ELSEVIER
Journal of Contaminant Hydrology 33 (1998) 347–359
JOURNAL OF Contaminant Hydrology

Limitations in recovering the history of a groundwater contaminant plume

Todd H. Skaggs [a,*], Z.J. Kabala [b]

[a] *U.S. Salinity Laboratory, 450 W. Big Springs Rd., Riverside, CA 92507, USA*
[b] *Department of Civil and Environmental Engineering, Duke University, Durham, NC 27708, USA*

Received 24 February 1997; revised 7 March 1998; accepted 18 March 1998

Abstract

Groundwater contaminant transport is a dispersive process and consequently there are limits to what may be learned about a contaminant's origins (history) from measurements of its present spatial distribution. The extent of these limitations in a particular case depends on a number of factors, including the dispersive properties of the transport medium, the accuracy of the measured plume, and the interval of time over which the plume has evolved. We propose to characterize the limitations that will be encountered in a given transport medium by performing numerical simulations that assess our ability to recover various 'test functions', i.e., hypothetical contaminant release functions that are designed to provide insight into the effect that a transport medium's dispersive properties have on the recoverability of a plume history. Knowing our ability to recover test functions that differ, for example, in the time span over which the release occurs or in the amount of mass released allows us to make inferences about our ability to recover the history of an arbitrary plume in a transport medium with dispersive properties similar to those used in the simulations. Our ability to recover a test function is quantified using a Monte Carlo methodology. The method is demonstrated for the case of Gaussian test functions.

Keywords: Ground water; Chemical dispersion; Monte Carlo analysis; Inverse problem

1. Introduction

In many groundwater pollution scenarios, environmental regulators are confronted with the task of distributing liability among parties who are potentially responsible for a contaminant plume whose spatial and temporal origins are unknown. Numerical models

* Corresponding author.

PII: S0169-7722(98)00078-3

of contaminant transport are sometimes used to assist in re-creating the history of a plume. This type of modelling has been limited primarily to the realm of environmental consultants, but some work does exist in the open literature (Gorelick et al., 1983; Uffink, 1989; Bagtzoglou et al., 1992; Wagner, 1992; Skaggs and Kabala, 1994, 1995; Woodbury et al., 1996; Snodgrass and Kitanidis, 1997).

In previous work (Skaggs and Kabala, 1994) we discussed the ill-posed nature of recovering a plume's history from measurements of its current spatial distribution. Issues related to the problem's ill-posedness include the nonuniqueness of the recovered plume history and the presence of numerical instabilities that are associated with a high sensitivity to measurement and round-off errors. In numerical simulations of one dimensional convective-dispersive transport, we used Tikhonov regularization to recover the temporal release history (source function) of a plume that had originated from a known location. From the simulations we concluded that as long as a plume is not significantly dissipated, it should be possible to recover its release history even when the plume measurements contain moderate random measurement errors; when a plume is more dispersed, however, the presence of even moderate noise in the data significantly reduces the accuracy of the recovered release history.

The simulations of Skaggs and Kabala (1994) illustrate that the dispersive nature of contaminant transport causes a loss of information about a plume's release history as the plume travels forward-in-time. The loss of information results in a recovered release history that does not contain all the details of the true release history. A release history that is recovered over a short interval backward-in-time generally smears only the higher frequency components of the true history, but a release history recovered over a long interval may smear even the lowest frequency components of the true release.

In general there is always a large (possibly infinite) number of solutions that will reproduce the data within measurement error. The Tikhonov regularization methodology seeks a solution that reproduces the observed data and, at the same time, satisfies a smoothness criterion. Since the Tikhonov procedure always produces a recovered release history that accurately reproduces the data, there is the practical problem of assessing whether the recovered history is accurate at the scale of interest, or whether the details of the true release have been lost. Looked at another way: how can we tell when a particular plume has evolved to the point where we can no longer recover accurately its history? Suppose, for example, we know that a plume has originated from a particular site within the last twenty years and that the site has had four different owners over that time. If our recovered release history indicates that there was a slow and steady contaminant release over most of the twenty years, how do we know that this is the true release history and not a smeared version of, say, a catastrophic release that occurred in the tenth year?

What is required is a method for determining the time scale that is resolved by a release history recovered over a given time interval in a given transport medium. Using the previous twenty year example, is the resolution of a regularized release history going to be adequate for estimating the contaminant release in a given 5 year period, or is the resolution, diminished because of plume dispersal, going to be insufficient for making this determination. The objective of this paper is to develop a methodology for addressing these questions.

2. Regularized release history of a contaminant plume

In this section we briefly review the Tikhonov regularization methodology used by Skaggs and Kabala (1994) to recover the release history of a contaminant plume.

The one-dimensional transport of a dissolved, conservative contaminant through a saturated homogeneous porous medium may be described the convection–dispersion equation (CDE). The CDE, subject to the initial and boundary conditions

$$c(x,0) = 0,\ x \geq 0,$$

$$c(0,t) = c_{\text{in}}(t),\ t \geq 0$$

$$c(\infty,t) = 0,\ t \geq 0$$

is expressible in integral form as

$$c(x,t) = \int_0^t c_{\text{in}}(\tau)\, f(x,t-\tau)\mathrm{d}\tau \tag{1}$$

where c is the solution concentration [ML^{-3}], $c_{\text{in}}(t)$ is the contaminant source function at the inlet boundary [ML^{-3}], and f is

$$f(x,t) = \frac{x}{\sqrt{4\pi D t^3}} \exp\left(-\frac{(x-vt)^2}{4Dt} \right)$$

where v [LT^{-1}] is the pore-water velocity and D [$\mathrm{L}^2\mathrm{T}^{-1}$] is effective diffusion–dispersion coefficient.

If a contaminant plume's spatial distribution is observed at time $t = T$, we may reformulate (Eq. (1)) as the system of equations

$$\hat{c}(x_j;T) = \int_o^T c_{\text{in}}(\tau)\, f(x_j,T-\tau)\mathrm{d}\tau + \epsilon_j,\ j = 1,\ldots,n \tag{2}$$

where $\hat{c}(x_j;\ T) = c(x_j;\ T) + \epsilon_j$ is the measured concentration at $x = x_j$ and $t = T$, ϵ_j is a random measurement error associated with the plume measurement at x_j, and n is the number of plume measurements. Applying numerical quadrature to the integral in (Eq. (2)) leads to:

$$\hat{\boldsymbol{c}} = F\boldsymbol{c}_{\text{in}} + \boldsymbol{e} \tag{3}$$

where the n-length vector $\hat{\boldsymbol{c}}$ has elements $\hat{c}_j \equiv \hat{c}(x_j;\ T)$, the m-length vector $\boldsymbol{c}_{\text{in}}$ has elements $c_{\text{in;k}} \equiv c_{\text{in}}(t_k)$, m is the number of grid points used in the quadrature, the $n \times m$ matrix $\mathbf{F}$ has elements $F_{jk} \equiv w_k f(x_j,\ T - t_k,\ w_k)$ is the quadrature weight, and the n-length vector $\boldsymbol{e}$ has elements $e_j \equiv \epsilon_j$. If estimates of the transport parameters v and D are available the contaminant release history at $x = 0$ over the time interval $0 \to T$ would be recovered by solving (Eq. (3)) for $\boldsymbol{c}_{\text{in}}$. However, (Eq. (3)) is known to be a very ill-conditioned system of equations (Skaggs and Kabala, 1994) and the solution is extremely sensitive to the random errors ϵ_j. In other words, slight changes in the measured plume may result in drastic changes in the recovered contaminant release history.

Skaggs and Kabala (1994) applied Tikhonov regularization (Tikhonov and Arsenin, 1977) to obtain an approximate solution to (Eq. (3)). Tikhonov regularization mitigates the solution's sensitivity to errors by requiring that the solution satisfy a smoothness criterion as well as reproduce the data. Specifically, for the plume history problem, the regularized release history, $c_{\rm in}^{\rm reg}$, is the source function that minimizes

$$V = \|\hat{\boldsymbol{c}} - \mathbf{F}\boldsymbol{c}_{\rm in} a\|^2 + \alpha^2 \|\mathbf{H}\boldsymbol{c}_{\rm in}\|^2 \tag{4}$$

subject to the constraint

$$\boldsymbol{c}_{\rm in} \geq 0$$

where $\|\cdot\|$ is the Euclidean norm, $\mathbf{H}$ is the regularization operator, and α is the regularization parameter. The matrix $\mathbf{H}$ is specified so that $\mathbf{H}\boldsymbol{c}_{\rm in}$ is equal to a divided difference approximation of the second time derivative of $c_{\rm in}(t)$.

Minimizing V requires finding a regularized release history $c_{\rm in}^{\rm reg}$ that optimally minimizes both terms on the right side of (Eq. (4)). The first term measures the accuracy with which $\boldsymbol{c}_{\rm in}^{\rm reg}$ reproduces the observed plume whereas the second term measures the smoothness of $\boldsymbol{c}_{\rm in}$ as quantified by its second derivative. The regularization parameter α determines the relative weight given to these two terms in minimizing the overall functional and is said to determine the regularization strength of the solution. When α is small, more weight is given to the fit of $\mathbf{F}\boldsymbol{c}_{\rm in}^{\rm reg}$ to the measured plume and the regularized release history is therefore more sensitive to the data (and to the errors contained in it); when α is large, more weight is given to the smoothness of $\boldsymbol{c}_{\rm in}^{\rm reg}$ and the regularized release history is less sensitive to the data. Finding an optimal value for the regularization parameter α implies finding an appropriate balance between the smoothness of the solution and its fit to the data. In this sense the regularized release history is a 'best possible solution' to a problem whose solution is non-unique. We refer the reader to Skaggs and Kabala (1994) and Provencher (1982a,b) for more details on the solution technique, as well as information on the procedure used to determine a value for the regularization parameter α.

3. A motivating example

We briefly review some results from Skaggs and Kabala (1994) that provide the motivation for the remainder of this paper. For details on the simulations from which Fig. 1a, b, and c were produced, we refer the reader to the discussion of Fig. 4 and Fig. 6 of Skaggs and Kabala (1994).

Fig. 1a shows two snapshots of a contaminant plume's spatial distribution taken at two different times, $t = T = 300$ and $t = T = 600$ (arbitrary units of time). The solid line

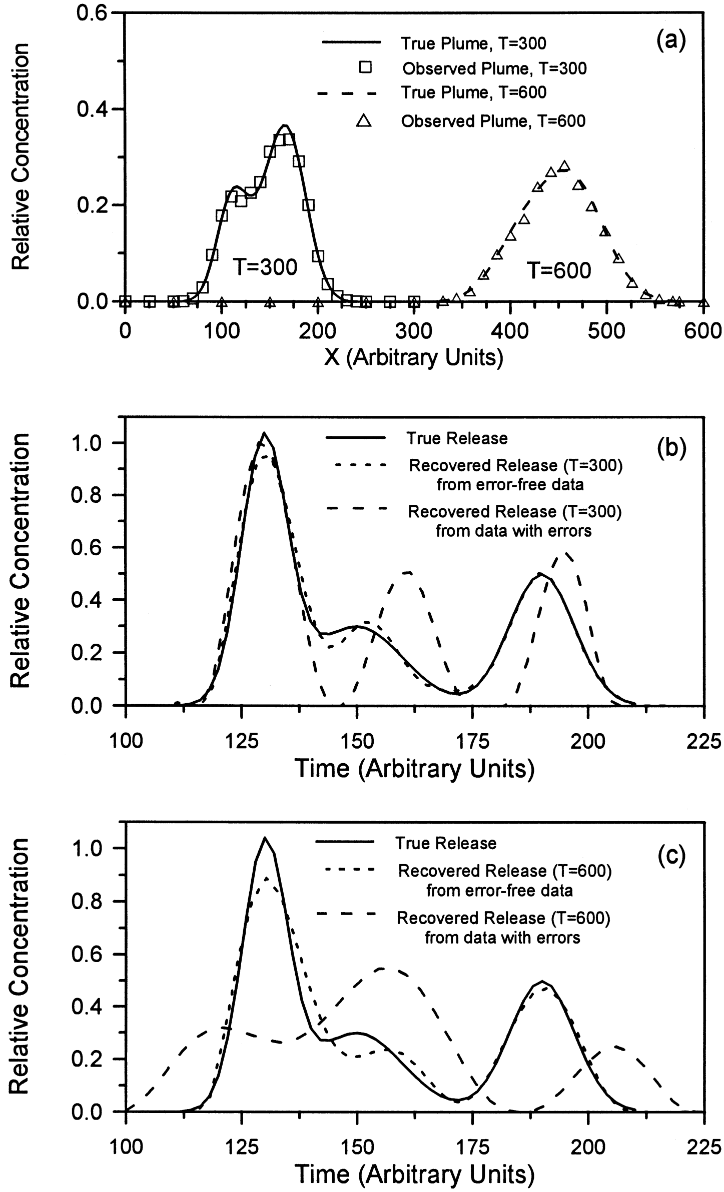

Fig. 1. (a) Observed plume at times $t = 300$ and $t = 600$; (b) release histories recovered from data at $t = 300$; (c) release histories recovered from data at $t = 600$.

is the true plume $c(x;\, T)$ at time $t = T = 300$, the squares are the plume measurements $\hat{c}(x_j;\, T) = c(x_j;\, T) + \epsilon_j$ at $t = T = 300$, the dashed line is the true plume at $t = T = 600$, and the triangles are the plume measurements at $t = T = 600$. Two additional data sets that we consider are not pictured; they are measurements of the same plumes with the same sampling grid, but they do not contain any random errors.

Fig. 1b contains the regularized release histories $\boldsymbol{c}_{\mathrm{in}}^{\mathrm{reg}}$ obtained from the two data sets measured at time $t = T = 300$. The dotted line is the regularized solution obtained from the error-free data set, the dashed line is regularized solution obtained from the data set containing errors, and the solid line is the true release history. The regularized release history obtained from the exact data is very close to the true release history. The regularized release history obtained from the data containing measurement errors is also reasonable—the first peak is recovered very accurately, the second and third peaks are overshot somewhat, and the interval of time over which contamination occurred is found almost exactly.

Fig. 1c shows the regularized release histories obtained from the data sets measured at time $t = T = 600$ using the same procedure as with the time $t = T = 300$ data. The regularized release history found from the error-free data set (dotted line) is again very close to the true release history (solid line). However, the regularized release history found from the data containing measurement errors (dashed line) is very poor, containing none of the details of the true release history. At the later time $T = 600$, the combined effects of dispersion and measurement errors have reduced the information content of the data to the point where it is no longer possible to recover the contaminant plume's history.

We may conclude, then, that under favorable physical conditions it should be possible to recover a plume history adequately, but if a plume is overly dissipated, or if the measurement errors are too large, it will not be possible. This raises the practical problem of identifying into which category a particular plume falls when we do not know the true release history, and this is the problem that we address in the remainder of the paper.

4. Monte Carlo methodology

To reiterate, as a plume evolves forward-in-time, it disperses and, as a result, information about the plume's release is lost. Consequently, there are limitations to what can be learned about a contaminant's history from observations of its present spatial distribution, and these limitations become more severe the further a plume travels forward-in-time.

The problem we wish to address may be stated as follows: over what interval of time is it reasonable to recover a contaminant release history in a given transport medium, and at what time-scale can we expect to recover the details of the release? For instance, is it reasonable to expect that we can recover a 5 year contaminant release history in such detail that we can report with confidence the amount of contaminant release in a particular month, or is the loss of information due to dispersion such that at best we can provide an estimate for the release in a particular year? Or has the plume dissipated to the point where we can say almost nothing at all about its release?

We propose to address this problem by performing numerical simulations of the transport medium and assessing our ability to recover various 'test functions', i.e., hypothetical contaminant release functions which are designed to provide insight into the effect that a transport medium's dispersive properties have on the recoverability of a plume history. Knowing our ability to recover test functions that differ, for example, in the time span over which the release occurs or in the amount of mass released allows us to make inferences about our ability to recover the history of an arbitrary plume in a transport medium with dispersive properties similar to those used in the simulations.

The test function simulations are carried out using the test release function in (Eq. (1)) to calculate a true plume $c(x_j;\, T)$. Random errors are added to create a measured plume $\hat{c}(x_j;$ T), and from this synthetic plume a regularized release history is calculated. The test functions we use are Gaussian releases, i.e.;

$$c_{\text{in}}^{\text{test}}(t) = \frac{M}{s(2\pi)^{1/2}} \exp\left(\frac{-(t-t_0)^2}{2s^2} \right) \tag{5}$$

where s is a measure of the spread of the release function about the mean release time t_0 and M is proportional to the amount of mass in the release. Varying the parameter s of the test function changes the time span over which the contaminant release occurs, with almost all of the release occurring within $t_0 \pm 2s$. We consider $4s$ to be the time scale that must be resolved to recover a particular test function. Varying M also changes the character of the test function. Eq. (5) is not the only possible test function. Depending on the situation, it may be of interest to consider other functions, such as two-peaked functions that differ in the separation distance between the peaks. For now we will consider only the Gaussian test functions. Fig. 2 shows four sample test functions with varying s and M.

In addition to s and M, the measurement errors and the particular error realization used to create $\hat{c}(x_j;\, T)$ influences our ability to recover a test function. We therefore

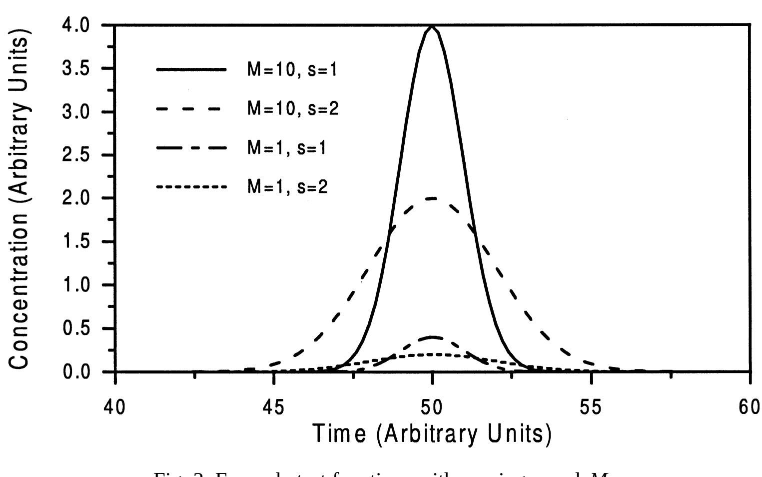

Fig. 2. Example test functions with varying s and M.

quantify our ability to recover a particular test function using a Monte Carlo methodology. For each test function, a true plume $\hat{c}(x_j; T)$ is calculated for the time of interest T. Fifty different measured plumes are created by adding a different realization of the random measurement errors to the true plume. We then calculate regularized release histories for each of the 50 synthetic plumes. The accuracy of each of the 50 trials is measured by comparing the regularized peak concentration with the true peak concentration—if it is within a specified tolerance (we arbitrarily use $\pm 25\%$ of the true peak concentration), we consider the test function to be accurately recovered.

The fraction of the 50 Monte Carlo simulations that is correctly recovered provides information about our ability to re-create the history of an arbitrary contaminant plume in a transport medium with dispersive characteristics similar to those used in the Monte Carlo analysis. Specifically, the percentage of the 50 trials that is accurately recovered is a measure of the likelihood that a regularized release history, calculated over the time interval $T - t_0$ for a plume with the same measurement error statistics as those used in the simulations, is accurate at the time scale characteristic of the test function. For instance, if the percentage of accurate recoveries is, say, 10, then we know that it is not reasonable to expect that a plume history may be recovered with the time scale resolution indicated by the test function; if the percentage is 90 or 100, however, we know that we may put a great deal of confidence in the regularized release history.

To illustrate the Monte Carlo methodology, we consider a transport medium described by (Eq. (1)) with $v = 1$ and $D = 1$ and implement the methodology using the test function (Eq. (5)) with $t_0 = 50$. The choice of centering the test function about $t_0 = 50$ is arbitrary and is made so that the boundary at $t = 0$ does not interfere in the recovery of the test function. The time that is of interest is the plume evolution time, i.e., the interval between the release of the plume and the observation of the plume, $T - 50$.

We will assume that the measurement errors ϵ_{j} are constant relative errors, i.e.:

$$\epsilon_j = \delta \xi_j c_T^{\mathrm{true}}(x_j) \tag{6}$$

where ξ is a normally distributed random deviate with unit variance ($\xi \sim N(0,1)$) and the scaler δ determines the magnitude of the errors. Eq. (6) may not be a reasonable structure for the error in a particular situation; indeed, it may be that the appropriate error structure or error magnitude is unknown. In this case, it would be necessary to perform separate Monte Carlo analyses for different error structures and error magnitudes and to consider these results when assessing the likelihood that an accurate release history can be obtained.

For the first test case, consider a test function with $s = 2$, $M = 1$ and plume measurement errors with magnitude given by $\delta = .01$. The spacing of the plume measurement grid is $x_{j+1} - x_j = 2.5$ units of space, with the total number of measurements being the number that is required to sample the whole plume. If we were performing these simulations to evaluate an actual plume, we would want to use a sampling grid that is similar to the one that would be used to measure the plume; here we have used an abundance of measurements so as to assure that the resolution of the sampling scheme is not a limiting factor in our solution procedure. Fifty Monte Carlo simulations were performed for plume evolution times $T - 50 = 25$, 50, 75, 100, 150, 200, and 250 and the resulting regularized release histories of the individual trials are plotted along with the true test function in Fig. 3. Fig. 3a is a plot of the true test function which has a peak concentration of 0.2; the horizontal lines at 0.25 and 0.15 indicate $\pm 25\%$ of the peak concentration. We consider any Monte Carlo trial that results in a regularized release function with a peak that falls in the region delineated by these lines to be a successful recovery of the test function. Fig. 3b is for the shortest plume evolution time $T-50=25$ and 100% of the Monte Carlo trials resulted in regularized release functions that fall within our tolerance interval; we may conclude, then, that in the considered transport medium, when plume measurement error statistics are similar to those used in these simulations, that it is possible to recover a plume history over an interval of 25 units of time and be accurate at a scale of $4s=8$ units of time. Fig. 3c is a plot of the results for the plume evolution time $T-50=50$. In this case, 98% of the trials were acceptable; in one of the 50 trials, the realization of the error was such that the regularized release function did not fall within the range defined by our tolerance level. Thus, we can still say with great confidence that a plume history can be accurately recovered for the plume evolution time $T-50=50$ at a scale of 8 units of time. As the plume evolution time $T-50$ is increased further (Fig. 3d, e, f, g, and h), the percentage of correctly recovered release histories decreases; at $T-50=250$ (Fig. 3h), only 40% of the trials result in accurate regularized release histories.

The preceding results are, of course, only relevant for the case of plume measurement errors described by Eq. (6) with $\delta=0.01$ and the plume sampling grid that was used. To get a more comprehensive picture of the recoverability of a plume in a system with $v=1$ and $D=1$, we must perform additional simulations while systematically varying the relevant test function parameters and measurement error statistics. Table 1 summarizes the results of such additional Monte Carlo simulations by giving the percentage of accurate recoveries for various test functions at different plume evolution times $T-50$. Each row of the table represents either a different test function, i.e., different combinations of s and M, or a different level of noise δ in the plume measurements. Each column is a different value of the plume evolution time $T-50$, with $T-50$ increasing from 25 to 550 moving left to right across the table. The first row is the same example that was plotted in Fig. 3 (i.e., $s=2$, $M=1$, and $\delta=0.01$). The second row is for $s=1$ and $M=1$ with $\delta=0.01$. The test function has the same amount of mass as before but the release occurs over a smaller time interval, $4s=4$ units of time. The percentage of accurate recoveries are lower at each plume evolution time, ranging from 78% at $T-50=25$ to less that 10% at the larger times. We see that under the same conditions it is more difficult to obtain a regularized release history that resolves the smaller time scale. The next row is for $s=1$ and $\delta=0.01$ again, but now there is 10 times more mass in the test function ($M=10$). The results are nearly the same as the preceding case indicating that when measurement errors are proportional to the measured concentration, the amount of mass in a test function does not affect our ability to recover it. If instead of Eq. (6) we generated measurement errors according to

$$\epsilon_j=\delta\xi_j,\ \xi\sim N(0,1) \tag{7}$$

where δ is again a scaler, then the amount of mass would be important. A test function with more mass would be easier to recover because the observed concentrations at time

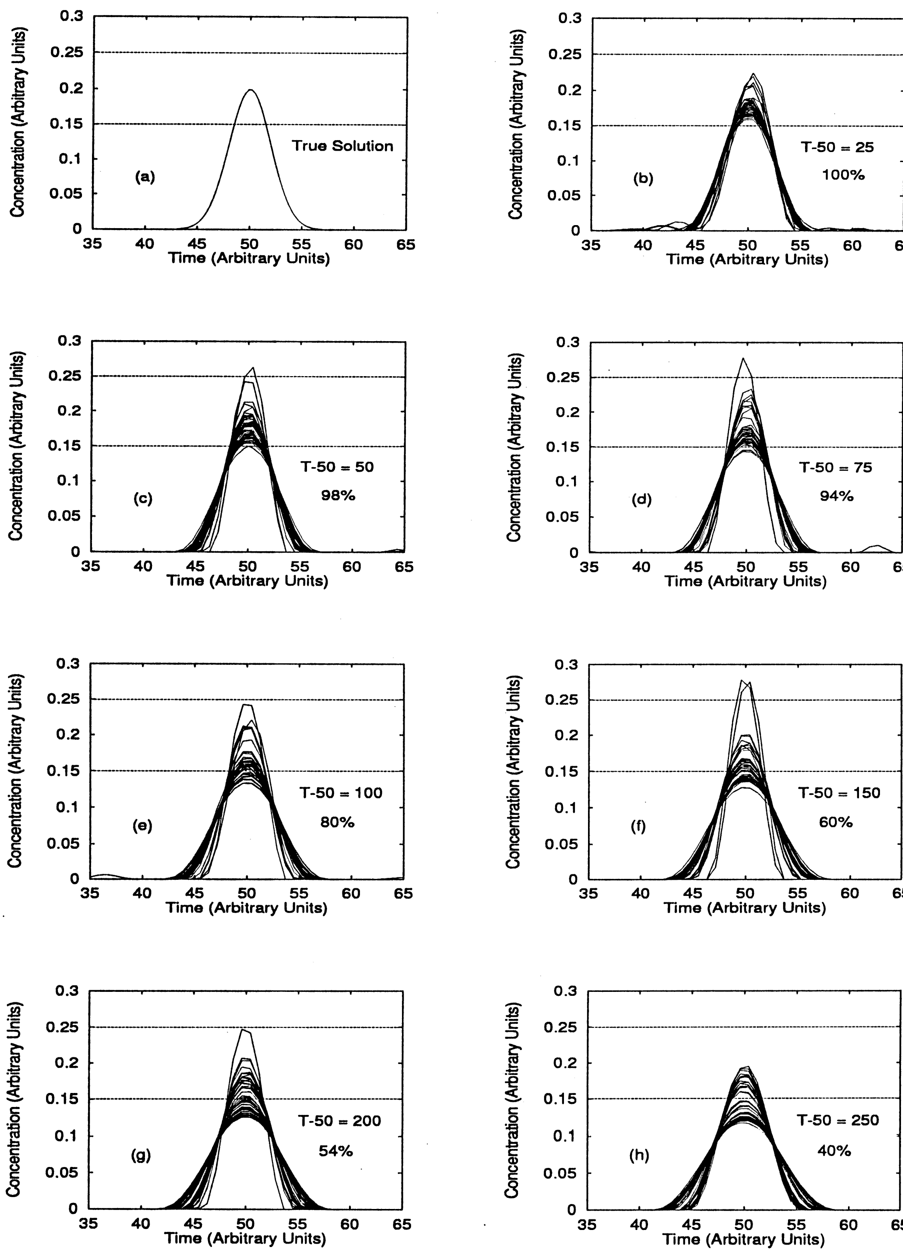

Fig. 3. (a) The true test function. If a recovered test function has a peak concentration in the range illustrated by the horizontal lines, it is considered an accurate recovery; $(b-h)$ results of individual Monte Carlo trials for different values of the plume evolution time T-50. The percentage indicated in the lower right corner of each plot is the fraction of Monte Carlo trials that resulted in an accurate recovery.

Table 1
Percentage of accurate test function recoveries

Parameter values	Time interval between contaminant release and plume observation												
	25	50	75	100	150	200	250	300	350	400	450	500	550
$M=1,\ s=2,\ \delta=0.01$	100	98	94	80	60	54	40	54	50	26	38	36	32
$M=1,\ s=1,\ \delta=0.01$	78	50	50	42	16	18	8	6	16	10	8	10	8
$M=10,\ s=1,\ \delta=0.01$	76	44	52	30	12	18	6	6	4	2	6	6	4
$M=1,\ s=2,\ \delta=0.05$	38	28	36	12	36	30	12	26	18	14	12	0	2
$M=1,\ s=4,\ \delta=0.01$	96	100	100	100	100	100	100	88	96	90	82	96	82
$M=10,\ s=4,\ \delta=0.01$	94	100	100	100	100	100	98	100	92	76	64	88	86
$M=1,\ s=4,\ \delta=0.05$	96	100	60	70	42	22	12	56	56	36	40	42	48

$t=T$ would be higher and the addition of an error to a high concentration has less impact than when the same size error is added to a low concentration. In the fourth row, the test function is the same as used in the first row ($s=2$, $M=1$), but now the magnitude of the measurement errors is 5 times greater ($\delta=0.05$). With this larger error, the results are under 50% even for the shorter plume evolution times. The next (fifth) row is with the smaller noise level again ($\delta=0.01$) but now $s=4$; the mass is still $M=1$. Attempting to resolve this larger time scale ($4s=16$), we obtain the best results so far. At plume evolution times up to $T-50=250$, we recover nearly 100% of plumes correctly. Beyond this, the percentage is still over 80. The second to last row is the same as the row above it only with more mass ($M=10$); again we see that the amount of mass does not affect the results. The last row is for $s=4$, $M=1$, and $\delta=0.05$. With this higher level of error, we still obtain over 90% accuracy at the shorter plume evolution times but the accuracy is significantly reduced at the longer times.

Producing a table such as Table 1 allows us to address the questions raised earlier. For example, for an arbitrary plume that is transported according to Eq. (1) with $v=1$ and $D=1$, what is the likely resolution of a recovered release history found over a time interval equal to 150 units of time? From Table 1 we conclude that it is: (i) highly likely (100% accurate recoveries) that the recovered release history correctly resolves the details of the true release at a time scale of $4s=16$ units of time (rows 5 and 6); (ii) somewhat likely (60% accurate recoveries) that the details are resolved at a time scale $4s=8$ (row 1); and (iii) unlikely (16 and 12 percent accurate recoveries) that the details are resolved at a time scale $4s=4$ (rows 2 and 3). If the plume measurement uncertainty is higher ($\delta=0.05$, rows 4 and 7), the chances of correctly resolving the release history at the scales $4s=8$ and $4s=16$ is reduced approximately by half (36 vs. 60% accurate recoveries and 42 vs. 100% accurate recoveries, respectively).

Alternatively we may ask, How far back-in-time can we recover a plume release history and still resolve the details of the true release at a time scale of $4s=8$? For $\delta=0.01$, we see from row 1 that there is: (i) a high likelihood ($>80\%$ accurate recoveries) that we can go back 100 units of time; (ii) a moderate likelihood ($\approx 50\%$ accurate recoveries) that we can go back 350 units of time; and (iii) a lesser likelihood that we can go back any farther.

5. Summary

Groundwater contaminant transport is a dispersive process and consequently there are limits to what may be learned about a contaminant's history from measurements of its present spatial distribution. The extent of the limitations in a particular case depends on a number of factors, including the dispersive properties of the transport medium, the accuracy of the measured plume, and the interval of time over which the plume has evolved. Since regularization methodologies for recovering plume histories always produce release scenarios that accurately reproduce the plume measurement data, there is the practical problem of determining whether the regularized solution is accurate, or whether the cumulative effect of the limiting factors have resulted in a regularized solution that does not contain the details of the true solution. We have presented a methodology that can be used to characterize the limitations that can be expected in recovering a plume history in a particular transport medium. The methodology uses a Monte Carlo approach to assess our ability to recover various hypothetical contaminant release functions that are designed to provide insight into the effect that various limiting factors have on the recoverability of a plume history. Determining our ability to recover the hypothetical releases allows us to make inferences about our ability to recover an arbitrary plume in a transport medium with dispersive characteristics similar to those used in the simulations.

The approach presented herein is general in the sense that transport models other than Eq. (1) and solution methodologies other than Tikhonov regularization can be used (see, e.g., Gorelick et al., 1983; Bagtzoglou et al., 1992; Wagner, 1992; Skaggs and Kabala, 1995; Woodbury et al., 1996; Snodgrass and Kitanidis, 1997). Other factors that may limit the accuracy of the recovered history can also be included in the analysis, such as transport parameter uncertainty and source location uncertainty.

Acknowledgements

This research was supported partially by USDA CSRS under grant 93-37102-9393.

References

Bagtzoglou, A.C., Dougherty, D.E., Tompson, A.F.B., 1992. Application of particle methods to reliable identification of groundwater pollution sources. Water Resour. Manage. 6, 15–23.

Gorelick, S.M., Evans, B., Remson, I., 1983. Identifying sources of groundwater pollution: An optimization approach. Water Resour. Res. 19 (3), 779–790.

Provencher, S.W., 1982a. A constrained regularization method for inverting data represented by linear algebraic or integral equations. Comput. Phys. Commun. 27, 213–227.

Provencher, S.W., 1982b. CONTIN: a general purpose constrained regularization program for inverting noisy linear algebraic and integral equations. Comput. Phys. Commun. 27, 229–242.

Skaggs, T.H., Kabala, Z.J., 1994. Recovering the history of a groundwater contaminant. Water Resour. Res. 30 (1), 71–79.

Skaggs, T.H., Kabala, Z.J., 1995. Recovering the history of a groundwater contaminant plume: method of quasi-reversibility. Water Resour. Res. 31 (11), 2669–2673.

Snodgrass, M.F., Kitanidis, P.K., 1997. A geostatistical approach to contaminant source identification. Water Resour. Res. 33 (4), 537–546.

Tikhonov, A.N., Arsenin, V.Y., 1977. Solutions of Ill-Posed Problems, Winston and Sons, New York.

Uffink, G.J.M., 1989. Application of Kolmogorov's backward equation in random walk simulations of groundwater contaminant transport, Contaminant Transport in Groundwater, Kobus, Kinzelbach (Eds.), Ashgate.

Wagner, B.J., 1992. Simultaneous parameter estimation and contaminant source characterization for coupled groundwater flow and contaminant transport modelling. J. Hydrol. 135, 275–300.

Woodbury, A.D., Ulrych, T.J., 1996. Minimum relative entropy inversion: theory and application to recovering the release history of a groundwater contaminant. Water Resour. Res. 32 (9), 2671–2681.