

10. CONDITIONAL SIMULATION OF UNSATURATED SOLUTE TRANSPORT

10.1 Introduction

The stochastic evaluation of solute transport is of interest where not only the expected behavior of contaminant movement but also the uncertainty associated with the mean concentration prediction must be evaluated. To obtain the statistical parameters of the input random field variables (RFVs, see section 2.5.1) K_s and α in (4-8), measurements must be taken to determine K_s and α at the site that needs to be evaluated (unless these data are available from similar or nearby sites). In many cases, measurements are also available that are related to those two parameters, although they represent a different physical quantity, for example soil tensiometer data or concentration measurements. Data of physical variables that are different from, but related to the constitutive parameters of unsaturated flow are often referred to as "indirect" information. The unconditional stochastic method presented in the previous two chapters ignores any available indirect data and considers only the statistical properties of the "direct" data. The approach is satisfactory in applications where either the lateral extent of the contamination source or the travel distance of interest is very large with respect to the correlation scale of the soil and if the soil is of only mild heterogeneity. In such cases, the actual solute plume is "ergodic" (see chapter 2) i.e., the stochastic mean concentration plume accurately predicts the actual plume and the concentration variance is zero. But for a point source or very localized contamination, the travel distance required for the plume to reach ergodicity may be exceedingly large. Dagan (1986) suggested that the ergodicity assumption is valid only after the plume has been displaced several hundred correlation scales. In the unsaturated zone, this may correspond to several tens of meters (see field studies referenced in the introduction to chapter 8). If soil heterogeneity is found on a number of distinct scales of increasing order, ergodicity may not be achieved at all, even in the deep unsaturated zones

encountered in semi-arid and arid environments (see chapter 9).

The non-ergodic mean plume concentration has a meaning much different from the ergodic mean plume concentration. It merely is a mass conservative, best estimate of the local, time dependent concentration probability. Unlike the pdf of other RFVs, the concentration pdf is difficult to determine due to the non-stationarity in space and time. Hence, the significance of the first two unconditional concentration moments is questionable if the variability is very large. To condition the stochastic evaluation of solute transport on all of the available information - including the deterministic value of single measurement data - is therefore a desirable approach not only to reduce the uncertainty of the concentration prediction but also to fully reflect the information content of the available field measurements.

Conditional stochastic analysis has been applied to a number of groundwater problems. Dagan (1982, 1984) derived analytical perturbation expressions for the conditional moments of the saturated hydraulic conductivity (input variable), the conditional head, and the spatial plume moments (output variables) in a Bayesian framework. The work accounted for local measurements of the hydraulic conductivity, of the head, and of the groundwater pore velocity. Delhomme (1979) used the geostatistical method to generate conditional random input fields of the saturated hydraulic conductivity (chapter 3; Journel, 1974). By generating random fields of K_s and solving the saturated flow equation numerically, he evaluated the conditional head moments through a Monte Carlo simulation. A similar approach was taken by Smith and Schwartz (1981) who not only analyzed the conditional head, but also the conditional solute arrival time to demonstrate the principal effect of conditioning. Binsariti (1980) and Clifton and Neuman (1982) used transmissivity and water table measurements to condition the transmissivity fields. They applied the statistical inverse method introduced by Neuman and Yakowitz (1978) to condition the hydraulic conductivity data on measurements of head. Clifton and Neuman (1982) reported a large decrease in prediction uncertainty with respect to the head moments, when head measurements are included in the conditional approach. The first effort to condition hydraulic conductivity and the velocity field on concentration data was made by

Graham and McLaughlin (1989). They presented a first order analytical stochastic solution based on spectral perturbation analysis and Kalman filtering. This work has to date been the only rigorous approach that allows for conditioning with concentration data. Indirect and direct information is used in the Lagrangian conditional transport analysis by Rubin (1991a), who uses a cokriging approach to obtain the conditional moments of the velocity from measurements of the saturated hydraulic conductivity and/or the head. The covariances and cross-covariances necessary for the cokriging are derived from a linear first order analysis based on Dagan's work (1984). Using the conditional velocity fields, the conditional spatial moments (center of mass and moment of inertia) of the contamination plume are evaluated in a Lagrangian framework by particle tracking. Zhang and Neuman (1994a,b,c,d) develop a new approach to obtain conditional concentration moments, conditional spatial moments of the mean concentration, and conditional solute flux moments based on the Eulerian-Lagrangian transport theory by Neuman (1993). Transmissivity and hydraulic head data are used in their work to condition the concentration moments.

To date, no attempt has been made to also analyze *unsaturated* transport with conditional stochastic methods. Recently, an exact formalism to predict the conditional moments of transient unsaturated flow (but not transport) in heterogeneous media has been suggested (Neuman and Loeven, 1994). In principle, all of the above approaches lend themselves for an analysis of the conditional plume and concentration moments under unsaturated conditions. The main difficulty encountered in the numerical (Monte Carlo) approach on one hand is the prohibitive amount of computation time needed to obtain just one steady-state velocity field from the conditional random input fields of K_s and α . The difficulty of the analytical approaches on the other hand is the derivation of covariance and cross-covariance functions necessary to obtain the conditional velocity moments.

The work presented in the previous chapters has overcome both limitations: An efficient numerical approach to compute steady-state unsaturated velocity fields, given a heterogeneous realization of K_s and α , has been introduced (chapter 7; Harter and Yeh, 1993).

This allows the efficient implementation of Monte Carlo simulations similar to the work by Smith and Schwartz (1981a,b). The stochastic moments of the unsaturated flow variables $f = \log K_s$, $a = \log \alpha$ (log: natural logarithm), and soil water tension (head) h have also been derived (chapter 4; Yeh et al., 1985a,b). With these theoretical moments, Rubin's (1991a) analysis of conditional plume moments can easily be extended to unsaturated flow. His semi-analytical approach, however, is limited to small perturbations. In this chapter the (nonlinear) numerical Monte Carlo technique is applied to derive various conditional stochastic transport parameters without having to linearize either the flow or the transport equation. Linearization is only used to generate conditional input random fields f and a given data of either f , a , or h . For the conditioning, a geostatistical inverse method called cokriging is applied (Myers, 1982; Kitanidis and Vomvoris, 1983).

Conditional simulation of unsaturated transport distinguishes itself from the conditional simulation of saturated transport not so much in the principle of the approach as in the interdependencies between input and output RFVs. The same measurement data play a different role depending on whether they are applied to saturated or unsaturated flow. In unsaturated flow two independent parameters (or more - depending on the choice of the constitutive relationship) define the actual local hydraulic conductivity. The unsaturated flow problem is inherently nonlinear i.e., head and conductivity are interdependent unlike in the saturated case, where the conductivity is independent of the head. It is therefore expected that the data measured in the field and used to condition the stochastic analysis have a relevant information content that is distinctly different from the saturated case. Much of the usefulness of one type of measurement will depend on the availability of other types of measurements. Measuring, for example, either the saturated hydraulic conductivity, or the soil pore size distribution parameter, or the soil water tension each by itself should result in much less conditioning than the combined effect of all three measurements.

The main objective of this chapter is to investigate the role of both indirect information (soil water tension data) and direct information (K_s and α data), and the role of their spatial

distribution (monitoring network or sampling strategy) on the uncertainty of the conditional stochastic prediction of non-reactive solute transport under variably saturated conditions in isotropic and anisotropic soils. A second objective is to discriminate the effect that conditioning has on the various measures of solute transport. Besides analyzing the local concentration moments (Rubin (1991a, Zhang and Neuman, 1994b), the conditioning effects on the spatial plume moments (Dagan, 1982; 1984), on the arrival time (Smith and Schwartz, 1981; Zhang and Neuman, 1994c), and on the integrated breakthrough (Zhang and Neuman, 1994c) at an arbitrary compliance surface are examined. The structure of this chapter is as follows: The theoretical background and the implementation of the conditional unsaturated flow and transport model is described in sections 10.2 and 10.3. The hypothetical field soil sites for the conditioning study are a subset of the example soils described in the previous chapter and are selected in section 10.4. The impact of different sampling strategies or monitoring network designs on the reduction in the spatial moments of solute transport is investigated in sections 10.5 through 10.8. Parameter uncertainty in the context of conditional simulation is addressed in section 10.9. Section 10.10 discusses the role of the spatial plume moments as a measure to judge the effect of conditioning. In many applications involving environmental compliance at a particular location or surface, the variable of interest is the solute arrival time or breakthrough curve and not the spatial plume distribution. In the two sections 10.11 and 10.12, the effect of conditioning on several local and integrated measures of solute travel time is studied. The conditional mean concentration prediction at a highly conditioned site is compared to the deterministic inverse modeling prediction in section 10.13. The chapter closes with a summary and conclusion.

10.2 Theory of Conditional Simulation by Cokriging

In chapter 3.3 a method was introduced to generate conditional random fields of the same random field variable (RFV) of which measurement data are available. In the context of this chapter the term conditioning is also applied to the process of generating random fields (and

their dependent functions) that are not only conditioned on data from the same RFV (direct data e.g., random K_s fields conditioned on K_s data) but also or even exclusively on data from other physically related RFVs (indirect data e.g., random K_s fields conditioned on head data). The conditional simulation technique used in this study is based on the same principles and numerical techniques as the conditional simulation algorithm described in chapter 3.3, equation 3-14 (Journel, 1974; Delhomme, 1979). The important difference is that cokriging rather than kriging is employed because of the multivariate nature of the problem. The kriging equations are given in (2-46) through (2-48). The cokriging equations are identical to the kriging equations (2-46), (2-47) in chapter 2 (Carr and Myers, 1985). However, the array of measured data \mathbf{X}_1 in (2-46) contains data from more than one RFV e.g., from saturated hydraulic conductivity data *and* head data, while the array of unknown data \mathbf{X}_2 is - as in kriging - comprised of data exclusively from one RFV e.g., the saturated hydraulic conductivity K_s . Λ_{12} in (2-46) is the weight matrix of the measured data \mathbf{X}_1 with respect to the estimate \mathbf{X}_2 , which is either of the same RFV as \mathbf{X}_1 (kriging) or of a different RFV (cokriging). In either case the kriging weight matrix Λ_{12} is computed by solving the covariance matrix equation (2-47). For the cokriging case, the cross-covariances between two RFVs must be known to determine the matrices \mathbf{C}_{11} and \mathbf{C}_{12} in (2-47). Note, that the covariance and cross-covariance functions must be positive definite, otherwise (2-47) has no general solution (Myers, 1982).

10.3 Conditional Monte Carlo Simulation: Methods

10.3.1 Principal Elements of the Monte Carlo Algorithm

The principal procedures in the conditional Monte Carlo simulation are identical to those of the unconditional Monte Carlo simulation introduced in chapters 8 and 9 (see Figure 8.2). Conditional realizations of f and a are generated and a conditional approximate solution h is computed explicitly. The realization of each of these three RFVs is passed to MMOC2, which computes the steady-state soil water tension through a finite element solution of Richards equation, the flux field through a finite element solution of Darcy's law, and the transient solute transport by using a modified method of characteristics (chapter 5). The procedure is repeated for 150 to 300 realizations (see below). Finally, the appropriate statistical sample parameters are computed from the output of the Monte Carlo simulation. The only difference between the conditional simulations in this chapter and the unconditional simulations in chapter 9 is the algorithm used to generate the random field realizations f and a and the approximate solution h , all of which must be conditioned on measurement data, which are provided as input. As in the previous chapter, the statistical parameters describing the RFVs f and a are assumed to be known. The next sections discuss the actual implementation of the conditional random field generator and the conditional extension of the ASIGN method described in chapter 7. A flow chart of conditional ASIGNing and Monte Carlo simulation is shown in Figure 10.1.

10.3.2 Generating Conditional Random Fields

The conditional random field generator developed for this study is an extension of the spectral random field generator described in chapter 3. Unconditional spectral representations dZ_f and dZ_a are generated. The unconditional realizations f and a are computed from their respective spectral representations via fast Fourier transform (FFT). The unconditional random field realizations are needed to obtain the conditional random field realizations from (3-14).

However, following the approach by Gutjahr et al. (1992), (3-14) is rearranged and the kriging equation (2-47) is used to explicitly write the kriged terms in (3-14). Then the algorithm for generating the conditional realization can simply be written as:

$$\mathbf{X}_{2s}^c = \mathbf{X}_{2s} + \Lambda_{21} (\mathbf{X}_1 - \mathbf{X}_{1s}) \quad (10-1)$$

where Λ_{21} is the kriging weight matrix, \mathbf{X}_{2s} is the unconditionally generated mean removed realization f or a . \mathbf{X}_1 is the array of field measured data (unconditional mean removed), which may include f , a , and h data. \mathbf{X}_{1s} is the array of unconditionally simulated data at the particular locations, where measurements of the same variable are available in the field site (also unconditional mean removed). \mathbf{X}_{2s}^c is the conditional mean removed realization f^c or a^c . (10-1) is computed once for each realization of each RFV.

If \mathbf{X}_1 contains any head measurements (soft conditioning) or if (10-1) is used to compute a linearized conditional solution $h_L^c = \mathbf{X}_{2s}^c$ (conditional ASIGNing, see below), the unconditional realization h must be computed from the unconditional realizations f and a to fill \mathbf{X}_{1s} or \mathbf{X}_{2s} or both. Conditional flow simulation therefore requires that the unsaturated flow equation be solved twice: once to obtain the unconditional random field h from the unconditionally generated realizations f and a , and a second time to obtain the conditional nonlinear solution h^c from the conditional realizations f^c and a^c .

In the classical conditional approach, the unconditional head solution h is computed using standard finite difference or finite element models. A more efficient method would be to use ASIGNing (chapter 7), which combines the spectrally derived first order, linear approximation h_L of the head with the finite element model MMOC2. This is still a computationally very expensive conditioning algorithm. For this study a much more efficient alternative is chosen: The computation of the "true" nonlinear unconditional head $h(f,a)$ with MMOC2 is omitted altogether. Instead the linear approximation $h_L(f,a)$ is used to fill either \mathbf{X}_{1s} or \mathbf{X}_{2s} or both on the right-hand side of (10-1). Recall that the linear head solution h_L is simply cogenerated with f and a by explicitly solving for $dZ_h(\mathbf{k}) = f(dZ_r(\mathbf{k}), dZ_a(\mathbf{k}), H, \Gamma)$ (4-26) and by

applying the FFT on dZ_h . Using a linearized unconditional solution h_L in the conditioning process is consistent with the linear estimation procedure (10-1) from which the conditional random fields f^c and a^c are obtained.

10.3.3 Conditional ASIGNing

The conditional random field realization h_L^c from (10-1) is used as initial approximation of the nonlinear finite element solution h_{FE}^c to the unsaturated flow equation given the conditional realizations f^c and a^c . This is equivalent to the ASIGNing procedure for the unconditional flow simulation (chapter 7), and is therefore called conditional ASIGNing. The conditional random fields h_L^c obtained from (10-1) are accurate enough to allow the numerical algorithm to converge very efficiently to the conditional finite element solution $h_{FE}^c(f^c, a^c)$. No rigorous study similar to that in chapter 7 has been implemented here to determine for which range of soil heterogeneity conditional ASIGNing leads to efficient numerical steady-state flow solutions. However, in this work conditional ASIGNing has successfully been applied to soils with $\sigma_y^2 \leq 3.2$.

It must be emphasized that the use of the spectrally derived linear head solution (4-26) leads to a double advantage in the conditioning algorithm (10-1): It allows for a very efficient evaluation of the unconditional head field (which reduces the CPU time by approximately three orders of magnitude compared to using MMOC alone, see Figures 7.6-7.8). The unconditional random head solution in turn is needed not only to condition f and a , but also to provide an initial approximation of the conditional head field such that the finite element solution converges approximately two orders of magnitude faster than without such an initial approximation (chapter 7). The computational savings achieved by using the linear, spectral head solution h_L in this context are so enormous that an entire conditional Monte Carlo simulation of unsaturated steady-state flow with several hundred realizations can be carried out as CPU-efficiently as a single conditional realization based on finite element solutions alone (i.e.

without using the linear, spectral head solution).

10.3.4 Covariances and Cross-covariances for the Cokriging Matrix Λ_{21}

To obtain the conditional random field \mathbf{X}_{2s}^c , the difference term $(\mathbf{X}_1 - \mathbf{X}_{1s})$ in (10-1) is cokriged using the kriging matrix Λ_{21} , which remains identical for all realizations. Due to storage space limitations, however, the entries to the kriging matrix Λ_{21} are actually recomputed for every realization. With a field size of over 10,000 nodes (size of \mathbf{X}_{2s} and \mathbf{X}_{2s}^c) and up to almost 1000 data measurement points (a maximum 320 measurements of each of the three variables f , a , and h) (size of \mathbf{X}_1 and \mathbf{X}_{1s}), the size of the Λ_{21} matrix would exceed 10 million entries for each of the three RFV fields, which adds to the equivalent of 3×80 Mb of memory when stored in double precision (8 byte per entry).

The cokriging matrix Λ_{21} is obtained by solving the covariance matrix equation (2-47). The solution is computed by inverting \mathbf{C}_{11} using Cholesky decomposition. The inverted matrix is then multiplied with \mathbf{C}_{12} . The subroutine SPPICD in ESSL (IBM, 1993) is used for the matrix inversion.

The cross-covariance functions C_{fh} and C_{ah} and the covariance function C_{hh} in the covariance matrices \mathbf{C}_{11} and \mathbf{C}_{12} can be computed from the analytical linear (cross-)spectral density functions S_{hh} , S_{fh} , and S_{ah} (chapter 4). In chapter 8 it was shown that these quasi-analytical, linearized (cross-)covariance functions are in qualitative agreement with the numerically determined, nonlinear sample functions, but differ in their absolute values if the perturbations are large. For the conditional simulation a modified quasi-analytical solution of the (cross-)covariances is developed based on a calibration of the quasi-analytical, linear (cross-)covariances against the numerically determined sample (cross-)covariance functions. The analytical functions are multiplied by a correction factor (defined separately for each (cross-)covariance function and for each example soil) such that the modified analytical solutions match the numerical, nonlinear sample (cross-)covariance functions with minimal

error. The calibration of the analytical cross-covariance functions is implemented by visual matching. The head covariance is calibrated such that the modified analytical solution for the variance exactly matches the numerically determined variance. Preliminary experiments were implemented and it was found that the results from the Monte Carlo simulation are very robust with respect to the potential error in the calibration procedure.

Calibrating the analytical (cross-)covariances rather than directly employing the empirical (cross-)covariance functions derived in chapter 8 has two advantages: The analytical cross-covariance functions are found to yield invertible covariance matrices C_{11} . However, no attempt has been made to rigorously prove that either the analytical (cross-)covariance functions or the calibrated (cross-)covariance functions yield positive definite matrix functions. This point needs to be further investigated. The second advantage is that the discretization and domain size of the empirical covariance function becomes irrelevant if the empirical functions are only used for calibration of the analytical functions. The transport simulations in the last and in this chapter, for example, are carried out in a domain roughly four times as large as the empirical unsaturated flow studies in chapter 8. Only the calibrated analytical (cross-)covariance functions provide values of C_m , C_{ah} , and C_{hh} at the large lag-distances needed in this conditional flow and transport study. Since the small lag-distances are the most important ones in terms of cokriging, it is sufficient to use the empirical cross-covariance solution at short lag-distances to obtain a reliable, calibrated analytical cross-covariance functions even at large lag-distances.

10.3.5 Nodal and Elemental Properties in the Finite Element Model vs. Grid Properties in the Spectral Random Field Generator

In the finite element realization that serves as hypothetical field site (see next section), head h and concentration c are nodal values while the saturated hydraulic conductivity f and the pore size distribution parameter a are element properties. In contrast, the spectral random field generator and conditioning algorithm assume an identical grid and support for all variables. For the purpose of the conditional simulation it is simply assumed that the support scale of the nodal and elemental properties that are "measured" are identical, and that the bottom left node of each element has the same support and location as the element itself. This introduces a small error in the computation of the cross-covariances (which are functions of distances between measurement points; see chapter 9). The error is negligible since the element discretization is rather small compared to the correlation scale. To be consistent, the assignment of nodal and elemental properties in the finite element model from the conditional random field realizations f^c and a^c and the initial head h^c follows the reverse order: The f , a , and h value at the i^{th} column in the j^{th} row of the conditional random fields are assigned to the i^{th} element in the j^{th} element row (f , a) and to the i^{th} node in the j^{th} nodal row (h), which is the bottom left node to the i^{th} element in the j^{th} element row.

10.4 "Field Test Sites" and Sampling Strategies: Methodology

10.4.1 "Field Test Sites"

The so-called "field sites" that are investigated here are computer-generated hypothetical soil cross-sections (see also comment in the introduction to chapter 8). Computer-generated field-sites allow a rigorous analysis of the information content of measurement data that can be retrieved through conditional stochastic simulation. In the artificial field sites "field" hydraulic properties and the movement of the contamination plume can be perfectly sampled. The

physical processes governing the flow and transport of the so-called "real" plume and the random properties of the soil are perfectly known. Measurement errors, parameter estimation errors, and upscaling problems can for the moment be neglected.

A number of field sites are studied in this chapter to evaluate the amount of uncertainty reduction achieved by measuring relevant data *in situ*. Each "field-site" is a single, randomly chosen realization from the unconditional Monte Carlo simulations of the previous chapter. Only a subset of the soil types and moisture conditions simulated in chapter 9 is chosen for the purpose of conditional simulation. The subset includes isotropic, anisotropic, wet, and dry soils, soils with high textural variability and soils with moderate textural variability, soils with correlated f and a , and soils with uncorrelated f and a (soil sites from simulations #3, #12, #15, #21, #22, #28 in chapter 9, see also Table 9.1). Independent of the mean moisture content or the textural variability of the soil, soils with similar unsaturated hydraulic conductivity statistics are expected to behave alike not only with respect to the unconditional concentration moments (see chapter 9) but also with respect to the conditional concentration moments.

Soil #3 is isotropic with $\sigma_y^2=0.86$ and a weakly anisotropic covariance in y . All other soils are strongly anisotropic with a vertical correlation scale of f , λ_{fz} , equal to 50 cm and a horizontal correlation scale of y , λ_{fx} , equal to 300 cm. Soil #12 (wet, correlated f and a) has the lowest $\sigma_y^2 = 0.53$. Soils #28 (wet, uncorrelated f and a) and #15 (dry, uncorrelated f and a) have similar $\sigma_y^2 = 1.76$ and $\sigma_y^2 = 1.47$, respectively, although their textural variability (σ_f^2 , σ_a^2) differs. Soils #21 (very dry, correlated f and a) and #22 (wet, uncorrelated f and a) both have a very high $\sigma_y^2 = 3.12$ and $\sigma_y^2 = 3.16$, respectively. These sites are grouped into four categories of soils: Isotropic soil with mild to moderate variability (#3), anisotropic soil with mild variability (#12), anisotropic soil with high variability (#28, #15), and anisotropic soils with very high variability (#21, #22).

10.4.2 Sampling Strategies

Two basic sampling networks were designed for each of the three parameters f , a , and h : a "sparse" and a "dense" sampling network. The sparse sampling network consists of measurement locations along three vertical columns near the plume source (e.g. tensiometer nests or boreholes) spaced 1 horizontal λ_{fx} apart with measurements at every $2\lambda_{fz}$ depth interval (e.g. Figure 10.5e). No data are sampled from an area within $2\lambda_{fz}$ of the bottom boundary of the simulation domain. The center column intersects the source area of the solute plume (see chapter 9). The total number of data points in the sparse network is 40.

A dense sampling network consists of double the data-density of the sparse sampling network i.e., $0.5 \lambda_{fx}$ in the horizontal and $1 \lambda_{fz}$ in the vertical. In the dense sampling network, the data are sampled throughout the entire simulation domain except the area within $2\lambda_{fx}$ of the vertical simulation domain boundaries (e.g. Figure 10.6a) resulting in a total of 320 data points per RFV. Measurements of K_s and α are obtained at identical locations. The sampling grid for the head measurements is shifted both in the vertical and horizontal direction such that a head measurement point is at the center between four adjacent K_s measurements (Figure 10.5i). The dense sampling network also includes measurements at all locations (nodes or elements) within and adjacent to the contamination source, which is defined on the nodal grid. Monte Carlo simulations are implemented with various combinations of f , a , and h sampling networks as listed in Table 10.1.

10.5 Conditional Simulation of Unsaturated Flow

The results of the conditional flow simulations are important to subsequently understand the behavior of the conditional solute plume, since conditioning directly affects the uncertainty about the prediction of the soil water tension h and the logarithm of the unsaturated hydraulic conductivity y . The reduction of the velocity and concentration variance is only an indirect consequence of the conditioning. The effect of conditioning on the statistical moments of y , h , and the pore velocity \mathbf{v} is demonstrated for the two conditional simulations A and G (Table 10.1)

of field site #28, a highly variable, wet, anisotropic soil with uncorrelated f and a . Monte Carlo simulations of other field sites give qualitatively similar results. Field-site #28 (like all other sites except #3) is a vertical cross-section that is 12.8m deep and 24m wide, which is approximately as wide but twice as deep as the experimental Las Cruces trench site (Wierenga et al., 1991). The variability of f and a is 2.25 and 0.04, respectively. The two parameters are considered to be independent of each other. The geometric mean of alpha, Γ , is 0.01 cm^{-1} . Recall, that the soil is anisotropic with correlation scales for f and a of 3m and 50 cm in the horizontal and vertical direction, respectively. The cross-section is therefore about $26 \lambda_{fz}$ deep.

Figure 10.2a-d shows the actual field values of the maps of y , h , and the pore velocity components v_x and v_z at site #28. The cross-section of y has the typical random character described in chapter 8. In the center of the cross-section a high conductivity lens is layered immediately above a relatively low conductivity area (Figure 10.2a), which correlates with a partially saturated, very wet lens overlying a dry area with a relatively high tension (Figure 10.2b). The horizontal velocities reach some of their highest absolute values in this large, wet region, because flow is around the low conductivity area. Another distinct feature of the horizontal velocity field at the site is a strong positively diagonal downward/sideward flow immediately beneath the contamination source (Figure 10.2c). The vertical velocity field has the typical pattern of broad low velocity areas interrupted by relatively narrow vertical bands of higher velocities (Figure 10.2d, see chapter 8).

The main features are preserved in the mean predictions of the two conditional simulations A (Figure 10.2e-h) and G (Figure 10.2i-m). A is based on a dense data network of all three variables f , a , and h , while G is based on only 40 soil water tension data (sparse network) from three tensiometer nests. As the number of conditioning data decreases, the mean prediction becomes more and more uniform, asymptotically approaching the stationary moments of the unconditional simulation. Since only head information is available to conditional simulation G, it is not surprising to find that the head data are those best preserved.

Figure 10.3a-d again shows the field site values of y , h , and \mathbf{v} together with the variance

distribution of the two simulation A (Figure 10.3e-h) and G (Figure 10.3i-m). In conditional simulation A, the variance of y (Figure 10.3e) decreases to less than 0.1 at the f and a measurement points (but not at the head measurement points) compared to an unconditional variance of 1.8. The extremely low variance is very local. Due to the anisotropic structure of the soil, the conductivity variance is reduced significantly stronger in the horizontal vicinity of the measurement points than in the vertical vicinity. Between rows of measurements, the conductivity variance increases to values of nearly 0.7 in conditional simulation A. In the sparse data conditional simulation G, the minimum local variance of y is several times larger than in conditional simulation A with values between 0.4 and 0.7 (Figure 10.3i).

The head variance in much of the area with tensiometer data is reduced to less than 150 cm^2 in conditional simulation A (Figure 10.3f) from 4900 cm^2 in the unconditional simulation (chapter 8). The head variance, however, is nowhere less than 120 cm^2 . In conditional simulation G where no other data are used besides head, the variance reduction is not as strong with a minimum variance of less than 300 cm^2 (Figure 10.3k). Despite the use of conditioning head data, the head variance does not become zero at the measurement points. Neither does the mean head at those locations always coincide with the measured value (Figure 10.4a). This is an artifact of the linear conditioning procedure (Kitanidis and Vomvoris, 1983; Peck et al., 1988; Yeh et al., 1993; Gutjahr et al., 1994): The conditional realizations f^c and a^c are obtained through *linear* estimation (10-1) from head measurements (among others). But the conditional realization h^c is computed by solving the *nonlinear* flow equation (5-1). Note that the head measurement data cannot be applied as internal boundary conditions in the finite element solution.

Although not entirely consistent, the conditioning technique generally gives satisfactory results. It is weakest in areas where steep head gradients exist e.g., in the center of the simulation domain between the very wet and very dry areas mentioned above (Figure 10.4a). In those areas the head changes rapidly with distance, and if the location of the steep head gradient is not predicted with a very high accuracy, large head variances and a significant

deviation of the mean head from the measured head are the result. At field sites with smaller σ_y^2 , the discrepancy between measured and mean conditional head decreases (Figure 10.4b,c). Future work will have to assess how much more consistency is achieved by using a numerical nonlinear solution instead of a first order perturbation solver within the conditioning algorithm.

In the vicinity of datapoints the horizontal velocity variance is reduced stronger than the vertical velocity variance: from an unconditional variance of 0.02 (cm/d)^2 to less than 0.004 (cm/d)^2 , a reduction of over 80% (Figure 10.3g,l). In contrast, the vertical velocity variance is reduced only to approximately 60%-80% of the unconditional variance (Figure 10.3h,m). The conditioning effect in the horizontal direction is very strong for the horizontal velocity, which can be particularly well seen in the results for simulation G. The vertical velocity components are well-conditioned by data in the immediate vertical vicinity, while conditional f , a , and h data have little effect on the vertical velocity in the nearby horizontal vicinity. This is not surprising since the vertical velocity covariance has a strong vertical correlation scale and a very short horizontal correlation scale (see chapter 8).

10.6 Sampling Network Design Impacts on Concentration Prediction

10.6.1 Organization of Graphical Output for Concentration Moments

The maps of the actual plumes, of the mean concentration m_c , and of the concentration coefficient of variation CV_c are plotted in Figures 10.5 - 10.22. The organization of each of these figures is identical: Each figure is divided into twelve panels plotted in four rows and three columns. Each panel shows a vertical cross-section of the soil site with the horizontal axis being the horizontal distance [cm] and the vertical axis being the vertical distance [cm]. Each vertical row represents the results of one Monte Carlo simulation (MCS) or of the actual plume movement. The variable that is mapped in the panels is indicated above the top panel of each column. Each panel in a column represents a different output time. It increases from top to bottom and is measured in dimensionless units $t' = V_z / \lambda_{fz}$. V_z is the (arithmetic) sample

mean vertical velocity computed from the unconditional MCS for the particular soil site (chapter 8). λ_{fz} is the vertical correlation scale of f . $t=0$ is the time of solute release. The initial area of uniform concentration is indicated by the small black box in each panel. The output times are identical to those of the unconditional MCSs in chapter 9 (see Table 9.2) and are indicated in the panels of the leftmost column of each figure. Each row corresponds to only one output time. The concentration maps are plotted with five gray-shaded contour levels, to which the labels are found at the right side of the rightmost column. The soil site number (#) to which the results belong is indicated in the bottom right corner. Note that the actual and mean concentration contour levels at a particular output time t' are identical for the maps of all MCSs of one site and correspond to the contour levels chosen for the actual plume. All concentration data are normalized with respect to the initial concentration c_0 . The contour levels have logarithmic intervals (log base 10) and range over two orders of magnitude such that the maximum contour level is at the most 15% below the peak concentration of the actual field solute plume at time t' (e.g. if the peak concentration is $9.45E-2$, the contour levels are from $9.00E-4$ to $9.00E-2$). In all panels showing CV_c maps, the CV_c contour levels are in increments of 0.3 in the interval from 0.5 to 2 as indicated by the labels on the right side of each figure. In addition, the minimum CV_c of each CV_c map is printed out explicitly and the location is indicated where necessary. The first panel in each column indicates the locations where conditional data are available for the MCS shown in that column. Open circles are datapoints, at which f and a data are measured. Black dots indicate soil water tension measurement points. Unless it is the map of the actual plume, the second panel in each column indicates the type of MCS with the amount of data used for conditioning.

10.6.2 Solute Plume Movement at the Field Site

Again, field site #28 and the results from its Monte Carlo simulations are discussed at length in this section. The Monte Carlo simulations of other example soils are discussed more

briefly in subsequent sections, where the emphasis is a comparative analysis. The plume dynamics at field site #28 are depicted in Figure 10.5a-d. Initially the plume moves diagonally downwards along a strong diagonal velocity field (Figure 10.2c,d), the tip of the plume splits into two at an early travel time and spreads horizontally as it reaches a large wet area located above a relatively dry lens in the center of the simulation domain (see previous section). At late time $t=31$ residual concentration is found primarily within and underneath the dry, low permeability area. The plume is distinctly non-Gaussian with no tendency towards a more Gaussian behavior even at late times (see also chapter 9).

10.6.3 Sensitivity of Concentration Moments to Sampling Networks (Site #28)

"Dense" sampling network for all parameters (f, a, h): simulation A

In this first example, data are available at a high sampling frequency (every five nodes in each direction or $0.5\lambda_r$ in the horizontal and $1\lambda_r$ in the vertical). Both independent RFVs f and a are sampled at identical locations. The soil water tension h is measured at nearby locations (conditional simulation A, Table 10.1). This is the highest density measurement grid used in any of the simulations (Figure 10.5e). From a practical point of view, such a high density of observation points cannot be achieved without partially removing or destroying the site (e.g., in the trench site experiment by Wierenga et al., 1989) But this type of conditional simulation serves as a benchmark test to illustrate by how much prediction uncertainty can be reduced in an optimally sampled field site.

As would be expected, the conditional mean concentration distribution is very similar but not identical to the actual concentration distribution (Figure 10.5e-h). The conditional solute plume shows many of the broader patterns of the actual field plume, but does not distinguish between some local random patterns. The length scale associated with the differences between the conditional plume and the actual plume is significantly larger than the scale of the sampling

intervals for the f , a , and h data (compare actual and mean plume at, for example, $t'=8$). There are several reasons to explain why it is possible that the conditional simulation uncertainties go beyond the confinement of the measurement grid:

- a. The available field measurements are only indirect pieces of information with respect to predicting the movement of the solute plume. Neither concentration nor velocities are measured directly to confine the predictions.
- b. Even the information about the hydraulic conductivity itself is not entirely certain anywhere because head measurements are taken at different points than f and a measurements. A small amount of uncertainty about y remains even at the f and a data locations, especially in this highly heterogeneous soil (see Figure 10.3e).
- c. The conditioning on the head data has been implemented only in a linear, approximate manner i.e., the conditioning algorithm does not yield a zero head variance at the head data locations (see section 10.3 and results in section 10.5).

The conditional concentration prediction from simulation A clearly shows the early diagonal displacement, the horizontal spreading along a low permeability zone in the center of the simulation domain, the breakthrough to the bottom boundary in the right half of the domain at $t'=16$, and the residual concentration below the low permeability zone at late time $t'=31$. The movement of the highest concentrations, or the plume center, is predicted very accurately. Error in predicting lower concentrations is indicated by the slightly larger vertical and horizontal spreading of the conditional mean plume. Conditioning at this density is very useful not only for predicting the movement of the center of a contaminant plume but perhaps more importantly the general patterns of the fringes of the contamination plume (indicated by the $c/c_{\max}(t') = 0.01$ contour line). The approach seems particularly useful to identify possible preferential transport paths on one hand and solute retention areas on the other hand.

To assess the difference between prediction and actual plume the absolute nodal error measure E is introduced:

$$E = \sum_{\mathbf{x} \in \Omega} |c_f(\mathbf{x}) - \langle c(\mathbf{x}) \rangle^c| \quad (10-2)$$

where Ω indicates the simulation domain, $\langle c(\mathbf{x}) \rangle^c$ is the conditionally simulated mean concentration, and $c_f(\mathbf{x})$ is the actual field-site concentration at node location \mathbf{x} . For the four output times $t' = 4, 8, 16,$ and 31 in Figure 4.5e-h, the absolute nodal error in the conditional simulation is $E = 5.4, 4.7, 3.9,$ and $2.8,$ respectively. This is approximately 40% less than in the unconditional simulation with $E = 9.0, 8.0, 8.7,$ and $4.0,$ respectively (Figure 10.7i-m). Note, that the nodal error measure emphasizes errors in the regions of high concentrations. Errors occurring away from the center of the conditional plume are small and therefore insignificant with respect to E and cannot be reflected by such a measure. Often, however, very low concentrations are of equal concern. Then a "success" measure like E can be very misleading.

A statistically important measure of the quality of the simulated prediction is the concentration coefficient of variation CV_c where $CV_c = \text{std}_c / m_c$, the ratio of the concentration sample standard deviation over the sample mean concentration. Figure 10.8e-h depicts the dynamics of the CV_c plume for conditional simulation A. The CV_c plume depicts areas of least uncertainty (darkest colors). At the $CV_c = 2$ contour line, it is generally as wide or wider in the horizontal direction as the mean concentration plume at the 1% $c/c_{\max}(t')$ contour line, but vertically less extensive than the mean plume. This is consistent with the $0.5\lambda_f$ horizontal sampling density vs. a $1.0\lambda_f$ vertical sampling density and the general anisotropic structure of the soil. It is also consistent with the fact that the head correlation scale is larger in the horizontal direction than in the vertical direction. Conditioning reduces concentration uncertainty stronger into the horizontal distances than into the vertical distances from the measurement point. The minimum coefficient of variation in the center of the CV_c plume increases with time from 0.35 at $t'=4$ to 0.56 at $t'=31$ indicating increased uncertainty near the center of the plume. The minimum CV_c location coincides with the location of the peak mean concentration. Like in the unconditional simulations of chapter 9 (see also Figure 10.7i-m), the

area of low uncertainty defined by the $CV_c=2$ contour increases with increasing mean plume size and time (compare Figure 10.8e-h with Figure 10.5e-h).

Numerical mass balance problems contribute approximately 0.1 to the CV_c . This value is estimated by computing the coefficient of variation of the mass balance variability between different realizations. Initially the numerical mass balance error $CV_{bal}(t) = \text{std}_{bal}(t) / \text{mass}_{tot}$ (standard deviation of the total mass balance in the domain divided by the total initial mass) is zero, then rapidly increases at early time and reaches a relatively stable plateau of 0.1. The mass balance error is inherent to the modified method of characteristics and must be attributed to the heterogeneous velocity field, for which the fourth order Runge-Kutta travel path integration is known not to be accurate (see chapter 5).

Not sampling α , dense grid

Simulation B (Figure 10.5i-m) uses the same f and h data as simulation A, but assumes that nothing is known in situ about $\log\alpha$. Only its unconditional mean and variance and the form of its pdf are given. The difference between the conditional concentration moments of simulation A and B is small, because the variance of $\log\alpha$ is not very large and the soil is rather wet. The information about K_s and h gives reasonably accurate estimates of the unsaturated hydraulic conductivity even without measurements of α . The mean plume is slightly larger with a smaller peak concentration. At early time ($t'=4$), simulation B has 25% higher CV_c than simulation A, but as time proceeds, the differences in CV_c decrease (Figure 10.8i-m).

Sampling saturated hydraulic conductivity only

The uncertainty reduction relative to the unconditional MCS that is achieved from saturated hydraulic conductivity measurements alone is considerably smaller than in the

previously discussed conditional simulations A and B. When sampled in a dense network (simulation C, Figure 10.6a-d), the minimum CV_c at $t'=4$ is 0.81, almost double as large as in simulation B, which includes the head measurements in addition to the data used in simulation C. The ratio of the minimum CV_c in simulation C over that in the unconditional simulation decreases only slightly with time. Since the variance of a is moderate, f data are helpful in discriminating the most probable fast flow paths from likely slow flow areas. The diagonal flow path near the source is obvious in the conditional mean prediction, and so is the low permeability area in the center of the cross-section. Clearly, the peak concentrations in C are lower than in simulations A and B due to the larger concentration variability. When f is sampled only on the sparse network (not shown) the results are almost identical to those shown in Figure 10.6e-h for conditional simulation D with f and a data on the sparse network. The indifference between the two simulations is again due to the moderate variability of a , its relatively small mean value, and the relative wetness of the soil (the unsaturated hydraulic conductivity differs little from the saturated hydraulic conductivity). Compared to the dense sampling network for f (simulation C), the conditional solute plume D is considerably more disperse, particularly at later time. The minimum $CV_c=0.88$ at $t'=4$ is a less than 10% increase over simulation C but $CV_c=1.29$ at $t'=16$ is a 30% increase over simulation C. These values approach those for the unconditional simulation, even exceed them at $t'=16$. The CV_c away from the plume center, however, is always less in conditional simulation D than in the unconditional simulation.

Sampling soil water tension only

For the conditional simulation G, three hypothetical tensiometer nests are installed three meters or one correlation length apart with tensiometers placed every one meter or two correlation lengths in the vertical (sparse sampling network). The forty head measurements are used for conditioning the f and a input random fields. Sparse sampling network G yields the

least conditioned simulation in this study. Nevertheless the tensiometer data cause a significant improvement in the mean concentration prediction compared to the unconditional mean concentration prediction (compare Figure 10.7e-h and Figure 10.7i-m). The mean plume movement indicates the initial diagonal movement, the split of the plume into two lobes and it hints at a low permeability zone in the center of the simulation domain. The peak concentrations are much closer to the actual peak concentration than in the unconditional simulation. The estimates of both the front and the tail of the plume are more realistic than in the unconditional simulation.

If the number of tensiometers is doubled in both the vertical and horizontal direction and extended over a larger cross-section (conditional simulation F, Figure 10.7a-d), the prediction of the expected concentration does not improve very much. This shows that additional measurements of the tension are not associated with an equal amount of uncertainty reduction. The minimum CV_c at $t'=4$ are 0.50 and 0.65, respectively, compared to 0.95 in the unconditional simulation. Similar ratios between the minimum CV_c of the different simulations are obtained at later times. The minimum CV_c in both simulations F and G are significantly lower than in simulation C (dense f data, direct conditioning) (0.50, 0.65 vs. 0.81). In terms of uncertainty it therefore appears that soil water tension data by themselves yield a greater improvement of the prediction than saturated hydraulic conductivity data. Only at $t'=31$, the minimum CV_c in simulation G is higher than in simulation C, while the minimum CV_c in simulation F (dense h data, indirect conditioning) remains below that in simulation C. This is in partial contrast to the mean concentration prediction, which - particularly at $t'=16$ and $t'=31$ - seems significantly better in simulation C than in either simulation F or G. The relatively lower minimum CV_c in simulations F and G is probably caused by the smaller horizontal spreading of the mean concentration plumes and its higher peak concentration.

Also, for both the sparse and the dense sampling networks, the conditional CV_c plumes seemingly "know" more about the actual plume behavior than the respective m_c plumes, since their spatial pattern better mimic the actual plume. This latter observation is partially

coincidence, partially due to the particularly strong horizontal spreading of the plume around $t'=16$. Horizontal velocity components have a stronger horizontal correlation than vertical velocity components. Similar observations are not made in MCSs of other field sites.

The results underline the importance of soil water tension data in the conditional simulation of transport in highly heterogeneous flow fields. They also indicate that the minimum CV_c alone can only serve as a guideline to measure reduction in prediction uncertainty.

Other sampling network combinations for f, a and h

Figure 10.6i-m shows the mean concentration results for conditional simulation E based on a sparse network of f and a measurements combined with a dense network of h measurements. Again, the results are almost identical to those with sparse f data alone together with a dense h measurement network (conditional simulation H, Figure 10.21a-d). The spatial concentration distribution is significantly better predicted than in either simulation D or F alone. The mean plume is much less dispersed resulting in higher concentrations at the center of the plume. The improvement is particularly visible at $t'=16$, when both the actual and the mean plume exhibit the strongest horizontal spreading. The visual information is supported by the results for CV_c (Figure 10.9i-m). The minimum CV_c for simulation E at $t'=4$, for example, is 0.43 compared to 0.88 and 0.50 in simulations D and F, respectively. The simulation results are also better than those obtained from conditioning on a dense f sampling network (simulation C) and very similar to simulation B which utilizes dense network data for both f and h . Compared with simulation F, the additional saturated hydraulic conductivity information particularly helps to outline the extremely high and extremely low permeability areas, since the spatial variability of α is not very strong. But the comparison between simulations B and E also points to the fact that there is no gain in increasing the number of f measurements from 40 to 320, when so many head data are already available for conditioning. This is particularly

important, since saturated hydraulic conductivity measurements are much more difficult to implement in situ than head measurements. The results from this simulation indicate that a combination of in situ h and f data, with more h measurements than f measurements, may be the most economical approach to design a monitoring or sampling network.

10.6.4 Comparison to a Dry, Anisotropic Field Site of Equivalent Variability in y

Field site #15 has a much smaller textural variability than field site #28: The variance of f and a are only 1 and 0.01, respectively, instead of 2.25 and 0.04 at site #28 (see chapter 8). However, the increased dryness (mean head $H = -1000$ cm) leads to a strong increase in the unsaturated hydraulic conductivity unconditional variance, which is 1.5 compared to an unconditional variance of 1.8 at field site #28. The head variances are also similar: Field site #15 has an unconditional head variance of 4400 cm^2 vs. 4900 cm^2 at field site 28.

Figure 10.11a-m shows the actual plume at the field site and the conditional mean concentration for the highest density and lowest density data simulations discussed above (simulations A and G). Not surprisingly the tortuosity of the spatial path taken by the plume is very similar to field site #28. The agreement between the conditional mean plumes A and G and the actual plume is comparable to the findings at field site #28. This visual impression is confirmed by the concentration coefficients of variation, which are very similar to those found in the equivalent conditional simulations of field site #28 (Note that the output times for field site 15 are slightly different: 5, 10, 20, and 40 instead of 4, 8, 16, and 31). At similar variances of the unsaturated hydraulic conductivity and soil water tension and for the same mean α and correlation structure, the effects of conditioning on a set of f , a , and h data (simulation A) or on h data alone (simulation g) are similar, regardless of the mean soil water tension and the variability of the soil saturated hydraulic conductivity.

However, conditioning on f alone (conditional simulation C) neither improves the mean concentration prediction, nor reduces the minimum CV_c as much as in the wet soil #28 when

compared to the unconditional simulation (Figure 10.11n-u; Figure 10.12n-u). Relative to simulation G, the mean concentration in simulation C has a much larger longitudinal and transverse extension indicating significantly more uncertainty about the actual travel velocity and the travel path. The higher uncertainty is caused by the weak correlation between saturated and unsaturated hydraulic conductivity in dry soils, if f and a are uncorrelated. In contrast, the soil water tension data as in site #28 provide information not so much on the unsaturated hydraulic conductivity but on the gradient field and hence the approximate travel path of the plume. If data on a and f are not available or if data on only one of the two parameters is available, soil water tension data must therefore be considered an important source of information for more accurate transport predictions.

10.7 Unsaturated Hydraulic Conductivity Variance and the Effect of Conditioning Data

Increasing the unsaturated hydraulic conductivity and head variance leads to increased solute spreading and increased uncertainty in the prediction of the mean travel path of a solute plume. This theoretical result (Russo, 1993a) has been confirmed in the previous chapter. Here, a qualitative analysis is given of the effect of soil variability on the *conditional* concentration moment prediction. The concentration moments of two field sites are compared with those discussed in the previous section. All field sites have the same anisotropic correlation structure for f and a and the same mean F and A (see previous section, chapter 8). Field site #12 is a moderately heterogeneous soil with correlated f and a , an unsaturated hydraulic conductivity variance $\sigma_y^2 = 0.53$ and a head variance of 1900 cm^2 with a mean head $H = -150 \text{ cm}$. Field site 21 is the same as field site 12, but in a very dry condition ($H = -3000 \text{ cm}$), resulting in $\sigma_y^2 = 3.2$ and a head variance 7600 cm^2 . In terms of σ_y^2 , field site 12 ranks lowest and field site 21 highest among the sites tested. Note that $\rho_{af} = 1$, which means that f data perfectly predict a at the same location.

The unconditional plume for field site #12 is much less dispersed than those at other field sites (Figure 10.13a-d). Consequently the unconditional simulation itself is a fairly good description of the actual plume (at least compared to the conditional simulation results at field site #28), although ergodicity (zero concentration variance) is not achieved even for this moderately heterogeneous soil. The most obvious difference between the unconditional mean plume and the actual plume is the rate of displacement. At $t' = 20$, for example, the center of the actual plume has traveled significantly further than the center of the unconditional plume (compare Figure 10.13a-d to Figure 10.15a-d). The conditional simulation A (high data density, Figure 10.15e-h) captures the actual rate of displacement of the field plume as well as its particular shape. In contrast, the conditional simulation G (sparse head data only, Figure 10.15i-m) offers little improvement over the unconditional simulation. This is again reflected in the concentration coefficient of variation (Figure 10.14a-d and Figure 10.16e-h,i-m). At $t' = 20$,

the minimum coefficient of variation in the unconditional simulation is 0.79, which decreases by less than 25% to 0.63 in conditional simulation G, but by more than 75% to 0.19 in conditional simulation A. Similar observations can be made at other output times. This is a much better improvement in conditional simulation A compared to the results for field site #28, but much less of an improvement in conditional simulation G. The difference is caused by the changing information content of the data that are used for conditioning: In the heterogeneous soil #28, the flow path of the plume is rather tortuous and hence dictated primarily by the spatial distribution of the soil water tension. In this soil (#12), flow is almost parallel, a situation which has been conceptualized in many stochastic soil flow and transport models as the "parallel column model" (c.f. Destouni, 1993). Here, the uncertainty is reduced to predicting the rate of solute displacement while the travel path is well-known. The vertical velocity is then primarily controlled by the saturated hydraulic conductivity and α . In simulation G, these values are conditioned indirectly through the h data. In simulation A both are known at a high density and combined with the lower variability of the soil result in a lower CV_c than the comparable simulation for site #28.

For the same reasons, h data are even more important in simulating site #21 than in simulating site #28. The unconditional mean concentration prediction (Figure 10.13e-h) has a very high variability (Figure 10.14e-h). Nevertheless, the characteristic features of the actual plume are well captured even by conditional simulation G (compare Figure 10.17i-m with Figure 10.17a-d): The initial diagonal displacement ($t'=5$), the characteristic s-shape at $t'=10$, and the residual concentration not far below the source. Although conditional simulation A (Figure 10.17e-h) offers considerable improvements over simulation G, the uncertainty associated with an unconditional simulation is so large that the three tensiometer nests for simulation G alone offer almost as much improvement in prediction accuracy as all the data in simulation A together. Again, the visual impression from the mean plume maps are confirmed quantitatively by the CV_c . At $t'=5$, the unconditional minimum CV_c is 1.51 (Figure 10.14e), which improves by over 50% in simulation G to 0.73 (Figure 10.18i), and by almost 70% to 0.47

in conditional simulation A (Figure 10.18e). Similar results occur at $t'=10$. At $t'=20$, the minimum CV_c of the unconditional simulation reduce to 1.37, while those for the conditional simulations further increase relative to earlier time.

The two field sites #12 and #21 are particularly educational in that they illustrate how the information content of field data with respect to the conditional moments of the concentration changes with soil water content. Both sites represent the same type of soil, but under different mean soil water tension conditions. Depending on H , the same amount of on-site field data yields distinctly different improvements in the conditional plume prediction relative to an unconditional stochastic plume prediction. Conditioning on tensiometer measurements is particularly useful in soils with highly heterogeneous flow paths i.e., in soils with a high degree of textural heterogeneity, in very dry soils, or in soils with a steep average slope α of the $\log K(h)$ function. In soils with almost exclusively parallel vertical flow and therefore only mildly heterogeneous unsaturated hydraulic conductivity fields, the same tensiometer measurements have almost negligible effects. In contrast, saturated hydraulic conductivity data and data defining α are important data to reduce uncertainty in soils with more or less vertical parallel flow. But they loose their information content (measured in terms of minimum CV_c reduction relative to the unconditional minimum CV_c) in soils with very tortuous flow paths.

10.8 Anisotropy Ratio and the Effect of Conditioning Data

The isotropic soil site 3 is chosen for comparison with the conditional simulation results of the previous two sections. Relative to the horizontal correlation scale, the horizontal plume spreading is much larger in the isotropic soil than in the anisotropic soil #28, even though σ_y^2 at site 3 is only half of that at site #28 (see chapter 9). Note, that the vertical to horizontal scale ratio for all site 3 maps is 3:1 (half the aspect ratio $\lambda_{rh}/\lambda_{rz}$ of the anisotropic soils) instead of 1:1 (as in all maps of the anisotropic soil sites). Thus, the horizontal correlation scale for the site

#3 maps appears as half the absolute length of the correlation scale for the maps of the anisotropic soils.

Since flow path tortuosity (Figure 10.19a-d) contributes considerably to the uncertainty in the concentration prediction, in situ head measurements significantly reduce the prediction uncertainty (compare Figures 10.13i-m, 10.14i-m to Figures 10.19i-m, 10.20i-m). The minimum CV_c reduction at $t'=5$ is more than 50% from 1.07 to 0.51 with simulation G, and more than 85% from 1.07 to 0.14 with simulation A. Curiously, however, in this particular case the tensiometer data increase the error in the mean concentration prediction near the source at $t' > 10$: In simulation G, a secondary concentration peak appears and remains immediately underneath the source. Since the CV_c in this area is very high, the anomaly is probably caused by an outlier and should be neglected.

10.9 Conditional Simulation under Parameter Uncertainty

In all of the previous simulations it is assumed that the stochastic parameters describing the first and second moment of the input parameters f and a are known with certainty. In actuality, these parameters must be derived from a sample population of field and laboratory measurements. Generally, these sample populations are very small and the estimated mean and covariance are themselves RFVs (see chapters 3 and 8) i.e., their actual value is associated with a degree of uncertainty that is best measured in terms of the theoretical sampling error. Note that sampling errors are not the same as measurement errors. The effect of measurement errors, although important, has not been considered here. Parameter uncertainty in a conditional stochastic framework has been addressed by Smith and Schwartz (1981b) who implemented a specific type of conditional Monte Carlo analysis of saturated flow and transport to assess the additional uncertainty introduced by the sample estimation of the saturated hydraulic conductivity. Their objective was to assess the difference in the moments of the solute flux and concentration distribution introduced to the unconditional stochastic analysis by parameter

uncertainty. Their approach was too CPU expensive to address the combined effect of parameter uncertainty and measurement network design. The conceptual limits imposed on their study are the same for this study. Therefore, an alternative method is implemented to understand - at least qualitatively - the effect of parameter uncertainty in the conditional framework presented here.

Unlike the method by Smith and Schwartz (1981b) it is not the objective of this exercise to define quantitatively the increase in concentration variance or the change in mean concentration due to parameter uncertainty (which is computationally not feasible for the conditional case). Instead the problem is tackled from the following point of view: How different is a MCS result, if the sample moments, which are used as input to the MCS, are "far off" the actual ensemble moments? As in Smith and Schwartz (1981b), the exercise here will be restricted to the assessment of the effect of parameter uncertainty in the mean and the variance of the input RFVs f and a and in the mean H of the soil water tension. It is still assumed that f and a are known to be Gaussian distributed, that they are independent of each other ($\rho_{af} = 0$), and that the correlation functions of f and a are known with certainty. For the purpose of this exercise, "far off" sample moments are defined by the 95% confidence interval of the sample moment distribution. For simplicity and without loss of generality, the distribution of the sample moments is assumed to also be Gaussian (instead of e.g. the t -distribution most commonly used). Then the standard deviation ϵ_G of the sample mean G of g ($g=f,a,h$) is the square root of (8-3) and the standard deviation $\epsilon_{s,g}$ of the sample standard deviation std_g is the square root of (8-5) (chapter 8). The "far off" sample moments are the means $m_{g,95} = \mu_g \pm 2\epsilon_G$ and the standard deviation $\text{std}_{g,95} = \sigma_g \pm 2\epsilon_{s,g}$, where μ_g is the ensemble mean and σ_g is the ensemble standard deviation of g . Theoretically, the sample moments are "worse" i.e., more different from the ensemble moments, in only one of twenty sampling campaigns.

Data for the conditional simulation H of soils site #28 are chosen to demonstrate the method. Conditional simulation H consists of 40 measurements of f , 40 measurements of a ,

and of 320 measurements of h (Table 10.1). Since the f and a measurements are taken at least one correlation length apart, one can reasonably assume independence. Using (8-3) and (8-5), $\epsilon_F = 0.24$, $\epsilon_A = 0.032$, $\epsilon_{s,f} = 0.17$, $\epsilon_{s,a} = 0.022$. Hence, with a 95% probability the sample mean estimates of f and a are within the intervals $[-0.5, 0.5]$ and $[-4.7, -4.5]$, respectively (two standard deviations about the mean). Note that the uncertainty about the mean of α is so small that it can be neglected. With the same probability, the sample standard deviations of f and a must be within the intervals $[1.2, 1.8]$ and $[0.16, 0.24]$, respectively. The soil water tension data exhibit a much stronger correlation and are available at a much denser grid. It is conservatively assumed that the 320 correlated head data are equivalent to only 50 independent head data. From chapter 8, the unconditional head standard deviation is known (70 cm). Then $\epsilon_H \approx 10$ cm.

Four Monte Carlo simulations are implemented. The first one (simulation I) is implemented with overestimated parameters for the variances of f and a ($\sigma_f=1.8$, $\sigma_a=0.24$). It also strongly overestimates the mean of a . Although A can be determined very accurately under the above assumptions, it is the most difficult one to estimate in the field, since it is generally derived from fitting theoretical equations such as Gardner's (chapter 4) to measurements of unsaturated hydraulic conductivity or the soil water retention curve. In the first simulation, mean $\log \alpha$ is therefore set to -4.0 , simulating a type of measurement error of half an order of magnitude. Thus the unsaturated hydraulic conductivity is expected to be lower than in the actual field site due to the steeper average slope of the unsaturated conductivity function. Figure 10.21e-h shows the conditional mean concentration for this case in comparison to the simulation with the correct parameters (Figure 10.21a-d). Due to the reduced mean vertical flux (lower mean conductivity), the plume moves much slower, but along the same path as in the perfect parameter case. The higher variance in f and a does not contribute significantly to the spreading of the mean concentration. The higher variance in f is offset by the conditioning effect of the actual field data (which are exactly the same as in the perfect parameter case). The higher variance in a is still relatively low and doesn't affect the overall result significantly.

In the opposite case with underestimated variances ($\sigma_f^2=1.2$, $\sigma_a^2=0.16$) and with a very small mean $\log\alpha=-5.2$ (simulation J, Figure 10.21i-m), the plume moves much faster than the actual plume. The unconditional mean vertical velocity is 6.8 times faster than in the previous case I and almost twice as large as at the actual site. But the plume moves again along the same travel paths and with only a small decrease in plume spreading. Since the travel paths are essentially the same in simulations H, I, and J, the differences in plume spreading are best compared for travel times that correspond to similar travel-distances: Conditional simulation H at $t'=8$ against I at $t'=31$ (to be accurate it should be $t'=28$) against J at $t'=4$. The differences in the plume spreading caused by erroneous assumptions about the soil variability are obvious. Note that the different output times use different contouring levels.

Figure 10.22e-h shows the mean plume prediction from a simulation that again overestimates the variances of f and a , but has the correct A (mean of a) and an overestimate of F , the mean of f (conditional simulation K). Due to the conditioning, the plume moves only slightly faster than in the perfect parameter case (Figure 10.22a-d) and again with little extra spreading. Similar results are found, if the f and a parameters are estimated correctly, but the mean soil water tension is too wet (conditional simulation L), resulting in a higher average conductivity (Figure 10.2i-m).

These results show that conditioning not only reduces the uncertainty due to spatial heterogeneity, but also reduces the unknown errors that arise from a limited knowledge of the overall soil properties. Conditioning data tend to neutralize the parameter estimation error. With a high amount of tension data and some conductivity data, mean and variance estimation becomes a relatively minor source of uncertainty compared to the uncertainty arising from the spatial variability of the parameters. The effect of uncertainty in the correlation function needs to be explored in a future study.

10.10 Conditional Mean Displacement Variance and Conditional Moment of Inertia

The concentration moments are a function of both space and time. The Monte Carlo analysis of concentration moments is therefore associated with large amounts of data storage and data handling. It is also cumbersome - particularly in three dimensions - to visualize the results in an efficient manner. The spatial moments of the solute plume i.e., the center of mass and the moment of inertia (chapter 9), are a much more concise measure of the concentration distribution and of the uncertainty about the plume location. They reduce the multiple plots of two- or three-dimensional concentration fields to a single-valued function of time. Dagan (1984), Rubin (1991c), and Zhang and Neuman (1994c) use the second spatial moments of the mean solute plume to illustrate the effect of conditioning in mildly heterogeneous porous media. They show that the second moment of the mean concentration plume decreases towards the size of the actual plume as the number of conditioning points increases, since the spreading of each individual plume is smaller than that of the mean plume.

In chapter 9, three measures related to the position and size of the solute plume were introduced: the mean vertical and horizontal spreading of each plume around its center, $\langle M_{zz} \rangle$ and $\langle M_{xx} \rangle$; the variance of the plume center displacement in the vertical and horizontal direction, $\text{var}(M_z)$ and $\text{var}(M_x)$; and the vertical and horizontal spreading of the mean concentration plume, X_{zz} and X_{xx} , which is computed as the sum of the two former measures (eqn. 9-3, see also Fisher et al., 1979). Figure 10.23 shows these moments for the actual plume at field site #28 as a function of dimensionless time $t' = tV_z/\lambda_{fz}$ (solid curve) with those from the corresponding Monte Carlo simulations A (dense grid of f, a, h data; long dashed curve), H (dense grid of h data, sparse grid of f data; short dashed curve), and G (sparse grid of h data; dash-dotted curve), and for the unconditional simulation of this site (dotted curve). The dimensionless moments (indicated by an apostrophe ') are normalized by dividing the actual moments with the product of the square of the vertical correlation scale of f and the variance of f, $(\lambda_{fz}\sigma_f^2)$ (see also chapter 9).

The average plume spreading $\langle M_{zz} \rangle$ and $\langle M_{xx} \rangle$ are determined primarily by the variability and correlation scale of the soil texture (Russo, 1993a). The moments of individual

plumes may vary in many different ways, particularly since the spatial distributions of the solute plumes are generally non-Gaussian for the hypothetical field sites studied here. At site #28, the horizontal spreading of the actual plume has a small step increase at early time t' (0.5 - 1.5) which stems from the diagonal plume movement, and a very large increase after $t'=8$ to almost 8 times the value for the vertical spreading at $t'=15$ (Figure 10.23a,b) indicating the horizontal plume movement observed in the map of the actual concentration distribution (Figure 10.5a-b). The vertical spreading indicates a strong expansion-contraction cycle between $t'=8$ and $t'=16$. Since the expansion is seen for both the horizontal and the vertical moment, there is likely an accelerated diagonal movement of parts of the plume after $t'=8$ (no concentration data are available for the time between $t'=8$ and $t'=16$ to exactly explain the anomaly in M_{zz}). Curiously, $\langle M_{zz} \rangle$ from simulation A shows exactly the opposite anomaly at the same time (Figure 10.23a): a strong contraction followed by some expansion. While the actual plume has a vertical spreading comparable to the unconditional plume, all conditional plumes significantly underpredict the vertical spreading of the solute. In contrast, the conditional data improve the prediction of horizontal spreading that occurs after $t'=8$ (Figure 10.23b). At earlier times, the horizontal spreading predicted is very similar for all three types of conditioning and for the unconditional simulation.

The variance $\text{var}(M_z)$ of the vertical displacement of the plume center decreases significantly as the number of data increases (Figure 10.23c). The unconditional simulation has a very large center displacement variance, much larger than the average spreading $\langle M_{zz} \rangle$ of the individual plumes. For $t' < 7$ the horizontal center displacement variance $\text{var}(M_x)$ also decreases (Figure 10.23d). However, due to the strong horizontal spreading of the mean plume after $t'=7$, the conditional horizontal plume spreading increases with the number of conditioning points, and so does $\text{var}(M_x)$. At $t' > 12$ $\text{var}(M_x)$ is larger in conditional simulation A than in the unconditional simulation.

From those results it follows that the vertical spreading X_{zz} of the mean plume (Figure 10.23e), which is the sum of the average plume spreading and the plume center displacement

variance, shows - for most parts - the expected decrease as the number of conditioning points increases. The decrease in X_{zz} stems mainly from the decrease in the plume center displacement variance $\text{var}(M_z)$. But due to the difference between $\langle M_{zz} \rangle$ and the actual vertical spreading of the field plume, X_{zz} becomes smaller than that of the actual plume for conditional simulation A at $t' > 8$. For the horizontal X_{xx} , the results are ambiguous at early time ($t' < 8$). Only at later times, X_{xx} comes closer to the actual horizontal plume spreading as the number of conditioning points increases. Due to the strong horizontal spreading of the actual plume, the horizontal spreading of the mean plume actually increases with the number of conditional data available ($t' > 8$). Although not shown, it is found that the spatial moments obtained from the conditional mean concentration plume are in good agreement with the sum $X_{ii} = \langle M_{ii} \rangle + \text{var}(M_i)$ (Figure 10.23e,f).

Overall the results indicate that the plume moments do not very accurately reflect the prediction improvement as demonstrated by the conditional concentration moments in previous sections. The discrepancy between the quality of the moment prediction and the amount of conditional data invested is caused primarily by the non-Gaussian shape of the plume. Multiple peaks, meandering, parting of the plume and many other particular features of solute plumes in moderately to highly variable soils cannot be characterized by the spatial moments of the plume. The larger the plume deviates from the Gaussian shape the less information is contained in the first and second moment, since it becomes very sensitive to the particular distribution of the plume i.e., to higher order spatial plume moments. This explains why the increase in conditioning points does not necessarily improve the accuracy of the moment prediction.

In soils with a less variable flow field, the spatial concentration distribution is much closer to the Gaussian form, and hence the above three measures of plume spreading are increasingly helpful in describing the actual contribution of conditioning i.e., in describing the actual plume movement. This is demonstrated in the spatial analysis of the concentration distribution at site #12 (Figure 10.24), which has an only moderate variability in the unsaturated hydraulic conductivity ($\sigma_y^2 = 0.5$). But even here it is seen that at different times, the accuracy

of the predictions do not necessarily reflect the degree of conditioning (e.g. the prediction of X_{xx} in Figure 10.24f at time $t=5$). This is consistent with findings of Zhang and Neuman (1994c).

In soils with flow fields that are even more variable than at site #28, the value of using spatial moments of the concentration distribution to assess the plume movement via conditional simulation becomes questionable due to the highly irregular shape of the actual solute plume, as demonstrated for site #21 (Figure 10.25), where $\sigma_y^2 = 3.2$.

10.11 Conditional Local Solute Travel Time

From a regulatory point of view, the spatial distribution of solute concentration is in many cases not as much of interest as the arrival time distribution of the solute at some compliance surface. For transport through the unsaturated zone, the compliance surface is mostly the aquifer water table, since the aquifer rather than the soil itself is the resource that is protected under many environmental regulations. In this section the focus is the solute flux breakthrough at a certain depth below surface as a function of time and location on the compliance surface. In the two-dimensional simulations the compliance surface reduces to a horizontal compliance axis (CL). The compliance axis in all simulations is located across the center of the two-dimensional, vertical simulation domain ($z/\lambda_{fz} = 11.6$). In chapter 9, two measures to characterize solute breakthrough at each horizontal location x of the CL were introduced: The time t_p of arrival of the peak solute flux and the time t_a of first exceedance of a certain compliance solute flux.

Figure 10.26 shows the peak time of solute flux at the field site and the mean and variance of the peak time obtained from the conditional and unconditional MCSs. The vertical time axis is normalized by the mean unconditional travel time $t = V_z/z$ where z is the distance from the source to the compliance surface directly underneath the source. At the field site, the main features of solute breakthrough are two areas of relatively fast breakthrough i.e., early peak

solute flux time, at $x \approx -200$ cm and at $400 \text{ cm} < x < 1000$ cm (horizontal distances are measured with respect to the point directly under the solute source). This corresponds to the two advancing fronts of the solute plume seen in Figure 10.5 at $t'=8$ and $t'=16$. Prediction of these two distinct and quickly advancing fronts would seem critical for regulatory purposes.

None of the conditional simulation predicts an advancement of peak travel time as fast as it actually occurs at the field site. The best predictions are by conditional simulation A, which predicts peak concentrations of the two advancing fronts to occur approximately 10% later. Peak times at other locations of the compliance axis are estimated conservatively by simulation A. Conditional simulation H (dense head data, sparse f data) predicts peak arrival generally later than A and shows much less distinction for the two advancing fronts. Conditional simulation G (sparse head data) makes a better distinction between the two advancing fronts and gives a fairly accurate prediction of the location of breakthrough of these fronts, but at later time than the other two conditional simulations. The variance of the peak time increases with less conditional data due to the increase in concentration variance seen previously. Between simulation A and the unconditional simulation, the difference in variance is approximately a factor 2 near the center of the compliance axis.

A more detailed picture of solute breakthrough is given by the arrival time t_a or first exceedance of the compliance solute flux s/s_0 . In Figure 10.27 these times are mapped for various compliance solute flux levels. The lowest s/s_0 levels are first exceeded i.e., arrive first. The white area outside the plume indicates that the compliance level was never exceeded (see chapter 9). The advancement of the two split fronts of the solute plume is again seen in the map for the field site: The arrival time of all but the highest s/s_0 levels is much earlier at $-200 \text{ cm} < x < 0$ cm and $400 \text{ cm} < 600$ cm than at $x = 200$ cm. Conditional simulation A is the only simulation that reproduces a similar split pattern. But even for this high data density, the arrival times of any particular s/s_0 level is underestimated by approximately 20%. Also, at most lower s/s_0 levels, the conditional simulation predicts an earlier breakthrough of the left front when compared to the right front, while in reality the opposite is true. The distinction between two

advancing fronts is entirely lost in the other conditional simulations, which show the arrival time of any s/s_0 level to be the shortest to the left of the center of the compliance axis, and then longer the further away from the center. The unconditional simulation has the largest error in predicting the arrival time of s/s_0 levels due to the fact that the actual plume moves overall faster than the average velocity. Again the variance increases for less conditioning data.

Much of the insensitivity of the travel time moments may be due to the particular location of the compliance surface at field site #28: It coincides with the region of strong horizontal movement, which is well predicted with simulation type A, but less explicit and at slightly different locations with simulation types G and H. At field site #21, the compliance surface coincides with strong vertical solute movement on the left side of the center of the compliance line. In this case, a much higher sensitivity of the solute arrival moments e.g., the mean arrival time t_a of s/s_0 is found, although the soil flux is much more variable.

In summary, the mapping of the two time parameters to characterize the solute flux breakthrough as a function of location on the compliance axis may not nearly be as sensitive to conditioning as the local concentration moments are. If strong horizontal flux occurs at or near the compliance surface, the results of the travel time analysis may be very vague. Ideally the compliance surface should be located such that solute flux is known to be predominantly normal to the compliance surface. Otherwise the effect of conditioning may only be weak in the arrival time of certain solute flux compliance levels or the peak concentration. While conditioning works well to predict the general spatial distribution of the concentration, this shows that it may be less effective for predicting solute flux arrival times or breakthrough curves, since the amount of data necessary to predict these accurately can under circumstances be enormous. And depending on the flux pattern across the compliance surface, the improvement in the mean time predictions due to sparse conditioning may be relatively small or very significant.

10.12 Conditional Integrated Solute Breakthrough Curves

Given that the local breakthrough curves are not necessarily sensitive to conditioning, the question arises whether the integrated solute breakthrough curve (BTC) is similarly insensitive. The solute breakthrough curve represents total mass flux across the compliance surface at any given time. The data are compiled not only at the compliance surface in the center of the domain, but also at every other quartile of the domain depth ($1/4$ depth or $5.4\lambda_{fz}$ travel distance, $3/4$ depth or $17.8\lambda_{fz}$ travel distance, and bottom boundary or $23.8\lambda_{fz}$ travel distance). The actual breakthrough curve at field site #28 is plotted against the mean breakthrough curves of the conditional simulations A, H, G, and the breakthrough curve of the unconditional simulation (Figure 10.28a-d). The mean breakthrough curves clearly show the effect of conditioning, particularly at the two top compliance axes (Figure 10.28a,b). At those two levels, the actual plume BTC has an almost Gaussian shape, and increasing the data density leads to less dispersed mean BTCs in the conditional simulation with the result that the peak concentration is much better predicted if more data are used for conditioning. At the $3/4$ depth and the bottom of the simulation domain, the BTC of the actual plume significantly deviates from the Gaussian shape due to the very heterogeneous shape of the solute plume. The conditional breakthrough curves give results of variable accuracy, which cannot be related to the number of conditioning points (Figure 10.28c,d). This is not surprisingly similar to the insensitivity of the spatial plume moments at later time, since the temporal plume moments are related to the spatial plume moments and suffer from similar disadvantages. Note that the variance of the breakthrough curves nevertheless decrease as the number of conditioning points increases (Figure 10.28e-h). Also, as the number of conditioning points increases, the time span of high standard deviations of the BTC decreases. This indicates that while the prediction about the average solute flux at a given time does not necessarily improve with the amount of data used for conditioning, the uncertainty about the prediction decreases almost always. Similar observations are made for field site #21 (Figure 10.30). In contrast, the effect of conditioning is very obvious in the mean BTCs for the moderately heterogeneous site #12 (Figure 10.29) as found by other researchers in the past (Smith and Schwartz, 1981b; Gutjahr

et al., 1994; Zhang and Neuman, 1994c).

10.13 A Deterministic Geostatistical Inverse Approach in Comparison

It may be argued that with a reasonable amount of information - such as that in sampling scheme A - the stochastic technique is superfluous and a reasonable prediction can be made through deterministic approaches alone. This argument may be appropriate if one is interested in a prediction of solute transport without an estimation of the associated uncertainty. With deterministic modeling uncertainty cannot be quantified. If a model is only needed to give an approximate prediction of the solute plume movement, methods other than the stochastic technique should be considered.

By generating a zero variance unconditional realization of f and a and by assuming that all other pertinent statistical moments are known ($F, \Gamma, \text{cov}_f, \text{cov}_a, \text{cov}_h, \text{ccov}_{fh}, \text{ccov}_{ah}$, where cov is the covariance and ccov is the cross-covariance), the conditional simulation algorithm of section 10.3 (Figure 10.1) reduces to the geostatistical inverse modeling approach described similarly for saturated groundwater flow by Neuman and Yakowitz (1979), Kitanidis and Vomvoris (1983): The measurements of f , a , and h are used to estimate the remaining unknown f and a data in the simulation grid through the linear, unbiased, cokriging estimator introduced earlier. The steady-state head solution and solute transport is then computed for the cokriged f and a field. The geostatistical inverse modeling technique is only one of several other indirect inverse modeling techniques (Schweppe, 1973; Neuman and Yakowitz, 1979; Carrera and Neuman, 1986; Peck et al., 1988).

The geostatistical inverse approach (cokriging) is applied to field site #28 given the same data as used for the conditional simulation A of that site. Since the measured data density is relatively exhaustive, the f and a parameter estimation is associated with only small errors (see section 10.5). Like any random realization of the conditional simulation A (section 10.6), the concentration distribution predicted from the geostatistical inverse model is a very good

approximation of the overall plume movement (compare Figures 10.31e-h, 10.31i-m). The solute plume predicted by the inverse model is less dispersed than the conditional mean solute plume since it is not an average concentration. It is also less dispersed than the actual plume, since the underlying parameter fields for f , a , and h are subject to minimal perturbation given the conditional data. A less tortuous travel path and a mass balance error in the transport simulation of up to +13% lead to higher predicted peak concentrations in the inverse model than observed at the field site.

10.14 Summary and Conclusions

A number of studies presented in the literature analyze the uncertainty associated with predicting transport of solutes in heterogeneous unsaturated soils. Both analytical and numerical models have been developed to address this issue. But without exception in situ measurement data have not been incorporated in the stochastic analysis except to determine the unconditional parameters of the statistical models that describe soil heterogeneity. Neither has any study to date taken advantage of the available in situ information for soil water tension or other data indirectly related to the soil textural heterogeneity (indirect data) to reduce the prediction uncertainty associated with the unconditional stochastic approach to modeling solute transport. Most recently, Neuman and Loeven (1994) have introduced a new approach that allows one to derive the conditional moments of the soil water tension, soil water content, and soil water flux. But the approach has not yet been applied to also derive conditional concentration moments. With the conditional approach developed in this work a model is provided to compute the spatial distribution of estimation errors associated with solute transport predictions subject to in situ data measurements of either direct data (K_s and/or α) or indirect data (soil water tension) or a combination of both.

The difficulties that have prevented conditional simulation of nonlinear unsaturated flow and transport in the past are overcome by introducing an approach called conditional ASIGNing. The method is based on the ASIGNing technique (Harter and Yeh, 1993; chapter 7), which generates not only unconditional random fields of f and a , but also an approximate linearized solution h_L to accelerate the CPU-time for the finite element solution h of Richards equation. In this chapter, the ASIGNing method is combined with the geostatistical approach in general (Matheron, 1971; Journel, 1974; Delhomme, 1979) and cokriging in particular (Myers, 1982; Carr and Myers, 1985). The key to the efficiency of the new conditional simulation algorithm is the use of the first order perturbation approach described in chapters 4 and 7 to compute an unconditional head random field h_L that is approximately consistent with

the unconditional random fields f and a . h_L is an intermediate result in the conditioning algorithm (Figure 10.1). The h_L field is necessary to condition the corresponding f and a realizations (eqn. 10-1) on soil water tension data and to construct a geostatistical estimate of the conditional soil water tension realization h_L^c as initial estimate to again accelerate the CPU-time for the finite element solution h^c of Richards equation. Although the linearized, first order head solution is only an approximation of h , the results indicate that its application in the conditioning process is justified since the conditional moments of f^c and a^c are also linear estimates (cokriged estimates). Like other (linear) conditional algorithms conditional ASIGNing does not lead to solutions h^c of the flow equation (given the conditional f^c and a^c) that perfectly honor the measured head data. But the conditional variance of the head at the measurement points are at the most 5%-10% of the unconditional head variance. Forcing the correct heads at the points of measurement by imposing internal boundary nodes would lead to ill-conditioned gradient and velocity fields.

In this chapter, conditional ASIGNing has successfully been applied to the Monte Carlo simulation of conditional stochastic transport in a number of hypothetical soil types with varying degrees of textural variability, anisotropy, and moisture content. Conditional moments were analyzed not only of the spatial distribution of the concentration mean and variance at time t , but also of the overall plume spreading and of the arrival time of the solute as measured by a number of different parameters such as the solute breakthrough curve at a hypothetical compliance surface. In summary of the Monte Carlo simulation (MCS) results I have the following conclusions:

From a numerical-technical point of view, conditioning even on a few indirect data is an important tool to eliminate some of the most unlikely possible plume travel dynamics in the Monte Carlo sampling procedure. Most importantly, conditioning on either direct or indirect data at the source removes the outlier problem in unconditional Monte Carlo simulations of highly heterogeneous soils. Outliers of concentration levels occur near the source due to sometimes extremely low conductivities at or near the solute source. Outliers of permeability

values far from the source are generally no problem, since the travel path of the plume avoids such stagnant areas. Conditioning on near source information greatly reduces the risk of outliers that may bias the sample statistics. In addition, less realizations are necessary per MCS to achieve the same level of sample moment accuracy due to the decrease in the ensemble variability. Conditional simulations are therefore computationally less expensive than unconditional simulations, although additional CPU time is needed for the conditioning of each realization. In this analysis, 150 conditional realizations gave very accurate sample estimates, if either h or f or both are measured in a dense grid. If only sparse sampling data are available, the number of realizations was increased to 300, the same number as in the unconditional simulations (chapter 9).

The most important difference between unsaturated and saturated conditioning is the physical nature of unsaturated hydraulic conductivity, which is not an independent property of the soil, but determined by the soil water content and soil water tension and by a number of textural properties of the soil. The unsaturated conductivity cannot be conditioned directly unless it is measured in situ. It is here assumed to be dependent on two parameters besides the soil water tension h , which are often measured at different locations. If both parameters $f = \log K_s$ and $a = \log \alpha$ of the $K(h)$ function (4-8) are known at one location, and if the soil water tension is measured nearby ($0.25 \lambda_f$), then the conditional variance of the unsaturated hydraulic conductivity reduces to almost negligible values. But if a is not measured, the uncertainty about K at the f measurement points may be significantly larger depending on the mean and variance of a and depending also on the mean soil water tension.

In this study it is assumed that unsaturated steady-state flow occurs under unit gradient conditions with the major anisotropy axis transverse to the mean flow direction. The steep mean gradient and the transverse anisotropy have a stabilizing effect on the flow pattern. Hence, for anisotropic soils with moderately heterogeneous flow ($\sigma_y^2 < 1$) the assumption of one-dimensional vertical flow (parallel column model) is mostly justified. In such soils the unconditional simulation approach yields results that are a relatively good approximation of the

actual plume even if the plume is from a small source (provided the source location is known). The uncertainty about the solute plume movement is reduced primarily to uncertainty about the vertical travel velocity. Since moderately variable flow is mostly restricted to wet soils, the travel velocity is strongly correlated with the saturated hydraulic conductivity. Conditioning on f will therefore reduce uncertainty more than conditioning on a or h . In moderately heterogeneous, anisotropic soils conditioning on head data alone will not significantly improve the unconditional prediction of solute transport.

In soils with strongly heterogeneous, anisotropic flow fields ($\sigma_y^2 \geq 1$) i.e., in very heterogeneous soils or in dry soils, the travel path significantly deviates from the vertical direction and is characterized by a significant amount of horizontal displacement and tortuosity. In isotropic soils, similar observations are made even for moderately heterogeneous flow fields ($\sigma_y^2 > 0.5$). Solute plumes of small initial lateral dimensions (0.3λ) are found to have multiple peaks, multiple fronts, and are generally of a very erratic shape. With the unconditional stochastic transport approach, the uncertainty about the plume movement in both the horizontal and vertical direction leads to very large mean concentration plumes (see chapter 9). While an unreasonable amount of data would have to be retrieved from the soil to accurately predict the solute movement in such highly heterogeneous soils, conditioning on either a few indirect or a few direct data will significantly improve the prediction of the mean concentration plume and reduce the prediction uncertainty as measured by the spatial distribution of concentration mean and concentration coefficient of variation.

The information content (i.e., the ability to reduce uncertainty in a conditional simulation) of f alone decreases not only with increasing heterogeneity but also as the soil dries out, particularly if the mean and variability of a is large and if a is not strongly correlated with f . Then the information content of the head becomes important for two reasons: The spatial distribution of the head carries information about the head gradient field in the soil and therefore about the travel path of a solute plume. Secondly, soil water tension data help to better estimate the unsaturated hydraulic conductivity, which controls both travel velocity and travel

path. Hence, in soils with highly variable flow fields, conditioning with head data significantly reduces transport prediction uncertainty despite the fact that the conditioning technique itself relies on a strong linearization of the physical process, which becomes less valid as the flux variability increases (Kitanidis and Vomvoris, 1983). In very heterogeneous soils it appears that soil water tension data reduce the solute movement prediction uncertainty (as measured by the minimum CV_c) more than the same amount of saturated hydraulic conductivity data. The positive effect of head conditioning in very heterogeneous porous media was also observed for saturated groundwater flow (Gutjahr et al., 1994).

This is a very encouraging result since the cost of equipment and labor associated with soil water tension data is generally lower than that associated with obtaining saturated hydraulic conductivity data. It is more likely to find in situ head data than to find in situ information about the saturated hydraulic conductivity. The simulations suggest that a combined network with a relatively high sampling/monitoring rate for soil water tension and a relatively sparse sampling/monitoring rate for saturated hydraulic conductivity leads to a significant decrease in prediction uncertainty about the concentration. From the examples in this study it appears that the combined uncertainty reduction due to a combination of h and f data is beyond the additive impact of head data by themselves and f data by themselves.

The minimum concentration coefficient of variation CV_c is used as a summary measure of the conditioning effect on uncertainty reduction. The changes in the minimum CV_c appear to be approximately consistent with the visual changes on the concentration maps. Future research must address the question of converting the large amount of information about $c(\mathbf{x},t)$ into other representative parameters. In this context I would also like to point out that an exact definition and quantification of the term "uncertainty" is very difficult, an issue that should be addressed in future research.

The spatial moments of the mean concentration plume are another example of a compact measure to study the impact of conditioning on solute transport in heterogeneous porous media (Dagan, 1982, 1984; Rubin, 1991a; Zhang and Neuman, 1994c). The analysis

of the spatial moments of solute plumes in moderately variable flow fields indicated that conditional data primarily reduce the uncertainty about the center of the plume. The uncertainty about the mean spreading of the solute plume is fairly small and the spatial distribution of the actual solute concentration is not much unlike a Gaussian plume. In mildly to moderately heterogeneous, anisotropic soils, conditional spatial moments of the mean solute plume and conditional mean breakthrough curves therefore accurately reflect the effect of conditioning on the solute transport prediction. In strongly heterogeneous flow fields, however, and in isotropic soils with moderate heterogeneities, the significance of the spatial moments of the conditional mean plume is strongly diminished due to the erratic (non-Gaussian) shape both of the individual solute plume realizations and of the mean solute plume. It becomes therefore difficult to quantitatively assess the effects of conditioning by analyzing the spatial solute plume moments alone.

Similarly, the effect of conditioning on the arrival time or breakthrough of a solute at a compliance surface or compliance point some distance away from the solute source is often felt much less direct than in the spatial pattern of the conditional mean concentration, if the flow patterns are strongly heterogeneous. While a significant decrease is found in the variance of these measures, the actual shape of the conditional mean breakthrough curve (locally and integrated) may or may not be similar to the actual breakthrough curve, even with a dense sampling grid for f , a , and h . If the solute flux across the compliance surface is not predominantly normal to the surface, the mean solute flux and arrival time become very sensitive to small changes in the conditional mean flow field, and the effect of conditioning becomes ambiguous. It therefore appears that the conditional simulation of breakthrough curves in highly heterogeneous porous media must be accompanied by the concentration mean and variance maps to help explain what is seen in the breakthrough curve and to evaluate the effect of conditioning on the breakthrough. Breakthrough curves of solute transport from small sources in strongly heterogeneous soils should therefore not be expected to be very accurate, even when the amount of conditional data is large.

Besides spatial heterogeneity, the estimation of the statistical parameters f , a , and h and the associated estimation error introduces additional uncertainty into the solute transport prediction. The impact of parameter uncertainty is found to diminish with the amount of data available in situ, because of reduced sampling error and more constraints on the stochastic simulation. Parameter uncertainty about the variance of f and a is mainly responsible for increasing the overall spreading of the mean plume. In contrast, errors in the sample mean of f , a , and h result in an error about the mean flux prediction and consequently the error will be not in the travel path of the plume, but merely in the travel time. Overall, however, it appears that the uncertainty of the solute transport prediction arising from soil heterogeneity is much more significant than the uncertainty arising from parameter uncertainty.

In this study, several simplifications are made not only to be able to compare numerical with analytical solutions (see chapters 8 and 9), but also to be able to establish some fundamental relationships between monitoring/sampling network and the heterogeneity of the soil. Future work must address the effect of variable moisture content and transient flow conditions. Measurement errors, parameter estimation errors, particularly about the correlation structure, and error in assuming the wrong models describing the $K(h)$ and $h(\theta)$ relationship further increase prediction uncertainty and should be addressed in future research. Thus, it may be expected that the effects of conditioning become smaller. The geostatistical conditional simulation model must be recognized not to be a perfect measure of uncertainty itself, because it is based on a linearization of a nonlinear physical problem (see section 10.2 and 10.3) and on several important assumptions about the concept of spatial variability in soils (see chapter 2). A field validation of the conditional stochastic approach suggested here is therefore necessary. However, the model is based on statistical concepts, and it seems at this time impossible to implement such a field validation rigorously, because many field experiments (samples) would be needed to judge about the goodness of the stochastic model.

From a practical point of view, the results are both encouraging and disappointing. They are encouraging in that they show that with less computational effort than in the classic

unconditional approach, and with data that are relatively simple to obtain in situ (soil water tension), the uncertainty about the predicted plume movement in space can be drastically reduced, particularly for applications to highly heterogeneous soils. It is encouraging also in that the conditional mean concentration predictions are pinpointing to areas where the plume displacement significantly differs from the typical downward movement. This helps to identify locations from which additional data may be taken. If the unsaturated flow field is very heterogeneous conditioning on a few indirect or direct data will greatly improve the stochastic predictions associated with unconditional simulation and with macrodispersion analysis (see chapter 9). But the results are discouraging in that the simulations have shown how difficult it is to describe the (conditionally simulated) plume movement in highly heterogeneous soils by simple measures such as the spatial moments of the mean solute plume or the minimum concentration coefficient of variation. The study has also underlined the difficulty of predicting solute breakthrough at some compliance depth even when conditioning on a high density of direct and indirect data. Further research needs to be done to address these disadvantages. It appears from the results presented in this chapter that even with an enormous amount of field sampling it will be very difficult to predict every detailed aspect of solute transport in moderately to highly heterogeneous soils, particularly the prediction of very low levels of solute concentration either in front of an advancing contamination plume or as residual.

Table 10-1

Classification of the conditional simulation types. Conditional simulations A through H are based on different sampling networks for the parameters f, a, and h. Simulations I through L are applied to field site 28 only to assess the effect of erroneous statistical input parameters.

conditional simulation type	sampling density f	sampling density a	sampling density h
A	dense	dense	dense
B	dense	-	dense
C	dense	-	-
D	sparse	sparse	-
E	sparse	sparse	dense
F	-	-	dense
G	-	-	sparse
H	sparse	-	dense
I (as H)	but: ($\sigma_f^2 = 1.8$, $\sigma_a^2 = 0.24$, $\langle \log \alpha \rangle = 4.0$)		
J (as H)	but: ($\sigma_f^2 = 1.2$, $\sigma_a^2 = 0.16$, $\langle \log \alpha \rangle = 5.2$)		
K (as H)	but: ($\sigma_f^2 = 1.8$, $\sigma_a^2 = 0.24$, F = 0.5		
L (as H)	but: (H = -140 cm)		

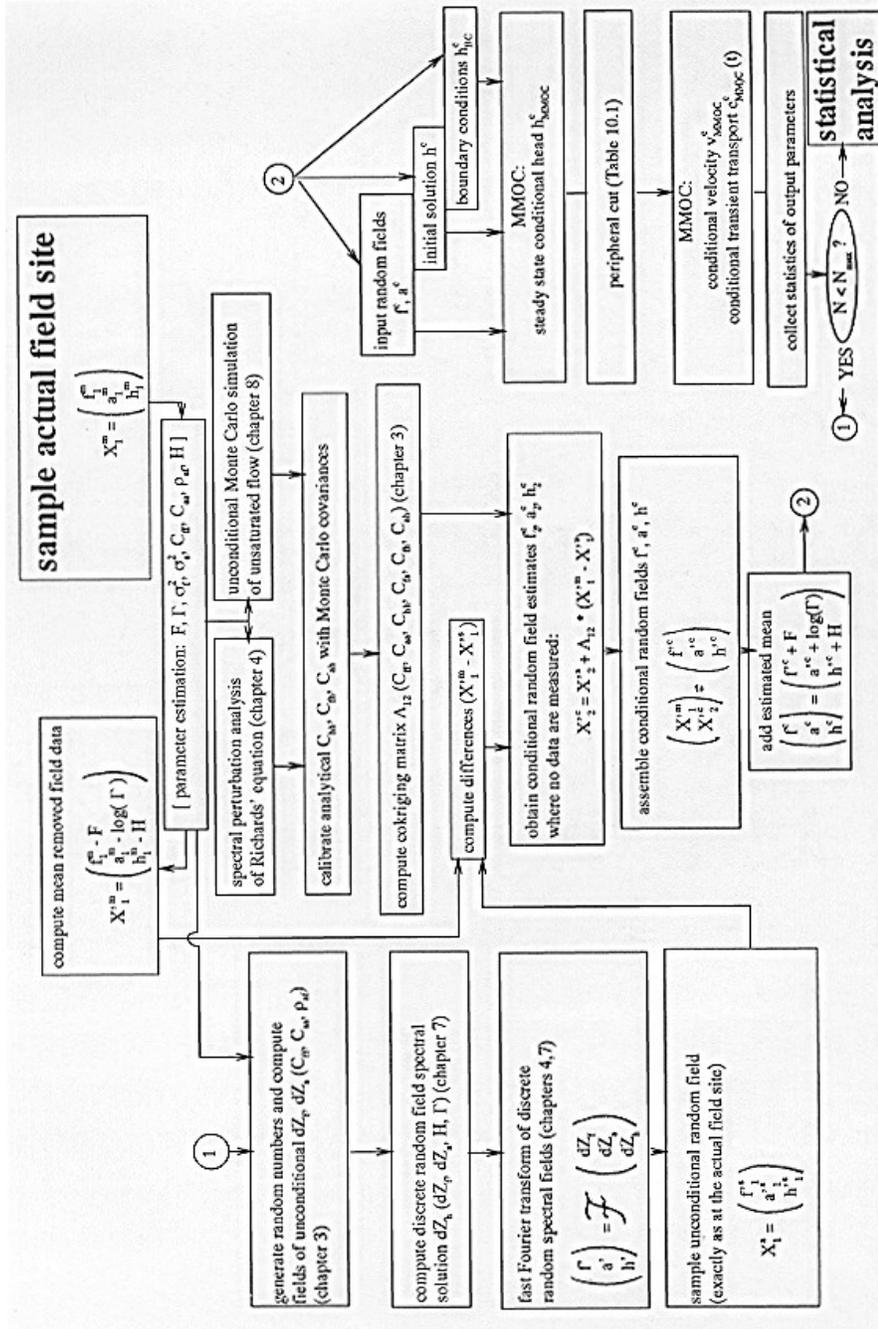


Figure 10.1

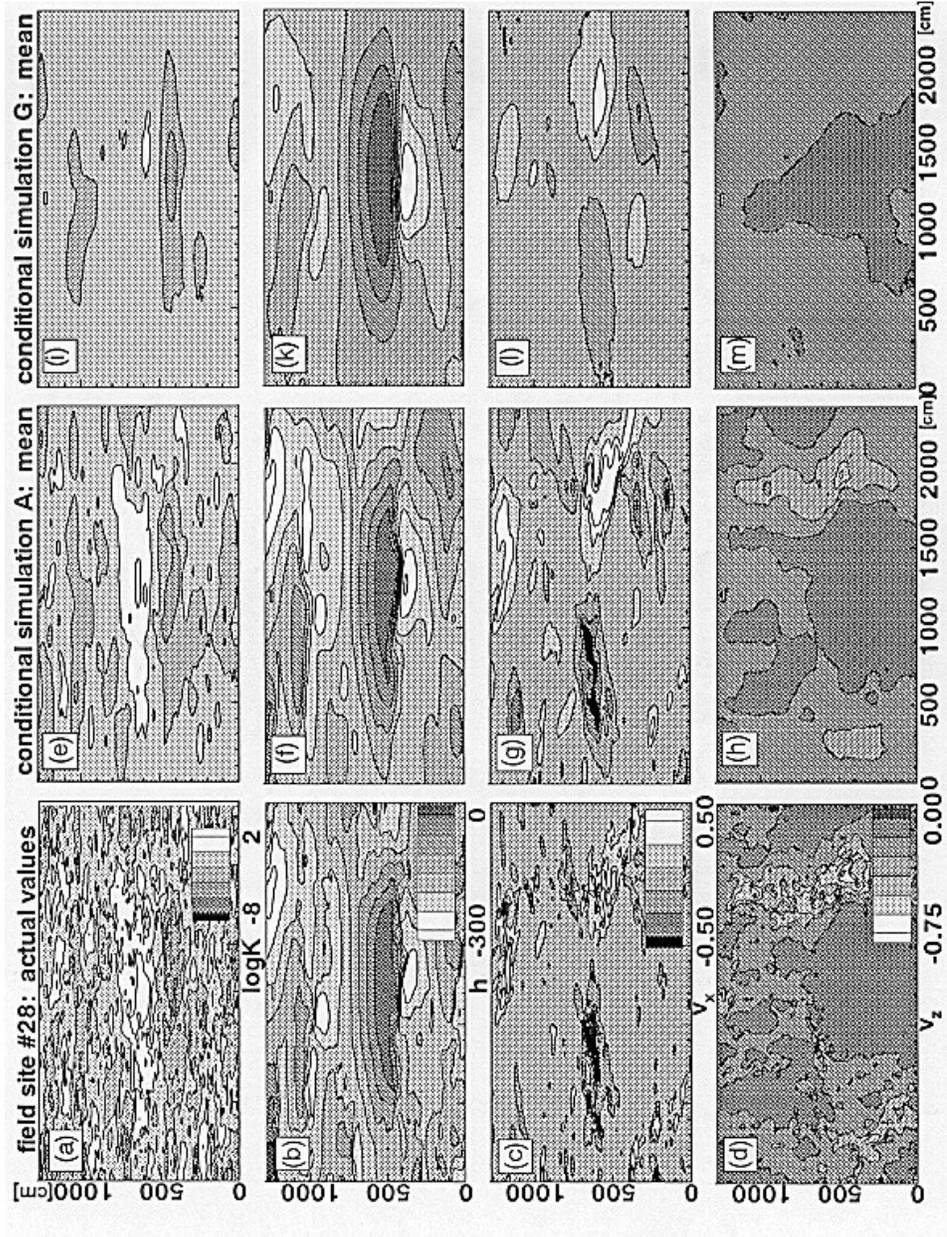


Figure 10.2
for the parameters $\log K$ (a,e,i), h (b,f,k), v
identical throughout each row.

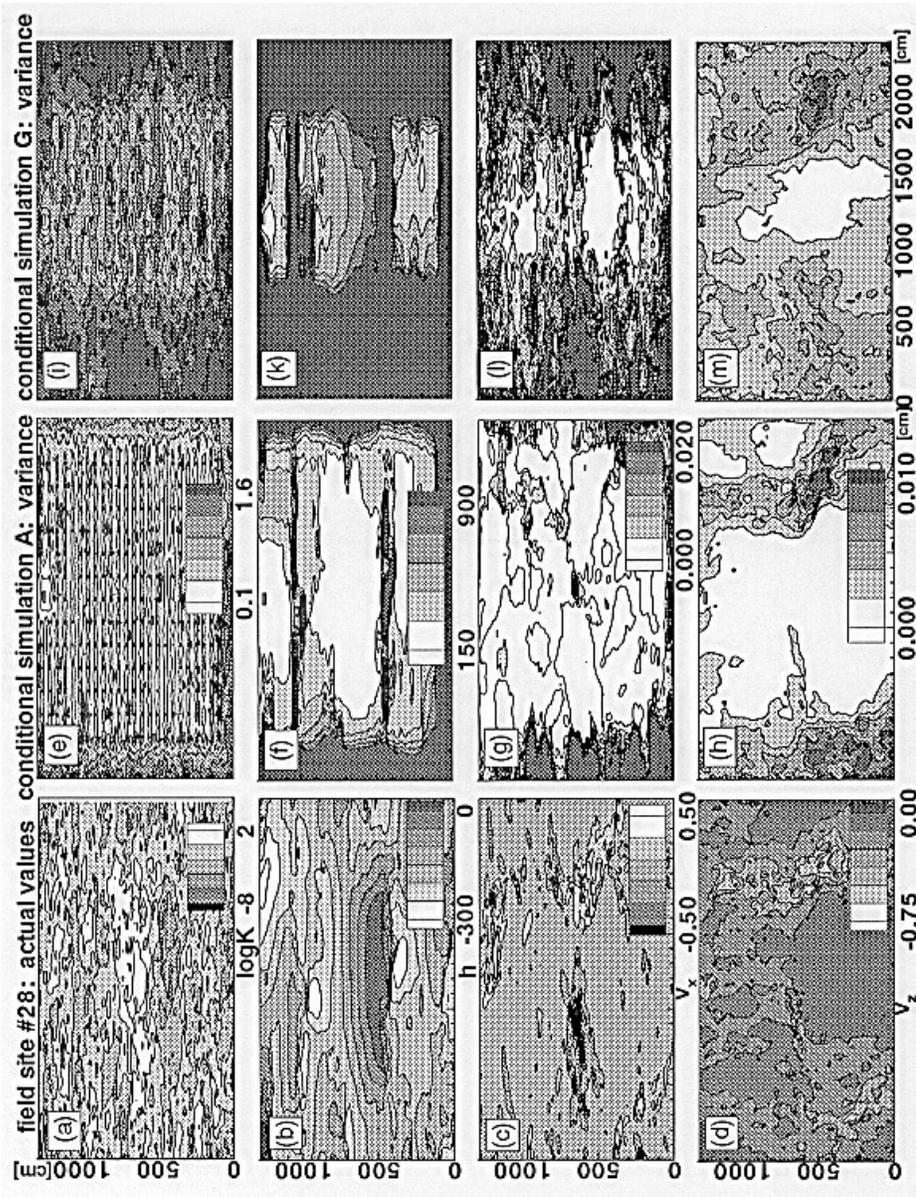


Figure 10.3

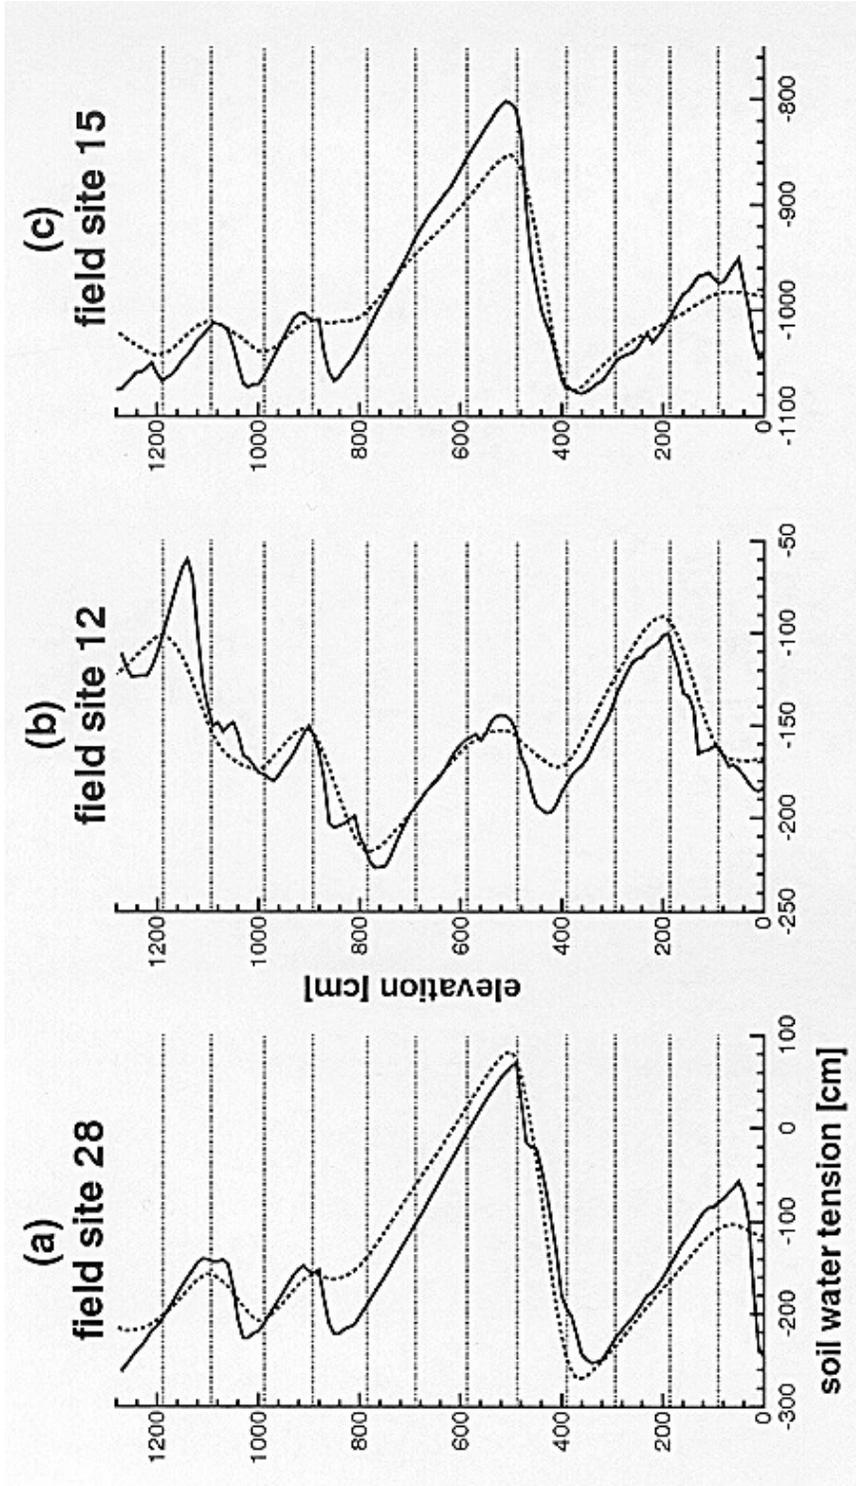


Figure 10.4 third of three tensiometer nests for three different field sites. The vertical location of the tensiometer data are indicated by dotted lines. Field site #28 and #12 are wet soils with high and moderate variability in logK respectively. Field site #15 is a dry soil with similar textural variability as site #12, but a head variance comparable to site #28.

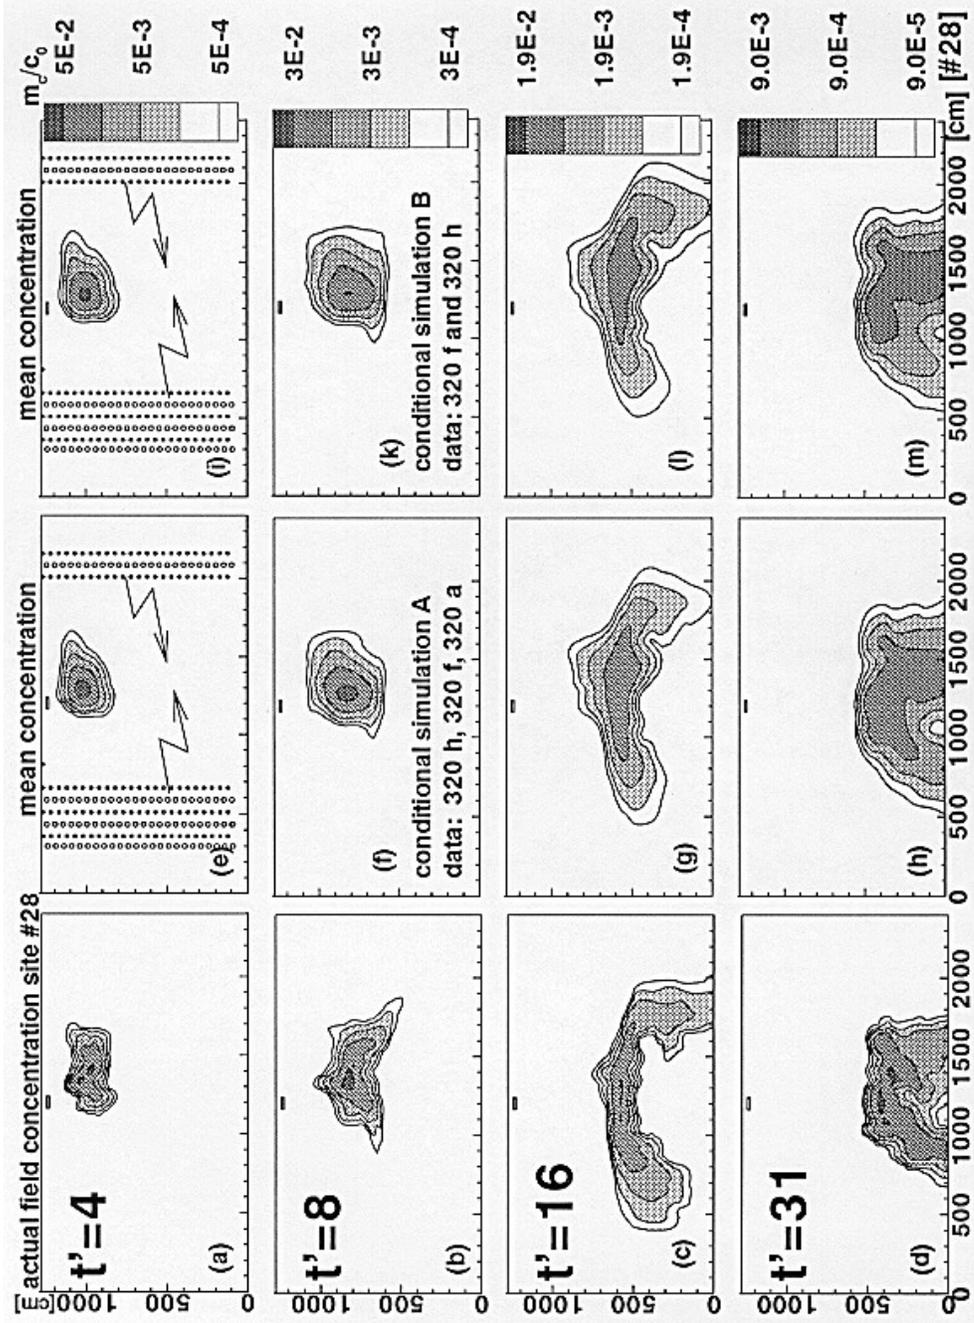


Figure 10.5

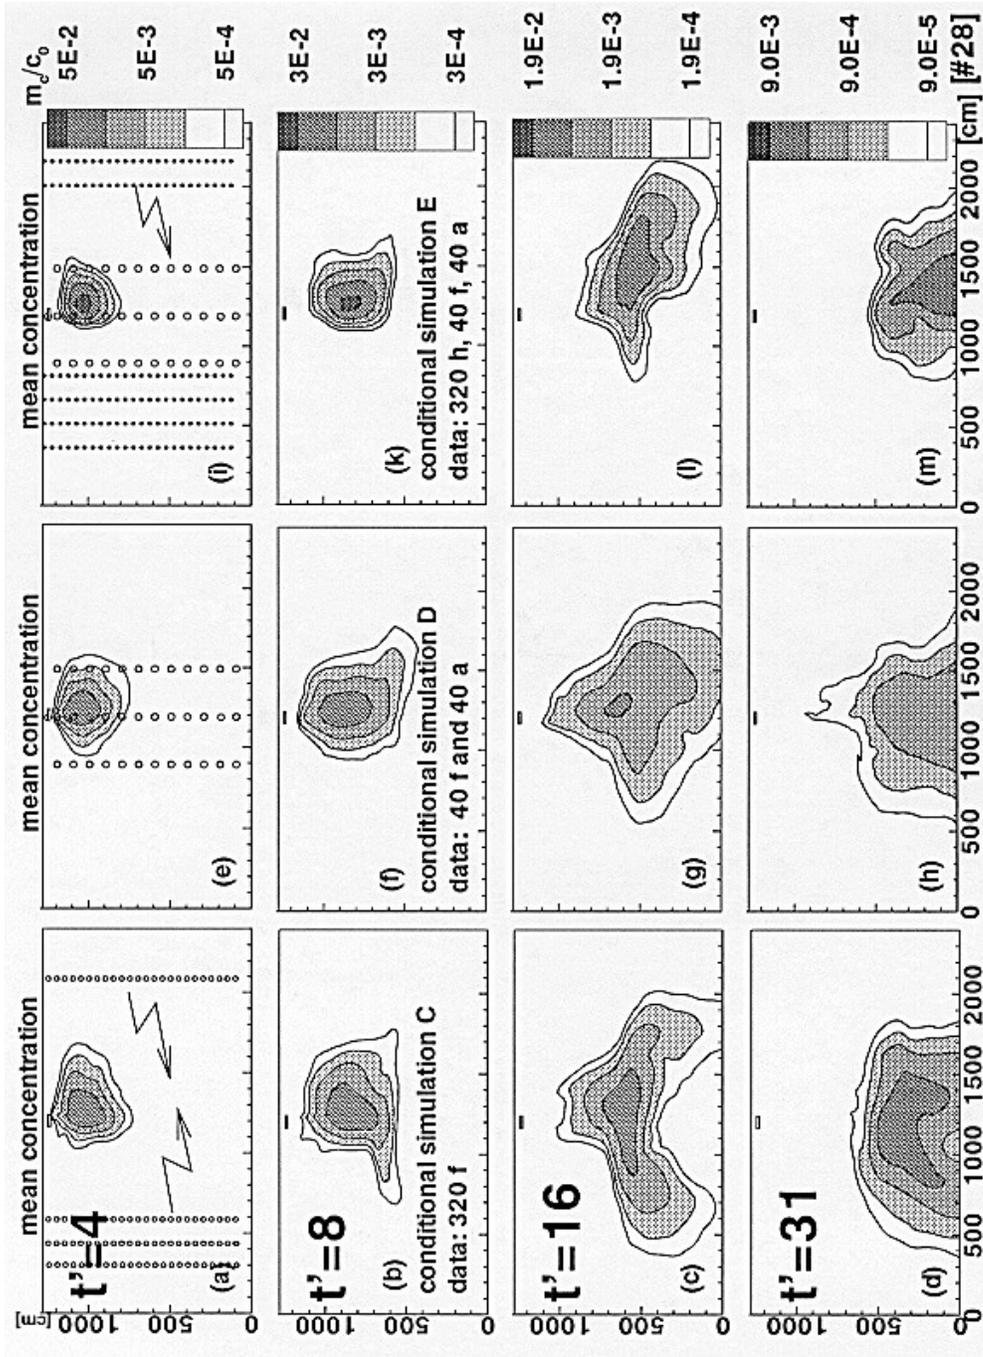


Figure 10.6

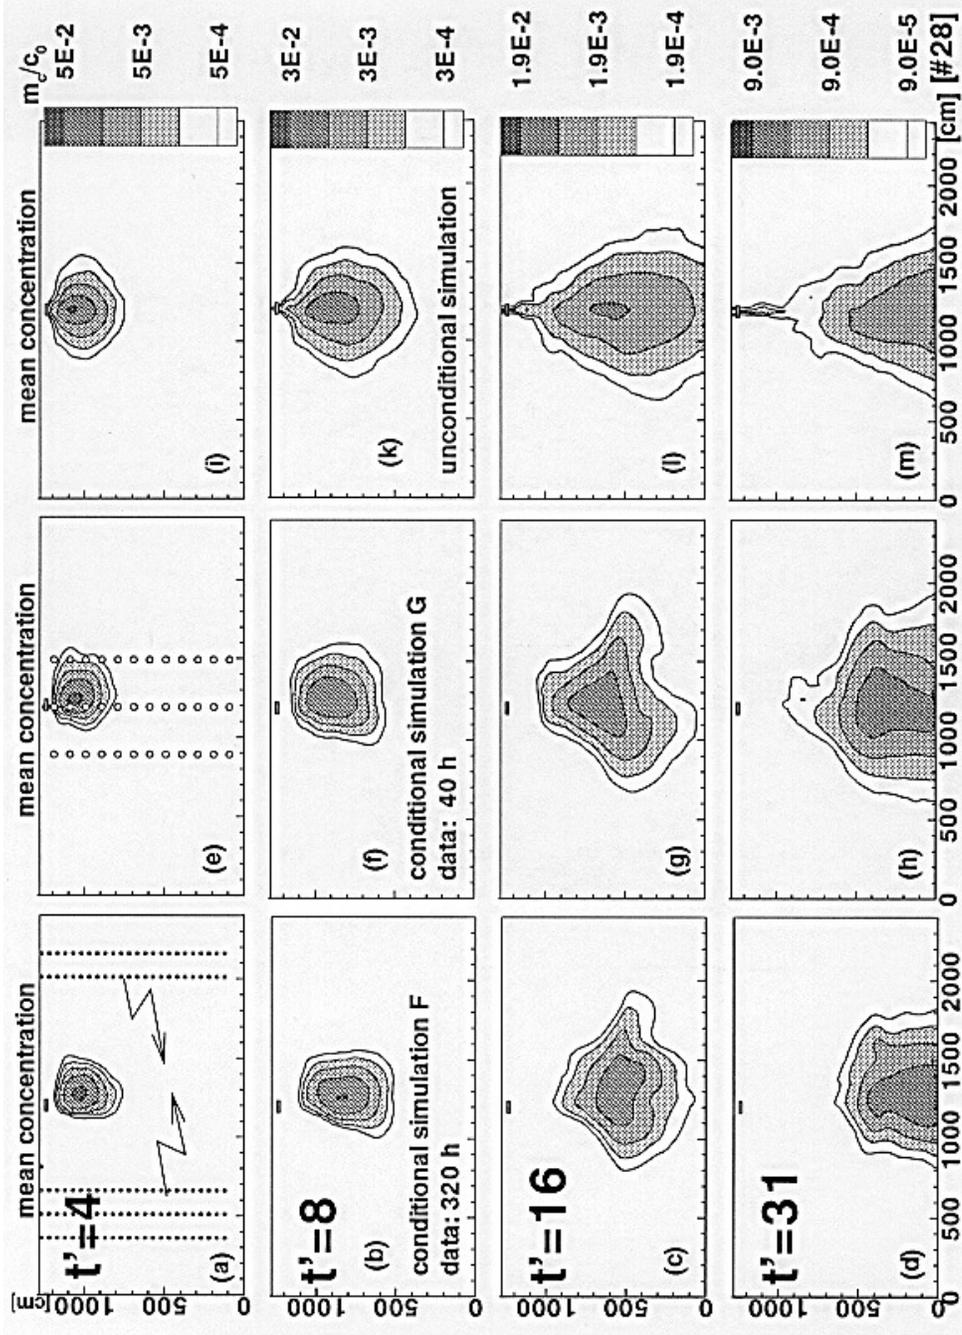


Figure 10.7

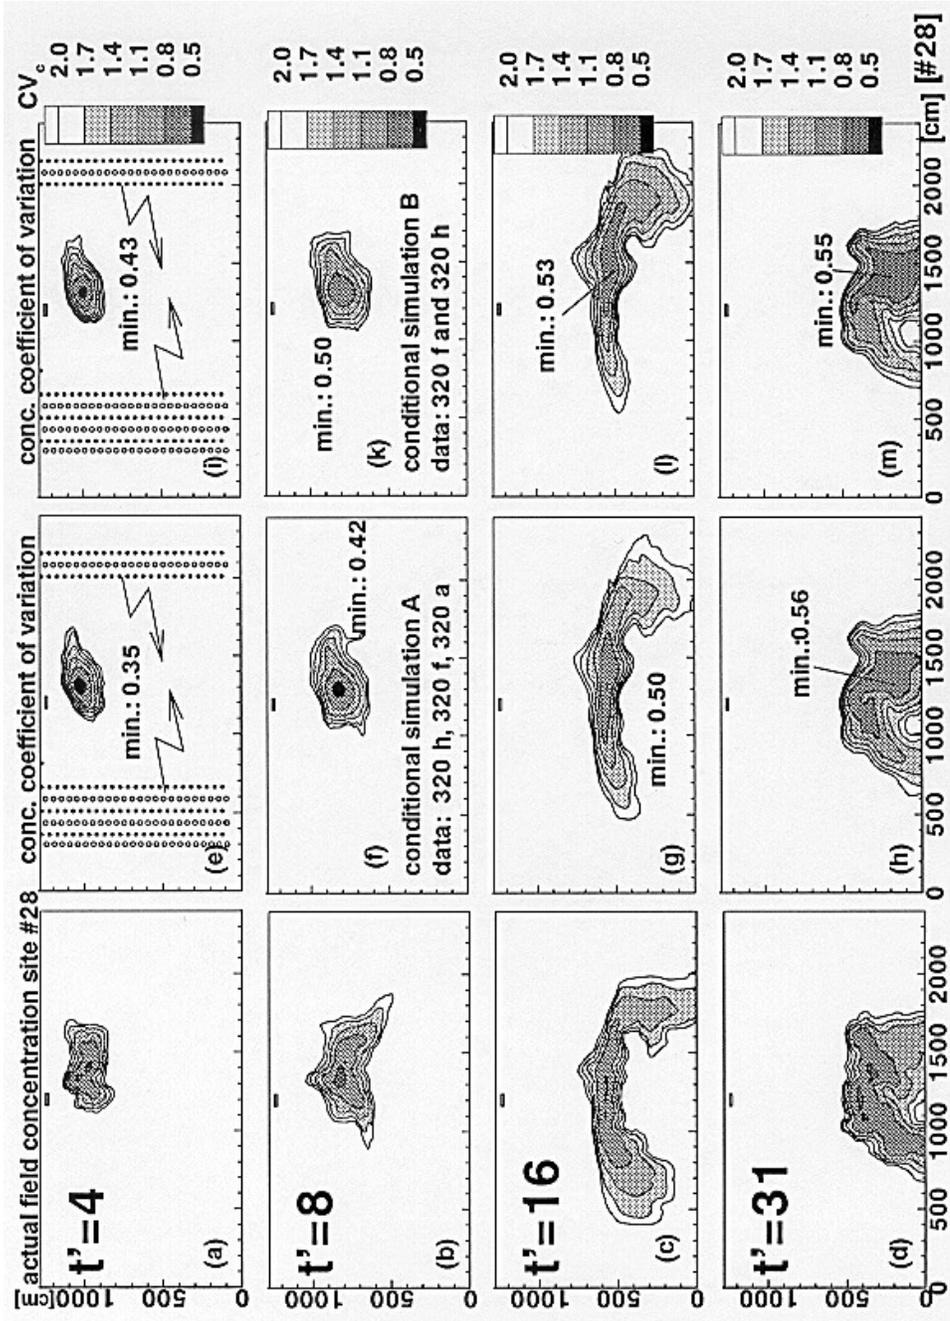


Figure 10.8

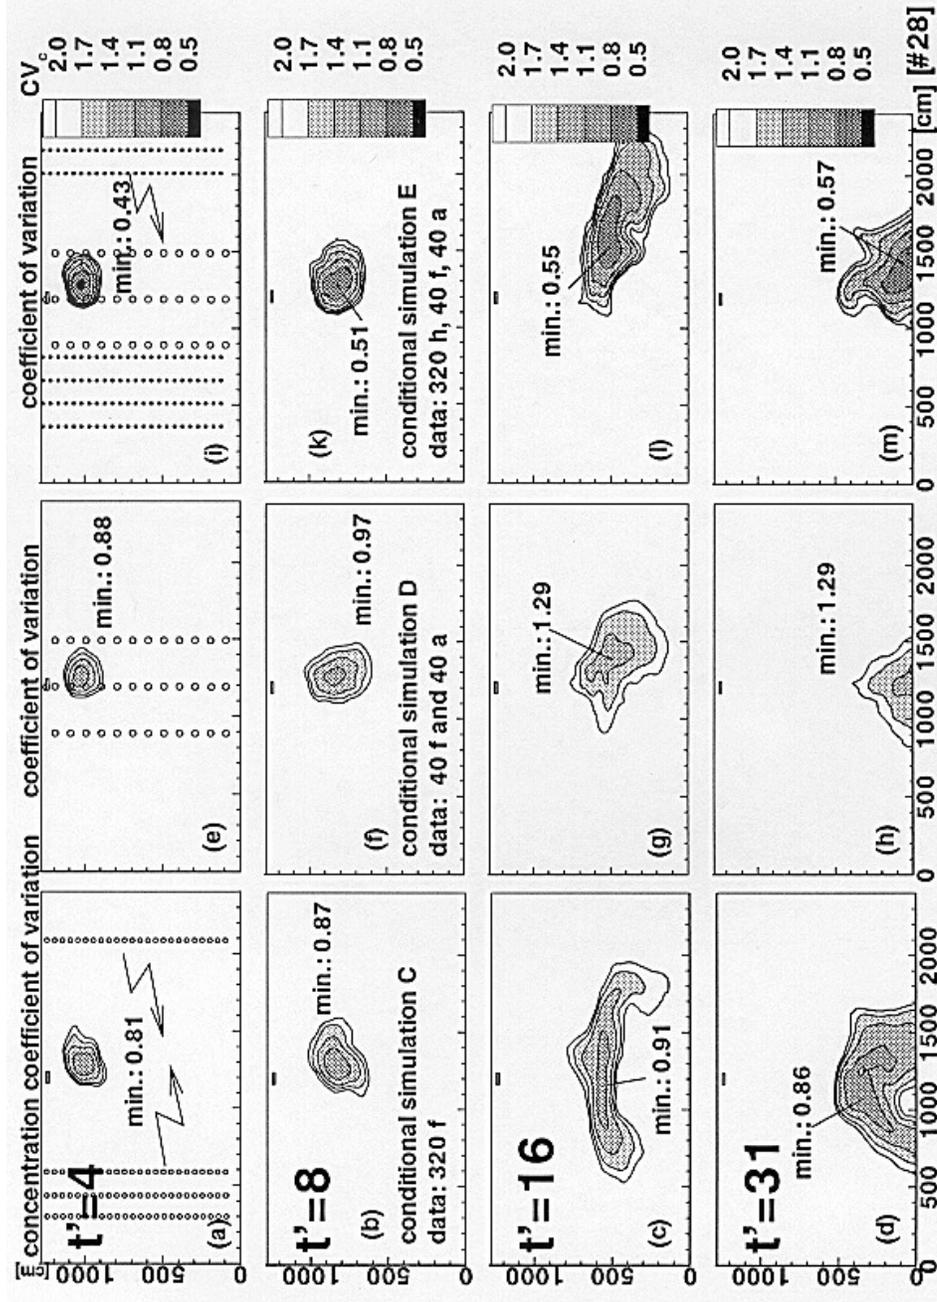


Figure 10.9

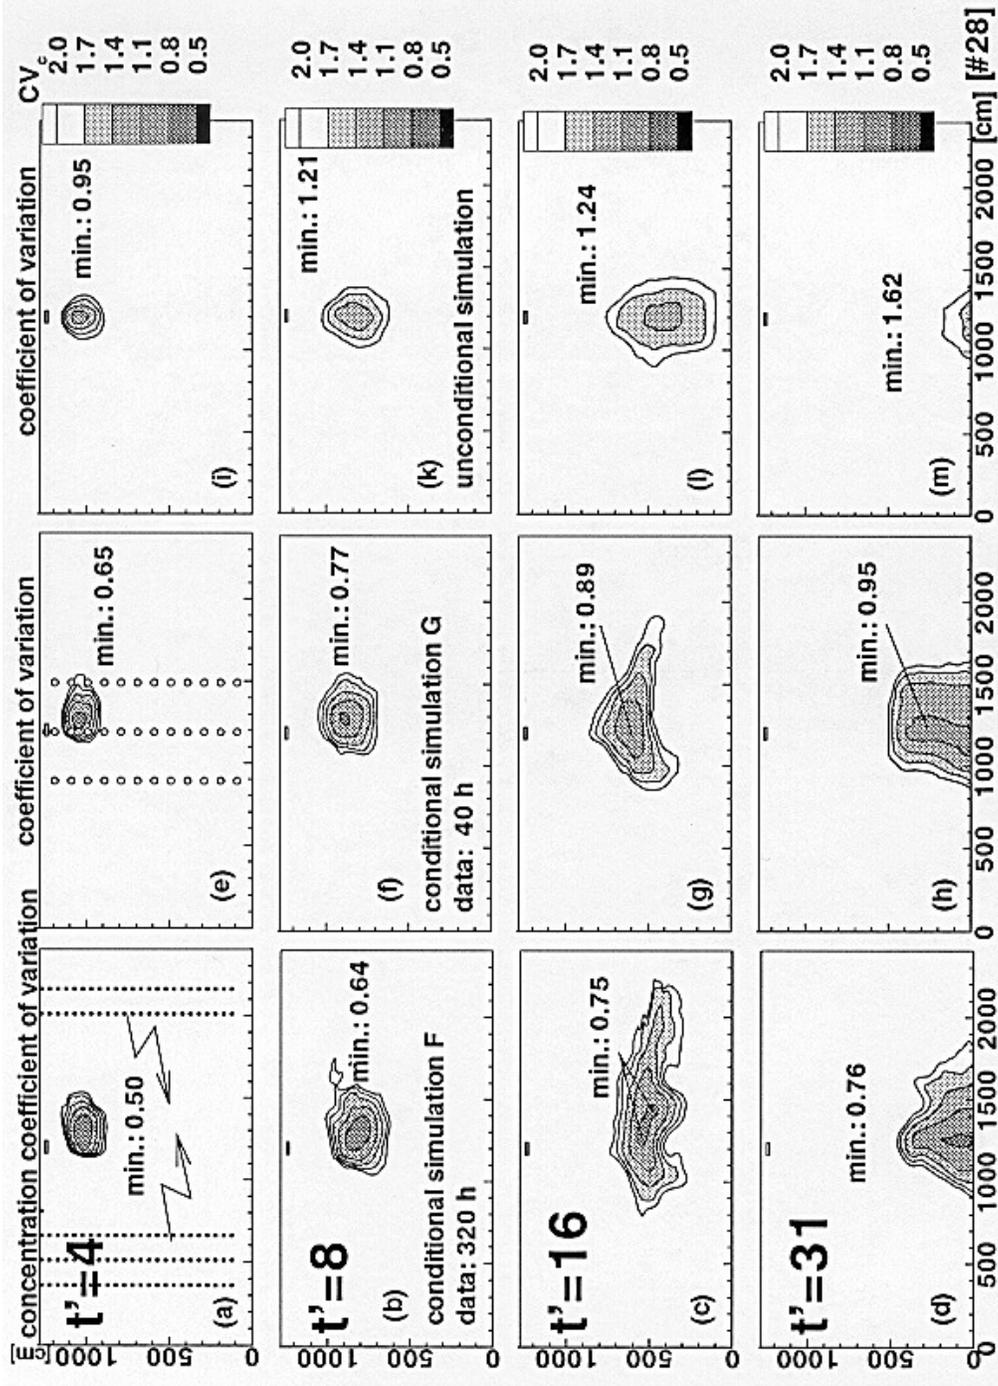


Figure 10.10

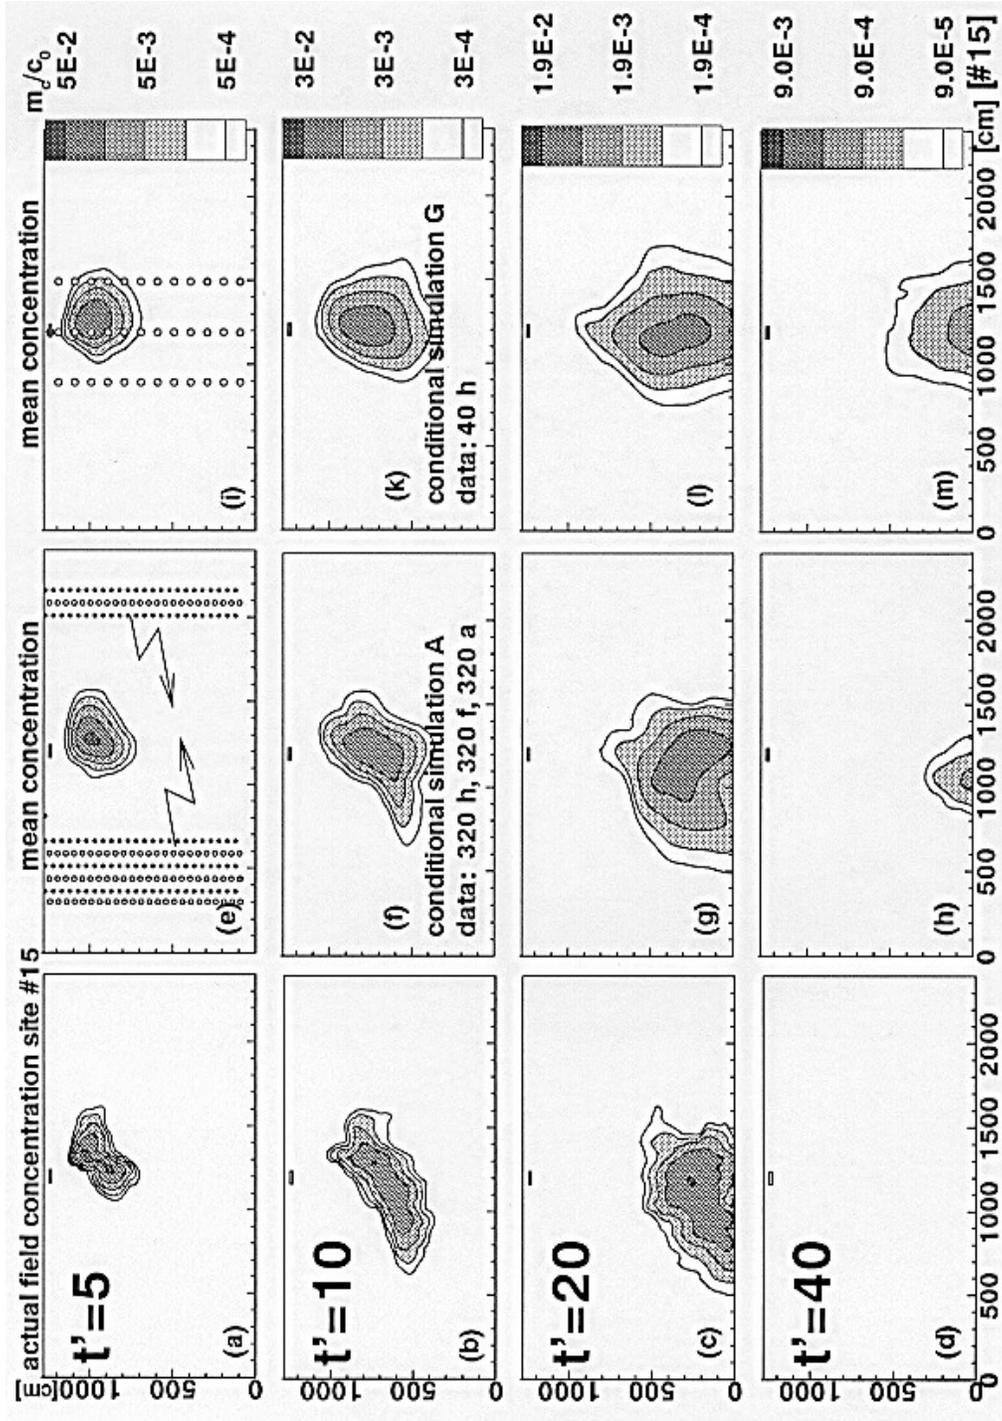


Figure 10.11

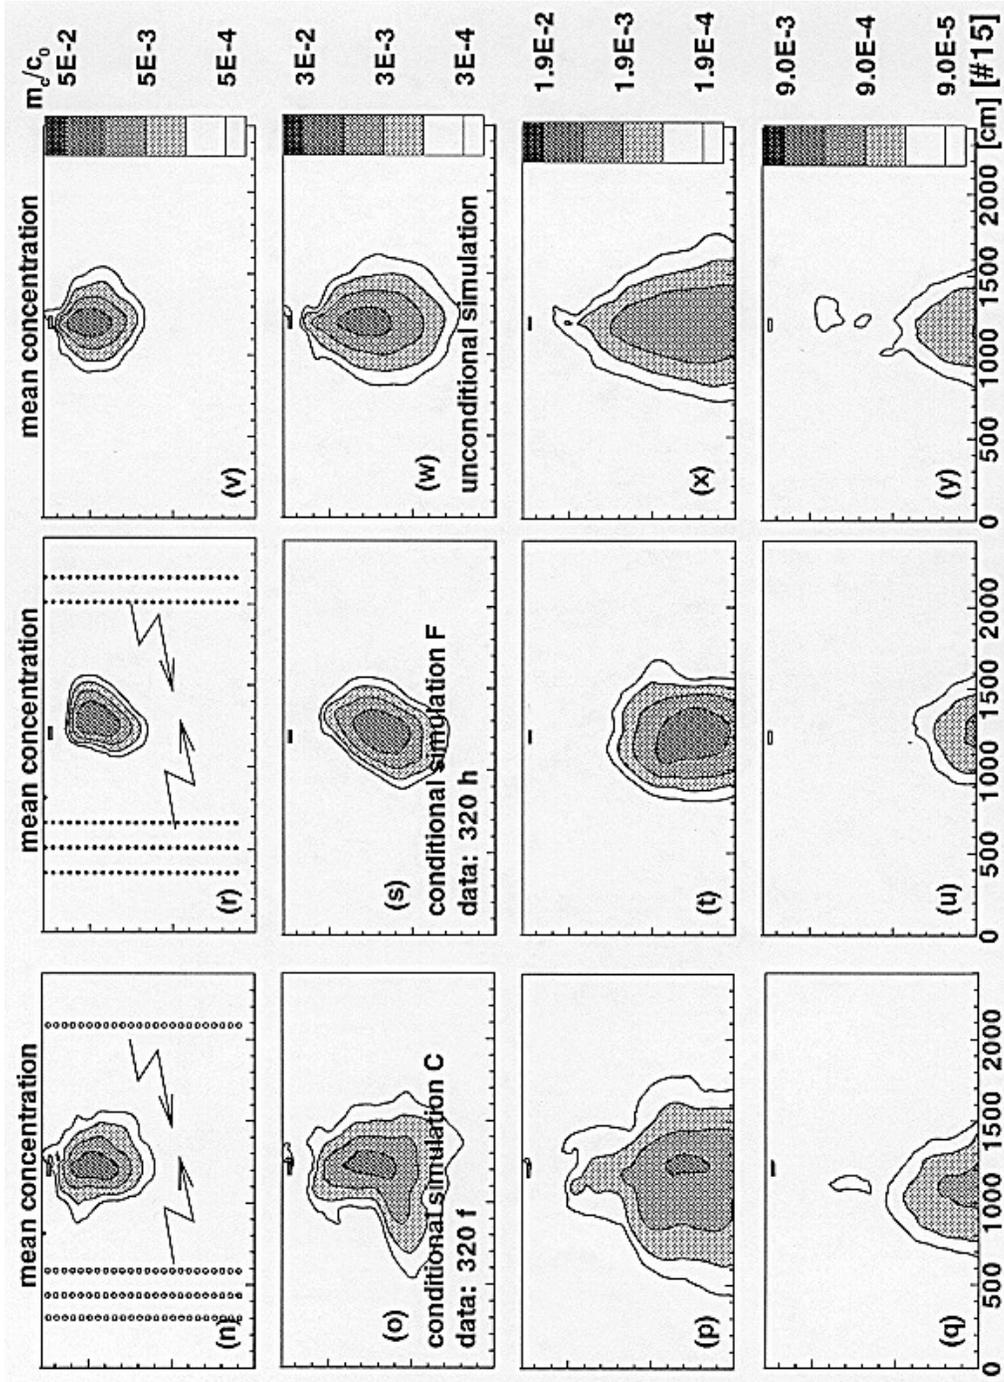


Figure 10.11

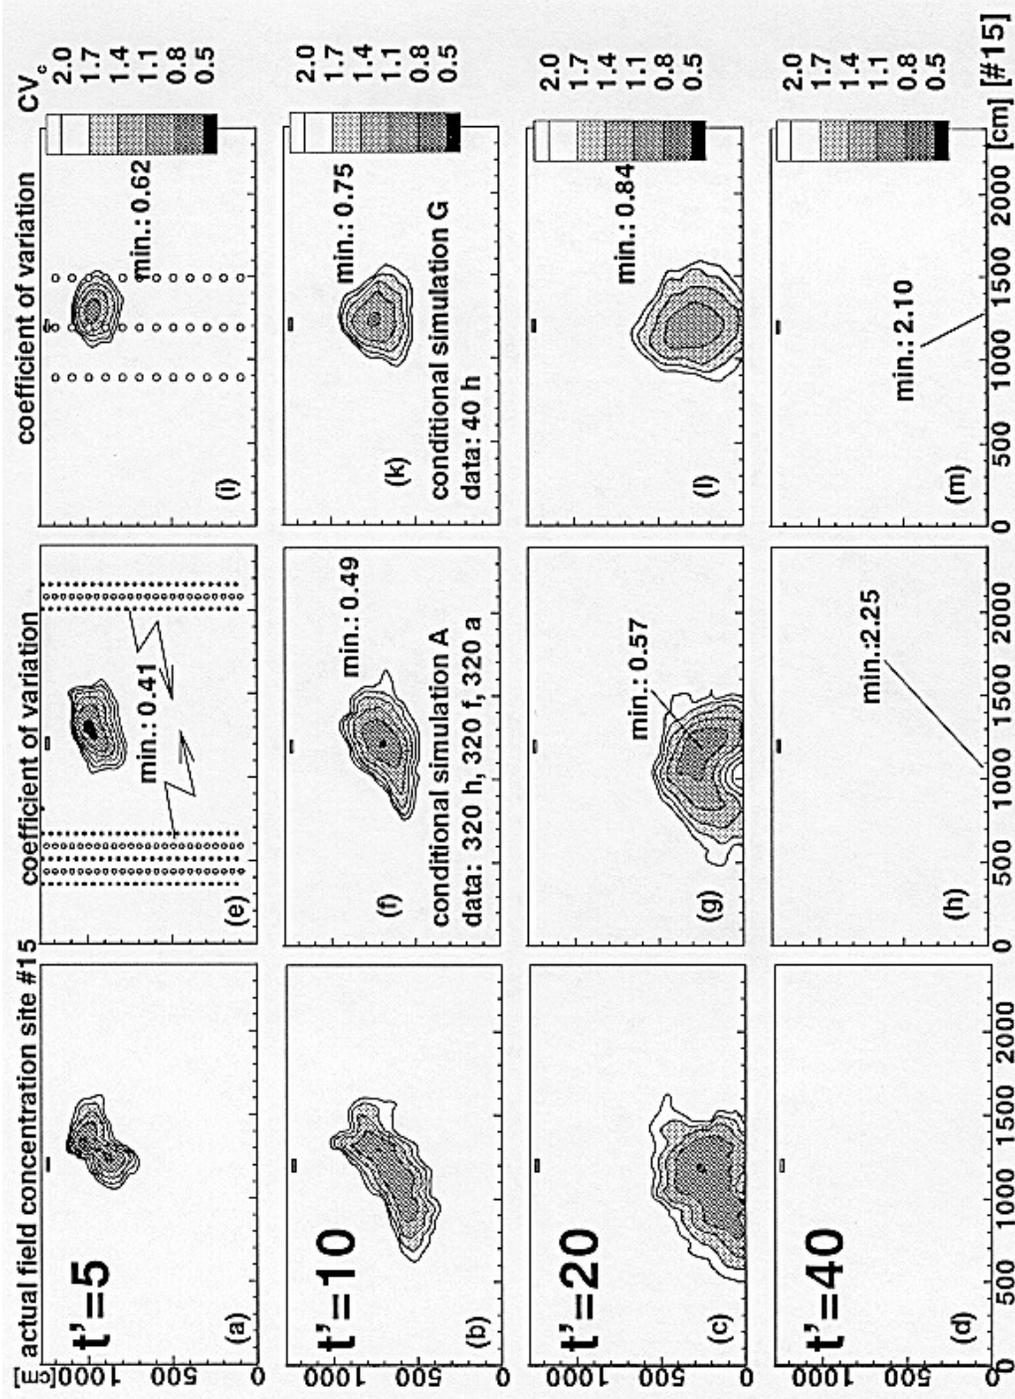


Figure 10.12

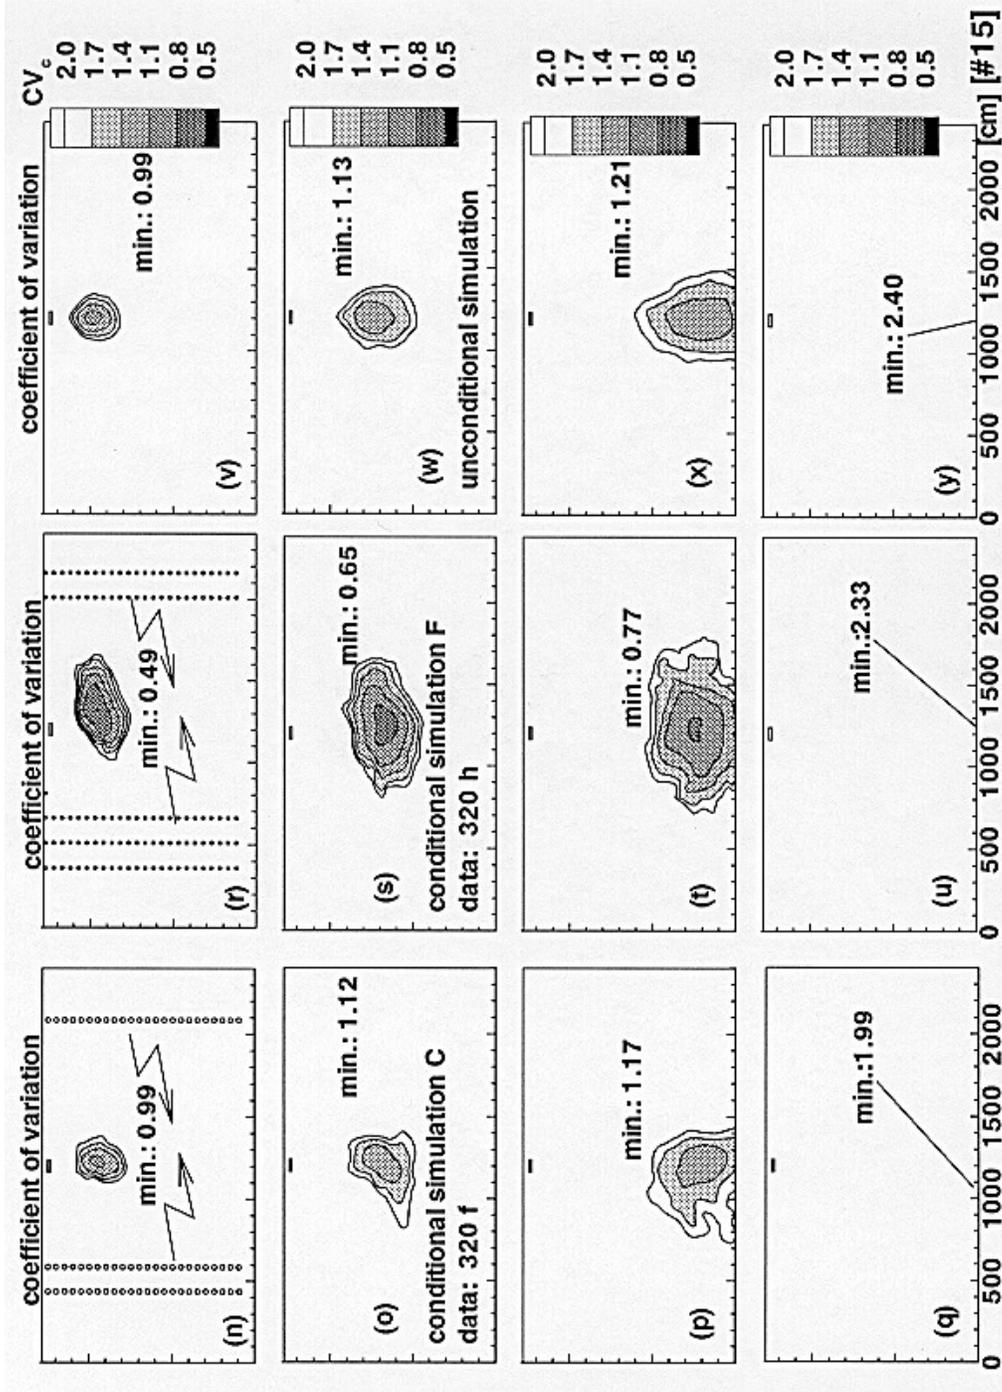


Figure 10.12

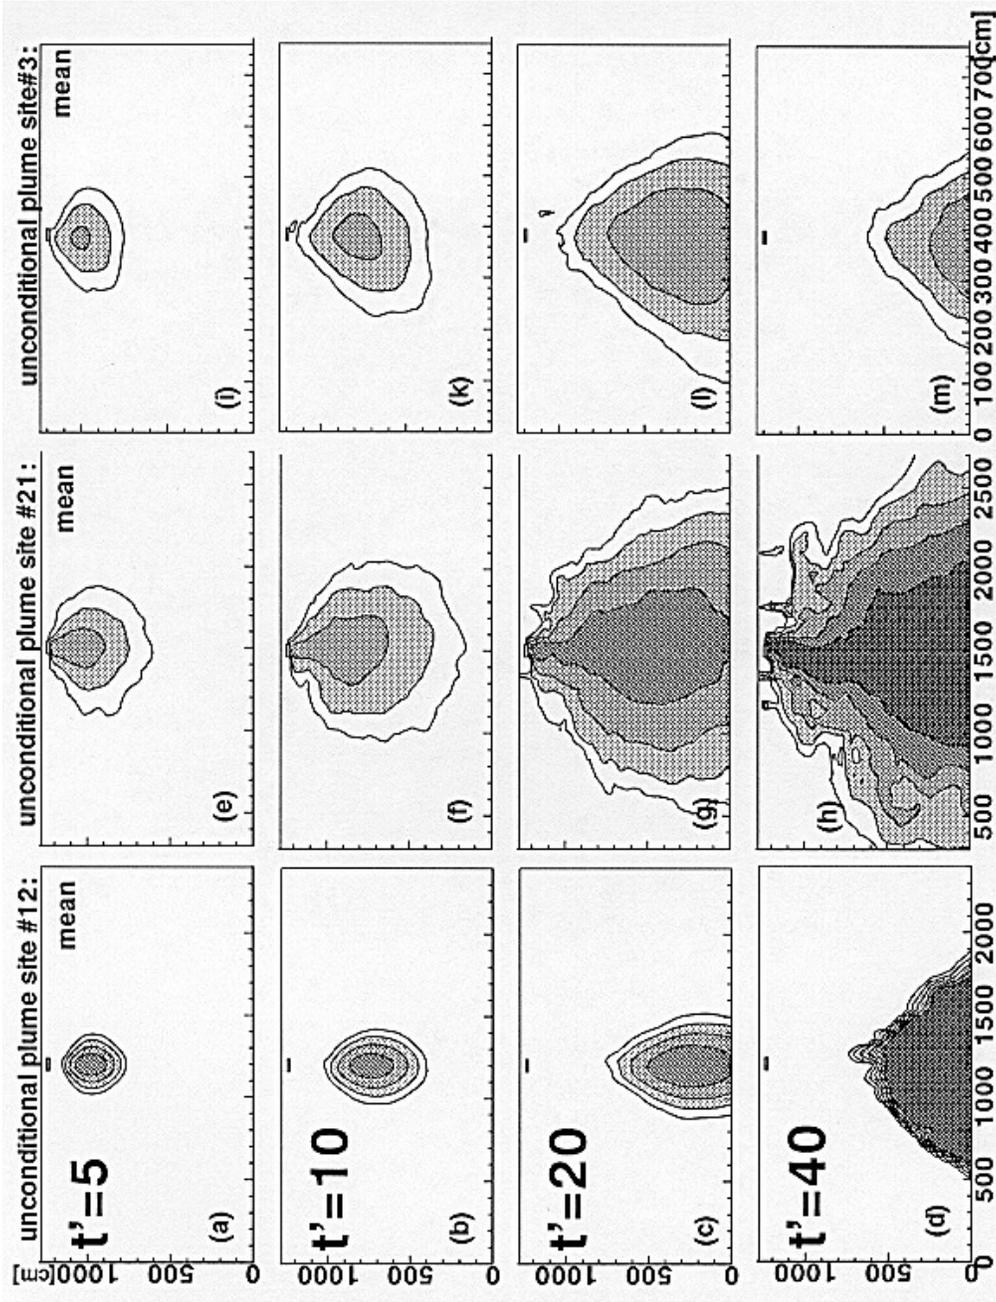


Figure 10.13

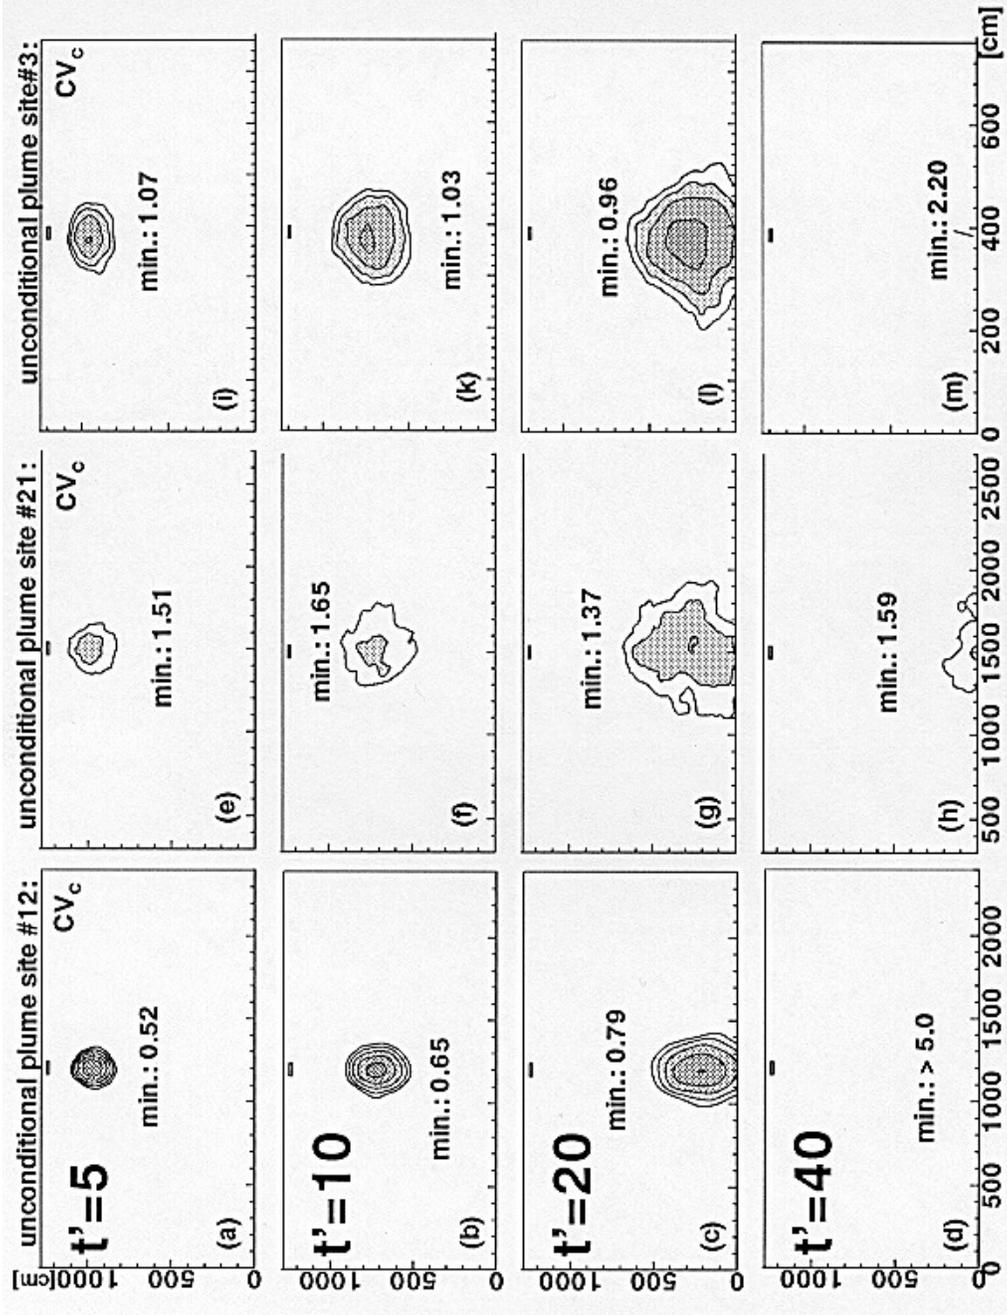


Figure 10.14

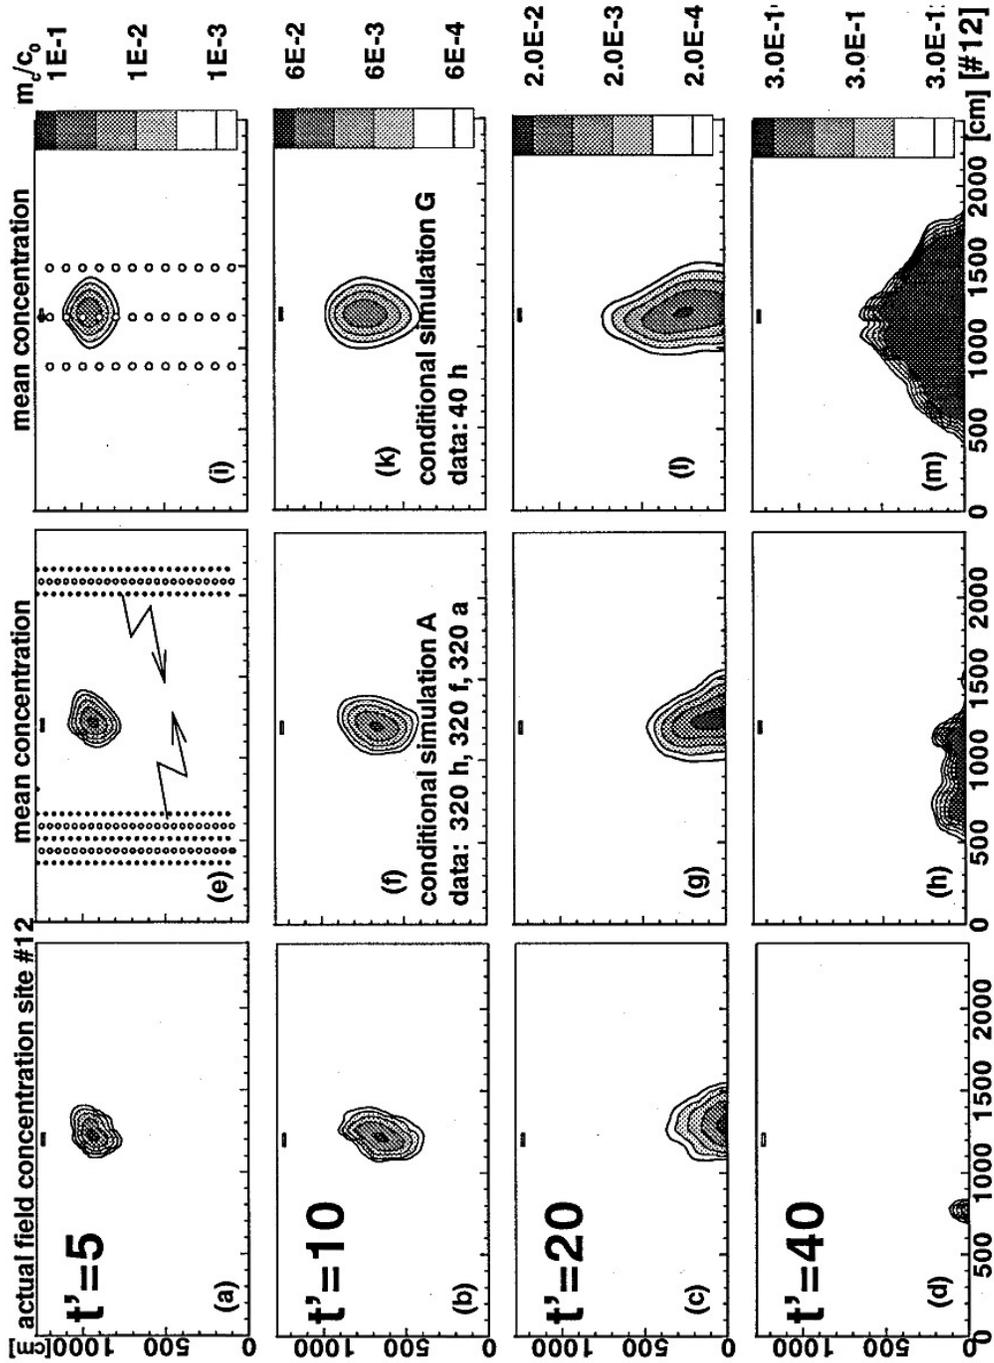


Figure 10.15

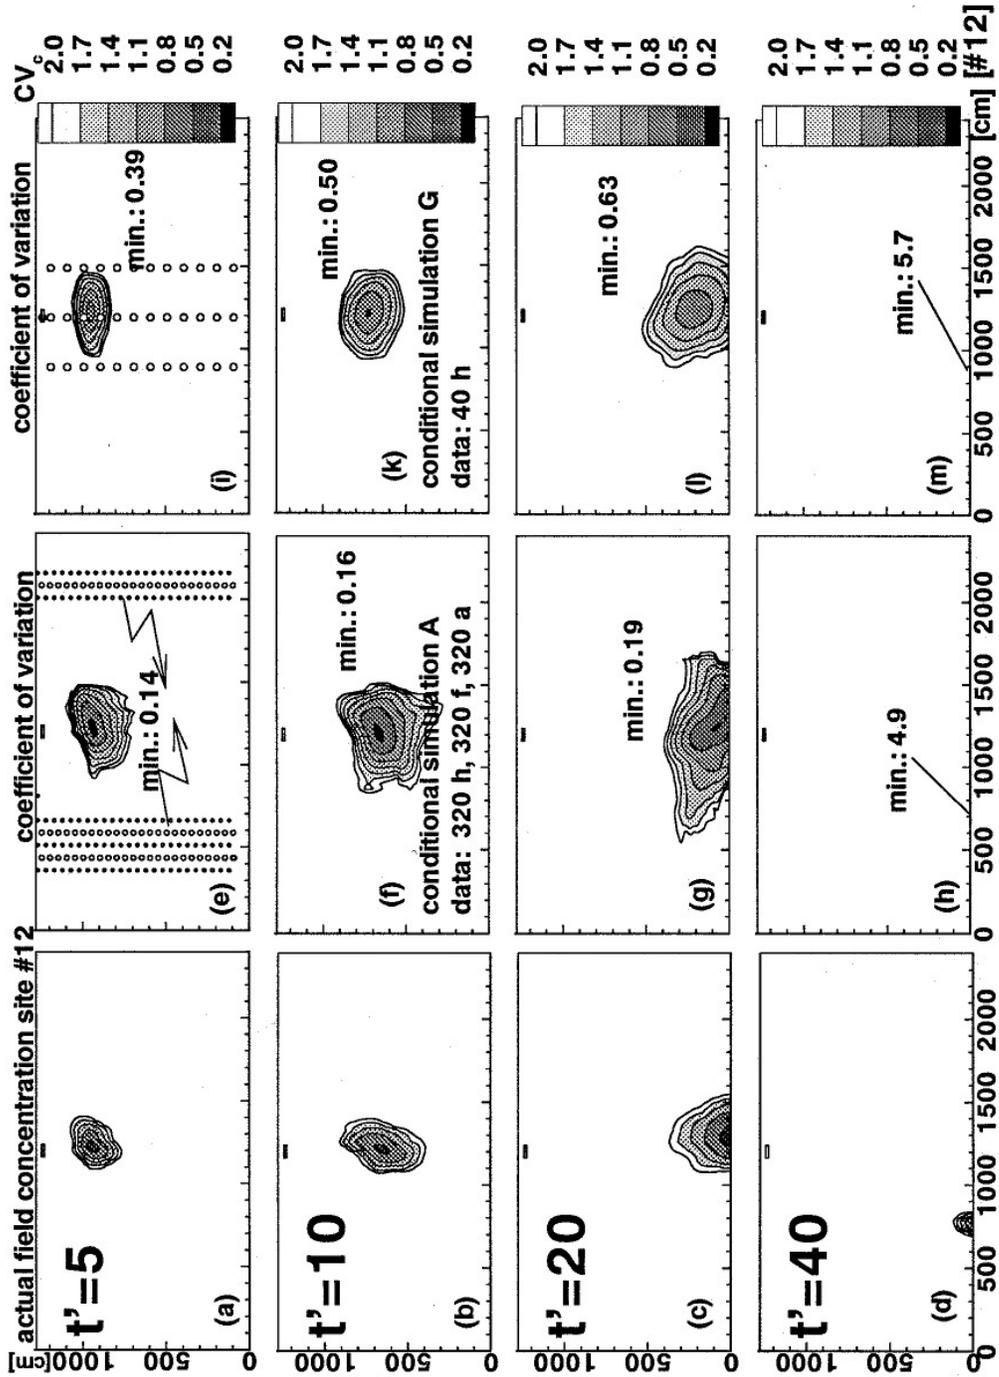


Figure 10.16

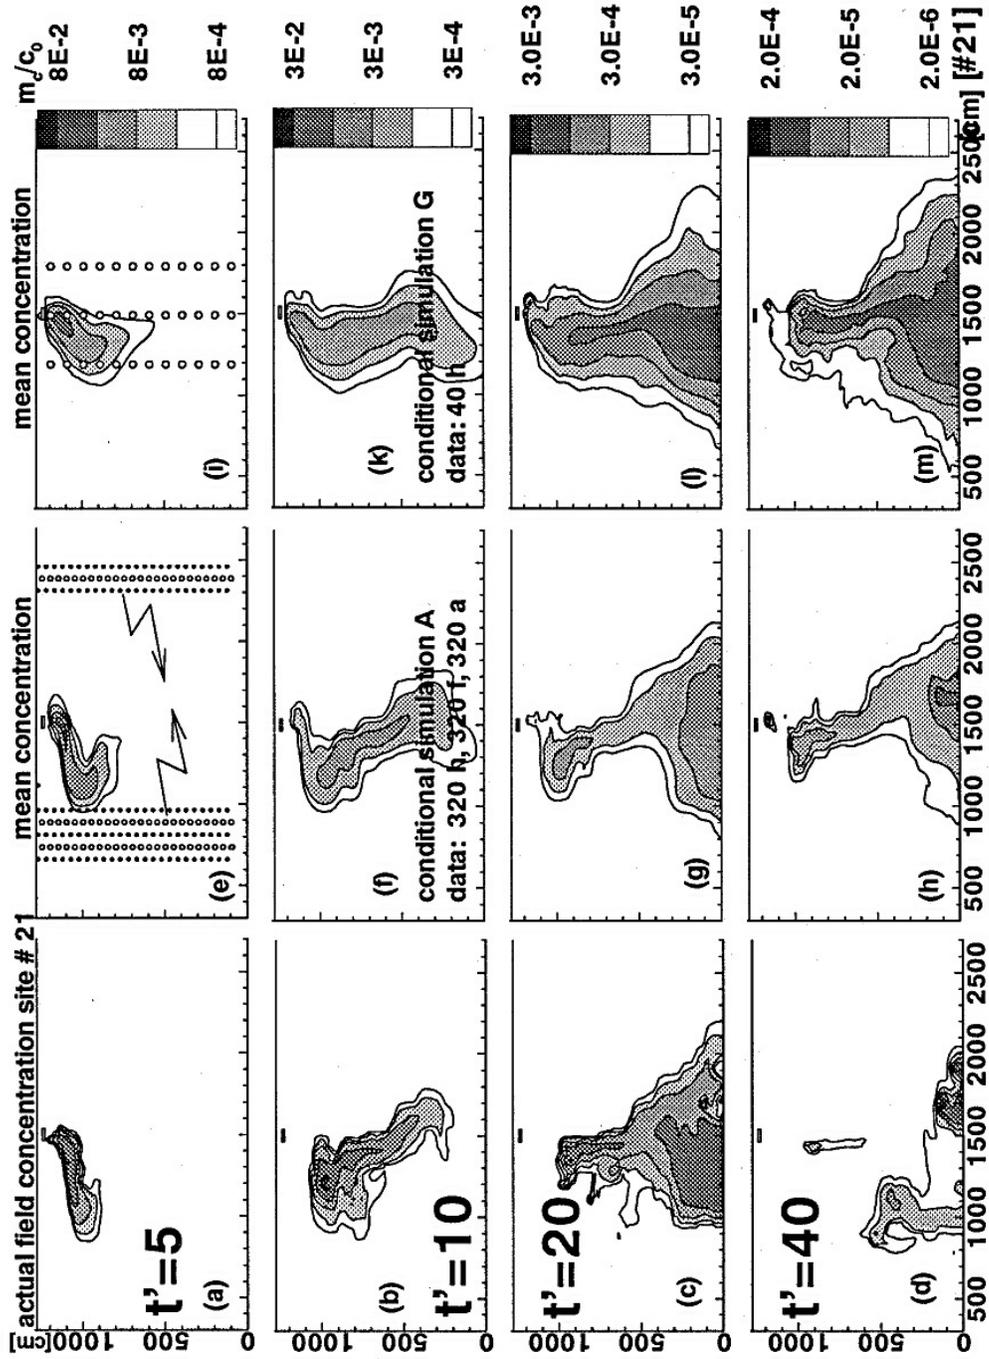


Figure 10.17

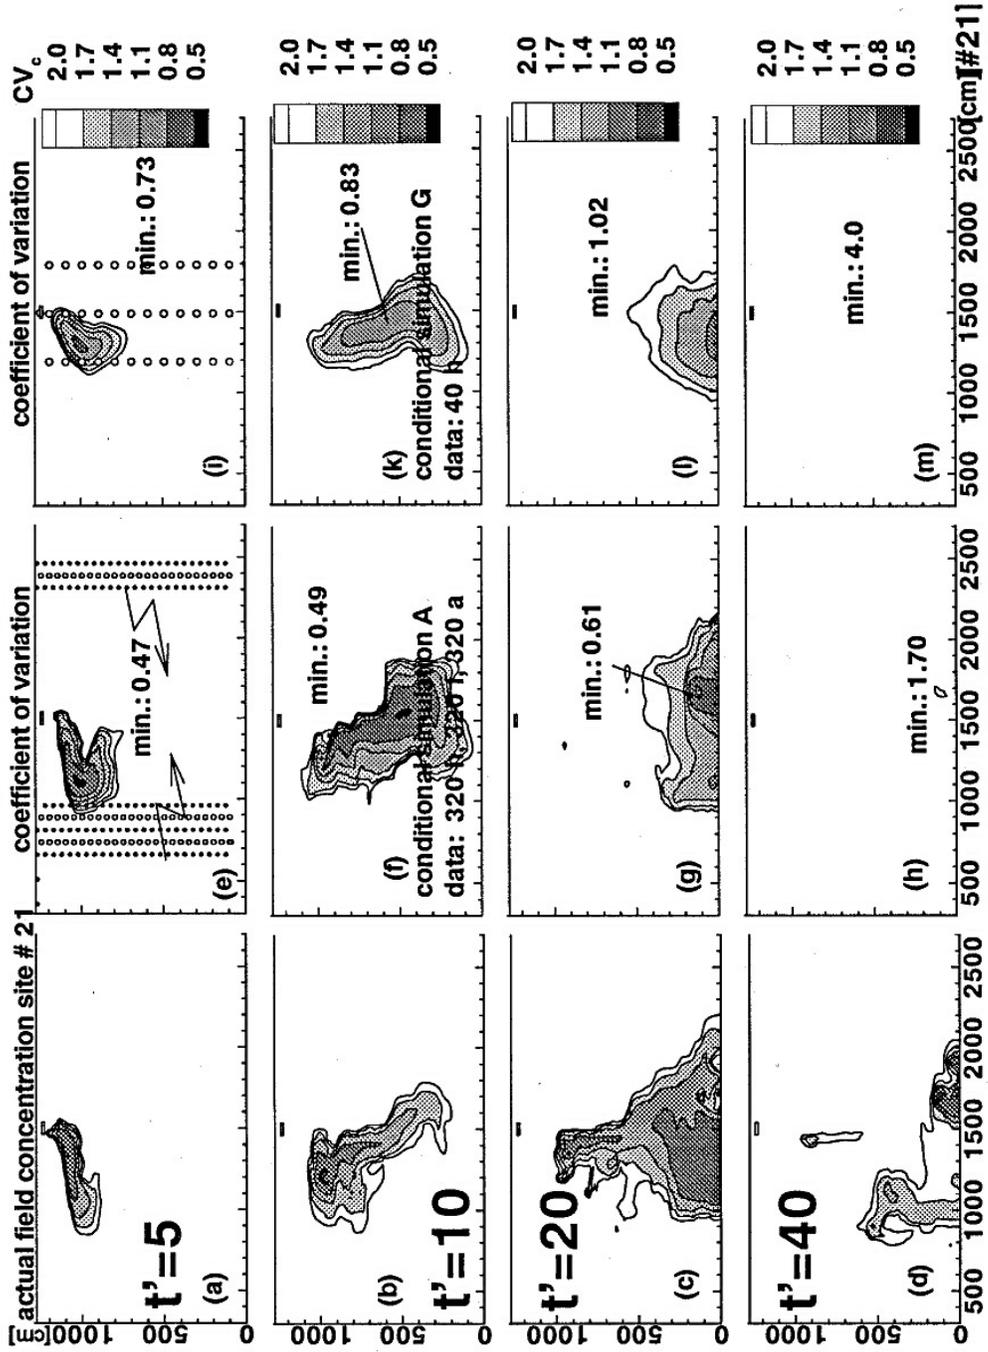


Figure 10.18

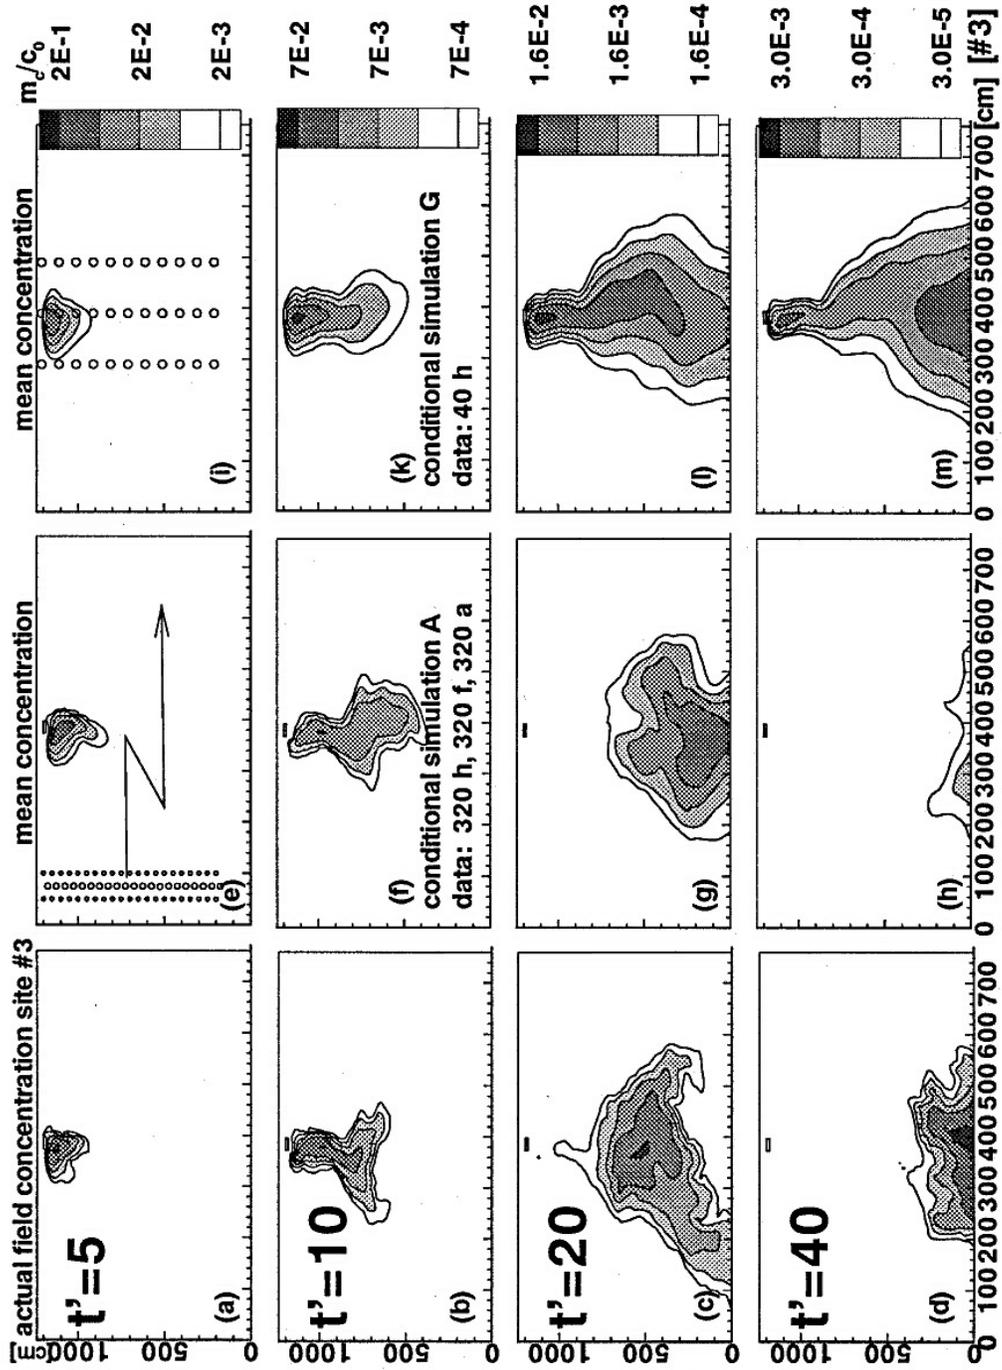


Figure 10.19

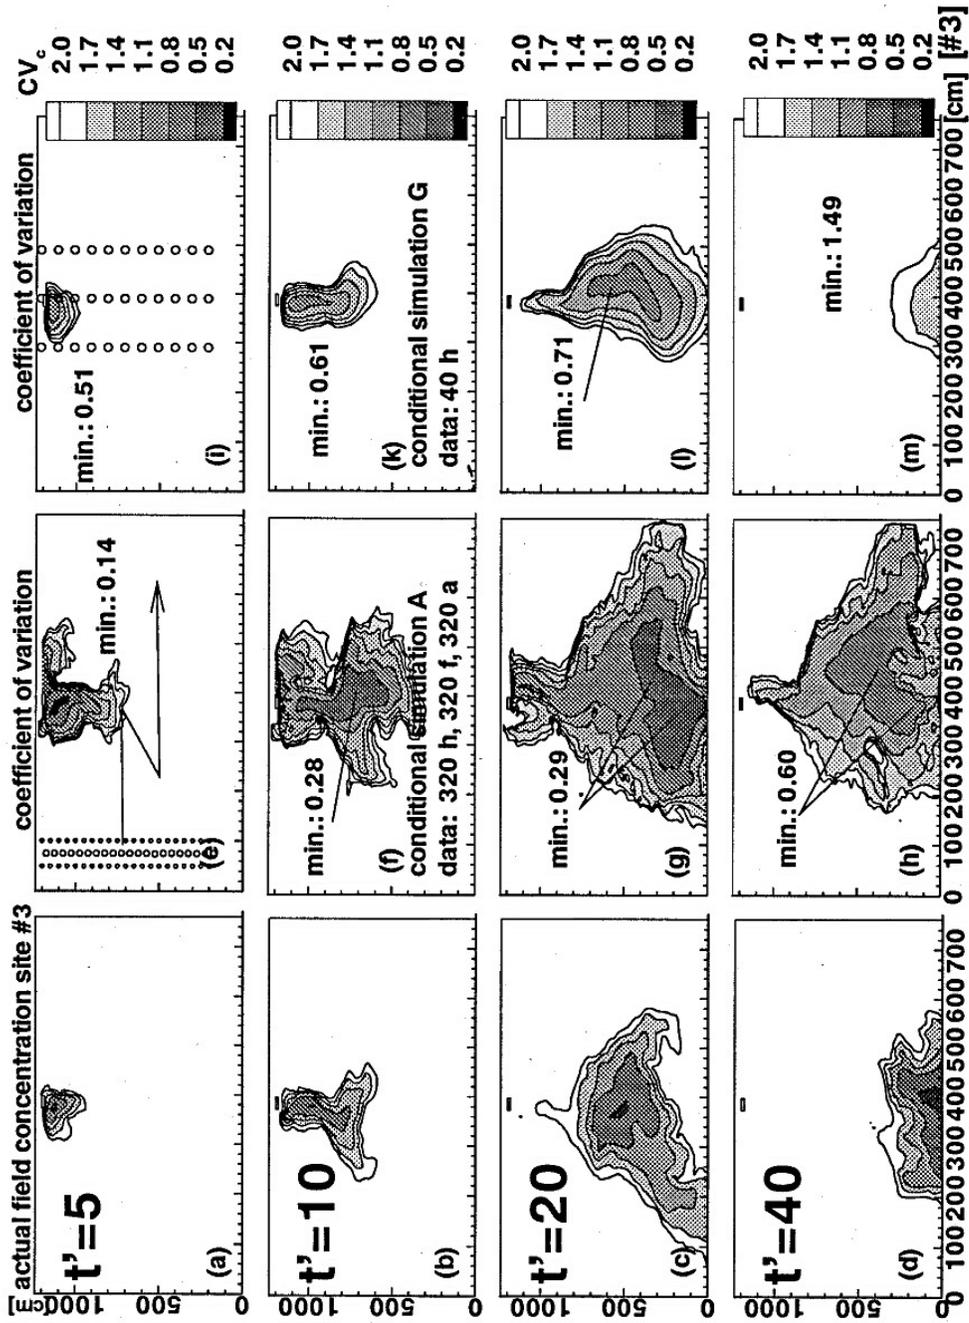


Figure 10.20

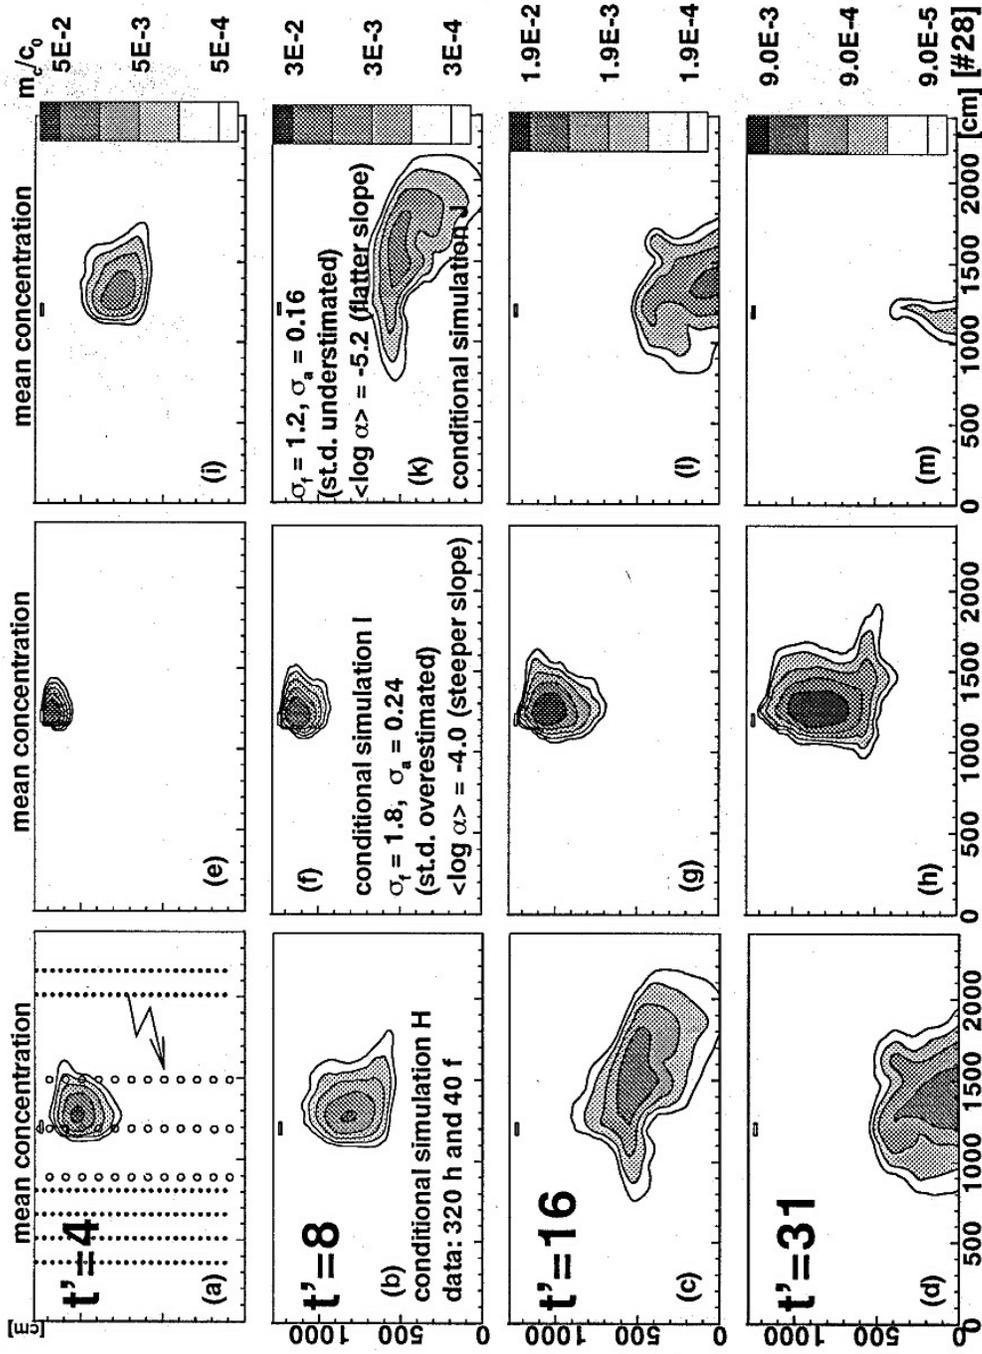


Figure 10.21

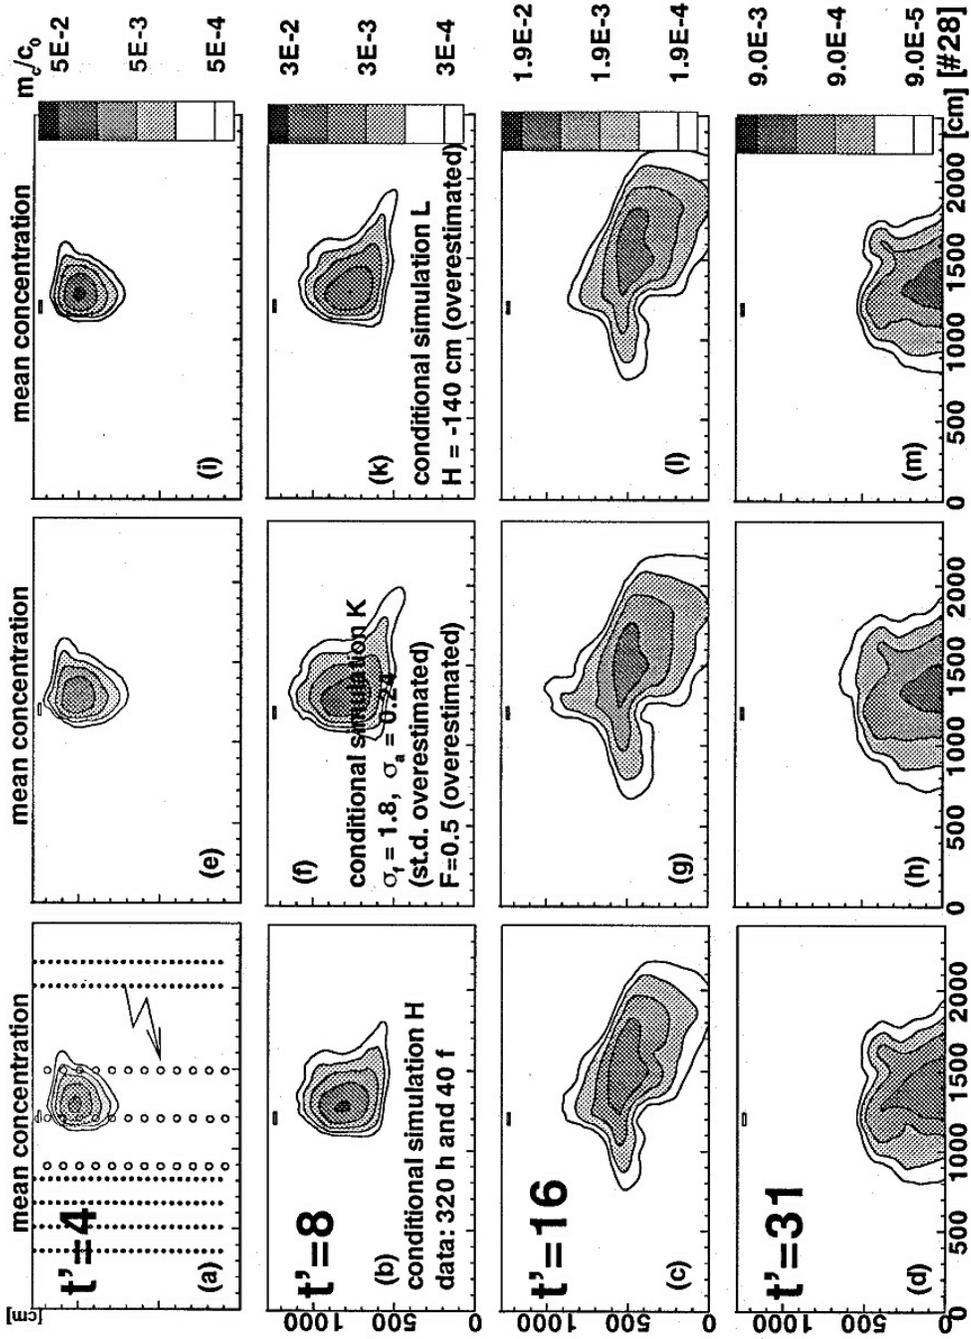


Figure 10.22

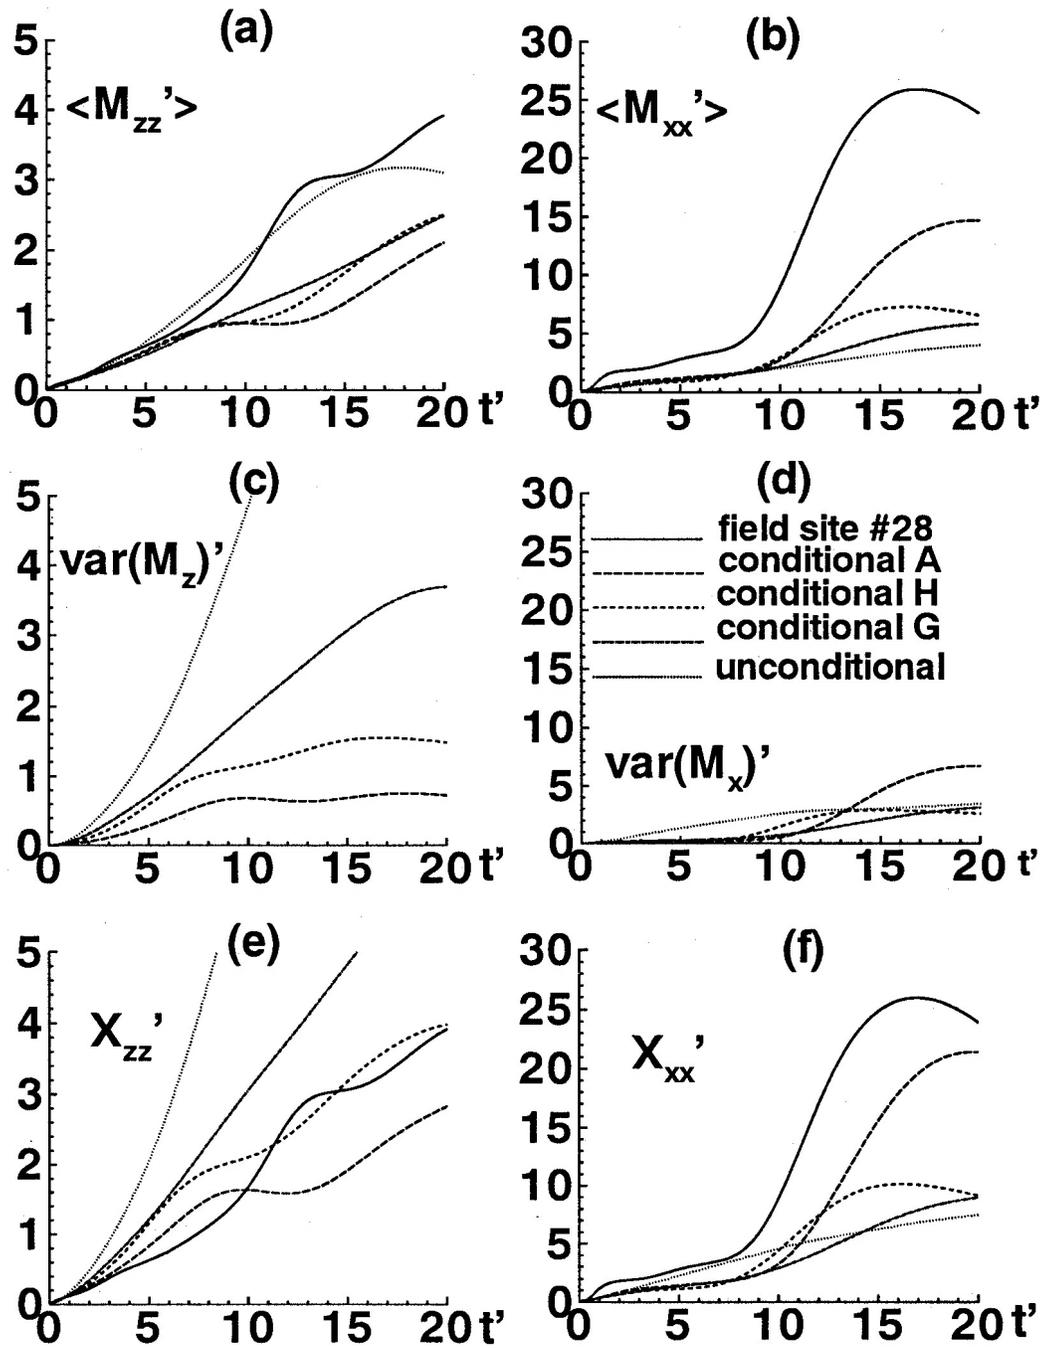


Figure 10.23 Average moment of inertia of individual plumes (a,b), variance of the plume center of mass (c,d), and moment of inertia of the mean plume (e,f) in the vertical (left column) and horizontal dimension (right column). All moments are normalized (') by $(\lambda_{iz} \sigma_f^2)$.

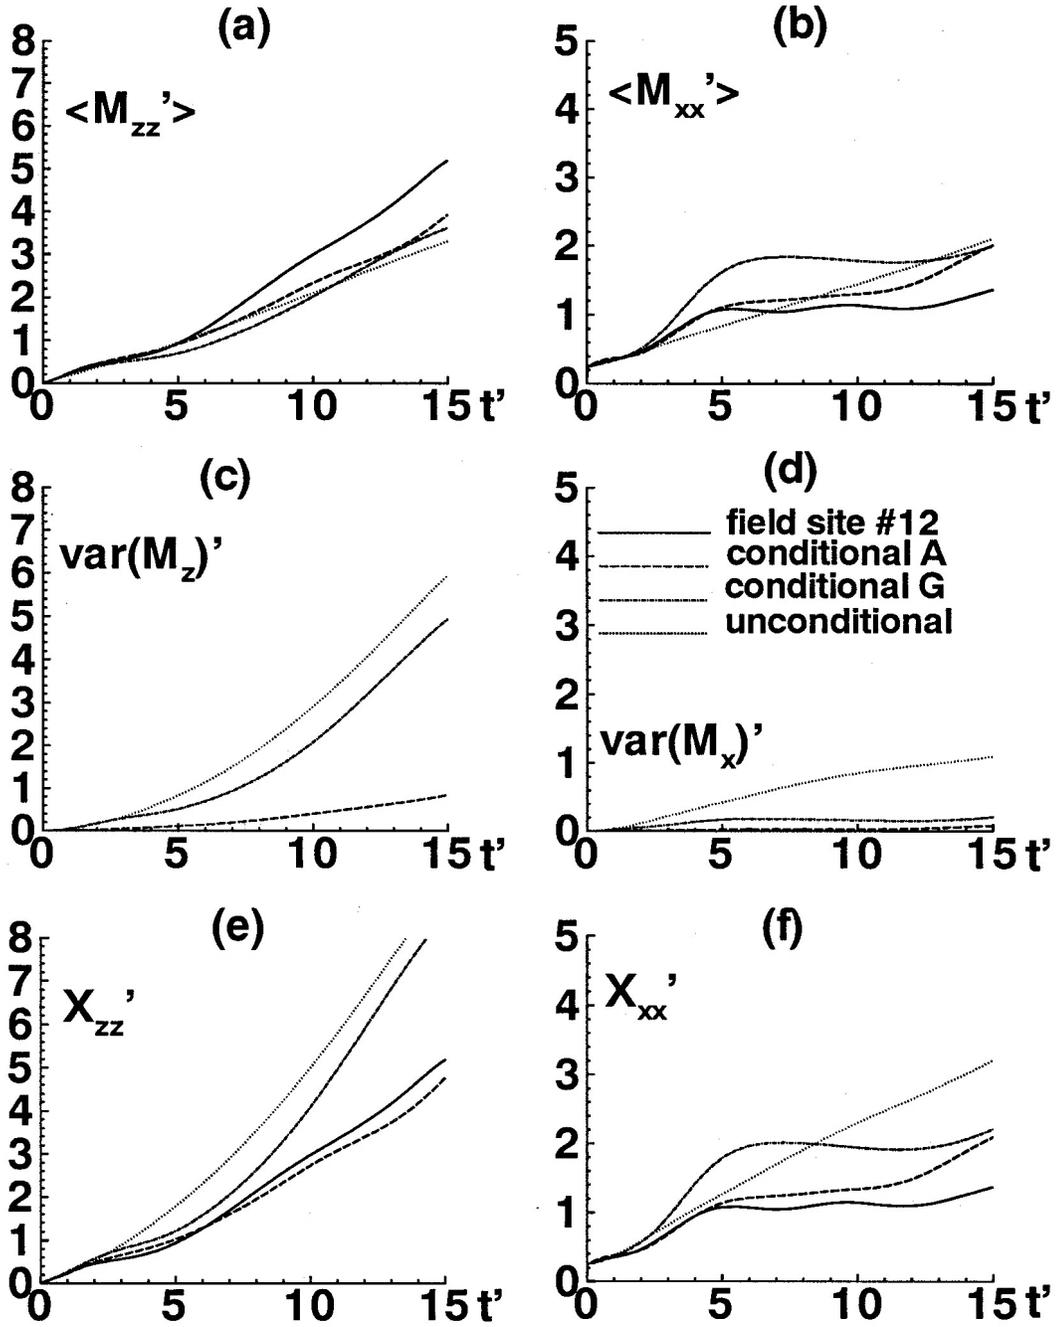


Figure 10.24 Average moment of inertia of individual plumes (a,b), variance of the plume center of mass (c,d), and moment of inertia of the mean plume (e,f) in the vertical (left column) and horizontal dimension (right column). All moments are normalized (') by $(\lambda_{iz} \sigma_f^2)$.

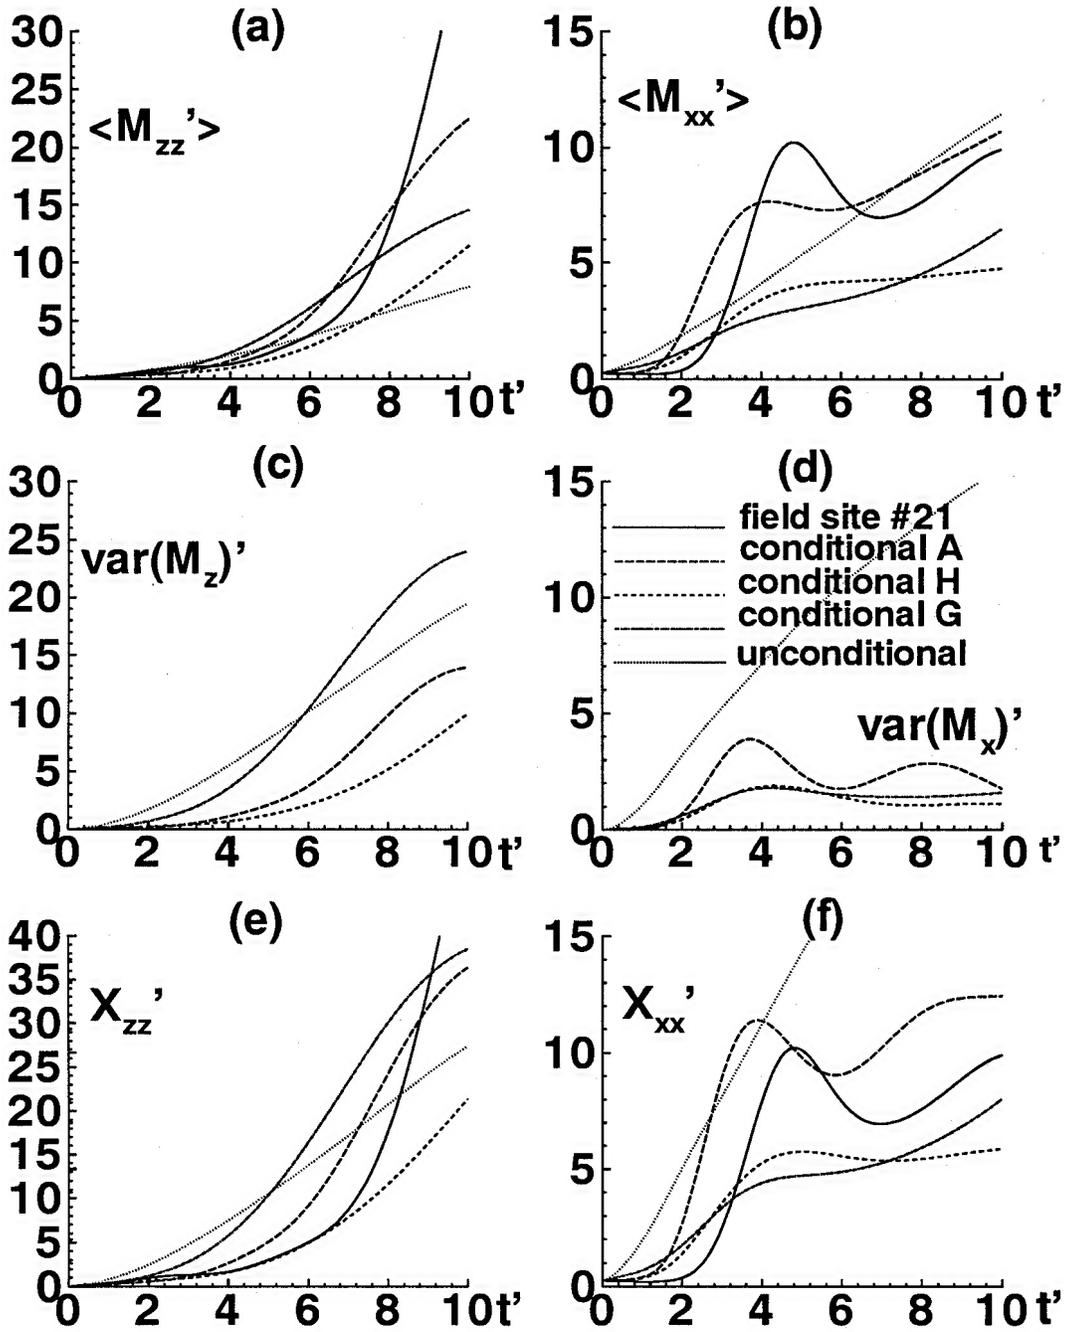
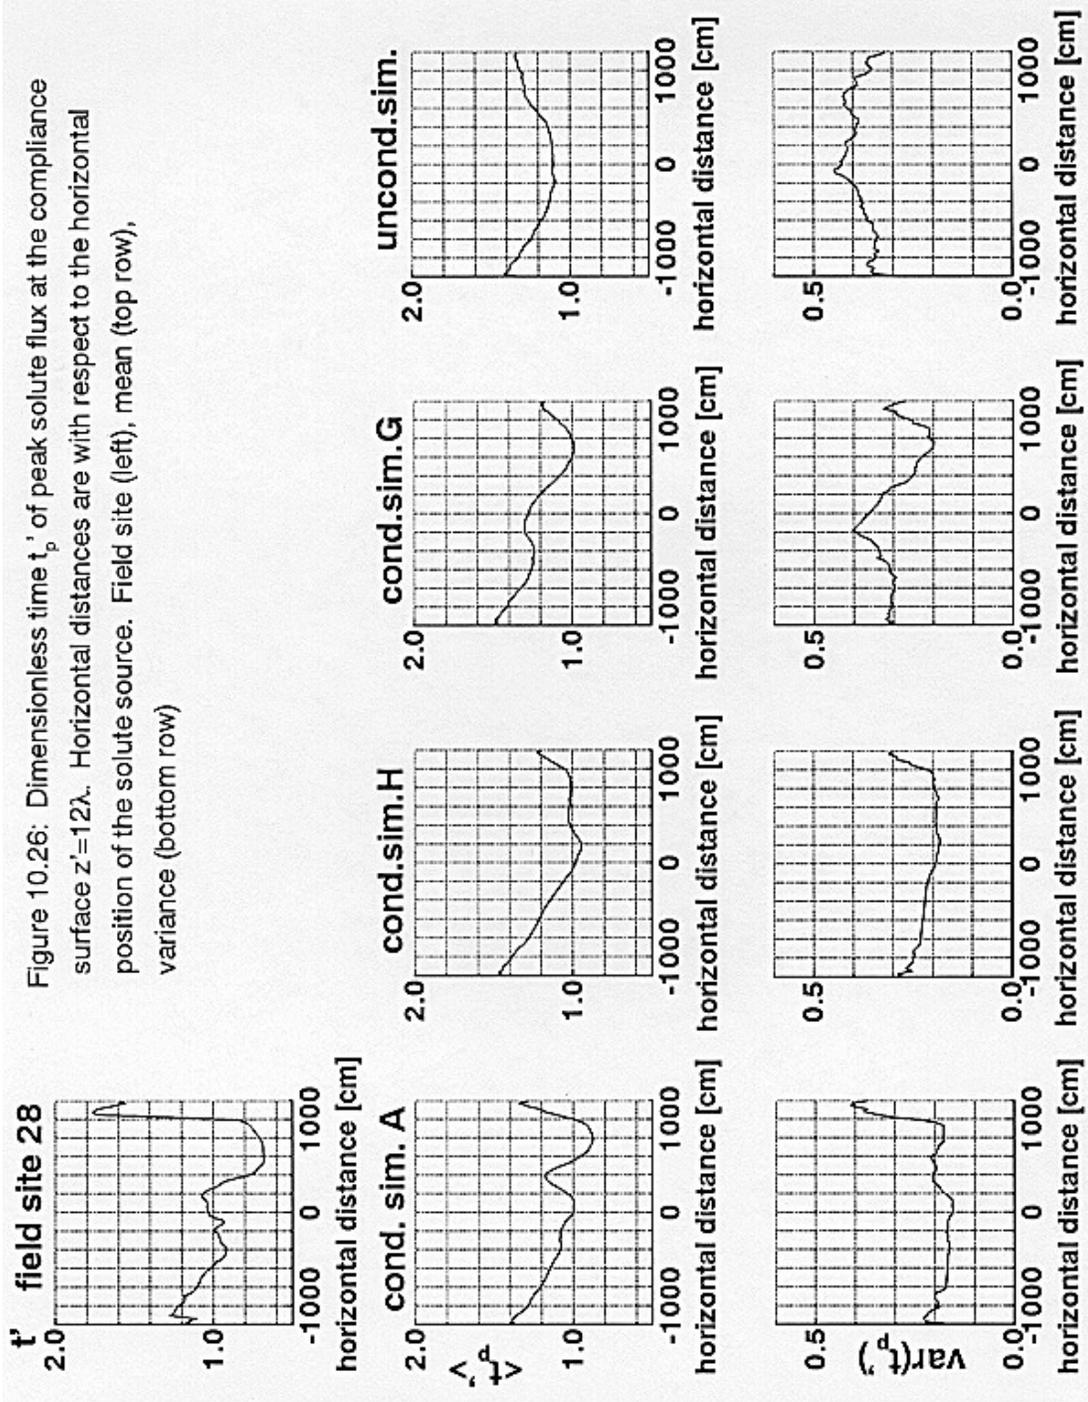


Figure 10.25 Average moment of inertia of individual plumes (a,b), variance of the plume center of mass (c,d), and moment of inertia of the mean plume (e,f) in the vertical (left column) and horizontal dimension (right column). All moments are normalized (') by $(\lambda_{iz} \sigma_f^2)$.



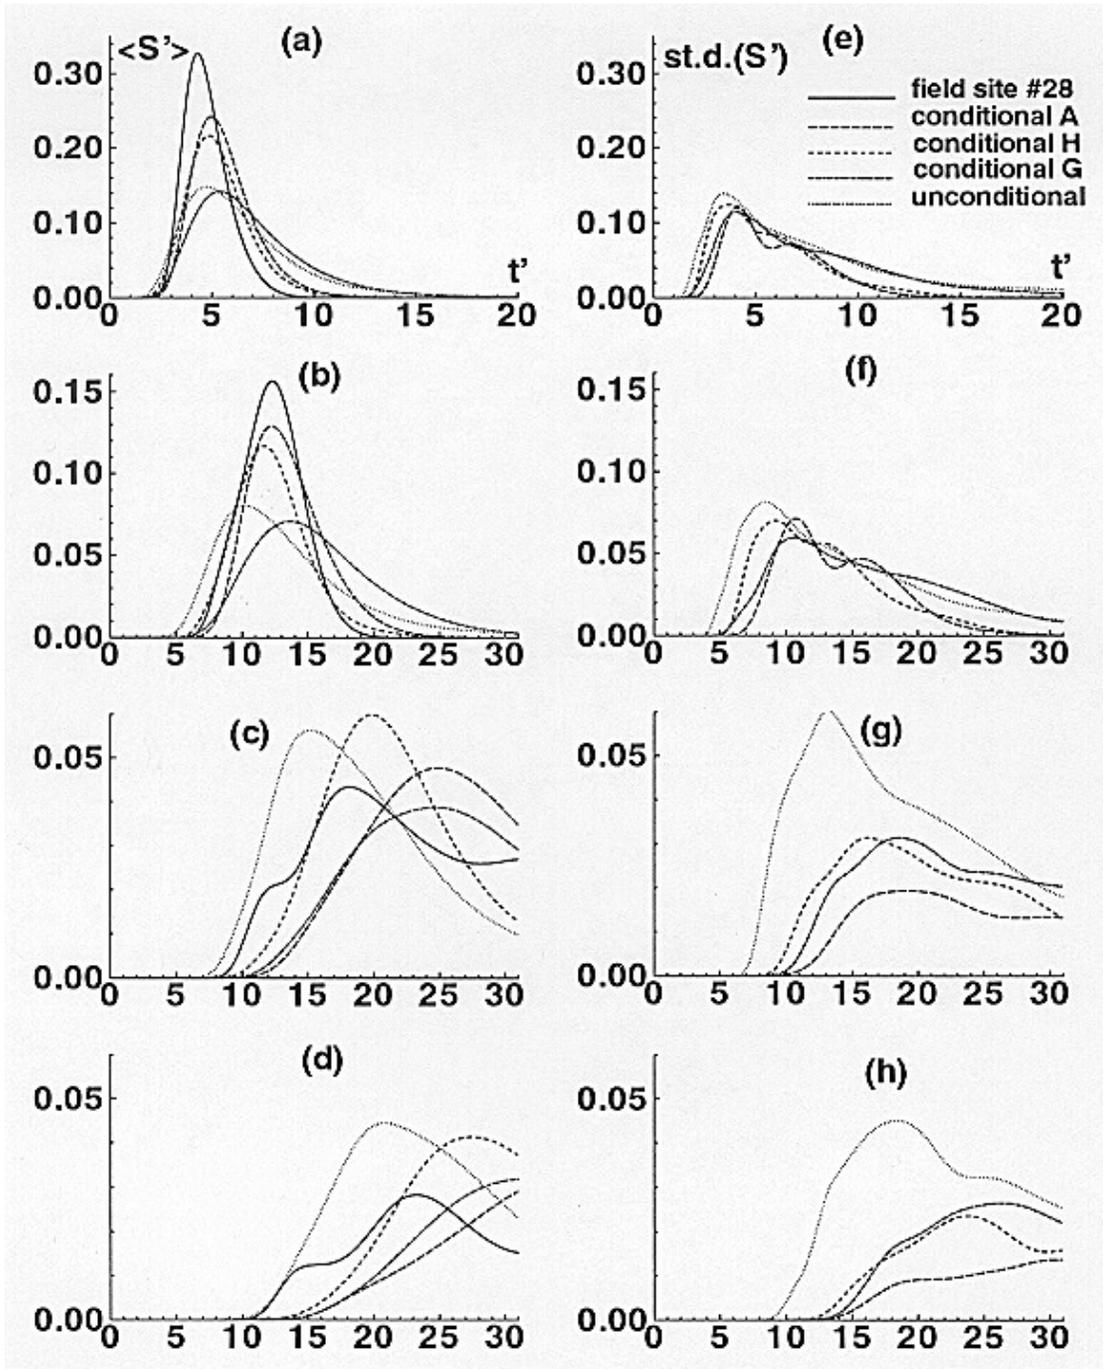


Figure 10.28 Normalized mean breakthrough curve $\langle S'(t') \rangle$, and standard dev. of the breakthrough curve $st.d.(S')$ at different depth: $5.4\lambda_{fz}$ (top row), $11.6\lambda_{fz}$ (second row), $17.8\lambda_{fz}$ (third row), and $23.8\lambda_{fz}$ (bottom row). The breakthrough curve for the field site is only plotted in the left column.

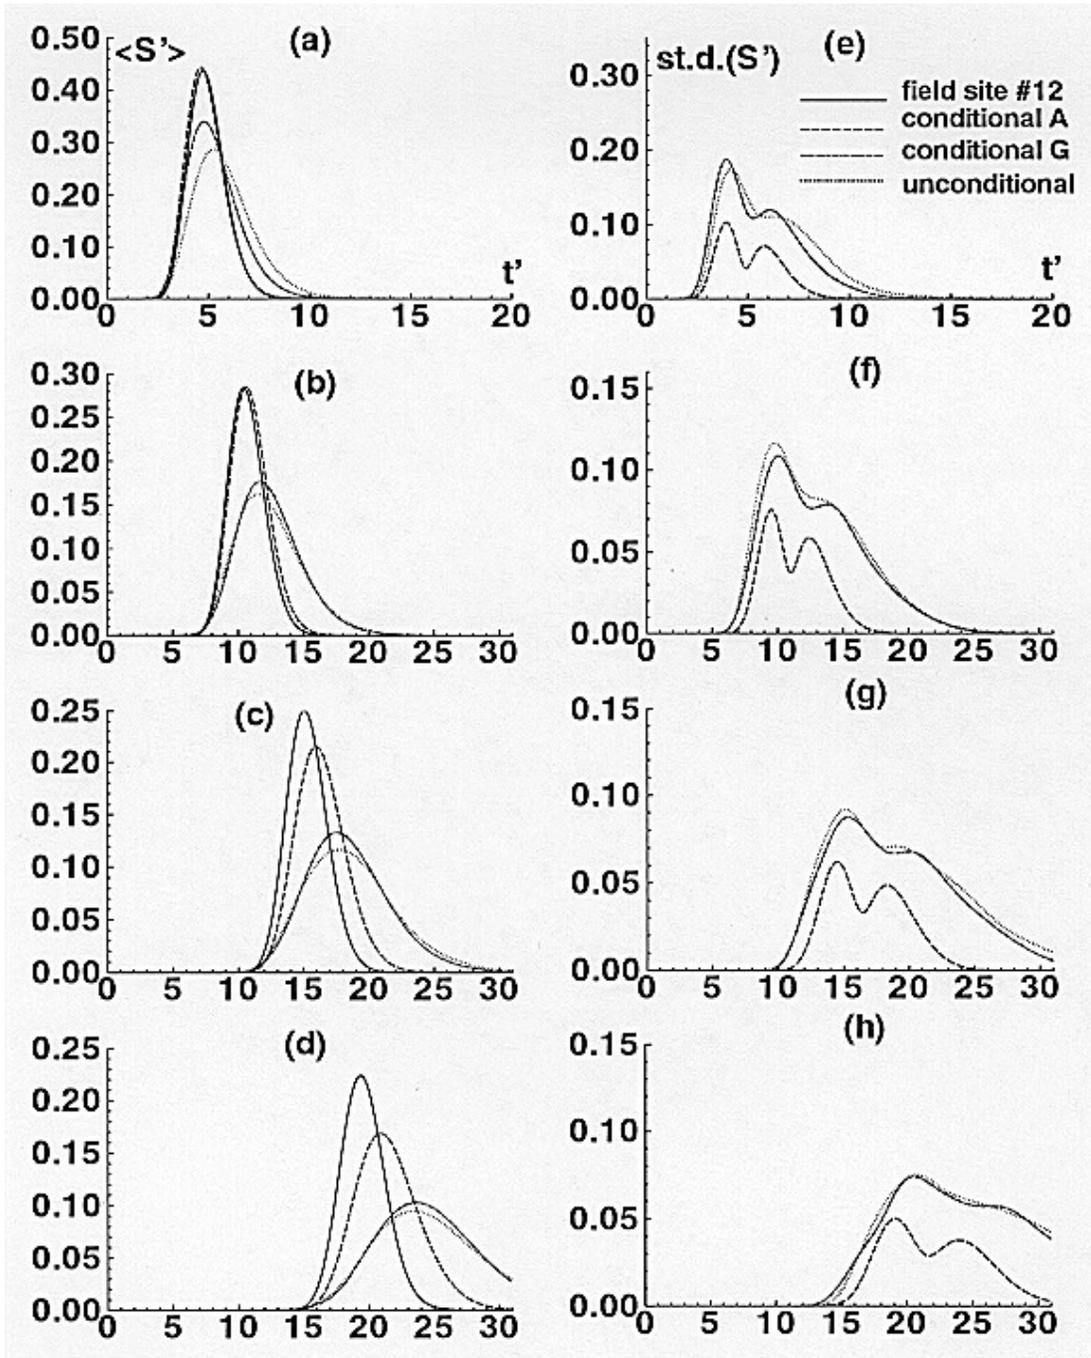


Figure 10.29 Normalized mean breakthrough curve $\langle S'(t') \rangle$, and standard dev. of the breakthrough curve $st.d.(S')$ at different depth: 5.4 λ_{fz} (top row), 11.6 λ_{fz} (second row), 17.8 λ_{fz} (third row), and 23.8 λ_{fz} (bottom row). The breakthrough curve for the field site is only plotted in the left column.

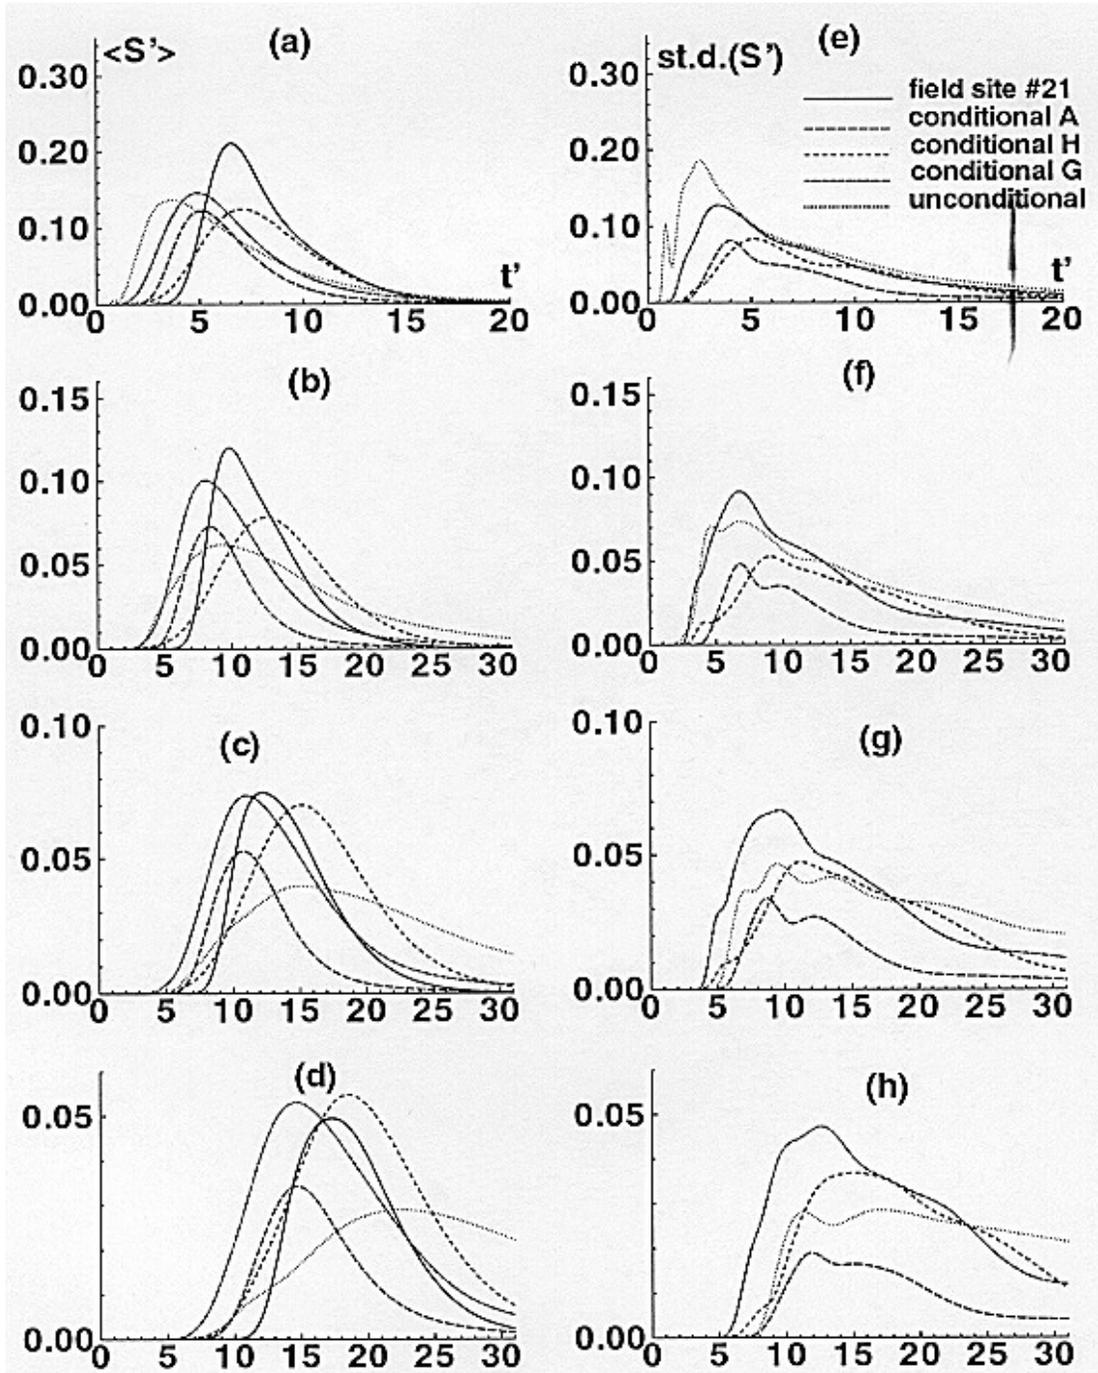


Figure 10.30 Normalized mean breakthrough curve $\langle S'(t') \rangle$, and standard dev. of the breakthrough curve $st.d.(S')$ at different depth: $5.4\lambda_{fz}$ (top row), $11.6\lambda_{fz}$ (second row), $17.8\lambda_{fz}$ (third row), and $23.8\lambda_{fz}$ (bottom row). The breakthrough curve for the field site is only plotted in the left column.

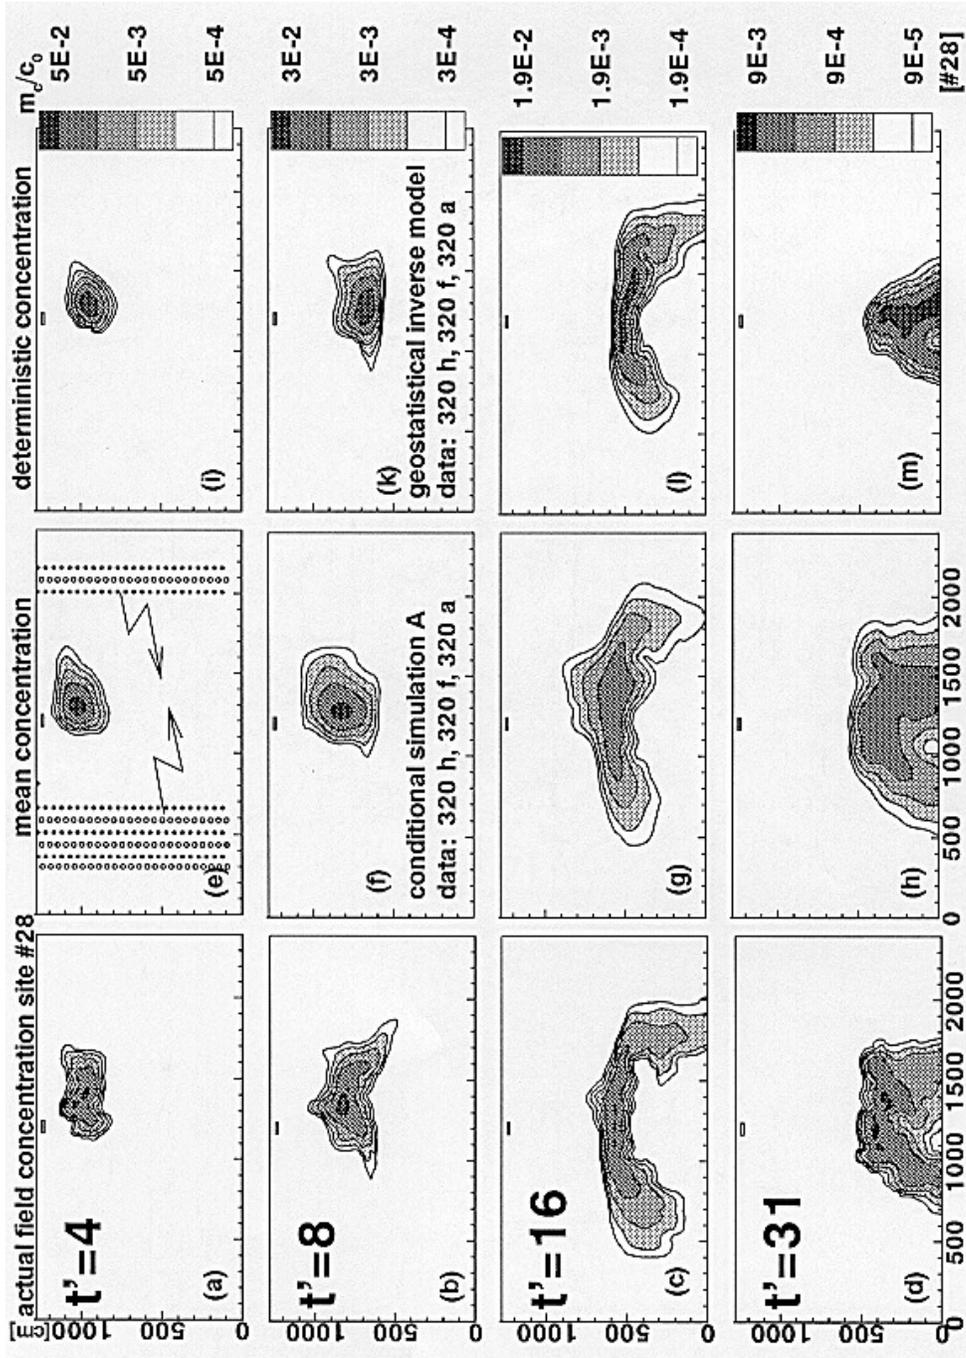


Figure 10.31