

International Journal of High Performance Computing Applications

<http://hpc.sagepub.com/>

A Numerical Simulation of Groundwater Flow and Contaminant Transport on the CRAY T3D and C90 Supercomputers

S. F. Ashby, W. J. Bosl, R. D. Falgout, S. G. Smith, A. F.B. Tompson and T. J. Williams
International Journal of High Performance Computing Applications 1999 13: 80
DOI: 10.1177/109434209901300105

The online version of this article can be found at:
<http://hpc.sagepub.com/content/13/1/80>

Published by:



<http://www.sagepublications.com>

Additional services and information for *International Journal of High Performance Computing Applications* can be found at:

Email Alerts: <http://hpc.sagepub.com/cgi/alerts>

Subscriptions: <http://hpc.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://hpc.sagepub.com/content/13/1/80.refs.html>

A NUMERICAL SIMULATION OF GROUNDWATER FLOW AND CONTAMINANT TRANSPORT ON THE CRAY T3D AND C90 SUPERCOMPUTERS

S. F. Ashby
W. J. Bosl
R. D. Falgout
S. G. Smith
A.F.B. Tompson

LAWRENCE LIVERMORE NATIONAL LABORATORY,
CALIFORNIA, U.S.A.

T. J. Williams

LOS ALAMOS NATIONAL LABORATORY, LOS ALAMOS,
NEW MEXICO, U.S.A.

Summary

Numerical simulations of groundwater flow and chemical transport through three-dimensional heterogeneous porous media are described. The authors employ two CRAY supercomputers for different parts of the decoupled calculation: the flow field is computed on the T3D massively parallel computer, and the contaminant migration is simulated on the C90 vector supercomputer. The authors compare simulation results for subsurface models based on homogeneous and heterogeneous conceptual models and find that the heterogeneities have a profound impact on the character of contaminant migration.

Address reprint requests to S. F. Ashby, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, P.O. Box 808, L-561, Livermore, CA 94551.



The International Journal of High Performance Computing Applications,
Volume 13, No. 1, Spring 1999, pp. 80-93
© 1999 Sage Publications, Inc.

1 Introduction

Numerical simulations of groundwater flow and chemical transport in subsurface formations play an increasingly important role in the field of environmental engineering. In the context of contaminated soil and groundwater systems, simulations are typically used to (1) determine the basic flow and chemical migration patterns at contaminated sites; (2) analyze, develop, or otherwise optimize remediation strategies for the removal of contaminants; (3) demonstrate compliance with regulatory cleanup standards; and (4) evaluate contaminant travel times and attenuation rates for use in environmental risk analyses (National Research Council, 1990).

Realistic simulations must account for the fact that subsurface media are naturally nonuniform and heterogeneous. One way to systematically analyze flow in these systems involves the application of homogenization procedures in which larger scale (1 km) model equations are developed to reflect the effects of smaller scale (10 m) heterogeneity (Dagan, 1989; Gelhar, 1993). In this sense, numerical simulations based on homogenized models seek to reproduce the bulk flow and transport behavior observed over relatively large spatial scales. This type of approach, however, cannot reveal the complicated preferential flow paths that are present in heterogeneous media, nor can it account for or fully represent the effects of nonlinear interactions that occur in localized regions. In addition, the idealizations used in the homogenized approaches to describe the character and form of subsurface heterogeneity may not be applicable in all systems.

Because of this, we are following a strategy of using detailed (local-scale) simulation to complement simpler, bulk-scale models and homogenization theories (Dagan, 1989; Gelhar, 1993). A typical application may involve a field site that is a kilometer square in areal extent and a hundred meters in depth. To represent the effects of medium heterogeneity on the scale of 1 to 10 m, spatial meshes with 10^8 or more nodes may ultimately be required.

In our simulations, heterogeneity will be represented with a random, spatially correlated model of property variability. In this sense, statistical properties of measured data (mean, variance, and spatial correlation or persistence) are reproduced in one or more representative realizations of the subsurface formation (Tompson, Ababou, and Gelhar, 1989). Hundreds of time-dependent simulations may be required as part of an exploratory investigation of various remediation options or to evaluate the impact of different model parameters (Bellin, Salandin,

and Rinaldo, 1992). Consequently, we need the capability to repeatedly run detailed simulations of large sites, and this need drives our use of supercomputers.

In this paper, a simple case study is used to illustrate these effects by comparing detailed flow and transport simulations (more than 8 million spatial zones) for two representations of a model subsurface formation. In each case, the formation is multilayered, consisting of seven distinct hydrostratigraphic units (which we call *geounits* for brevity): five alluvial layers (composed of variably mixed gravels, sands, silts, and clays), a clay region, and a fault zone. In the first simulation, we assume the geounits have homogeneous (uniform) properties; in the second simulation, we assume they have heterogeneous (nonuniform) properties, as described by the random model referred to above.

Saturated (single-phase) flow in the system is computed by first solving an elliptic partial differential equation (PDE) for the hydraulic head (similar to a flow potential), from which velocities are later computed. Contaminant migration is simulated using the computed velocity field and an independent particle-in-cell code. Because contaminant concentrations are typically dilute, the transport and flow calculations can be considered decoupled. The flow computation is done on a 256-node CRAY T3D massively parallel computer using the simulation code PARFLOW, and the transport is done on a CRAY C90 vector supercomputer using the particle-in-cell code SLIM.

2 Simulation Methodology

In this section, we describe our simulation strategy. Specifically, we discuss problem definition, groundwater flow modeling, and contaminant transport.

2.1 SUBSURFACE MODEL DEFINITION

The subsurface model is the key to problem specification. Geologists and hydrologists use hard data (gathered from monitoring wells and core samples) and soft data (via indirect measurements and inferred knowledge about a given site) to develop a conceptual model of the subsurface. This usually results in the identification of distinct hydrostratigraphic units (geounits), each having a specific distribution of soils and associated properties, such as hydraulic conductivity. This process is often aided by the use of a front-end graphical user interface (e.g., Groundwater Modeling System) (U.S. Department of Defense, 1995), which allows the engineer or scientist to define and adjust the modeled geounits interactively.

“We need the capability to repeatedly run detailed simulations of large sites, and this need drives our use of supercomputers.”

Flow in subsurface media is usually slow enough to be described by Darcy's law (Bear, 1972),

$$\varepsilon \mathbf{v} - K \nabla h = 0, \quad (1)$$

where \mathbf{v} is the average groundwater velocity (L/T), ε is the local medium porosity, h is the hydraulic head (L), and K is the hydraulic conductivity of the medium (L/T).

In a homogenized model, an effective mean value for the hydraulic conductivity and porosity is assigned to each geounit (Gelhar, 1993), even though K may vary locally by several orders of magnitude. In our detailed models, stochastic simulation is used to reproduce variabilities (i.e., heterogeneities) within the layers (see Figure 1). Specifically, we use Tompson's turning bands algorithm (Tompson, Ababou, and Gelhar, 1989) to generate a spectral random field for the log-hydraulic conductivity ($\ln K$) distribution. This distribution often is modeled by the following (Gelhar, 1993; Tompson, 1993; Tompson, Ababou, and Gelhar, 1989; Tompson, Vomvoris, and Gelhar, 1988):

$$\ln K(\mathbf{x}) = F + f(\mathbf{x}),$$

where F is a constant and $f(\mathbf{x})$ is a realization of a stationary, zero-mean, spatially correlated random field. The correlation function of f is anisotropic and assumed to be of the form

$$C(\mathbf{r}) = \sigma^2 \exp(-(r_x/\lambda_x)^2 - (r_y/\lambda_y)^2 - (r_z/\lambda_z)^2)^{1/2},$$

where $\mathbf{r} = (r_x, r_y, r_z)$ is a spatial separation lag vector. Input parameters to this model include the geometric mean K_G (L/T) for K ($K_G = \exp(F)$), the variance σ^2 of f , and correlation lengths (L) λ_x , λ_y , and λ_z . Each geounit may have its own set of geostatistics; the turning bands algorithm is run within each. In this way, we represent both the gross geologic features and the fine-scale heterogeneities. As we will see in the numerical simulations, these heterogeneities strongly influence contaminant migration (see Ashby et al., 1995, for a discussion of the parallel implementation and performance of turning bands).

2.2 FLOW FIELD COMPUTATION

Our mathematical model of groundwater flow is derived by substituting Darcy's law (equation 1) into the conservation of fluid mass in a nondeforming medium

$$\nabla \cdot (\varepsilon \mathbf{v}) - Q = 0$$

to obtain

$$-\nabla \cdot (K \nabla h) - Q = 0. \quad (2)$$

Here, Q is a source term (T^{-1}) used, for example, to represent pumping wells. The hydraulic conductivity realization is obtained from a turning bands algorithm (applied to each layer) as described above. At present, the problem domain is assumed to be a parallelepiped; the boundary conditions may be Dirichlet, Neumann, or mixed.

We use a standard 7-point finite volume spatial discretization on a uniform mesh. After discretization, we obtain a large system of linear equations, $Ah = q$. The coefficient matrix A is symmetric positive definite and has the usual seven-stripe pattern. The matrix has order $N = n_x \times n_y \times n_z$, where the n_i are the number of grid points in the x , y , and z directions, respectively. For problems of interest, N is in the millions; the large number is dictated by the size of the physical site and the need to resolve heterogeneities adequately. Once the hydraulic head is computed, the velocity field can be calculated easily using a simple differencing scheme. This field is then passed to a transport code to simulate contaminant migration (see below).

The solution of the large linear system is computationally intensive and must be done efficiently and accurately. We are using a multigrid-preconditioned conjugate gradient (MGCG) solver. An attractive feature of good multigrid algorithms is that the rate of convergence is independent of problem size, meaning that the number of iterations remains fairly constant. Although one could use multigrid as a stand-alone solver, we prefer to use it as a preconditioner because it is easier to implement in this fashion. The outer CG iteration provides robustness, especially when the underlying PDE has a nearly discontinuous coefficient function, as in our case. In the current implementation of MGCG, the preconditioning step consists of a single standard V-cycle, with either damped Jacobi or red/black Gauss-Seidel smoothing options (pre- and post-smoothing is done). Since our problems tend to be anisotropic (due to a skewed grid cell aspect ratio), we employ a semicoarsening strategy in which the grid is coarsened in one spatial direction at a time. Because of the discontinuous nature of K , we use operator-induced prolongation and restriction, as well as an algebraic definition of the coarse grid operators. See Ashby and Falgout (1996) for a complete description of this algorithm and its parallel performance for a variety of test problems.

To parallelize the MGCG algorithm, we distribute the problem data across a logical 3-D process grid consisting

of $P = p \times q \times r$ processes. The data within a process are viewed as a three-dimensional subgrid of grid points (as defined by the discretization of equation (2)). For example, vector element data and matrix row data (viewed as stencils) are associated with grid points in the owning process's subgrid. Note that although we distribute the problem data by decomposing the problem domain, we are not doing domain decomposition in the algorithmic sense. We are solving the full problem rather than independent subproblems.

When possible, the communications and computations in PARFLOW are scheduled so that they overlap, thereby enhancing parallel efficiency on machines with the appropriate hardware support. All message-passing primitives are localized within a machine-dependent library called AMPS, which has been layered on top of several message-passing systems, including the Reactive Kernel, PVM, Chameleon, MPI, and CRAY SHMEM. The results in this paper correspond to the SHMEM port.

2.3 CONTAMINANT TRANSPORT SIMULATION

We now consider the transport of one or more reactive chemical contaminants within a steady-state groundwater flow field (as calculated by PARFLOW, for example). Specifically, we consider chemical constituents that exist (at low concentration) within the flowing aqueous phase and as sorbed phases on the solid matrix of the porous medium. Mass transfer between the liquid and solid phases is assumed to be governed by reversible, equilibrium, sorption reactions (de Marsily, 1986). This means that the chemical mass is partitioned between the liquid and solid phases instantaneously, or at least much faster than the time scale of advective flow (equilibration time is commonly on the order of a few minutes compared to a time scale of days for our problem).

Because the flow and transport processes can be decoupled, we may use the computed steady-state velocity field directly within the total mass balance for each constituent i ,

$$\frac{\partial \rho_i}{\partial t} + \nabla \cdot \left(\frac{\mathbf{v} \rho_i}{\mathcal{R}_i} \right) - \nabla \cdot \left(\boldsymbol{\varepsilon} \mathbf{D}_i \cdot \nabla \left(\frac{\rho_i}{\boldsymbol{\varepsilon} \mathcal{R}_i} \right) \right) = \frac{-\rho_i}{\boldsymbol{\varepsilon} \mathcal{R}_i} Q. \quad (3)$$

Here, $\rho_i(\mathbf{x}, t)$ is the total (aqueous plus solid) mass density of constituent i at a point in the subsurface formation (m/L^3), $\mathbf{v}(\mathbf{x})$ is the fluid velocity computed by PARFLOW (L/T), and $\mathbf{D}_i(\mathbf{v})$ is a hydrodynamic dispersion tensor (L^2/T). This tensor is typically described by



“The solution of the large linear system is computationally intensive and must be done efficiently and accurately.”

$$\mathbf{D}_i \approx \alpha_T |\mathbf{v}| \mathbf{I} + (\alpha_L - \alpha_T) \frac{\mathbf{v} \otimes \mathbf{v}}{|\mathbf{v}|},$$

where α_L and α_T are the local longitudinal and transverse dispersivities (L), respectively.

In equation (3), we have assumed that sorption onto the solid phase occurs such that the aqueous concentration, c_i (m/L^3), may be inverted from

$$\rho_i = \varepsilon c_i \mathcal{R}_i. \quad (4)$$

The dimensionless quantity \mathcal{R}_i represents the local partitioning or retardation capacity of the soil and may generally be concentration dependent. In the experiments below, we consider a nonreactive tracer ($i = 1$) and a linearly sorbing compound ($i = 2$) such that $\mathcal{R}_1 = 1$ and $\mathcal{R}_2 = 1 + \varepsilon k_d / (1 - \varepsilon)$, where k_d is a dimensionless sorption coefficient.

In our simulations, the soil sorptivity will be specified from a direct correlation with hydraulic conductivity through $\ln \varepsilon k_d \approx -0.86 - 0.32 \ln K$ (Tompson, 1993). The quantities ε , α_L , and α_T will be held constant and set to representative local values of 0.3 ft, 1.0 ft, and 0.1 ft, respectively.

The right-hand side in equation (3) accounts for a loss of chemical mass from a pumping sink of strength Q . Equations (3) and (4) are the fundamental transport equations we wish to solve (for ρ_i and c_i).

We use an efficient particle-grid method (Tompson, 1993; Tompson and Dougherty, 1992; Tompson, Vomvoris, and Gelhar, 1987; Uffink, 1983) to solve these equations. The total concentration distribution for constituent i , ρ_i , is represented by a collection of N_i particles (each of mass m). The concentration at a given point in the subsurface is determined from the mass density of particles in a neighborhood about that point. The number of particles per unit mass specifies the particle resolution. The greater this resolution, the more accurate the simulation and the greater the computational expense. A coupled streamline advection and random walk procedure is used to move and diffuse the particles through the subsurface. Particles may be removed from the system as a result of pumping within the domain or if they cross a domain boundary (Tompson, 1993; Tompson et al., 1994).

In the SLIM computer code (Tompson, Vomvoris, and Gelhar, 1987), the continuum variables are represented as cell-centered quantities on the grid used in PARFLOW (which computes the hydraulic head at nodes). During particle displacement, these quantities are gathered to particle positions using an interpolation formula

(Tompson and Dougherty, 1992). Following particle displacement, the particle masses are scattered back to the grid cells using another interpolation formula, leading to a distribution of $\rho_i(\mathbf{x}, t)$ on the grid. Some particles may be removed to account for boundary conditions and mass sinks.

An initial point source is represented by assigning an appropriate number of particles to the desired cells to represent a specified amount of total (aqueous plus sorbed) mass. Boundary conditions at the edges of the grid domain may allow for outflow, reflection (no flux), or Dirichlet (constant concentration) conditions. Isolated pumping wells located within single-grid cells remove particles at random as a function of the time step and pumping rate.

3 Simulation Results

In this section, we describe the results of two numerical simulations, the purpose of which is to study the influence of fine-scale heterogeneities on contaminant migration (described above) through a multilayered subsurface. Our test problem corresponds to a physical domain of size $12,700 \times 12,700 \times 630$ ft³ with $N = 257 \times 257 \times 129$ grid points (over 8 million spatial zones). This implies an approximate grid spacing of $\Delta_x = 49.6$ ft, $\Delta_y = 49.6$ ft, and $\Delta_z = 4.9$ ft, yielding an approximate grid cell aspect ratio of 10:10:1. The subsurface models are shown in Figure 1. The geostatistical parameters used in each layer of the heterogeneous model are given in Table 1. In the homogeneous model, the geometric means, K_G , are used throughout the layers (i.e., $\sigma = 0$). We impose no-flow conditions on the top and bottom of the domain and piecewise linear Dirichlet boundary conditions on the four vertical sides (hydraulic head values ranged from 510 to 670 ft, depending on the side).

We study the impact of a single pumping well (located in the center of the domain) on the migration of the two chemical species of interest for each of the homogeneous and heterogeneous models. The pumping action was simulated by fixing the hydraulic head at the appropriate grid cell to 530 ft. We simulate 411 years of chemical transport, turning the well on at time $t = 103$ years. To do this, we computed two flow fields: one without the well and one with the well. (Because the fluid and soil matrices are incompressible, there are no transient effects associated with starting the well.) These flow fields were passed to SLIM, which made the switch at the prescribed time step. We remark that this investigation is similar to the one

described in Tompson et al. (1994), but there we modeled only a 320-ft deep two-layer subsurface medium on a coarser computational grid ($\Delta_x = \Delta_y = 100$ ft).

3.1 FLOW VELOCITIES

The MGCG algorithm was halted once the C -norm of the relative residual was less than 10^{-9} (where C denotes the multigrid preconditioning operator). We used one step of symmetric red/black Gauss-Seidel as the smoother (both pre- and post-smoothing). The resulting hydraulic heads were then used to compute the velocity components needed to simulate contaminant migration with SLIM. The hydraulic heads were computed on a $257 \times 257 \times 129$ node-centered grid, and the velocity components were computed on the corresponding $256 \times 256 \times 128$ cell-centered grid.

The head contours for both the homogeneous and heterogeneous subsurface models (with the pumping well) are illustrated in Figure 2. The xy plots correspond to slice-plane data taken midway in the z -direction. The xz plots correspond to data taken one-fifth of the way in the y -direction. From the figure, we see that the general direction of flow is southeast to northwest in both subsurface models (flow runs perpendicular to the contour lines from high head to low head). However, the wavy contours in the heterogeneous model indicate that the direction of flow varies considerably more than in the homogeneous model. Furthermore, the conductivities in the heterogene-

Table 1
Geostatistical Parameters Used in Each of the Heterogeneous Layers Shown in Figure 1

Geounit	K_G	σ	λ_x	λ_y	λ_z
Top layer	4.56	2.08	400.0	400.0	20.0
Second layer	5.92	1.55	400.0	400.0	20.0
Third layer	1.98	1.43	400.0	400.0	20.0
Fourth layer	2.55	1.45	400.0	400.0	20.0
Fifth layer	1.39	1.98	400.0	400.0	20.0
Clay layer	0.01	1.46	400.0	400.0	20.0
Fault zone	0.02	—	—	—	—

NOTE: K_G units are ft/day, and λ units are ft.

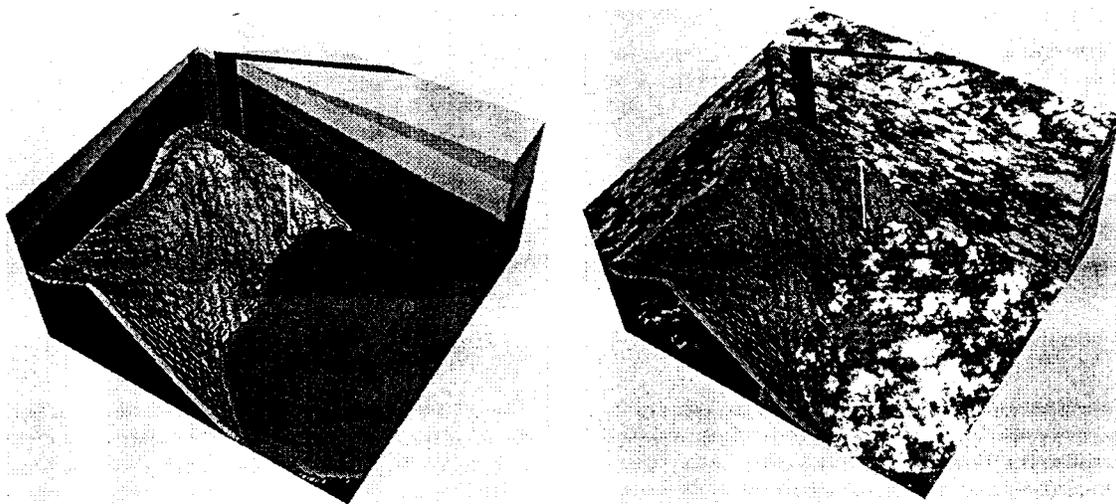


Fig. 1 Subsurface model with homogeneous layers (left) and heterogeneous layers (right). The dark gray region represents a clay layer

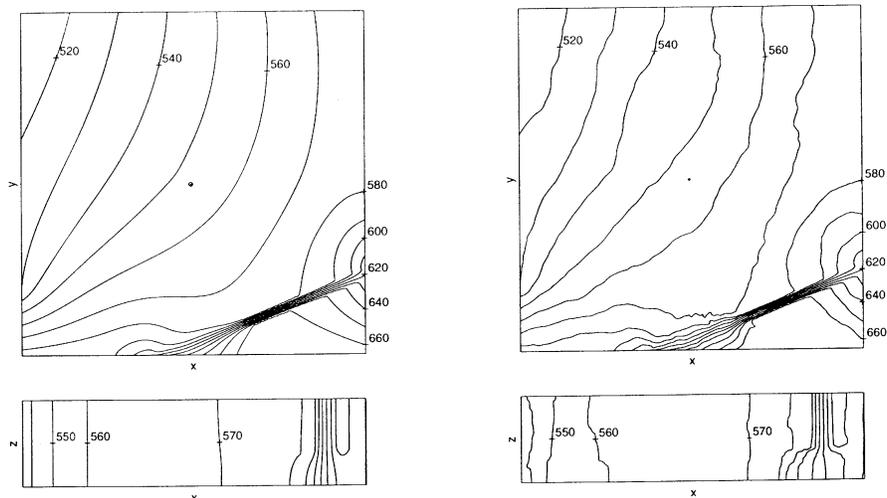


Fig. 2 Head contours for a homogeneous model (left) and a heterogeneous model (right) for selected xy and xz slice planes

Table 2
MGCG Iterations and Wall-Clock Time (in seconds) Required to Solve for the Hydraulic Head on a 257 257 65 Grid (on $P = 4 \ 4 \ 4$ processors)

Subsurface Model	Multigrid-Preconditioned Conjugate Gradient			
	Well	Iterations	Time	Total Time
Homogeneous	No	14	54.5	232.1
Homogeneous	Yes	14	54.5	231.4
Heterogeneous	No	27	99.1	394.5
Heterogeneous	Yes	27	99.1	400.7

ous model result in greatly varying velocity magnitudes. This variation of both flow direction and magnitude produces a large dispersive effect in chemical transport that cannot be accurately modeled with homogeneous models.

Another point of interest is illustrated in the xz contour plots of Figure 2. Many groundwater codes are only two-dimensional simulators that assume flow in the vertical is negligible. From the figure, we see that vertical flow is significant in both of our subsurface models, particularly near the interface of the fault and clay regions. However, as mentioned above, the vertical flow is much more varied in the heterogeneous model.

3.2 PARALLEL PERFORMANCE ON THE T3D

The hydraulic heads for both subsurface models were computed on a 256-node CRAY T3D (each node consists of a 150 MHz DEC Alpha processor and 64 MB of memory). The MGCG iteration counts and timings (in seconds) required for each of the head computations are summarized in Table 2. We give the MGCG time for solving the large linear system, as well as the total time spent in the simulator. This latter time includes input/output (I/O), which is expensive because of the sheer size of the output files (about 68 MB), and problem setup, which is expensive in the case of the heterogeneous model because each processor must call the turning bands algo-

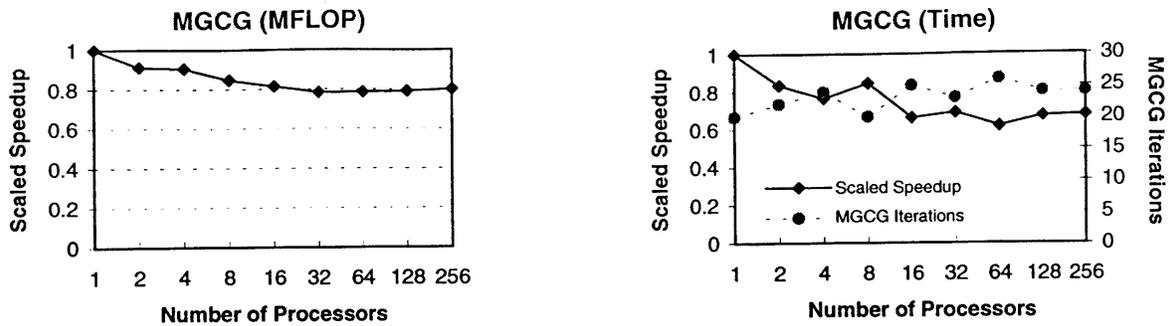


Fig. 3 Scaled speedup on the CRAY T3D

NOTE: MGCG = multigrid-preconditioned conjugate gradient. MFLOPs = million floating-point operations per second.

rithm for every layer that intersects its subgrid data. The process grid topology was $P = 4 \times 4 \times 4$. Notice that the heterogeneous problems take longer, which is due to ill-conditioning induced by the variability in K . The well has a negligible effect on the flow calculation but, of course, affects the resulting head values.

In Figure 3, we report scaled speedup for the MGCG on the T3D. Here, each processor is given a $64 \times 64 \times 32$ subgrid, so that the total problem size on $P = p \times q \times r$ processors is $N_p = 64p \times 64q \times 32r$. In other words, we allow the total problem size to grow with P . Moreover, the shape of the problem domain is determined by the process grid topology $p \times q \times r$. The point of this study is to see how well the routines make use of additional processors. Our goal is to obtain nearly flat curves (good scalability) that are near one (good scaled efficiency). For a more complete discussion of the parallel performance of MGCG, see Ashby and Falgout (1996). For information on the parallel performance of other PARFLOW routines, see Ashby et al. (1995).

In the first graph, scaled speedup is defined to be $M_p / (P M_1)$, where M_p is the MFLOPs (million floating-point operations per second) achieved on P processors. Here, we measure the scalability of our implementation of MGCG. As we see, the graph is fairly flat, indicating good scalability. In the second graph, we define scaled speedup to be T_1 / T_p , where T_p is the time required to execute the MGCG algorithm (to convergence) on P processors. Since the number of iterations required for convergence fluctuates with P , this graph measures the combined scalability of the algorithm itself and our implementation

“This variation of both flow direction and magnitude produces a large dispersive effect in chemical transport that cannot be accurately modeled with homogeneous models.”

of it. We also plot the MGCG iteration count, which varies between 20 and 26 iterations (using the C norm stopping criterion). Notice the inverted relationship between scaled speedup and iteration count (as one would expect). These two graphs demonstrate that both the MGCG algorithm and its implementation are scalable.

3.3 CHEMICAL TRANSPORT

After computing the flow fields, they were passed to SLIM for the contaminant migration study. Recall that in each simulation (homogeneous layers vs. heterogeneous layers), we used two flow fields: one without the well and one with the well. The fields were switched at a specified time to simulate the effect of an undetected contaminant plume spreading for a period of time before activating pump-and-treat remediation. The initial constant concentration was established by assigning particles to 128 contiguous grid cells. We used approximately 781 particles per cell with a total of 10^5 particles per contaminant species in the system.

We simulated 411 years of chemical transport of the sorbing and nonsorbing chemical species. For the homogeneous model, $\Delta t = 25$ days. In the heterogeneous model, the presence of preferential flow channels dictated a smaller time step of $\Delta t = 5$ days. In each case, the ambient (no-well) P_{ARFLOW} velocity field was used to drive the contaminant transport for 103 simulated years. At this point, the pumped (with the well) P_{ARFLOW} velocity field was substituted and used to drive the transport for another 308 simulated years. In each time step, a probabilistic model was used to remove contaminant particles from a hydraulic capture zone surrounding the well.

Figures 4 and 5 show the nonsorbing contaminant plume at four different times in the simulations, in the homogeneous and heterogeneous flow fields, respectively. The times chosen are not all the same for the homogeneous and heterogeneous cases because the plume moves much faster in the heterogeneous flow field, due to the effect of preferential flow channels. The first frame in each figure is $t = 21$ years, before the plume has advanced far from its initial loading in a localized region of space. The second frame in each figure is $t = 103$ years, at the last time step before the well flow field and contaminant sink model were turned on. The third frame for the heterogeneous case shows absorption of the plume both upstream and downstream of the well. The fourth frames for both cases show that nearly all of the contaminant mass upstream of the well has been cleaned up.

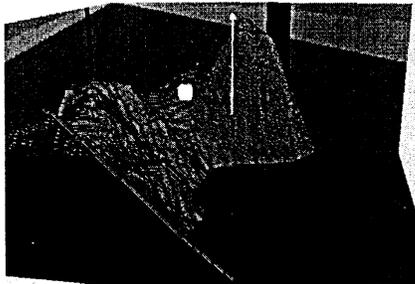
Figures 6 and 7 show the chemically sorbing contaminant plume. The sorbing contaminant migrates noticeably slower than the nonsorbing contaminant because of the retarding effect of the chemical reaction. This is clearly seen in the last frame of the uniform conductivity case, in which the nonsorbing contaminant (see Figure 4) has already been completely cleaned up by the well, while the sorbing plume (see Figure 6) has not yet reached the well's capture zone. The impact of chemical reaction is also clear in the second frame of the heterogeneous conductivity case, in which the sorbing contaminant plume still has much more mass upstream of the well at the initiation of pumping.

3.4 PERFORMANCE ON THE C90

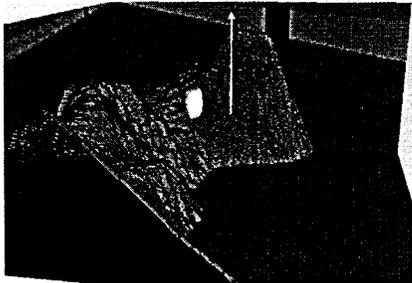
We ran the SLIM simulations on a single processor of a CRAY C90. The uniform conductivity simulation took about 1 hour of CPU time, ran at an average of 136 MFLOPs, and showed an average vector length of 88 (the ideal length is the length of the vector unit, 128). The heterogeneous simulation took about 3.3 hours of CPU time, ran at an average of 146 MFLOPs, and showed an average vector length of 80. Particle-grid simulations are not generally known to achieve high MFLOP rates; they are bound by memory-access speed. These SLIM runs, for example, executed about 132 and 148 million vector memory references per second for the uniform and heterogeneous cases, respectively. Note that the heterogeneous simulation ran at a higher MFLOP rate than the homogeneous simulation but showed a smaller average vector length. The reasons for this are too complex to examine in detail here, but the different profiling measurements given below are a contributing factor.

Profiling measurements on the SLIM for the uniform case reveal that about 34% of the time went toward computing the total mass density values, which includes the accumulation (scatter-add) from the particles. This time can be strongly dependent on the method of density calculation (Hockney and Eastwood, 1981). Moving the particles, including computation of the displacement values, took 25% of the time. Modeling the pumping well, a sort of specialized boundary condition, took 22% of the time. Particle boundary condition enforcement and diagnostic breakthrough counting (not discussed here) took another 18%; the remaining time was spent in miscellaneous routines.

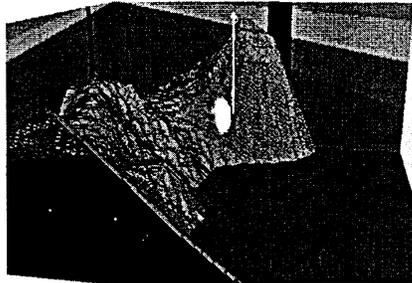
Profiling measurements for the heterogeneous case show that about 35% of the time went toward computing the total mass density values, which includes the accumu-



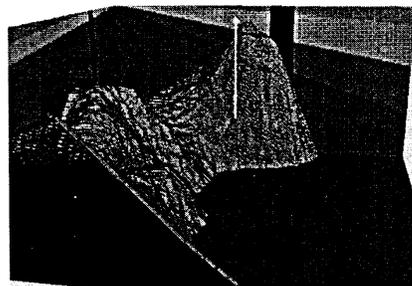
$t = 21$ years



$t = 103$ years

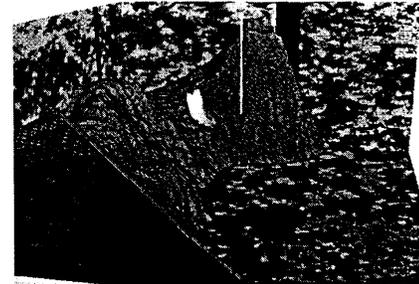


$t = 274$ years

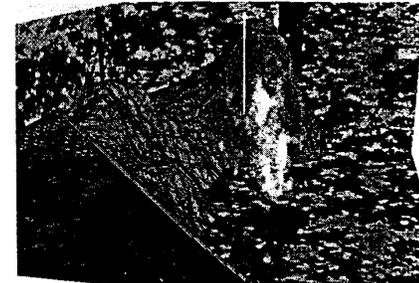


$t = 411$ years

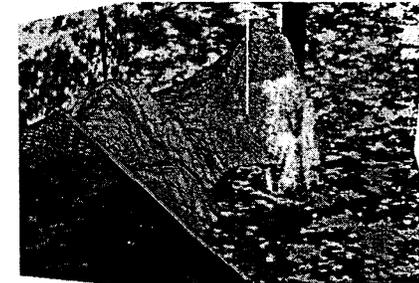
Fig. 4 Snapshots in time: Dispersion of nonsorbing contaminant assuming homogeneous layers



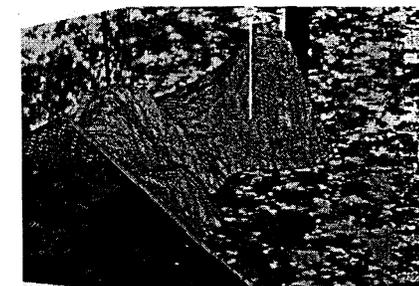
$t = 21$ years



$t = 103$ years

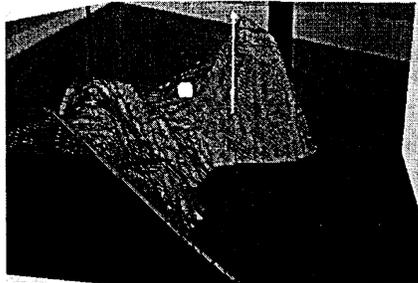


$t = 137$ years

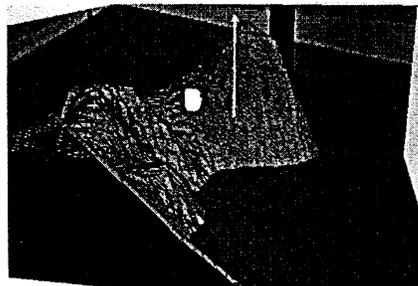


$t = 342$ years

Fig. 5 Snapshots in time: Dispersion of nonsorbing contaminant assuming heterogeneous layers



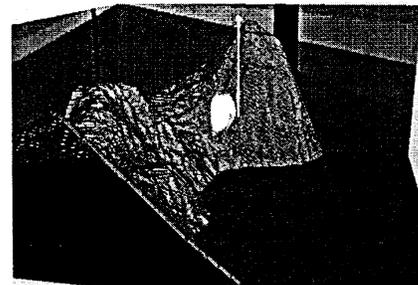
$t = 21$ years



$t = 103$ years

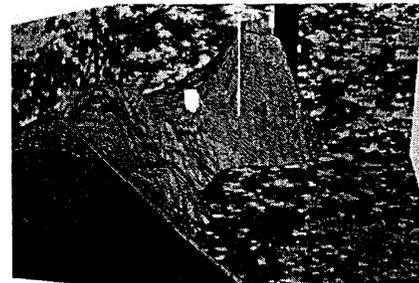


$t = 274$ years

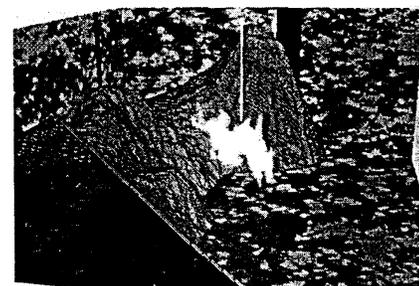


$t = 411$ years

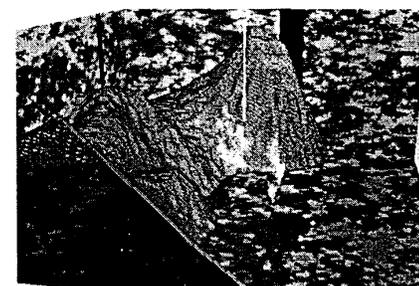
Fig. 6 Snapshots in time: Dispersion of sorbing contaminant assuming homogeneous layers



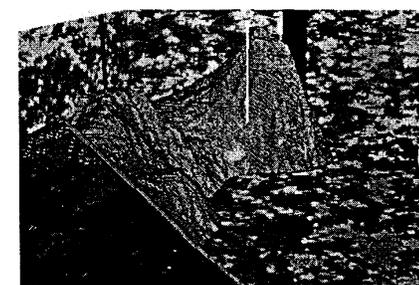
$t = 21$ years



$t = 103$ years



$t = 137$ years



$t = 342$ years

Fig. 7 Snapshots in time: Dispersion of sorbing contaminant assuming heterogeneous layers

lation (scatter-add) from the particles. Moving the particles took 21% of the time. Modeling the pumping well took 16% of the time. Determining the liquid mass densities from the total mass densities took 7.3% of the time. Particle boundary condition enforcement and diagnostic breakthrough counting took another 18%; the remaining time was spent in miscellaneous routines.

The reason that the heterogeneous simulation used 5 times as many time steps as the homogeneous simulation, yet took only 3.3 times as much total CPU time, is not easily explained. Early tests led us to speculate that this was partly accounted for by the greater dispersion in the heterogeneous case—the plume extends to many more grid cells than in the homogeneous case, which leads, in principle, to greater vector length in operations over all grid cells containing any contaminants. A closer look revealed that in the parts of the code operating on only grid cell data, the average vector length measured was actually shorter in the heterogeneous case. Aside from this curious behavior, more careful measurements also showed that the floating-point operation rates in all parts of the code are the same or higher in the heterogeneous case than the uniform case; this at least accounts empirically for the greater execution efficiency of the random case.

4 Conclusions

Our results indicate that `PARFLOW` can readily handle very large and highly resolved steady flow problems in realistic, nonuniform porous media. Our particular application has illustrated the importance and impact of nonuniformity in soil properties on solute migration and spreading phenomena. Performance statistics of `PARFLOW`'s MGCG scheme on the T3D have shown excellent scalability properties. The companion particle-in-cell code, `SLIM`, works very well on the C90 for simulating the migration and spreading of small, well-contained solute pulses but suffers when the mass distribution is widespread and concentrations are large.

Future communications will report on improvements to `PARFLOW` that include the addition of a more general transient flow formulation, allowing for multiple-phase flow problems to be treated, an internal (grid-based) transport scheme, and an adaptive grid capability for modifying the spatial grid resolution in arbitrary portions of the simulation domain. We will also report on the specific application of `PARFLOW` to problems at several actual field locations.

ACKNOWLEDGMENTS

We acknowledge the valuable participation of Chuck Baldwin in the `PARFLOW` project and Barry Becker for help with graphics.

The CRAY T3D experiments described in this paper were run on the 256-node machine located at Lawrence Livermore National Laboratory (LLNL) as part of the H4P Industrial Computing Initiative funded by the U.S. Department of Energy (DOE) defense programs. The CRAY C90 experiments were run on the machine located at the National Energy Research Supercomputer Center at LLNL.

This work was performed at LLNL under contract W-7405-ENG-48 for the DOE. It was supported in part by the DOE Defense Programs Technology Transfer Initiative; the Laboratory Directed Research and Development Program at the LLNL; the Mathematical, Information, and Computational Sciences Division of the Office of Energy Research, DOE; and by the Strategic Environmental Research and Development Program of the U.S. Department of Defense through the cooperation of the USAE Waterways Experiment Station (Vicksburg, Mississippi).

BIOGRAPHIES

Steven F. Ashby, Ph.D., is Director of the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). CASC was formed in March 1996 to conduct research in computational mathematics, computer science, and scientific computing in support of LLNL programs. Dr. Ashby's research interests include numerical linear algebra (iterative methods and adaptive preconditioning techniques for large linear systems), algorithm design and analysis, and scientific computing. He has collaborated with colleagues on a number of scientific applications, including the numerical simulation of groundwater flow and contaminant transport through heterogeneous porous media (the subject of this paper). Dr. Ashby earned his Ph.D. in Computer Science at the University of Illinois in 1987, and joined LLNL in 1988. He was nominated for the Computerworld Smithsonian Information Technology Leadership Award for Breakthrough Computational Science (one of only six national nominees) for his leadership of the `ParFlow` subsurface simulation project. He is a member of the Program Committee for Colorado Conferences on Iterative Methods, the Phi Beta Kappa and Sigma Xi academic honor societies, and the Society for Industrial and Applied Mathematics. He has co-authored numerous papers in a variety of areas, and has been the principal investigator on several large research projects.

William J. Bosl, Ph.D., is a computer scientist at the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). He joined LLNL in 1992, and was one of the Laboratory scientists named in the Computational Directorate's Extraordinary Programmatic Team Award presented in 1996. His research interests include computer simulation of coupled hydro-elastic processes, the role of pore fluids in crustal processes (induced seismicity and aftershocks), geostatistics (the characterization and hydraulics of fractures), and parallel algorithms (particularly lattice Boltzmann methods). Dr. Bosl holds a Ph.D. in Geophysics from Stanford University, granted in 1998. He earned his M.S. in Mathematics at the University of Pittsburgh in 1986, and an M.S. in Atmospheric Science at the University of Michigan in 1984. He is a member of the American Geophysical Union, the Seismological Society of America, the American Association of Petroleum Geologists, and the Society for Industrial and Applied Mathematics. He also holds memberships in the Phi Beta Kappa, Phi Kappa Phi, and Sigma Phi Sigma academic honor societies.

Robert D. Falgout, Ph.D., is a mathematician at the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). He also serves CASC as the Scalable Linear Solvers project leader. Through the Scalable Linear Solvers project, CASC develops algorithms and software for solving large, sparse linear systems of equations on parallel computers. Particular problems of interest arise in the simulations codes being used to study physical phenomena in the defense, environmental, energy and biological sciences. Dr. Falgout's research interests include scientific computing with emphasis on numerical algorithms and their implementation on high-performance computers, especially multigrid methods. He earned a Ph.D. in Applied Mathematics at the University of Virginia (UVA) in 1991, and is a member of the Society for Industrial and Applied Mathematics.

Steven G. Smith is a computer scientist at the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). He has been affiliated with LLNL since 1991. His research interests include scientific computing on distributed and highly parallel architectures, software tools and software engineering techniques for high-performance computing, and scientific data visualization. Mr. Smith earned his B.S. in Computer Science at the University of Illinois at Urbana-Champaign in 1991. He has co-authored more than a dozen publications on topics as varied as numerical simulation of groundwater flow, the multicomputer toolbox, and high-level message passing constructs. Prior to joining LLNL, he was employed by the Cooperative Education Program located at IBM Federal Systems Division in Manassas, Virginia.

Andrew F. B. Tompson, Ph.D., is a hydrologist in the Geosciences and Environmental Technologies Division of Lawrence Livermore National Laboratory (LLNL). His research interests include physics of fluid flow, chemical transport, and chemical

transformation in porous media; fluid mechanics, groundwater supply and contamination, and the mathematical modeling of these processes; computational solution methods and numerical analysis; geostatistics, and stochastic processes. Recent project activity at LLNL has focused on a variety of subjects including modeling analyses of novel in-situ contaminant remediation technologies, radionuclide migration away from underground nuclear tests, artificial recharge in large heterogeneous aquifers, and multiphase flow, as predicted from new balance equations that more rigorously account for the effects of microscale interfacial processes.

Dr. Tompson earned his Ph.D. from Princeton University (Civil Engineering) in 1985 and his Sc.B. (Civil Engineering) from Brown University in 1980. He joined LLNL in 1986 after a two-year postdoctoral appointment at MIT. He is a member of the American Geophysical Union and the Society for Industrial and Applied Mathematics and is on the editorial board of *Advances in Water Resources* and *Computational Geosciences*.

Timothy J. Williams, Ph.D., is a computational physicist at the Advanced Computing Laboratory (ACL) at Los Alamos National Laboratory (LANL). He joined LANL in 1995 as part of the Parallel Object-Oriented Methods and Applications (POOMA) team. Prior to LANL, Dr. Williams was at Lawrence Livermore National Laboratory from 1989, where his work included particle-in-cell (PIC) simulation of plasmas and of contaminant migration in groundwater flows. His research interests include developing an object-oriented framework for scientific computing on parallel architectures, and applying it to computational physics simulations including multimaterial compressible hydrodynamics and PIC simulation of tokamak plasma microturbulence. Dr. Williams holds a Ph.D. in Plasma Physics from the College of William and Mary, granted in 1988. He earned his B.S. in Physics and Mathematics from Carnegie-Mellon University in 1982.

REFERENCES

- Ashby, S. F., and Falgout, R. D. 1996. A parallel multigrid preconditioned conjugate gradient algorithm for groundwater flow simulations. *Nuclear Science and Engineering* 124:145-159. (Also available as LLNL Technical Report UCRL-JC-122359.)
- Ashby, S. F., Falgout, R. D., Smith, S. G., and Tompson, A.F.B. 1995. The parallel performance of a groundwater flow code on the CRAY T3D. In *Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing*, Society for Industrial and Applied Mathematics, February, San Francisco. (Also available as LLNL technical report UCRL-JC-118604.)
- Bear, J. 1972. *Dynamics of fluids in porous media*. New York: Dover.
- Bellin, A., Salandin, P., and Rinaldo, A. 1992. Simulation of dispersion in heterogeneous porous formations: Statistics,

- 1st order theories, convergence of computations. *Water Resources Res.* 28:2211-2227.
- Dagan, G. 1989. *Flow and transport in porous formations*. New York: Springer Verlag.
- deMarsily, G. 1986. *Quantitative hydrogeology: Groundwater hydrology for engineers*. New York: Academic Press.
- Gelhar, L. W. 1993. *Stochastic subsurface hydrology*. Englewood Cliffs, NJ: Prentice Hall.
- Hockney, R. W., and Eastwood, J. W. 1981. *Computer simulation using particles*. New York: McGraw-Hill.
- National Research Council. 1990. *Ground water models: Scientific and regulatory applications*. Washington, DC: National Academy Press.
- Tompson, A.F.B. 1993. Numerical simulation of chemical migration in physically and chemically heterogeneous porous media. *Water Resources Res.* 29:3709-3726.
- Tompson, A.F.B., Ababou, R., and Gelhar, L. W. 1989. Implementation of the three-dimensional turning bands random field generator. *Water Resources Res.* 25:2227-2243.
- Tompson, A.F.B., Ashby, S. F., Falgout, R. D., Smith, S. G., Fogwell, T. W., and Loosmore, G. A. 1994. Use of high performance computing to examine the effectiveness of aquifer remediation. In *Proceedings of the International Conference on Computational Methods in Water Resources*, edited by A. Peters, G. Wittum, B. Herrling, U. Meissner, C. Brebbia, W. Gray, and G. Pinder, vol. 2. Dordrecht, the Netherlands: Kluwer Academic. (Also available as LLNL technical report UCRL-JC-115374.)
- Tompson, A.F.B., and Dougherty, D. E. 1992. Particle-grid methods for reacting flows in porous media with application to Fisher's equation. *Applied Mathematical Modelling* 16:374-383.
- Tompson, A.F.B., Vomvoris, E., and Gelhar, L. W. 1988. Numerical simulation of solute transport in randomly heterogeneous porous media: Motivation, model development, and application. Technical Report 316, R. M. Parsons Laboratory, Department of Civil Engineering, Massachusetts Institute of Technology. (Also available as LLNL technical report UCID-21287.)
- Uffink, G. 1983. A random walk model for the simulation of macrodispersion in a stratified aquifer. In *Proceedings of the IAHS Symposia*, vol. HS2, 18th General Assembly, August, Hamburg, Germany.
- U.S. Department of Defense. 1995. *Groundwater modeling system: Reference manual*. Provo, UT: Engineering Computer Graphics Laboratory, Brigham Young University.