Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Hydrology 281 (2003) 281–295
Journal of Hydrology
www.elsevier.com/locate/jhydrol

Coupled inverse modelling of groundwater flow and mass transport and the worth of concentration data

Harrie-Jan Hendricks Franssen*, Jaime Gómez-Hernández, Andrés Sahuquillo

Departamento de Ingeniería Hidráulica y Medio Ambiente, Universidad Politécnica de Valencia, 46071 Valencia, Spain

Received 11 September 2001; revised 15 June 2002; accepted 8 May 2003

Abstract

This paper presents the extension of the self-calibrating method to the coupled inverse modelling of groundwater flow and mass transport. The method generates equally likely solutions to the inverse problem that display the variability as observed in the field and are not affected by a linearisation of the state equations. Conditioning to the state variables is measured by an objective function including, among others, the mismatch between the simulated and measured concentrations. Conditioning is achieved by minimising the objective function by gradient-based methods. The gradient contains the partial derivatives of the objective function with respect to: log conductivities, log storativities, prescribed heads at boundaries, retardation coefficients and mass sources. The derivatives of the objective function with respect to log conductivity are the most cumbersome and need the most CPU-time to be evaluated. For this reason, to compute this derivative only advective transport is considered. The gradient is calculated by the adjoint-state method. The method is demonstrated in a controlled, synthetic study, in which the worth of concentration data is analysed. It is shown that concentration data are essential to improve transport predictions and also help to improve aquifer characterisation and flow predictions, especially in the upstream part of the aquifer, even in the case that a considerable amount of other experimental data like conductivities and heads are available. Besides, conditioning to concentration data reduces the ensemble variances of estimated transmissivity, hydraulic head and concentration.

Keywords: Inverse modelling; Self-calibrated algorithm; Stochastic transport; Geostatistics; Data worth; Heterogeneity

1. Introduction

Flow and transport predictions are always uncertain due to an imperfect knowledge of the aquifer properties. It is of crucial importance in practical applications to reduce the uncertainty of these predictions. The conditioning to measurement data by inverse modelling algorithms offers means to reduce prediction uncertainty. A large number of inverse modelling techniques is available in the literature (e.g. Kitanidis and Vomvoris, 1983; Dagan, 1985; Carrera and Neuman, 1986a,b,c; Sahuquillo et al., 1992; RamaRao et al., 1995; Oliver et al., 1997; Hanna and Yeh, 1998; Hu, 2000). However, the conditioning to concentration (c) data is less widely studied, in spite of the fact that c data, for example from controlled tracer tests, may give

* Corresponding author. Address: Institut für Hidromechanik und Wasserwirtschaft, ETH-Hönggerberg, HIL G33.3, Zürich CH-8093, Switzerland. Tel.: +41-1-6333-074; fax: +41-1-6331-061.
E-mail address: hendricks@ihw.baug.ethz.ch (H.-J.H. Franssen).

0022-1694/03/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/S0022-1694(03)00191-4

important information on spatially variable aquifer properties like hydraulic conductivities (K).

Graham and McLaughlin (1989a,b) condition to concentration measurements by propagating and updating conditional concentration ensemble moments. In their approach, heterogeneous transmissivity is the only source of uncertainty and modelled as a, possibly non-stationary, random space function. First, a methodology is presented to calculate the temporal evolution of non-conditional ensemble concentration moments, in combination with a numerical method to solve the moment propagation equations (Graham and McLaughlin, 1989a). Next, the formulation is extended for conditional concentration ensemble moments. The concentration ensemble moments are updated by the Kalman filter approach at times that new concentration measurements become available (Graham and McLaughlin, 1989b). The methodology does not allow the updating of the heterogeneous transmissivity field, the source of uncertainty. Other limitations are the restriction to moderately heterogeneous transmissivity fields and the CPU-intensivity of the method. Rubin (1991a,b) derives, in two papers, an approach that allows the conditioning of concentration ensemble moments on hydraulic head, conductivity, velocity and concentration measurements. Concentration residuals (deviations between predicted concentrations and measured concentrations) are interpolated in space and time by cokriging so that at each location and time an expected deviation is obtained. The covariances of the time–space concentration field are obtained from tracking a sufficiently large number of particles through a velocity field that may be conditioned to conductivity, head and velocity data. The approach of Rubin is also limited to moderately heterogeneous transmissivity fields and does not update the transmissivity field. Medina and Carrera (1996) solve the fully coupled inverse problem and their methodology allows the estimation (updating) of many flow and transport parameters. The methodology is an extension of the zonation approach of Carrera and Neuman (1986a,b,c). The approach does not have the limitations of the previous methods, and its main disadvantage is the limitation in representing heterogeneity by the number of zones in which the aquifer has to be partitioned, a number that cannot be too large to avoid unstable solutions. Some applications of the methodology of Medina and Carrera can be found in Medina (1993), Wagner (1992), Iribar et al. (1997), Sonnenborg et al. (1996) and Mayer and Huang (1999).

None of the previous studies generates multiple equally likely aquifer models as solution to the fully coupled inverse problem. In this paper an extension of the self-calibrating method is presented that generates multiple equally likely aquifer models, all of them consistent with the experimental information that may consist of K data, indirect information on K, storativity coefficient (S) data, hydraulic head (h) data (steady and/or transient), c data and data on retardation coefficients (R). The parameters that can be generated are spatially variable K, spatially variable S, prescribed boundary heads, mass injection and/or extraction and spatially variable R. The method is not restricted to a small variance of K, and since it provides multiple realisations for all parameters and state variables, local probability distribution functions are obtained.

This paper presents the extension of the self-calibration method to the coupled inverse modelling of flow and transport. Previously, the method was developed for the inverse modelling of 2D steady-state flow (Sahuquillo et al., 1992; Gómez-Hernández et al., 1997; Capilla et al., 1997), 2D transient flow (Hendricks Franssen et al., 1999) and 3D flow in fractured media (Hendricks Franssen and Gómez-Hernández, 2002).

This paper consists of two main parts. In the first part the methodology is presented. In the second part a synthetic study is discussed that illustrates the methodology and investigates the worth of concentration data.

2. Theory

Saturated constant-density groundwater flow in a fully confined aquifer is considered

$$\nabla \cdot K \nabla h = S \frac{\partial h}{\partial t} + W$$

where K is hydraulic conductivity tensor, S is storativity, W represents sinks and sources, h is hydraulic head and t is time.

The following equation for the transport of solutes is considered

$$\frac{1}{R}\nabla\cdot(\phi D\cdot\nabla c - \phi v c) = \phi\frac{\partial c}{\partial t} + q$$

where R is the retardation coefficient, ϕ is the aquifer porosity, D is the hydrodynamic dispersion tensor, c is the mass concentration, v is the pore velocity vector and q are mass sinks and sources.

The spatially variable K is modelled by a multi-Gaussian Random Stochastic Field (RSF) and also S and R may be modelled as RSF. Conditioning data may be available for the following parameters and state variables: K, S, R (also indirect or soft information on these parameters could be used), h (steady and/or transient) and c. The outcome sought is the generation of a large number of realisations of the spatially variable parameters (mainly K, but also possibly other parameters like S, R, boundary conditions and mass source information) that are all conditioned to the available data.

2.1. Conditioning to K, S and R

The first step in the conditioning is to generate a field $Y = \log_{10}K$, and possibly also $Z = \log_{10}S$ (in case of transient groundwater flow) and R (in case of linear reactive transport) conditional to Y, Z and R data. This is achieved by geostatistical simulation, like, for instance, MultiGaussian sequential co-simulation (Gómez-Hernández and Journel, 1993). A variogram function has to be specified for each of the three parameters and in case there is an evidence for cross-correlation, also cross-variograms between parameters have to be specified. Conditioning to soft information of K, S or R could be done by sequential indicator simulation (Gómez-Hernández and Srivastava, 1990). The generated realisations of the RSF's display spatial variability as observed in the field and modelled by the variograms.

2.2. Solving groundwater flow and mass transport

The groundwater flow and the mass transport equation are solved for each of the generated sets of Y, Z and R. A block-centred finite difference approach is used to solve the equations numerically. Boundary and initial conditions have to be specified, together with information on water injection (recharge) or extraction (pumping), the contaminant concentration of the water injected and the dispersivity coefficients. After solving the state equations, the simulated h and c are compared with the measured values at the same locations and times.

The following objective function is defined that measures the mismatch between simulated and measured head and concentration values

$$J=\sum_{t=1}^{N_t}\sum_{i=1}^{N_h}\xi_{it}(h_{i,t}^{\mathrm{SIM}}-h_{i,t}^{\mathrm{MEAS}})^2+\psi\sum_{t=1}^{N_{\mathrm{TC}}}\sum_{i=1}^{N_{\mathrm{C}}}\zeta_{it}(c_{i,t}^{\mathrm{SIM}}-c_{i,t}^{\mathrm{MEAS}})^2$$

where the first term corresponds to the head discrepancies at the different time steps (if only steady-state flow is simulated this corresponds to the first and only time step) and the second term to the discrepancies between measured and simulated concentrations. Additional terms that penalise the differences between initial parameter values and updated parameter values can be included, but were set to zero in the study that is presented in this paper. N_{h} is the number of head measurement locations, N_{t} the number of time steps with head measurements, $h_{i,t}$ the heads, N_{C} is the number of concentration measurement locations, N_{TC} the number of time steps with concentration measurements, $c_{i,t}$ the concentrations and the superscripts SIM and MEAS refer to 'simulated' and 'measured', respectively. The weights ξ_{it} and ζ_{it} are chosen inverse-proportional to the average estimated measurement errors. The value of the trade-off coefficient ψ should be chosen considering which type of data is more important to reproduce and their overall variability within the area of study.

The objective function J is minimised in order to achieve a satisfactory reproduction of the head and concentration data. When J reaches a value below a user-defined one, it is considered that the head and concentration data are reproduced sufficiently close. The minimisation process is also terminated in case the number of iterations exceeds a user-specified maximum number of iterations or in case the objective function reduction is very small during at least 10 succesive iterations. It is possible that the optimisation finds a local minimum, but in case of a sufficient close reproduction of the measurement data this is not considered a problem. Nevertheless,

the implementation of faster converging optimisation algorithms that are less sensitive to local minima is subject to future research.

2.3. Minimisation of the objective function

If J is not sufficiently low, a gradient-based minimisation process is started. The gradient of the objective function with respect to those parameters deemed most relevant to the flow and transport processes is calculated. These parameters include the log conductivity (or log transmissivity) and possibly prescribed boundary heads, storativity coefficients, retardation coefficients and mass sinks and sources. Since these parameters are spatially variable and display a spatial correlation, the gradient is not computed with respect to the parameter values at each discretisation cell, but only at each of the so-called master blocks, imposing that the parameter variations in between master blocks can be obtained by interpolating the master blocks variations. The master blocks can also be seen as a way to simplify the parameterisation of the perturbation applied to the seed parameter field(s) in order to minimise the objective function. The number of master blocks is always much smaller than the total number of cells, and they are overlaid on a regular grid. The interpolation in between master block locations is carried out by ordinary kriging (or cokriging if several parameter fields are modified simultaneously). In case for example only log hydraulic conductivity is perturbed the hydraulic conductivity at a grid cell is updated applying the following formula

$$\Delta Y_{ij} = \sum_{k=1}^{N_P} \lambda_{ij}^{k} \Delta Y_k$$

where N_P is the number of master blocks, ΔY_k the log conductivity perturbations at the master blocks, λ_{ij}^{k} the kriging weights for the interpolation of the perturbation at location ij from the master location perturbations. The calculated kriging weights are based on an interpolation using a variogram that is equal to the initially specified variogram.

The determination of the optimal perturbations ΔY_k is not straightforward. Since the objective function includes the state variables, and these depend in a complex way, through the flow and transport equations, on the parameters, determining the gradient of the objective function with respect to the parameter perturbations at master locations is cumbersome. The details in determining this gradient can be found in Gómez-Hernández et al. (1997), for the parameters related to the flow equation, and in Hendricks Franssen (2001) for the parameters related to the transport equation. In both cases, adjoint state formulations can alleviate the computations. The calculation of the derivatives of the objective function with respect to the parameters is especially complicated due to the complex relationship between parameters and state variables, especially concentrations with respect to conductivities, since a local change in conductivity induces a global change in the spatial distribution of piezometric heads, which in turn modify the seepage velocities and the elements of the hydrodynamic dispersion tensor. This complexity requires that an approximation be made and, for the purpose of computing the gradient of the objective function, the coefficients of the dispersion tensor are considered constant (an approximation that is found to be adequate in the tests carried out). The above approximation affects only the computation of the gradient, and not the computation of the state of the system, which is always determined solving the full flow and transport equations.

Non-linear optimisation algorithms are used to calculate an updating direction as a function of the gradient vector; these algorithms include steepest descent or conjugate gradient. Then, linear search is carried out to optimise the step size along the updating direction. Details are given in Hendricks Franssen et al. (1999) and Hendricks Franssen (2001). The resulting perturbation fields are added to the fields from the last iteration and the groundwater flow and mass transport equations are solved for the new, updated fields.

3. Synthetic study

The study domain has extension of 500 m by 500 m and is divided in 50 by 50 square grid cells of 10 m in the side.

A reference log transmissivity field is generated with an average transmissivity of -6.0 $\log_{10}(m^2/s)$. The variogram of log transmissivity is spherical with

a range of 125 m, zero nugget and sill of 0.25 $(\log_{10}(m^2/s))^2$.

Steady-state groundwater flow is simulated with prescribed head boundaries on the south and north and impermeable boundaries on the west and east. The prescribed head value is 5 m along the southern boundary and 0 m along the northern boundary. These prescribed head values force a flow from south to north.

On the southern boundary, a line contamination source is present. The contaminants are introduced in the system by a constant concentration of 1.0 unit (standardised concentration) along a 100 m long transect centered on the southern boundary. For the rest of the southern boundary the prescribed concentration values are equal to zero. On the northern boundary also prescribed concentrations equal to zero are imposed. These boundary conditions are justified because during the time of the transport simulation the northern boundary is far enough from the contaminant plume. The contaminants are subject to hydrodynamic dispersion and do not show retardation or any chemical reaction. The longitudinal and transversal dispersion coefficients are set equal to 10 m. This value is large to avoid numerical instabilities (Peclet number equal to one). However, the assumption that the gradient could be estimated adequately neglecting the partial derivatives of the dispersive part of the transport equation with respect to the log conductivity perturbations was still acceptable. The mass transport equation is

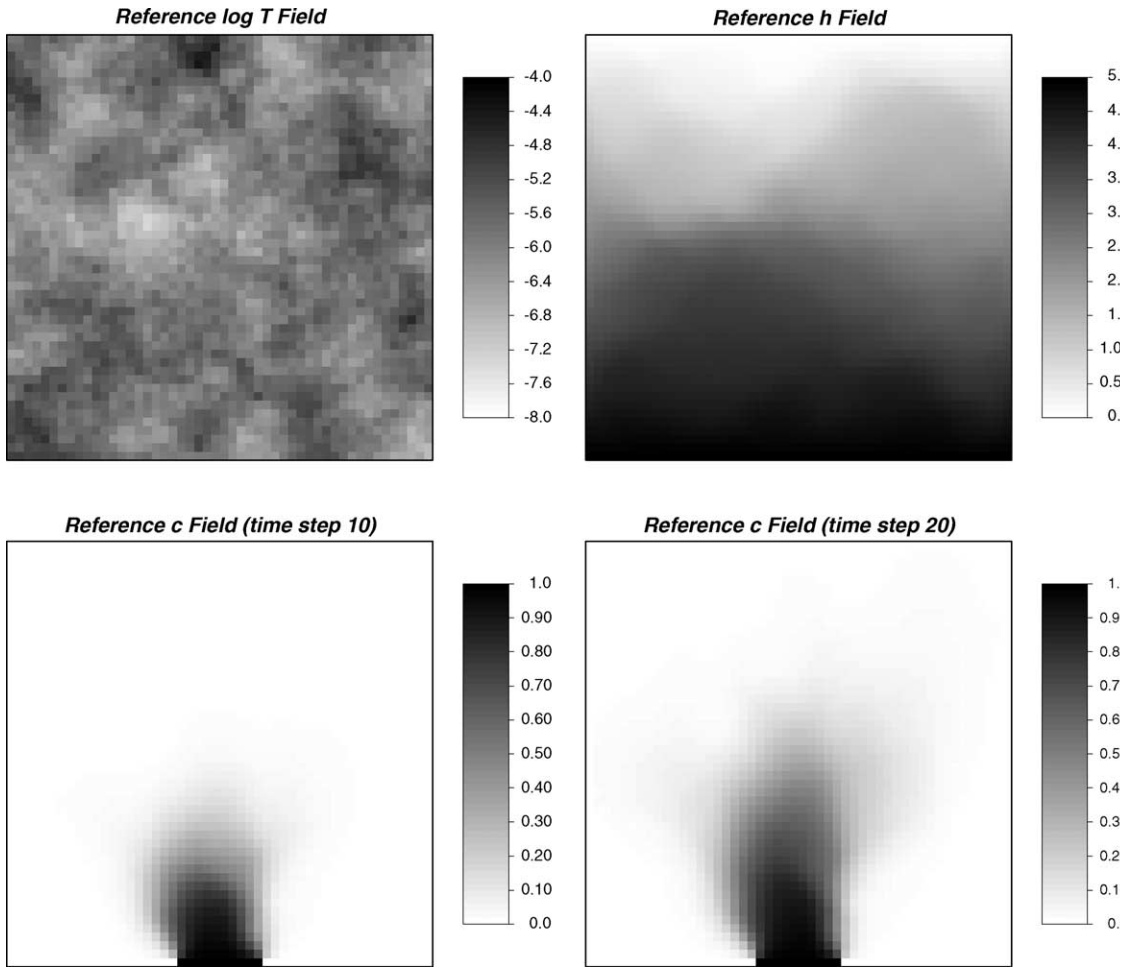

Fig. 1. Reference log transmissivity field, steady-state head field and concentration fields for time step 10 and 20.

Table 1
Conditioning data sets used in the different scenarios analysed

	20 Y data	20 h data	13 c data
Scenario 1	No	No	No
Scenario 2	Yes	No	No
Scenario 3	No	Yes	No
Scenario 4	Yes	Yes	No
Scenario 5	No	No	Yes
Scenario 6	Yes	No	Yes
Scenario 7	No	Yes	Yes
Scenario 8	Yes	Yes	Yes

solved for time steps of 10^9 s (31.69 years) until 2×10^{10} s (633.76 years).

Fig. 1 shows the reference transmissivity, steady-state hydraulic head, and concentration (at two different time steps) fields.

3.1. Scenarios studied

One hundred equally likely solutions to the inverse problem are calculated for eight different scenarios. The scenarios differ in the kind of conditioning data. Table 1 illustrates the types of data used in each of the eight scenarios. In four of the eight scenarios (scenarios 2, 4, 6 and 8) 20 Y data are used in the conditioning procedure. The Y data are obtained by random sampling the reference Y field. Also in four of the eight scenarios (scenarios 3, 4, 7 and 8) 20 steady-state head data are used as conditioning information. The steady-state head data are sampled at the same locations as the transmissivity data. Finally, in four of the studied scenarios (scenarios 5, 6, 7 and 8) concentration data are used in the conditioning procedure and for these scenarios a coupled inversion procedure was done. The concentration data are sampled from the reference concentration field for the last time step (2×10^{10} s) at 13 locations along a monitoring line perpendicular to the mean flow direction, 100 m from the source. No concentration data from earlier time steps were used. However, the method allows to consider time series of concentration measurements and it is expected that time series of concentration measurements (instead of one single measurement) would yield a better estimation of the unknown parameters and a larger uncertainty reduction. Fig. 2 shows the location of the measurement data.

For the eight different scenarios the aim was to generate realisations of log transmissivity with patterns of variability similar to those of the reference case and reproducing the h and c data as close as possible (no measurement errors are modelled). The Y data are reproduced exactly by construction of the seed fields and the restrictions of the self-calibrating algorithm. This means that at the locations were Y data are available this measurement is reproduced (in case of no measurement error), but that at the rest of the grid cells there is uncertainty with respect to Y. In case that both h and c data are used as conditioning information it is necessary to define a trade-off value (the parameter ψ in the expression given in Section 2.2) in order to weight the two pieces of information. It was found that a value equal to 1.0 yields good results in this study in terms of reproducing both measured hydraulic head and concentration data closely. Larger trade-off values

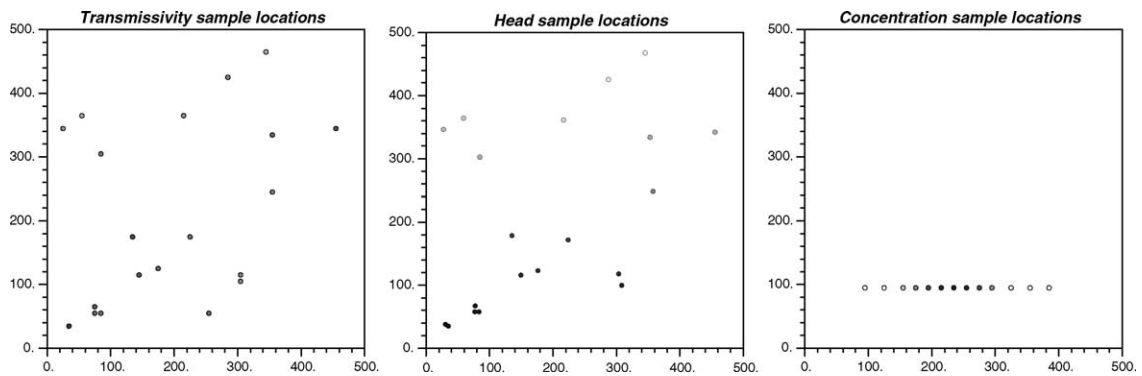

Fig. 2. Sample locations of transmissivity, steady-state piezometric head, and concentration at time step 20.

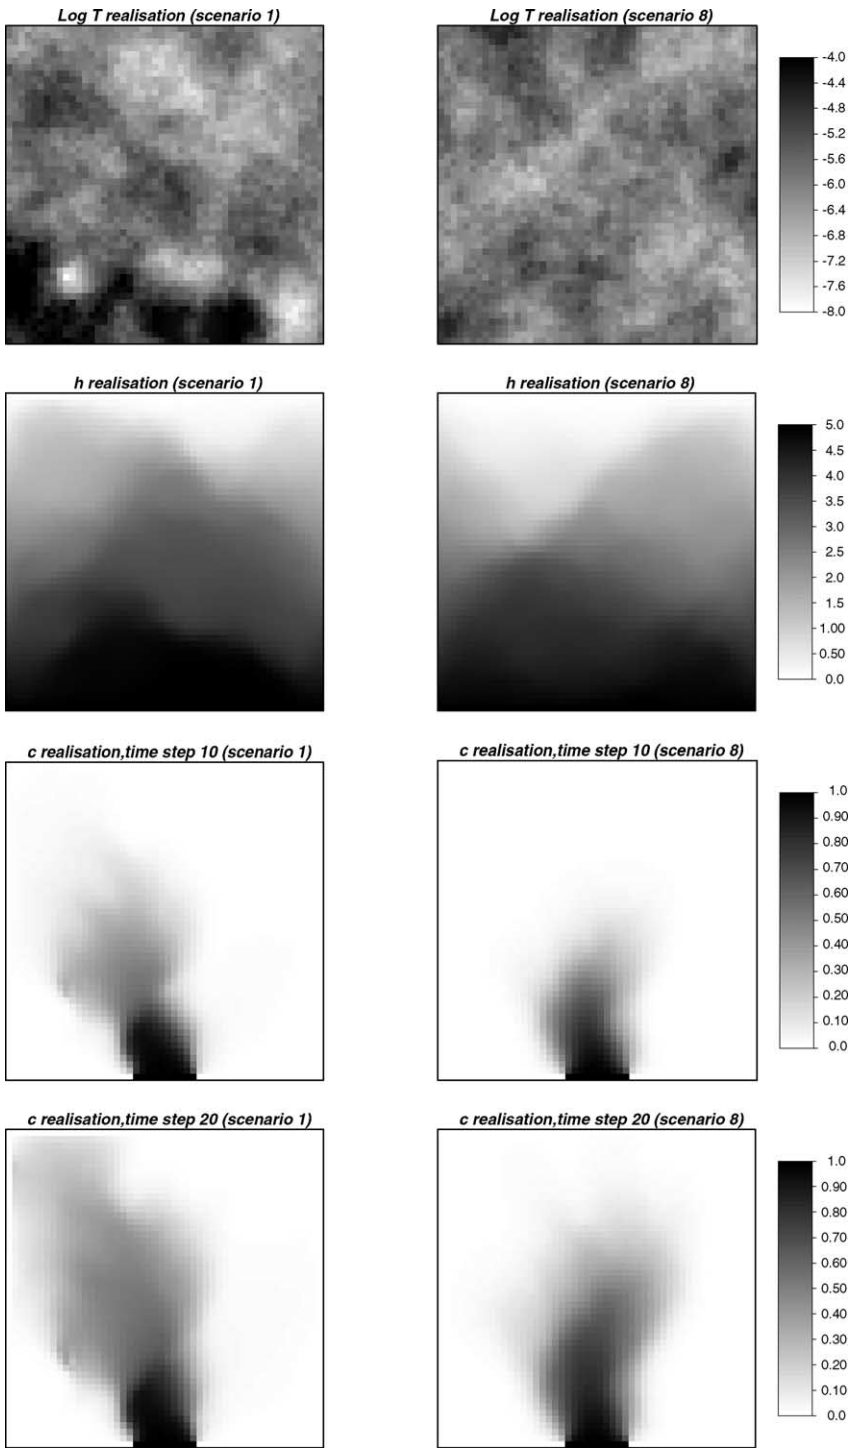

Fig. 3. On the left column an unconditional realisation (scenario 1), on the right column a conditional one using the logtransmissivity, steady-state head and concentration data sets (scenario 8).

did not result in a sufficiently close reproduction of the hydraulic head data.

One hundred master blocks are located within the simulation domain on a square regular grid.

Fig. 3 shows, as an example, a realisation not conditioned to any information (scenario 1) and a realisation conditioned to Y, steady-state h and c data (scenario 8). The figure illustrates how the realisation conditioned to all the information is closer to the reference fields than the unconditional realisation. In the unconditional realisation the contaminant plume extends much further than in the reference field.

3.2. Evaluation of results

For each of the eight scenarios, 100 equally likely realisations are generated by the sequential self-calibrated method, as outlined in Section 2. Each of the realisations is compared with the reference fields and the following performance measures are defined for each of the scenarios

$$\mathrm{AAE}(X) = \frac{1}{N}\sum_{i=1}^{N} |\bar{X}_{\mathrm{SIM},i} - X_{\mathrm{REF},i}|$$

$$\mathrm{AESD}(X) = \frac{1}{N}\sum_{i=1}^{N} \sigma_{X_i}$$

where AAE is the average absolute error, AESD the average ensemble standard deviation, N the number of discretisation grid cells, and i is a grid cell index, X represents either decimal log transmissivity, steady-state hydraulic head or mass concentration at a certain time step, the over bar indicates ensemble average, the subscript SIM refers to the realisations, and the subscript REF to the reference values; finally, σ_{X_i} is the ensemble standard deviation of X at a given node.

3.3. Results

Table 2 shows the calculated AAE for log transmissivity, hydraulic head and concentration for the eight scenarios. The results are standardised so that AAE is equal to 100 for the scenario with no conditioning data. Fig. 4 shows the ensemble averaged log transmissivity, hydraulic head and concentration fields for some of the scenarios.

The AAE for log transmissivity, hydraulic head and concentration are below 100 for all the scenarios with conditioning data. It indicates that the presence of conditioning data results in all cases in an improved characterisation of the transmissivity, hydraulic head and concentration fields.

Table 2
Standardised average absolute error (scenario 1 = 100) for the log transmissivity, steady-state head and concentration (averaged over 20 time steps) fields

	AAE(Y)	AAE(h)	AAE(c)
Scenario 1	100	100	100
Scenario 2	**92**	93	68
Scenario 3	91	**49**	97
Scenario 4	**81**	**34**	70
Scenario 5	96(80)	92(54)	**54**(**29**)
Scenario 6	**91**(**72**)	90(73)	**58**(**17**)
Scenario 7	88(74)	**52**(**73**)	**43**(**31**)
Scenario 8	**79**(**65**)	**35**(**55**)	**37**(**28**)

Bold script indicates conditioning to this type of data. Between brackets are given the statistics calculated over the upstream part of the aquifer.

3.3.1. Results when a single type of data are used for conditioning

In this section, scenarios 2, 3 and 5 are analysed, that is, those scenarios in which only one type of conditioning data is considered. In all cases, conditioning to a single type of data helps in improving the characterisation of the three attributes; log transmissivity, steady-state piezometric head and concentration although in different ways. When only transmissivity data are used, the characterisation of the transmissivity improves (AAE(Y)-reduction of 8% with respect to the scenario with no conditioning data, i.e. scenario 1), but also the characterisation of the steady-state head field (AAE(h)-7% reduction) and it is particularly important the improvement of the characterisation of the concentration field (AAE(c)-average reduction of 32% over the 20 time steps). It shows how conditioning to transmissivity data helps in producing a better description of the seepage velocities resulting in a noticeable improvement of the characterisation of the concentrations. Conditioning to hydraulic head data yields a similar reduction of AAE(Y) as before (9%), a much more pronounced AAE(h) reduction (51%) and just a small AAE(c) reduction (an average of 3% over the 20 time steps). It is not straightforward to explain why head and

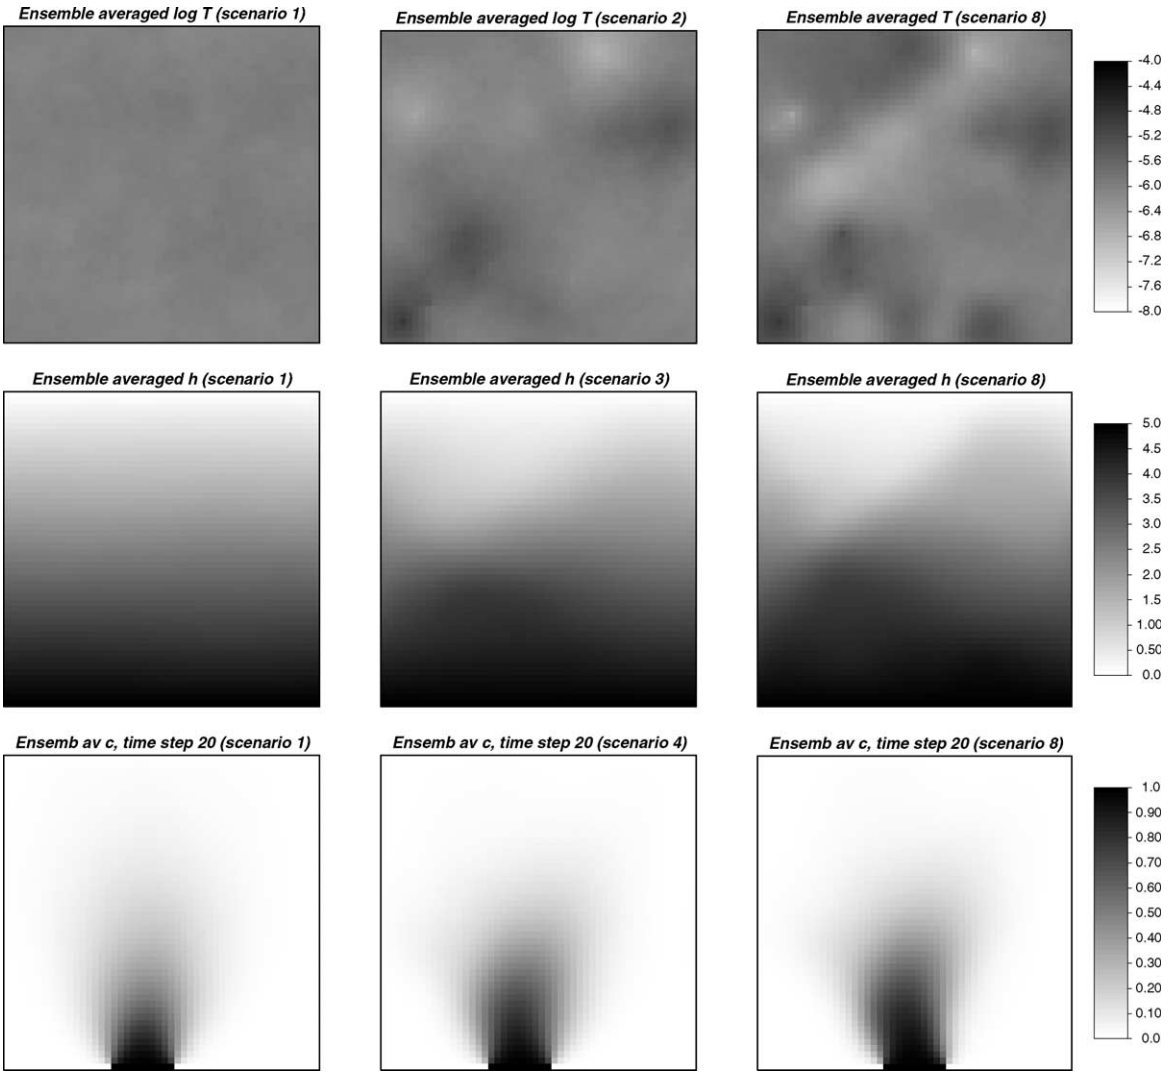

Fig. 4. Ensemble averages of log transmissivity for scenarios 1, 2 and 8; ensemble averages of steady-state head for scenarios 1, 3 and 8; ensemble averages of concentration at time step 20 for scenarios 1, 4 and 8.

transmissivity data had such a different impact on the characterisation of the concentration field in this case, particularly considering that the improvement of the characterisation of the head field should result in a much better description of hydraulic head gradients, and thus of seepage velocities. Finally, the use of concentration data only yields the smaller reduction AAE(Y) (4%), some AAE(h) reduction (8%) and, as would be expected, a large AAE(c) reduction (46%).

The impact of concentration conditioning is more pronounced if the aquifer is divided in two by a horizontal line at the level of the monitoring location. In such case, the values of the AAE computed only on the upstream zone (the one between the source and the monitoring line) are much smaller than when computed over the entire aquifer. Indeed, AAE(Y) shows a 20% reduction with respect to the unconditional case, AAE(h) shows a 46% reduction and AAE(c) shows a 71% reduction. This illustrates that the AAE(Y) and AAE(h) reductions are small for the downstream zone of the aquifer, but very important for the upstream zone of the aquifer. This is logical as the concen-

trations are much more sensitive to groundwater flow velocities (and thus the transmissivities and hydraulic heads) in the upstream zone of the aquifer than in the downstream zone. Finally, the characterisation of the concentration field improves both for the upstream zone of the aquifer and the downstream zone, but also in this case, the improvement is more significant for the upstream zone.

In the previous discussion the value of AAE(c) was reported on average over the 20 time steps. Fig. 5 shows the temporal evolution of the AAE(c) for scenario 1 (no conditioning data) and the temporal evolution of the AAE(c) for scenario 5 (13 concentration data). From Fig. 5 also the difference between the previous two values can be deduced, that is, the reduction on AAE(c) with respect to the unconditional scenario introduced by conditioning to concentration data only. For the whole aquifer, the AAE(c) increases throughout the simulation time and reaches its maximum at the last simulation time step (time step 20). This is the case for both scenario 1 (no conditioning data) and scenario 5 (conditioning to concentration data). The concentration data at time step 20 reduce the AAE(c) for all time steps, but do not result in a lower AAE(c) for time step 20 as compared to earlier time steps. It is interesting to compare the values of AAE(c) for the entire aquifer and for the upstream zone. In both cases, and for both scenarios, they start at zero (perfect prediction of no concentration at time zero); however, the evolution of AAE(c) is increasing when computed over the entire aquifer and tends towards zero when computed only on the upstream zone. At large times, the concentration measurements at the monitoring locations will provide little information about the distribution of downstream concentrations; however, it will provide almost exact information on whether the plume is fully developed and the area in between source and monitoring is fully contaminated. In any case and for both evaluations of the AAE(c), the influence of the conditioning data is clearly noticeable by the lower values of the scenario 5 with respect to scenario 1.

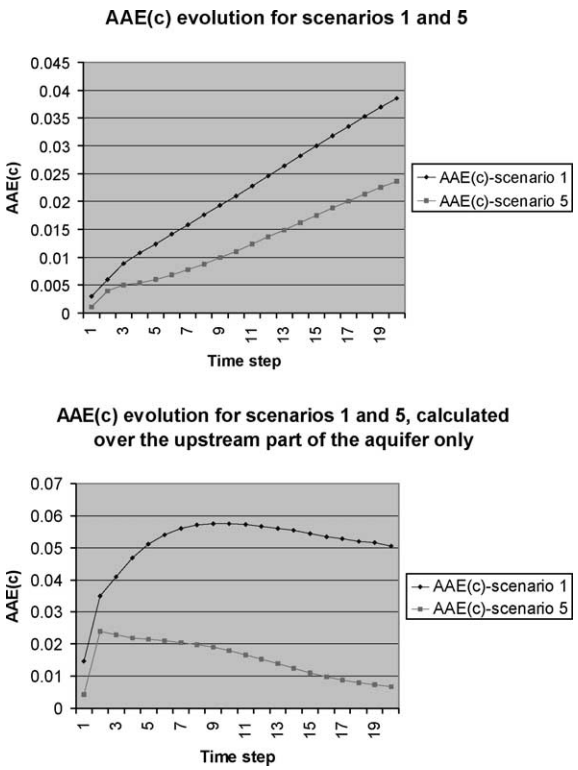

Fig. 5. Evolution of AAE(c) for scenarios 1 and 5 (in absolute value, without standardisation), computed over the entire aquifer (top) and over the upstream zone only (bottom).

3.3.2. Results for different kinds of conditioning information

In case both transmissivity and hydraulic head data are available (scenario 4) the AAE reductions are larger and close to the sum of the AAE-reductions for the scenarios 2 (just transmissivity data) and 3 (just hydraulic head data). The AAE(Y) reduction is 19%, the AAE(h) reduction 66% and the AAE(c) reduction on average over the 20 time steps is 30%.

The combination of concentration and transmissivity data (scenario 6) results in a stronger AAE-reduction than only concentration data or only transmissivity data. However, in this case, the AAE-reduction is smaller than the sum of the AAE-reductions of scenario 2 (just transmissivity data) and scenario 5 (just concentration data). The AAE(Y) reduction is 9% (the sum would have been 12%), the AAE(h) reduction 10% (the sum would have been 15%) and the AAE(c) reduction 42%. It should be noticed that the use of just concentration data yielded a stronger AAE(c) reduction.

The combination of concentration and hydraulic head data (scenario 7) yields important AAE(c) reductions. Whereas the AAE(Y) and AAE(h) reductions are close to the sum of the AAE(Y) and

AAE(h) reductions for scenario 3 (just hydraulic head data) and scenario 5 (just concentration data), the AAE(c) reduction (57%) is even larger than the sum of the AAE(c) reductions for scenario 3 (just hydraulic head data; reduction 3%) and scenario 5 (just concentration data; reduction 46%).

Finally, in case all conditioning information is used (scenario 8) the best results are obtained. The AAE(Y) reduction is 21%, the largest reduction of all scenarios studied. As found in other synthetic studies (e.g. Hendricks Franssen et al., 1999), the AAE(h) reduction is usually much bigger than the AAE(Y) reduction, but it is the improved characterisation of the transmissivity field that guarantees that under other flow regimes a better characterisation of the groundwater flow will be obtained. The AAE(Y) reduction for scenario 8 is nearly equal to the sum of the AAE(Y) reductions for just transmissivity data (8%), just hydraulic head data (9%) and just concentration data (4%). Fig. 4 shows that the conditioning data are able to characterise the principal zones of enhanced and reduced transmissivity. The AAE(h) reduction is 65%, which is basically the same reduction (66%) obtained for scenario 3, with just hydraulic head and transmissivity data. Hence, the concentration data did not help to reduce further the AAE(h), which may be due to a slightly less accurate reproduction of the hydraulic head data, when both types of data are used in the benefit of reproducing the concentration data. Nevertheless, the final AAE(h) reduction for scenario 8 is again very close to the sum of the AAE(h) reductions for just transmissivity data (7%), just hydraulic head data (51%) and just concentration data (8%). Finally, the AAE(c) reduction (again averaged over the 20 time steps) reaches also its maximum in case all the conditioning information is used (63%). The reduction is around 66% for the time steps 9 until 20, but for the first time steps the reduction is smaller. Fig. 4 shows that the conditioning data are able to characterise approximately the contaminant plume, while in the unconditional case the contaminant plume covers a smaller part of the aquifer than in the reference. The AAE(c) reduction for this case is smaller than the sum of the AAE(c) reductions for just transmissivity data, just hydraulic head data and just concentration data.

For scenario 5 (conditioning to only concentration data) it was found that the AAE-reductions are much larger for the upstream part of the aquifer. Is this still the case if besides concentration data also hydraulic head and/or transmissivity data are used in the inverse modelling? The AAE(Y) reductions are larger in the upstream part of the aquifer for all cases where concentration measurement data are available. We saw that for scenario 5 the overall AAE(Y) reduction was 4%, while in the upstream part it was 20%. For scenario 6 (concentration and transmissivity data) the overall reduction is 9%, and the reduction for the upstream part is 28%. For scenario 7 (concentration and head data) the contrast in AAE(Y) reduction between the upstream and downstream part is reduced: the overall reduction is 12% and the reduction for the upstream part is 26%. For scenario 8 (concentration, head and transmissivity data) the overall reduction is 21% while it is 35% for the upstream part. It can be concluded that the head and transmissivity data give an additional AAE(Y) reduction, but the additional value of the concentration data is evident, especially from the ensemble statistics of the upstream part of the aquifer. This is not the case for the AAE(h) reduction. While for scenario 5 (concentration data only) there was a strong contrast with an overall AAE(h) reduction of 8% and a reduction in the upstream part of 46%, for all other scenarios with concentration data the contrast is less or even inexistent. For scenario 6 (concentration and transmissivity data) the overall AAE(h) reduction was 10% and for the upstream part 28%. In case also hydraulic head data are available (scenarios 7 and 8) the AAE(h) reduction is not stronger for the upstream part than for the rest of the aquifer. Finally, the AAE(c) reduction is in all cases larger in the upstream part of the aquifer than in the rest of the aquifer. For scenario 5 (just concentration data) we saw already that the overall AAE(c) reduction was 46% and the reduction for the upstream part was 71%. For scenario 6 (concentration and transmissivity data) the contrast is even bigger: 42% AAE(c) reduction for all the aquifer and 83% for the upstream part only. The head data reduce the contrast in AAE(c) reduction as they were more capable to improve the characterisation of the concentration field downstream of the monitoring locations. For scenario 7 (concentration and head data) the overall AAE(c) reduction was 57% and for the upstream part it was 69%. In case head, concentration and transmissivity data are available

(scenario 8) the overall AAE(c) reduction was 63% and for the upstream part it was 72%.

The behaviour of the AAE(c) as function of the simulation time step is the same for the scenarios 6, 7 and 8 as for the scenario that just concentration data were available (scenario 5). The AAE(c), when calculated over the whole aquifer, increased continuously in time. On the contrary, the AAE(c) calculated over the upstream part of the aquifer reached its maximum very fast after the start of the transport simulations (time step 1 or 2) and decreased afterwards.

3.3.3. Average ensemble standard deviations

Table 3 shows the calculated AESD for each of the scenarios. The AESD can be considered a measure of local accuracy of the ensemble-averaged conditional fields with respect to the reference field. The general conclusion is that the conditioning data result in a reduction of the AESD, therefore improved local accuracy. It means that the conditional ensemble averaged, transmissivity, hydraulic head and concentration fields are not only closer to the reference fields (AAE-reduction), but also that the uncertainty on these estimates is smaller (AESD-reduction). Fig. 6 shows the ensemble-averaged standard deviations of the transmissivity, steady-state head and concentration fields for some of the scenarios.

Looking into more detail we see that the AESD-reductions with respect to the unconditional case are smaller than the corresponding AAE-reductions. Nevertheless, the AESD-reductions are also considerable. The AESD(Y) reduction is 19% for scenarios 4 and 8. This is similar to the maximum AAE(Y) reduction (for scenario 8) of 21%. It is found that transmissivity and head data help to reduce the uncertainty on the transmissivity field. Fig. 6 clearly illustrates the impact of the conditioning transmissivity data. However, when conditioning to concentration data, the surprising result is that the AESD(Y) increases. If only concentration data are used for conditioning (scenario 5) the AESD(Y) increase is 7% as compared with the scenario with no conditioning data. In case that also transmissivity data are used (scenario 6) an AESD(Y) decrease of 12% is observed; this decrease is smaller than the decrease of 15% observed when only transmissivity data were used in the conditioning.

For the uncertainty on the head field, conditioning to concentration data also results in an AESD(h) increase as compared to the unconditional case. However, the increases are smaller than for AESD(Y). For scenario 5 the AESD(h) is 5% larger than for the case with no conditioning data. Nevertheless, the maximum AESD(h) reduction is achieved in case all conditioning data (including concentration data) are used. The steady-state head data contribute the most to the AESD(h) reduction: in case that only steady-state head data are available, an AESD(h) reduction of 41% is already achieved. Fig. 6 shows how the zone with local AESD(h) above 0.15 reduces in case conditioning data are used. If the maps are compared with the head sample locations (Fig. 2) it is clear that the larger AESD(h) correspond to zones with few experimental head data.

The uncertainty on the concentration field, characterised by means of the AESD(c) also shows an increase in case that only concentration data are used in the conditioning. AESD(c) for the scenario in which only concentration data are used is 8% larger than the scenario with no conditioning data. This increase can be explained by the fact that the ensemble averaged contaminant plume covers a smaller part of the aquifer in the unconditional case than in the reference field. The part that is not covered by the contaminant plume has a low ensemble standard deviation. For the scenario that also concentration data are used in the conditioning, the position of the contaminant plume is better characterised and its location is closer to the reference

Table 3
Standardised average ensemble standard deviation (scenario 1 = 100) for the log transmissivity, steady-state head and concentration (averaged over 20 time steps) fields

	AESD(Y)	AESD(h)	AESD(c)
Scenario 1	100	100	100
Scenario 2	85	79	95
Scenario 3	96	59	97
Scenario 4	81	51	82
Scenario 5	107	105	108
Scenario 6	88	82	90
Scenario 7	98	61	84
Scenario 8	81	50	69

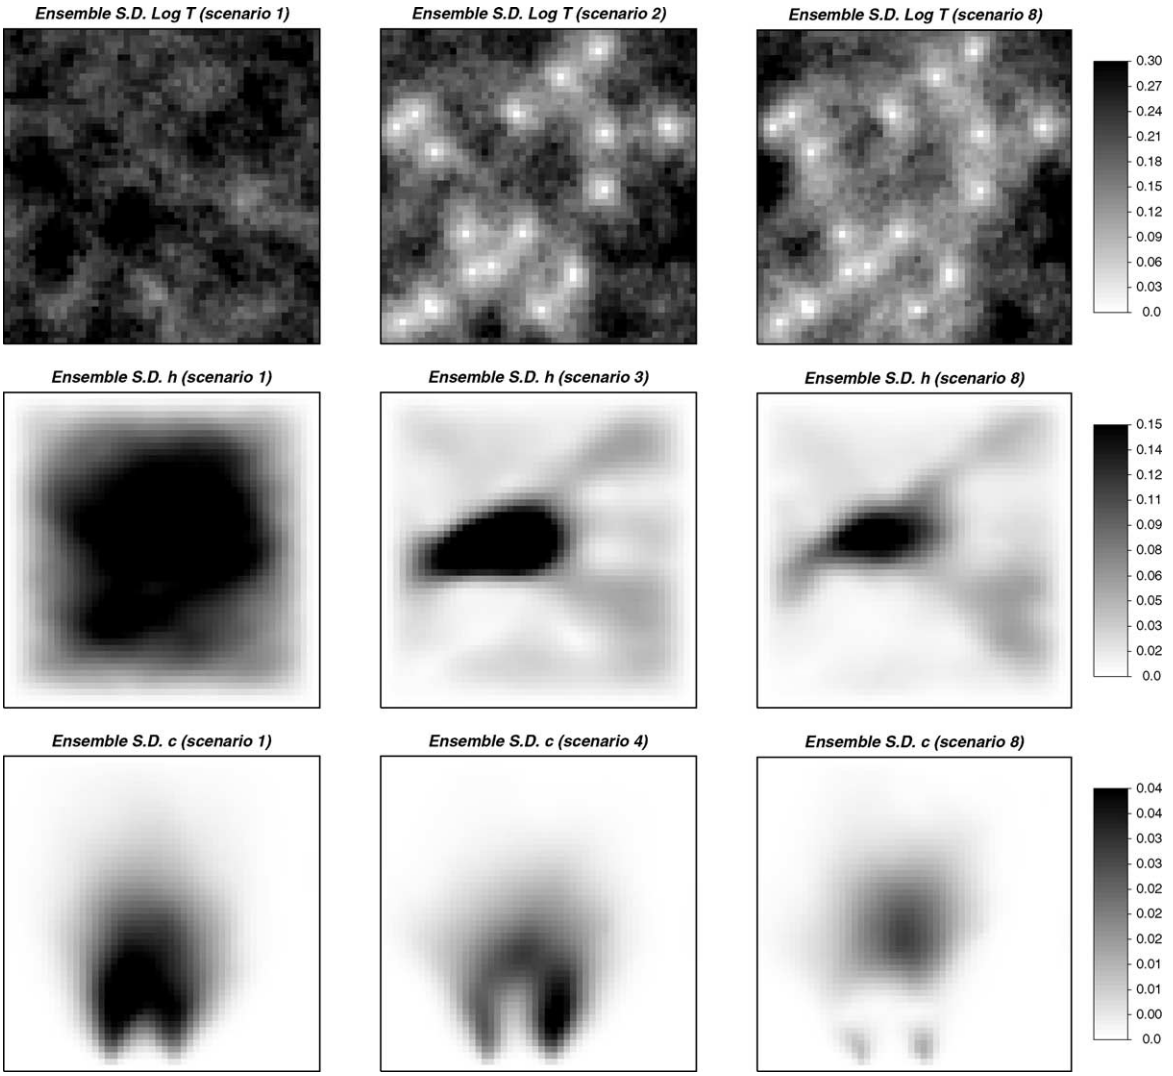

Fig. 6. Ensemble standard deviations of log transmissivity for scenarios 1, 2 and 8; of steady-state head for scenarios 1, 3 and 8; and of concentration at time step 20 for scenarios 1, 4 and 8.

plume. However, this also means that the uncertainty on the contaminant concentration increases in those parts of the aquifer in which for the case of no conditioning data the contaminant concentration was close to zero. In case more conditioning data are available the AESD(c) finally decreases and the lowest AESD(c) is found for the scenario where transmissivity, head and concentration data are available; the AESD(c) reduction is 31%. The AESD(c) reduction is clearly smaller than the AAE(c) reduction. It is thought that this is also—in part—due to the mismatch on the position of the contaminant plume in the unconditional case. Finally, it can be concluded that the concentration data contribute the most to the AESD(c) reduction (in case all the conditioning data are available), but head and transmissivity data also have an important contribution to reduce the uncertainty. Fig. 6 illustrates that the AESD(c) reduce in case more conditioning data are used and shows also that the largest AESD(c) correspond to the zones with an elevated concentration gradient.

4. Discussion and conclusions

The synthetic study illustrates that when concentration data are also used for conditioning, the characterisation of the transmissivity, steady-state head and concentration fields are improved. The impact of the concentration data on the characterisation of the upstream part of the aquifer is especially noticeable. The synthetic study also illustrates that the best results are obtained in case transmissivity, head and concentration data are used together. The uncertainty also reduces in the conditioning process; the lowest variances are found for the cases that transmissivity, steady-state head and concentration data are used as conditioning information. However, in this specific case the presence of just concentration data yielded an average ensemble standard deviation larger than in the unconditional case. The fact that the average plume in the unconditional case is much smaller than the reference plume makes that the ensemble standard deviation of the concentration field first increases when conditioning data become available; the first conditioning data correct the position of the plume and the associated underestimation of the variance of the concentration field.

This study shows the importance of the conditioning to concentration data. However, in this synthetic study just one source of uncertainty—the transmissivities—was considered. In reality many more sources of uncertainty should be considered, like boundary conditions, porosity, local dispersivities and, in some cases, the information on the mass sources. In controlled conservative tracer experiments the amount of injected mass is nearly exactly known and the additional uncertainty may be limited. However, the application of this methodology on contaminated sites with an uncertain mass release history is more complicated. In those studies, the spatial–temporal characterisation of the contamination source is of especial interest. A main question is to what extent it is possible to characterise the aquifer transmissivities and the contaminant source altogether, given a limited amount of conductivity data, a strongly spatially variable hydraulic conductivity and limited information on the contaminant release history. As the methodology allows calibrating mass release information, porosity and retardation, it should be applied in more complicated synthetic cases with multiple sources of uncertainty and in a real-world case study in order to find out what is the worth of concentration data in those cases. Yet, this synthetic study gives promising results with respect to the worth of concentration data.

Acknowledgements

The financial support of the Spanish Nuclear Waste Management Company (ENRESA) and CICYT (project number HID99-0481) is gratefully acknowledged.

References

Capilla, J.E., Gómez-Hernández, J.J., Sahuquillo, A., 1997. Stochastic simulation of transmissivity fields conditional to both transmissivity and piezometric data. 2. Demonstration on a synthetic aquifer. Journal of Hydrology 1–4 (203), 175–188.

Carrera, J., Neuman, S.P., 1986a. Estimation of aquifer parameters under transient and steady state conditions: 1. Maximum likelihood method incorporating prior information. Water Resources Research 22 (2), 199–210.

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

Carrera, J., Neuman, S.P., 1986c. Estimation of aquifer parameters under transient and steady state conditions: 3.Application to synthetic and field data. Water Resources Research 22 (2), 228–242.

Dagan, G., 1985. Stochastic modelling of groundwater flow by unconditional and conditional probabilities: the inverse problem. Water Resources Research 21 (1), 65–72.

Gómez-Hernández, J.J., Srivastava, R.M., 1990. ISIM3D: an ANSI-C three dimensional multiple indicator conditional simulation program. Computer and Geosciences 16 (4), 395–400.

Gómez-Hernández, J.J., Journel, A.G., 1993. In: Soares, A., (Ed.), Joint Simulation of MultiGaussian Random Variables, Geostatistics Tróia '92, vol. 1. Kluwer Academic Press, Dordrecht, pp. 85–94.

Gómez-Hernández, J.J., Sahuquillo, A., Capilla, J.E., 1997. Stochastic simulation of transmissivity fields conditional to both transmissivity and piezometric data. 1. Theory. Journal of Hydrology 1–4 (203), 162–174.

Graham, W.D., McLaughlin, D., 1989a. Stochastic analysis of nonstationary subsurface solute transport. 1. Unconditional moments. Water Resources Research 25 (2), 215–232.

Graham, W.D., McLaughlin, D., 1989b. Stochastic analysis of nonstationary subsurface solute transport. 2. Conditional moments. Water Resources Research 25 (11), 2331–2355.

Hanna, S., Yeh, T.C.J., 1998. Estimation of co-conditional moments of transmissivity, hydraulic head and velocity fields. Advances in Water Resources Research 22 (1), 87–95.

Hendricks Franssen, H.J.W.M., 2001. Inverse stochastic modelling of groundwater flow and mass transport. PhD thesis Technical University of Valencia.

Hendricks Franssen, H.J.W.M., Gómez-Hernández, J.J., 2002. 3D Inverse modelling of groundwater flow at a fractured site using a stochastic continuum model with multiple statistical populations. Accepted for publication in Stochastic Environmental Research and Risk Assessment 16 (2), 155–174.

Hendricks Franssen, H.J.W.M., Gómez-Hernández, J.J., Sahuquillo, A., Capilla, J.E., 1999. Joint simulation of transmissivity and storativity fields conditional to hydraulic head data. Advances in Water Resources Research 23, 1–13.

Hu, L.Y., 2000. Gradual deformation and iterative calibration of Gaussian-related stochastic models. Mathematical Geology 32 (1), 87–108.

Iribar, V., Carrera, J., Custodio, E., Medina, A., 1997. Inverse modelling of seawater intrusion in the Llobregat delta deep aquifer. Journal of Hydrology 198, 226–244.

Kitanidis, P.K., Vomvoris, E.G., 1983. A geostatistical approach to the inverse problem in groundwater modelling (steady state) and one-dimensional simulations. Water Resources Research 19 (3), 677–690.

Mayer, A.S., Huang, C., 1999. Development and application of a coupled-process parameter inversion model based on the maximum likelihood estimation method. Advances in Water Resources Research 22 (8), 841–853.

Medina, A., 1993. Joint estimation of the parameters of the flow and transport equation (in Spanish). PhD thesis, Technical university of Catalonia.

Medina, A., Carrera, J., 1996. Coupled estimation of flow and solute transport parameters. Water Resources Research 32 (10), 3063–3076.

Oliver, D.S., Cunha, L.B., Reynolds, A.C., 1997. Markov Chain Monte Carlo methods for conditioning a permeability field to pressure data. Mathematical Geology 29 (1), 61–91.

RamaRao, B.S., LaVenue, A.M., de Marsily, G., Marietta, M.G., 1995. Pilot point methodology for automated calibration of an ensemble of conditionally simulated transmissivity fields. 1. Theory and computational experiments. Water Resources Research 31 (3), 475–493.

Rubin, Y., 1991a. Prediction of tracer plume migration in disordered porous media by the method of conditional probabilities. Water Resources Research 27 (6), 1291–1308.

Rubin, Y., 1991b. The spatial and temporal moments of tracer concentration in disordered porous media. Water Resources Research 27 (11), 2845–2854.

Sahuquillo, A., Capilla, J.E., Gómez-Hernández, J.J., Andreu, J., 1992. Conditional simulation of transmissivity fields honouring piezometric data. In: Blain, W.R., Cabrera, E. (Eds.), Hydraulic Engineering Software IV, Fluid Flow Modeling, Kluwer, Dordrecht, pp. 201–214.

Sonnenborg, T.O., Engesgaard, P., Rosbjerg, D., 1996. Contaminant transport at a waste residual deposit. 1. Inverse flow and nonreactive transport modeling. Water Resources Research 32 (4), 925–938.

Wagner, B.J., 1992. Simultaneous parameter estimation and contaminant source characterization for coupled groundwater flow and contaminant transport modelling. Journal of Hydrology 135, 275–303.