

Assimilation of Simulated Doppler Radar Observations with an Ensemble Kalman Filter*

CHRIS SNYDER

National Center for Atmospheric Research, Boulder, Colorado

FUQING ZHANG

Department of Atmospheric Science, Texas A&M University, College Station, Texas

(Manuscript received 11 November 2002, in final form 30 January 2003)

ABSTRACT

Assimilation of Doppler radar data into cloud models is an important obstacle to routine numerical weather prediction for convective-scale motions; the difficulty lies in initializing fields of wind, temperature, moisture, and condensate given only observations of radial velocity and reflectivity from the radar. This paper investigates the potential of the ensemble Kalman filter (EnKF), which estimates the covariances between observed variables and the state through an ensemble of forecasts, to assimilate radar observations at convective scales. In the basic experiment, simulated observations are extracted from a reference simulation of a splitting supercell and assimilated using the EnKF and the same numerical model that produced the reference simulation. The EnKF produces accurate analyses, including the unobserved variables, after roughly 30 min (or six scans) of radial velocity observations. Additional experiments, in which forecasts are made from the ensemble-mean analysis, reveal that forecast errors grow significantly in this simple system, so that the ability of the EnKF to track the reference solution is not simply because of stable system dynamics. It is also found that the covariances between radial velocity and temperature, moisture, and condensate are important to the quality of the analyses, as is the initialization chosen for the ensemble members prior to assimilating the first observations. These results are promising, especially given the ease of implementing the EnKF. A number of important issues remain, however, including the initialization of the ensemble prior to the first observation, the treatment of uncertainty in the environmental sounding, the role of error in the forecast model (particularly the microphysical parameterizations), and the treatment of lateral boundary conditions.

1. Introduction

A long-standing problem in meteorology is to estimate, or retrieve, the state of the atmosphere in some domain given observations of reflectivity and radial velocity from one or more Doppler radars. (Alternatively, one may speak of assimilating these observations.) Doing this requires an algorithm that infers from the observations those variables not directly observed, such as vertical velocity and temperature, and this in turn requires additional information beyond the observations themselves. Such information is available, in principle, from our knowledge of the governing equations, and various techniques have been devised to utilize the governing equations in assimilation or retrieval.

This paper explores for the first time the use of an ensemble Kalman filter (EnKF) to assimilate single-Doppler radar observations in a cloud-scale model. The EnKF is a novel and flexible technique for data assimilation, first proposed in the geophysical literature by Evensen [1994; but see also Leith (1983, 375–377)]. The EnKF uses forecast covariances between observed and unobserved variables to spread information from the observations both spatially and between variables. These covariances are estimated from an ensemble of prior forecasts initialized when observations were last available. Section 2 provides further background on the EnKF.

Other approaches for analyzing the atmospheric state given radar observations fall into two categories: retrieval algorithms and four-dimensional variational assimilation (4DVAR). Retrieval algorithms typically begin by estimating the wind field given observations of radial velocity from one or more radars. Since only certain components of the velocity are observed, the entire velocity field is estimated using additional information or constraints (such as continuity, the prognostic

* The National Center for Atmospheric Research is sponsored by the National Science Foundation.

Corresponding author address: Dr. C. Snyder, NCAR, P.O. Box 3000, Boulder, CO 80307-3000.
E-mail: chriss@ucar.edu

equation for radial velocity, or the evolution of the reflectivity field). Given knowledge of the wind field, the thermodynamic variables are then estimated from the vertical momentum equation after first calculating pressure perturbations at each level from the horizontal momentum equations. Recent results from different algorithms, as well as earlier references, can be found in Xu et al. (2001), Montmerle et al. (2001), and Weygandt et al. (2002).

Four-dimensional variational schemes seek to fit a numerical simulation to observations spread over a time interval by adjusting the state at the beginning of the interval. Like retrieval algorithms, 4DVAR has shown practical value in estimating unobserved variables given radar observations (Sun and Crook 1997, 1998). Unlike the more empirical retrieval algorithms, however, 4DVAR analyzes all variables in a unified fashion. 4DVAR also allows the systematic treatment of observation error and information from recent forecasts, but covariance matrices for both the observation and forecast error must be specified. Further background on 4DVAR appears in Sun and Crook (1997), while Rabier et al. (2000) discuss the implementation of 4DVAR for an advanced numerical weather prediction model.

The EnKF has several appealing properties that motivate its consideration as an alternative to retrieval algorithms or 4DVAR. Like 4DVAR, it is a statistical scheme that handles uncertainty in the observations and the prior forecast gracefully and approximates, subject to certain assumptions, the Bayesian update for the forecast state given new observations. The EnKF also provides direct estimates of the forecast covariances from the forecast ensemble and then explicitly updates that ensemble to be consistent with the uncertainty of the analysis. Thus, because it produces not only an analysis (the ensemble mean) but also an ensemble of initial conditions, the EnKF is a natural foundation for ensemble forecasting schemes. This is particularly important at convective scales, where probabilistic forecasts may be warranted even at lead times of an hour. Finally, the EnKF is both relatively simple to implement, as tangent linear and adjoint versions of the forecast model are not required, and relatively straightforward to parallelize.

While applications of the EnKF to large-scale flows have progressed nearly to the point of operational testing (Mitchell et al. 2002; further references appear in section 2), similar successes are not guaranteed at convective scales, where the motions are fully three dimensional, are driven by distinctly nonlinear microphysical processes, and lack the approximate balances between the mass and wind fields, such as geostrophy, that pertain at larger scales. In order to test the EnKF at convective scales and with radar observations, we extract simulated radial-velocity observations from a numerical simulation of a splitting supercell and then use the EnKF to assimilate those observations. A detailed description of the numerical experiments appears in section 3.

The results of these experiments show that the EnKF

has strong potential for convective-scale assimilation; specifically, the state of the simulated supercell can be accurately estimated given an initial sounding for the domain and several volume scans of radial velocity. Section 4 presents these results, together with some discussion of the characteristics of the forecast covariances. Since supercells are robust solutions, in the sense that a variety of initial conditions can and will produce qualitatively similar supercells given an appropriate environmental sounding, one might worry that prediction and data assimilation for a system of supercells are anomalously easy. We show in section 5 that forecast errors in fact grow significantly even in our idealized system. Section 6 addresses tuning the filter, the importance of the cross covariances between radial velocity and temperature or moisture or rainwater, and the role of the initialization of the ensemble. We summarize in the final section and outline a number of interesting issues that are not addressed in this initial study.

2. The ensemble Kalman filter

This section provides further background on the ensemble Kalman filter. We begin with a heuristic look at the central ideas in section 2a, then informally review the Kalman filter (section 2b), and finally detail the specific ensemble technique employed here (section 2c). Readers familiar with the EnKF may wish to skip directly to section 3, while those interested in a more complete and rigorous introduction to estimation theory and the Kalman filter are referred to Cohn (1997). Further discussion of aspects of the EnKF, together with results for large scales in the atmosphere, can be found in Houtekamer and Mitchell (1998), Hamill and Snyder (2000), Anderson (2001), Whitaker and Hamill (2002), and Mitchell et al. (2002). There is also a substantial oceanic literature; see Brusdal et al. (2003), Keppenne and Rienecker (2002), and references therein.

Except for the dimensions of the state and observation vectors, our notation will follow that of Ide et al. (1997).

a. Basic notions

Suppose that at some time, $t = t_k$, and over some region of interest, we possess an ensemble of forecasts that are representative, in a sense to be made more precise below, of the forecast uncertainty. Typically, the ensemble-mean forecast is the best prediction of the atmospheric state at t_k . Now suppose we receive an observation valid at the same time t_k . Our problem is then to assimilate the observation or equivalently to update our prediction of the state and its uncertainty given this new observation. While the observation obviously carries information about the measured quantity, we wish also to extract information about the unobserved portion of the state. How can this be done?

Figure 1a portrays the situation schematically. For definiteness, suppose that the observation is of radial

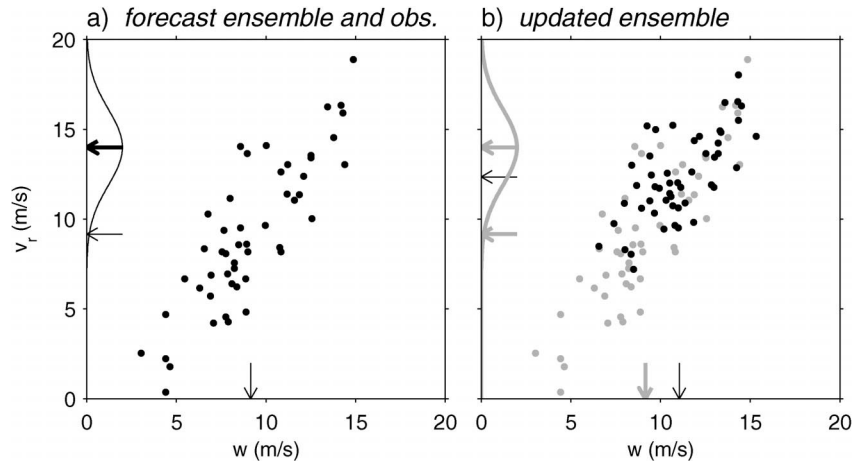


FIG. 1. Schematic depiction of the EnKF update (or analysis). (a) Forecast values of w and v_r for an ensemble of 50 members (black dots). Thin arrows on each axis show the ensemble means, each roughly 9 m s⁻¹. A hypothetical observation of v_r equal to 14 m s⁻¹ is indicated by the bold arrow on the v_r axis, while the accompanying curve is meant to suggest the observational uncertainty. (b) As in Fig. 1a but for the updated ensemble given the observation of v_r . The forecast ensemble and observation are reproduced in gray for reference.

velocity v_r at some location and the unobserved variable is vertical velocity w at another location. The forecasts from each ensemble member for v_r and w are shown as a scatterplot; in this case, the sample of ensemble members suggests that v_r and w are positively correlated in the forecast. Our best forecasts of each variable, the ensemble means, are indicated on each axis by the thin arrows. The observed value of v_r is assumed to be 14 m s⁻¹ and is marked on the vertical axis by a thick arrow, along with the specified observational error distribution.

Intuitively, we expect that the best estimate for v_r should lie between the observed value and the mean forecast, and that the uncertainty in this updated, or posterior, estimate should be smaller than that of either the observation or the forecast. In addition, it is clear that refining our estimate of v_r provides information on w ; the updated estimate of w should increase, since v_r and w were positively correlated in the forecast and the observation indicated that v_r was larger than was forecasted.

The updated situation given the observation is shown schematically in Fig. 1b. The updated ensemble is again displayed as a scatterplot and the updated means indicated by arrows on each axis, while the forecast and observed quantities from Fig. 1a are shown in gray. As intuition suggested, the updated ensemble has less spread in both variables, reflecting the additional information provided by the observation; the updated estimate of v_r lies between the forecast mean and the observation; and the mean of w has increased from the forecast value. The EnKF, which provided the updated ensemble in Fig. 1b, will be discussed below after a brief review of the Kalman filter.

b. The Kalman filter

Let \mathbf{x} of dimension N_x be the state of the system in some discrete representation, such as values on a regular grid. Thus, \mathbf{x} consists of all gridpoint values for all variables concatenated into a single vector of length N_x . For the purposes of this paper, we will ignore the often prickly issues that surround projecting the continuous atmosphere onto a discrete representation and accounting for the uncertainties associated with that projection (see Cohn 1997).

Since we have a limited set of imperfect observations, the true state of the system, denoted by \mathbf{x}' , cannot be determined precisely. It is therefore convenient to consider \mathbf{x}' to be a random variable; the most that can be known about the system is then $p(\mathbf{x}')$, the probability distribution function (pdf) of \mathbf{x}' . Our goal becomes to estimate and forecast $p(\mathbf{x}')$ given the available observations.

Now suppose, as in the section above, that we have made a forecast of $p(\mathbf{x}')$ at a time t_o and possess a set or vector \mathbf{y}^o of N_y observations also valid at t_o .¹ Subject to two assumptions, the Kalman filter provides formulas for calculating $p(\mathbf{x}' | \mathbf{y}^o)$, the pdf of \mathbf{x}' given the observations \mathbf{y}^o . The first assumption is that the observations are linearly related to \mathbf{x}' :

$$\mathbf{y} = \mathbf{H}\mathbf{x}' + \boldsymbol{\epsilon}, \quad (1)$$

where \mathbf{H} is a $N_y \times N_x$ matrix mapping the state variables

¹ To be more precise, $p(\mathbf{x}')$ is the pdf for \mathbf{x}' at $t = t_o$ conditioned on all observations prior to \mathbf{y}^o ; this is why we refer to $p(\mathbf{x}'(t = t_o))$ as a forecast. Since all the equations in this section are valid at $t = t_o$ and $p(\mathbf{x}')$ is always conditioned on the prior observations, we will suppress explicit references to t_o or the observations prior to t_o .

onto the observations, and ϵ is a random error vector of dimension N_y that is independent of \mathbf{x}^t . The second assumption is that both the prior (or background) forecast of $p(\mathbf{x}^t)$ and the pdf of ϵ are Gaussian, where the forecast has mean $\bar{\mathbf{x}}^f$ and covariance \mathbf{P}^f and the observation error has zero mean and covariance \mathbf{R} .

Under these assumptions, $p(\mathbf{x}^t | \mathbf{y}^o)$ is also Gaussian, with mean $\bar{\mathbf{x}}^a$ and covariance \mathbf{P}^a given by the Kalman filter analysis equations,

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \mathbf{K}(\mathbf{y}^o - \mathbf{H}\bar{\mathbf{x}}^f), \quad (2a)$$

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}^f, \quad (2b)$$

where

$$\mathbf{K} = \mathbf{P}^f\mathbf{H}^T(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}. \quad (3)$$

In order to assimilate subsequent observations, a method is also required to predict the forecast covariances at later times given \mathbf{P}^a ; we will discuss the covariance propagation used in the EnKF in section 2c below. Since

$$\mathbf{P}^f\mathbf{H}^T \equiv \text{Cov}(\mathbf{x}^t, \mathbf{x}^t)\mathbf{H}^T = \text{Cov}(\mathbf{x}^t, \mathbf{H}\mathbf{x}^t),$$

$\mathbf{P}^f\mathbf{H}^T$ is the forecasted covariance of the state and observed variables. We will define $\mathbf{P}_{xy}^f = \mathbf{P}^f\mathbf{H}^T$ to reflect this.

Interpretation of (2) is simplest for a single, scalar observation y^o . In that case, \mathbf{P}_{xy}^f is a column vector \mathbf{c} of dimension N_x , $\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R}$ is a scalar d , and the update of the i th element of the mean becomes

$$\bar{x}_i^a = \bar{x}_i^f + c_i(y^o - \mathbf{H}\bar{\mathbf{x}}^f)/d. \quad (4a)$$

This result corresponds heuristically to Fig. 1 (setting x_i to be the vertical velocity and y to be the radial velocity): the updated estimate \bar{x}_i^a differs from the prior estimate \bar{x}_i^f by an amount proportional to the product of $y^o - \mathbf{H}\bar{\mathbf{x}}^f$ and $c_i = \text{Cov}(x_i^f, \mathbf{H}\mathbf{x}^t)$. For example, if the observation y^o is greater than its forecast value $\mathbf{H}\bar{\mathbf{x}}^f$ and if x_i^f is positively correlated with $\mathbf{H}\mathbf{x}^t$ ($c_i > 0$), then the analyzed estimate \bar{x}_i^a should be greater than the forecast \bar{x}_i^f .

The update of the covariance matrix also simplifies for a single observation, with (2b) becoming

$$\mathbf{P}^a = \mathbf{P}^f - \mathbf{c}\mathbf{c}^T/d. \quad (4b)$$

If we think of the forecast uncertainty of a state variable as being measured by its variance (the corresponding diagonal element of \mathbf{P}^f), then the observation reduces the uncertainty most in state variables that have large covariance with the observed quantity.

We will use one further property of the Kalman filter to simplify the implementation of the EnKF below. If the observations (1) have independent errors, then the update (2), which treats all observations (i.e., each element of \mathbf{y}) simultaneously, is equivalent to serial assimilation of individual observations through repeated application of (4). To be more precise, one can assimilate the first element y_1^o of \mathbf{y}^o using (4) and then assimilate subsequent y_j^o again using (4) but with \mathbf{P}^f replaced in the definition of \mathbf{c} by the \mathbf{P}^a calculated for y_{j-1}^o , or

equivalently using the expectation conditioned on y_1^o, \dots, y_{j-1}^o to calculate $\mathbf{P}_{xy}^f = \text{Cov}(\mathbf{x}^t, \mathbf{H}\mathbf{x}^t)$. We emphasize that independent observation errors and serial assimilation of observations are not required for the EnKF but do result in a simpler algorithm.

Finally, it is important to emphasize that the Kalman filter update is optimal only in the case of linear observation operators and Gaussian errors [see (1) and the associated discussion]. Convective-scale dynamics, however, are often nonlinear, in part because of the importance of latent heat release and other microphysical processes. Thus, forecast pdfs will potentially be non-Gaussian and the Kalman filter update will at best approximate the mean and covariance of $p(\mathbf{x}^t | \mathbf{y}^o)$.

c. The algorithm

As discussed above, the Kalman filter is a scheme for updating $p(\mathbf{x}^t)$ given observations \mathbf{y}^o . The idea underlying ensemble filtering techniques is to work with a sample, or ensemble, drawn from $p(\mathbf{x}^t)$ rather than with $p(\mathbf{x}^t)$ itself. This notion was first proposed for data assimilation in the geophysical literature by Evensen (1994) but is of course at the heart of ensemble forecasting as well.

The algorithm proceeds in two steps, an analysis or update step and a forecast or propagation step, both of which are described in detail below. The analysis step begins with an ensemble drawn from $p(\mathbf{x}^t)$ and converts this ensemble into a sample from the updated distribution $p(\mathbf{x}^t | \mathbf{y}^o)$ conditioned on the most recent observations \mathbf{y}^o . The analysis proceeds according to the Kalman filter equations (2)–(4) but with the required means and covariances replaced by the sample values

$$\hat{\mathbf{x}}^f = N_e^{-1} \sum_{n=1}^{N_e} \mathbf{x}_n^f, \quad (5)$$

$$\hat{\mathbf{P}}_{xy}^f = (N_e - 1)^{-1} \sum_{n=1}^{N_e} (\mathbf{x}_n^f - \hat{\mathbf{x}}^f)(\mathbf{H}\mathbf{x}_n^f - \mathbf{H}\hat{\mathbf{x}}^f)^T. \quad (6)$$

In the forecast step, forecasts are made from each ensemble member produced in the analysis step; this is a Monte Carlo approximation to propagating the $p(\mathbf{x}^t | \mathbf{y}^o)$ forward to the time of the next observation. (Imperfections of the forecast model should also be accounted for in this step, although this issue will not arise in the present experiments.) The forecast ensemble is then used in the next analysis step and the algorithm continues.

As an aside, we note that this algorithm unifies ensemble forecasting and data assimilation. The ensemble of forecasts provides the statistical information, required for data assimilation, concerning the uncertainty at a given time and location and the relations between observations and state variables. The analysis step, in turn, explicitly constructs an appropriate analysis ensemble to serve as initial conditions for subsequent forecasts.

1) THE ANALYSIS STEP

In the analysis step, observations are processed serially, as described above. For each observation, the algorithm first calculates the mean of the prior ensemble from (5) and the difference of each member from the mean. The ensemble mean is then updated according to (4a), but replacing c_i and d with sample covariances calculated from the ensemble.

In addition to updating the ensemble mean, the scheme must also update the difference of each member from the mean, which corresponds to the covariance update (2b) of the Kalman filter. Following Whitaker and Hamill (2002), the update of the n th member from the mean is given by

$$\mathbf{x}_n^a - \hat{\mathbf{x}}^a = [\mathbf{I} - \beta(\hat{\mathbf{c}}/\hat{d})\mathbf{H}](\mathbf{x}_n^f - \hat{\mathbf{x}}^f), \quad (7)$$

where $\beta = [1 + (\mathbf{R}/\hat{d})^{1/2}]^{-1}$, and $\hat{\mathbf{c}}$ and \hat{d} are sample estimates, as in (5) and (6), of \mathbf{c} and d defined prior to (4). [See Anderson (2001) for a different, but mathematically equivalent, version of (7).] The ensemble $\{\mathbf{x}_n^a\}$ calculated in (7) then becomes the prior ensemble for the assimilation of the next observation and the algorithm continues until all observations are processed.

In practice, the update is restricted to those state variables that are within a certain radius from the observation location. We do this because state variables far from the observation location typically have small covariances c_i with the observation variable; at large distances, the sampling error (i.e., the error incurred by estimating covariances from a finite sample) then becomes comparable to or larger than c_i unless the ensemble is very large (Houtekamer and Mitchell 1998; Hamill et al. 2001). Most previous applications of the EnKF have restricted the influence of observations in the horizontal but not in the vertical; here, for convective-scale motions that are not quasi-two dimensional, we allow observations to influence only state variables within a sphere of given radius. Besides the substantial computational savings, the resulting algorithm also performs better than if each observation influenced all state variables (Houtekamer and Mitchell 1998; Anderson 2001; Hamill et al. 2001). Houtekamer and Mitchell (2001) present a more sophisticated approach to restricting the influence of an observation on distant state variables; this will be discussed further in section 6b.

Two additional points deserve comment. First, this ensemble-based algorithm asymptotically approaches the Kalman filter update in the limit of a large ensemble and Gaussian distributions; differences arise only from the approximation of covariances by their sample values. Second, since (7) is deterministic, the ensemble produced by this scheme is not strictly a random sample from $p(\mathbf{x}' | \mathbf{y}^o)$, and the ensemble is perhaps better thought of a set of states whose sample covariance approximates \mathbf{P}^a .

2) THE FORECAST STEP

Given the analysis ensemble, the forecast step simply involves forecasting each member forward to the time of the next available observations. This procedure is a Monte Carlo approximation to the computationally overwhelming propagation of the full pdf $p(\mathbf{x}')$ forward in time, at least to the extent that the analysis ensemble is a random sample from $p(\mathbf{x}')$. Moreover, if the forecast dynamics is nonlinear, this procedure generalizes the covariance propagation step of the extended Kalman filter, in which the analysis error covariance matrix is propagated by the tangent linear dynamics based on a linearization about the nonlinear forecast trajectory from the analysis mean. If the forecast dynamics is linear, the forecast step, like the analysis step, approaches that of the Kalman filter in the limit of a large ensemble.

In practice, forecast models are typically imperfect and the forecast step should account for such imperfections, either through the addition of noise to each member after the forecast (see, e.g., Mitchell and Houtekamer 2000) or through the incorporation of stochastic noise terms in the model itself. The simplified experiments presented here, however, all assume a perfect forecast model.

3. Description of the assimilation experiments

We will test the EnKF using simulated observations of radial velocity from an isolated supercell thunderstorm. In the experiments, a numerical model first produces a reference solution of the supercell. Simulated observations are then constructed by adding random observational error to the radial velocity from the reference solution, and those observations are assimilated using the EnKF and the same numerical model. These experiments suffice to demonstrate the feasibility and potential of the EnKF for convective-scale data assimilation, but it is important to note that, in any practical implementation, we will have neither a perfect forecast model nor complete knowledge of the observation errors.

a. The reference solution

The reference solution begins with a warm, moist bubble in a horizontally uniform environment; that is, u , v , θ , and q , vary only with height outside the bubble, and w is zero. The environmental sounding (Fig. 2) is based on the Oklahoma City sounding from 0000 UTC 25 July 1997, where 7 m s^{-1} has been subtracted from the zonal wind in order to minimize the movement of the right-moving supercell through the domain. The warm bubble initiates a convective cell, which first forms rain after 20 min of simulation. As is common for soundings as in Fig. 2, the initial cell splits into a strong primary supercell that moves to the right of the environmental shear and a weaker, secondary supercell




FIG. 2. Skew T diagram for the environmental sounding. Temperature and dewpoint ($^{\circ}\text{C}$) profiles are indicated by thick solid and thick dashed lines, respectively. Wind vectors are shown at the right (half barbs, 2.5 m s^{-1} ; full barbs, 5 m s^{-1} ; flags, 25 m s^{-1}).

moving to the left of the shear. Splitting of the initial cell occurs at about 55 min, and the left-moving cell passes from the computational domain after 100 min. Snapshots of the reference solution will be shown in section 4.

The numerical model used in the reference simulation (and in the assimilation experiments) is that of Sun and Crook (1997) and is documented in detail there.² Briefly, the model solves the nonhydrostatic equations using θ_l , the liquid-water potential temperature, as the thermodynamic variable and including only warm-rain microphysics. The equations are discretized spatially using second-order centered differences and a second-order Adams-Bashforth time step is used. The lateral boundary conditions depend on the sign of the normal component of velocity at the boundary. Where there is flow into the domain, gradients normal to the boundary are computed by assuming that, outside the domain, each variable is given by the environmental sounding, while at out-flow boundaries normal gradients are computed using one-sided differences.

For all simulations, the computational domain is a $70\text{-km} \times 70\text{-km}$ square in the horizontal with 2-km grid

² In a more recent version of the model, developed after the bulk of the results reported here, the numerical smoothing algorithms have been modified and the effective dissipation in the model has been significantly reduced (J. Sun 2002, personal communication). We have repeated a limited number of experiments using the newer version of the model and find no qualitative change in our results, although the new version produces only a single, right-moving supercell when initialized as in our reference simulation.

spacing and extends from surface to 17 km in the vertical with 500-m grid spacing. The origin of the Cartesian coordinates (x, y, z) is taken to lie at the lower left (southwest) corner of the domain.

b. The experiments

Simulated Doppler radar wind observations are extracted from the reference solution as follows. We assume that 1) the radar is located at the southwest corner of the computational domain, at $(x, y, z) = (0, 0, 0)$; 2) it measures v_r , the radial velocity in a spherical coordinate system centered on the radar; 3) the observations have independent, Gaussian random errors of zero mean and variance $R = 1 \text{ m}^2 \text{ s}^{-2}$; and 4) observations are available at 5-min intervals and at each grid point where the rainwater $q_r > 0.13 \text{ g kg}^{-1}$. While observations of radar reflectivity undoubtedly will provide useful information in practice, they are not assimilated in the present experiments for simplicity.

The observations are thus related to the reference state by

$$v_r = (x/r)u + (y/r)v + (z/r)w + \epsilon, \quad (8)$$

where $r = (x^2 + y^2 + z^2)^{1/2}$ and ϵ is drawn from $N(0, R)$. Note also that the dependence of v_r on the fall speed of rain has been neglected; its inclusion has no qualitative influence on the results below. Given velocities on the computational grid, v_r is calculated by first averaging u , v , and w to q_r grid points from the two adjacent, staggered grid points for each velocity component, and then using the averaged velocities in (8).

The assimilation experiments begin at $t = 20$ min, when rain first begins to form in the reference simulation. Observation sets, consisting of v_r at all points with q_r exceeding the threshold given above, are then assimilated every 5 min thereafter. The analysis variables are the same as the forecast model's prognostic variables (u , v , w , water vapor, and q_r). A typical observation set includes $O(10^3)$ individual observations.

The initialization of the ensemble begins with the environmental sounding shown in Fig. 2, which is assumed known. (Estimation of the environmental sounding is an outstanding problem for convective-scale forecasting and will be addressed in a subsequent study.) Each ensemble member is initialized at $t = 0$ by adding realizations of Gaussian noise to the environmental sounding. This noise is independent at each grid point and for each variable, has zero mean, and has standard deviation 3 m s^{-1} for each velocity component and 3 K for θ_l . Water vapor and cloud water are initialized using the environmental sounding at each level.

Our choice of these statistics for the initial ensemble is motivated entirely by simplicity. As will be discussed in sections 6c and 7, more sophisticated initializations are possible and, because of the relatively short duration of the experiments, will influence the performance of the EnKF. Except where noted, the EnKF uses 50 mem-

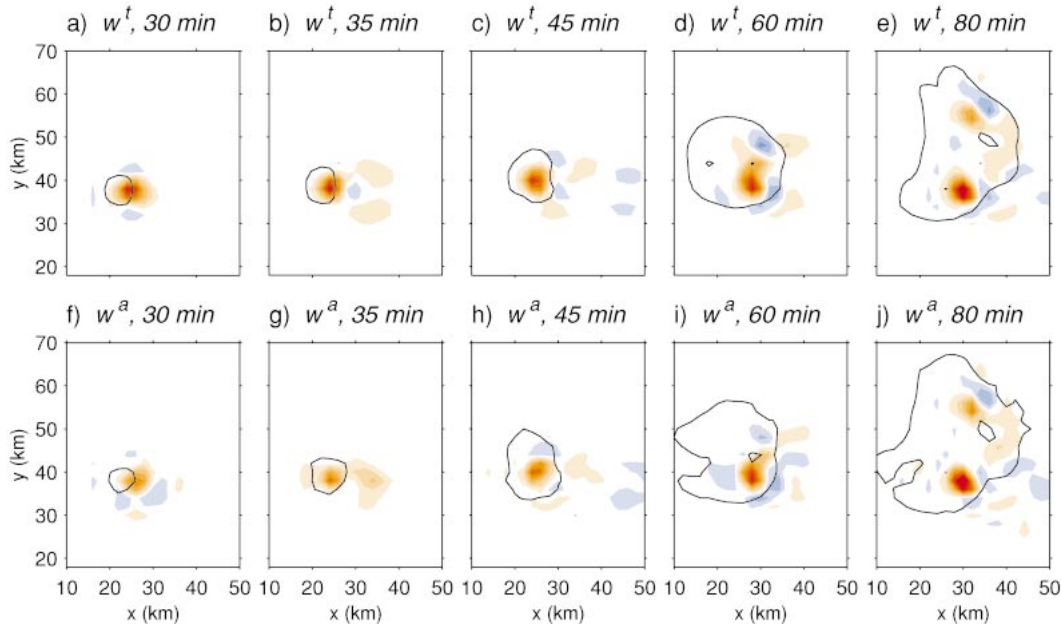


FIG. 3. Vertical velocity at $z = 6$ km in (a)–(e) the reference simulation (w^t) and (f)–(j) the EnKF analysis (the ensemble mean). Shades of red and blue indicate upward and downward motion, respectively, with gradations of color every 2.5 m s^{-1} beginning at $\pm 1.25 \text{ m s}^{-1}$ and up to a maximum of 26.25 m s^{-1} . Contours of the -0.75-K temperature perturbation at $z = 1$ km are also displayed (black lines). Fields are shown at $t = 30, 35, 45, 60,$ and 80 min, as marked on each panel.

bers in all experiments, and each observation is allowed to influence state variables within a sphere of radius 4 km.

4. Results

This section presents results from the basic assimilation experiment outlined above. We first show that the EnKF analyses closely approximate the reference solution after a few assimilation cycles and then discuss the characteristics and quality of the ensemble covariances.

a. Ensemble-mean analyses

Figure 3 compares the ensemble-mean analysis of vertical velocity (\bar{w}^a) with the reference solution (w^t) at several times. The reference solution evolves as described in section 3a, with the initial cell splitting into long-lived left- and right-moving supercells (Figs. 3a–e). At $t = 30$ min (after three assimilation cycles; Fig. 3f), the analysis suggests an updraft in rough agreement with the reference solution. The structure of the analyzed updraft and its relation to the buoyancy and rainwater, however, are sufficiently in error that the cell decays during the next 5-min forecast and remains too weak in the analysis at $t = 35$ min (Fig. 3g). By $t = 45$ min (Fig. 3h), the analysis approximates the location, size, and strength of the main updraft. The analysis continues to improve beyond this time and after an hour of

assimilation ($t = 80$ min; Fig. 3j) captures much of the detailed structure of the reference solution.

The analysis also faithfully approximates the thermodynamic variables, which unlike w have no direct influence on the observations. Figure 3 also shows the -0.75-K contour of the temperature perturbation at a height of 1 km, which broadly outlines the low-level cold pool. As for the vertical velocity, the analysis contains some information after a few assimilation cycles and then beyond $t = 60$ min becomes increasingly detailed and accurate.

The overall quality of the ensemble-mean forecasts and analyses can be obtained from Fig. 4, where the rms errors for the horizontal wind, θ_t , w , and rainwater are shown as a function of time. Errors are averaged only over the portion of the domain where $q_r > 0.1 \text{ g kg}^{-1}$ in order to provide a more accurate and sensitive measure of the analysis quality near the cells. (Over the rest of the domain, errors tend to be small simply because the variability of the reference simulation about the initial sounding is small.) As was evident in Fig. 3, the analyses improve rapidly over the first 20–30 min of assimilation. The errors in all fields then level off at a magnitude that is small compared to the $O(10 \text{ m s}^{-1}, 10 \text{ K})$ variations typically found near the convective cells in the reference solution.

The results shown to this point are all based on a single set of initial ensemble members. Since the initial ensemble is drawn randomly with a specified pdf, as are the observation errors, one may expect some random




FIG. 4. The rmse of the ensemble mean from the EnKF, averaged over all points at which $q_r > 0.1 \text{ g kg}^{-1}$, for four quantities: horizontal wind (vector magnitude in m s^{-1} ; thick solid lines), w (m s^{-1} ; gray), θ_i (K; dotted), and rainwater (g kg^{-1} ; thin solid). Errors for both the forecast and analysis means are shown at each analysis time, producing the “sawtooth” appearance of the curves. Results for the first analysis (at $t = 20$) are omitted because only 12 observations are available at that time.

variation of the results for different realizations of these random quantities. To quantify this variation, the rms error of the ensemble-mean analysis of w is shown in Fig. 5 for 12 different realizations. Variations are most significant over the first four cycles (up to $t = 40$ min), after which all realizations reach similar error levels.

b. Individual members

Before turning to the covariance information in the ensemble, it is useful also to examine how individual members behave. The vertical velocity from the first and second members of the ensemble is shown in Fig. 6 at $t = 80$ min. As would be expected given the excellent agreement between the ensemble mean and the reference solution at this time, each member is similar to the reference solution in the region near the actual convective cells where observations are available. They differ elsewhere. In particular, a line of spurious convective cells is evident in the second member along the northern boundary of the domain ($y = 70$ km).

Somewhat less than one-half of the members possess such spurious cells at $t = 80$ min. These cells can be traced to the initialization of the members, which typically excites a few weak cells in each member spread throughout the domain. The spurious cells that survive do so because they are located away from v_r observations, so their evolution is not altered during analysis updates. Even the surviving cells exhibit slow decay, however, since convective cells prefer to be widely separated, as their narrow plume of ascent and broad region of subsidence suppress nearby cells. Since the observed cells are continually reinforced during the assimilation, their subsidence gradually weakens spurious cells.




FIG. 5. The rmse of the ensemble-mean analysis of w at various times and for 12 realizations of the initial (random) ensemble perturbations and the observation errors. For clarity, four of the realizations are indicated by thin black lines, four others by gray lines, three others by dotted lines; the realization shown in Figs. 2–4 is indicated by a thick black line.

Finally, we note that although many members have spurious cells, their positions and strength are largely a random consequence of the ensemble initialization. Thus, the spurious cells average to (nearly) zero in the ensemble mean.

c. Ensemble covariances

Within the EnKF, the forecast ensemble provides the covariance information required in the assimilation of observations. This section discusses two issues related to the ensemble covariances: their consistency with the statistics of the error of the ensemble-mean forecasts and analyses, and their form and characteristics.

The issue of consistency arises because we are using the variations of the ensemble about its mean to estimate the statistics of the error of the ensemble mean. To see this, recall from section 2 that the KF assumes $p(\mathbf{x}^f)$, the forecast pdf for the true state, to have mean $\bar{\mathbf{x}}^f$ and covariance \mathbf{P}^f . Thus, the error of the mean, $\mathbf{x}^f - \bar{\mathbf{x}}^f$, also has covariance matrix \mathbf{P}^f . The EnKF estimates $\bar{\mathbf{x}}^f$ by the ensemble mean $\hat{\mathbf{x}}^f$ and \mathbf{P}^f by the sample covariance $\hat{\mathbf{P}}^f$ [defined as in (6)]. A simple consistency relation then follows: the ratio of the expected total variance of the ensemble, $E[(N_e - 1)^{-1} \sum |\mathbf{x}_n^f - \hat{\mathbf{x}}^f|^2]$, to the expected squared error of the ensemble mean, $E(|\mathbf{x}^f - \hat{\mathbf{x}}^f|^2)$, is equal to $N_e/(N_e + 1)$ (Murphy 1988).

Figure 7 shows this ratio in the present experiments, where we have replaced the expected values in both the numerator and denominator with averages over the 12 realizations shown in Fig. 5. In addition, both the squared error and the ensemble variance are again (as in Figs. 4, 5) summed over only those grid points where $q_r > 0.1 \text{ g kg}^{-1}$. The ensemble variance at the first analysis time is a factor of 2–4 smaller than the error of the ensemble mean. Their ratio then increases steadily

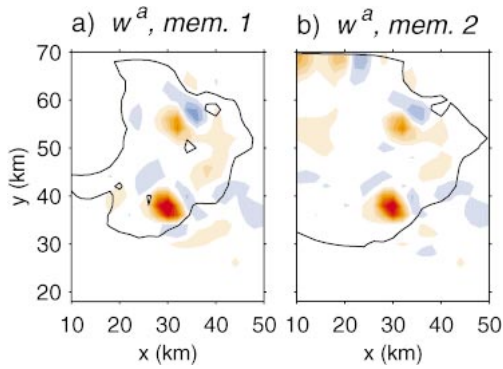


FIG. 6. As in Fig. 3, but showing the vertical velocity for the first and second members of the ensemble at $t = 80$ min.

with time, although it remains generally less than $N_e / (N_e + 1)$.

The results of Fig. 7 suggest that there is scope to improve the performance of the EnKF. For example, the ratio of variance to squared error at early times is determined (and at later times, influenced) by our choice of the ensemble's initial variance. Although we have not tested this possibility, the ratio shown in Fig. 7 could likely be improved by increasing the initial variance. Further discussion of tuning the EnKF to optimize its performance appears in section 6.

The increase of the ratio with time differs from the behavior found in other implementations of the EnKF. In those implementations, the ratio typically decreases through successive assimilation cycles (e.g., Houtekamer and Mitchell 1998) owing to a systematic underestimation of the analysis variance by the EnKF.³

To understand why the ensemble variance steadily grows relative to the error of the ensemble mean, recall from Fig. 6 that, throughout the experiment, many members retain spurious convective cells in unobserved portions of the domain. During the forecast for a given member, these spurious cells (if present) interact with the “observed” cells and increase the rate at which the forecast of the observed cells diverges from the reference simulation. The ensemble-mean forecast, however, diverges more slowly from the reference solution, since the spurious cells are spread randomly through the unobserved areas and their averaged effect on the mean is small (except perhaps for some spatial smoothing of the observed cells). Thus, consistent with the shape of the curves in Fig. 7, the ensemble variance increases more rapidly than the squared error during forecasts and

³ One reason for this can be seen from (4b), the scalar-observation update for the KF covariances, which shows that the total analysis variance $\text{tr}(\mathbf{P}^a)$ is reduced by an amount proportional to $\text{tr}(\mathbf{c}\mathbf{c}^T) = \mathbf{c} \cdot \mathbf{c}$ relative to the forecast variance. When \mathbf{c} is estimated from a finite sample, sampling error biases the estimate of $\mathbf{c} \cdot \mathbf{c}$ to be too large, so that the estimate of \mathbf{P}^a is correspondingly too small. See van Leeuwen (1999) for a more rigorous analysis.




FIG. 7. Ratio of the ensemble variance to the squared error of the ensemble mean for horizontal wind (thick solid lines), w (gray), and θ_l (dotted). Both the numerator and denominator of the ratio are averaged over 12 realizations of the experiment. The horizontal line indicates $N_e / (N_e + 1)$, the ratio of the expected values of sample variance and squared error for an ensemble drawn from the same distribution as the reference solution.

overcomes the tendency for the EnKF update to underestimate the analysis variance.

Consideration of the ratio of variance to squared error over the entire domain (not shown) also supports this view. The ensemble variance is initially uniform throughout the domain, yet away from the observed cells the errors of the ensemble mean are small since the motions themselves are weak. The ratio of variance to squared error is then larger than $N_e / (N_e + 1)$ (typically between 2 and 4); this “extra” uncertainty outside the observed regions contaminates forecasts in the observed regions and leads to more rapid growth of variance than squared error there.

We emphasize that the spurious cells in some members are a direct consequence of the initialization of the ensemble with spatially white noise throughout the domain. Thus, the choice of the initial ensemble strongly influences the results of diagnostics, such as shown in Fig. 7, even at $t = 100$ min. Indeed, as will be shown in section 6c, the initial ensemble also affects the performance of the EnKF throughout the assimilation experiments.

We now turn to the ensemble estimates of the forecast covariances. These are of interest both because little is known about them for convective-scale prediction and because they provide some justification for the radius of influence assumed in our implementation of the EnKF, although that justification is by no means complete.

The left-hand panels of Fig. 8 display variances (i.e., diagonal elements of $\hat{\mathbf{P}}^f$) for the ensemble of 5-min forecasts valid at $t = 80$ min. The variances are shown in a vertical cross section through the updraft of the right-moving supercell. The variance of both w^f (Fig. 8a) and θ_l^f (Fig. 8e) coincides with the updraft and its




FIG. 8. Variances and correlations estimated from the ensemble at $t = 80$ min in the x - z plane along $y = 36$ km, which passes through the maximum updraft. (left) Variances of (a) w^f , (c) u^f , and (e) θ_f^f are indicated by black contours (at 0.2, 0.4, and 0.8 of the maximum in the cross section); (right) the correlation (black lines, negative values dashed; contours at ± 0.3 and ± 0.6) between the forecast value of v_r at the point $x = 30$ km, $z = 5$ km (indicated by a black dot) and (b) w^f , (d) u^f , and (f) θ_f^f at each point in the cross section. Shades of red and blue indicate positive and negative values, respectively, of fields from the reference solution: (a), (b) w^f , with shading increments every 4 m s^{-1} beginning at $\pm 2 \text{ m s}^{-1}$; (c), (d) u^f , with shading as in Figs. 8a,b; and (e), (f) θ_f^f , with shading increments every 4 K beginning at $\pm 2 \text{ K}$.

accompanying θ_f deficit, while $\text{Var}(u^f)$ has two maxima near $z = 10$ km, one just upstream of the updraft and the other extending downstream from the updraft. None of the variance fields has pronounced maxima near the surface.

As might be expected given that the EnKF analyses estimate unobserved variables, the ensemble reveals significant correlations between v_r and the state variables. Three examples appear in the right-hand panels of Fig. 8, which show the correlation of w^f , u^f , or θ_f^f at each point in the cross section with the forecast value of v_r at a point at the base of the updraft. Like the variances, the correlations reflect the structure of the reference solution: for w^f and θ_f^f , significant correlations (negative for w and positive for θ_f) extend along the height of the updraft and its accompanying buoyancy deficit. For u^f , strong positive correlations coincide with the region of

weak flow from the base of the updraft upward and downstream.

Figure 8 provides justification for our choice of a 4-km radius of influence in the assimilation, at least in its order of magnitude: that radius is comparable to the scale of variation of the covariances. Moreover, the largest covariances are found within that radius of the observation point. It is clear, however, that the 4-km radius does not include all locations with significant covariance with v_r at the observation point, nor does it exclude all locations with small covariances that are likely strongly contaminated by sampling error.

Figure 8 also demonstrates that covariances have complex structure that is highly inhomogeneous and anisotropic, and are flow dependent in that their structure is related to position and form of the convective cell in reference state. It would likely be difficult to model the

covariance structure and its relation to reference state with just a few parameters.

5. Forecast-error growth

In many geophysical systems, the accuracy of state estimates is limited in part by forecast-error growth. It is not obvious, however, that the present system of two isolated supercells will behave similarly. In particular, each supercell is long lived and quasi-steady, which indicates that they are at least structurally stable to small perturbations. This in turn raises the possibility that the decrease of analysis error with time shown in Figs. 4 and 5 arises mainly because forecast errors grow slowly or decay. This section examines the forecast-error growth in our experiments.

Perhaps the simplest diagnostic is to compare the error of the ensemble-mean 5-min forecast to that of the preceding analysis. As can be seen in Fig. 4, errors do typically grow over the course of the forecast, although there are instances in the initial few cycles in which decay occurs.

Errors also grow during longer forecasts. Figure 9a shows rms errors for w for forecasts beginning from the ensemble-mean analysis at various times. (Although they are not shown, forecast errors for other variables behave similarly. Also note that in Fig. 9 the errors are averaged over the entire domain.) The time required for the error to double is roughly 10–20 min in all cases, consistent with the dynamical timescale of the supercells. Moreover, it is clear that the assimilation of observations with the EnKF provides a significantly better state estimate than would be available given only a forecast from an earlier time.

Although it does not prevent error growth, the structural stability of the supercells manifests itself in certain characteristics of the forecast errors. Specifically, initial errors can have two effects on the forecast: they may alter the position and the instantaneous intensity of the supercells, or they can extinguish one or both cells. At $t = 100$ min, for example, the forecasts shown in Fig. 9a each contain a right-moving supercell, displaced by a few grid points from that in the reference solution and typically weaker, while none of the forecasts, except that from $t = 80$ min when splitting has already begun, capture the left-moving cell.

In addition, the structural stability of the supercells means that smaller initial errors should produce smaller forecast errors, even beyond the time at which error growth begins to saturate. To test this possibility, we have performed additional forecasts whose initial conditions (at each of the times shown in Fig. 9a) were created by scaling the initial error by a factor of 2, 0.4, or 0.1 and then adding that rescaled perturbation to the reference solution. Figure 9b illustrates the results of these experiments in terms of the rms error of w . As expected, smaller errors at any initial time give rise to smaller forecast errors through the entire length of the

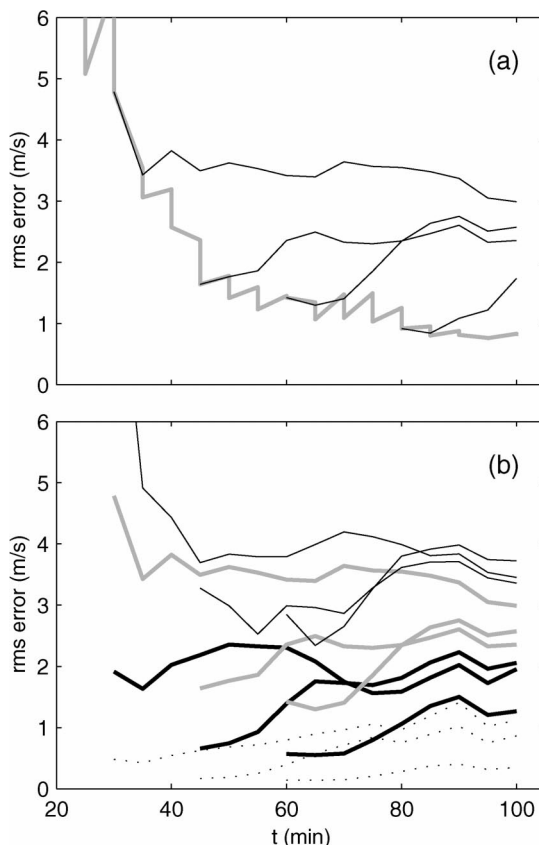


FIG. 9. (a) The rmse for forecasts of w (black lines), starting from the ensemble-mean analysis at $t = 30, 45, 60, 80$ min and averaged over the entire domain. Errors for the forecast and analysis means of w are indicated by the gray lines as in Fig. 4 (b) As in Fig. 9a, but for forecasts starting from initial conditions whose error has been rescaled by factors of 2 (thin lines), 0.4 (thick), and 0.1 (dotted). To be more precise, initial conditions for each forecast are created by calculating the error of the ensemble-mean analysis at the appropriate time, multiplying the error field by a constant scalar factor (2, 0.4, 0.1), and adding the rescaled error to the reference solution. The rmse for the original forecast (scale factor of 1) is shown in gray.

forecast. In a chaotic system, error growth would proceed from even very small initial errors until the forecast was completely decorrelated from the reference solution, and the curves in Fig. 9b would, in contrast, all approach a limiting value with time, regardless of the size of the initial error. Presumably, this would also occur in a more complex convective situation with multiple cells and outflows.

6. Other issues

a. Importance of covariance information

We have suggested that the EnKF is particularly appealing for use at convective scales because of its ability to estimate forecast covariances with a minimum of prior assumptions, and thereby infer state variables that are unobserved. Here, we test this assertion with an experiment in which the EnKF updates only the three




FIG. 10. As in Fig. 5, but for the experiment in which θ_i , q_i , and q_r are not updated in the analysis; i.e., observations of v_r influence only the velocities.

components of the wind given observations of v_r , and does not use the information available in cross covariances between v_r and θ_i , q_r , and q_i .

The rms error of the ensemble-mean analysis of w is shown as a function of time in Fig. 10 for 10 realizations of the experiments (i.e., for 10 realizations of the initial ensemble and 10 realizations of the observation errors). The error is typically more than three times what it would be if all variables were updated during the analysis (see Fig. 5), and even in the best realization the error is still doubled. The error in variables other than w behaves similarly.

Except in the single, anomalously accurate realization evident in Fig. 10, the analyses (not shown) fail to capture either the updraft or surface cold pool of the observed cell at least through $t = 45$ min. Instead, the analyses (and the individual members) contain a noisy mix of updrafts and downdrafts spread over the region of observations, with a variety of spurious cells outside that region in individual members. Beyond an hour, the analyses begin to exhibit narrow updrafts in the locations of the observed cells and small cold pools beneath these updrafts. The analyzed updrafts then decay markedly during the course of the 5-min forecasts to the next analysis time.

We have also performed experiments in which the observed variable is u and only u is updated by the EnKF, so that cross covariances between components of velocity are also ignored in the assimilation. Analysis errors are even larger and the analyses never capture the updrafts of the observed cells (not shown). Together, these two sets of experiments show that covariances estimated by the EnKF contain significant information about relations among the state variables, despite the sampling errors associated with an ensemble of 50 members.

b. Tuning of algorithm

To implement EnKF as described here, one must choose the ensemble size N_e and the radius of influence for the observations. Although we have not systematically tuned the algorithm (by determining optimal radius for a given N_e), we have performed a number of experiments with values other than $N_e = 50$ and $r = 4$ km.

It is clear for these additional experiments that the results in section 4 depend quantitatively on the choice of N_e and r . In particular, decreasing N_e to 25 increases the rms error of the ensemble-mean analyses and also makes the results less robust, in the sense that there will be realizations of the initial ensemble for which the analyses are always a poor approximation to the reference simulation, as was the case for the experiments described in section 6a above that did not update θ_i or the moisture variables. (Increasing N_e has smaller, but opposite, effects.) For sufficiently large r (roughly 20 km or larger, using $N_e = 50$), the ensemble-mean analysis noticeably deteriorates and “outlying” realizations again become frequent, although more modest variations of r , from 2 to 6 km, have little effect on the analysis quality. Such dependence on N_e and r is broadly consistent with other examples of the EnKF (e.g., Houtekamer and Mitchell 1998; Anderson 2001; Whitaker and Hamill 2002).

The results of our experiments also depend on how the analysis update is performed within the radius r . A more sophisticated approach is to decrease the influence of an observation on the state with increasing distance from the observation location by multiplying c_i in (4a) by a correlation function with local support, as in Houtekamer and Mitchell (2001). In this way, the influence of an observation on the analysis decreases smoothly to zero at finite radius, rather than jumping discontinuously to zero. Using such an approach improves the results presented here (A. Caya 2002, personal communication).

Another aspect of the algorithm that can likely be improved is our choice of the probability distribution from which the initial ensemble is drawn. In the results just presented, the initial ensemble is centered on the environmental sounding and each member is initialized with independent Gaussian noise at each grid point and in each variable. Although the initial variance in this simple scheme could presumably be tuned, we will show in section 6c below that a careful choice of that initial distribution is likely more important.

A final aspect of ensemble filters that is often subject to tuning is the covariance inflation (Anderson 2001; Hamill et al. 2001), in which the ensemble covariances are multiplied by a scalar factor slightly greater than 1 to compensate for the usual bias of the EnKF to underestimate the analysis uncertainty. We have performed experiments with several inflation factors and always find that inflation degrades the results. There appears to




FIG. 11. As in Fig. 5, but for the local perturbations experiment in which the initial ensemble perturbations are nonzero only in a $20 \text{ km} \times 20 \text{ km}$ box centered on the location of the first echoes.

be two related reasons for this: first, the tendency for the ensemble forecasts to overestimate, because of the presence of spurious cells, the growth of variance in the region of the observed cells; and second, the fact that the inflation enhances the spurious cells by increasing deviations from the ensemble mean. Since the spurious cells are an artifact of the ensemble initialization, it seems likely that covariance inflation may yet prove beneficial given a more sophisticated initialization. We leave a more systematic exploration of tuning the EnKF for subsequent work. The following section provides further discussion of the role of the initial ensemble.

c. Dependence on the statistics of the initial ensemble

Sections 4b,c discuss how the initialization of each member with noise throughout the domain leads to spurious convective cells in many members, which in turn degrade the forecasts from those members and affect the properties and performance of the EnKF. A natural question is then the extent to which the result might be improved with a different initialization of the ensemble.

One easy modification, which should reduce the spurious cells, is to restrict the initial noise in each member to the vicinity of the radar echoes. To be more specific, we have performed experiments in which the members are initialized at $t = 0$ with noise confined to a $20 \text{ km} \times 20 \text{ km}$ box centered on the location of the first echoes (i.e., the first nonzero values of q_r) at $t = 20$ min. Except for this restriction of the initial noise to a portion of the domain, these experiments are identical to those discussed in section 4. Our use of information from $t = 20$ min to initialize the ensemble at $t = 0$ is akin that of 4DVAR or the Kalman smoother (Cohn 1997), in which observations influence the state estimate at earlier times as well as the present time. Note also that this initialization is feasible in practice, as one would simply wait until observations were available and then initialize the ensemble at a somewhat earlier time.




FIG. 12. As in Fig. 7, but for the local perturbations experiment.

The rms error for w is shown in Fig. 11. Analyses are significantly improved over those based on the original initialization of the ensemble. Averaging over 12 realizations of the experiments, error at $t = 50$ min is reduced by a factor of 2, while that at $t = 100$ min is reduced by a factor of 3.

In the individual members (not shown), some spurious cells are still excited by the initial noise, but these are of course located much closer to the observed cell on average. The subsidence surrounding the observed cell then suppresses the spurious cells more strongly than in the previous experiments, so that they are weaker throughout the simulation and almost all have disappeared by $t = 70$ min. The fact that the forecasts from individual members are not degraded by the presence of spurious cells is at least one factor contributing to the improved performance of the EnKF.

In section 4c, we also argued that the presence of spurious cells in some members led to a continual increase in the ensemble variance relative to the squared error of the ensemble mean. Figure 12 shows the ratio of these quantities for the present “local perturbation” experiment and should be compared to Fig. 5. The ratio of variance to error increases through the first half of both experiments. By $t = 70$ min, the ratio stops its systematic increase in the local perturbation experiment; this coincides with the time at which most spurious cells have been suppressed.

7. Summary and discussion

In the experiments presented here, we have used the EnKF to assimilate simulated Doppler radar observations of radial velocity in a nonhydrostatic, cloud-scale model. The observations are taken (together with random observational noise) from a reference simulation of a splitting supercell storm, produced with the same numerical model. We assume that observations are available every 5 min, but only in the small fraction of the domain where the rainwater in the reference simulation exceeds a threshold.

These experiments demonstrate the potential of the EnKF for assimilation of radar data at the convective scale. Using an ensemble of 50 members, the EnKF produces analyses that accurately approximate the true state (i.e., that from the reference simulation) after about seven assimilation cycles and a half hour of observations. Variables not directly observed, including the vertical velocity and temperature, are accurately estimated in the analyses. Examination of forecasts also shows that the simple supercell simulation considered here supports significant growth of forecast error and, thus, that the ability of the EnKF to track the growth and splitting of the cell is not simply a consequence of stable system dynamics.

In principle, the crucial element of the EnKF is its direct estimate of forecast covariances between radial velocity (or other observed quantities) and the state variables. To test the importance of such covariances, we also performed experiments in which the covariances of v_r with the thermodynamic and moisture variables were set to zero in the assimilation, so that the observations did not influence the analyses of temperature, moisture, and cloud water. The rms analysis error more than tripled in these experiments.

It is worth emphasizing that the dynamics of moist convection represents a significant test for the EnKF. Unlike the large-scale flows in the atmosphere or basin scales in the ocean, which have been the setting of all previous implementations of the EnKF, convective-scale motions generally lack approximate, static balances, such as geostrophy, that link the velocity to thermodynamic fields. Our results indicate that the lack of such balances is not a fundamental obstacle to the use of the EnKF at convective scales, as the dynamically produced relations among state variables (i.e., those that arise through the evolution of the flow) are sufficiently strong and develop rapidly. In addition, moist convection is driven by distinctly nonlinear microphysical processes, such as condensation, and tends to form discrete, coherent structures, such as supercells, whose dynamics are nonlinear; both of these facts call into question the Gaussian assumptions that underlie the EnKF analysis step. The success of the EnKF in the present problem suggests that the nonlinearity inherent in moist convection is also not a fundamental obstacle to the EnKF. Nevertheless, continued development of the EnKF for convective scales will undoubtedly require further consideration of nonlinear and non-Gaussian effects, particularly in relation to reflectivity observations.

Our experiments also differ from those of previous studies with the EnKF in that they cover a limited time interval, spanning only a few dynamical timescales. The quality of the analyses, and even diagnostics of the performance of the scheme such as the ratio of variance to error, are therefore strongly influenced throughout our experiments by the choice of the initial ensemble. A particular example appears in section 6c. (In contrast, other studies have collected results over long simula-

tions of systems that possess a statistically steady state, so that the initial ensemble is not important.) Although the initial conditions for the ensemble should clearly reflect our best knowledge of $p(\mathbf{x}')$ prior to any observations, considerable latitude remains for choosing them reasonably. We expect that improving the initial ensemble and diagnosing its role are likely to be persistent issues for the EnKF at convective scales.

The performance of the EnKF relative to retrieval techniques or 4DVAR is an important question that we have not addressed in this paper. In fact, our use of the model of Sun and Crook (1997) was motivated by the possibility of comparing against the existing 4DVAR scheme for that model. Comparisons are under way and will be reported elsewhere (A. Caya and J. Sun 2003, personal communication).

There are a number of other important issues that were beyond the scope of this initial paper. These include the use of reflectivity observations; estimation of the environmental sounding and its uncertainty; accounting for imperfections in the forecast model, particularly the microphysical parameterizations; the treatment of lateral boundary conditions and their uncertainty; and quantification of errors in real radar observations and in the forward operators for both radial velocity and reflectivity, as well as quality control for radar observations. Progress on all of these issues is likely important if the EnKF, or indeed another technique, is to be applied routinely and successfully to the assimilation of radar observations.

Acknowledgments. We are particularly indebted to Juanzhen Sun for the use of her cloud model in this study. Both Alain Caya and David Dowell have shared with us their results related to various refinements of the algorithm used here. It is a pleasure to acknowledge helpful discussions with them and with William Skamarock and Jeff Anderson. This research was supported at NCAR by the U.S. Weather Research Program and by NSF Grant 0205655.

REFERENCES

- Anderson, J. L., 2001: An ensemble adjustment filter for data assimilation. *Mon. Wea. Rev.*, **129**, 2884–2903.
- Brusdal, K., J. M. Brankart, G. Halberstadt, G. Evensen, P. Brasseur, P. J. van Leeuwen, E. Dombrowsky, and J. Verron, 2003: A demonstration of ensemble based assimilation methods with a layered OGCM from the perspective of operational ocean forecasting systems. *J. Mar. Syst.*, **40–41**, 253–289.
- Cohn, S. E., 1997: An introduction to estimation theory. *J. Meteor. Soc. Japan*, **75**, 257–288.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99** (C5), 10 143–10 162.
- Hamill, T. M., and C. Snyder, 2000: A hybrid ensemble Kalman filter/3D-variational analysis scheme. *Mon. Wea. Rev.*, **128**, 2905–2919.
- , J. S. Whitaker, and C. Snyder, 2001: Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.*, **129**, 2776–2790.

- Houtekamer, P. L., and H. L. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique. *Mon. Wea. Rev.*, **126**, 796–811.
- , and —, 2001: A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, **129**, 123–137.
- Ide, K., P. Courtier, M. Ghil, and A. C. Lorenc, 1997: Unified notation for data assimilation: Operational, sequential, and variational. *J. Meteor. Soc. Japan*, **75**, 181–189.
- Keppenne, C. L., and M. M. Rienecker, 2002: Initial testing of a massively parallel ensemble Kalman filter with the Poseidon isopycnal ocean general circulation model. *Mon. Wea. Rev.*, **130**, 2951–2965.
- Leith, C. E., 1983: Predictability in theory and practice. *Large-Scale Dynamical Processes in the Atmosphere*, B. J. Hoskins and R. P. Pearce, Eds., Academic Press, 365–383.
- Mitchell, H. L., and P. L. Houtekamer, 2000: An adaptive ensemble Kalman filter. *Mon. Wea. Rev.*, **128**, 416–433.
- , —, and G. Pelerin, 2002: Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Mon. Wea. Rev.*, **130**, 2791–2808.
- Montmerle, T., A. Caya, and I. Zawadzki, 2001: Simulation of a midlatitude convective storm initialized with bistatic Doppler radar data. *Mon. Wea. Rev.*, **129**, 1949–1967.
- Murphy, J. M., 1988: The impact of ensemble forecasts on predictability. *Quart. J. Roy. Meteor. Soc.*, **114**, 463–493.
- Rabier, F., H. Järvinen, E. Klinker, J.-F. Mahfouf, and A. Simmons, 2000: The ECMWF operational implementation of four-dimensional variational assimilation. Part I: Experimental results with simplified physics. *Quart. J. Roy. Meteor. Soc.*, **126**, 1143–1170.
- Sun, J., and N. A. Crook, 1997: Dynamical and microphysical retrieval from Doppler radar observations using a cloud model and its adjoint. Part I: Model development and simulated data experiments. *J. Atmos. Sci.*, **54**, 1642–1661.
- , and —, 1998: Dynamical and microphysical retrieval from Doppler radar observations using a cloud model and its adjoint. Part II: Retrieval experiments of an observed Florida convective storm. *J. Atmos. Sci.*, **55**, 835–852.
- van Leeuwen, P. J., 1999: Comments on “Data assimilation using an ensemble Kalman filter technique.” *Mon. Wea. Rev.*, **127**, 1374–1377.
- Weygandt, S. S., A. Shapiro, and K. K. Droegemeier, 2002: Retrieval of model initial fields from single-Doppler observations of a supercell thunderstorm. Part I: Single-Doppler velocity retrieval. *Mon. Wea. Rev.*, **130**, 433–453.
- Whitaker, J. S., and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924.
- Xu, Q., H. D. Gu, and S. Yang, 2001: Simple adjoint method for three-dimensional wind retrievals from single-Doppler data. *Quart. J. Roy. Meteor. Soc.*, **127**, 1053–1067.