

Demonstration of Optimization Techniques for Groundwater Plume Remediation

Stefan Finsterle

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

September 2000

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.

Ernest Orlando Lawrence Berkeley National Laboratory
is an equal opportunity employer.

Executive Summary

We examined the potential use of standard optimization algorithms for the solution of aquifer remediation problems. Costs for the removal of dissolved or free-phase contaminants depend on aquifer properties, the chosen remediation technology, and operational parameters (such as number of wells drilled and pumping rates). A cost function must be formulated that may include actual costs and hypothetical penalty costs for incomplete cleanup; the total cost function is therefore a measure of the overall effectiveness and efficiency of the proposed remediation scenario.

In this study, the cost function is minimized by automatically adjusting certain operational parameters. The impact of these operational parameters on remediation is evaluated using a state-of-the-art three-phase, three-component flow and transport simulator, which is linked to nonlinear optimization routines. The report demonstrates that methods developed for automatic model calibration are capable of minimizing arbitrary cost functions.

Two illustrative examples are presented. While hypothetical, these examples demonstrate that remediation costs can be substantially lowered by combining simulation and optimization techniques. The second example on co-injection of air and steam also make evident the need for coupling optimization routines with an accurate state-of-the-art process simulator. Simplified models are likely to miss significant system behaviors such as increased downward mobilization due to recondensation of contaminants during steam flooding, which can be partly suppressed by the co-injection of air.

Table of Contents

Abstract	iii
1. Introduction	1
2. Modeling Multiphase Contaminant Transport Using T2VOC	2
3. Optimization algorithms implemented in iTOUGH2	3
4. Example No. 1: Remediation of dissolved TCE plume	6
5. Example No. 2: NAPL removal by combined air-steam injection	13
6. Summary and conclusions	18
Acknowledgment	18
References	19

List of Figures

Figure 1. Log-permeability field of synthetic, two-dimensional, confined aquifer. 100 kg of TCE are spilled at the location marked by the square. The locations of six remediation wells are also indicated.	7
Figure 2. Distribution of dissolved TCE four months after the spill.	7
Figure 3. Reduction of cost function for four different minimization algorithms. Each square represents one T2VOC simulation.	9
Figure 4. Mass of dissolved TCE left in the aquifer as a function of pumped volume for the initial and optimized designs. The pumping schedule was optimized to minimize the total pumped volume when 5 kg of TCE are left in the aquifer.	10
Figure 5. Distribution of dissolved TCE after 10, 30, 100, and 240 days of continuous pumping with the initial pumping schedule (left column) and the optimized pumping schedule (right column). The pumping rate is indicated at each well.	12
Figure 6. Layout of test cell experiment for studying thermally enhanced soil vapor extraction.	14
Figure 7. Distribution of free-phase TCE after 12, 30, and 60 minutes after injection of pure steam (left column) and a mixture of steam and air mixture (right column).	15
Figure 8. Contour plots of cost terms as a function of steam enthalpy and air injection rate. The optimal design parameters minimize the total costs. The minimum was identified by the Levenberg-Marquardt algorithm within a few iterations, regardless of the starting point.	17

List of Tables

Table 1. Minimization algorithms implemented in iTOUGH2	4
Table 2. Optimized pumping schedules, cleanup times, and costs	8
Table 3. Input parameters for simulation of thermally enhanced soil vapor extraction experiment	13

1. Introduction

The design of a cleanup operation for a contaminated aquifer involves environmental, hydrological, technical, and economic issues. The main task is to locate and characterize the contaminant plume and to select an effective and efficient remediation technology. The suitability of a proposed method depends on the location and geometry of the plume, the hydrogeologic characteristics of the aquifer, the chemical properties of the contaminant, and the overall remediation goal. Once a technology has been chosen, the operational scheme and its parameters (for example, the number of wells and their pumping schedule) can be optimized to reduce remediation costs.

Management of groundwater remediation by means of optimization techniques seeks to maximize contaminant removal while minimizing capital, operating, and maintenance costs. Furthermore, technical constraints and regulatory cleanup standards must be observed. Using this approach to design cleanup operations for the remediation of dissolved or free-phase contaminant plumes in the subsurface requires predictive modeling capabilities in combination with nonlinear optimization techniques. Operational parameters can be determined by iteratively minimizing an objective function that involves actual or hypothetical cleanup costs. The procedure is similar to that used for automatic calibration of a process model against field data (referred to as “inverse modeling”), where hydrogeologic parameters are determined by minimizing the differences between the observed and calculated system response.

The purpose of this study is to examine whether the combination of a state-of-the-art simulation model and standard optimization techniques can be used to improve the design of a remediation system. The intent is to provide a proof of principle rather than to solve an actual remediation problem. Consequently, the study focuses on the numerical performance of the simulation and optimization tools; only a limited effort has been made to develop complex scenarios and realistic cost functions.

Computer programs have been developed at the Lawrence Berkeley National Laboratory (LBNL) to solve forward and inverse problems in groundwater hydrology. We make use of the iTOUGH2 code (Finsterle, 1999abc), which provides inverse modeling capabilities for the nonisothermal, three-phase, three-component flow and transport simulator T2VOC (Falta et al., 1995), which is based on the general-purpose integral finite difference code TOUGH2 (Pruess, 1991; Pruess et al., 1999).

A summary description of the T2VOC simulator is given in Section 2, followed by a brief discussion of the optimization algorithms implemented in iTOUGH2 (Section 3). Two illustrative examples are presented in this report. The first example (Section 4) simulates the remediation of a large contaminant plume through an array of extraction wells. The pumping rate in each well is optimized to reduce cleanup costs. The second example (Section 5) examines the performance of thermally enhanced soil vapor extraction, where the effectiveness of NAPL removal depends on steam temperature and the amount of co-injected air. Conclusions from this preliminary demonstration can be found in Section 6.

2. Modeling Multiphase Contaminant Transport Using T2VOC

The computer code used in this study is T2VOC (Falta et al., 1995) for modeling three-phase (gas, water, non-aqueous phase liquid (NAPL)), nonisothermal flow of water, air, and a water-soluble volatile organic compound (VOC) in three-dimensional heterogeneous porous media. T2VOC is an extension of the TOUGH2 general-purpose simulation program (Pruess, 1991; Pruess et al., 1999), which uses the integral finite difference formulation for solving fully coupled mass and energy balance equations.

Each of the three phases flows in response to pressure and gravitational forces according to the multiphase version of Darcy's law, which includes the effects of relative permeability and capillary pressure between the phases. Water properties in the liquid and vapor state are calculated using steam table equations given by the International Formulation Committee (IFC, 1967). Thermophysical properties of the NAPL phase (such as saturated vapor pressure and viscosity) are calculated as functions of temperature, while specific enthalpy and density are computed as functions of both temperature and pressure. Gas phase thermophysical properties (including molecular diffusivities) are considered to be functions of temperature, pressure, and gas phase composition. A general equation-of-state is provided to compute the necessary NAPL/VOC thermophysical and transport properties. Based on semi-empirical corresponding-states methods, the chemical parameters are calculated as functions of critical properties such as the critical temperature and critical pressure, which are readily available for many substances (Reid et al., 1987). This approach makes T2VOC applicable to a variety of contamination problems involving different NAPLs that are either denser or lighter than water and that have different vapor pressures and solubilities.

For numerical solution, the balance equations are discretized in space based on an integral finite difference formulation (Narasimhan and Witherspoon, 1976). Time is discretized fully implicitly using first-order backward finite differences. Discretization results in a set of nonlinear coupled algebraic equations which are solved simultaneously by means of Newton-Raphson iterations. A conjugate gradient algorithm is used to solve the linear equations arising at each iteration step (Moridis and Pruess, 1995).

For a detailed description of the physical processes, governing equations, and numerical schemes employed in T2VOC, see Falta et al. (1992, 1995) and Pruess et al. (1999).

3. Optimization Algorithms Implemented in iTOUGH2

As discussed in the introduction, iTOUGH2 combines the simulation capabilities of T2VOC with optimization techniques to perform model calibration (inverse modeling) and to support the design of remediation systems. In both applications, a performance measure is either minimized or maximized by adjusting certain input parameters or design variables. For example, model calibration consists of reducing the differences between the simulation results and measured data (such as water potentials or temperatures) by adjusting the model input parameters (such as the absolute permeability or thermal conductivity of the porous medium). A cleanup operation can be improved by minimizing, for example, the remediation time, which can be achieved by increasing pumping rates or steam temperature. Increasing pumping rates or steam temperature, however, leads to higher energy costs. Therefore, the objective function to be minimized should also reflect these costs in order to obtain an optimal remediation design.

Solving an optimization problem thus requires two steps. First, a cost function has to be defined as a function of T2VOC output variables, which in turn depends on certain T2VOC input parameters. This function is usually highly nonlinear and can even be discontinuous. Moreover, it contains contributions from different sources, which have to be appropriately weighted against each other. Secondly, a minimization algorithm is needed, capable of updating T2VOC input parameters in order to reduce the value of the cost function. In the remainder of this section, we present the methods used to minimize the cost function.

The cost function is denoted by Z and consists of a sum of cost contributions, $Z = \sum z$. The vector \mathbf{p} contains all operational parameters (i.e., T2VOC input parameters) that are automatically adjusted to minimize Z . The iterative minimization of Z involves computing a correction vector $\Delta\mathbf{p}_k$ such that the new parameter set $\mathbf{p}_{k+1} = \mathbf{p}_k + \Delta\mathbf{p}_k$ leads to a reduction in the cost function, $Z_{k+1} < Z_k$, at each iteration k . A number of algorithms are available to solve nonlinear minimization problems (see, for example, Gill et al. (1981)). They differ in their strategy to calculate $\Delta\mathbf{p}_k$. Some methods require the evaluation of partial derivatives of the cost function with respect to the parameters. Many of the methods are specialized to efficiently deal with nonlinear least-squares problems, where the objective function is a sum of squared weighted residuals. The cost functions considered here are likely to be nonquadratic, and different, less efficient methods may need to be applied. Table 1 summarizes the minimization algorithms implemented in iTOUGH2. As they are able to track more general problems, the algorithms become less efficient, i.e., they require more solutions of the forward problem, which involves simulating the cleanup operation by solving the corresponding multiphase flow and transport equations. For a detailed description of the algorithms and their implementation, the reader is referred to Finsterle (1999a).

Many of the forward runs conducted to evaluate the cost function are independent from each other and can thus be performed in parallel. With the exception of Simulated Annealing, all methods shown in Table 1 are implemented for parallel execution using message passing based on the Parallel Virtual Machine concept (PVM; Geist et al., 1994). For details, see Finsterle (1998).

Table 1. Minimization Algorithms Implemented in iTOUGH2

Minimization algorithm	Description	Advantages and disadvantages
Gauss-Newton*	$\Delta \mathbf{p} = -(\mathbf{J}^T \mathbf{W} \mathbf{J})^{-1} \mathbf{J}^T \mathbf{W} \mathbf{z}$	Efficient for linear least-squares problems only.
Levenberg-Marquardt*	$\Delta \mathbf{p} = -(\mathbf{J}^T \mathbf{W} \mathbf{J} + \lambda \mathbf{I})^{-1} \mathbf{J}^T \mathbf{W} \mathbf{z}$	Efficient for nonlinear least-squares problem; requires derivatives.
Downhill Simplex	Minimum approach through sequence of reflections, expansions, and contractions of an $n+1$ simplex.	No assumptions made about form of cost function; relatively inefficient.
Simulated Annealing	Take random steps of decreasing average size; accept uphill steps with certain, decreasing probability.	No assumptions made about cost function; may escape local minima; inefficient.
Grid Search	Evaluate cost function in entire parameter space.	Complete information about cost function; very inefficient.

* \mathbf{J} Jacobian matrix; \mathbf{W} : Weighting matrix; \mathbf{I} : Identity matrix; \mathbf{z} : Residual/cost vector

Optimization problems are susceptible to being ill-posed, which leads to nonunique or unstable solutions. If multiple parameter sets yield nearly identical values of the performance measure at or near the minimum, the solution is nonunique. While nonuniqueness is a serious deficiency if estimating hydrogeologic parameters (inverse modeling), the fact that multiple solutions may exist is of little practical concern for cost-minimizing optimization problems such as the ones discussed here. A cost-effective remediation design is an acceptable solution even if additional, equally favorable configurations may exist. This must be distinguished from the potential presence of (multiple) local minima, representing sub-optimal solutions. Note that the presence or absence of local minima is a characteristic of the cost function and the choice of the parameters subjected to optimization. However, the presence of local minima does not mean that the optimization problem is ill-posed. The user must simply select an appropriate minimization approach to either detect, avoid, or escape local minima.

As mentioned above, nonuniqueness is not considered problematic by itself. However, nonuniqueness is often accompanied by instability, where vastly different solutions are obtained as a result of small changes in the definition of the cost function. Such instabilities should be avoided by a careful formulation of the cost function and deliberate selection of the design variables to be optimized.

The most important element of the optimization approach discussed here is the numerical process model itself, because the optimal design is calculated based on a numerical prediction of the cleanup operation. This prediction is based on a model that is uncertain. Sources for errors include: (1) unknown or uncertain aquifer properties (e.g.,

spatial distribution of permeability and adsorption coefficients); (2) incomplete process description (e.g., neglect of multicomponent diffusion in gas phase); (3) model simplifications (e.g., simulating three-dimensional flow using a two-dimensional model; specifying homogeneous model domains despite the presence of multi-scale spatial variability of formation properties); and (4) uncertainty in initial and boundary conditions (e.g., uncertain recharge rates; uncertain location and extent of contaminant plume). The solution to the optimization problem is likely to be (significantly) affected by a change in the underlying flow and transport model. Therefore, site characterization is a very important aspect of the overall optimization problem. Sufficient data must be collected to be able to determine key hydrogeologic properties. These are the properties that most strongly affect the predicted fate of the contaminants during the simulated cleanup operation.

Site characterization through monitoring or active aquifer testing is in itself an optimization problem, in which data of sufficient sensitivity and quality must be obtained to yield parameter estimates of acceptably low uncertainty. In summary, test design, monitoring, inverse modeling, and remediation design are strongly related and pose an overall optimization problem. Recall that this study deals only with the computational aspects of finding cost-effective design variables for a given cleanup operation and a given model.

4. Example No. 1: Remediation of Dissolved TCE Plume

Consider a confined aquifer of 10 m uniform thickness. The aquifer is heterogeneous with a mean permeability of 10^{-11} m^2 , a standard deviation of one order of magnitude, a correlation length of 30 m along the main west-east (X) flow direction, and a correlation length of 10 m in the north-south (Y) direction. The log-permeability field shown in Figure 1 was generated using Sequential Indicator Simulations (Deutsch and Journel, 1992). Effective porosity is 0.4. A natural hydraulic gradient of 0.01 is imposed across the model domain of 200 m length; no-flow boundaries are prescribed along the northern and southern sides of the model. The model domain is discretized into $40 \times 25 \times 1 = 1,000$ gridblocks with dimensions $5 \text{ m} \times 5 \text{ m} \times 10 \text{ m}$.

We simulated a potential contamination of the aquifer by releasing 100 kg of trichloroethylene (TCE) at $X = 47.5 \text{ m}$ and $Y = 67.5 \text{ m}$ (the spill location is indicated by a square in Figure 1). TCE has a water solubility of 1.51×10^{-4} [mole fraction]. Consequently, the spill is assumed to be completely dissolved as the contaminant plume spreads. Figure 2 shows the concentration of TCE dissolved in the aqueous phase four months after the spill. The bulk of the TCE mass follows a relatively narrow high-permeability channel, but significant amounts of TCE disperse in longitudinal and transverse directions as a result of aquifer heterogeneity. The contaminant plume eventually bifurcates into multiple fingers.

In a very simplified scenario, we assume that the contamination will be remediated by drilling six wells approximately aligned with the plume axis; the pattern is shown in Figure 2. This requires that some (albeit incomplete) information about the location of the contaminant plume is available either through sampling or preliminary modeling. We further assume that the costs for remediating the site are directly proportional to the total amount of contaminated water being extracted from the subsurface. In other words, the cost function to be minimized is simply given by

$$Z = \sum_{i=1}^6 q_i t \quad (1)$$

where q_i is the constant pumping rate in Well i , and t is the time required to extract a pre-defined amount of TCE, for example 95% of the total spill.

It is obvious that this scenario requires significant refinement in order for it to be realistic. For example, the cost function could be expanded to include (1) drilling costs, whereby a well should not be drilled if the pumping rate is less than a certain threshold value, (2) concentration-dependent treatment costs, (3) capital costs as a function of overall remediation time. In a variation of the cleanup operation, one could consider time-dependent pumping rates, whereby pumping is reduced or wells are shut down as soon as the TCE concentration reaches a certain low value. Treated waste water could be reinjected to impose hydraulic controls on the contaminant plume. Furthermore, different criteria for stopping remediation could be applied. Instead of requesting that a certain (large) amount of TCE be removed (as in this study), one could require that the TCE concentration at a nearby drinking water well is not to exceed a certain level. Even more stringent would be the requirement to comply with drinking-water standards everywhere within the modeled domain. The optimal solution is likely to (potentially strongly) depend on the chosen remediation target, the factors entering the cost function, and the details of the proposed cleanup operation (i.e., the design variables).

If each well is pumping at a constant rate of 1 kg s^{-1} , an initial simulation shows that 95% of the spilled TCE would be recovered within $t = 240$ days. Based on Eq. (1), the “cost” for this initial scenario is calculated to be 124,800 units.




Figure 1. Log-permeability field of synthetic, two-dimensional, confined aquifer. 100 kg of TCE are spilled at the location marked by the square. The locations of six remediation wells are also indicated.




Figure 2. Distribution of dissolved TCE four months after the spill.

Using iTOUGH2, the pumping rates in each of the 6 wells were automatically adjusted in an attempt to reduce the cost function. The results obtained with four different minimization algorithms are summarized in Table 2. All algorithms were able to reduce the cost function, albeit to various degrees and coming up with different optimal pumping schedules—the cost function obviously exhibits many local minima and the solution is nonunique, as expected. The Gauss-Newton algorithm was quickly trapped in a local minimum. Nevertheless, requiring only a few simulations, the cost function was considerably reduced, yielding a suboptimal solution with a relatively low total pumping rate and a correspondingly long cleanup time. The generality of the Levenberg-Marquardt algorithm (even though designed to minimize nonlinear least-squares functions) was able to further reduce costs to a value that is approximately half of that obtained with the original scenario. The Downhill Simplex algorithm performs well in this problem, arriving at even lower costs using a relatively small number of forward simulations. Finally, Simulated Annealing was able to come up with a pumping schedule that yields the smallest costs. However, a very large number of simulations was required to arrive at the solution. It is interesting to note that the cleanup time with the optimized pumping schedule is approximately the same as with the original design, with a 60% reduction in the total amount of water being pumped.

Table 2. Optimized Pumping Schedules, Cleanup Times, and Costs

MA*	# of runs	Time [day]	Cost function	Pumping rate [kg/s]						Total
				Well 1	Well 2	Well 3	Well 4	Well 5	Well 6	
-	1	240	124,800	1.00	1.00	1.00	1.00	1.00	1.00	6.00
GN	15	340	78,800	-	1.25	0.63	0.13	1.11	0.89	4.01
LM	133	250	61,400	0.49	1.52	-	0.14	0.64	-	2.79
DS	89	190	58,700	0.21	1.95	0.27	-	1.13	-	3.56
SA	1224	250	53,300	0.15	1.20	-	-	0.91	0.17	2.43

* Minimization algorithm:

GN = Gauss-Newton; LM = Levenberg-Marquardt; DS = Downhill Simplex; SA = Simulated Annealing

The strategy and performance of the four minimization algorithms considered here is visualized in Figure 3. Each square represents one T2VOC simulation. Note that the Gauss-Newton and Levenberg-Marquardt methods require n simulations at each iteration to numerically evaluate the sensitivity matrix, where n is the number of adjustable parameters. The cost function may temporarily increase as unsuccessful uphill steps are proposed by the Levenberg-Marquardt algorithm. The Downhill Simplex algorithm shows cost increases prior to a contraction of the simplex. The random nature of Simulated Annealing leads to many runs with higher costs than previous simulations. These higher costs are sometimes accepted with a decreasing probability, allowing the algorithm to escape local minima.

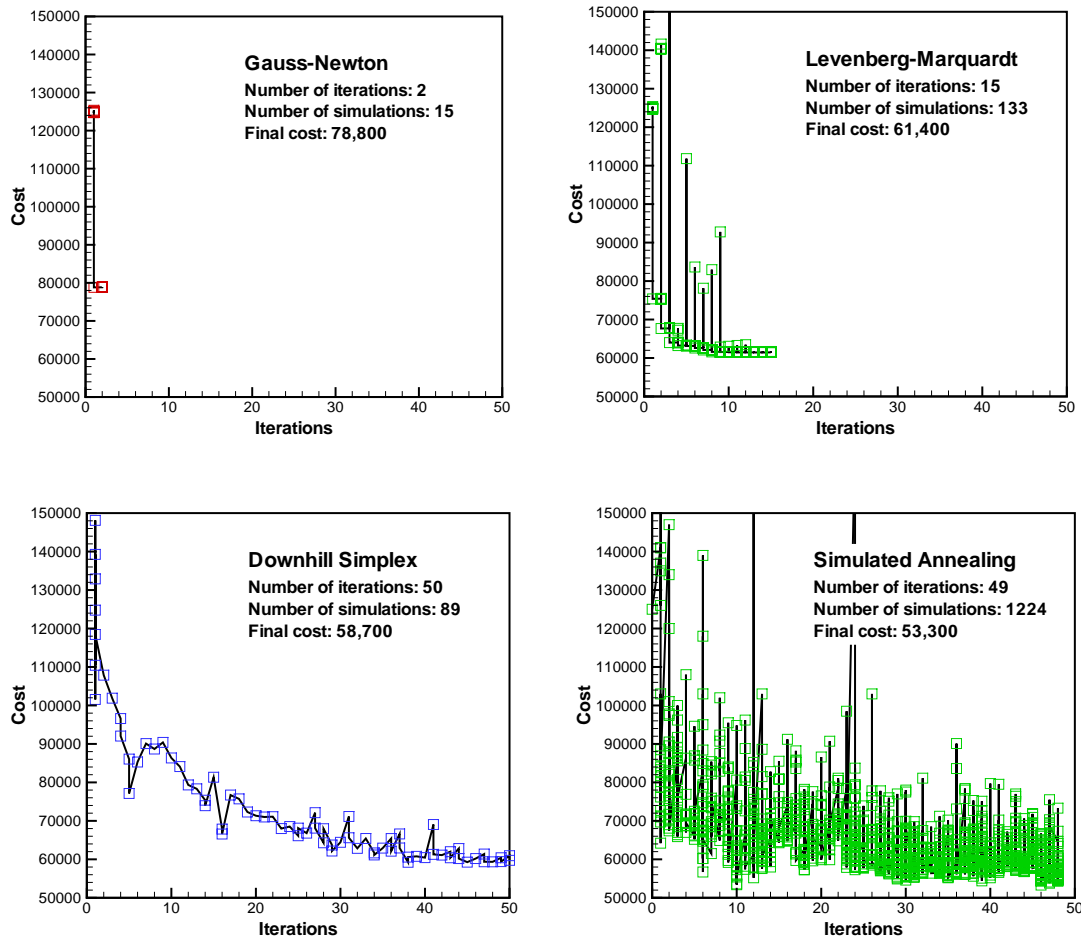


Figure 3. Reduction of cost function for four different minimization algorithms. Each square represents one T2VOC simulation.

Figure 4 shows the amount of dissolved TCE removed from the aquifer as a function of pumped groundwater volume for the initial design (with a rate of 1 kg s^{-1} in each of the six wells) and for the optimum design determined by Simulated Annealing. Recall that optimization occurred by minimizing the total volume pumped at the time when 95 kg of TCE have been removed. The cost reduction is indicated by the horizontal arrow in Figure 4. Using the pumping rates proposed by the optimization model, more TCE is removed with each cubic meter of groundwater as compared to the initial design. However, the changing width of the gap between the two curves indicates that a different optimal solution would be obtained if the remediation criterion were changed. However, it appears that the automatically determined pumping schedule remains superior to the initial design even if the threshold of acceptable residual TCE is reduced.




Figure 4. Mass of dissolved TCE left in the aquifer as a function of pumped volume for the initial and optimized designs. The pumping schedule was optimized to minimize the total pumped volume when 5 kg of TCE are left in the aquifer.

Figure 5 shows the remediation process using the initial pumping schedule (left column) in comparison with the optimized design (right column). The uniform pumping rates of the initial design lead to an area of relative stagnation between Wells 1, 2, 4, and 5. Moreover, pumping in Well 4 has the effect of pulling the plume into a previously uncontaminated region of the aquifer. The symmetry is broken in the optimized design, where the largest rates are assigned to the wells in the center of the plume, i.e., Wells 2 and 5. Wells 1, 3, 4 and 6 could probably be shut down without significant loss in overall remediation effectiveness. This would save additional drilling costs, which were not considered in the cost function (1), making the design even more favorable.

It should be noted that the optimized design targets the bulk of the TCE plume, i.e., it efficiently remediates the high-concentration area while neglecting the low-concentration, downstream tail of the plume. This is a consequence of the chosen remediation criterion, which only requires a reduction of the total TCE in the aquifer. Additional terms would have to be added to the cost function in order to ensure that no contaminated water is able to escape downstream.

This first example illustrates that the optimization algorithms implemented in iTOUGH2 can be used to improve remediation designs by automatically minimizing a cost function. While the cost function chosen here is extremely simple, its complexity can be increased arbitrarily. A key feature of this approach is that the costs are calculated based on an accurate process model simulating flow and transport in the aquifer for each remediation scenario. One should realize, however, that the optimization relies on accurate characterization data. If the permeability structure of the aquifer or the initial contaminant distribution are unknown, the optimal pumping schedule is not guaranteed to be as successful when implemented in the field.

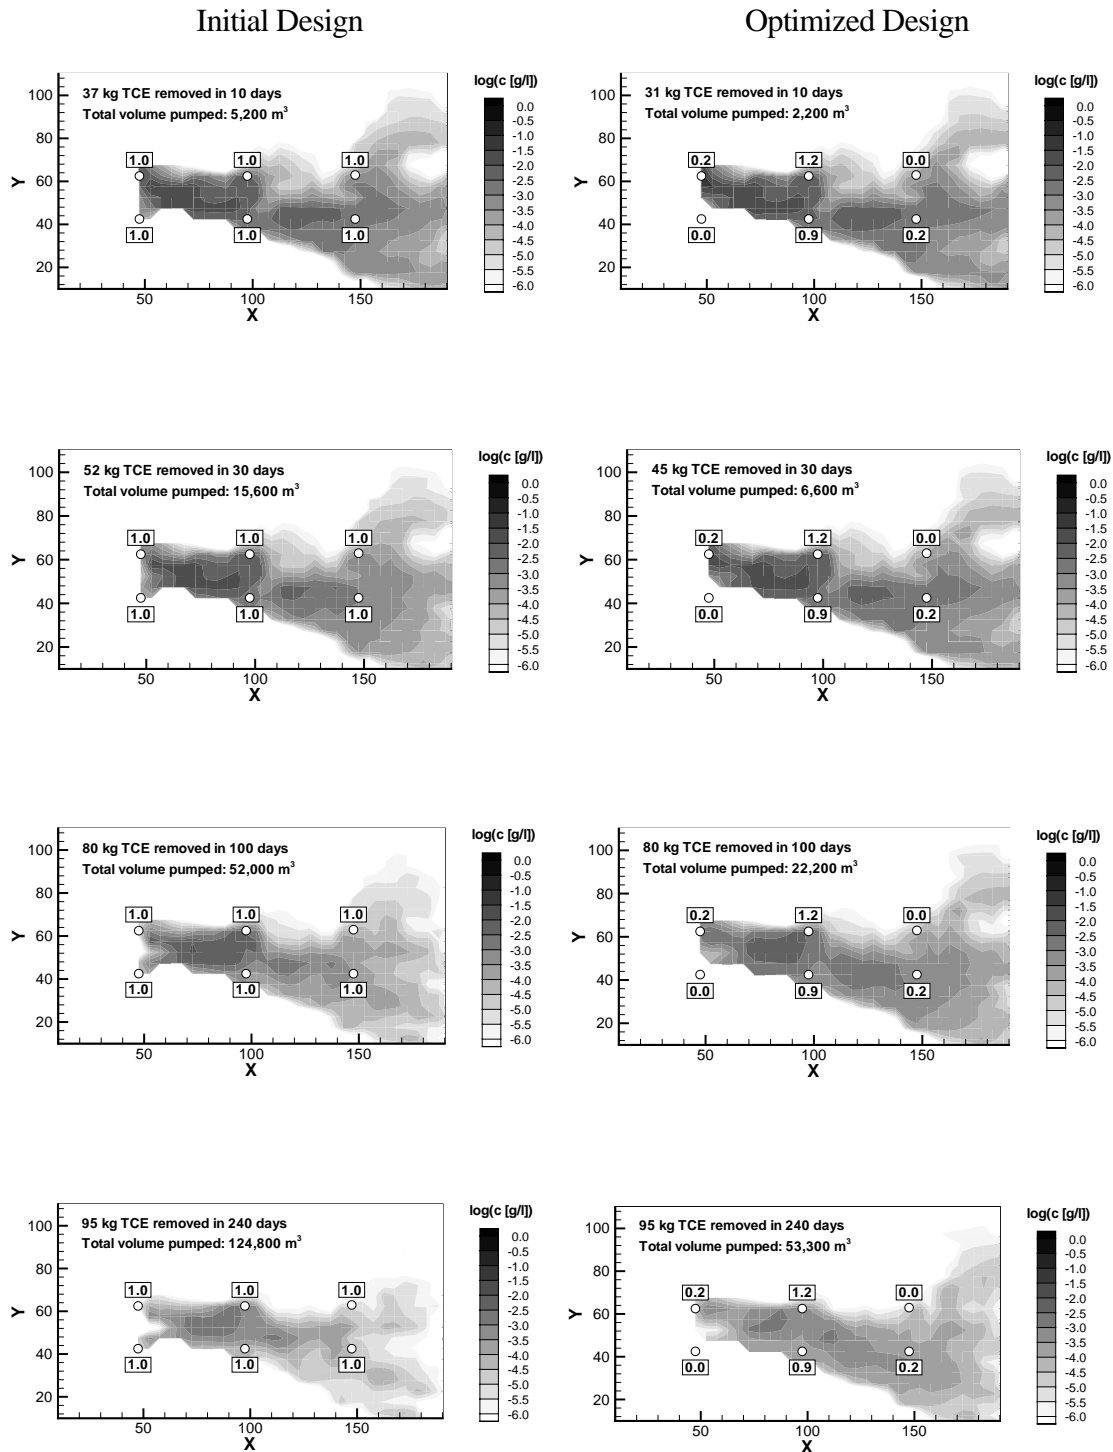


Figure 5. Distribution of dissolved TCE after 10, 30, 100, and 240 days of continuous pumping with the initial pumping schedule (left column) and the optimized pumping schedule (right column). The pumping rate is indicated at each well.

5. Example No. 2: NAPL Removal by Combined Air-Steam Injection

The widely used soil vapor extraction method for the removal of volatile NAPLs from the unsaturated zone is often inefficient and requires long remediation times. Steam injection is a means to increase NAPL volatilization, thus increasing the mobility of the contaminant by inducing a phase change. Even residual NAPL contamination in low-permeability lenses can be mobilized and successfully transported to an extraction well. However, since the propagation of the thermal front is retarded compared to the steam front, the volatilized contaminant often recondenses as it encounters cooler regions. High NAPL saturations may build up at these condensation fronts, making the NAPL sufficiently mobile for it to be transported downwards by gravity. This process reduces the efficiency of thermally enhanced soil vapor extraction systems and may even lead to an unwanted mobilization of contaminants towards the water table.

Recondensation of volatilized NAPL can be partly suppressed by co-injection of a noncondensable gas such as air, increasing advective transport of volatile contaminants to the extraction well. A laboratory experiment has been performed in the VEGAS facility at the University of Stuttgart, Germany, to study thermally enhanced soil vapor extraction using steam/air mixtures (Betz, 1998; Betz et al., 1998). We use a configuration similar to that of the VEGAS experiment to perform our synthetic optimization study.

The layout is shown in Figure 6. The model domain of dimensions $1.0 \text{ m} \times 0.1 \text{ m} \times 0.75 \text{ m}$ represents a vertical test cell filled with coarse sand with a lens of finer sand built into it (solid box). A stochastic, spatially correlated, anisotropic permeability field is generated to induce heterogeneity. Initially, the sand is at residual water saturation. TCE at residual NAPL saturation is emplaced (dash-dotted box), partly penetrating the sand lens. Steam and air are injected at the left-hand side boundary between a depth of $Z = -0.3$ and -0.2 m , and a vertical extraction well is installed at $X = 0.875 \text{ m}$ to a depth of $Z = -0.5 \text{ m}$. Input parameters are summarized in Table 3.

Table 3. Input Parameters for Simulation of Thermally Enhanced Soil Vapor Extraction Experiment

Parameter	Coarse sand	Fine sand
Mean log-permeability [m^2]	-9.0	-11.0
Porosity	0.42	0.42
Rock grain density [kg m^{-3}]	2650	2650
Rock grain specific heat [$\text{J kg}^{-1} \text{ }^\circ\text{C}^{-1}$]	840	840
Heat conductivity [$\text{W m}^{-1} \text{ }^\circ\text{C}^{-1}$]	wet dry	1.60 0.35
<i>Relative permeability: Stone (1970); Capillary pressure function: Parker et al. (1987)</i>		
Residual water saturation	0.1	0.1
Residual NAPL saturation	0.1	0.1
Residual gas saturation	0.01	0.01
n -parameter	4.0	4.0
$\alpha_{\text{gas-NAPL}}$ [m^{-1}]	15.0	5.0
$\alpha_{\text{NAPL-water}}$ [m^{-1}]	150.0	50.0
<i>Initial conditions</i>		
Initial gas pressure [Pa]	101300	101300
Initial water saturation	0.1	0.1
Initial NAPL saturation in dash-dotted box	0.1	0.1




Figure 6. Layout of test cell experiment for studying thermally enhanced soil vapor extraction.

Figure 7 shows the NAPL saturation after 12, 30 and 60 minutes of pure steam injection (left column) and combined injection of steam and air (right column). If only steam is used to thermally enhance volatilization and contaminant removal, recondensation of TCE at the steam front leads to the appearance of a free NAPL phase near the bottom of the test cell. To avoid this potentially dangerous side-effect of steam flooding, a penalty term is added to the cost function. The appearance of TCE as a free NAPL phase near the bottom of the test cell will lead to an appropriately weighted increase in the cost function. Any parameter combination that leads to downwards NAPL mobilization is therefore rejected by the minimization algorithm.

Figure 7 also shows that the co-injection of air enables increased removal of TCE from the low-permeable sand lens, whereas the scenario with steam injection simply volatilizes and recondenses TCE within the lens.

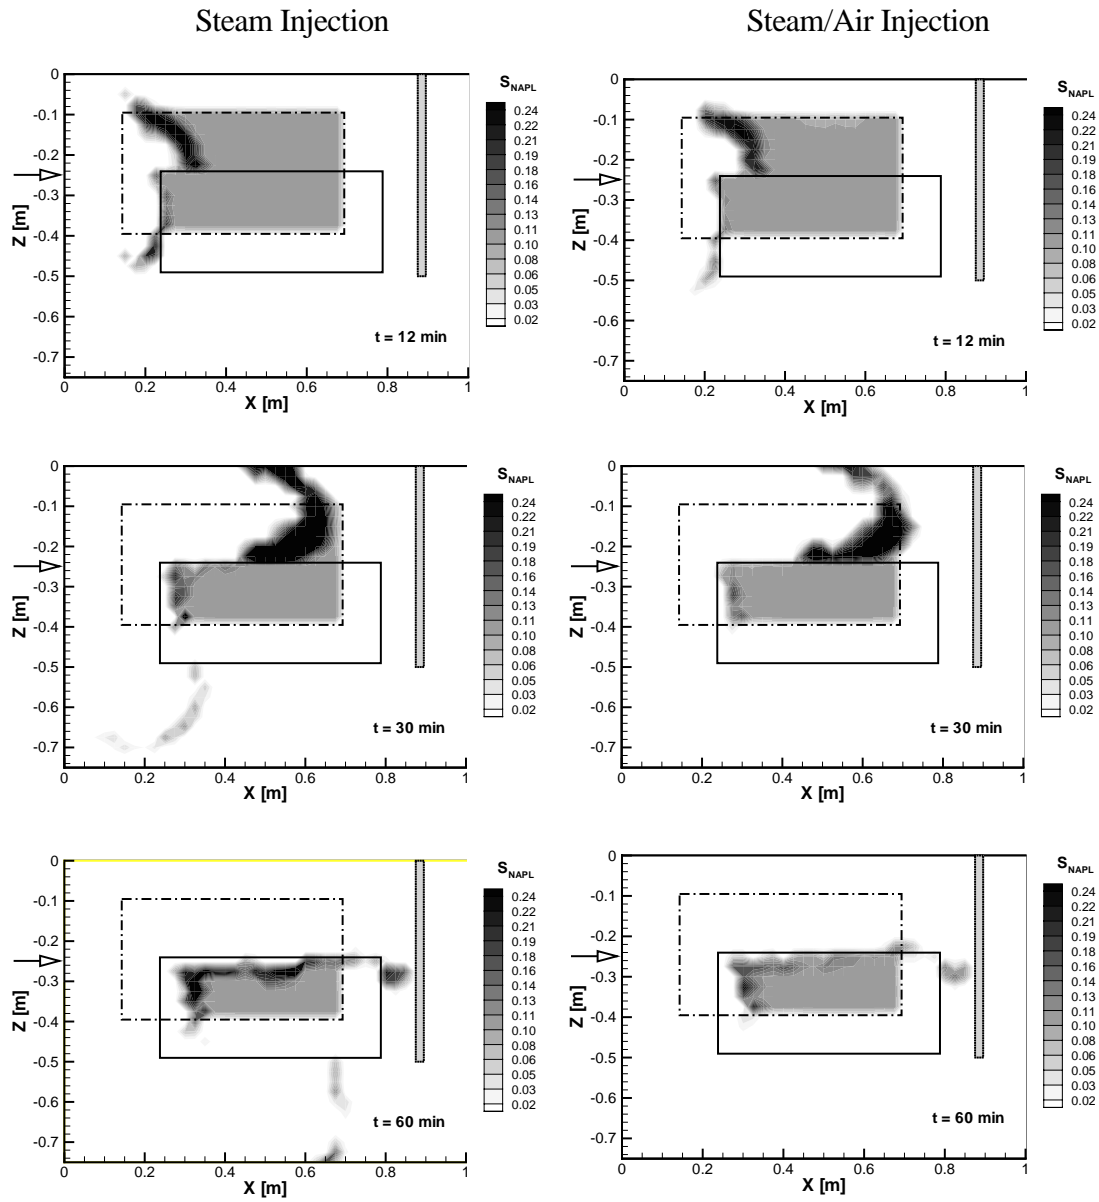


Figure 7. Distribution of free-phase TCE after 12, 30, and 60 minutes after injection of pure steam (left column) and a mixture of steam and air mixture (right column).

The overall approach to such an optimization problem is demonstrated using the following, generic cost function:

$$Z = a \cdot M_{TCE} + b \cdot M_{NAPL} + C_{steam} + C_{air} \quad (2)$$

The first term in (2) represents residual contamination, where M_{TCE} is the total mass of TCE (volatilized, dissolved, or in NAPL phase) left in the test cell after 1 hour of soil vapor extraction. The second term is a penalty term for NAPL that potentially reaches the water table as a result of mobilization and recondensation effects. The variable M_{NAPL} is the NAPL volume accumulated in the bottom 10 cm of the test cell. C_{steam} and C_{air} are the costs to generate high-enthalpy steam and to inject air, respectively. Finally, the coefficients a and b are used to appropriately weigh the different contributions to the total cost function.

The weighting coefficients a and b must (and can) be determined in any real application by means of an economic model, where monetary values are assigned not only to actual costs but also to potential benefits or damages. In case they remain uncertain, a sensitivity study can be performed to examine the robustness of the design to changes in the assumptions of the economic model. In the current demonstration, no effort has been made to find a realistic basis for the choice of these coefficients. The values used are: $a = 10$ cost units per kg of TCE left in the test cell, $b = 500$ cost units per liter of NAPL encountering the bottom of the test cell, $C_{steam} = 0.0014$ cost units for increasing specific steam enthalpy by 1 kJ, and $C_{air} = 200$ cost units for increasing the air injection rate by 1 g/s.

The two design parameters adjusted are the enthalpy of the injected steam, e_{steam} , and the air injection rate, q_{air} . Note that specific steam enthalpies of 2677 and 2875 kJ kg⁻¹ refer to injection of pure (dry) steam at 100 °C and 200 °C, respectively. Enthalpies between 419 and 2677 kJ kg⁻¹ indicate that a certain fraction of the injected water is in liquid form (wet steam) at a temperature of 100 °C. The enthalpy is therefore a measure of the so-called steam quality. In these simulations, the steam injection rate is held constant at 0.8 g s⁻¹.

We examined the cost function in the two-dimensional parameter space in the ranges $0 < q_{air} < 0.1$ g s⁻¹ and $2200 < e_{steam} < 2875$ kJ kg⁻¹. Figure 8 shows contour plots of the cost function and its individual contributions. The simple linear cost model for the production of high-enthalpy steam and for injecting air is shown in Figure 8a. In reality, this cost function is likely to be nonlinear and may even be discontinuous (for example., if a second pump must be purchased to reach a higher injection rate). Figure 8b shows the term that penalizes downward migration of NAPL. The plot shows that injection of air at a rate higher than about 0.25 g s⁻¹ successfully prevents any significant recondensation of TCE at the steam front that bypasses the fine sand lens near the bottom of the test cell. Increasing steam enthalpy leads to higher NAPL volatilization and thus requires slightly higher air injection rates to avoid NAPL accumulation at the bottom. Figure 8c shows that more TCE is removed from the test cell with increasing steam enthalpy and increasing air injection rate. The relationship is almost linear. Finally, Figure 8d shows the total cost function (see Equation 2).

The combination of all terms leads to a low-cost region in the parameter space near $e_{steam} = 2550$ kJ kg⁻¹ and $q_{air} = 0.35$ g s⁻¹. This minimum is surrounded by regions where steam enthalpy is too low for efficient NAPL volatilization, air injection rate is too low to prevent unwanted recondensation, and enthalpy and injection rate are too high, producing generation costs that cannot be compensated by a corresponding increase in remediation efficiency. It is obvious that the location of this minimum depends on the form of the cost function and the weighting coefficients.

The contour plot shown in Figure 8d (generated using the grid search option of iTOUGH2) provides complete information about the optimization problem at hand. However, grid search is usually prohibitive if more than two parameters are to be determined. We therefore examined the performance of the Levenberg-Marquardt algorithm, which calculates the cost function (and its gradient) only at a few discrete points in the parameter space. Starting from three different points in the parameter space, the minimum was identified within a few iterations, requiring significantly fewer simulation runs than the grid search method.

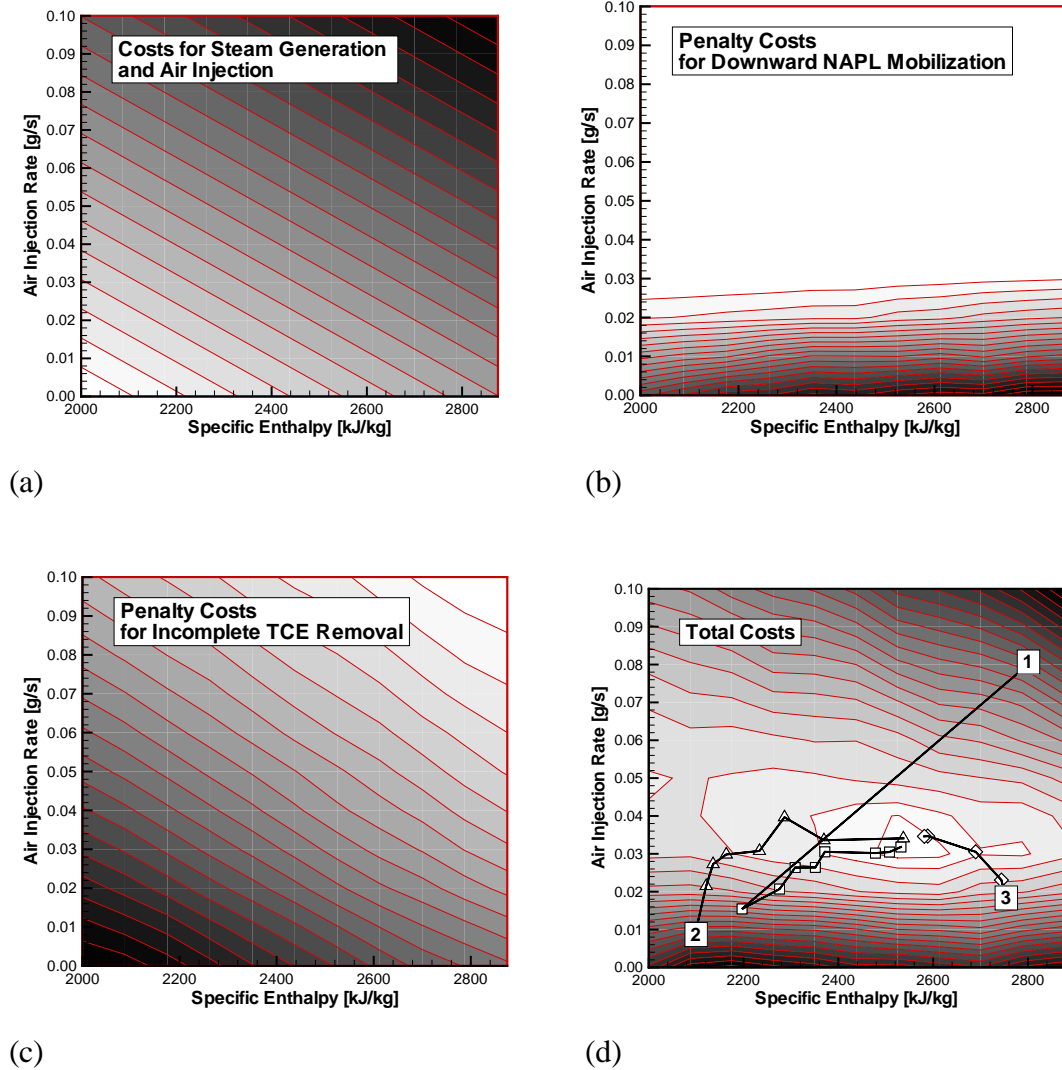


Figure 8. Contour plots of cost terms as a function of steam enthalpy and air injection rate. The optimal design parameters minimize the total costs. The minimum was identified by the Levenberg-Marquardt algorithm within a few iterations, regardless of the starting point (indicated by numerals).

6. Summary and Conclusions

We examined the potential use of standard optimization algorithms for the solution of groundwater management problems. A cost function measures the effectiveness and efficiency of a proposed remediation scenario. The cost function is minimized by automatically adjusting certain operational parameters (such as pumping rates). The impact of these operational parameters on remediation is evaluated using a sophisticated three-phase, three-component flow and transport simulator (T2VOC). Changes in operational parameters are propagated through the process model, affecting the fate of contaminants and eventually the cost function. The minimum of the cost function is identified by iTOUGH2 using strategies that are characteristic of the chosen minimization algorithm.

We demonstrated that the minimization algorithms implemented in iTOUGH2 are able to determine optimal remediation parameters. The combination of simulation and optimization techniques provides the capability to simulate the complex processes governing the remediation of NAPL-contaminated aquifers, and the flexibility to specify and minimize arbitrarily complex cost functions.

The robustness of the optimal solution remains to be investigated. Specifically, the amount of characterization data required to perform reliable predictions must be determined. Uncertainties in hydrologic parameters as well as incomplete knowledge about the source and distribution of the contaminants may strongly affect the predicted performance of a proposed cleanup operation and thus affects the accuracy of the calculated cost function. Note, however, that these difficulties are common to all methods that try to optimize aquifer remediation

Acknowledgment

This work was supported by the U.S. Department of Energy, Office of Environmental Management, Office of Science and Technology, under Contract Number DE-AC03-76SF00098. The author would like to thank C.A. Doughty and C.M. Oldenburg (LBNL) for their reviews of the manuscript.

References

- Betz, C., A. Färber, R. Schmidt, H. Kobus, and K. Pruess, Verfahren zur thermisch-unterstützten in-situ Reinigung von Böden in der ungesättigten Zone ohne Verfrachtung flüssigen Schadstoffes in tiefer gelegene Bodenbereiche, Patent Nr. DE 197 07 096, Munich, Germany, 1998.
- Betz, C., A. Färber, C.M. Green, H.-P. Koschitzky, and R. Schmid, Removing volatile and semi-volatile contaminants from the unsaturated zone by injection of steam/air-mixture, *Proceedings: Contaminated Soil '98*, 6th International FZK/TNO Conference, May 17–21, Edinburgh, UK, 575–584, 1998.
- Deutsch, C.V., and A.G. Journel, *GSLIB: Geostatistical Software Library and User's Guide*, Oxford Univ. Press, New York, 1992.
- Falta, R.W., K. Pruess, I.Javandel, and P.A. Witherspoon, Numerical modeling of steam injection for the removal of nonaqueous phase liquids from the subsurface, *Water Resour. Res.*, 28(2), 433–465, 1992.
- Falta, R.W., K. Pruess, S. Finsterle, and A. Battistelli, *T2VOC User's Guide*, Report LBL-36400, Lawrence Berkeley Laboratory, Berkeley, Calif., 1995.
- Finsterle, S., *Parallelization of iTOUGH2 Using PVM*, Report LBNL-42261, Lawrence Berkeley National Laboratory, Berkeley, Calif., 1998.
- Finsterle, S., *iTOUGH2 User's Guide*, Report LBNL-40040, Lawrence Berkeley National Laboratory, Berkeley, Calif., 1999a (see also <http://www-esd.lbl.gov/iTOUGH2>).
- Finsterle, S., *iTOUGH2 Command Reference*, Report LBNL-40041 (updated), Lawrence Berkeley National Laboratory, Berkeley, Calif., 1999b.
- Finsterle, S., *iTOUGH2 Sample Problems*, Report LBNL-40042 (updated), Lawrence Berkeley National Laboratory, Berkeley, Calif., 1999c.
- Geist, A., A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam, *PVM: Parallel Virtual Machine—A User's Guide and Tutorial for Networked Parallel Computing*, MIT Press, Cambridge, 1994.
- Gill, P.E., W. Murrey, and M.H. Wright, *Practical Optimization*, Academic Press, Inc., London, 1981.
- International Formulation Committee, *A Formulation of the Thermodynamic Properties of Ordinary Water Substance*, IFC Secretariat, Düsseldorf, Germany, 1967.
- Moridis, G.J., and K. Pruess, *T2CG1: A Package of Preconditioned Conjugate Gradient Solvers for the TOUGH2 Family of Codes*, Report LBL-36235, Lawrence Berkeley Laboratory, Berkeley, Calif., 1995.
- Narasimhan, T.N. and P.A. Witherspoon, An integrated finite difference method for analyzing fluid flow in porous media, *Water Resour. Res.*, 12(1), 57–64, 1976.
- Parker, J.C., R.J. Lenhard, and T. Kuppusamy, A parametric model for constitutive properties governing multiphase flow in porous media, *Water Resour. Res.*, 23(4), 618–624, 1987.
- Pruess, K., *TOUGH2—A General-Purpose Numerical Simulator for Multiphase Fluid and Heat Flow*, Report LBL-29400, Lawrence Berkeley Laboratory, Berkeley, Calif., 1991.

- Pruess, K., C. Oldenburg, and G. Moridis, *TOUGH2 User's Guide, Version 2.0*, Report LBNL-43134, Lawrence Berkeley National Laboratory, Berkeley, Calif., 1999.
- Reid, R.C., J.M. Prausnitz, and B.E. Poling, *The Properties of Gases and Liquids*, McGraw-Hill, New York, 1987.
- Stone, H.L., Probability model for estimating three-phase relative permeability, *Trans. SPE of AIME*, 239,214-218, 1970.