

# Simultaneous estimation of transmissivity values and zonation

Margaret J. Eppstein and David E. Dougherty

Research Center for Groundwater Remediation Design, Department of Civil and Environmental Engineering  
University of Vermont, Burlington

**Abstract.** The extended Kalman filter (EKF) has long been recognized as a powerful, yet computationally intensive, methodology for stochastic parameter estimation. Three improvements to traditional algorithms are presented and applied to heterogeneous transmissivity estimation. First, the costly EKF covariance updates are replaced by more efficient approximations. Second, the zonation structure of the distributed parameter field being estimated is dynamically determined and refined using a partitioning clustering algorithm. Third, a new method of merging first and second moments of random fields that have heterogeneous statistics is introduced. We apply this method, called random field union, as an alternative to conventional random field averaging for the systematic shrinking of covariance matrices as the dimensionality of the parameter space is reduced. The effects of these three improvements are examined. In applications to steady state groundwater flow test problems, we show that the first and second improvements reduce the computational time requirements dramatically, while the second and third can improve the accuracy and stability of the results. The resulting integrated method is successfully applied to a larger, more realistic calibration test case under steady and cyclostationary flow conditions (similar to regular seasonal fluctuations). When flow is steady, the method can be viewed as iterative; when flow is transient, the method is fully recursive.

## 1. Introduction

The usefulness of numerical groundwater models as a decision support tool for long-term remediation and management problems is dependent on the accuracy of the predictions. Prediction errors are introduced through several sources, for example, simplifications of the physics that are built into the mathematical model, errors in assessment of initial and forcing conditions, and errors in values of underlying transport parameters, such as transmissivity. Progress on parameter estimation has continued to lag behind advances in the other two areas. Reviews of groundwater parameter estimation methodologies are presented by *Yeh* [1986] and *Sun* [1994].

### 1.1. Kalman Filtering in Subsurface Hydrology

One method that has received sporadic attention in the groundwater modeling community is the Kalman filter (KF), which has potential applications to both forward and inverse groundwater modeling. The statistical elegance of the filter has prompted continued interest from researchers, while the computational complexity of its implementation has so far limited its usefulness in large-dimensionality problems.

In the 35 years since the Kalman filter was derived [*Kalman*, 1960] this powerful recursive estimating algorithm has been widely and successfully applied to problems in a variety of fields, especially in control systems engineering. The filter acts to automatically calibrate state variables in linear systems by optimally weighting the value of incoming time series measurements versus current state estimates, on the basis of their relative covariances. The covariance time and measurement updates are computationally intensive. The extended Kalman filter (EKF) is a nonlinear adaptation of the technique; the

state vector is augmented with uncertain parameters which are then recursively estimated along with the state. The state transitions, which may be nonlinear in the estimated parameters, are linearized about the current estimates using a first-order Taylor series approximation after each measurement update. EKF is even more computationally demanding than KF: not only are the covariance matrices larger, but the Jacobian (sensitivity), transition, and Kalman gain matrices must be recomputed after each measurement update because of the nonlinearities. A presentation of the EKF algorithm is given in section 2.1 of this paper. The reader is referred to *Gelb* [1974] for an excellent overview of the theory and example applications of both KF and EKF.

In the mid-1970s, researchers began exploring the applicability of Kalman filtering to groundwater modeling problems. *McLaughlin* [1976, 1978] identified four major areas in which Kalman filtering might prove useful in hydrogeologic applications: (1) application of the linear KF for improved descriptions of distributed state variables, (2) exploitation of the uncertainty estimates maintained by the recursive filter for real-time decision support for optimizing groundwater management procedures and/or (3) location of wells for monitoring networks, and (4) the potential use of the nonlinear EKF for hydrologic parameter estimation. *McLaughlin* [1976] was successful in implementing a 63-state linear KF to improve estimates of hydraulic head fluctuations in a transient flow problem but noted that the improvements in the head predictions caused by the measurement updates were short-lived. He recognized that long-term flow and transport predictions would be improved only through improved estimation of the underlying model parameters and that although this would theoretically be possible with EKF, it would not be computationally feasible for large-dimensionality systems unless efficient suboptimal approximations to the filter were developed. About the same time, *Kitanidis* [1976] also proposed use of

Copyright 1996 by the American Geophysical Union.

Paper number 96WR02283.  
0043-1397/96/96WR-02283\$09.00

EKF for hydrologic parameter estimation, and *Wilson et al.* [1978] clearly detailed how EKF can be applied to distributed transmissivity estimation based on head measurements. They applied this model to a synthetic two-dimensional transient flow domain and accurately estimated anisotropic, distributed transmissivities at 42 nodes, based on frequent head measurements at every node; initial state error, measurement error, and system noise were all set to zero. Like McLaughlin, these authors also noted the need for increased computational efficiency in the algorithm.

Despite 2 decades since these early investigations, subsequent publications regarding applications of Kalman filtering to estimation problems in subsurface hydrology have been surprisingly sparse. Some success has been achieved in using the KF to improve current state estimates of head [*Schmidtke et al.*, 1982; *Zhou et al.*, 1991; *Graham and Tankersley*, 1993; *Hantush and Mariño*, 1994] and contaminant concentrations [*Jinno et al.*, 1989; *Yu et al.*, 1989; *Graham and McLaughlin*, 1989, 1991; *Zou and Parr*, 1995]. The filter continues to show promise in the optimal design of monitoring networks [*Zhou et al.*, 1991; *Andricevic*, 1993] and in determining optimal pumping rates [*Andricevic and Kitanidis*, 1990; *Lee and Kitanidis*, 1991; *Andricevic*, 1993]. *Şen* [1984] used a linear filter to estimate storativity and transmissivity on the basis of drawdown data from a pumping test as a means to avoid type-curve matching.

Use of the nonlinear EKF for distributed hydrologic parameter estimation has continued to be thwarted by the computational demands of the algorithm. Most applications in this area tend to force low dimensionality of the parameter space (through assumptions of homogeneity or a priori zonation), utilize small synthetic domains, and have a large number of high-quality observations available for conditioning. *Schmidtke et al.* [1982] used EKF to estimate homogeneous permeability, fillable porosity, and recharge rate in a small two-dimensional transient flow domain based on filtered head measurements. *Van Geer and Van Der Kloet* [1985] estimated the single lumped ratio of transmissivity over storativity in a one-dimensional transient flow problem by filtering perfectly known head measurements from two observation wells; they found that a suboptimal version of the filter, in which head estimates were not conditioned on measurements, introduced little error into the transmissivity estimation. Similarly, *Jinno et al.* [1989] estimated one- and two-zone lumped values for dispersion and adsorption coefficients, velocity, and the Fourier coefficients of a series expansion for the process model in a one-dimensional transport problem by filtering simulated noisy concentration measurements. *Andricevic and Kitanidis* [1990] incorporated estimation of homogeneous transmissivity (conditioned on hydraulic head levels) and longitudinal dispersivity (conditioned on solute concentration levels) into an optimal control policy for determining pumping rates in a contaminated one-dimensional aquifer. This work was extended to two dimensions by *Lee and Kitanidis* [1991]; these authors estimated three uncorrelated transmissivity zones from noisy head and concentration measurements at seven wells in a transient flow domain with two contamination plumes.

*Graham and McLaughlin* [1989, 1991] have applied EKF to larger problems. *Graham and McLaughlin* [1989] proposed a "partitioned" filtering method for greater computational efficiency and tested it on synthetic two-dimensional examples. In the partitioned algorithm the two-dimensional steady state velocity field is conditioned on noise-free head and log conduc-

tivity measurements from a dynamically growing pool of observation wells, while the concentration estimates are simultaneously conditioned on noise-free concentration measurements. *Graham and McLaughlin* [1991] use real concentration data from the Borden tracer test site to improve estimation of the velocity and concentration fields; the computational burden of the EKF was reduced by employing a "dual-grid" approach in combination with efficient square root formulations.

Suboptimal approximations to the filter, such as the dual-grid approach [*Graham and McLaughlin*, 1991] and the elimination of state conditioning [*Van Geer and Van Der Kloet*, 1985] mentioned above, are computationally necessary for large-dimensionality problems. Research in suboptimal efficient implementations of KF and EKF has been especially active in the fields of meteorological data assimilation [see *Todling and Cohn*, 1994] and electrical engineering [e.g., *Gelb*, 1974; *Burl*, 1993; *Carlson*, 1990; *Chui et al.*, 1990; *Roy et al.*, 1991]. Dynamics simplification, reduced resolution, and local filter approximations have been employed.

## 1.2. Parameterization of Transmissivity

Although the true dimension of the spatial domain of transmissivity is infinite, inverse transmissivity estimation requires the modeler to limit the dimensionality of the parameter space. This is generally done by predetermining the parameter structure in one of three ways: (1) the transmissivity field is considered to be homogeneous [*Schmidtke et al.*, 1982; *Van Geer and Van Der Kloet*, 1985; *Jinno et al.*, 1989; *Andricevic and Kitanidis*, 1990], (2) heterogeneous transmissivities values are considered distributed over all the nodes or elements in the discretized domain [*Wilson et al.*, 1978], or (3) the transmissivity field is parameterized by a priori "zonation" of the transmissivity field into piecewise constant partitions [*Jinno et al.*, 1989; *Lee and Kitanidis*, 1991] or interpolated from a few basis nodes [*Yeh and Yoon*, 1981; *M. C. Hill et al.*, A controlled experiment in groundwater flow modeling, 1, Calibration using nonlinear regression to estimate parameter values, submitted to *Water Resources Research*, 1994, hereinafter referred to as *Hill et al.*, submitted manuscript, 1994]. Each of these three methods has its advantages and disadvantages. Assumptions of homogeneity are computationally most efficient but obviously prohibit the resolution of spatially heterogeneous transmissivity fields. Nodally (or elementally) distributed parameter systems have maximum flexibility in modeling heterogeneities but are computationally prohibitive for large domains. Furthermore, high dimensionality of the parameter space can lead to instability in the solution [*Yeh and Yoon*, 1981] unless accompanied by a similar complexity of observations (as was done by *Wilson et al.* [1978]), although this is not realistic for real-world modeling problems. Zonation methods attempt to strike a balance between these two extremes, but the modeler must make a priori assumptions on the number and shapes of zones despite insufficient geological data. *Yeh and Yoon* [1981] clearly demonstrate the critical need to properly dimension transmissivity fields. They show that while the output least squares objective function decreases with increasing dimensionality, the uncertainty of the transmissivity estimates increases, and that there is an optimum dimensionality associated with this trade-off that coincides with the minimum nodal transmissivity error. However, as was noted by *Sun et al.* [1995], this optimum dimension can never be exactly known for real domains, since the true observation errors and transmissivity covariances can

only be estimated. Moreover, even if the optimum transmissivity dimension can be determined, the question of the zone shapes and locations must still be resolved. Incorrect assumptions on parameter structure can lead to large errors in parameter estimates, and unless the structure and values are determined simultaneously, it is impossible to estimate either correctly [Sun and Yeh, 1985]. Sun and Yeh [1985] proposed a nested iterative combinatorial optimization approach to simultaneous structure and value estimation of two-dimensional transmissivities and applied it successfully to small synthetic domains. An alternative approach based on a linear functional strategy is suggested by R. S. Anderssen and coworkers [Anderssen and Dietrich, 1987; Anderssen and Chow, 1992] for transient flow domains. We were unable to locate references to any subsequent applications of these methods. Currently, state-of-the-art inverse models (e.g., the three-dimensional model recently proposed by Sun *et al.* [1995]) continue to impose some form of a priori transmissivity structure despite the recognized associated risks.

### 1.3. Contributions of This Paper

In this paper, we present contributions in three areas. First, we propose a fairly radical, yet demonstrably effective, approximation to the full-order EKF for improved computational efficiency. Second, starting with constant elemental transmissivity estimates, we alternate transmissivity value estimation with data-driven zonation (DDZ) using a partitioned cluster analysis methodology, thereby causing a progressive reduction in the number of zones. In this way we simultaneously resolve the transmissivity zonation structure and zone values, avoiding the problems of both overparameterization and a priori zonation, and further reduce the computational requirements of the filter. Third, we introduce an alternative method for merging heterogeneous random fields, a process necessitated by the dynamic merging of zones. We demonstrate the separate effects of each of these three new methodologies on four small synthetic test cases. Finally, the integrated algorithm is used to successfully estimate transmissivities in "Synthetic Valley," a larger, more realistic synthetic domain developed by the United States Geological Survey (USGS) for the purpose of comparing model calibration techniques (Hill *et al.*, submitted manuscript, 1994).

## 2. Parameter-Estimating Algorithm

The algorithm we employ for transmissivity estimation is based on the extended Kalman filter. We offer a brief overview of the basic EKF algorithm, and then describe our modifications to it.

### 2.1. Extended Kalman Filtering

The linear Kalman filter [Kalman, 1960] is the statistically optimal solution to the stochastic estimation of a random process based on measurements over time. Provided that the error covariances of the system states, process noise, and measurement noise are Gaussian and known, the KF can be shown to be an unbiased, consistent, minimum-variance filter for optimally weighting the effect of noisy, sparse, irregular time series measurements on estimates of state variables. The extended Kalman filter [Gelb, 1974] is a nonlinear adaptation of the KF in which the state vector is augmented with uncertain parameters to be estimated, and the covariance matrices are similarly augmented. The EKF algorithm can be divided into two major

steps: a time update, in which the augmented state error covariance matrix is evolved to the next time step, followed by a measurement update, in which the augmented state vector and the augmented state covariance matrix are adjusted in accordance with any available measurements for that time step. Wilson *et al.* [1978] have clearly detailed the required EKF updates for a system in which head measurements are used to estimate distributed transmissivity values; we briefly reiterate a similar presentation below.

**2.1.1. EKF time update.** Time advancement of the state vector of nodal heads ( $h$ ) from time  $t$  to time  $t + \delta t$  in a stochastic transient flow model can be expressed in state-space notation as follows:

$$h_{t+\delta t} = \Phi_t h_t + \Lambda_t f_t + w_t \quad (1)$$

where  $\Phi$  and  $\Lambda$  are transition matrices;  $f$  is the forcing function vector incorporating sources, sinks, and boundary conditions; and  $w_t$  is a random vector of zero mean, time-independent, system noise. The noise term accounts for model error due to a variety of sources, some of which may be spatially correlated [Van Geer *et al.*, 1991]. Note the nonlinearity that is introduced (even for confined, steady flow) by recursive distributed transmissivity estimation because the transition matrices are a function of the dynamic estimates of the stochastically modeled zonal log transmissivities  $y_t$ .

Initial estimates for the state covariance matrix  $\text{Cov}(h)$ , state-parameter cross-covariance matrix  $\text{Cov}(h, y)$ , and parameter covariance matrix  $\text{Cov}(y)$  must be provided. The state and state-parameter covariance matrices are advanced through time as follows:

$$\text{Cov}(h_{t+\delta t}, y_{t+\delta t}) \leftarrow \Phi_t \text{Cov}(h_t, y_t) + J_t \text{Cov}(y_t) \quad (2)$$

$$\begin{aligned} \text{Cov}(h_{t+\delta t}) &\leftarrow \Phi_t \text{Cov}(h_t) \Phi_t^T + J_t \text{Cov}(y_t) J_t^T \\ &+ J_t \text{Cov}(h_t, y_t)^T \Phi_t^T + \Phi_t \text{Cov}(h_t, y_t) J_t^T + Q \end{aligned} \quad (3)$$

where the time-varying Jacobian (sensitivity) matrix  $J_t$  is

$$J_t = \left. \frac{\partial}{\partial y} [\Phi_t h_t + \Lambda_t f_t] \right|_t \quad (4)$$

and  $Q$  is the time-independent covariance matrix of the system noise  $w_t$ . The leftward arrows are used to denote replacement; this notation simplifies the presentation of EKF by eliminating the need for cumbersome premeasurement and postmeasurement update symbols and implies sequentiality in the equations. Equations (2)–(3) (and possibly (4), if flow is nonlinear or forcing conditions are transient) must be performed during each time step of the simulation, whether or not measurements are available for that time step.

The fully explicit state-space formulation in (1) is a convenient way to express the time advancement mathematically but is rarely the form in which the actual computation is performed. Typically, fully or partially implicit schemes are employed to improve the stability of the results.

**2.1.2. EKF measurement update.** During time steps in which state measurements are available, the parameter and state vectors, and their associated covariance matrices, incorporate this new information in an optimal manner. The Kalman gain matrices ( $K^y$  and  $K^h$ ), which weight the relative uncertainties of incoming measurements against the uncertainty of the current state estimate, are computed as follows:

$$K_{t+\delta t}^y \leftarrow \text{Cov}(h_{t+\delta t}, y_{t+\delta t})^T M_{t+\delta t}^T \cdot [R_{t+\delta t} + M_{t+\delta t} \text{Cov}(h_{t+\delta t}) M_{t+\delta t}^T]^{-1} \quad (5)$$

$$K_{t+\delta t}^h \leftarrow \text{Cov}(h_{t+\delta t}) M_{t+\delta t}^T [R_{t+\delta t} + M_{t+\delta t} \text{Cov}(h_{t+\delta t}) M_{t+\delta t}^T]^{-1} \quad (6)$$

where  $M$  is the observation matrix and  $R$  is the covariance matrix of the measurement white noise.

Updates of the parameter and state vectors are then performed using incoming (probably sparse and intermittent) state measurements  $Z_t$ ; the influence of these measurements on the current estimates is optimally weighted (in a Bayesian sense) via the Kalman gain matrices:

$$y_{t+\delta t} \leftarrow y_t + K_{t+\delta t}^y (Z_{t+\delta t} - M_{t+\delta t} h_{t+\delta t}) \quad (7)$$

$$h_{t+\delta t} \leftarrow h_{t+\delta t} + K_{t+\delta t}^h (Z_{t+\delta t} - M_{t+\delta t} h_{t+\delta t}) \quad (8)$$

The reduction in uncertainty due to the measurement update is recorded in the parameter, state, and cross-covariance matrices as follows:

$$\text{Cov}(y_{t+\delta t}) \leftarrow \text{Cov}(y_t) - K_{t+\delta t}^y M_{t+\delta t} \text{Cov}(h_{t+\delta t}, y_{t+\delta t}) \quad (9)$$

$$\text{Cov}(h_{t+\delta t}) \leftarrow (I - K_{t+\delta t}^h M_{t+\delta t}) \text{Cov}(h_{t+\delta t}) \quad (10)$$

$$\text{Cov}(h_{t+\delta t}, y_{t+\delta t}) \leftarrow (I - K_{t+\delta t}^h M_{t+\delta t}) \text{Cov}(h_{t+\delta t}, y_{t+\delta t}) \quad (11)$$

## 2.2. Approximate Extended Kalman Filtering

Unfortunately, the statistical simplicity of the EKF does not translate into computational simplicity; the nonlinearities and very large covariance matrices inherent in EKF make it computationally prohibitive on large-dimensionality systems, such as useful field-scale groundwater flow and transport simulations. Although EKF is frequently referred to as “optimal,” it is actually “suboptimal” owing to required linearizations and the realities of uncertain statistics and simplifications in the process model. This observation has led various researchers to search for other, more efficient, suboptimal approximations to the filter which still yield useful results. *Dee* [1991] has proposed dynamics simplification for improved efficiency in the linear KF as applied to large systems. We take this concept further for nonlinear EKF and eliminate the explicit use of the transition matrix from the EKF covariance updates altogether. Our approximate EKF (dubbed AEKF) is detailed below.

**2.2.1. AEKF time estimates.** In full EKF the state prediction error is assumed to be a function of the current state error, the error in stochastically modeled parameter estimates, and the system noise, which includes errors due to deterministically modeled parameters (equation (3)). If the effects of errors in parameter estimates and system noise overwhelm the residual effects of the current state error, then the state covariance matrix  $\text{Cov}(h)$  and the state-parameter cross-covariance matrix  $\text{Cov}(h, y)$  can be estimated as needed, rather than updated each time step as in full EKF. AEKF is based on this assumption; the degree of suboptimality of this approximation to EKF will depend on the level to which the assumption is not satisfied. For steady state groundwater flow models in which boundary conditions are known, the underlying transmissivity field drives the system to a much larger extent than current head estimates, especially if the time step is reasonably large, so AEKF is a reasonable approximation for such systems.

In AEKF, equations (2) and (3) are replaced by the follow-

ing simpler approximations, in which all terms (except  $J_t$ ) containing the transition matrix  $\Phi$  have been eliminated:

$$\text{Cov}(h_{t+\delta t}, y_{t+\delta t}) \approx J_t \text{Cov}(y_t) \quad (12)$$

$$\text{Cov}(h_{t+\delta t}) \approx \text{Cov}(h_{t+\delta t}, y_{t+\delta t}) J_t^T + Q \quad (13)$$

Although  $\text{Cov}(y)$  still requires a prior estimate,  $\text{Cov}(h)$  and  $\text{Cov}(h, y)$  do not. Equations (12) and (13) are first-order approximations with model error included.

The computationally intensive calculation of the Jacobian (equation (4)), as well as equations (12) and (13), need to be performed only during time steps for which measurements are available. Hence although AEKF is still a recursive calibration and estimation algorithm (in the sense that time series information is incorporated sequentially into the model), there is no additional computational burden added to the simulator for time steps for which no measurements are available, in distinct contrast to full EKF.

**2.2.2. AEKF measurement update.** In the AEKF algorithm, equations (5) and (6) are unchanged but use the approximations for  $\text{Cov}(h, y)$  and  $\text{Cov}(h)$  from equations (12) and (13). Equations (10) and (11) are eliminated; only equation (9) is required for the measurement update of covariances, since only the  $\text{Cov}(y)$  covariance matrix is maintained between time steps. The parameter vector  $y$  and state vector  $h$  are updated by equations (7) and (8) in AEKF exactly as in EKF, but using the approximated Kalman gains from (5) and (6) (based on the approximations (12) and (13)).

**2.2.3. Summary of AEKF equations.** The computations required by EKF are expressed in equations (2)–(11). In AEKF, these are reduced to equations (4)–(9) and (12)–(13). For the reader’s convenience, the AEKF equations are repeated in summary form here.

The AEKF covariance matrix approximations for each time step in which measurements are available are equations (4), (12), and (13):

$$J_t = \frac{\partial}{\partial y} [\Phi h_t + \Lambda f_t] \Big|_t$$

$$\text{Cov}(h_{t+\delta t}, y_{t+\delta t}) \approx J_t \text{Cov}(y_t)$$

$$\text{Cov}(h_{t+\delta t}) \approx \text{Cov}(h_{t+\delta t}, y_{t+\delta t}) J_t^T + Q$$

The AEKF measurement updates of states, parameters, and covariance matrices are equations (5), (6), (7), (8), and (9):

$$K_{t+\delta t}^y \leftarrow \text{Cov}(h_{t+\delta t}, y_{t+\delta t})^T M_{t+\delta t}^T$$

$$\cdot [R_{t+\delta t} + M_{t+\delta t} \text{Cov}(h_{t+\delta t}) M_{t+\delta t}^T]^{-1}$$

$$K_{t+\delta t}^h \leftarrow \text{Cov}(h_{t+\delta t}) M_{t+\delta t}^T [R_{t+\delta t} + M_{t+\delta t} \text{Cov}(h_{t+\delta t}) M_{t+\delta t}^T]^{-1}$$

$$y_{t+\delta t} \leftarrow y_t + K_{t+\delta t}^y (Z_{t+\delta t} - M_{t+\delta t} h_{t+\delta t})$$

$$h_{t+\delta t} \leftarrow h_{t+\delta t} + K_{t+\delta t}^h (Z_{t+\delta t} - M_{t+\delta t} h_{t+\delta t})$$

$$\text{Cov}(y_{t+\delta t}) \leftarrow \text{Cov}(y_t) - K_{t+\delta t}^y M_{t+\delta t} \text{Cov}(h_{t+\delta t}, y_{t+\delta t})$$

## 3. Parameter Structure Detection Algorithm

In order to avoid incorrect a priori zonation, our algorithm starts conservatively by assuming each element in the finite element mesh (or each cell in a finite difference model) is a

distinct log transmissivity zone. These small zones are then recursively merged into larger zones using a process we refer to as data-driven zonation. Some results from an early implementation of this idea are presented by *Eppstein and Dougherty* [1994].

During time steps in which measurements are available, the estimates of zonal log transmissivities are updated by some parameter estimation method (we use EKF or AEKF). On the basis of the updated zone values, a cluster analysis is performed on the zones. For each cluster of zones thus determined, adjacent zones in the same cluster are merged. The statistics of the log transmissivities of the new zone (mean, variance, and covariance with other zones) are estimated by systematically merging the statistics of the clustered zones. In the current implementation, zones are never subdivided; the dimension of the parameter space is strictly nonincreasing.

Three choices must be made in implementing DDZ: (1) which clustering algorithm to implement, (2) what clustering criteria to utilize, and (3) how to systematically merge the statistics of clustered zones.

### 3.1. Clustering Algorithm

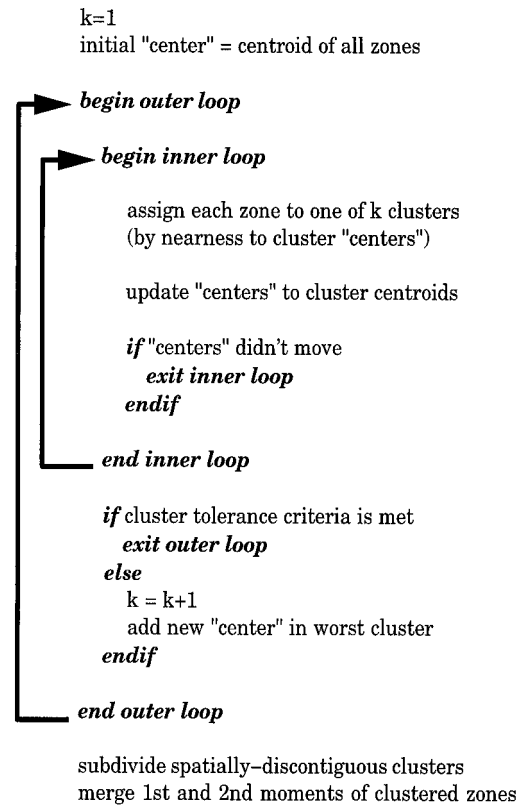
We employ an iterative partitional clustering algorithm [*Jain and Dubes*, 1988] in order to parsimoniously select the zonation structure with the minimum number of zones that meets our stopping criterion. The algorithm works by iteratively applying a standard  $k$  means clustering algorithm (which is itself iterative); the number of clusters  $k$  is increased each iteration until the cluster tolerance is met within all clusters. Physically discontinuous clusters are subsequently partitioned into smaller continuous clusters. The algorithm is depicted in Figure 1.

### 3.2. Clustering Criteria

In this study, clustering is performed along a single dimension: zone transmissivity (not zone log transmissivity). Cluster centroids are trivially calculated as the arithmetic means of the transmissivities of the member zones. Zones are thus clustered solely on the basis of similarity of the current zonal transmissivity estimates. Spatial contiguity of clusters is enforced in a postprocessing manner, as shown in Figure 1.

Cluster tolerance is defined as 5% of the maximum distance between all cluster centroids in the current estimate. The stopping criterion for the partitional clustering algorithm requires that zone transmissivities of all members of each cluster must be within their cluster centroid  $\pm$  the cluster tolerance. The resultant clusters are validated using the modified Hubert  $\Gamma$  statistic [*Dubes*, 1987], which is a measure of correlation between the matrix of interzone Euclidean distances and intercentroid Euclidean distances of the zone clusters. We require this correlation to be greater than 0.95 to ensure our stopping criterion is conservative, favoring too many clusters to too few. Basing the cluster criterion on zone transmissivities, rather than zone log transmissivities, was found to be more effective because it resulted in more liberal zonation at the low end of the log transmissivity range.

In all of our reported experiments, the transmissivity estimates for all elemental zones were initialized to the same a priori "average" estimate of transmissivity. The use of a constant relative cluster tolerance measure thereby results in an absolute cluster tolerance measure that increases asymptotically as the transmissivity estimates in the domain diverge owing to incoming measurements. Hence more liberal zonation



**Figure 1.** The clustering algorithm used in the current implementation of data-driven zonation (DDZ). The outer loop represents the partitional aspect of the algorithm in which the number of clusters  $k$  is increased as needed; the inner loop represents the clustering about  $k$  means.

is allowed (to a point) as the information content of the model increases, when errors in zonation are presumably less likely to occur.

The simple cluster criterion adopted in this study was largely due to the simplistic framework of the artificial test domains employed. In a real domain, other types of available data (such as soil type) could be used to help determine cluster membership. In addition, spatial relationships between zones and clusters could be exploited to a much greater degree than what we have currently implemented. For example, one could introduce anisotropy into the tolerance, since transmissivity differences transverse to the primary direction of flow are detected more easily than those that are parallel to the flow. Similarly, one could favor less complicated zone shapes. Although we have not yet exploited these options, the use of cluster analysis provides a natural mechanism for incorporating a wide variety of types of hard and soft data into the recursive zonation process.

### 3.3. Merging Statistics of Clustered Zones

After the zones have been partitioned into clusters, the member zones of each cluster are merged. Each log transmissivity zone is being modeled as a two-dimensional Gaussian stochastic process, or random field. The question of how to merge the first and second moments of two (or more) random fields, in order to estimate the statistics of the merged zone, is nontrivial. In this paper we examine two methods: (1) the traditional method known as random field averaging (RFA),

which is based on an assumption of statistical homogeneity among zones, and (2) a new method we refer to as random field union (RFU), which allows for statistical heterogeneity.

**3.3.1. Random field averaging.** The well-known random field averaging method [Vanmarcke, 1983] is based on integrating a statistically homogeneous random function over its two space dimensions and dividing by the area. In the specific case of merging two log transmissivity zones (modeled by the random fields  $X$  and  $Y$ , of areas  $n$  and  $m$ , and correlated with a third normally distributed random field  $Z$ ) into a single new random field ( $W$ ), this reduces to an area-weighted averaged of  $X$  and  $Y$ , as follows:

$$W \triangleq \frac{nX + mY}{n + m} \quad (14)$$

$$\hat{\mu}_W = \frac{n\hat{\mu}_X + m\hat{\mu}_Y}{n + m} \quad (15)$$

$$\hat{\sigma}_{ZW} = \frac{n\hat{\sigma}_{ZX} + m\hat{\sigma}_{ZY}}{n + m} \quad (16)$$

$$\hat{\sigma}_W^2 = \frac{n^2\hat{\sigma}_X^2 + m^2\hat{\sigma}_Y^2 + 2nm\hat{\sigma}_{XY}}{(n + m)^2} \quad (17)$$

Note that by RFA, the estimated variance of the merged zone  $\hat{\sigma}_W^2$  is always less than (or equal to, if  $X = Y$ ) the area-weighted average of the variance estimates for the component zones  $\hat{\sigma}_X^2$  and  $\hat{\sigma}_Y^2$ ; it is dependent on the cross correlation of the zones  $\hat{\sigma}_{XY}$  but is independent of the mean estimates  $\hat{\mu}_X$  and  $\hat{\mu}_Y$ . This makes sense if  $X$  and  $Y$  are assumed to be local average processes of the statistically homogeneous random field  $W$ ; since variance is treated as uncertainty, it is intuitive that this uncertainty should decrease as the sample size increases.

**3.3.2. Random field union.** The assumption of statistical homogeneity is not met by the log transmissivity zones to be merged according to our data-driven zonation algorithm. In DDZ we combine two (or more) log transmissivity zones having similar yet distinct estimates that may reflect real differences in physical properties. The zones are combined for purposes of computational efficiency and increased sensitivity despite potential statistical heterogeneity. Intuitively, when merging zones that have identical mean estimates, we expect the combined variance estimate to be the weighted average of the component variance estimates. In contrast, when the mean estimates of the component fields are different, we expect that the variance of the merged field should similarly increase. In other words, we wish to view variance as a measure of variability, as opposed to uncertainty, in this context. We define a new method for merging statistics of statistically heterogeneous random fields, called random field union, which exhibits these properties.

It is well known from estimation theory [Stark and Woods, 1986] that given  $k$  independent and identically distributed (i.i.d.) observations of a normally distributed random variable  $W$ , the consistent, maximum-likelihood estimators for the mean  $\hat{\mu}_W$ , covariance  $\hat{\sigma}_{WZ}$  (with another normally distributed random variable  $Z$  with estimated mean  $\hat{\mu}_Z$ ), and variance  $\hat{\sigma}_W^2$  are given by

$$\hat{\mu}_W = \frac{1}{k} \sum_{i=1}^k W_i \quad (18)$$

$$\hat{\sigma}_{WZ} = \frac{1}{k} \sum_{i=1}^k (W_i - \hat{\mu}_W)(Z_i - \hat{\mu}_Z) \quad (19)$$

$$\hat{\sigma}_W^2 = \frac{1}{k} \sum_{i=1}^k (W_i - \hat{\mu}_W)^2 \quad (20)$$

Equations (19) and (20) are slightly biased, although the bias tends to 0 as  $k \rightarrow \infty$ .

The RFU of two (or more) fields is an area-weighted union of the fields. The statistics of the combined fields can be estimated by simulating a large number of samples from the combined field and applying the maximum-likelihood estimators of equations (18)–(20), as follows.

Assume that  $W$  is a random field that is partitioned into two normally distributed but statistically heterogeneous subfields  $X$  and  $Y$  (with areas  $n$  and  $m$ ) for which we have estimated moments. Further, assume that we have  $j$  observations per unit area of  $W$ , which are i.i.d. within each of the subfields  $X$  and  $Y$ , and assume that  $X$  and  $Y$  are correlated with a third normally distributed random field  $Z$  with mean estimate  $\hat{\mu}_Z$ . The distribution of  $W$  can be approximated by a normal distribution with the following statistics:

$$W \triangleq X \cup Y \quad (21)$$

$$\begin{aligned} \hat{\mu}_W &= E \left[ \frac{1}{j(n + m)} \left( \sum_{i=1}^{jn} X_i + \sum_{i=1}^{jm} Y_i \right) \right] \\ &= \frac{n\hat{\mu}_X + m\hat{\mu}_Y}{n + m} \end{aligned} \quad (22)$$

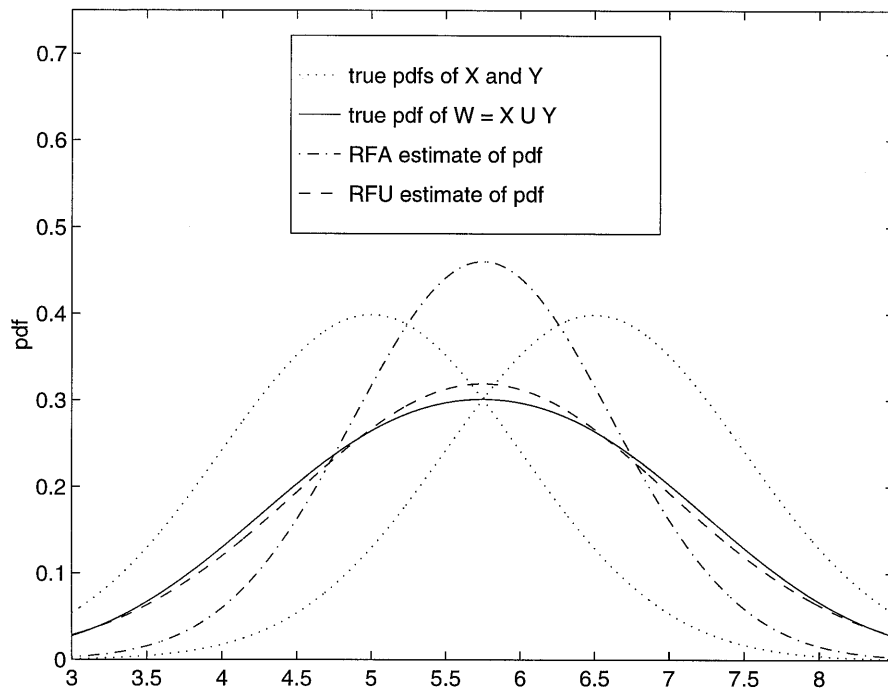
$$\begin{aligned} \hat{\sigma}_{ZW} &= E \left[ \frac{1}{j(n + m)} \left( \sum_{i=1}^{jn} (X_i - \hat{\mu}_W)(Z_i - \hat{\mu}_Z) \right. \right. \\ &\quad \left. \left. + \sum_{i=1}^{jm} (Y_i - \hat{\mu}_W)(Z_i - \hat{\mu}_Z) \right) \right] \\ &= \frac{n\hat{\sigma}_{ZX} + m\hat{\sigma}_{ZY}}{n + m} \end{aligned} \quad (23)$$

$$\begin{aligned} \hat{\sigma}_W^2 &= E \left[ \frac{1}{j(n + m)} \left( \sum_{i=1}^{jn} (X_i - \hat{\mu}_W)^2 + \sum_{i=1}^{jm} (Y_i - \hat{\mu}_W)^2 \right) \right] \\ &= \frac{n\hat{\sigma}_X^2 + m\hat{\sigma}_Y^2}{n + m} + \frac{n(\hat{\mu}_W - \hat{\mu}_X)^2 + m(\hat{\mu}_W - \hat{\mu}_Y)^2}{n + m} \end{aligned} \quad (24)$$

where the symbol  $E$  is used to denote expected value. Equations (21)–(24) are trivially extended to the union of more than two zones by adding an additional term into the summations for each zone being merged.

The derived formulas for estimators for both the merged mean  $\hat{\mu}_W$  (equations (15) and (22)) and for the covariance of the merged field with another field  $\hat{\sigma}_{ZW}$  (equations (16) and (23)), are identical for both RFA and RFU.

The estimator for the merged variance  $\hat{\sigma}_W^2$  differs markedly between RFA and RFU (compare (17) and (24)). Consider the two terms of (24). The first term is simply the area-weighted average of the component variances. The second term takes into account the disparity between the means of the compo-



**Figure 2.** Probability density functions (pdfs) of merge of two heterogeneous random fields as estimated by random field averaging (RFA) and random field union (RFU).

nent fields; the greater this difference, the more the variance will increase. Thus in RFU the variance of a merged field is always greater than (or equal to, if  $\hat{\mu}_X = \hat{\mu}_Y$ ) the area-weighted average of the component variances, in direct contrast to the lower variance estimate by RFA. The variance estimate by RFU, though dependent on the means of the merged zones, is independent of the cross correlation between the merged zones because it treats variance as a measure of variability of the statistically heterogeneous media, as opposed to uncertainty of moment estimates of a statistically homogeneous medium.

The difference between RFA and RFU is easily visualized with an example. Suppose we merge two normally distributed random fields  $X$  and  $Y$ , each with unit area and unit variance but with means  $\mu_X = 5.0$  and  $\mu_Y = 6.5$  and covariance  $\sigma_{XY} = 0.5$ . The probability density function (pdf) for the combined field  $W$  can be determined exactly by summing the pdfs of  $X$  and  $Y$  and dividing by 2. The true pdfs of  $X$ ,  $Y$ , and  $W$  are shown in Figure 2, along with estimated pdfs assuming a normal distribution of  $W$  and using the estimates for mean and variance from both RFA and RFU. Both RFA and RFU accurately estimate  $\mu_W$ , but the variance estimates differ significantly. By RFA,  $\hat{\sigma}_W^2 = 0.75$ , which can be seen to be low. By RFU,  $\hat{\sigma}_W^2 = 1.6$ , which agrees more closely with the true  $\sigma_W^2$ . The slight error evident in the normal RFU estimate shown in Figure 2 is due to the slight nonnormality of the true pdf of  $W$  caused by the relatively large difference between  $\mu_X$  and  $\mu_Y$ . When  $\mu_X$  and  $\mu_Y$  are reasonably close to each other (as in zones to be merged) the pdf of the union  $W$  is approximately normal; by the central limit theorem, this tendency toward normality will increase with the number of zones merged. By RFA,  $X$  and  $Y$  are taken to be biased estimates of the larger homogeneous field  $W$ ; by RFU,  $X$  and  $Y$  are treated as unbiased estimates of subfields of the heterogeneous field  $W$ .

In DDZ we apply RFU (and, for comparison, RFA) for the

purpose of shrinking the parameter covariance matrix  $\text{Cov}(y)$  as the dimensionality of the parameter space is reduced through merging zones in the same cluster. In this way, information accrued in the covariance matrix as a result of the filtering process is not lost as a result of the zonation process.


## 4. Experiments

### 4.1. Numerical Model

Four estimation algorithms were compared on four specified reference problems, with known solutions, designed to test the performance of the methods under well-controlled conditions. On the basis of simulated head measurements from six regularly spaced observation wells, the two-dimensional heterogeneous transmissivity fields were estimated.

In order to isolate the contribution made by each of the three major improvements we offer to traditional extended Kalman filtering, experiments using four combinations are presented: full extended Kalman filtering (EKF), approximate extended Kalman filtering (AEKF), AEKF with data-driven zonation using random field averaging (AEKF, DDZ, RFA), and AEKF with data-driven zonation using random field union (AEKF, DDZ, RFU). The abbreviations defined above are used to identify the methods in several figure legends and captions.

All four reference problems are steady flow in an isotropic confined aquifer. Although contaminant transport was also modeled (for inclusion in the computational requirements), it was not used in the model calibration. Flow occurs in a 100-m-square domain, with  $h = 10$  m specified at one side,  $h = 0$  m specified on the opposing side, and no normal flow conditions on the other two sides. A Galerkin finite element method [Wang and Anderson, 1982] (using 100 bilinear square elements, 10 in each horizontal dimension) with fully implicit (backward) Euler time stepping ( $\delta t = 200$  days) and a direct



**Figure 3.** True elemental transmissivities of the four 100-element test cases: (a) z2, (b) ramp, (c) z3hml, and (d) z3mhl. The dark gray represents a transmissivity of 1 m/d, the medium gray represents 5.5 m/d, and the white represents 10 m/d.

solver was implemented in Matlab (version 4.2c.1, The Math Works, Natick, Massachusetts, 1994) to model the problems. As a measure of computational effort, Matlab's "flops" command was utilized to count the number of floating point operations used in each simulation.

#### 4.2. Sample Transmissivity Domains

Transmissivities were considered to be constant within each element in the finite element mesh, but could vary between elements. The four test domains differed only in the underlying transmissivity fields. Descriptions of these four fields (and their identifying abbreviations) are as follows: (1) z2, two rectangular transmissivity zones (10 m/d and 1 m/d), with zone boundary transverse to the primary direction of flow (Figure 3a), (2) ramp, a 10-zone "ramp" field (10 m/d down to 1 m/d) of transmissivities, which decreases in the primary direction of flow (Figure 3b), (3) z3hml, a three-zone (10 m/d to 5.5 m/d and 1 m/d) field with boundaries transverse, parallel, and diagonal to the primary direction of flow (Figure 3c), and (4) z3mhl, a three-zone (5.5 m/d to 10 m/d and 1 m/d) field with boundaries transverse, parallel, and diagonal to the primary direction of flow (Figure 3d). These four transmissivity fields yield four head fields with subtle but detectable differences.

#### 4.3. Parameter and Covariance Initialization

Since the intent of this study is to compare the behavior of algorithmic variations with the basic EKF method, as opposed to examining the performance of EKF itself under uncertainty, measurements at the six observation wells were assumed to be perfect; hence the measurement error covariance matrix  $R$  was set to zeros. To apply the EKF and AEKF parameter estima-

tion methods, the system head noise covariance matrix  $Q$  was assumed uncorrelated in time, but exponentially correlated in space with a correlation length of four elements and a constant variance of 0.2 (variance of system noise on known heads was set to 0). For full EKF the state covariance matrix  $Cov(h)$  was initialized to be  $2Q$  and the state-parameter cross-covariance matrix  $Cov(h, y)$  was initialized to zeros. For all methods the log transmissivity covariance matrix  $Cov(y)$  was based on an exponential model with a correlation length of 10 elements and an initial variance of 1.0. Although the four test domains do not exhibit an exponential log transmissivity correlation structure, our intent was to see how well the algorithm could recover the true structure despite this incorrect (but commonly applied) assumption. For all estimation experiments on these four domains, elemental transmissivities were initialized to 5.5 m/d for all elements in the finite element mesh. Storativity  $S$  is 0.001 and is considered known.

All simulations were run for 15 time steps, with perfect observations from all six observation wells available at each time step. Since the measurements are static for these test problems and there is no measurement error, one could consider these as 15 iterations; in the more general framework of these recursive estimators, this is not the case.

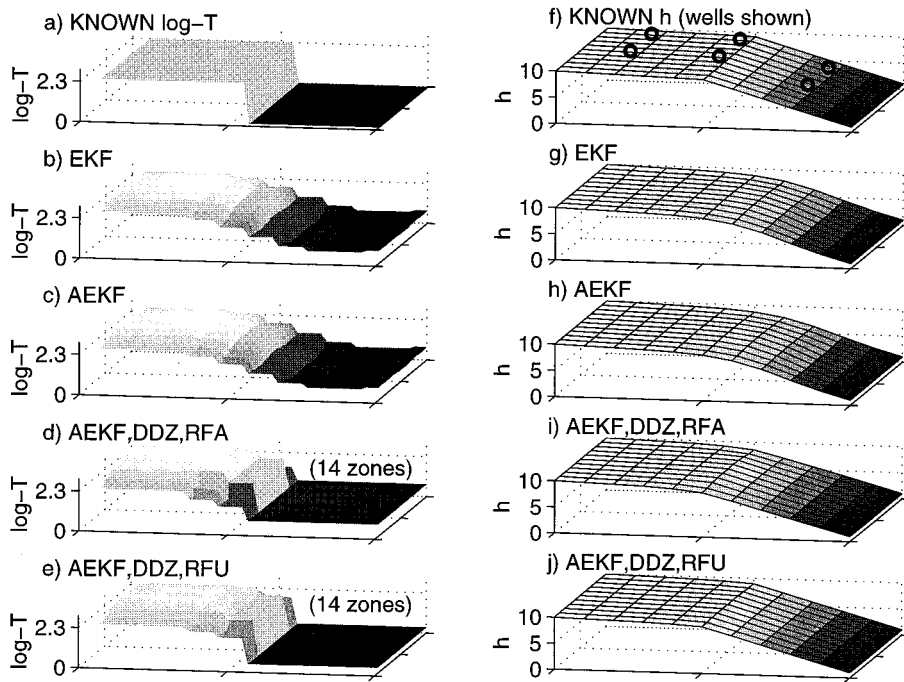
### 5. Results From the Four Test Domains

The two nonlinear filtering algorithms, EKF and AEKF, yield essentially identical estimates on all four test domains whether or not used in conjunction with data-driven zonation; all results of both methods (except for computational effort) are indistinguishable from each other on all plots. Apparently, for these test cases the effects of the parameter estimates so dominate the head prediction that the approximations of AEKF result in no significant loss of information. In order to avoid redundancy, we have not presented results of the DDZ methods used with EKF, as these were identical to the results using DDZ with AEKF.

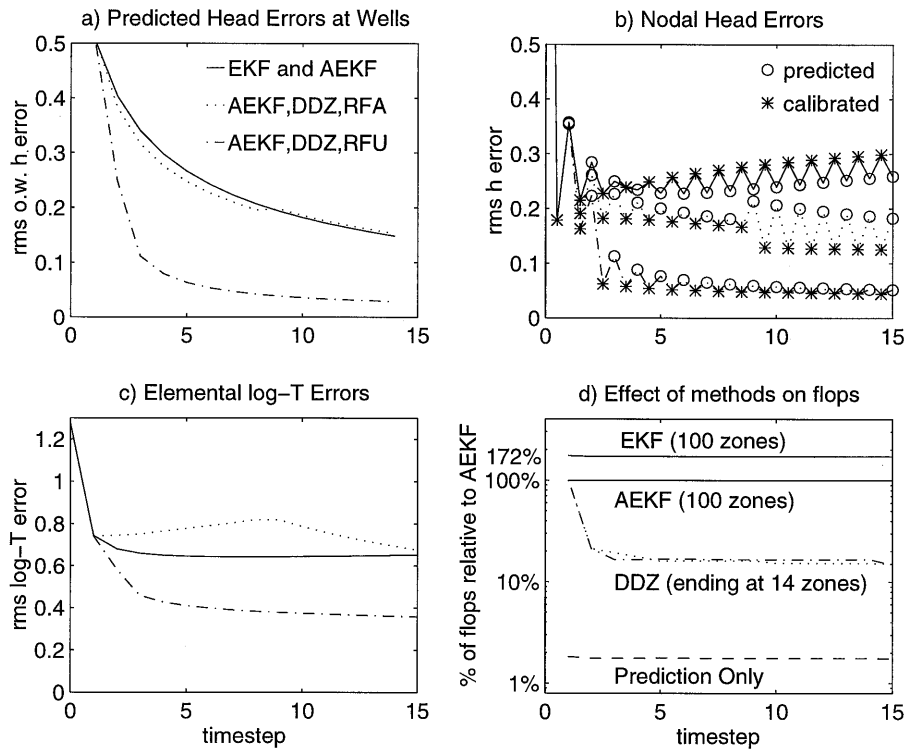
The most interesting differences are apparent between methods on the z2 domain (see Figure 4), and we present this case in the most detail. While the unzoned methods detect the correct general "trend" of the transmissivity field, they are unable to resolve the sharp zone boundary, and both ends of the transmissivity field are a bit high (Figures 4b and 4c). The methods that employ data-driven zonation are able to resolve the zone boundary between low and high transmissivity (Figures 4d and 4e). The estimated transmissivity field using random field averaging to merge the zones results in values that are still too high at both ends of the domain (Figure 4d). The best estimate is achieved using random field union to dynamically merge zones (Figure 4e), where zones evolve increasing independence through merging and are thus more free to converge on correct values. The resulting head predictions (Figures 4g to 4j) clearly show that the model based on AEKF with DDZ using RFU generates the most accurate estimate of the head field.

Errors of estimates were quantified by taking the root-mean-squares (rms) of three types of differences: (1) between the predicted heads at the observation wells and known values (also known as the "innovations"  $Z - Mh$ ) which drive the filtering process), (2) between both predicted and calibrated nodal heads across the whole domain, and (3) between the elemental log transmissivity estimates and known values.


Error plots for the z2 domain are shown in Figure 5. The



**Figure 4.** Known and estimated (left) log transmissivity and (right) head fields for the z2 test case. Transmissivity estimates are based on 15 measurement updates using perfect head measurements from the six observation wells shown in Figure 4f. Head fields are predicted using the final transmissivity fields. The two DDZ methods coincidentally generate the same zonation structure at the end of 15 time steps on this domain, although the evolution of this structure varies slightly between RFA and RFU.



**Figure 5.** Performance of the four methods on the z2 domain. (a) Root-mean-square error of the predicted head at the six observation wells (the “innovations”). (b) Root-mean-square nodal head errors over the entire domain: at each time step the error of the predicted heads is marked with a small circle (the time update), followed by the error of the calibrated heads (the measurement update). (c) Root-mean-square elemental log transmissivity error: for the zoned methods, each element is assigned its zone log transmissivity value. (d) Number of flops required per time step, shown as percent relative to those required by AEKF. Note that on the error plots (Figures 5a to 5c) the curves for EKF and AEKF are completely coincident.



**Figure 6.** Sensitivity of the root-mean-square log transmissivity errors on the z2 test domain to the size of the time step ( $dt$ , in days). (a) The results using unzoned EKF become increasingly inaccurate as the time step is shortened. (b) The results using AEKF with DDZ and RFU are relatively robust despite large changes in time step.

innovations decrease nearly monotonically for all methods (Figure 5a), which is to be expected because the filter is designed to minimize this error. Dynamic zonation occasionally briefly disrupts the monotonicity of this decline by a very small amount (the largest observed increase was only 0.016 m using RFA on the ramp domain). The EKF is “optimal” only in the sense that it minimizes the sum of the estimated variances based on the innovations; in reality, errors in state and parameter estimation are introduced by a variety of sources. For example, information is lost by the required linearization of the transition matrix, errors in state estimates or measurements may not be normally distributed, and the true correlation structure between head nodes and log transmissivity zones is generally unknown. Thus while the filter will act to decrease the innovations in the measurements, overall state and parameter estimates are not guaranteed to decrease. Furthermore, overparameterization can result in instability [Yeh and Yoon, 1981]. All methods show an initial decrease in log transmissivity error on the z2 domain (Figure 5c). However, the unzoned methods begin to exhibit very slightly divergent behavior in log transmissivity estimates over time, and in Figure 5b it can be seen that the head calibrations imposed by the filter (shown by the asterisks) actually begin to increase the error in the overall head field starting at the fifth “time step” (iteration); presumably, this is caused by the overparameterization of transmissivity into 100 zones (one per element) relative to only six observation wells. Application of data-driven zonation prevented this instability and resulted in convergent behavior in all cases tested, with calibration continuing to improve the estimates. Zone merging can cause discontinuities in the error curves, but the new information that can be extracted from new zonations ultimately leads to better estimates. Despite the fact that for the z2 domain, data-driven zonation using both RFA and RFU resulted in the same zonation pattern, the value estimates achieved using RFU were considerably better than those resulting from RFA (Figures 5a, 5b, and 5c).

The dynamically zoned methods, especially using RFU, are also more robust in the face of errors in the size of the time step and/or estimation of storativity. The relatively high value of storativity (0.001) required a fairly long time step to achieve steady state; yet AEKF using DDZ with RFU produced nearly identical estimates when run with a time step of 200 days or a time step of 1 day (Figure 6b). In contrast, EKF became increasingly unstable and divergent as the time step was short-

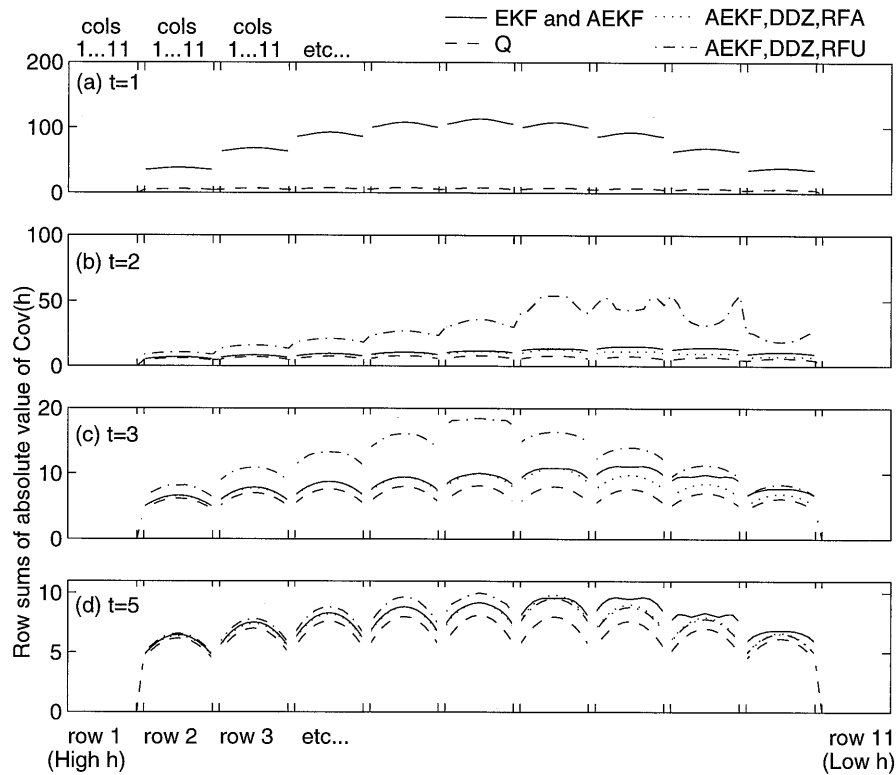
ened (Figure 6a). This trend was most noticeable on the z2 domain, but was also observable on the other three test domains.

We use floating point operation (flop) counts as a machine-independent indication of computational time requirements. Although our current Matlab implementations are not tuned for efficiency (e.g., factored forms are not employed, matrix symmetry is not taken advantage of, and overhead is present for data visualization, calculation of matrix norms and error estimates, etc.), comparing flop counts can still shed light on the relative performance of the various algorithms. The relative computational efficiency of the methods on the z2 domain is shown in Figure 5d. On all four 100-element domains, AEKF requires only 58% of the flops required by full EKF. Relative to AEKF, the addition of data-driven zonation decreases the computational effort by roughly an order of magnitude after the first time step, although the exact amount is dependent of the amount of zone merging incurred. For comparison, the number of flops required by prediction only is also shown.

Results of the four methods on the other three domains are less distinct from each other; all methods show stable behavior and error results are very similar. As was previously mentioned, AEKF performs identically to EKF in these test cases in all ways except computation time. As with the z2 domain, on the three-zone domains (z3hml and z3mhl) the data-driven zonation methods are able to detect the zone boundaries better than the unzoned methods, with RFU doing slightly better than RFA in this regard. Not surprisingly, the unzoned methods give slightly better results on the ramp domain, but although the number of zones is reduced to 16 using RFA and to 9 using RFU, these zones tend to form “steps,” and no spurious sharp boundaries are introduced.

It is worth noting that we also ran Monte Carlo experiments using varying amounts of simulated white noise on the measurements. Convergence is generally slowed by increasing the measurement error, but the results are not qualitatively affected in other ways; the relative behaviors of the methods remain consistent with what we report here.

Recall that the measurement update of the transmissivity zones is driven by the Kalman gain matrix  $K^y$  (equation (7)), which is constructed based on  $\text{Cov}(h, y)$  and  $\text{Cov}(h)$  (equation (5)). In EKF  $\text{Cov}(h, y)$  and  $\text{Cov}(h)$  are updated just prior to construction of the gain matrix via equations (2)–(3), whereas in AEKF  $\text{Cov}(h, y)$  and  $\text{Cov}(h)$  are approximated



**Figure 7.** Row sums of the absolute value of  $\text{Cov}(h)$  at all nodes in the finite element mesh for selected time steps on the z2 test domain, for all four estimating algorithms and in comparison to the system noise covariance  $Q$ . Note the change in vertical scaling of the four graphs as head variance decreases over time. Curves for all four methods coincide in Figure 7a; curves for EKF and AEKF coincide in all panels. The other three domains showed similar patterns.

from equations (12)–(13). For the test cases currently under consideration, the approximate covariance matrices of AEKF are essentially identical to the updated ones of EKF, and hence the algorithms perform identically. For example, Figure 7 shows the row sums of the absolute value of the  $\text{Cov}(h)$  matrix for four time steps (just prior to the computation of the gain matrix); the curves for EKF and AEKF are completely coincident. This results from the fact that the terms containing the transition matrix  $\Phi$ , which are eliminated in the AEKF approximations, are orders of magnitude smaller than the retained terms.

Despite the fact that data-driven zonation yields similar or identical zonation patterns whether used with RFA or RFU, the  $\text{Cov}(h)$  matrix differs markedly between the two DDZ methods, thus accounting for their differing performance. While DDZ using RFA yields a  $\text{Cov}(h)$  matrix that is nearly coincident (slightly smaller in some areas) with that of the unzoned methods, DDZ with RFU has a much larger  $\text{Cov}(h)$  in the early time steps once zonation has commenced (Figures 7b and 7c). It is also interesting to note how the uncertainty represented by  $\text{Cov}(h)$  starts out highest near the downstream end of the domain in RFU (Figure 7b) and is then driven farther upstream (Figure 7c) as the filtering progresses. By about the fifth time step the head covariances for all methods have nearly converged to the level of the system noise  $Q$  (Figure 7d). The patterns depicted in Figure 7 for the z2 domain were similar for the other three 100-element test domains.

In order to understand the similarities and differences in the

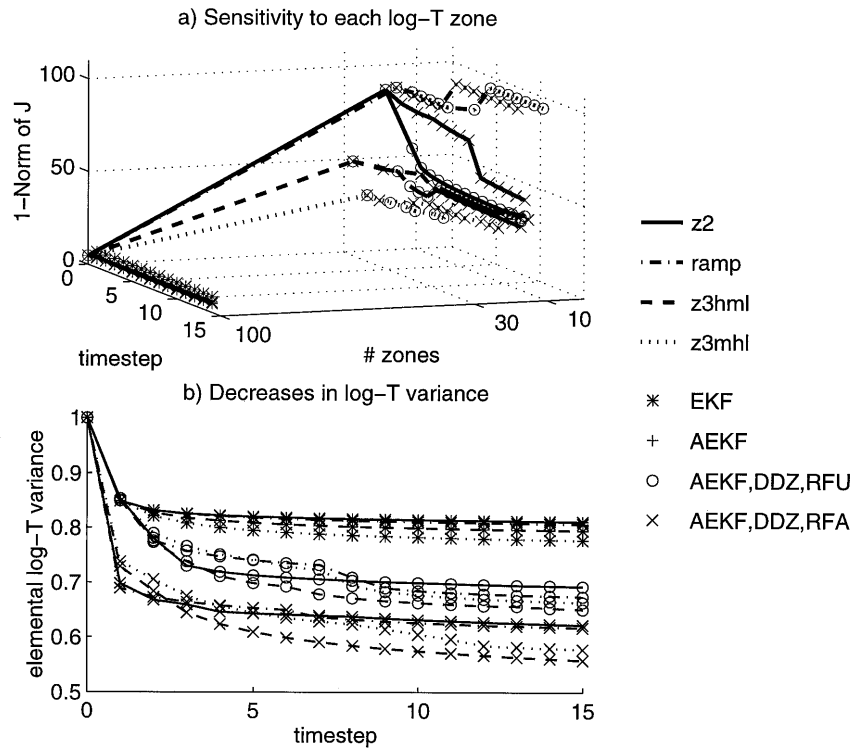
$\text{Cov}(h)$  matrices between the unzoned methods, DDZ with RFA and DDZ with RFU, it is instructive to examine the two major inputs into the construction of  $\text{Cov}(h)$ : the sensitivity matrix  $J$  and the log transmissivity covariance matrix  $\text{Cov}(y)$ .

It is intuitive that we should be able to increase sensitivities of heads to log transmissivity zones by using larger zones, and our data support this. This sensitivity is quantified using the 1-norm  $\|J\|_1$  of the sensitivity matrix  $J$  as defined below:

$$\|J\|_1 \triangleq c \sum_{r=1}^{\max \text{ nodes}} |J_{rc}| \quad (25)$$

In Figure 8a we show the changes in  $\|J\|_1$  for all four methods on all four domains as a function of time step and number of zones. Although the exact sensitivity varies widely with the particular domain and zonation pattern, the  $\|J\|_1$  typically increases by an order of magnitude with the first zonation from 100 zones down to less than 30 zones. Thereafter, subsequent zonations may initially increase or decrease the sensitivity, but when a zonation pattern remains the same over successive time steps, there is a tendency for the sensitivity to decrease very gradually as zone values converge on better estimates. Increased sensitivity alone cannot explain the higher  $\text{Cov}(h)$  of DDZ with RFU, since it is apparent from Figure 8a that the  $\|J\|_1$  using DDZ with RFA is at least as high as that using DDZ with RFU.

As was mentioned previously, the covariance of the log transmissivity zones  $\text{Cov}(y)$  is also an important component in the construction of the filter gain matrix. It is difficult to di-



**Figure 8.** The causes of the differing performance between methods. (a) The effect of data driven zonation on the 1-norm of the sensitivity matrix,  $\|J\|_1$ . Overall sensitivity is increased by about an order of magnitude for the zoned methods relative to the unzoned methods. (b) Elemental log transmissivity variance over time for the four recursive estimating algorithms on the four sample domains, as estimated by an area-weighted average of the diag (Cov ( $y$ )). Note that variance is lower using RFA than using RFU, but covariance between zones is calculated in the same way for these two methods; the result is that zones are more highly correlated in RFA than in RFU.

rectly compare Cov ( $y$ ) for different zonations. We compare the area-weighted average (“elemental”) log transmissivity variances (Figure 8b). On all test domains, DDZ with RFA has the lowest average log transmissivity variance, and the unzoned methods have the highest variance. The variance of DDZ with RFU is lower than with the unzoned methods because the sensitivities in  $J$  are higher using DDZ, resulting in a higher gain matrix  $K^y$  and a correspondingly greater reduction in uncertainty. The variance of DDZ using RFA is lower than that of DDZ using RFU because the variances computed using equation (17) are always less than or equal to the variances computed using (24), given the same initial zone statistics. However, the zone cross covariances computed for RFA by (16) are identical to those computed for RFU by (23), resulting in a higher correlation between zones merged by RFA than by RFU. Average interzone correlations evolved similarly when using DDZ with RFU as when using the unzoned methods. Apparently, the lower log transmissivity variances (relative to the cross covariances) with RFA largely cancel the effect of the increased sensitivities from the zonation, resulting in a state covariance matrix Cov ( $h$ ) nearly identical to that of the unzoned methods (Figure 7). In DDZ with RFU the increased sensitivities together with the maintenance of higher log transmissivity variances translate into higher state covariance estimates (Figure 7).

In summary, on four 100-element steady-flow test cases, AEKF yields estimates virtually identical to those from EKF at significantly less computational cost. Data-driven zonation re-

duces the cost even further and yields estimates that are as good as or better than those generated by the unzoned methods. Dynamic zonation also introduces a stabilizing effect by reducing the dimensionality of the parameter space and makes the method more robust in the presence of errors in estimates of other physical or numerical parameters, such as storativity and/or time step. RFA and RFU tend to produce a similar zonation pattern, but these zones are less correlated in RFU than in RFA, resulting in better estimates of zone values when RFU is used to merge zone statistics.

## 6. USGS Synthetic Valley Test Case

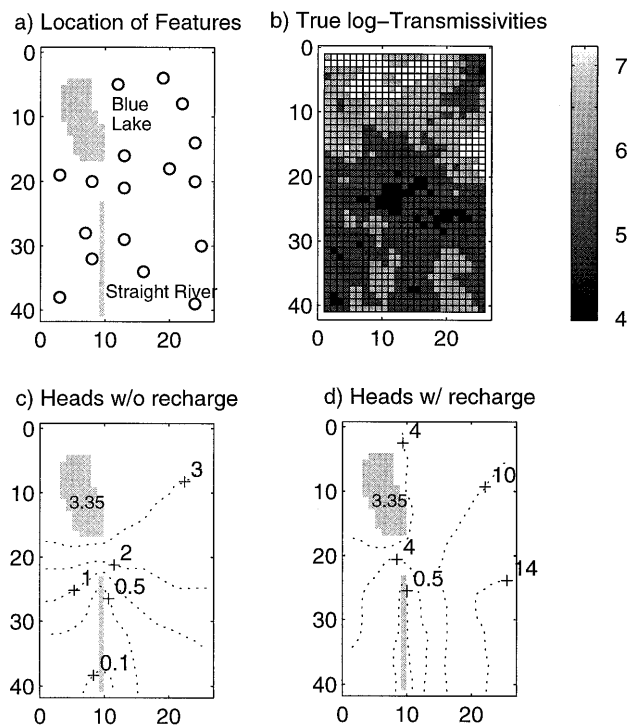
The computational efficiency of the integrated method of AEKF with DDZ using RFU has enabled us to test it on “Synthetic Valley” (Figure 9), a larger, more realistic synthetic domain than the four test cases presented earlier. Synthetic Valley was developed by the United States Geological Survey (Hill et al., submitted manuscript, 1994) for the purpose of comparing model calibration techniques. The model simulates an undeveloped alluvial valley in hydrologic steady state surrounded by low-permeability bedrock; no flow enters or leaves the valley boundaries except via precipitation, evaporation, and river discharge. The domain is discretized into a 40 by 25 element finite element mesh (1000 elements, 1066 nodes). Two surface water features are evident, Blue Lake and Straight River, both with known surface elevations. Seventeen monitoring wells are in place the first season (Figure 9a). (Although

more wells are added in subsequent years, we simulated head measurements from only the original 17 wells. Simulated driller's logs are also available, although we did not use that information in our simulations.) Distributed hydraulic conductivities for all 1000 elements of Synthetic Valley are known (they were reportedly generated stochastically using the turning bands method), and the depth to the bottom of the aquifer is given throughout. In this study, hydraulic conductivities were transformed into approximate transmissivities by multiplying by the aquifer depth (ignoring transient changes in depth); the resulting log transmissivity field is shown in Figure 9b. The head field was simulated under two different steady state conditions; one with no forcing conditions in which only the known values of lake and river surfaces were used to establish the head field (Figure 9c) and one in which a recharge term was present (including the effects of rainfall, evaporation from the lake, evapotranspiration from the land, and discharge from the river) as shown in Figure 9d. Note that these two head fields are dramatically different from each other, with the directions of flow nearly perpendicular to each other in places, and that the overall head gradients are much steeper in the recharge case.

We present results from three experiments on the Synthetic Valley domain. All 1000 elemental transmissivities were initialized to a value of 563 m/d (the mean of the true known elemental transmissivities). All other statistics were initialized exactly as was described in section 4.3 for the four smaller test cases. Using noise-free head "measurements" from the 17 wells each time step, we used AEKF with DDZ using RFU to estimate the log-transmissivities under the following conditions: case A, using the steady-state head field from Figure 9c with no recharge simulated for six time steps; case B, using the steady state head field from Figure 9d with the recharge term present for six time steps; and case C, using steady state head measurements from the field shown in Figure 9d (the recharge case) and from the field shown in Figure 9c (the no-recharge case) in alternate time steps for a total of six time steps, starting with the recharge case.

The final log transmissivity estimates of Synthetic Valley for the three cases at the end of 6 time steps are shown in Figure 10a. We assess the quality of these estimates in two ways. First, the error root-mean-square of the differences between elemental log transmissivities of each of the three estimates and the true elemental log transmissivities as shown in Figure 9b were computed and are displayed by the pluses in Figure 11c. Second, the true log transmissivity values were zoned into each of the three zonation patterns resulting from cases A, B, and C. Effective transmissivities of two-dimensional fields have been shown to be approximated by the geometric means of the distributed transmissivities [Gutjahr *et al.*, 1978]. Accordingly, we determined the "true" zone values by taking the logs of the geometric means of the true elemental transmissivities over each of the identified zones and then assigning each zone value back to the elemental log transmissivities of zone members. These true zoned log transmissivities are shown in Figure 10b. The error root-mean-square of each of the elemental differences between the estimates in Figure 10a and the true zone values in Figure 10b were computed and are displayed by the crosses in Figure 11c.

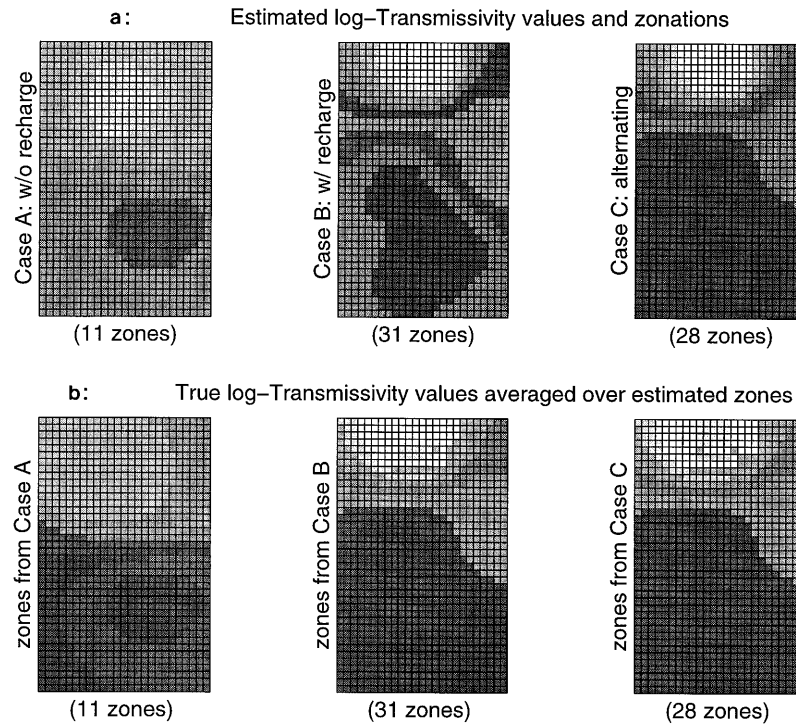
Despite the low head gradient in case A (and concomitantly small innovations, as shown in Figure 11a), log transmissivity error is reduced (Figure 11c), and the general high and low regions of the log transmissivity field are detected (Figure 10a,



**Figure 9.** USGS Synthetic Valley calibration test domain: (a) Locations of the lake, river, and the 17 observation wells. (b) "True" elemental log transmissivity field in 1000 elements. (c) "True" steady-state head field with no-flow boundaries, fixed surface elevations at the lake and river, and no recharge term. (d) "True" steady-state head field with no-flow boundaries, fixed surface elevations at the lake and river, and a recharge term present accounting for the effects of rainfall, evaporation from the lake, evapotranspiration from the land, and discharge from the river.

case A). In case B, where the head gradient is much higher, the innovations are also higher (Figure 11a), and the resulting zonation pattern captures the primary features of the true field much more closely (compare case B of Figure 10a with Figure 9b). In case C, where the recharge conditions were alternated with the no-recharge conditions, more information is available to the filter, and the best results are achieved. Although the zonation pattern is very similar to that of case B (Figure 10a, case C), since it is largely established by the same first dynamic zonation, the zone value estimates are improved (Figure 11c) owing to the fluctuating increases in the innovations (Figure 11a). (We also performed experiments alternating the no-recharge case with the recharge case for a total of six time steps, starting with the no-recharge case. Since in this case the zonation pattern is determined primarily (in the first zonation) by the weaker conditions of the no-recharge case, the errors of the resulting estimates lie between those of case A and case B; this highlights the importance of using the strongest forcing conditions in the initial pass of the filter before applying dynamic zonation.)

The improvement in estimates from case A to case B to case C is apparent by noting the increasing similarity between Figures 10a and 10b from left to right. The root-mean-square error plots shown in Figure 11c confirm the relative quality of the estimates in these three cases; the general reduction in log transmissivity error (the pluses) indicates that the evolving



**Figure 10.** USGS Synthetic Valley test domain zoned log transmissivity fields. (a) Final log transmissivity estimates using AEKF with DDZ and RFU after six time steps under three forcing conditions: case A, no recharge term; case B, with a recharge term; and case C, alternating between recharge and no recharge. (b) True zoned log transmissivities, using the zonations generated by the three test cases. Zone values were computed by taking the logs of the geometric means of the true elemental transmissivities within each identified zone. In both Figures 10a and 10b the gray scale is the same as indicated in Figure 9b.

zonation patterns are consistent with the true field, and the larger reductions in the errors of zone values relative to the true “zoned” values (the crosses) indicate that the values continue to improve within the identified zones. The convergent behavior between the error root-mean-squares of the predicted and calibrated head fields (Figure 11b) shows that all cases are stable.

In the 1000-element Synthetic Valley test case, DDZ reduced the number of flops required per time step by 2 orders of magnitude as compared with the unzoned AEKF (Figure 11d).

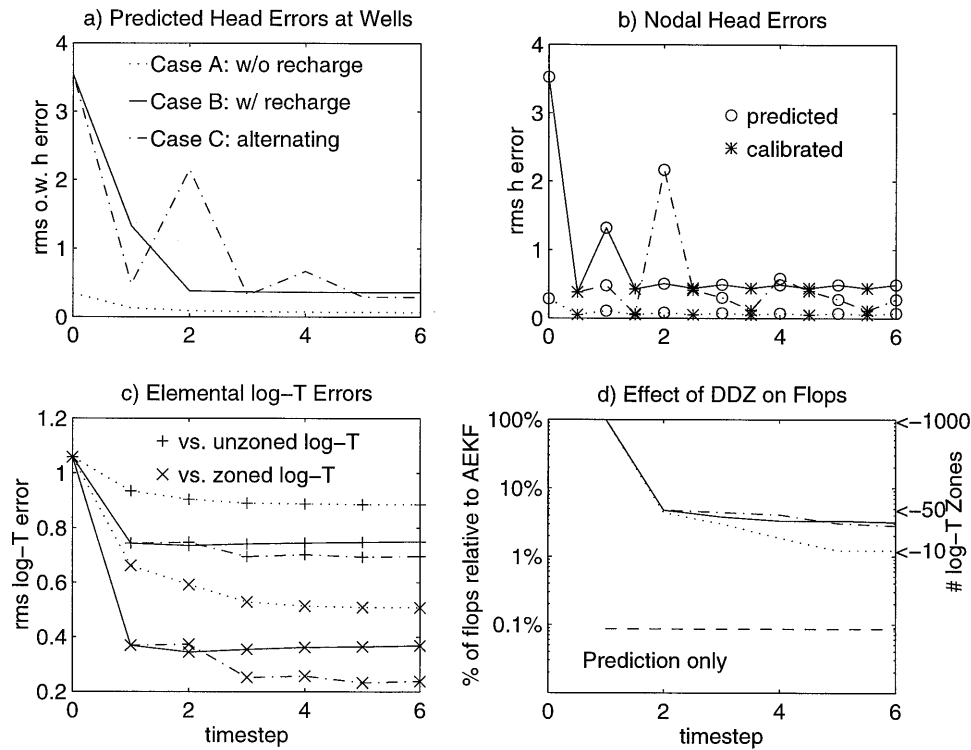
## 7. Conclusions and Future Research

The extended Kalman filter can be efficiently approximated by replacing the covariance time updates with approximations that do not use the transition matrix. This approximate extended Kalman filter is valid where the errors in the parameters being estimated and the system error dominate the current state error, such as with steady or long-term cyclic groundwater flow where distributed transmissivities are being estimated. Experiments are currently underway to assess the usefulness of this approximation under more general transient conditions. Although the EKF was developed for real-time recursive calibration of models in control systems engineering, in the groundwater estimation problem it can also be exploited as an iterative method for stationary flow processes or as an iteratively recursive method for cyclostationary (seasonal) flow.

The bulk of the computational effort of approximate extended Kalman filtering is in the repeated and expensive calculation of the Jacobian sensitivity matrix. We are currently

researching potential ways to reduce the time required for calculating the Jacobian. First, we are exploring the possibility of reducing the number of rows in the Jacobian from the number of nodes to the number of observation wells and then estimating (or eliminating, as suggested by *Van Geer and Van Del Kloet* [1985]) the measurement update of the state variables. A reduced Jacobian may be computed quickly and accurately using the adjoint (variational) method [*Sun*, 1994] and would also result in smaller  $Cov(h)$  and  $Cov(h, y)$  matrices. Preliminary work on this approach has been very encouraging. Second, we are assessing the use of an update for the Jacobian for successive time steps in which the zonation structure remains stable. The very low sensitivity of the Jacobian to changes in zone values (Figure 8a) is suggestive that such updates may be feasible, or may possibly even be eliminated following *Sun and Yeh* [1985].

The use of cluster analysis techniques for data-driven zonation is shown to be an effective method for automatic parameterization. By simultaneously estimating the values and structure of heterogeneous transmissivity fields, the advantages of zonation are obtained (stability, low computational complexity) without the risks of improper zonation associated with a priori determination of the transmissivity structure. Although the current implementation is restricted to cluster analysis along only one dimension (transmissivity value), we are exploring the expansion of this to multiple dimensions, which could take into account both parametric properties (such as zone shape and orientation relative to flow) and nonparametric properties (such as soil type, where known) of the zones.



**Figure 11.** Performance of AEKF with DDZ using RFU on the Synthetic Valley domain under three different forcing conditions: case A, no recharge vector; case B, with recharge; and case C, alternating between recharge and no recharge. (a) Root-mean-square error of the predicted head at the 17 observation wells (the “innovations”). (b) Root-mean-square nodal head errors over the entire domain. At each time step the error of the predicted heads is marked with a small circle (the time update), followed by the error of the calibrated heads (the measurement update). (c) Root-mean-square elemental log transmissivity error for the three cases as compared with the true unzoned log-transmissivities shown in Figure 9b, indicated here by pluses, and as compared to the true zoned log transmissivity estimates as shown in Figure 10b, indicated here by the crosses. (d) Number of flops required per time step, shown as percent relative to those required by unzoned AEKF. The approximate number of zones can be inferred from the numbers shown on the right, but these do not correspond to the logarithmic tick marks on the vertical axis.

Heterogeneous random fields can be merged using a random field “union,” in which the merged variance takes into account differences in the means of the merged fields but is independent of the covariance between fields being merged. This contrasts with the more traditional random field “averaging,” which is based on an assumption of statistical homogeneity; when this assumption is violated, random field averaging underestimates merged variances and inflates interfield correlations. Random field union has proven useful in data-driven zonation and may be applicable to other areas in stochastic subsurface modeling, such as the determination of initial estimates of transmissivity variance based on heterogeneous strata from well-borehole data.

The integrated method of approximate extended Kalman filtering with data-driven zonation using random field union has been shown to yield reasonable estimates (consistent with the amount of information available) of the heterogeneous transmissivity domains for four small synthetic test cases as well as the larger, more realistic domain of the USGS Synthetic Valley calibration test case run under two different steady state conditions and a cyclicly steady condition. In all cases, the only a priori transmissivity value that is required is a single estimate of the average transmissivity in the domain. The method has been applied to two-dimensional domains only. Three dimen-

sional domains will also yield to this strategy, requiring only greater computational resources.

The methods presented here apply to the (essentially) elliptic groundwater flow equation. Although solute concentration measurements are available at many contaminated facilities, these data will not be readily assimilated using the techniques reported. In many granular and fractured systems, advective transport dominates, and hence the solute transport equation is almost hyperbolic. This results in concentration information being propagated along real characteristics in finite time, whereas the saturated groundwater flow equation provides for global propagation at infinite speed. An alternative strategy for incorporating concentration measurements into estimation of parameter structure and value is currently being considered.

**Acknowledgments.** The research reported in this paper was supported in part by National Science Foundation grants ASC-91000226, MSS-9214698, and OSR-9350540; the Government may retain some rights to this work. M.J.E.’s participation in this project was made possible in part by a grant from the AAUW Educational Foundation.

## References

Anderssen, R. S., and S.-S. Chow, Resolving the transmissivity zonation in a confined aquifer, in *Mini-Conference on Inverse Problems in*

- Partial Differential Equations, Canberra, 1990, Proc. Cent. Math. Its Appl.*, vol. 31, edited by A. K. Pani and R. S. Anderssen, pp. 247–261, Aust. Natl. Univ., Canberra, 1992.
- Anderssen, R. S., and C. R. Dietrich, *The Inverse Problem of Aquifer Transmissivity Identification, Res. Rep. CMA-R01-87*, Cent. for Math. Anal., Aust. Natl. Univ., Canberra, 1987.
- Andricevic, R., Coupled withdrawal and sampling designs for groundwater supply models, *Water Resour. Res.*, 29(1), 5–16, 1993.
- Andricevic, R., and P. K. Kitanidis, Optimization of the pumping schedule in aquifer remediation under uncertainty, *Water Resour. Res.*, 26(5), 875–885, 1990.
- Burl, J. B., A reduced order extended Kalman filter for sequential images containing a moving object, *IEEE Trans. Image Process.*, 2(3), 285–295, 1993.
- Carlson, N. A., Federated square root filter for decentralized parallel processes, *IEEE Trans. Aerosp. Electron. Syst.*, 26(3), 517–525, 1990.
- Chui, C. K., G. Chen, and H. C. Chui, Modified extended Kalman filtering and a real-time parallel algorithm for system parameter identification, *IEEE Trans. Autom. Control*, 35(1), 100–104, 1990.
- Dee, D. P., Simplification of the Kalman filter for meteorological data assimilation, *Q. J. R. Meteorol. Soc.*, 117, 365–384, 1991.
- Dubes, R. C., How many clusters are best?—An experiment, *Pattern Recognit.*, 20(6), 645–663, 1987.
- Eppstein, M. J., and D. E. Dougherty, Parameter estimation with data-driven zonation, in *Computational Methods in Water Resources X*, vol. 1, edited by A. Peters et al., pp. 727–734, Kluwer Acad., Norwell, Mass., 1994.
- Gelb, A. (Ed.), *Applied Optimal Estimation*, MIT Press, Cambridge, Mass., 1974.
- Graham, W., and D. B. McLaughlin, Stochastic analysis of nonstationary subsurface transport, 2, Conditional moments, *Water Resour. Res.*, 25(11), 2331–2355, 1989.
- Graham, W. D., and D. B. McLaughlin, A stochastic model of solute transport in groundwater: Application to the Borden, Ontario, tracer test, *Water Resour. Res.*, 27(6), 1345–1359, 1991.
- Graham, W. D., and C. D. Tankersley, Forecasting piezometric head levels in the Floridian aquifer: A Kalman filtering approach, *Water Resour. Res.*, 29(11), 3791–3800, 1993.
- Gutjahr, A. L., L. W. Gelhar, A. A. Bakr, and J. R. MacMillan, Stochastic analysis of spatial variability in subsurface flows, 2, Evaluation and application, *Water Resour. Res.*, 14(5), 953–959, 1978.
- Hantush, M. M., and M. A. Mariño, Two-dimensional analysis and optimal estimation in aquifers: Random recharge, *Water Resour. Res.*, 30(2), 559–569, 1994.
- Jain, A. K., and R. C. Dubes, *Algorithms for Clustering Data*, Prentice Hall, Englewood Cliffs, N. J., 1988.
- Jinno, K., A. Kawamura, T. Ueda, and H. Yoshinaga, Prediction of the concentration distribution of groundwater pollutants, in *Groundwater Management: Quantity and Quality, LAHS Publ. 188*, 131–142, 1989.
- Kalman, R. E., A new approach to linear filtering and prediction problems, *J. Basic Eng.*, 82D, 35–45, 1960.
- Kitanidis, P. K., A unified approach to the parameter estimation of groundwater models, M.S. thesis, Dep. of Civ. Eng., Mass. Inst. of Technol., Cambridge, 1976.
- Lee, S., and P. K. Kitanidis, Optimal estimation and scheduling in aquifer remediation with incomplete information, *Water Resour. Res.*, 27(9), 2203–2217, 1991.
- McLaughlin, D. B., Application of Kalman filtering to groundwater basin modeling and prediction, in *Real-Time Forecasting/Control of Water Resource Systems, ILASA Proc. Ser.*, edited by E. F. Wood, pp. 109–123, Pergamon, New York, 1976.
- McLaughlin, D. B., Potential applications of Kalman filtering concepts to groundwater basin management, in *Applications of Kalman Filter to Hydrology, Hydraulics, and Water Resources*, edited by C. Chiu, pp. 639–655, Univ. of Pittsburgh, Pittsburgh, Pa., 1978.
- Roy, S., R. H. Hashemi, and A. J. Laub, Square root parallel Kalman filtering using reduced-order local filters, *IEEE Trans. Aerosp. Electron. Syst.*, 27(2), 276–289, 1991.
- Schmidtke, K. D., E. A. McBean, and J. F. Sykes, Stochastic estimation of states in unconfined aquifers subject to artificial recharge, *Water Resour. Res.*, 18(5), 1519–1530, 1982.
- Şen, Z., Adaptive pumping test analysis, *J. Hydrol.*, 74, 259–270, 1984.
- Stark, H., and J. W. Woods, *Probability, Random Processes, and Estimation Theory for Engineers*, Prentice-Hall, Englewood Cliffs, N. J., 1986.
- Sun, N.-Z., *Theory and Applications of Transport in Porous Media*, vol. 6, *Inverse Problems in Groundwater Modeling*, edited by J. Bear, Kluwer Acad., Norwell, Mass., 1994.
- Sun, N., and W. W.-G. Yeh, Identification of parameter structure in groundwater inverse problem, *Water Resour. Res.*, 21(6), 869–883, 1985.
- Sun, N.-Z., M.-C. Jeng, and W. W.-G. Yeh, A proposed geological parameterization method for parameter identification in three-dimensional groundwater modeling, *Water Resour. Res.*, 31(1), 89–102, 1995.
- Todding, R., and S. E. Cohn, Suboptimal schemes for atmospheric data assimilation based on the Kalman filter, *Mon. Weather Rev.*, 122, 2530–2557, 1994.
- Van Geer, F. C., and P. Van Der Kloet, Two algorithms for parameter estimation in groundwater flow problems, *J. Hydrol.*, 77, 361–378, 1985.
- Van Geer, F. C., C. B. M. TeStroet, and Zhou, Y., Using Kalman filtering to improve and quantify the uncertainty of numerical groundwater simulations, 1, The role of system noise and its calibration, *Water Resour. Res.*, 27(8), 1987–1994, 1991.
- Vanmarcke, E., *Random Fields: Analysis and Synthesis*, MIT Press, Cambridge, Mass., 1983.
- Wang, H. F., and M. P. Anderson, *Introduction to Groundwater Modeling: Finite Difference and Finite Element Methods*, W. H. Freeman, New York, 1982.
- Wilson, J., P. Kitanidis, and M. Dettinger, State and parameter estimation in groundwater models, in *Applications of Kalman Filter to Hydrology, Hydraulics, and Water Resources*, edited by C. Chiu, pp. 657–679, Univ. of Pittsburgh, Pittsburgh, Pa., 1978.
- Yeh, W. W.-G., Review of parameter identification procedures in groundwater hydrology: The inverse problem, *Water Resour. Res.*, 22(2), 95–108, 1986.
- Yeh, W. W.-G., and Y. S. Yoon, Aquifer parameter identification with optimum dimension in parameterization, *Water Resour. Res.*, 17(3), 664–672, 1981.
- Yu, Y.-S., M. Heidari, and Wang Guang-Te, Optimal estimation of contaminant transport in ground water, *Water Resour. Bull.*, 25(2), 295–300, 1989.
- Zhou, Y., C. B. M. Te Stroet, and F. C. Van Geer, Using Kalman filtering to improve and quantify the uncertainty of numerical groundwater simulations, 2, Applications to monitoring network design, *Water Resour. Res.*, 27(8), 1995–2006, 1991.
- Zou, S., and A. Parr, Optimal estimation of two-dimensional contaminant transport, *Ground Water*, 33(2), 319–325, 1995.

D. E. Dougherty and M. J. Eppstein, Research Center for Groundwater Remediation Design, Department of Civil and Environmental Engineering, University of Vermont, Burlington, VT 05405-0156.

(Received August 24, 1995; revised May 9, 1996; accepted July 23, 1996.)