° and dips 78° SE (U. Frick, Nagra, 1989, written communication). Small fractures are inferred to branch from the main fracture (Figure 4.10). An alternative interpretation is that the boreholes intersect a braided series of subparallel fractures that is about 0.5 m thick. This alternative interpretation is perhaps more consistent with the appearance of the MI-zone fractures in the AU tunnel (Figure 4.6a) and avoids the problem of the best-fit plane lying as much as 0.4 m from some borehole/fracture intersections.

The NW margin of the MI zone contains abundant biotite at both the AU and laboratory tunnel exposures, and an Alpine tension fissure extends from the NW edge of the zone in the laboratory tunnel (Figure 4.6). A similar feature occurs in the AU tunnel at A148 near the biotite-rich SE edge of the shear zone. Alpine tension fissures might occur elsewhere along the MI zone.



Figure 4.8. Photograph showing braided structure of fractures in a NE-striking fracture zone exposed at the surface above the Grimsel laboratory. Fifteen-cm ruler for scale.



XBL 903-782

Figure 4.9. Diagram showing the three-dimensional braided fracture structure of the NE-striking fracture zones at Grimsel. The braided pattern appears more drawn out in vertical cross sections than in plan view.




Figure 4.10. Diagram of the MI "fracture" from Frick and others, 1988.

No major structures are likely to intersect the MI zone within a few meters of the AU tunnel. The nearest shear zone in the AU tunnel is 11m south of the MI zone, and it nearly parallels the MI zone. A prominent fracture and a fracture zone are exposed approximately 5 m north and south of the MI zone, respectively, in the laboratory tunnel. They are also nearly parallel to the MI zone. A SE-striking lamprophyre exposed in the AU tunnel 26 m north of the MI zone probably intersects the MI zone about 15 m NE of the AU tunnel. Some chlorite-bearing fractures that strike to the south are exposed a few meters north of the MI zone in the AU tunnel, and they may intersect the MI zone a few meters west of the AU tunnel. However, these fractures do not appear to be part of a prominent fracture zone. Quartz veins are intersected in boreholes BOMI 87.007 at 11 m and BOMI 87.009 at 9.5 m, but the significance of these veins is uncertain. The combined borehole and tunnel data suggest that the MI zone probably is not hydraulically well-connected to other major fracture zones in the immediate vicinity of the migration tests.

4.3. Hydrogeologic Implications of Fracture Structure at the MI Site

The anisotropic fracture structure of the NE-striking fracture zones at Grimsel indicates that the MI zone would be hydrologically anisotropic. The rock along the MI zone is strongly foliated and most macroscopic fractures in the zone probably are subparallel to the foliation. As a result, the average hydraulic conductivity along the zone probably is greater than that across it. The fractures along NE-striking shear zones at Grimsel consistently display a braided pattern, with the fractures having a more tortuous appearance in plan view than in cross section (Figure 4.9). As a result, the hydraulic conductivity is likely to be greater along dip than along strike.

Even though the borehole data suggests the portion of the MI zone within several meters of the AU tunnel may have a fairly uniform structure, the marked change in structure between the AU and laboratory tunnels indicates the hydrologic properties may be more variable along the zone than the structural data from the boreholes might suggest. Mylonite, cataclasite, and other fault-filling materials probably are not uniformly distributed along the zone, and accordingly permeability along the MI zone could also vary.

5.0. Hydraulic Tests at MI

The hydrological data is crucial to our modelling effort. This data is used to "anneal" the model, to validate the model building process, and to assess the predictive abilities of the solutions we find.


In this chapter we will describe the well tests conducted at the MI site. In Chapter 6 we show that these well test data are used in an example application of the annealing technique. In the example, we will anneal the hydraulic heads by forward-modeling well tests. Although only the hydraulic test data will be used in the example, any test data can be used as long as the test can be forward-modeled. In the following section we will describe the hydraulic tests conducted at MI. We will also explain how we chose the boundary conditions for the model.

5.1. Hydraulic Tests

At the MI site Solexperts (1988) conducted several constant pressure tests and a constant injection test in the Migration fracture zone while the pressure was measured at a number of observation points. A total of eight boreholes are drilled from the Laboratory Drift to intersect the Migration fracture zone (Figure 4.4). The intervals in the boreholes that intersect the fracture are isolated with packers. The packed-off intervals are numerous and closely located in a relatively simple geometry. This data presents a unique opportunity to apply and improve the simulated annealing technique.

The first group of tests was conducted in late September, 1987 (Solexperts, 1988). There were six constant pressure tests where a different interval was chosen as the pumping well among the eight packed-off intervals in each test. Figure 5.1 shows the pressure vs. time data observed in these intervals. As can be seen from the figure, the pressure in the pumping well was not entirely constant. This presents a problem in the forward modeling because the boundary

Hydrotests Migration 26/9-1/10/87 Constant head/flow interference tests



XBL 899-6296

Figure 5.1. Pressure data observed during six constant pressure tests (Solexperts, 1988).

condition at the pumping well is difficult to specify. In the first and the third test the pressure was relatively constant. However, both tests were rather short and as a result the volume covered by these tests is relatively small. For these reasons, these well test data were not used for annealing.

The second group of tests conducted from late December 1988, through mid-February, 1989, (Solexperts, 1989) consisted of a constant flow withdrawal test from well 9 followed by a buildup test (Figure 5.2). The flow period was maintained for more than 35 days and the buildup was monitored for about 20 days. The data from these tests are considered to contain more information on the Migration fracture zone because of the long test duration. The pressure data in the flow period was used in the annealing application. Unfortunately, the flow rate fluctuated in the first few minutes. A stable flow rate could not be maintained at the initial flow rate of 150 ml/min. Subsequently the rate was changed to 340 ml/min at 4 minutes into the test.

In the forward model for the transient case, the flow rate is assumed to have been constant at 340 ml/min from the onset of the test. To incorporate the discrepancy in the rate the starting time of the test is adjusted so that the cumulative flow is the same (Figure 5.3). Alternatively, the actual change in flow rate could be modeled. However, the flow rate in the first four minutes is very unstable and considered to be inaccurate. Furthermore, the difference is believed to be so small that it does not warrant the additional complication in the model. The steady state case was modelled using the steady state pressure data observed before pumping.

5.2. Boundary Conditions

Because we are using a finite size model, the boundary conditions along the sides of the model have to be prescribed. However, they are almost as difficult to estimate as the patterns of flow channels within the fracture. This is because the medium is not homogeneous. Faults and interconnected fractures cause the pressure field to deviate from that of a homogeneous medium. The presence of multiple tunnels further complicates the problem. If we choose to model a much larger block of rock, we could use the hydrostatic boundary condition by making the model boundary virtually unaffected. However, this is difficult because there is an upper limit in the size we can feasibly model. Also, the water table level is not known. Fortunately, our experience has

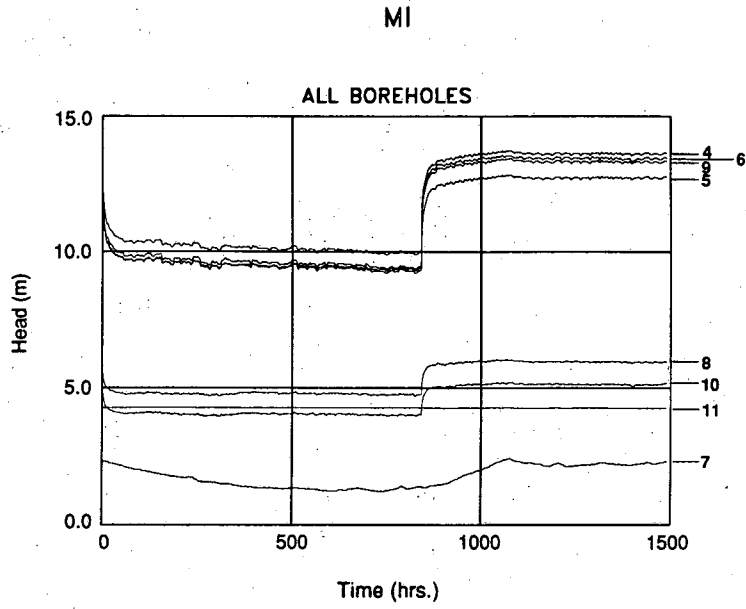


Figure 5.2. Pressure records from the constant flow withdrawal test followed by a buildup test (Solexperts, 1989).

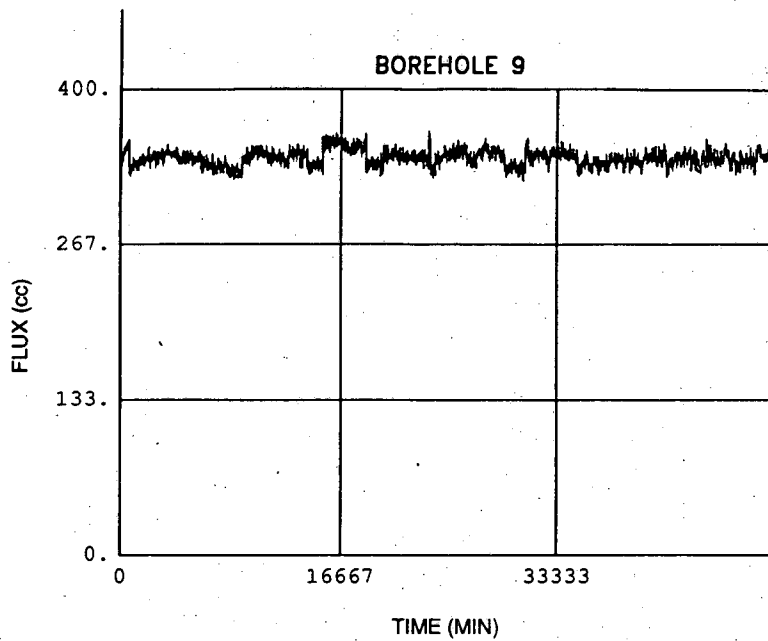


Figure 5.3. The adjusted flow rate in BOMI 87.009.

shown that the annealing procedure is virtually insensitive to the flow patterns outside of the area where observations are made, so it would be a waste to construct an unnecessarily large model (Long, et al., 1990). The size and boundary conditions of our model are chosen based on these factors.

There are three openings in the plane of the MI fracture, i.e., the access tunnel and two laboratory tunnels. These openings act as sinks with unknown strength. The undisturbed water table was probably more than 400m above the openings initially. However, the drainage to the tunnels may lower the water table depending on the recharge rate. The pressures measured in EB-1 through 6 before the laboratory tunnels were excavated indicate that the pressure field was unevenly disturbed (Figure 5.4) by the access tunnel (Nagra, 1985). This is presumably due to the heterogeneities such as faults and sheer zones in the mountain. Therefore, it is expected that the hydraulic head distribution is still very irregular at this time.

In spite of all the above, a crude estimate of the model outer boundary condition can be obtained using the method of images if the medium is assumed to be homogeneous and the water table is assumed to be flat and constant. The head at any point outside of the tunnels can then be expressed as:


$$h(x,z) = \sum_j \frac{Q_j}{2\pi T} \left[\ln r_{Sj} - \ln r_{Ij} \right] + Z_0 \quad (5.1)$$

where Q_j denotes the flow rate into the j -th tunnel and r_{Ij} and r_{Sj} are the distances from the image and the actual center location of the j -th tunnel, respectively. The unknowns are Q_j , T and the height of the water table above the origin, Z_0 . However, after the cases where Z_0 ranged from 50 m to 400 m were investigated, it was concluded that under the current assumptions the head distribution near the tunnels is not very sensitive to Z_0 . Therefore, the only unknowns are Q_j/T because Q_j and T can not be obtained separately using this method.

The steady state pressure data observed in the intervals in MI fracture as well as the head at the tunnel walls can be used in equation 5.1 to estimate for the values of Q_j/T . Because there are more equations than the number of unknowns, least square optimization is used to determine the

coefficients. It was found that all values of Q_j/T are close and that the pressure distribution near the tunnels are nearly identical to that of the case where all values of Q_j are equal (Figure 5.5). As can be seen from the figure, the boundary condition along the perimeter of the modeled area is relatively constant except for the left and right sides where the effect of the nearby drifts are felt. Therefore we chose to use a constant head for outer boundary conditions.

A modification of the steady state head difference energy function was defined for this case because there is some uncertainty about the boundary conditions (see Section 2.1). We are not very confident of the outer boundary conditions. The relative head structure observed away from the boundaries is considered to contain more information. So, the outer boundary conditions were shifted up and down such that we could minimize the energy function relative to the observed values. In this way, we are putting all the emphasis on the pattern of head values, not on the absolute value of the head.



XBL 905-1793

Figure 5.4. Pressures in the exploratory boreholes configuration before the laboratory tunnel excavation (after Nagra, 1985).




Figure 5.5. Estimated pressure distribution in meters around the tunnels. The rectangle represents the region modeled.

6.0. The Migration Site Case

Once we have developed a set of possible conductors in the fracture flow system, and estimated the boundary conditions, we are ready to use the inversion algorithm. The hydrologic tests we use control the final result of the model building procedure. The hydraulic inversion technique, using a simulated annealing algorithm, finds discontinuous flow geometries which behave like the real system. Since the problem is ill-posed, we find multiple solutions. The resulting solutions are assessed using statistical methods and expert opinion. The MI study has allowed us to validate our inverse models and identify future improvements. We have modeled the MI zone as a two dimensional system. We first developed a template which contains all the channels which could possibly be in our fracture flow model. We then use the annealing algorithm to develop channel models for the steady state and the transient case. Five solutions were found for the steady state pressure case. A cross-validation study was done to estimate the prediction error associated with using the steady state solutions to predict the steady state head value at an arbitrary point in the channel model. For the transient case, only a single solution was found.

The template for the MI site was based on the geology of the MI zone. Geologic field investigations indicated that the hydraulic conductivity of the MI zone was probably greatest in the dip direction of the zone. Accordingly, the template for the migration zone was constructed with straight channels in the dip direction. This assumption is based on geologic evidence for hydrologic anisotropy of the fracture zone (see Section 4.3). The grid has a coarse border that surrounds a fine mesh in the vicinity of the eight wells and the drift. The central region with the fine mesh is where we expect annealing to be able to resolve details in the pattern of conductance. The border region moves the boundary outside the region we modeled hydrologically and diminishes perturbation from the boundary effects. No meaningful resolution of pattern is expected in the


border region. The template is a 50×42 meter grid with a 30×22 meter center region. The template has three fracture channel sets: the inclination from horizontal (as measured in the plane of the zone) are 30° , 90° , and 150° . Figure 6.1 shows the MI template. The spacing between fractures is 0.8 m in the center and 2.4 m in the border region. The eight boreholes are represented by marked nodes in the grid, and the drift is an open circle in the mesh.

We tried first to build the simplest possible model. Therefore, each channel in the template has the same conductance. In Section 5, we explained how we estimate the value of this conductance at each iteration, using a curve shifting technique. The specific storage of the channels in the border region is $3.0 \times 10^{(-5)}$ 1/m, and the center channels have a specific storage of $1.0 \times 10^{(-5)}$ 1/m, so that the total storativity in each region is the same.

The steady pressure case and the transient pressure case were studied. In the steady flow case, head measurements at the various wells were measured in response to drawdown to the drift running through the fracture. This is the initial condition for the transient case. Figure 5.2 shows the head record in the wells during a pumping test (see Section 5). The heads at the end of the recovery period were used for the steady state case and the transient drawdown data were used in the transient case.

6.1. Steady State Results

The annealing algorithm was used five times to find configurations which matched the observed data for heads at the end of the recovery period. Five configurations are given in Figure 6.2, case 1 through case 5. The first solution was found starting from a configuration with 60% of the elements present, the second solution was found starting from a configuration with 70% of the elements present, and the third, fourth, and fifth were found starting from configurations with 80, 90, and 100% of the elements present. We refer to these as case 1 through case 5. We can see that the annealed models show a lack of connection between wells 7, 11 and wells 4, 6, 9. Also, annealing has found lack of connection between well 11 and the boundary. This is happening because well 11 had a very low head, close to zero. As the drift boundary is at atmospheric pressure and the outer boundaries are held at a constant head estimated using the heads shown in



Key to well numbers.




Figure 6.1. The MI template.





Figure 6.2a. The first of five solutions which have similar geometries, case 1.



XBL 905-6386

Key to well numbers.






Figure 6.2b. The second of five solutions which have similar geometries, case 2.



XBL 905-6388

Figure 6.2c. The third of five solutions which have similar geometries, case 3.



XBL 905-6390

Key to well numbers.





Figure 6.2d. The fourth of five solutions which have similar geometries, case 4.



XBL 905-6391

Figure 6.2e. The fifth of five solutions which have similar geometries, case 5.

Figure 5.5, steady state annealing encourages a connection to the drift. Transient annealing might be used to see if this well is simply not connected to anything if the data is sufficient. This data would have to sequentially use different holes as different sources.

The energy of the solutions versus iteration number are given in Figure 6.3. Each of the solutions was found using the same temperature schedule, but starting from different configurations. This means that the number of changes the algorithm accepted at each temperature in the schedule is the same for all the solutions. The solutions are within the estimated measurement error of the observations, which we assumed was on the order of 0.3 m, based on the tolerance for similar instruments.

The prediction error for the steady state head response is calculated using the statistical method called cross-validation as described in Section 2.3. At the MI site, we have 8 wells and a drift. We use the steady state observed pressure response H_{obs}^i , at each well, i , under conditions of steady flow to the drift. We would like to know the error associated with using our annealing model to predict the steady state head response at an arbitrary nearby point. We do the following:

1. Leave the steady state head value for well i out of the energy function.
2. Five annealing solutions were found starting from cases 1 through 5. Each initial configuration, or case, had a different percentage of elements turned on.
3. For each end configuration, $C_1^{-i} \cdots C_5^{-i}$, calculate a predicted steady state value for well i . These predicted head values are $H_1^{-i} \cdots H_5^{-i}$.
4. Calculate the mean squared prediction error for well i :

$$PE^2(i) = \frac{\sum_{j=1}^5 |H_j^{-i} - H_{obs}^i|^2}{5}$$

The estimated prediction error is:

$$\hat{PE} = \left[\frac{\sum_{i=1}^8 PE(i)}{8} \right]$$

\hat{PE} is then an estimate of the error involved in using one annealing model to predict the head

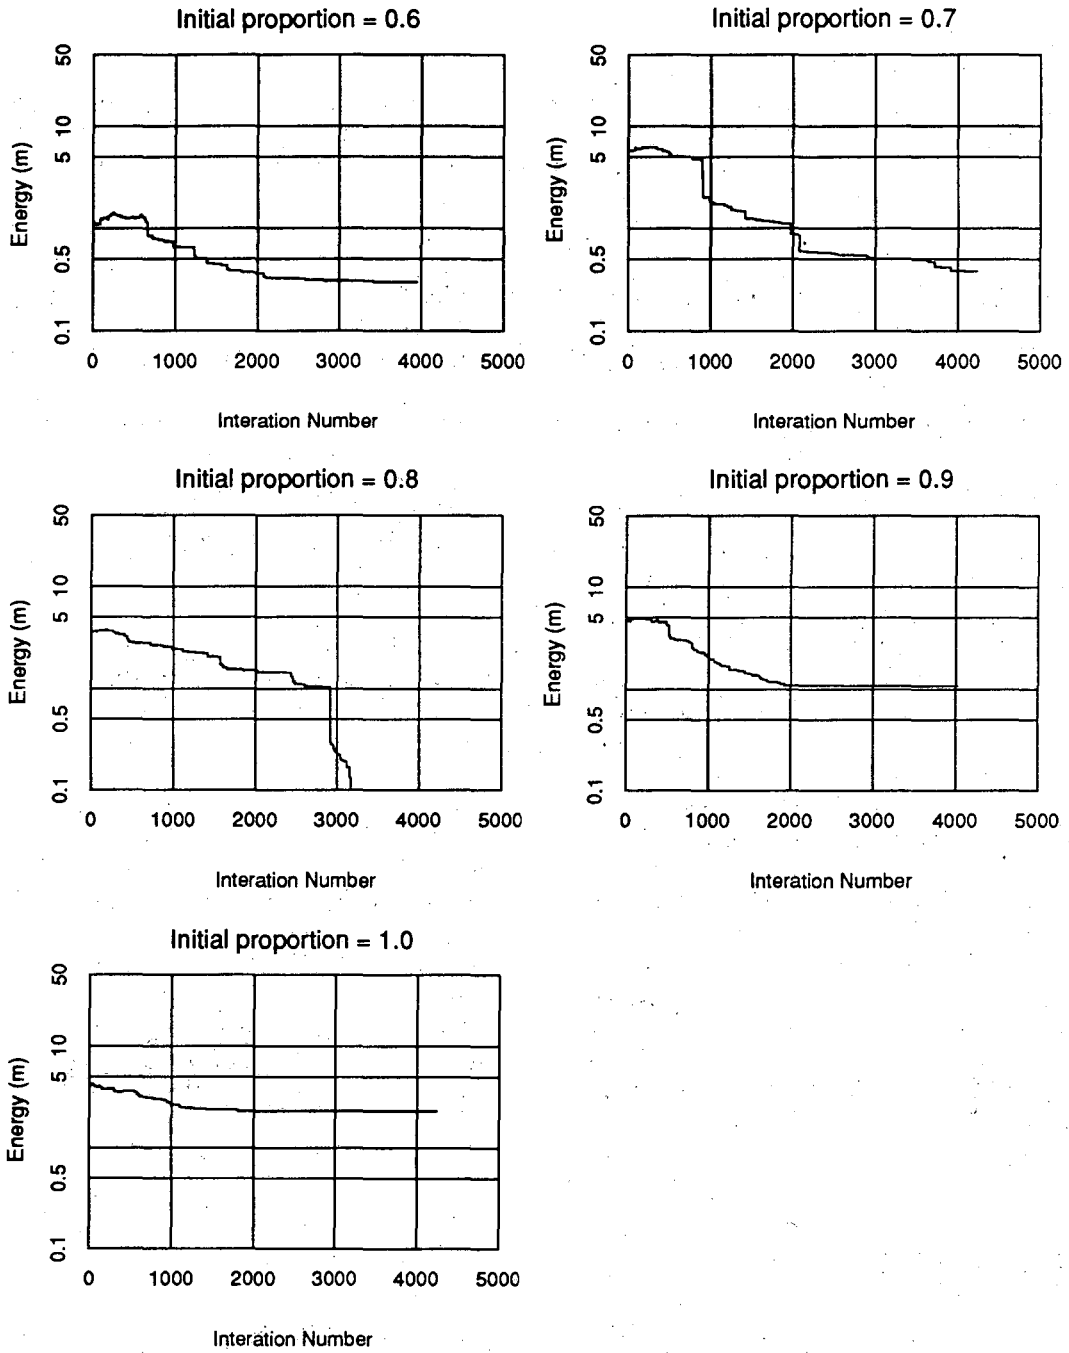


Figure 6.3. Energy versus iteration curves for the five steady state annealing solutions.

response of any other well in the vicinity.

Similarly, we can compute \hat{PE} for the mean or median predicted head response of the five solutions for each well left out. We used the loss function to compare the mean and median predicted value of the five solutions and found that the median was a slightly better predictor of steady state pressure at a given point on the grid. We then compared \hat{PE} for predictions made using each solution independently with \hat{PE} for predictions made by generating 5 solutions and using the median value as your prediction. The estimated prediction error found using a single solution was 4.3 m, and the estimated prediction error for using the median of five solutions was 3.3 m (Table 6.1). If we ignore well 11, the estimated prediction error using the median is 2.3 m. The estimated density of prediction error for the median of 5 solutions also shows the median is expected to give a lower prediction error. We can make better predictions if we base them on multiple solutions instead of a single solution. The prediction error for well 11 was very large, and tends to have a big effect on the prediction error (see Table 6.1, Appendix B).

Table 6.1. The observed steady state head values at each well and the predicted head values found using the median value for five annealing solutions. In each case the steady state head at the indicated well was left out of the energy function.

Well Left Out	Obs. Head	Median (\hat{h}) (m)	PE (i) (m)	PE (i, \hat{h}) (m)
4	9.97	6.64	3.8	3.3
5	10.95	5.95	5.0	5.0
6	10.22	7.72	4.0	2.5
7	0.64	0.988	3.2	0.3
8	3.37	0.96	2.5	2.4
9	8.07	9.99	2.1	1.9
10	4.0	5.07	2.8	1.1
11	1.04	11.37	11.1	10.3

6.2. Transient Results

The transient data were derived from a pumping test from well 9. Annealing this data required that we construct an energy function that compares observed and simulated head measurements taken at different times and locations.


In the case of MI, the flow rate was initially very unstable and a few minutes into the test the flow rate doubled. To forward-model this flow rate we assumed the flow rate to be constant at the higher level and adjusted the time at the observation wells downward to reflect the decrease in cumulative flow. The pressure data observed at various intervals were also "noisy" in the early time. We did not use this portion of the data for the energy calculation. To calculate the energy with equal weights in log time scale, and to smooth the data we used a window averaging scheme. The window averaging scheme used a moving averages algorithm.

The pressure in well 11 did not respond to the pumping of well 9. We assumed the pressure change in well 11 to be absolute zero in this case. Also, data from different intervals had a different "good" portion. Energy was calculated for the time duration that corresponds to the "good" portion of each well curve. Figure 6.4 shows the portion of the well test curves which were used in the energy function.

Because it was found that the results are not very sensitive to the boundary conditions prescribed at over a certain distance, we assumed the head was constant at 100 m away from the laboratory tunnel. The head value was assumed to be 100 m. By changing the mesh size, we found that the distance beyond which the results were fairly insensitive to small changes in the boundary conditions was approximately 80 meters. This was estimated from the measurements in the exploratory boreholes drilled prior to excavation of the laboratory tunnels. The boundary condition at the drift forced a positive flow into the drift during the well test.

Transient annealing was started with a triangular lattice consisting of two separate densities. The high density inner region reflects the possibility of high resolution near the well locations, while the lower density outer region minimizes the number of elements in the mesh to keep the computational time low. The inner core of elements are spaced at 0.8 m. This region encompasses all the wells of the MI fracture and the AU drift. This dense region extends five meters beyond the wells both vertically and horizontally. This produced an inner region of 30 m by 22 m.

The observation wells were located exactly at a node of the template if the distance to the nearest node was less than 0.1 meter otherwise a new node was created. If a new node was



XBL 9010-5942

Figure 6.4. The "good" portion of well test curves used to model the system.

created it was connected to the nearest three nodes with new elements which have the same conductance as the other elements in the inner region (wells 6 and 8). For improved accuracies in early time forward modeling, the elements connected to the pumping well (well 9) were subdivided. These elements were discretized from lengths of 0.8 m to lengths of 0.1, 0.2, and 0.5 m successively with increasing distance from the well.

The outer region extends another ten meters out from each side of the inner region, containing only a third of the fractures of the inner region, with element spacings of 2.4 m. The apertures and transmissivities were increased in the outer region to maintain consistent permeability between the two regions. Finally, as described in Section 5.2, the boundary conditions varied along the outer boundary and the drift wall to account for drawdown to 3 drifts.

Using the two level mesh proved to be difficult because the well response was seen at the boundary very quickly. So, we decided to create a new four level mesh that would increase the distance from the wells to the boundary. We embedded the mesh from the first run into another larger mesh, displayed in Figure 6.5. The old mesh is surrounded by another mesh half as dense as the previous outer mesh which encompasses the ventilation tunnel and the 5 m diameter access tunnel. This is enclosed by a network one-fourth as dense which intersects a fourth drift. Overall, the mesh extends 180 m vertically and horizontally. The conductances were also adjusted proportionally to maintain a constant average permeability throughout the network. The drawdown curves, calculated with the full grid are shown in Figure 6.6.

The annealing program was run on this mesh for a total of 15,510 iterations and the inner portion of the resulting mesh is displayed in Figure 6.7. The minimum energy configuration pressure drop curves are plotted in Figure 6.8 and the energy versus iteration number is shown in Figure 6.9. As in the steady-state case, the annealing procedure attempted to disconnect the wells from the boundary. There are only three connections left to the outer boundary from wells 4, 9, 6, 8, 10, and 5. This may not be significant because we do not expect good resolution outside the vicinity of the wells. The minimum energy mesh also shows that most of the vertical elements between wells 9 and 5 were deleted. This has the effect of removing much of the storage capacity




Figure 6.5. Template used for transient annealing of MI data.




Figure 6.6. Simulated well test response on the full template before annealing.




Figure 6.7. The annealed configuration at iteration number 15510.




Figure 6.8. Simulated well test response on the annealed configuration.




Figure 6.9. The energy versus iteration number for the transient annealing case.

between the wells while maintaining a strong hydrological connection. This is necessary to minimize the energy since well 5 is such a great distance from the pumping well but behaves similarly to wells 4 and 6 which are much closer. It seems that the low head at well 11 has been taken care of by isolating this well from the rest of the network, a solution quite different from the steady state response.

Another interesting feature of this network is that annealing was unable to match the response of well 7. Even though well 7 is physically close to the pumping well, 9, it is hydraulically insulated. Annealing is not able to find a tortuous enough path to account for this. This can be considered a problem of mesh refinement. However, it could also be that the real problem is three-dimensional and we have restricted this analysis to two-dimensions. Some of this difficulty might also be solved by allowing mesh elements to have different conductances. We are optimistic that annealing is able to identify much about the connectivity in the vicinity of the wells. A verification of this would require a substantial validation effort.

7.0. Conclusions and Recommendations

Three major features of the modeling approach described here are

- (1) The approach is focused on finding simplified equivalent models for a discontinuous systems where continuum approaches are not appropriate.
- (2) The model can be constructed iteratively, incorporating new data as it becomes available.
- (3) The approach allows flexible incorporation of many different types of data to constrain the inverse problem.

7.1. Modelling Approach and Conclusions

Our equivalent discontinuum models are designed to reproduce the hydrologic behavior of fracture networks in a simplified lattice model. Our model needs to be discontinuous because we can not predict the behavior of interest in a fractured flow system with continuous models. The model is mainly based on the known hydrologic response at the site, since the model will be used to predict hydrologic behavior.

The information obtainable at a given site is sketchy at first, over time more data is gathered. Even though the data at any given point in time is deficient, decisions have to be made using what is known, such as deciding what further data should be collected. At later stages of the effort, one has to decide how to change the model using new information. We believe that the process of building models and assessing their capabilities should be contemporaneous with the collection of data. By iterating between building models used to predict the system response and collecting data, one can better guide both activities. A conscientious use of the new data collected will allow model assessment, and the predictions made by the model will identify important questions which can be resolved by gathering new data. One should be able to learn much

about important statistical parameters, such as the spatial correlation and how to extrapolate model results to a larger region.

We may also incorporate geological and geophysical data into the model, to constrain the possible geometries which satisfy the observed hydrologic behavior. The solutions we obtain are a simple representation of our current state of all the different types of knowledge we have. However, representing an earth system is an ill-posed inverse problem. Many different mathematical models could be found that behave as the system behaves. The best we can do is to find some examples of models which will give the observed system response.

Developing an inverse model of this type requires a great deal of information, data, and expert opinion. The quality and availability of good data is crucial to the model, as is the scientific analysis. Eventually a model of this type should be able to fully integrate all the data available from hydrology, geology, and geophysics.

Within this philosophical framework there are practical problems which need to be addressed. These include the fact that the data varies in quality and there is a general problem with forward modeling of insitu experiments. Problems with data often occur if the data was collected under different conditions from those you expect to model, or is not measured in the location of interest. The so called "bad" data may still contain important information. For example, a ball-park estimate of the inflow rate that could be as much as 50% off might be more useful than a very accurate measurement of the pressure observed at a non-strategic observation point. Some values may be missing altogether, such as all the measurements in the third dimension. However the missing measurements may be essential input for the forward modeling. One such example is often the boundary conditions. A hydrologist may literally have to make up boundary conditions using expert reasoning. The discretion of a scientist is a necessary part of this approach.

Using transient data creates some difficulties and practical problems with forward modeling of field hydraulic tests. For instance, when there are multiple measurement points, all the recorded data do not always have a common "good" portion. One of the gauges may drift or pick up noise, or the pump may work erratically at times. The "bad" portion of the data induced

by these causes should not be matched against the forward modeling result. The "goodness" of data should be determined by the expert judgement of a hydrologist who is well-informed with the details of the field operation.

The steady state results show that we can easily match the MI case data and demonstrate the utility of multiple solutions to the inverse problem. The difference between the observed heads and those found numerically in each solution is very small (see Figure 6.3). We found using cross-validation that multiple solutions can be used to find a better estimator than a single solution. In addition, the solutions found after leaving out one of the steady state head values show that certain regions can be predicted using data points outside the region and others cannot. The annealing algorithm seems to 'smear' the nearby measured flow response over regions with no data available. However, unlike kriging, the nearby measurements are not linearly interpolated over these regions. The algorithm finds a random flow geometry which works; this will vary in each solution.

The steady state and transient solutions are different. As mentioned before the lack of continuity in some regions of the transient case flow geometry may indicate a low permeability zone. Similar regions in the steady state solution, such as the region between well 11 and the boundary, are harder to interpret. The large effect of the boundary conditions and the drift make it difficult to know if this could indicate a low permeability zone, or if it is a modelling artifact. However, both solutions clearly show a break between the region surrounding wells 7 and 11, and the region around wells 4, 6, and 9. The transient solution also shows that the region between well 9 and well 5 has low storativity relative to the rest of the modelled area.

The transient response of a system is more sensitive to the distribution of permeability than the steady state response. Theoretically, this makes transient annealing more attractive, but in the problems studied here we could not match the results from well 7 which is apparently in a low permeability zone. This is a significant result since it illustrates the limits of a model with uniform conductance for all channels. The model can not find a path long enough or tortuous enough to account for the delayed response at well 7. In the framework of this modelling technique, a

long tortuous path or a short lower permeability path will have the same effect. However, a short low permeability path is easier to represent.

After we have constructed a model and found solutions, we need to validate the model building process and assess the predictive performance of the results. The problems we solve have many acceptable solutions with different flow geometries. If the same generalized geometry is found in a certain region for many solutions, we may believe that it is likely to be real. However, we must rely on expert opinion to tell us if the flow geometry is reasonable, or if it is an artifact of the process. An example of such an artifact may be the steady state solution around well 11 discussed above. However, it is important to keep in mind that we are fundamentally trying to build a model of hydrologic behavior, not fracture geometry. Therefore, the best way to assess the validity of the model is by estimating the error associated with using the model to make hydrologic predictions. Estimating the prediction error requires that one use the model to predict some observed data that one did not use to build the model. We have shown that one method for estimating prediction error, cross-validation, can also be used to choose a single predictor if we have multiple solutions.

We believe that the model should reflect the inherent uncertainty in the problem. If the data is not adequate to give small uncertainty in one's solutions and predictions, then one should not find that they have small uncertainty. In the MI study, the prediction error for the steady state case was fairly high especially in the region around well 11 (see Table 6.1). This is a consequence of the extreme change in permeability between this region and wells 4, 9, and 6. In this case, given the data available, it would be unrealistic to expect a low prediction error. The data is not sufficient to allow us to know where the low permeability zone begins: it could be a small area around wells 11 and 7, or it could extend several meters in each dimension. Therefore, the solution has a high prediction error. In the transient case, we found that two flow tests were predicted fairly well, and one was not. This is partly a consequence of the fact that we could not match the data from well 7 very well.

7.2. Recommendations for Future Study

The MI experiment and the data provided by NAGRA enabled testing, assessment, some validation, and suggested improvements in the hydrologic simulated annealing fracture flow model.

Our long term goal is to develop an efficient method to build models which accurately predict hydrologic behavior. To accomplish this goal we need to integrate geological, geomechanical, geophysical, and hydrologic information to constrain the inverse site characterization problem. Further, any practical method of doing this must be faster than the current algorithm and allow us to represent a wider range of permeabilities. We are considering two paths to this end; using clusters and variable aperture, or using fractals with small scale tortuosity to represent the system. Finally, we need to better define the sensitivity and the predictive capabilities of the model. Below we first discuss approaches to the problem of integration, then efficiency and some comments about sensitivity studies.

An Integrated Model. The hydrogeological problem has many inverse solutions. We are learning how to use physical information to constrain the solution set. Geophysical, geomechanical, and geological data can be used to find an equivalent system with all known characteristics. At present, the physical data is used to generate a template which comprises some of the physical features. Fully incorporating this information is our final goal. We intend to use this physical information to identify fractured regions or major features, estimate their continuity and whether infilling is likely. We also hope to use radar to identify likely flow paths. Ultimately some of this information could be incorporated into an energy function. Hydrological data of different types may also be incorporated into an energy function to find solutions which match all the data. For example, we could use both steady state and transient data, or pressure and flow data to build the model. Tracer test data is also a very good candidate. In fact, we expect that tracer test data will greatly improve the inversion result. This is because tracer transport is advection dominated, whereas hydraulic pressure test is diffusion dominated.

Efficiency. Developing a practical algorithm requires operating on a larger scale than one element at a time, and representing a wider range of permeabilities. We build a model by examining one small channel at a time. The level of detail is too great since it takes many operations to resolve a big structure. We believe that the annealing algorithm could be accelerated by attempting to resolve bigger structures at each step. We plan to investigate using fractals and clusters of elements to work at a larger scale. The variation in permeability can be accommodated by using fractals with small scale tortuosity in some regions, or by allowing clusters to have a wider range of permeabilities.

Sensitivity. We also wish to investigate the sensitivity of our model to boundary conditions, conductance, and the template geometry. This may be done using synthetic cases. The geometry of the annealing solutions should also be studied using pattern comparisons such as the shortest path between two points, the spatial density of elements or permeability, etc.

This future research will involve both experimental and numerical work, and should enable us to build a good and efficient fracture hydrology model at a given site.

8.0. References

- Billiaux, D. Chiles, J. P., Hestir, K., and Long, J., 1989. Three-Dimensional Statistical Modeling of a Fractured Rock Mass - An Example for the Fanay-Augeres Mine, *International Journal of Rock Mechanics and Mining Science*, special issue on Forced Fluid Flow Through Fractured Rock Masses, in press.
- Black, J., Holmes, D., and Brightman, M., 1987. Crosshole Investigations, SKB Stripa Project Technical Report 87-18, Dec.
- Carrera, J. and Neuman, S. P., 1986. Estimation of Aquifer Parameters Under Transient and Steady State Conditions: 2. Uniqueness, Stability and Solution Algorithms, *Water Resources Research*, 22 (2) 211-227.
- Dougherty, D. and Marryett, R. A., 1990. Optimal Groundwater Management: 1. Simulated Annealing. Draft submitted for publication, February.
- Frick, U., Baertschi, P., and Hoehn, E., 1988. Migration Investigations, Nagra Special Bulletin, 23-34.
- Geier, J., Dershowitz, W. and Sharp, G., 1990. Prediction of Inflow into the D-holes at Stripa Mine, SKB Stripa Project Technical Report 90-06, April.
- Geman, S., and Geman, D., 1984. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6 (6) 721-741.
- Hajek, B. 1988. Cooling Schedules for Optimal Annealing, *Mathematic of Operations Research*, 13 (2) 311-329.
- Huber, P. J., 1981. *Robust Statistics*, John Wiley and Sons, New York.
- Karasaki, K., 1987. A New Advection-Dispersion Code for Calculating Transport in Fracture Networks. Lawrence Berkeley Laboratory Report LBL-22090.
- Keusen, H. R., Ganguin, J., Schuler, P., and Buletti, M., 1989. Felslabor Grimsel: Geologie, Nagra Technical Report 87-14, 120 pp.
- Kilpatrick, S., Gelatt, D. C., Vecchi, M. P., 1983. Optimization by Simulated Annealing, *Science*, 220.
- Long, J. C. S. 1983. Investigation of Equivalent Porous Medium Permeability in Networks of Discontinuous Fractures, Ph.D. Thesis, University of California at Berkeley.
- Long, J. C. S., Karasaki, K., Davey, A., Peterson, J., Landsfeld, M., Kemeny, J. and Martel. S., 1990. Preliminary Prediction of Inflow into the D-holes at the Stripa Mine, SKB Stripa Project Technical Report 90-04, February.
- Long, J. C. S., Remer, J. S., Wilson, C. R. and Witherspoon P. A. 1982. Porous Media Equivalents for Networks of Discontinuous Fractures, *Water Resources Research*, 18 (3) 645-658.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., and Teller, M., 1953. Equations of State Calculations by Fast Computing Machines, *Journal Chemical Physics*, 21.
- Nagra, 1987. Preliminary Volume of Nagra Technical Report 87-14.
- Olsson, O., Black, J., Gale, J., and Holmes, D., 1988. Site Characterization and Validation, Stage 2 - Preliminary Predictions, Swedish Geological Co. Report 88.

- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T., 1986. *Numerical Recipes*, Cambridge University Press.
- Robinson, P. C., 1984. Connectivity, Flow and Transport in Network Models of Fractured Media, Ph.D. thesis, St. Catherine's College, Oxford University, October.
- Solexperts, 1989. 566 Nagra Migration Long Term Constant Q Test in BOMI 87.009, Test Analysis and Description, 2/20/1989.
- Stone, M. 1974. Cross-Validatory Choice and Assessment of Statistical Predictions, *Journal Royal Statistical Society, Series B* (2) 111-147.
- Tarantola, A., 1987. *Inverse Problem Theory*, Elsevier Science Publishing Company Inc., New York, NY.
- Vomvoris, S., 1989. Written communication, 6/24/1989.

Appendix A

Implementation of the Simulated Annealing Algorithm

In the following pages, we give 10 of the 50 solutions found in our study of the effect of the starting point on the annealing solution (see Section 3.3). We have tried to assemble a set of configurations which are representative of the annealing solutions found for this case. Each annealing run began from a random configuration of elements with the density given. The synthetic and the full template are shown in Figures 2.1 and 2.2. The scale for all 10 solutions is shown in Figure A.1.




Figure A.1. An annealing solution found starting from a configuration with 60% of the possible elements.




Figure A.2. An annealing solution found starting from a configuration with 60% of the possible elements.




Figure A.3. An annealing solution found starting from a configuration with 70% of the possible elements.




Figure A.4. An annealing solution found starting from a configuration with 70% of the possible elements.




Figure A.5. An annealing solution found starting from a configuration with 80% of the possible elements.




Figure A.6. An annealing solution found starting from a configuration with 80% of the possible elements.

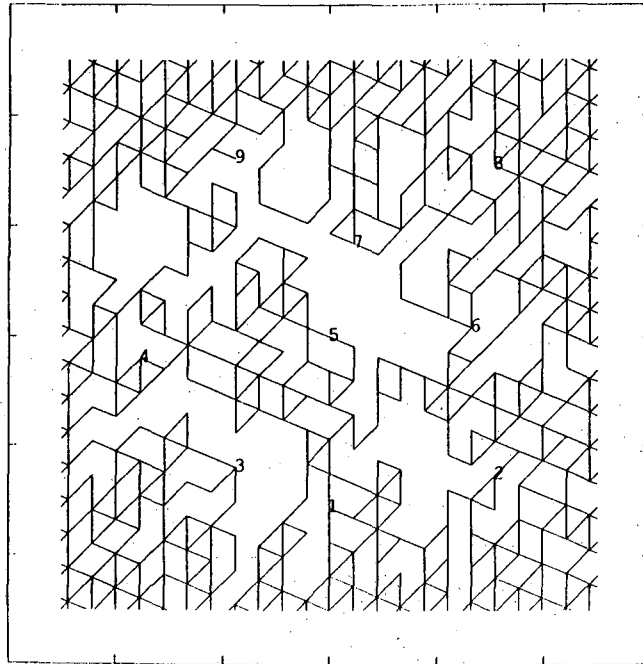


Figure A.7. An annealing solution found starting from a configuration with 90% of the possible elements.




Figure A.8. An annealing solution found starting from a configuration with 90% of the possible elements.




Figure A.9. An annealing solution found starting from a configuration with 100% of the possible elements.

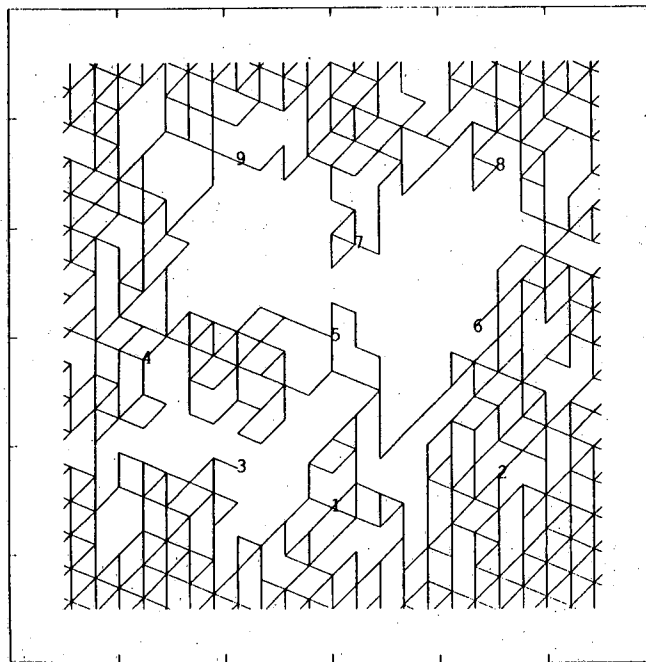


Figure A.10. An annealing solution found starting from a configuration with 100% of the possible elements.

Appendix B

Cross-Validation Studies

The steady state head value for each well was in turn left out of the energy function. Five annealing solutions were then found for each well left out and a flow calculation was done to find the predicted value for head at the well which was left out. The cross-validation flow geometries and a table giving the prediction error for each configuration follow.

The five solutions we found for each well used case 1 through case 5 as the starting configuration. As mentioned before, case 1 had an initial density of 60% and case 2 through case 5 had initial densities of 70%, 80%, 90% and 100%. We did this to see if some particular case had superior predictive qualities. As shown in the table, no case was outstanding. The scale for the flow geometries is given in Figure B.1.

Table B-1. Prediction error in meters. Each well was left out of the energy function.

Wells left out	Prediction Error (in meters)				
	Case 1 60%	Case 2 70%	Case 3 80%	Case 4 90%	Case 5 100%
4	-4.27	-3.33	-8.95	1.34	1.21
5	-4.62	-4.99	-5.73	-4.36	-5.20
6	-2.50	-4.61	-1.27	-2.15	9.48
7	10.24	0.34	0.34	4.67	0.34
8	-2.02	-3.37	-3.00	1.74	-2.41
9	2.20	1.91	1.88	2.34	1.92
10	4.12	-0.40	-1.47	1.07	6.89
11	10.67	10.36	11.39	11.79	11.37




Figure B.1. Cross-Validation Study. Annealing solution case 1, with well 6 left out.




Figure B.2. Cross-Validation Study. Annealing solution case 2, with well 6 left out.




Figure B.3. Cross-Validation Study. Annealing solution case 3, with well 6 left out.




Figure B.4. Cross-Validation Study. Annealing solution case 4, with well 6 left out.




Figure B.5. Cross-Validation Study. Annealing solution case 5, with well 6 left out.




Figure B.6. Cross-Validation Study. Annealing solution case 1, with well 7 left out.




Figure B.7. Cross-Validation Study. Annealing solution case 2, with well 7 left out.




Figure B.8. Cross-Validation Study. Annealing solution case 3, with well 7 left out.

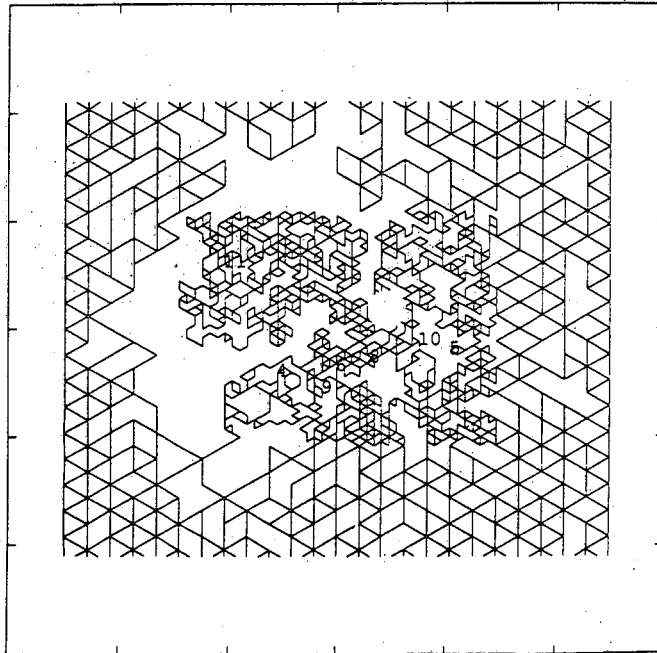


Figure B.9. Cross-Validation Study. Annealing solution case 4, with well 7 left out.




Figure B.10. Cross-Validation Study. Annealing solution case 5, with well 7 left out.




Figure B.11. Cross-Validation Study. Annealing solution case 1, with well 8 left out.




Figure B.12. Cross-Validation Study. Annealing solution case 2, with well 8 left out.




Figure B.13. Cross-Validation Study. Annealing solution case 3, with well 8 left out.




Figure B.14. Cross-Validation Study. Annealing solution case 4, with well 8 left out.




Figure B.15. Cross-Validation Study. Annealing solution case 5, with well 8 left out.




Figure B.16. Cross-Validation Study. Annealing solution case 1, with well 10 left out.




Figure B.17. Cross-Validation Study. Annealing solution case 2, with well 10 left out.




Figure B.18. Cross-Validation Study. Annealing solution case 3, with well 10 left out.




Figure B.19. Cross-Validation Study. Annealing solution case 4, with well 10 left out.




Figure B.20. Cross-Validation Study. Annealing solution case 5, with well 10 left out.




Figure B.21. Cross-Validation Study. Annealing solution case 1, with well 11 left out.




Figure B.22. Cross-Validation Study. Annealing solution case 2, with well 11 left out.




Figure B.23. Cross-Validation Study. Annealing solution case 3, with well 11 left out.




Figure B.24. Cross-Validation Study. Annealing solution case 4, with well 11 left out.




Figure B.25. Cross-Validation Study. Annealing solution case 5, with well 11 left out.

INFORMATION RESOURCES DEPARTMENT
LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720