Efficient Stochastic Simulation of Groundwater Contaminant Transport

T.H. Skaggs and D.A. Barry

Department of Environmental Engineering
Centre for Water Research
University of Western Australia

Abstract. Due to the heterogeneity of natural groundwater systems, any quantitative description of aquifer hydraulic properties is subject to uncertainty. Consequently, prediction of groundwater contaminant transport is also subject to uncertainty. Stochastic approaches to transport simulation quantify this uncertainty in terms of random variables and processes. An important practical consideration in the application of such methods is their large computational cost. In recent years, the first-order reliability method (FORM) has been introduced as a possible technique for obtaining stochastic results with low computational expense. Specifically, the implementation of FORM known as advanced FORM (AFORM) has been shown to produce reasonably accurate results when applied to simple problems. However, recently published results indicate that the computational burden of AFORM can equal, or even exceed, that of Monte Carlo simulation when applied to groundwater contamination problems with a large number of variables. If FORM is to be a viable alternative, the computational costs of the method must be lowered. In this work, we propose a more efficient implementation of FORM. The primary numerical difficulty that arises in AFORM is locating the linearization point, a procedure that requires the solution of a non-linearly constrained optimization problem. We reduce the number of variables in the constraint by zoning the aquifer parameters during this stage of the calculations, resulting in an algorithm with lower computational costs. The new approach is shown to produce results that are nearly identical to those obtained with AFORM when applied to a one-dimensional transport problem. Future work will be aimed at generalizing the procedure described herein.

1 INTRODUCTION

Contaminant transport simulation is an important tool for managing groundwater resources. Transport modelling is used to predict the effects that future management practices and remediation efforts will have on water quality.

The geohydraulic and geochemical properties of an aquifer dictate how a contaminant will move in groundwater. Geohydraulic properties, such as hydraulic conductivity, are known to be spatially variable in natural groundwater systems. Since an exact description of this spatial heterogeneity is not possible, transport modelling is always subject to some uncertainty. In stochastic approaches to transport simulation, this uncertainty is quantified in terms of random processes and variables.

An important practical consideration in stochastic transport modelling is the associated computational costs. For many problems, the calculations are feasible only with the use of supercomputers. In this paper, we review some approaches to stochastic transport simulation and suggest an alternative implementation of one method, the first-order reliability method (FORM), that can lead to important computational savings.

2 STOCHASTIC TRANSPORT SIMULATION

Most groundwater contaminant transport models are of the form [*McLaughlin and Wood*, 1988]

$$\frac{\partial c}{\partial t} + a(c, \mathbf{x}, \mathbf{p}) = 0, \tag{1}$$

where c is the contaminant concentration, $\mathbf{x}$ is the spatial coordinate, t is time, and $\mathbf{p}$ is a vector of uncertain parameters that characterize the hydraulic and chemical properties of the aquifer and contaminant. The operator a is constructed from the spatial derivatives of c and possibly additional chemical reaction, sink, or source terms.

The uncertainty in $\mathbf{p}$ may arise from a number of sources including imperfect field measurements and spatially varying aquifer properties. So that the uncertainty may be treated quantitatively, we regard $\mathbf{p}$ as a random vector that is characterized by the probability density $f_{\mathbf{P}}(\mathbf{p})$.

Let $G(\mathbf{p}; \mathbf{x}', t')$ be the concentration predicted by the transport model at time t' and location $\mathbf{x}'$. Since $\mathbf{p}$ is random, G is also a random variable. The randomness of G suggests that any prediction we make about concentrations at t' and $\mathbf{x}'$ should be made in probablistic rather than deterministic terms. This requires knowledge of not only the mean value of G, but the higher moments as well. Ideally one could compute the entire cumulative distribu-

tion function (cdf) of G. Formally this is defined by

$$Pr[G \leq g] = \int_{G \leq g} f_{\mathbf{P}}(\mathbf{p}) d\mathbf{p}, \qquad (2)$$

where the integration is defined over the region of p-space where $G \leq g$. In general, the integral in (2) is difficult if not impossible to compute because it is hard to make explicit the area of integration and because of numerical difficulties associated with multifold integrals.

One approach to evaluating the integral in (2) is Monte Carlo simulation. In this method random samples are drawn from $f_{\mathbf{P}}(\mathbf{p})$ and for each realization G is computed. The probability $Pr[G \leq g]$ is simply the fraction of realizations that resulted in a computed concentration G less than or equal to g. For a sufficiently large number of realizations, the computed probability converges to the true solution. Although it is possible to construct an accurate cdf in this way, the large computational cost of repeatedly computing G makes implementation of the approach impractical without supercomputing facilities and, consequently, Monte Carlo simulation is used solely as a research tool and is not used for engineering or regulatory purposes [*Barry*, 1990].

As an alternative to Monte Carlo simulation, we may attempt to find an approximate solution of the integral in (2). To begin we assume that a hypersurface defined by $G = g$ divides the parameter space into two regions: region $\mathcal{L}$ where $G < g$ and region $\mathcal{G}$ where $G > g$. The hypersurface, referred to as the limit state surface, is again difficult to make explicit. Now, to motivate an approximate solution methodology, consider a special case when the solution of (2) is possible; namely, when $f_{\mathbf{P}}(\mathbf{p})$ is the multivariate standard normal distribution ($\mathbf{p} \sim N(\mathbf{0}, \mathbf{I})$) and the limit state surface is a hyperplane. Under these condtions $Pr[G \leq g]$ is $\Phi(\beta)$ [*Madsen et al.*, 1986], where $\beta = \alpha^T \mathbf{p}^*$, α is a unit vector normal to the hyperplane limit state surface and directed toward region $\mathcal{G}$, $\mathbf{p}^*$ is the point on the limit state surface closest to the origin, and $\Phi(\cdot)$ is the univariate standard normal cdf. This situation is illustrated in Figure 1 for the case of $\mathbf{p}$ being a 2×1 vector with elements p_1 and p_2 and $f_{\mathbf{P}}(\mathbf{p})$ being the bivariate standard normal distribution. The figure is a contour plot of G with only the $G = g$ contour being shown; for the case of two parameters, the limit state 'surface' is this contour. The point $\mathbf{p}^*$ is referred to as the design point. It is easy to show that $|\beta| = (\mathbf{p}^{*T}\mathbf{p}^*)^{1/2}$; that is, the absolute value of β is the distance from the limit state surface to the origin.

Given the relative ease with which the integral in (2) can be evaluated for $\mathbf{p} \sim N(\mathbf{0}, \mathbf{I})$ and a hyperplane limit state surface, a reasonable approach to more general problems is to transform (at least approximately) to this special case. In other words, transform $\mathbf{p}$ to the space of uncorrelated standard normal variables and then approximate the limit state surface with a first-order series expansion.

The parameter transformation is accomplished with the mapping

$$\mathbf{u} = T(\mathbf{p}), \qquad (3)$$

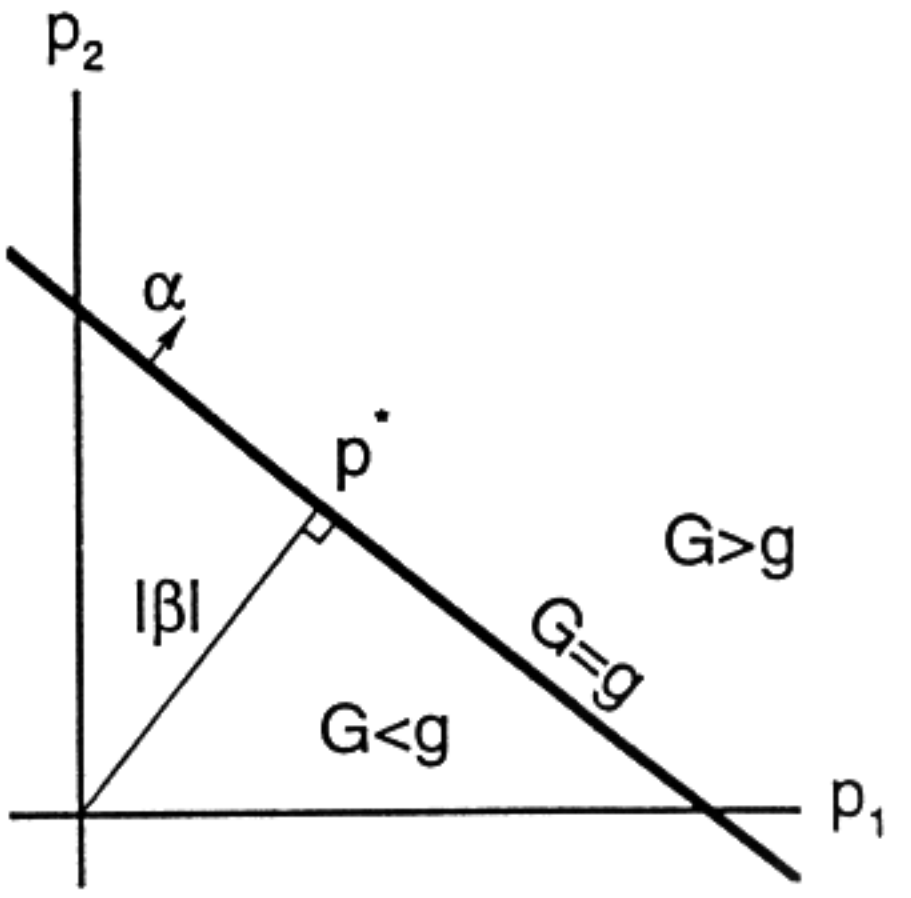

Figure 1: Illustration of a hyperplane limit state surface for the case of two parameters. Here, the 'surface' is simply the $G = g$ contour.

where the exact nature of T will depend on the distribution $f_{\mathbf{P}}(\mathbf{p})$ (for details, see *Madsen et al.* [1986]) and $\mathbf{u}$ is a vector of uncorrelated standard normal variates. This transformation also maps the limit state surface into u-space,

$$G(T^{-1}(\mathbf{u})) \equiv \widehat{G}(\mathbf{u}) = g. \qquad (4)$$

Likewise, the regions $\mathcal{L}$ and $\mathcal{G}$ are mapped into $\mathcal{L}_u$ and $\mathcal{G}_u$, respectively. In u-space, a first-order approximation to the limit state surface is

$$\widehat{G}(\mathbf{u}_0) + \nabla\widehat{G} \cdot (\mathbf{u} - \mathbf{u}_0) = g, \qquad (5)$$

where $\mathbf{u}_0$ is the point where $\widehat{G}$ is linearized (with corresponding point in p-space given by $\mathbf{p}_0 = T^{-1}(\mathbf{u}_0)$), $\nabla\widehat{G} = [\partial\widehat{G}/\partial u_1, \ldots, \partial\widehat{G}/\partial u_N]$ evaluated at $\mathbf{u}_0$, and N is the number of parameters. The gradient is evaluated as

$$\nabla\widehat{G} = \nabla G \cdot \frac{\partial T^{-1}}{\partial \mathbf{u}}, \qquad (6)$$

where $\nabla G = [\partial G/\partial p_1, \ldots, \partial G/\partial p_N]$ and $\partial T^{-1}/\partial \mathbf{u}$ is the Jacobian of the inverse transformation. The point on the linearized limit state surface closest to the origin (i.e., the design point) is

$$\mathbf{u}^* = [\nabla\widehat{G} \cdot \mathbf{u}_0 - \widehat{G}(\mathbf{u}_0) + g](\nabla\widehat{G} \cdot \nabla\widehat{G}^T)^{-1}\nabla\widehat{G}^T. \qquad (7)$$

The probability $Pr[G \leq g]$ is then computed as before, i.e. $Pr[G \leq g] = \Phi(\beta)$ with $\beta = \alpha^T \mathbf{u}^*$ and α directed towards $\mathcal{G}_u$.

The approximate method outlined above is known as the first-order reliability method (FORM). It was first used in structural engineering [e.g., *Madsen et al.*, 1986] and has more recently been applied to hydrologic problems [e.g., *Sitar et al.*, 1987; *Melching et al.*, 1990; *Cawlfield and Wu*, 1993; *Jang et al.*, 1994]. The accuracy of FORM is dependent on the suitability of the first-order approximation of the limit state surface.

There remains the question of choosing the linearization point $\mathbf{u}_0$. One possibility is to linearize about the mean value of the parameters, $\mathbf{u}_0 = T(\bar{\mathbf{p}})$. This is referred to as mean value FORM (MVFORM) and is expected to produce accurate probabilities near $G = G(\bar{\mathbf{p}})$ but will be less accurate near the tails of the distribution. Another approach is to linearize at the point on the limit state surface that is closest to the origin. This means that the linearization and design points will coincide ($\mathbf{u}_0 = \mathbf{u}^*$). A nonlinear limit state surface in u-space and its first-order approximation about the design point are shown Figure 2 for the case of $N = 2$. This approach is referred to

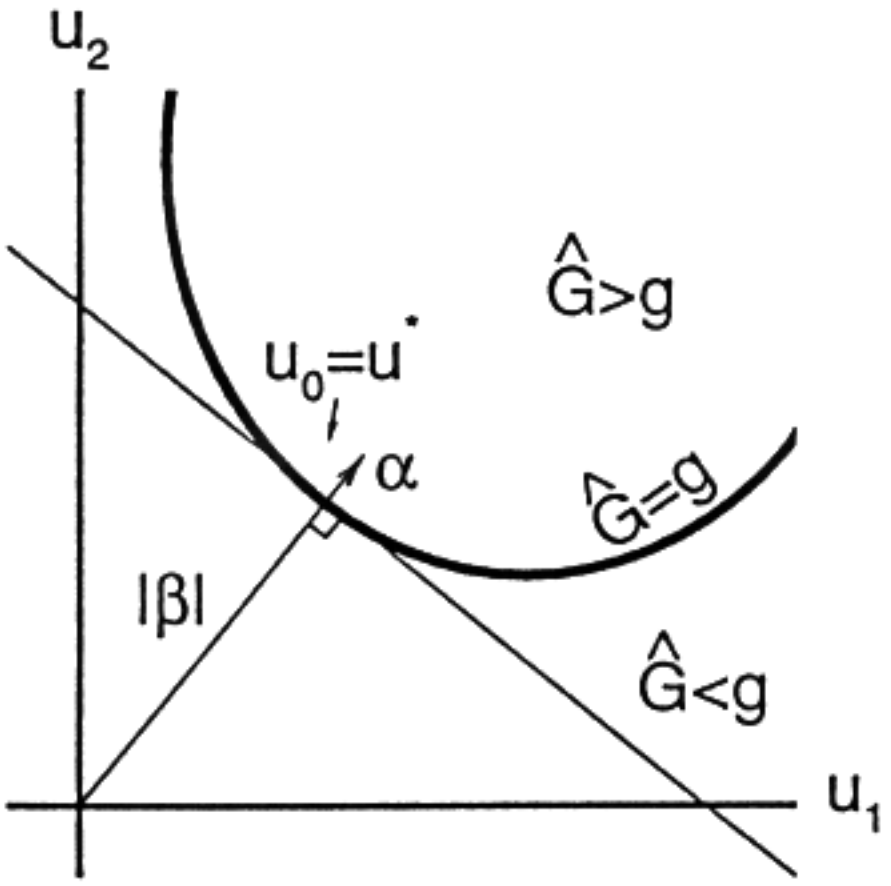

Figure 2: Limit state surface ($\hat{G} = g$) in u-space and the first-order approximation at $\mathbf{u}_0 = \mathbf{u}^*$.

as advanced FORM (AFORM) and is known to be more accurate than MVFORM over the entire distribution of G. If the limit state surface is a hyperplane in u-space, AFORM will produce exact results. However, this approach introduces additional computational requirements because for each value of g the design point must be located by solving the N-variable, nonlinearly constrained minimization problem

$$\min \mathbf{u}^T\mathbf{u} \tag{8a}$$

subject to

$$\hat{G}(\mathbf{u}) = g. \tag{8b}$$

Various iterative algorithms are available for solving such problems. However, all gradient-based algorithms require calculating $\nabla\hat{G}$ at each iteration and this is computationally burdensome (this contrasts with MVFORM which requires only a single calculation of $\nabla\hat{G}$). Recall from (6) that calculating $\nabla\hat{G}$ requires calculating of ∇G. The most straightforward way to compute the elements of ∇G is to perform repeated simulations while slightly perturbing one parameter at a time. The results of the simulations can then be used to compute finite difference approximations to the derivatives. If two-point difference quotients are used, a single calculation of ∇G requires $N+1$ runs of the simulation model. Since a minimization algorithm will take several iterations to converge, it is easy to see that for large N the solution of (8) will require many simulation runs and will be computationally expensive. *Jang et al.* [1994] analyzed a two-dimensional contaminant transport problem with a large number of parameters and reported that AFORM was at least as computationally expensive as Monte Carlo simulation. Special sensitivity methods can be used to improve the efficiency of computing ∇G [*Ahlfeld et al.*, 1988; *Skaggs and Barry*, 1995], but even then it is expected that AFORM will remain computationally expensive. The utility of using AFORM for stochastic analysis of moderate- to large-scale groundwater contaminant transport problems is questionable since it is not significantly more efficient than Monte Carlo simulation.

3 AN ALTERNATIVE FORM IMPLEMENTATION

The purpose of the present paper is to suggest an alternative implementation of FORM that aims to produce results that are comparable to those that are obtained with AFORM, but at less computational cost. The basic idea is to linearize about a point that is 'close' to the design point but which can be located with less effort than is required to solve the minimization problem (8). The new implementation is intended for problems that have a large number of parameters that arise from discretizing aquifer properties as part of a numerical scheme. When numerical methods such as finite differences or finite elements are used for transport simulation, the spatial domain is discretized creating a number of 'nodal' or 'elemental' parameters that represent spatially variable aquifer properties. The discretization, and consequently the resulting number of model parameters, is normally dictated by numerical requirements for convergence and stability. The implementation presented here is an AFORM-type algorithm in which a linearization point is found after groups of nodal and elemental variables have been collected together and set equal to one another. In essence, this 'zoning' allows us to find a linearization point using a much coarser grid than is required for the transport simulation. This leads to a minimization problem similar to (8), but with a constraint that is a function of fewer variables and, consequently, a constraint gradient that is easier to compute.

Although it is not necessary, for simplicity we assume in the discussion that follows all parameters in $\mathbf{p}$ originate from discretization of the aquifer (i.e., any other model parameters are assumed to be sure variables). We proceed by dividing the aquifer into zones in which all nodal and elemental variables in a given zone are taken to be equal to one another. Again for simplicity, we assume that each zone encompasses the same number of parameters (and nodes), $M = N/N_z$, where N_z is the number of zones and M is the number of parameters in each zone. Zoning the aquifer in this way is equivalent to partitioning the parameter vector such that

$$\mathbf{p} \equiv \mathbf{p}_z = [\mathbf{z}_1|\mathbf{z}_2|\dots|\mathbf{z}_{N_z}]^T,$$

where

$$\mathbf{z}_1 = [p_1, p_2, \cdots, p_M],$$
$$\mathbf{z}_2 = [p_{M+1}, p_{M+2}, \cdots, p_{2M}],$$
$$\vdots$$
$$\mathbf{z}_{N_z} = [p_{(N_z-1)M+1}, p_{(N_z-1)M+2}, \cdots, p_{N_zM}],$$

and setting

$$p_1 = p_2 = \ldots = p_M,$$
$$p_{M+1} = p_{M+2} = \ldots = p_{2M},$$
$$\vdots$$
$$p_{(N_z-1)M+1} = p_{(N_z-1)M+2} = \ldots = p_{N_zM}.$$

We can now define a linearization point in p-space as the the solution of a minimization problem that is constrained by a function of $\mathbf{p}_z$. Because this constraint is effectively a function of N_z variables, its gradient will be easier to compute than that of a constraint that is a function of $N > N_z$ variables. The linearization point $\mathbf{p}_0$ is defined as the solution of

$$\min\ (\mathbf{p} - \bar{\mathbf{p}})^T \mathbf{C}^{-1} (\mathbf{p} - \bar{\mathbf{p}}) \tag{9a}$$

subject to

$$G(\mathbf{p}_z) = g, \tag{9b}$$

where $\mathbf{C}$ is the parameter covariance matrix. The linearization point in u-space is then given by $\mathbf{u}_0 = T(\mathbf{p}_0)$, the design point is computed from (7), and the probability computed as before. Note that this is not the only possible way to define a linearization point using the zoning concept. For example, one could alternatively zone the parameters in u-space and proceed from there. Additional work is required to determine which is the preferred approach.

As $N_z \to N$, the zoning method outlined above becomes equivalent to AFORM. From a computational point of view, the smaller N_z the better. How small N_z can be made without significantly changing the results will depend on the problem. In the example presented below, $N_z = 1$ is used which reduces the number of variables in the constraint gradient from 41 to 1. This means that if a finite difference approximation is used as described above, the constraint gradient in (9) can be computed from 2 simulation runs whereas the gradient in (8) requires 42 runs.

4 EXAMPLE

To illustrate the new approach (referred to hereafter as ZFORM) and test whether it produces results that are consistent with AFORM, we consider the following one-dimensional problem in which non-reactive chemical transport is described by

$$\frac{\partial c}{\partial t} = D\frac{\partial^2 c}{\partial x^2} - v(x)\frac{\partial c}{\partial x}, \tag{10}$$

subject to

$$c(x,0) = 0, \qquad x \in [0,1], \tag{11a}$$
$$c(0,t) = 1, \qquad t \geq 0, \tag{11b}$$
$$\frac{dc(1,t)}{dx} = 0, \tag{11c}$$

where c is the concentration $[ML^{-3}]$, D is the effective dispersion coefficient $[L^2T^{-1}]$ (assumed to be a known constant), and v is the pore-water velocity $[LT^{-1}]$. The velocity is taken to be a lognormally distributed stationary random process with mean (of the log process) μ and covariance $\sigma^2 R$, where $R = \exp(-h/\lambda)$ is the correlation coefficient, λ is the correlation length ($\lambda << 1$), and h is the separation distance between any two points.

Equation 10 is evaluated numerically using a fully implicit finite difference scheme with spatial and temporal discretizations of $\Delta x = 0.025$ and $\Delta t = 0.001$, respectively. The discretized nodal velocities make up the elements of the parameter vector $\mathbf{p}$ with $N = 1/\Delta x + 1 = 41$. The distribution $f_{\mathbf{P}}(\mathbf{p})$ is the multivariate lognormal distribution with mean $\mu = [\mu, \ldots, \mu]$ and covariance matrix $\mathbf{C} = \sigma^2\mathbf{R}$, where $\mathbf{R}$ is the correlation matrix for $\mathbf{p}$ constructed using the negative exponential correlation coefficient R. Using $D = 0.02$, $\mu = 1.0$, $\sigma^2 = 0.09$, and $\lambda = 0.1$ we compute the concentration cdf at $x = x' = 0.7$ and time $t = t' = 0.2$. In the notation used above, $G(\mathbf{p}; x', t')$ is $c(0.7, 0.2)$.

As noted above, in implementing ZFORM we use $N_z = 1$. This means that the constraint (9b) is effectively a function of only one variable and will be satisfied at only one point. Thus, the 'minimization' problem (9) is reduced to simply solving

$$G(\rho) = g, \tag{12}$$

where $\rho \equiv p_1 = p_2 = \ldots = p_N$. Equation (12) can be solved using Newton-Raphson iteration. In this problem, the key computational difference between ZFORM and AFORM is that ZFORM requires repeated calculation of the derivative $\partial G/\partial\rho$ to solve (12), whereas AFORM requires repeated calculation of the 1×41 gradient vector $\nabla\widehat{G}$ to solve (8).

In Figure 3, the computed concentration cdf using MVFORM, AFORM, and ZFORM are shown along with the results of 10,000 Monte Carlo simulations (which may be considered the true solution for this problem). For reference, the computed concentration using the mean velocity at each node is $G(\bar{\mathbf{p}}) = 0.075$. For small g, the AFORM and ZFORM results are in good agreement with the Monte Carlo results whereas MVFORM predicts a higher probability. Near $g = 0.06$, all four methods are in agreement. Between $g = 0.06$ and $g = 0.1$, the three FORM methods are in agreement but their computed probability is somewhat lower than the Monte Carlo results. Beyond this, MVFORM deviates from AFORM and ZFORM and approaches the asymptotic $Pr[G \leq g] = 1$ much faster. The curvature of the AFORM and ZFORM results beyond $g = 0.1$ is similar to the Monte Carlo results, but the AFORM and ZFORM computed probabilities are consistently lower than those of the Monte Carlo simulations. Most important to the present study is the very close

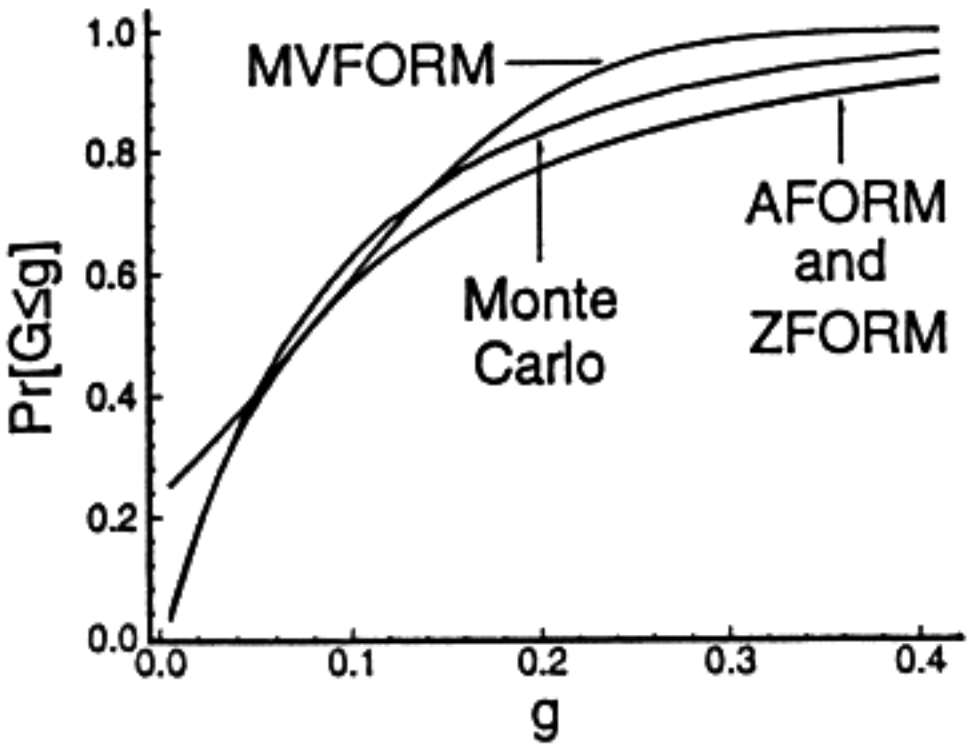

Figure 3: Computed concentration cdf using Monte Carlo simulation and three FORM implementations

agreement of the AFORM and ZFORM results across the distribution. This suggests that ZFORM may be a computationally efficient alternative to AFORM.

5 SUMMARY AND FUTURE DIRECTIONS

An important practical issue in stochastic modelling of groundwater contaminant transport is the high computational costs. For example, the high cost of Monte Carlo simulation makes its use impractical for many applications. This paper has reviewed some approaches to stochastic simulation and described a new computationally efficient implementation of one method, the first-order reliability method (FORM). The new approach involves zoning aquifer parameters during the first stage of the FORM calculations, creating a significant computational advantage over the standard FORM implementation (AFORM) when there is a large number of model parameters. In a simple example transport problem, the new method was shown to produce results that are nearly identical to those obtained with AFORM. This suggests the improved computational efficiency does not come at the expense of accuracy.

Further work remains to be done in generalizing the new method described herein. The results for the simple one-dimensional example are encouraging, but clearly more rigorous testing needs to be done, including two- and three-dimensional simulations. Additionally, we assumed zones that contain an equal number of aquifer parameters, suggesting an arbitrary zoning process where the only consideration was the number nodes in each zone. This seems appropriate for problems where water flow is uniform in the average, but for problems with complicated geometries and flow fields, it may be that an alternative zoning is preferred. These issues are the subject of our ongoing research.

6 REFERENCES

Ahlfeld, D.P., J.M. Mulvey, G.F. Pinder, and E.F. Wood, Contaminated groundwater remediation design using simulation, optimization, and sensitivity theory, 1. Model development, *Water Resour. Res.*, 24(3), 431-441, 1988.

Barry, D.A., Supercomputers and their use in modeling subsurface solute transport, *Rev. Geophys.*, 28(3), 277-295, 1990.

Cawlfield, J.D., and M.-C. Wu, Probabilistic sensitivity analysis for one-dimensional reactive transport in porous media, *Water Resour. Res.*, 29(3), 661-672, 1993.

Jang, Y.-S., N. Sitar, and A. Der Kiureghian, Reliability analysis of contaminant transport in saturated porous media, *Water Resour. Res.*, 30(8), 2435-2448, 1994.

Madsen, H.O., S. Krenk, and N.C. Lind, *Methods of Structural Safety*, Prentice-Hall, Inc., 403 pp., New Jersey, 1986.

McLaughlin, D., and E.F. Wood, A distributed parameter approach for evaluating the accuracy of groundwater model predictions, 1. Theory, *Water Resour. Res.*, 24(7), 1037-1047.

Melching, C.S., B.C. Yen, and H.G. Wenzel, Jr., A reliability estimation in modeling watershed runoff with uncertainties, *Water Resour. Res.*, 26(10), 2275-2286, 1990.

Sitar, N., J.D. Cawlfield, A. Der Kiureghian, First-order reliability approach to stochastic analysis of subsurface flow and contaminant transport, *Water Resour. Res.*, 23(5), 794-804, 1987.

Skaggs, T.H., and D.A. Barry, Sensitivity coefficients of a time-continuous groundwater contaminant transport model, in *Computational techniques and applications, CTAC-95; Proceedings of the seventh biennial conference on computational techniques and applications*, R. May and A. Easton eds., in press.