RANDOM POROUS MEDIA FLOW ON LARGE 3-D GRIDS: NUMERICS, PERFORMANCE, & APPLICATION TO HOMOGENIZATION *

RACHID ABABOU†

Abstract. Subsurface flow processes are inherently three-dimensional and heterogeneous over many scales. Taking this into account, for instance assuming random heterogeneity in 3-D space, puts heavy constraints on numerical models. An efficient numerical code has been developed for solving the porous media flow equations, appropriately generalized to account for 3-D, random-like heterogeneity. The code is based on implicit finite differences (or finite volumes), and uses specialized versions of preconditioned iterative solvers that take advantage of sparseness. With Diagonally Scaled Conjugate Gradients, in particular, large systems on the order of several million equations, with randomly variable coefficients, have been solved efficiently on Cray-2 and Cray-Y/MP8 machines, in serial mode as well as parallel mode (autotasking). The present work addresses, first, the numerical aspects and computational issues associated with detailed 3-D flow simulations, and secondly, presents a specific application related to the conductivity homogenization problem (identifying a macroscale conduction law, and an equivalent or effective conductivity). Analytical expressions of effective conductivities are compared with empirical values obtained from several large scale simulations conducted for single realizations of random porous media.

Key words. Porous media; Random media; Stochastic flow; Random fields; Effective conductivity; Homogenization; Conjugate gradients; Diagonal scaling; Parallel computing; Parallel speed-up; Amdahl's law; Autotasking; Convergence rate.

0. Introduction. Porous geologic media are inherently three-dimensional and heterogeneous over a wide range of scales. These features must be taken into account for accurate modeling of fluid flow and contaminant transport in the subsurface. For instance, in the case of buried nuclear waste sites, many scales of natural heterogeneity must be taken into account in order to capture all possible radionuclide pathways. Likewise, three-dimensionality is required in order to avoid artificial constraints on the degree of freedom and mechanical dispersion of contaminants. In addition, the space-time scales of interest are rather large, typically kilometers and thousands of years or more for high-level radioactive contaminants. This puts heavy constraints on numerical models. Detailed spatial vari-

* This paper was written while its author was at the Commissariat à l'Energie Atomique, CEN-Saclay, France. However, parts of this work are based on previous research conducted while the author was at the Center for Nuclear Waste Regulatory Analyses, San Antonio, Texas. Logistic help for some of the computations was provided by the NASA Ames Research Center, and Cray Research Inc. The 3-D graphics views of potential surfaces were produced with Dynamic Graphics IVM package. The author is indebted to C. Hempel for assistance with Cray timing analyses, to L.W. Gelhar for advice on several aspects of stochastic modeling, and to A.C. Bagtzoglou, G.W. Wittmeyer, T.J. Nicholson, B. Sagar, and M. Durin for help or encouragements on various aspects of this work. The views expressed here are solely those of the author.

† Commissariat à l'Energie Atomique Saclay, DRN/DMT/SEMT/TTMF, 91191 Gif-sur-Yvette Cedex, France.

ability must be incorporated in the coefficients of the governing equations, which must then be solved numerically on large grids. The reader is referred to reference [2] for a comprehensive review of conceptual models, field heterogeneity, and related issues, particularly in the context of nuclear waste disposal.

Here, we investigate computational issues associated with direct simulations of highly heterogeneous flows, and present a particular application related to the homogenization problem (i.e., the question of finding equivalent macroscale coefficients, or constitutive laws, given microscale data). We present a general-purpose numerical code (BIGFLOW) which can efficiently model large three-dimensional (3-D) flow systems in unsaturated, partially saturated, or saturated, heterogeneous geologic media. The code was initially developed at MIT as a research tool to investigate stochastic flow processes [1], and has undergone several enhancements since then. It is being used as a tool for research on subsurface transport of contaminants and tracers, environmental impacts of geologic disposal of nuclear waste, and other environmental problems that depend on a realistic and efficient representation of natural heterogeneity.

The main outlines of this article are as follows. First, a brief description of the flow code and its capabilities is given in Section 1. Secondly, the governing equations, discretization, and algebraic systems are presented in Section 2. Sections 3 through 6 develop an analysis of computational performance of the code for large and 'noisy' matrix systems. Since the matrix solver constitutes essentially the 'computational kernel' of the total code, we focus on the Diagonal Scaling Conjugate Gradient solver, which is especially sparse as well as highly vectorizable (see Section 3). The computational performance of the solver, and of the total code itself, are assessed in several ways: convergence rate analyses (theoretical in Section 4, empirical in Section 5); serial and parallel timings of the code; and evaluation of speed-ups due to coarse-grained parallelization (Section 6). In Section 7, we develop a specific application related to the conductivity homogenization problem. Empirical effective conductivities are 'measured' on numerical flow systems corresponding to large single realizations of random porous media. The numerical grids are typically on the order of one to ten million nodes and the degree of heterogeneity is quite large (Appendix A). The numerical results are compared, briefly, to certain analytical solutions of the conductivity homogenization problem (Appendix B).

1. The numerical code (BIGFLOW). For simulations, we use the BIGFLOW code. In general, BIGFLOW can solve linear as well as nonlinear flow equations for saturated as well as variably saturated porous media. The code is based on mass conservation and Darcy's law for saturated media, or Darcy-Buckingham's law for variably saturated media. In the former case, the dependent variable is the hydraulic potential, i.e. hydraulic head, or equivalently 'total pressure'. In the latter case, the

dependent variables are water content and pressure head (mixed variable formulation). In both cases, the governing equations and constitutive laws are appropriately generalized to account for fully three-dimensional, and possibly random, heterogeneity.

BIGFLOW is written in ANSI Standard Fortran 77, is free of machine-dependent directives, and is portable without modifications to a variety of computer systems, mainframes, and workstations. An implicit finite difference scheme is used for discretization. Optionally, a modified Picard scheme is used for linearizing unsaturated flow equations (outer iterations), and a preconditioned iterative method is used for solving the resulting matrix systems (inner iterations). Iterative matrix solvers which have been extensively used so far include the Strongly Implicit Procedure 'SIP' ([1],[3]), and Diagonally Scaled Conjugate Gradient 'DSCG' [see following sections]. As will be seen, the solution modules were especially coded to take advantage of sparseness and symmetry of the finite difference systems. A data processor (DATAFLOW) also allows interactive entry and analysis of 3-D numerical datasets. This set of codes, simulator and processor, constitutes the BIGFLOW package. See [8] for more details.

Figure 1: 3-D Views of Hydraulic Equipotential Surfaces for Groundwater Flow in Randomly Heterogeneous, Isotropic, Saturated Porous Medium (1 million nodes).

Figure 1: Top:Full View

Figure 1: Bottom:Cut-Off View of a Single Equipotential Slice

To illustrate the type of flow problems of interest here, we display in *Figure 1* the Numerical solution for steady flow in a randomly heterogeneous, isotropic, saturated porous medium. This figure gives a 3-D view of hydraulic equipotential surfaces (P). The grid size for this problem was $(101)^3$, or one million nodes. A number of other similar problems have been solved using either the SIP solver [1], or, more recently, the DSCG solver [see Table 1 further below]. The numerical solutions have been exploited to study effective conductivity of heterogeneous media [3], and to simulate stochastic contaminant transport for analyses of macrodispersion [17]. The reader will find, in Section 7, a study of effective conductivity that includes the random flow problem depicted in Figure 1, among others.

2. Equations and discretized system. For steady-state, saturated flow systems, combining mass conservation with Darcy's law yields a linear self-adjoint elliptic partial differential equation:

$$\frac{\partial}{\partial x_i}\left[K(x)\frac{\partial P}{\partial x_i}\right] = 0 \quad (\text{sum over } \; i = 1,2,3) \tag{1}$$

where $K(x)$ is the spatially variable hydraulic conductivity (m/s), and P is a total pressure, or hydraulic potential, expressed here in terms of the equivalent height of a water column (meters). We will refer to P as 'pressure' for short.

Solving equation (1) for pressure $P(x)$ also gives the water flux vector through Darcy's equation, $Q_i = -K(x)\ \partial P/\partial x_i$. Note that equation (1) is of the form $\mathrm{Div}(Q) = 0$, which expresses mass conservation. The water velocity vector, V, may be computed by dividing the flux vector by the porosity of the medium, ϕ. The latter quantity is a priori spatially variable as well. This may be important for tracer transport, but has no direct effect on the macroscale hydraulic conduction mechanism, which is our sole concern here.

Using two-point centered finite difference schemes along each of the 3 dimensions yields the discretized system of equations:

$$
\begin{aligned}
&-\tfrac{K[i-1/2]}{(\Delta x_1)^2}\ P[i-1] - \tfrac{K[j-1/2]}{(\Delta x_2)^2}\ P[j-1] - \tfrac{K[k-1/2]}{(\Delta x_3)^2}\ P[k-1] \\
&+\left\{\tfrac{K[i-1/2]+K[i+1/2]}{(\Delta x_1)^2} + \tfrac{K[j-1/2]+K[j+1/2]}{(\Delta x_2)^2} + \tfrac{K[k-1/2]+K[k+1/2]}{(\Delta x_3)^2}\right\} P[0] \\
&-\tfrac{K[i+1/2]}{(\Delta x_1)^2}\ P[i+1] - \tfrac{K[j+1/2]}{(\Delta x_2)^2}\ P[j+1] - \tfrac{K[k+1/2]}{(\Delta x_3)^2}\ P[k+1] \\
&= 0
\end{aligned}
\tag{2}
$$

where we used the short hand notation: $[0] = (i,j,k)$, $[i \pm 1/2] = (i \pm 1/2, j, k)$ and so on. The midnodal conductivity $K[i+1/2]$ is estimated by the geometric mean of nodal conductivities $K(i,j,k)$ and $K(i+1,j,k)$.

Once the discrete pressure field $P(i,j,k)$ is known, the discrete flux or velocity field may be computed from a discretized form of Darcy's law. This is done in a consistent manner by using two-point centered finite difference approximations, of the kind that led to eq. (2) in the first place. Furthermore, Neumann boundary conditions, in the form of prescribed pressure gradients arising from prescribed flux conditions, are also handled by the same two-point centered FD scheme. Incidentally, this requires assuming that a Neumann-type boundary is located at mid-nodal points, just next to the nodes bordering the grid. In this way, the 'second order accuracy' of the discretization scheme, as well as the seven-diagonal symmetric structure of the coefficient matrix, are preserved after implementation of boundary conditions and matrix condensation.

In the forthcoming test problems, the computational domain will be a 3-D rectangular or cubic prism, and flow will be driven by a 'regional' pressure gradient, obtained by imposing different pressures on two opposite boundaries (Dirichlet conditions). All other boundaries will be impervious, with zero flux, or equivalently zero pressure gradient (Neumann conditions).

Now, the finite difference system (2) may be equivalently expressed in

matrix-vector notation as:

$$Ay = b \tag{3}$$

where y is the vector of nodal pressures ($P_{i,j,k}$), b is the vector containing boundary terms from boundary conditions, and A is the seven-diagonal, heterogeneous conductivity matrix. Note that we use uppercase bold symbols for matrices, and lowercase bold symbols for vectors. The conductivity values in A are assumed to be either constant or randomly heterogeneous. Whatever the case may be, it can be shown that A is symmetric positive-definite and weakly diagonal dominant. The latter property holds provided that a Dirichlet condition be specified on at least one of the boundary nodes. For completeness, note also that A is an M-matrix, having strictly positive diagonal elements and negative off-diagonal elements.

In the BIGFLOW code, a diagonal-by-diagonal storage scheme is used for the sparse symmetric matrix A. For symbolic manipulation purposes, A can be represented in terms of four single-diagonal matrices(D_0, D_1, D_2, D_3), as follows:

$$A = D_0 + D_1 + D_2 + D_3 + {D_1}^T + {D_2}^T + {D_3}^T .$$

To each of the four diagonal and off-diagonal matrices D_i corresponds a vector d_i ($i = 0, 1, 2, 3$). Such vectors are the only entities that are actually manipulated in the numerical code. Thus, instead of a full $N \times N$ matrix, only four vectors of length N need to be stored, corresponding to the main diagonal and three non-zero off-diagonal lines in the lower half of the matrix. Each of the matrix vectors d_i, as well as the solution vector y, is represented by a triple-indexed array variable (one index per spatial dimension).

The specialized algebra and data structure described above minimizes both storage and CPU time. As a consequence, the total physical memory required for solving heterogeneous 3-D systems with DSCG and other sparse solvers is modest: about 13 words per node for saturated flow, and up to twice as much for unsaturated flow.

3. Implementation of diagonally scaled conjugate gradients (DSCG). The DSCG algorithm is implemented in two steps, first by applying symmetric diagonal scaling (DS) to the original system, and secondly by solving the scaled symmetric system using the Conjugate Gradient (CG) method. Symmetric diagonal scaling is implemented as shown in *Box 1*.

- Define diagonal preconditioner: $D_0 = \mathrm{diag}(A)$
- Compute scaled coefficient matrix: $A_* = {D_0}^{-1/2}\ A\ {D_0}^{-1/2}$
- Compute scaled right-hand side: $b_* = {D_0}^{-1/2}\ b$
- Define scaled system: $A_*\ y_* = b_*$
- After solving, get unscaled solution $y = D_0^{-1/2}\ y_*$

Box 1: Symmetric Diagonal Scaling.

0. Initialize parameters: $\mu_{\mathrm{old}} = \mu_{\mathrm{new}} = \omega = 0$
   Initialize residual vector: $r = b - A\ y$
   Initialize search vector: $p = 0$
1. Update μ-parameter: $\mu_{\mathrm{old}} = \mu_{\mathrm{new}}$
   $\mu_{\mathrm{new}} = 1/(r^{\mathrm{T}} r)$
2. Update search vector: $p = r + (\mu_{\mathrm{old}}/\mu_{\mathrm{new}})\ p$
3. Compute auxiliary vector: $z = A\ p$
4. Compute ω-parameter: $\omega = [\mu_{\mathrm{new}}(p^{\mathrm{T}} z)]^{-1}$
5. Update solution vector: $y = y + \omega p$
   Update residual vector: $r = r - \omega z$
6. Go to step 1 if stopping criterion not satisfied, else stop.

Box 2: Conjugate Gradient Iterations.

Given an initial guess for the (scaled) solution vector, the conjugate gradient algorithm iteratively solves the (scaled) symmetric system as indicated in *Box 2*. To obtain the full DSCG solver, the reader should replace the A, y, b, of Box 2 by the A_*, y_*, b_* quantities defined in Box 1. Note that D_0 is the main diagonal of the unscaled matrix A.

The stopping criterion invoked in Step 6 of Box 2 could be a maximum number of iterations, or an error norm criterion (ε) to be compared to the

L_2- or L_∞-norm of the error vector $\boldsymbol{\delta} = y_{\rm new} - y_{\rm old}$. In Step 3 of Box 2, the matrix-vector product $z = Ap$ is computed as a sum of seven shifted dot products, one for each non-zero diagonal line of A. All other array operations are straight dot products, except for the L_∞-norm of error.

4. Theoretical estimates of solver performance. For a wide class of iterative solvers that includes CG and DSCG, the number of iterations required to decrease the error by (say) six orders of magnitude is known to be approximately proportional to the square-root of the condition number of the coefficient matrix. In the case at hand, the condition number is typically $O(n^2)$, where n represents the uni-directional size of the grid along its largest side (see [1] and [12]). For each iteration, the computational work, or number of operations, is proportional to the multi-dimensional number of nodes (N). Multiplying by the estimated number of iterations yields a total work on the order $O(N^p)$, with exponent $p = 4/3$ for a 3-D cubic grid [$p = 3/2$ for a 2-D square grid; $p = 2$ for a 1-D grid].

These approximate 'order of magnitude' relations give useful indications on the way computer time (total work) increases with grid size. Note that the rate of increase is superlinear in all cases, but comes closer to linear as the 'effective' number of spatial dimensions increases from 1 to 3. The shape of the grid is also important; in terms of numbers of nodes, a narrow rectangular prism should be viewed, in effect, as a 1-D rather than a 3-D grid.

However, these approximate relations have several shortcomings:

- The convergence rate estimate does not indicate the influence of conductivity heterogeneity and spatial structure;
- It is only a 'worst case' estimate obtained from an approximate error upper bound;
- This 'worst case' estimate must break down as the number of iterations approaches the number of equations (N), since the Conjugate Gradient method gives the exact solution in no more than N iterations (within machine precision); and
- the assumption that, for each iteration, computational work is proportional to grid size (N), does not take into account non-proportional speed-up effects due to vector and parallel processing. The latter effects depend on machine architecture, and on the way algorithms like those shown in Boxes 1 and 2 are programmed.

5. Observed convergence rates. *Figure 2* shows the number of DSCG iterations (I) *vs.* uni-directional grid size (n) for the case of constant conductivity ($\sigma = 0$). In this special case, the equation being solved is actually a Poisson-type equation with constant coefficients. Therefore, diagonal scaling has no effect and the DSCG solver is equivalent to the straight CG solver. The numerical grids used in this series of tests were cubic lattices ranging from $(8)^3$ to $(128)^3$ nodes. Whence, the largest grid

has over two million nodes. The number of iterations (I) was defined as that required to decrease the L_∞-norm of error by six orders of magnitude. The approximately linear increase of $I(n)$, with respect to *uni-directional*, grid size, is in agreement with the theory presented above. Thus, it takes 105 iterations to solve the $(32)^3$ problem, and 'only' 375 iterations to solve the $(128)^3$ problem, which is 64 times larger.

Figure 2 Number of DSCG Iterations (I) vs. Uni-Directional Grid Size (n) in the Case of Constant Conductivity ($\sigma = 0$).

The convergence behavior of the DSCG solver is, however, more complicated than that suggested by Figure 2 alone (see Section 4). To illustrate this, we present in *Figure 3* several curves depicting the L_∞-norm of error *vs.* iteration count, for test problems with different degrees of heterogeneity, grid sizes, and log-conductivity structures. These test problems are summarily described in *Table 1*. The degree of heterogeneity is represented by σ, the standard deviation of log-conductivity (ℓnK). Note that grid sizes range from a few thousand nodes up to 7.6 million nodes. Descriptions of these or similar problems are also given in Section 7 and in *Appendix A (Table A1)*.

Figure 3 Convergence of DSCG Iterations (L_∞-Norm of Error *vs.* Iteration Count) for the Test Problems Described in Table 1.

In Figure 3, note the singular behavior exhibited by curve number 600B. Initially, convergence is slow as expected, due to the large unidirectional size of this grid ($n = 1001$). However, after about 700–800 iterations or so, the error drops quickly to machine precision and cannot decrease further. This illustrates the fact that the CG method always yields the exact solution, within machine precision, after N iterations *at most* (N is not very large for this problem). More generally, we expect that, for grids with very large aspect ratio, the number of iterations to achieve essentially exact solution is on the order of the number of nodes along the largest dimension of the grid. For instance, Test 600B has a quasi one-dimensional grid of size $N = 1001 \times 5 \times 5$, and it requires on the order of 1000 iterations (roughly) to achieve essentially exact solution.

In addition, comparing curves labeled 510 ($\sigma = 1$) and 520 ($\sigma = \sqrt{3}$) in Figure 3 indicates the influence of degree of heterogeneity (slower convergence). Comparing curves labeled 520 (gaussian distribution) and B521 (symmetric binary distribution) demonstrates the equally important influence of spatial structure. And finally, comparing curves number 520

(1 million nodes) and number 72A (7.6 million nodes) shows the influence of grid size (slower convergence). See also Section 7 for a brief discussion of the mass balance error incurred in each test problem, given the solution obtained at the last iterate.

The empirical convergence behavior reported in Figure 3 complements the theoretical estimate of convergence given earlier (Section 4), and can be used in assessing the number of iterations required for other, analogous flow problems. Let us summarize our observations in that respect. In some cases, one can hope to obtain essentially exact solutions within machine precision (Test 600B). For very large systems, the solution will usually be less accurate than machine precision. In the case of constant or mildly variable conductivity, the theoretical estimate may give sufficient indication on how CPU time grows with grid size for a given accuracy. For highly variable coefficients, the theory must be corrected based on empirical tests like those of Figure 3. For instance, it can be seen on Figure 3 that after 1000 iterations, the error was decreased by 12 orders of magnitude in a problem of moderate variability (Test 510), compared to 'only' 9 order of magnitude in a similar problem of larger variability (Test 520).

TABLE 1: BRIEF DESCRIPTION OF 'DSCG' TEST PROBLEMS

Test Number	**Ln(K) Spatial-Statistical Distribution**	**Grid size** $N = n_1 \times n_2 \times n_3$
600B	constant, $\sigma = 0$	$N = 1001 \times 5 \times 5$
510	Gaussian pdf, isotropic, $\sigma = 1$	$N = 101 \times 101 \times 101$
520	Gaussian pdf, isotropic, $\sigma = \sqrt{3}$	$N = 101 \times 101 \times 101$
521	Binary pdf, isotropic, $\sigma = \sqrt{3}$	$N = 101 \times 101 \times 101$
72A	Gaussian pdf, anisotropic, $\sigma = \sqrt{3}$	$N = 178 \times 120 \times 357$

6. Serial and parallel timings of BIGFLOW with DSCG solver. The performance of the DSCG-based code, expressed in CPU seconds, depends on grid size and number of iterations, and on machine-dependent additive and multiplicative factors. Following [1], timings can be expressed approximately in the form:

$$T(I, N) = (aI + b)N \tag{4}$$

where T is the total CPU time (seconds), "a" represents specific iterative work (seconds/iteration/million nodes), and "b" represents work spent outside the iterative solution process, or overhead (seconds/million nodes). As before, "I" is the number of iterations, and N is the number of nodes in multi-dimensional space. Note that "I" may be a pre-selected number of iterations, or alternatively, the number of iterations to decrease the error by a certain amount (say, six orders of magnitude). With the latter choice,

we have seen earlier that "I" is proportional to $N^{1/3}$ in the case of a cubic grid.

Serial Timings: For the DSCG-based code running serially on Cray-2 and Cray-YMP machines, we found empirically:

- $a_{(\text{Serial Cray 2})} = 0.48$ seconds per iteration per million nodes
- $a_{(\text{Serial Cray Y/MP})} \approx 0.20$ seconds per iteration per million nodes.

These constants were obtained from timings of several large test problems with randomly heterogeneous conductivities, most but not all of them involving cubic grids. The Cray dependency analyzer and optimizer, "fpp", was used on both machines. The "aggressive" optimization option was used on the Cray Y/MP. Note that the Cray Y/MP machine is faster than Cray-2 by a factor around 2.5 for these types of problems (in serial mode).

The serial Cray-2 timings were analyzed in detail using the "flowtrace" utility. It was found that compiling with the aid of the dependency analyzer decreased "a" by just a few percent, and that all inner loops vectorized with or without it. The overhead constant "b" was found to be sensitive to I/O formats: 28 seconds per million nodes with unformatted I/O's, compared to 116 seconds per million nodes with formatted I/O's.

Parallel Timings: Coarse-grained parallelization was studied on the Cray Y/MP8 by allowing the BIGFLOW code to run concurrently on k processors in dedicated mode ($1 \leq k \leq 8$). Again, the DSCG solver was used for solving random conductivity problems involving one to several million grid points. The BIGFLOW source code was not modified for multiprocessing. Instead, we let the Cray autotasking software perform the necessary code modifications and enhancements (we also used a compiler option to in-line the CG solver module). Estimates of speed-ups and of parallelizable fraction of code were obtained by comparing cumulated CPU times to wallclock times, and by applying Amdahl's law, as explained below. It is emphasized that our analysis concerns the total BIGFLOW code (main program and all modules).

Let k denote the number of processors, T_k the parallel CPU time or wallclock time for k concurrent processors running in dedicated mode, and T_1 the serial CPU time for a single processor. Define "f" as the fraction of parallelizable code, measured in serial CPU time units. Decomposing T_1 into parallelizable and non-parallelizable parts yields $T_1 = fT_1 + (1-f)T_1$. Neglecting any overheads, the parallel processing time is given by $T_k = fT_1/k + (1-f)T_1$. Hereafter, we refer to this parallel CPU time as "wallclock time". Now, the serial to parallel speed-up ratio is given by $r(k) = T_1/T_k$, and substituting T_k yields the following relation, known as Amdahl's law:

$$r(k) = \frac{k}{k(1-f)+f} \tag{5}$$

Note that the speed-up ratio is always greater than unity, at least in the absence of multitasking overhead (which we neglected).

Amdahl's law can be used ot obtain speed-ups when the parallelizable fraction "f" is known. Alternatively, Amdahl's law can also be inverted to obtain the fraction "f" from experimentally observed speed-ups. From the latter point of view, it is useful to introduce a new quantity $\rho(k) = r/k$, the average speed-up per active processor. Remarkably, it turns out that "f" can be expressed in the simple form:

$$f = \rho(k-1)/\rho(k). \tag{6}$$

In words, eq. (6) says that the parallelizable fraction of code can be expressed as the ratio of 'per processor' speed-ups obtained with consecutive values of k, that is, with $(k-1)$ and (k) processors, respectively.

Figure 4 depicts two speed-up curves $r(k)$ obtained for a 1 million node test problem (lower curve) and for a 7.6 million node test problem (upper curve). The test problems were described earlier in Table 1 (test numbers 510 and 72A, respectively). Figure 4 shows both actual speed-ups (circles), and analytical curves $r(k)$ (solid lines) from Amdahl's law (5). The latter curves were obtained after evaluation of the parallelizable fraction "f" defined above. The dashed straight line represents the ideal case $f = 100\%$, corresponding to a fully parallelizable code.

Under the Cray autotasking utility, we found $f = 82.5\%$ for the 1 million node problem (number 510), and $f = 89.1\%$ for the 7.6 million node problem (number 72A). The corresponding speed-up ratios when using all 8 processors are 3.59 and 4.53, respectively. The apparent sensitivity of speed-up curve to grid size and grid geometry may be due to trade-offs between vector and multiprocessing speed-ups (and, to a lesser extent, to the neglected multitasking overhead). The largest problem, with 7.6 million grid points, executed at about 750 MFLOPS (wallclock). Having identified certain ambiguities in the DSCG solver and the norm calculation modules, we expect to achieve faster rates in future, possibly well over 1 GFLOPS, by simple modifications of these modules.

We are now in a position to give a parallel timing for the BIGFLOW code running on Cray Y/MP8. Applying a speed-up ratio of approximately 3.5–4.5 to the Serial Cray Y/MP timings given earlier, we have, for 3D heterogeneous test problems of one million nodes or more:

$$a_{(\text{Parallel Cray Y/MP8})} \approx 0.04 - 0.06 \text{ secs per iteration per million nodes.}$$

Figure 4 Speed-Up Curves: Parallel/Serial Speed-Up Ratio (r) *vs.* Number of Cray Y/MP8 Processors Running Concurrently (k)

Comparisons with Connection Machine timings: It is interesting to note that the parallel performance of BIGFLOW on Cray is on the same order as that achieved by other DSCG-based, porous media flow codes running on the Connection Machine CM-2 ([9],[10]). Thus, for a homogeneous $(128)^3$ problem, a performance of 0.070 sec/iter/million nodes was achieved on 16K processors [9], and 0.043 sec/iter/million nodes on 32K processors [10]. However, it should be realized that timings like these can only be indicative. To put these code-to-code comparisons in perspective, the following should be noted:

- The CM-2 tests of references [9] and [10] were conducted on less than the full 64K processor machine;
- The CM-2 timings reported above were obtained for grid size equal to a power of two in each direction; CM-2 timings for grid sizes not exactly equal to powers of two were significantly slower than those reported here (a problem not encountered on Cray);

- The BIGFLOW code was written in Fortran 77, while the CM-2 codes were written in CM Fortran (a parallel language);
- The BIGFLOW code has somewhat more general functions, options, and structures than the CM-2 flow codes; the reader is referred to [8] for a complete documentation on BIGFLOW.

7. Homogenization study through large scale numerical experiments. In this section, we investigate the homogenization problem for randomly conductive porous media, based on synthetic flow fields simulated with the BIGFLOW code. The simulations were conducted as described in the previous sections, using mostly the DSCG solver (and a different solver in some cases). The 'experimental' approach can be summarized as follows. The idea is to generate a single realization of a random conductivity field in a large 3-D domain, to solve for total pressure and velocity on a finely discretized grid, and to recover the equivalent or effective conductivity (macroscale Darcy law) by 'measuring' spatial averages of flux and hydraulic gradient through the domain. The 'single realization - large domain' approach is particularly appropriate for statistically homogeneous flows, which are the object of this study.

Since the goal is to compare numerical 'experiments' with analytical solutions, it is important to preserve the correct statistical structure of the flow through the numerical solution process. The accuracy and performance of the iterative matrix solver were analyzed earlier. Concerning discretization errors, the influence of correlation scales, variability, and mesh size on truncation errors has been studied in detail [1]. Stochastic truncation errors were analyzed based on perturbations of the original stochastic PDE and of its discretized counterpart. See Chap. 5 of reference [1] for details, or reference [4] for an abstract. The reader may also consult reference [3] for empirical rules on spatial resolution requirements.

The flow problems used for studying the conductivity homogenization problem are listed in Appendix A (Table A1). Each flow experiment is conducted as follows. First, the 3-D rectangular or cubic domain (box) is dimensioned so that it contains many fluctuation scales of the pressure and velocity fields. It is then discretized into a regular mesh $(\Delta x, \Delta y, \Delta z)$ such that truncation errors and sampling errors remain reasonably low by all available estimates [see above-cited references]. Boundary conditions are chosen so that the mean flow is driven by a 'global' pressure gradient aligned with one of the axes of the box, and with mean flow velocity also aligned with that axis (X_i). These conditions are obtained by imposing two different values of total pressure on a pair of opposite faces of the box (fixed pressure - Dirichlet conditions), while keeping all other boundaries impervious to flow (zero flux - Neumann conditions).

The following method is used to 'measure' effective conductivity at the scale of the domain. The total discharge rate is computed through several

cross-sections orthogonal to the mean flow axis X_i (its value is independent of cross-section if mass is conserved). Dividing by cross-sectional area gives the mean flux density Q_i along X_i. The mean pressure gradient is calculated from the pressure imposed at the inlet and outlet boundaries and the length of the domain in that direction (L_i). Finally, these quantities are inserted in a hypothetical equation relating linearly the macroscale flux and the driving force (global pressure gradient), that is:

$$Q_i = -K_{ii}\{P(L_i) - P(0)\}/L_i \quad \textit{(no summation on i)} \tag{7}$$

In other words, we postulate the existence of macroscale Darcy law, and we use the 'measurements' to recover, from the postulated law, an effective conductivity coefficient (K_{ii}).

While the uniqueness of relation (7) remains to be verified, nevertheless, one can directly compare the single 'measured' value of (K_{ii}) to that predicted theoretically. Indeed, the postulated relation (7) is analogous to:

$$\langle Q_i \rangle = -K_{ij}\,\partial\langle P\rangle/\partial x_j \tag{8}$$

where Einstein's implicit summation rule on repeated indices is used ($j = 1, 2, 3$). This equation expresses an equivalent Darcy law in terms of ensemble means $\langle\bullet\rangle$ of flux and pressure gradient. It was developed in reference [11], based on first-order perturbation solutions of the original stochastic equations (Darcy's law and mass conservation at the local scale).

Now, the analogy between formulas (7) and (8) becomes clear in the case where the axes of statistical anisotropy of the random conductivity field are aligned with the natural axes of the box-shaped domain (case of Tests No. 72 A,B,C). In this situation, we are guaranteed by symmetry that the mean flux and mean gradient are both parallel to axis X_i, at least in the limit of infinite domain. With this proviso, the theoretical relation (8) becomes:

$$\langle Q_i \rangle = -K_{ii}\langle P\rangle/\partial x_i \quad \textit{(no summation on i)}, \tag{9}$$

which is clearly analogous to the measurement scheme (7). This argument demonstrates that equation (7) provides a scheme to measure the directional effective conductivity along axis "X_i".

Moreover, the notation "K_{ii}" also suggests that the quantity being measured is the i-th component of a tensorial effective conductivity in its principal system. Indeed, according to the theory developed in [11], the "K_{ij}" of equation (8) is a second rank symmetric tensor to first order in the log-conductivity variance. If this first-order theory were exact, then the result of our measurement would be exactly the i-th principal component of the tensor. Also, if the porous medium is statistically isotropic, and in the limit of infinite domain, the tensor must be isotropic by reason of statistical symmetry.

To give an example, consider the 3-D anisotropic problem labeled '72' (Appendix A). Three flow systems were simulated (72A, 72B, 72C), differing only by the direction of flow. The mean flow was parallel to the first, second and third principal axis (*resp.*). However, we did not attempt to test more completely the tensorial character of K. One way to do this would be to vary the direction of flow, not restricting it to be parallel to the principal axes. Similarly, for the isotropic tests No. 510, etc... (Appendix A), we assumed that only one direction of flow was needed since K_{11}, K_{22}, K_{33}, should all be equal in theory. That is, we did not attempt to test the isotropic character of the experimental flow field. This could be done by measuring K along more than one direction. Finally, it should be cautioned that multiple experiments may be needed in order to verify the uniqueness of equations (7)–(9) for each given type of porous medium, flow domain, and flow direction.

With these provisions, the results of effective conductivity 'measurements' are displayed graphically in *Figure 5* (top and bottom), including some comparisons with theoretical-analytical relations ([1],[6]) which are presented and discussed in *Appendix B*. The numerical flow problems have grid sizes ranging from 1 to 7.6 million nodes. The top part of Figure 5 displays results for statistically isotropic media, including one for which ℓnK is not gaussian. It can be seen that the theory of Appendix B is quite robust, being in agreement with measurements up to large degrees of heterogeneity (the largest ℓnK standard deviation is $\sigma = \ell n10 \approx 2.30$). In the bottom part of Figure 5, where theoretical K_{ii}'s are plotted directly against experimental ones, both isotropic and anisotropic problems were included. The only cases that exhibit significant discrepancies are the three anisotropic experiments corresponding to tests No. 72A, 72B, 72C. All other points correspond to isotropic problems. The special point (1,1) corresponds to the trivial case of zero variance, i.e. $K(x) = 1$.

Note that all simulations were conducted on Cray computers using, for most, the DSCG-based code (see previous sections and Appendix A, Table A1). However, we also added results from earlier simulations conducted for the isotropic test problems 510, 520, 530. They had been solved on a Cray-2 using the iterative Strongly Implicit Procedure (SIP) instead of DSCG. The random flow fields were fully analyzed in Chap. 6 of reference [1]. The plots of Figure 5, top and bottom, include isotropic effective conductivity results from the "SIP" simulations, along with those obtained with DSCG solver (the differences are minor).

Figure 5 Comparison of Numerical and Analytical Effective Conductivities. Top: Effective Conductivity K vs. Variance of ℓnK for Isotropic Media (Numerical = Symbols; Analytical = Solid Curve). Bottom: Analytical K_{ii} vs. Numerical K_{ii} for Isotropic and Anisotropic Media (Straight Line Represents Ideal Case of Perfect Fit).

Finally, note that the different test problems were not all solved with the same accuracy (this is briefly discussed in Appendix A). Since the net discharge rate through the domain must be zero in theory, we computed a relative mass balance error criterion as the ratio of net discharge rate to average ingoing–outgoing discharge rate. The relative error on the total mass rate was significant only for the anisotropic problem No. 72C, with flow orthogonal to strata. The relative error was about 2% or so. However, it was only 0.1% or less for the other flow directions (No. 72A & No. 72B). And for all other problems the error was insignificant, typically in the range 10^{-6} to 10^{-12} or even less.

From these results, we conclude that the theory described in Appendix B is probably most robust when the porous medium is statistically isotropic. The discrepancies observed in the anisotropic case may be due in part to the combined influence of insufficient domain size (a computational limitation) and relatively large degree of heterogeneity ($\sigma = \sqrt{3}$ for tests No. 72). The large heterogeneity, combined with anisotropy, may have caused non-negligible truncation errors in the numerical scheme, whence non-negligible local mass balance errors, although global mass balance error remains small. On the other hand, the observed discrepancies may also be due to shortcomings of the approximate theoretical formulas for such degrees of heterogeneity and anisotropy. The different possible causes for the discrepancies observed in the anisotropic case deserve to be examined more carefully, particularly in light of the remarkable fit obtained in the isotropic case.

8. Conclusions. The numerical experiments presented in this study demonstrate the feasibility of large, detailed 3-D simulations of heterogeneous porous media flow using an efficient numerical code (BIGFLOW) based on a sparse discretization and a dedicated, sparse, vectorizable iterative matrix solver (DSCG). Given that the code was ported without prior modification to Cray Y/MP8 under autotasking, the observed performance and speed-ups due to parallel processing are significant and encouraging. The ability to develop realistic flow simulations on large 3-D domains, and with almost arbitrary type of heterogeneity, should help make progress in the area of transport modeling. For instance, in applications such as disposal of high-level nuclear waste in geologic repositories, there is a need for accurate modeling of radionuclide transport over thousands of years or more, under realistic conditions taking into account natural heterogeneity. The predicted transport phenomena depend, of course, very strongly on the model used for producing velocity field(s).

Based on large simulations of flow through random porous media, empirical 'measurements' of effective macroscale conductivities were performed. The results were found to be in relatively good agreement with some proposed analytical solutions of the conductivity homogenization problem, given in closed form in Appendix B. The agreement was in fact

quite good for statistically isotropic media, less so for anisotropic media. More simulations are needed in order to test the uniqueness and tensorial character of the postulated macroscale Darcy law.

Finally, note that similar analyses have been undertaken for variably saturated or unsaturated flow processes through heterogeneous geologic media, including numerical studies with the flow code at hand ([1], [16]). However, because of the nonlinearity of Darcy's law in unsaturated media, the homogenization of conductivity, and the upscaling of Darcy's law, are even more difficult problems in this case. For data reviews, numerical results, and theoretical analyses on the macroscale conductive behavior of unsaturated media, see [1], [2], [7], [14], and [16], among others.

REFERENCES

[1] Ababou, R., 1988. *Three-Dimensional Flow in Random Porous Media*. Ph.D. Thesis, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts: 833 pp.

[2] Ababou, R., 1991. *Approaches to Large Scale Unsaturated Flow in Heterogeneous Stratified, and Fractured Geologic Media*. Report NUREG/CR-5743, U.S. Nuclear Regulatory Commission, Washington D.C.

[3] Ababou, R., D. McLaughlin, L.W. Gelhar, and A.F.B. Tompson, 1989. *Numerical Simulation of Three Dimensional Saturated Flow in Randomly Heterogeneous Porous Media*. Transport in Porous Media, 4: 549–565.

[4] Ababou, R., 1992. *Direct Simulation of Stochastic Darcy Flow and Perturbation Analysis of Finite Difference Errors*. Mini-Symposium on "*Darcy Flow in Composite and Random Porous Media*", ICIAM'91 Conf., Washington D.C., July 8–12, 1991. Abstract: *ICIAM'91 Proceedings, 2nd Internat. Conf. on Industr. & Appl. Math.*, R.E. O'Malley ed., SIAM, Philadelphia, 1992, Chap. 27, pp. 300–302.

[5] Ababou, R., B. Sagar, and G. Wittmeyer, 1992. *Testing Procedures for Spatially Distributed Flow Models*. Advances in Water Resources, 15, 181–198.

[6] Ababou, R., 1990. *Identification of Effective Conductivity Tensor in Randomly Heterogeneous and Stratified Aquifers*, in *Parameter Identification and Estimation for Aquifers and Reservoirs*, S. Bachu editor. Water Well Journal Publishing Company, National Water Well Association, Dublin, Ohio, pp. 155–157.

[7] Ababou, R. and T.-C.J. Yeh, 1992. *Random and Effective Conductivity Curves in Unsaturated Media: Anisotropy, Crossing Points, and Bounds*. Am. Geophys. Union, Fall Meeting, Special Session "*Effective Constitutive Laws for Heterogeneous Porous Media*", San Francisco, Dec. 1992. Abstract in EOS Transactions, Am. Geophys. Union, Oct. 27, 1992, p. 196.

[8] Ababou, R. and A.C. Bagtzoglou, 1993. *BIGFLOW: A Numerical Code for Simulating Flow in Variably Saturated Heterogeneous Geologic Media (Theory and User's Manual)*, Report NUREG/CR-6028, U.S. Nuclear Regulatory Commission, Washington D.C.

[9] Bagtzoglou, A.C., G.W. Wittmeyer, R. Ababou, and B. Sagar, 1992. *Application of a Massively Parallel Computer to Flow in Variably Saturated Heterogeneous Porous Media*, Proceedings, IXth Internat. Conf. *Computational Methods in Water Resources*, Denver, Colorado, June 9–12: 9 pp.

[10] Dougherty, D.E., 1991. *Hydrologic Applications of the Connection Machine CM-2*, Water Resources Research, 27(12): 3137–3147.

[11] Gelhar, L.W., and C.L. Axness, 1983. *Three Dimensional Stochastic Analysis*

of Macrodispersion in Aquifers, Water Resour. Res., 19(1), 161–180.

[12] Golub, G.H., and C.F. Van Loan, 1989. *Matrix Computations.* 2nd Edition. The John Hopkins University Press, Baltimore, Maryland: 642.

[13] Kohler, W. and G.C. Papanicolaou, 1982 *Bounds for the Effective Conductivity of Random Media*, pp. 111–130, in Macroscopic Properties of Disordered Media, Lecture Notes in Physics, No. 154, Springer-Verlag, 307 pp.

[14] Mantoglou, A. and L.W. Gelhar, 1987. *Effective hydraulic conductivities of transient unsaturated flow in stratified soils.* Water Resources Research 23(1): 57–67.

[15] Matheron, G., *Eléments Pour un Théorie des Milieux Poreux*, Masson et Cie, Paris, 164 pp., 1967.

[16] Polmann, D.J., D. McLaughlin, L.W. Gelhar, and R. Ababou, *Stochastic Modeling of Large-Scale Flow in Heterogeneous Unsaturated Soils*, Water Resources Research, 27(7), 1447–1458, 1991.

[17] Tompson, A.F.B., and L.W. Gelhar, 1990. *Numerical Simulation of Solute Transport in Three-Dimensional, Randomly Heterogeneous Porous Media.* Water Resources Research, 26(10): 2541–2562.

APPENDIX A:

Description of DSCG Random Flow Problems (Table A1)

Table A1 below gives a relatively detailed description of several test problems, most of which are examples of large, single realization flow simulations in 3-D random porous media. A brief summary of these problems was given in the text (Table 1), where they were used for convergence analysis of the DSCG solver (Figure 3).

Furthermore, these and other test problems were used in Section 7 (Figure 5) for studying effective macroscale conductivity, that is, the homogenization problem. Note that the flow problems of Table A1 were solved iteratively, using generally up to 1000 DSCG iterations. The latter choice was useful for studying convergence behavior (Figure 3). Additional random flow problems were used to produce the effective conductivity results of Figure 5. They were solved using the iterative SIP solver (Strongly Implicit Procedure) rather than DSCG, and the resulting flow fields were fully analyzed in reference [1].

All taken together, the test problems were not all solved in the same fashion or with the same accuracy. The initial guesses were different, in terms of their 'distance' to the final solution. Nevertheless, in almost every case, the solution obtained at the last iterate was deemed satisfactory in terms of error norm (the numerical noise was small compared to physical noise, in the root-mean-square sense), and in terms of total mass balance (the net discharge rate through the domain is very close to zero, as it should). The iteration errors depicted in Figure 3 are indicative of the accuracy with which most test problems were solved, including those not shown in the figure. Global mass balance errors, in terms of total mass rates, were discussed in Section 7.

Finally, it may be useful to comment briefly on the largest test problems, labeled (72A, 72B, 72C) in Table A1. These three tests correspond to three different flow experiments, where only the direction of mean flow is modified (A = flow parallel to X; B=flow parallel to Y; C = flow parallel to Z). That is, only boundary conditions were changed, while the same realization of random conductivity was used. The 3D grid comprised 7,625,520 nodes (roughly 7.6 million nodes), and each simulation required 98.0 Mega-Words of physical memory (roughly 100 MWords, or 800 MBytes).

TABLE A1:
Description of Some Flow Problems Solved With DSCG.

Test Pbs	*Ln(K) statistics*	*Grid & Mesh*	*Initial guess*
600 B	Constant, $\sigma = 0.$	N = $1001 \times 5 \times 5$ $\Delta x_i =$ (1,1,1)	Interior: P=0 Boundaries: P =+1 (left) P=-1 (right)
510	Gaussian p.d.f., isotropic $\lambda_i = (1, 1, 1)$, $\sigma = 1.$	N= $101 \times 101 \times 101$ $\Delta x_i =$ (1/3,1/3,1/3)	Linear P(x), solution of case $\sigma = 0.$
520	Gaussian p.d.f., isotropic, $\lambda_i = (1, 1, 1)$, $\sigma = \sqrt{3}.$	N= $101 \times 101 \times 101$ $\Delta x_i =$ (1/3,1/3,1/3)	Random P(x), solution of case $\sigma = 1.$
530	Gaussian p.d.f., Isotropic, $\lambda_i = (1, 1, 1)$, $\sigma = \ell n 10.$	N = $101 \times 101 \times 101$ $\Delta x_i =$ (1/3,1/3,1/3)	Random P(x), solution of case $\sigma = \sqrt{3}.$
521	Two-phase medium, symmetric binary pdf, indicator of gaussian isotropic field, $\lambda_i = (1, 1, 1)$, $\sigma = \sqrt{3}.$	N = $101 \times 101 \times 101$ $\Delta x_i =$ (1/3,1/3,1/3)	Linear P(x), solution of case $\sigma = 0.$
72 A,B,C (3 tests)	Gaussian p.d.f., anisotropic, $\lambda_i = (8/4, 4/4, 1/4)$ $\sigma = \sqrt{3}.$	N= $178 \times 120 \times 357$ $\Delta x_i =$ (1/3,1/2,1/6)	Linear P(x), solution of case $\sigma = 0.$

APPENDIX B:

Analytical Expressions for Effective Conductivity Tensor

Section 7 compares analytical and numerical effective conductivities for 3-D random porous media, the numerical results being obtained from certain 'measurements' on the simulated flow fields. We present here the analytical formulas used in the comparisons.

The proposed analytical model ([1],[6]) gives the components of a macroscale conductivity tensor in terms of the microscale conductivity field $K(x_1, x_2, x_3)$, under certain conditions of randomness, statistical homogeneity, and statistical anisotropy. The proposed relation is empirical, although specialized forms of it are confirmed by other, more fundamental results, including: (i) exact bounds; (ii) exact solutions in special cases involving lower dimensionality, statistical isotropy, symmetric distributions, and binary distributions; and (iii) approximate analytical solutions based on linearization and/or perturbation of the governing equations. Some of the latter results, which can be compared to the present analytical model, will be found in references [1], [11], [13], and [15], among others.

Underlying the proposed model is the assumption that spatial variability can be represented by a random function of space. Imperfectly stratified and anisotropic structures are described by means of directional fluctuation scales or correlation scales, while other features such as degree of variability, bimodality, etc, are conveyed by a probability distribution. The dimensionality of the flow system is also an important factor, therefore, the general case of a D-dimensional flow system will be considered ($D = 1, 2,$ or 3). The proposed model, then, postulates that effective conductivity is a second rank symmetric tensor, and expresses the principal components of this tensor by means of a power-average operator (somewhat similar to a Hölder norm):

$$\widehat{K}_{ii} = \langle K^{p_i} \rangle^{1/p_i} \qquad (i = 1, \ldots, D) \tag{1}$$

where the angular brackets $\langle \quad \rangle$ designate the operation of averaging. In this equation, the p_i's are directional averaging exponents. They are expressed in terms of the directional fluctuation scales ℓ_i, as follows:

$$p_i = 1 - \frac{2}{D} \frac{\ell_h}{\ell_i} \qquad (i = 1, \ldots, D) \tag{2}$$

where ℓ_h is the D-dimensional harmonic mean fluctuation scale:

$$\ell_h = \left[\frac{1}{D} \sum_{i=1}^{i=D} \ell_i^{-1} \right]^{-1} \tag{3}$$

Note that the averaging components are constrained to lie within the interval $[-1, +1]$, and that they sum up to $(D - 2)$.

To summarize, equations (1)-(2)-(3) give an analytical relationship for the D-dimensional effective conductivity tensor in terms of the single-point probability distribution, the principal directions, and the directional fluctuation scales of the microscale log-conductivity field. Note that the microscale data required for implementation of equations (1)-(3) are all of a statistical nature. For technical reasons, we prefer to use here the statistics of log-conductivity rather than conductivity.

The power-average effective conductivity tensor (1)-(3) can be expressed in closed form for several usual types of log-conductivity distributions, such as gaussian, binary, etc. In the case of a 'gaussian medium' with normally distributed ℓnK, applying equations (1)-(3) leads to:

$$\widehat{K}_{ii} = K_g \exp\left\{\frac{1}{2}\sigma^2\left[1 - \frac{2}{D}\,\frac{\ell_h}{\ell_i}\right]\right\} \qquad (i = 1\ldots, D) \tag{4}$$

when σ^2 is the log-conductivity variance, and K_g is the geometric mean conductivity. This relation was initially developed in reference [1] (Chap. 4, Eq. 4.48), in the equivalent form:

$$\widehat{K}_{ii} = (K_a)^{\alpha_i}\,(K_h)^{1-\alpha_i} \qquad (i = 1\ldots, D) \tag{5}$$

where $\alpha_i = (D - \ell_h/\ell_i)/D$, and K_a and K_h represent the arithmetic and harmonic mean conductivities, respectively.

Equations (4)–(5) were used in Section 7 to compute the 'theoretical' effective conductivities for all test problems having a gaussian ℓnK distribution (that is, all flow problems except for test No. 521). It should be emphasized that these equations yield effective conductivities identical with the extrapolated perturbation solutions of reference [11], except in the case of full 3-D anisotropy ($\ell_1 \neq \ell_2 \neq \ell_3$) where the expressions are close, but not identical in form. There are two main advantages to the present analytical model: it leads to closed form relations even in the case of 3-fold anisotropy; and it exhibits an explicit dependence on the (single-point) probability distribution of ℓnK.

For instance, a non-gaussian distribution of particular interest is the binary one. It can be used to represent a binary medium, made up of a mixture of two distinct conductive phases α and β, present in the proportions (ρ) and $(1-\rho)$ respectively. Phase α could be a sandstone porous matrix, and phase β a set of shale lenses or shale clast inclusions, for example. The single-point distribution of conductivity for this composite medium is of the form:

$$\begin{aligned} &\text{Prob}\,\{K(x_1, x_2, x_3) = K_\alpha\} = \rho \\ &\text{Prob}\,\{K(x_1, x_2, x_3) = K_\beta\} = 1 - \rho \end{aligned} \tag{6}$$

As before, we will assume as a first approximation that the spatial anisotropy of the random structure can be defined by three fluctuation scales ℓ_1, ℓ_2, ℓ_3.

Thus, specializing equations (1)-(3) for the binary distribution (6) gives:

$$\widehat{K}_{ii} = \{\rho K_\alpha^{p_i} + (1-\rho) K_\beta^{p_i}\}^{1/p_i} \qquad (i = 1, \ldots, D) \tag{7}$$

with averaging powers (p_i) as given previously in equation (2). In the case of a 3-D isotropic binary medium, let $D = 3$ and $\ell_1 = \ell_2 = \ell_3$. This yields $p_i = 1/3$ $(i = 1, 2, 3)$. Thus, inserting $p_i = 1/3$ in (7) gives the equation that was used in Section 7 to compute effective conductivity for the binary isotropic distribution (Test Problem No. 521).

For completeness, let us also examine the very special case of a 2-D, isotropic, binary medium. Let $\ell_1 = \ell_2$ for horizontal isotropy, and $D = 2$ for restriction to two-dimensional space, or equivalently $D = 3$ with $\ell_3 \to +\infty$ for horizontal flow through a vertically homogeneous medium. Either case yields $p_i \to 0$ for $i = 1$ and 2. Inserting this in (7) and using Taylor developments leads to:

$$\widehat{K}_{ii} = (K_\alpha)^\rho (K_\beta)^{1-\rho} \qquad (i = 1 \text{ and } 2) \tag{8}$$

where ρ represents the concentration of phase α, and $1-\rho$ the concentration of phase β.

Reference [6] discusses a hydrogeologic application of the proposed effective conductivity model, assuming a gaussian anisotropic medium. Briefly, the field study involved identifying hydraulic parameters from incomplete data collected in a heterogeneous aquifer. Given macroscale measurements of effective conductivity components, and partial measurements of microscale conductivities, the closed form relations (1)-(5) were used in an inverse fashion to evaluate some missing microscale statistics (the fluctuation lengths).