Demonstration of optimization techniques for groundwater plume remediation using iTOUGH2

Stefan Finsterle[*]

Earth Sciences Division, Lawrence Berkeley National Laboratory, CA 94720, USA

Abstract

We examined the potential use of standard optimization algorithms as implemented in the inverse modeling code iTOUGH2 (Finsterle, 1999abc) 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. The cost function is then minimized by automatically adjusting certain decision or operational parameters. We evaluate the impact of these operational parameters on remediation using a three-phase, three-component flow and transport simulator, which is linked to nonlinear optimization routines. We demonstrate that the methods developed for automatic model calibration are capable of minimizing arbitrary cost functions. An example of co-injection of air and steam makes 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.

Keywords: Optimization, remediation, multiphase flow, contamination, iTOUGH2

---

[*] Tel.: +510-486-5205; fax: +510-486-5686; e-mail address: SAFinsterle@lbl.gov (Stefan Finsterle)

1. Introduction

The design of a cleanup operation for a contaminated aquifer involves environmental, hydrological, technical, and economic issues (as well as legal and institutional aspects and demands from pressure groups) (U.S. Environmental Protection Agency, 1997). 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.

Optimization of groundwater remediation activities 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.

Simulation and optimization techniques have been combined to solve single- and multi-objective subsurface remediation management problems (see Freeze and Gorelick (1999) for a review). Remediation strategies are usually optimized by minimizing a cost function, in which multiple objectives are expressed in actual or equivalent monetary units. In this approach, the cost function to be minimized contains appropriately weighted terms representing, for example, drilling costs as a function of the number of wells, treatment costs as a function of pumping rates, capital costs as a function of cleanup time, and penalty terms to enforce remediation standards. This combined cost function is then minimized using linear or nonlinear single-objective optimization methods (see Culver and Shoemaker (1992) for a review). An alternative to optimizing a composite criterion is to determine a set of "Pareto optimal" designs. A remediation strategy is considered Pareto optimal if it dominates all other designs, i.e., if no part of the solution can be improved without making another part worse than a competing design. Consideration of the Pareto optimal set (or Pareto frontier) provides a better understanding of the trade-offs necessary to obtain greater relative efficiency on a given objective. An example of this approach is given in Erickson et al. (2002).

The purpose of this study is to examine whether the combination of a nonisothermal multiphase flow and transport 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.

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, see also http://www-esd.lbl.gov/iTOUGH2). iTOUGH2 provides inverse modeling capabilities for the non-isothermal, 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 review of iTOUGH2 applications can be found in Finsterle (2004). The FORTRAN source code of iTOUGH2 is distributed through the Energy Science and Technology Software Center (http://www.osti.gov/estsc/) and (for customers outside the U.S.) the Nuclear Energy Agency (http://www.nea.fr/html/dbprog/). Machine-dependent functions are provided for compilation and installation on various platforms (PC, Unix, Linux, supercomputers). Individual forward simulations can be executed in parallel on a heterogeneous cluster of workstations using Parallel Virtual Machine (PVM) (Geist et al., 1994; Finsterle, 1998).

In principle, any forward model can be linked to a general, model-independent, nonlinear parameter estimation package such as PEST (Doherty, 1994; http://www.sspa.com/pest/) or UCODE (Poeter and Hill, 1998; http://water.usgs.gov/software/ucode.html), thus providing optimization capabilities with various degrees of flexibility and efficiency.

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 extrac-

tion, 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 (Falta et al., 1995). 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, $Z = \Sigma z$. 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. (An exception to this rule is implemented in the Simulated Annealing optimization method, which allows uphill steps with a generally low and decreasing probability as a means to escape from local minima.) 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 and provides the reference to the original publication. All methods are also described in Press et al. (1992) and Finsterle (1999a).

The Gauss-Newton algorithm approximates the curvature of the cost function by a positive-definite Hessian matrix and then performs a step to its minimum. This means that if the cost function is quadratic, the Gauss-Newton algorithm identifies the optimum in a single step. For nonlinear models or non-quadratic cost functions, the algorithm (even if applied iteratively) may diverge. This problem is partly addressed by the Levenberg-Marquardt modification of the Gauss-Newton algorithm. After local linearization of the cost function with respect to the operational parameters, the Levenberg-Marquardt algorithm performs initially small, but robust steps along the steepest descent direction, and switches to more efficient quadratic Gauss-Newton steps as the minimum is approached. Both the Gauss-Newton and Levenberg-Marquardt algorithms require the calculation of first derivatives, which are evaluated numerically in iTOUGH2 using the perturbation method. This evaluation can be performed in parallel (Finsterle, 1998).

The downhill simplex method requires only function evaluations (i.e., no derivatives) and is therefore a robust but relatively inefficient minimization method. Starting with a simplex consisting of $n+1$ points in the n-dimensional parameter space, a series of steps is taken, most of which consist of moving the point of the simplex with the highest cost function through the opposite face of the simplex to a lower point. Other search directions are generated by reflection, expansion, and contraction of the simplex from the

8

previous step. The simplex algorithm should be used in iTOUGH2 when minimizing discontinuous cost functions or if the numerical evaluation of the derivatives is unstable.

A continuous version of the method of Simulated Annealing has been implemented in iTOUGH2. Simulated Annealing is a technique to find the (ideally global) minimum of the cost function in the presence of many local minima. Random steps in the parameter space are performed. A step is always accepted if the cost function is lowered, and it is sometimes accepted with a certain, decreasing probability, if an uphill step is taken. This scheme allows the algorithm to escape from local minima. Simulate Annealing requires many forward runs and is thus only applicable to small problems.

The global minimum can always be identified by systematically evaluating the cost function in the entire parameter space. Moreover, contouring the cost function reveals the potential presence of local minima, non-uniqueness problems, the correlation structure, stability problems, and nonlinearity effects. Evaluating the cost function on a grid in the entire parameter space is prohibitively expensive for higher-dimensional parameter spaces. This option is used for illustrative purposes (see Figure 8 below).

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

9

message passing based on the Parallel Virtual Machine concept (PVM; Geist et al., 1994); for details, see Finsterle (1998).

Optimization problems are susceptible to being ill-posed, which leads to nonunique or unstable solutions (Carrera and Neuman, 1986; Yeh, 1986; McLaughlin and Townley, 1996). 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 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 as well as the presence of local minima 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—in addition to the definition of a meaningful cost function—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 estimation uncertainty. In summary, test design, monitoring, inverse modeling, and remediation design are strongly related and pose an overall optimization problem. Specifically, a unified approach would need to include into the cost function the substantial costs for data collection, which will then need to be optimized to arrive at an acceptable model application reliability. An outline

of such an overall approach is given, for example, by Sun and Yeh (1990) and Wagner (1995). The incremental worth of data and the optimization of site characterization strategies is discussed in Freeze et al. (1992), James and Freeze (1993), and James and Gorelick (1994). The scope of this study is limited to the computational aspects of finding cost-effective design variables for a given cleanup operation and a given model.

The examples discussed in Section 4 below illustrate certain aspects of the simulation-optimization capabilities of iTOUGH2. The first example mainly demonstrates the performance of the code's different minimization algorithms for the solution of a standard remediation problem. It also shows the flexibility in formulating the optimization problem. For example, the (unknown) time for contaminant removal from an aquifer is part of the cost function, i.e., the simulation time is automatically adapted to the imposed remediation target, which in turn depends on the operational parameters to be optimized. The second example highlights the process simulation capabilities of the general-purpose reservoir simulator T2VOC, and how subtle physical processes can affect the optimal solution. Moreover, this case demonstrates that unconventional terms can be included in the cost function.

In iTOUGH2, the user has complete control and flexibility in specifying cost functions of arbitrary complexity. An interface subroutine is provided in which the user can program the cost as a function of any T2VOC input parameter and output variable. Similarly, an interface is provided for the user to specify non-standard operational parameters. Given the flexibility provided by these interfaces, the iTOUGH2 code (which was designed for the solution of parameter estimation problems) can be used without

modification (except for the user-specified interface functions) for the solution of cost minimization problems.

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 based on a spherical semi-variogram. The log-permeability field shown in Figure 1 was generated using Sequential Indicator Simulations (Deutsch and Journel, 1992). The effective porosity is assumed to have a constant value of 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 m $\times$ 5 m $\times$ 10 m.

We simulated a potential contamination of the aquifer by releasing 100 kg of trichloroethylene (TCE) at X = 47.5 m and Y = 67.5 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. Note that these initial conditions do not imply that TCE dissolves instantaneously. However, an implicit assumption is made that TCE has spread uniformly over the size of a single gridblock during the early stages of the contamination. The initial distribution used in the model reflects the uncertainty in location and release history of the contaminant.

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 site remediation 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

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

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. In the simulation, the cleanup time is calculated by linear interpolation between the two calculation times that bracket the remediation goal of 95% TCE removal. The width of this time bracket depends on the time step size, which in T2VOC is variable and automatically adjusted based on the convergence behavior of the Newton-Raphson iteration. Optimization of this cost function requires the calculation of the cleanup time as a function of pumping rates, either as a simple evaluation or as part of the calculation of a sensitivity coefficient. Calculation of derivatives is performed numerically by perturbing the pumping rates by 1% of their respective values.

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,

15

whereby a well should not be drilled if the pumping rate is less than a certain threshold value, (2) concentration-dependent treatment costs, and (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 depend strongly 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, 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.

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 consid-

16

erably 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.

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.

The computational efficiency of a minimization algorithm is determined by the number of forward simulations required; the overhead needed to calculate orientation and

17

step size to update the parameter vector is negligible. In this illustrative example, a single simulation of contaminant migration and removal takes on average approximately 9 CPU seconds (note that cleanup time and thus the CPU time depend on the parameter set being examined for each simulation). As indicated in Figure 3, the Gauss-Newton algorithm required about 2 minutes for the 15 simulations, whereas it took approximately 3 hours to arrive at the optimal solution proposed by Simulated Annealing.

Figure 4 shows the amount of dissolved TCE removed from the aquifer as a function of pumped groundwater volume for the initial design and for the optimum design determined by Simulated Annealing. The total amount of TCE in the aquifer (which is a direct output from the T2VOC simulator) is plotted as a function of the pumped volume, which is the total pumping rate of all six wells multiplied by time. Curves are produced for the initial and optimized designs with pumping rates of 6 kg/s and 2.43 kg/s, respectively (see Table 2). 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. Recall that the pumping schedule has been optimized for a target residual TCE mass of 5 kg. If optimized for a much smaller residual mass, the optimal pumping schedule would focus on the removal of a diffuse residual plume at late time rather than on the removal of bulk contaminants at early and intermediate times. This is reflected in the changing width of the gap between the two curves, indicating that a different optimal solution would be obtained if the remediation criterion were changed. Nevertheless, it appears that the automatically determined pumping schedule remains

superior to the initial design even if the threshold of acceptable residual TCE is reduced. Note that the solution is only optimal for the given conceptual model. Uncertainties in this model were discussed above; the solution to alternative models of similar occurrence probability would have to be evaluated to arrive at a robust optimal solution.

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, 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 amount of 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. Note that the total amount of contaminants escaping across the downstream model boundary is insignificant and does not affect the results.

This first example illustrates that the optimization algorithms implemented in iTOUGH2 can be used to improve remediation designs by automatically minimizing a

19

cost function. While the cost function chosen here is extremely simple, its complexity can be increased arbitrarily. The properties of the cost function may determine the choice of the minimization algorithm. For example, if the cost function is discontinuous (e.g., by containing capital cost terms for a well that may or may not be drilled), the derivative-based algorithms cannot be used; the Downhill Simplex algorithm was found to be a good compromise between efficiency and robustness for the minimization of most cost functions.

A key feature of this approach is that the costs are calculated based on a 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. Since the permeability structure of the aquifer or the initial contaminant distribution is uncertain, the optimal pumping schedule will not be as successful when implemented in the field.

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

The widely used soil vapor extraction method (U.S. Environmental Protection Agency, 1991) 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 (Looney and Falta, 2000). 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 noncondensible 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 m $\times$ 0.1 m $\times$ 0.75 m represents a vertical test cell filled with coarse sand with a lens of finer sand built into it (solid box). The domain is uniformly discretized into 1,200 grid blocks of dimensions 0.025 m $\times$ 0.1 m $\times$ 0.025 m. A spatially correlated, anisotropic permeability field is generated to induce heterogeneity. A spherical semi-variogram was used with horizontal and vertical correlation lengths of 0.2 and 0.06 m, respectively, and a sill value of 0.5 log cycles. 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 $X = 0.875$ m to a depth of $Z = $ -0.5 m. Input parameters are summarized in Table 3.

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, recondensa-

tion 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 downward 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.

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

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

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 refer to injection of pure (dry) steam at 100 °C and 200 °C, respectively. Enthalpies between 419 and 2677 kJ/kg 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.

We examined the cost function in the two-dimensional parameter space in the ranges $0 < q_{air} < 0.1$ g/s and $2200 < e_{steam} < 2875$ kJ/kg. 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 0.02 g/s successfully prevents any significant recondensation of

TCE at the steam front that bypasses the fine sand lens near the bottom of the test cell. Figure 8b also demonstrates that injecting even small amounts of air significantly reduces the costs related to the NAPL mobilization risk, outweighing the costs associated with air injection (see Figure 8a). Increasing steam enthalpy leads to higher NAPL volatilization and thus requires slightly higher air injection rates to avoid NAPL accumulation at the bottom. In other words, the enthalpy does not affect downward migration significantly, so poor steam quality would be sufficient to satisfy this partial performance criterion. However, high steam enthalpy improves overall remediation (see discussion of Figure 8c) and is thus desirable despite higher steam production costs (Figure 8a). 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 and q_{air} = 0.035 g/s. 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 8 iterations (requiring 27 simulations), i.e., substantially fewer than the 400 simulation runs performed with the grid search method.

Co-injection of air and steam for the enhanced removal of NAPL highlights both the usefulness of simulation-optimization techniques as well as the related challenges. As demonstrated in the example, subtle but complex physical effects (such as multiphase flow and transport associated with vaporization and recondensation) greatly impact the remediation success. Understanding and predicting the system behavior requires a sophisticated numerical simulator to increase the fundamental understanding of the effects governing co-injecting projects (for example, through laboratory experiments as described above) and to predict and optimize site-specific field applications. For the latter, a substantial characterization effort is required for these predictions to be reliable. Site-specific multiphase flow parameters, thermal properties, and information about multiscale heterogeneity need to be obtained. iTOUGH2 supports these multiphase parameter estimation problems (see Finsterle (2004) for an overview). Finally, a cost function needs to be developed that captures and quantifies the salient ecological and economic factors comprising the remediation goal. iTOUGH2 is a tool that addresses the first two issues (site characterization and predictive simulation), and it allows the user to examine and optimize arbitrary cost functions. Balancing the effort spent on each of these related areas (characterization and monitoring, hydrogeologic modeling, economic modeling, and optimization) remains the subject of an overall management and optimization problem.

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.

Two illustrative examples are provided. The purpose of the first example was to examine the relative performance of the minimization algorithms implemented in iTOUGH2 for a standard pump-and-treat cleanup operation. The total pumping cost to be minimized (see Equation (1)) was calculated as the product of the pumping rates and the required cleanup time to a certain standard (e.g., 95% contaminant removal). The decision variables are the pumping rates in six wells; they were constrained to nonpositive values (i.e., injection was not allowed as part of the solution). The purpose of the second example was to demonstrate that subtle physical effects determine the optimal solution in a thermally enhanced soil vapor extraction scheme with the co-injection of air. The cost function (see Equation (2)) included terms for incomplete remediation, a penalty term for NAPL mobilization, and costs for steam generation and air injection. The design variables are the steam enthalpy and air injection rate; they were contained to positive values.

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. While the illustrative examples discussed here used oversimplified cost functions and process models, the experience with iTOUGH2 for the solution of complex parameter estimation problems (Finsterle, 2004) provides confidence that the code is flexible enough and capable of optimizing realistic groundwater management problems. Note, however, that a reasonable solution to an optimization problem requires a good understanding of the underlying physics and the cost function, as well as a critical examination of the caveats of the problem formulation. Finally, computational limitations must be considered.

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. Stochastic methods may be a means to quantify the remaining uncertainties and to assess the related risks.

Acknowledgment

References

Betz, C., A. Färber, R. Schmidt, H. Kobus, and K. Pruess, 1998. 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.

Betz, C., A. Färber, C.M. Green, H.-P. Koschitzky, and R. Schmid, 1998. Removing volatile and semi-volatile contaminants from the unsaturated zone by injection of steam/air-mixture. In: Proceedings: Contaminated Soil '98, 6th International FZK/TNO Conference. May 17–21, Edinburgh, UK, 575–584.

Carrera, J., and S.P. Neuman, 1986. Estimation of aquifer parameters under transient and steady state conditions: 1. Uniqueness, stability, and solution algorithms. Water Resour. Res. 22 (2) 211–227.

Culver, T.B., and C.A. Shoemaker, 1992. Dynamic optimal control for groundwater remediation with flexible management periods. Water Resour. Res. 28 (3) 629–641.

Deutsch, C.V., and A.G. Journel, 1992. GSLIB: Geostatistical Software Library and User's Guide. Oxford Univ. Press: New York.

Doherty, J. 1994. PEST: Model-Independent Parameter Estimation. Watermark Numerical Computing: Australia.

Erickson, M., A. Mayer, and J. Horn, 2002. Multi-objective optimal design of groundwater remediation systems: Application of the niched Pareto genetic algorithm (NPGA). Adv. Water. Resour. 25 51–65.

Falta, R.W., K. Pruess, I. Javandel, and P.A. Witherspoon, 1992. Numerical modeling of steam injection for the removal of nonaqeous phase liquids from the subsurface. Water Resour. Res. 28 (2) 433–465.

Falta, R.W., K. Pruess, S. Finsterle, and A. Battistelli, 1995. T2VOC User's Guide. Report LBL-36400, Lawrence Berkeley Laboratory: Berkeley, Calif.

Finsterle, S., Parallelization of iTOUGH2 Using PVM, 1998. Report LBNL-42261, Lawrence Berkeley National Laboratory: Berkeley, Calif.

Finsterle, S., 1999a. iTOUGH2 User's Guide. Report LBNL-40040, Lawrence Berkeley National Laboratory: Berkeley, Calif., (see also http://www-esd.lbl.gov/iTOUGH2).

Finsterle, S., 1999b. iTOUGH2 Command Reference. Report LBNL-40041 (updated), Lawrence Berkeley National Laboratory: Berkeley, Calif.

Finsterle, S., 1999c. iTOUGH2 Sample Problems. Report LBNL-40042 (updated), Lawrence Berkeley National Laboratory: Berkeley, Calif.

Finsterle, S., 2004. Multiphase inversion modeling: Review and iTOUGH2 applications. Vadose Zone J. 3 747–762.

Freeze, R.A., B.R. James, J. Massman, T. Sperling and L. Smith, 1992. Hydrogeological decision analysis: 4, The concept of data worth and its use in the development of site investigation strategies Ground Water 30 (4) 574–588.

Freeze, R.A., and S.M. Gorelick, 1999. Convergence of stochastic optimization and decision analysis in the engineering design of aquifer remediation. Ground Water 37 (6) 934–954.

Gauss, C.F., 1821. Theoria Combinationis Observationum Erroribus Minimis Obnoxiae. In: G.W. Stewart (trans.), 1995. Theory of the Combination of Observations Least Subject to Errors, SIAM: Philadelphia, PA.

Geist, A., A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam, 1994. PVM: Parallel Virtual Machine—A User's Guide and Tutorial for Networked Parallel Computing. MIT Press: Cambridge.

Gill, P.E., W. Murrey, and M.H. Wright, 1981. Practical Optimization. Academic Press, Inc.: London.

International Formulation Committee, 1967. A Formulation of the Thermodynamic Properties of Ordinary Water Substance. IFC Secretariat: Düsseldorf, Germany.

James, B.R. and R.A. Freeze, 1993. The worth of data in predicting aquitard continuity in hydrogeological design. Water Resour. Res. 29 (7) 2049–2065.

James, B.R. and S.M. Gorelick, 1994. When enough is enough: The worth of monitoring data in aquifer remediation design. Water Resour. Res. 30 (12) 3499–3513.

Looney, B.B., and R.W. Falta (eds.), 2000. Vadose Zone Science and Technology Solutions. Battelle Press: Columbus, OH.

Marquardt, D.W., 1963, An algorithm for least squares estimation of nonlinear parameters. SIAM J. Appl. Math. 11 431–441.

McLaughlin, D., and L.R. Townley, 1996. A reassessment of the groundwater inverse problem. Water Resour. Res. 32 (5) 1131–1161.

Metrolpolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. 1953. Equation of state calculations by fast computing machines. J. Chem. Phys. 21 1087–1092.

Moridis, G.J., and K. Pruess, 1995. T2CG1: A Package of Preconditioned Conjugate Gradient Solvers for the TOUGH2 Family of Codes. Report LBL-36235, Lawrence Berkeley Laboratory: Berkeley, Calif.

Narasimhan, T.N. and P.A. Witherspoon, 1976. An integrated finite difference method for analyzing fluid flow in porous media. Water Resour. Res. 12 (1) 57–64.

Nelder, J.A., and R. Mead, 1965. A simplex method for function minimization. Comput. J. 7 307–313.

Parker, J.C., R.J. Lenhard, and T. Kuppusamy, 1987. A parametric model for constitutive properties governing multiphase flow in porous media. Water Resour. Res. 23 (4) 618-624.

Poeter, E.P., and M.C. Hill. 1998. Documentation of UCODE, a Computer Code for Universal Inverse Modeling. U.S. Geological Survey Water-Resources Investigations Report 98-4080.

Press, W.H., S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, 1992. Numerical Recipes in FORTRAN: The Art of Scientific Computing. 2nd edition, Cambdridge University Press: Cambridge.

Pruess, K., 1991. TOUGH2—A General-Purpose Numerical Simulator for Multiphase Fluid and Heat Flow. Report LBL-29400, Lawrence Berkeley Laboratory: Berkeley, Calif.

Pruess, K., C. Oldenburg, and G. Moridis, 1999. TOUGH2 User's Guide, Version 2.0. Report LBNL-43134, Lawrence Berkeley National Laboratory: Berkeley, Calif.

Reid, R.C., J.M. Prausnitz, and B.E. Poling, 1987. The Properties of Gases and Liquids. McGraw-Hill: New York.

Stone, H.L., 1970. Probability model for estimating three-phase relative permeability. Trans. SPE of AIME 239 214-218.

Sun, N.-Z., and W W.-G. Yeh, 1990. Coupled inverse problems in groundwater modeling, 2. Identifiability amd experimental design. Water Resour. Res. 26 (10) 2527–2540.

U.S. Environmental Protection Agency, 1991. Soil Vapor Extraction Technology: Reference Handbook. Office of Research and Development, EPA/540/2-91/003: Cincinnati, OH.

U.S. Environmental Protection Agency, 1997. Cleaning Up the Nation's Waste Sites: Markets and Technology Trends, 1996 Edition, Executive Summary. Technology Innovation Office: Washington, D.C.

Wagner, B.J., 1995. Recent advances in simulation-optimization groundwater management modeling. Rev. of Geophysics, Supplement 1021–1028.

Yeh, W. W.-G, 1986. Review of parameter identification procedures in groundwater hydrology: The inverse problem. Water Resour. Res. 22 (2) 95–108.

Table 1. Minimization algorithms implemented in iTOUGH2

Minimization algorithm Reference	Description	Advantages and disadvantages
Gauss-Newton[a] (Gauss, 1821)	$\Delta\mathbf{p} = -\left(\mathbf{J}^T\mathbf{W}\mathbf{J}\right)^{-1}\mathbf{J}^T\mathbf{W}\mathbf{z}$	Efficient for linear least-squares problems only; requires derivatives.
Levenberg-Marquardt[a] (Marquardt, 1963)	$\Delta\mathbf{p} = -\left(\mathbf{J}^T\mathbf{W}\mathbf{J} + \lambda\mathbf{I}\right)^{-1}\mathbf{J}^T\mathbf{W}\mathbf{z}$	Efficient for nonlinear least-squares problem; requires derivatives.
Downhill Simplex (Nelder and Mead, 1965)	Approaches minimum through sequence of reflections, expansions, and contractions of an $n+1$ simplex.	No assumptions made about form of cost function; relatively inefficient.
Simulated Annealing (Metropolis et al., 1953)	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.

[a] **J** Jacobian matrix; **W**: Weighting matrix; **I**: Identity matrix; **z**: Residual/cost vector

Table 2. Optimized pumping schedules, cleanup times, and costs for Example 1

MA[a]	# of runs	Time [day]	Cost function (in thousands)	Pumping rate [kg/s]						
				Well 1	Well 2	Well 3	Well 4	Well 5	Well 6	Total
-	1	240	125	1.00	1.00	1.00	1.00	1.00	1.00	6.00
GN	15	340	79	-	1.25	0.63	0.13	1.11	0.89	4.01
LM	133	250	61	0.49	1.52	-	0.14	0.64	-	2.79
DS	89	190	59	0.21	1.95	0.27	-	1.13	-	3.56
SA	1224	250	53	0.15	1.20	-	-	0.91	0.17	2.43

[a] Minimization algorithm:

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

SA = Simulated Annealing

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 [J/(kg °C)]		840	840
Heat conductivity [W/(m °C)]	wet	1.60	1.60
	dry	0.35	0.35
Relative permeability: Stone (1970); Capillary pressure function: Parker et al. (1987)			
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
$\alpha_{gas\text{-}NAPL}$ [m^{-1}]		15.0	5.0
$\alpha_{NAPL\text{-}water}$ [m^{-1}]		150.0	50.0
Initial conditions			
Initial gas pressure [kPa]		101.3	101.3
Initial water saturation		0.1	0.1
Initial NAPL saturation in dash-dotted box		0.1	0.1
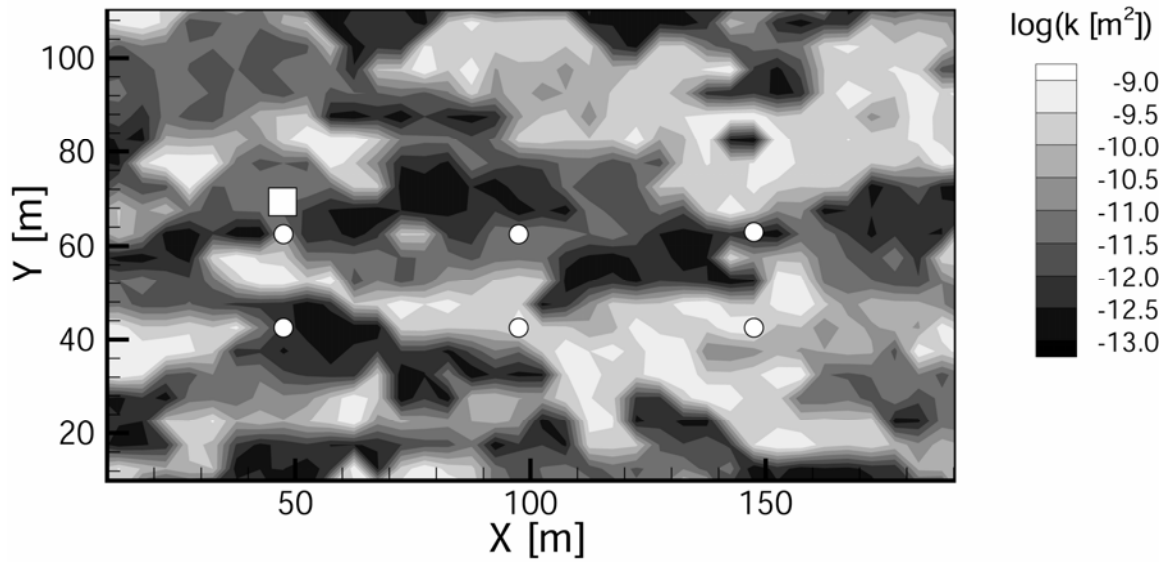
Figure 1. Log-permeability field of a synthetic, two-dimensional, confined aquifer. The square indicates the location where 100 kg of TCE are spilled. The locations of six remediation wells are also indicated.
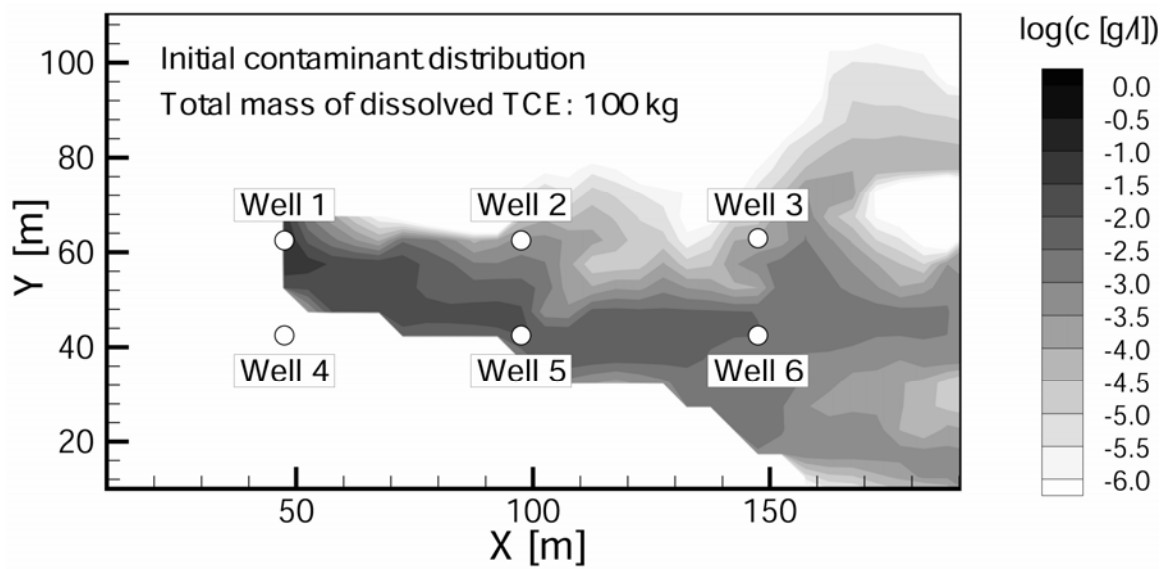
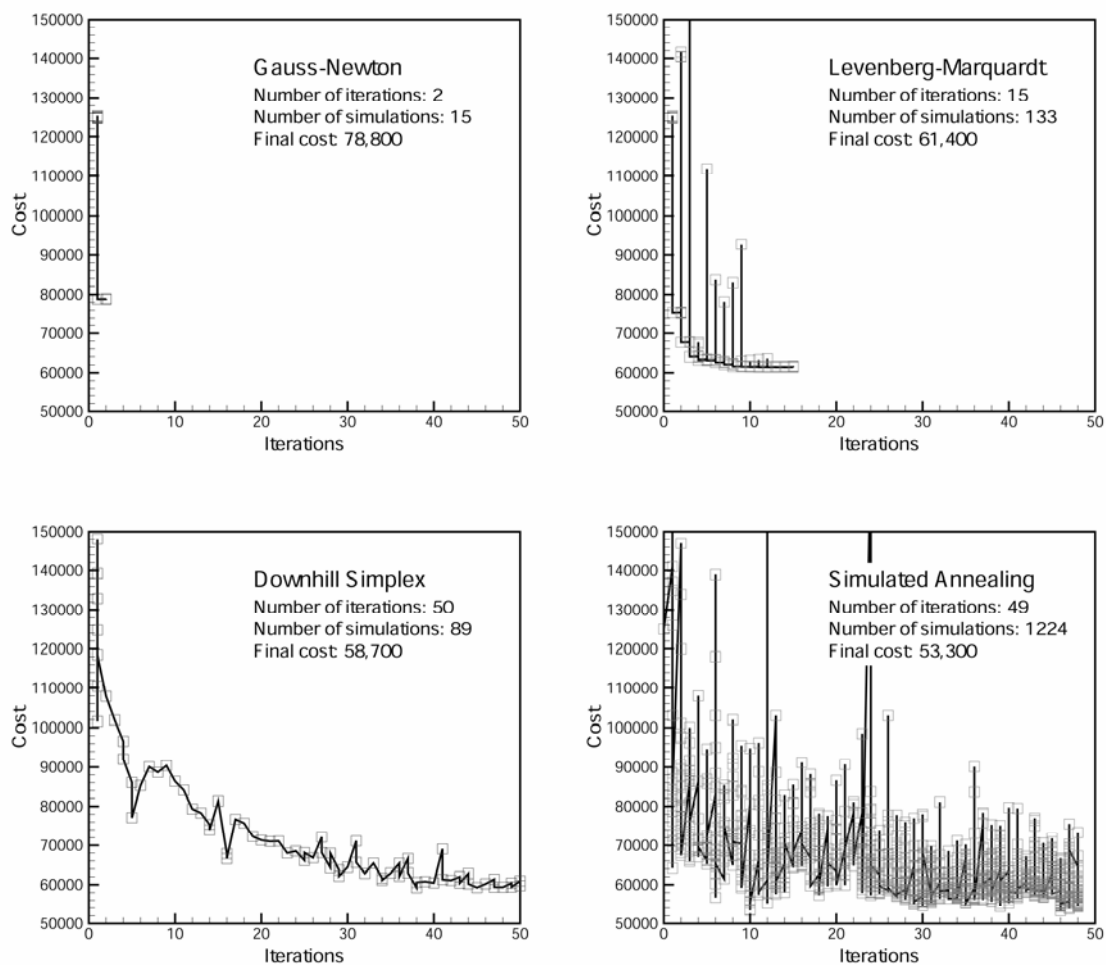**Figure 2.**   Distribution of dissolved TCE four months after the spill.

Figure 3. Reduction of cost function for four different minimization algorithms. Each square represents one T2VOC simulation.
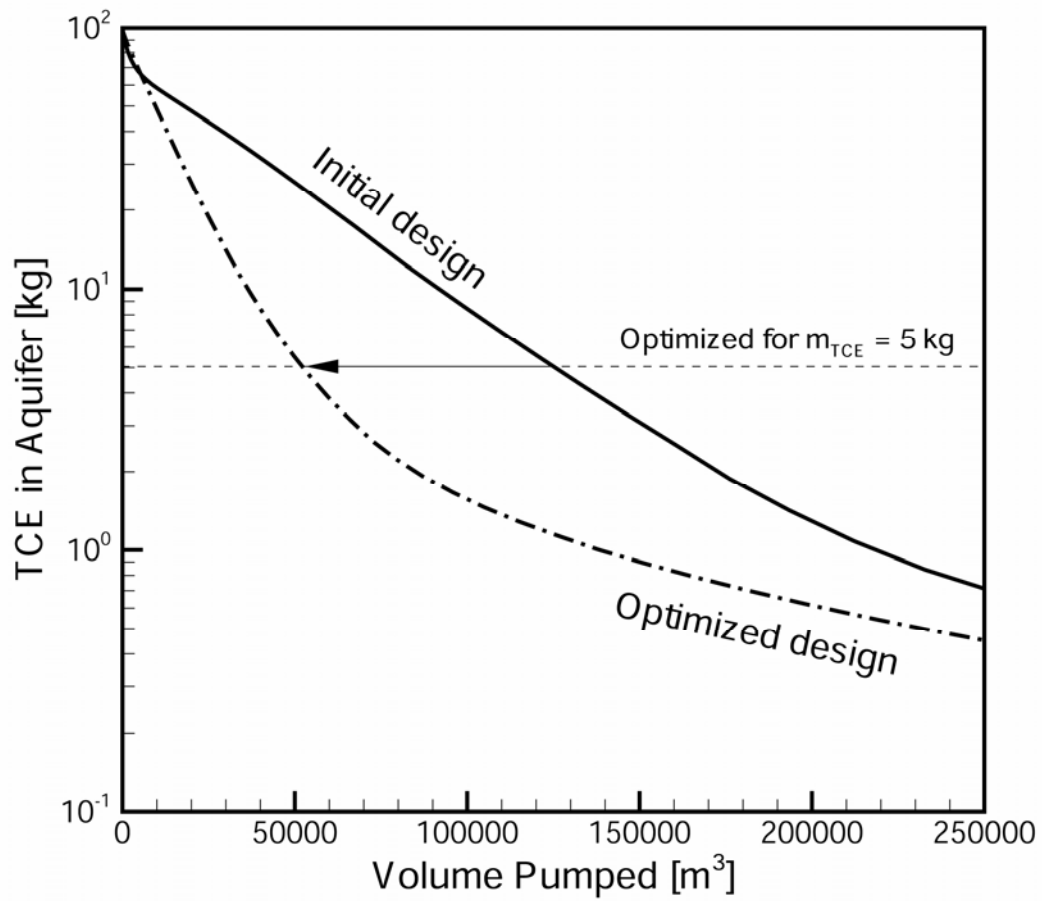
Figure 4. Mass of dissolved TCE left in the aquifer as a function of pumped volume for the initial and optimized designs. The pumping schedule was optimized to minimize the total pumped volume when 5 kg of TCE are left in the aquifer.
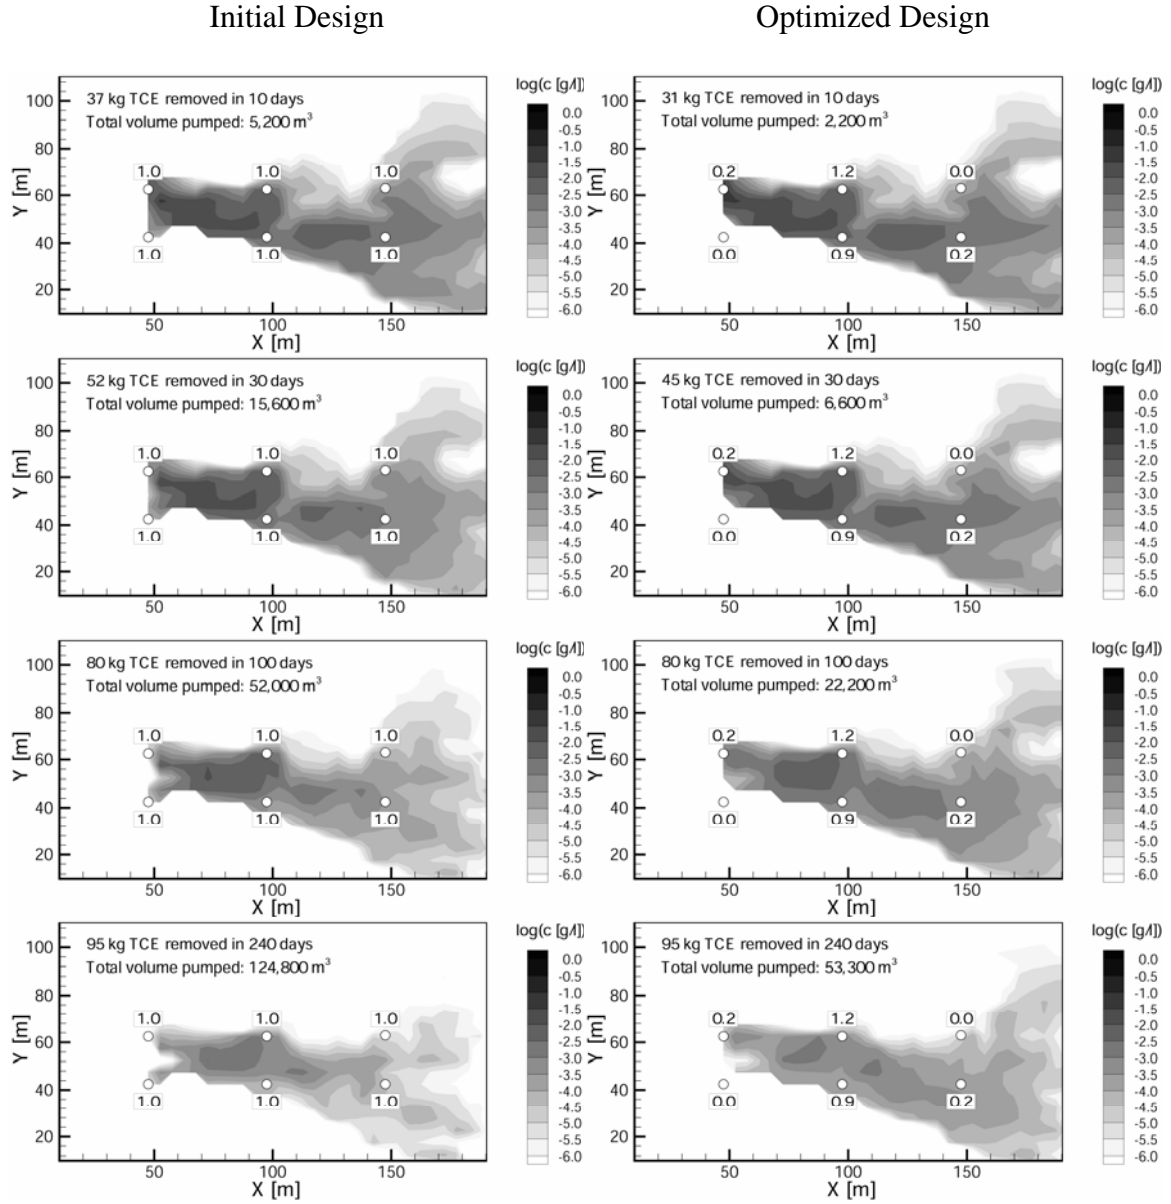
Initial Design   Optimized Design

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 in kg/s is indicated at each well.
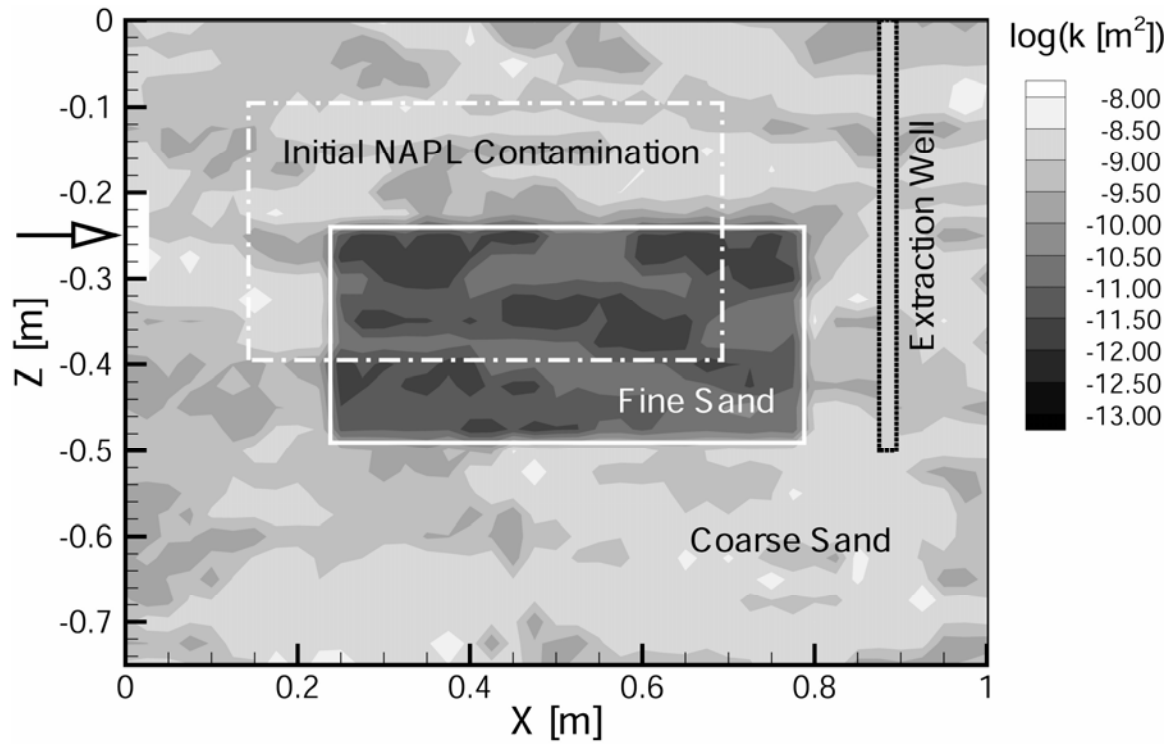
Figure 6. Layout of test cell experiment for studying thermally enhanced soil vapor extraction.

Steam Injection   Steam/Air Injection
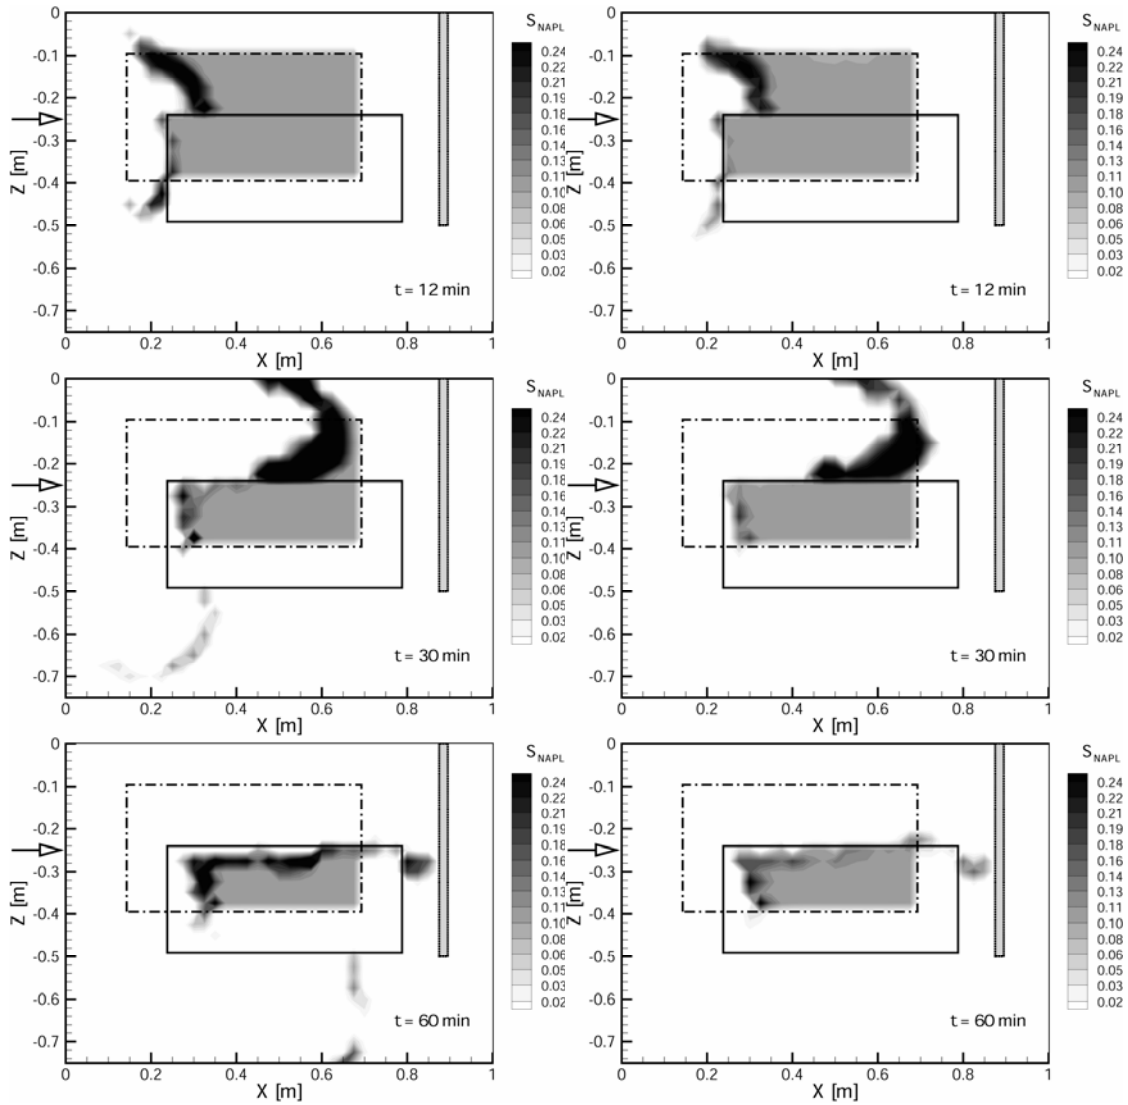

Figure 7. Distribution of free-phase TCE after 12, 30, and 60 minutes after injection of pure steam (left column) and a mixture of steam and air mixture (right column).
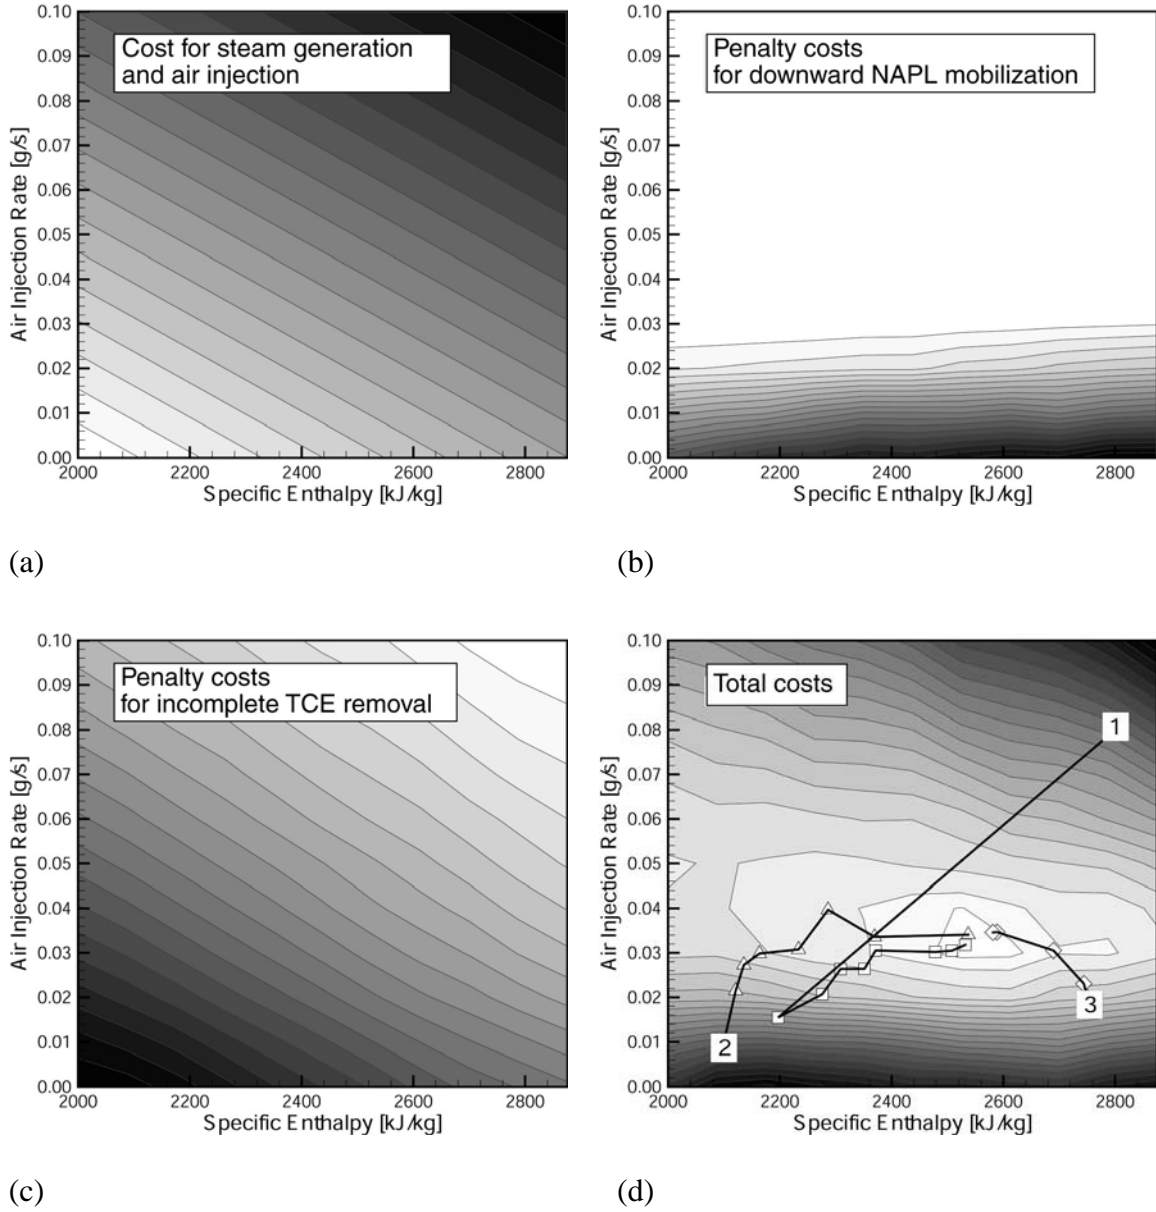
43

(a)   (b)

(c)   (d)

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).

44