

Atmospheric Data Assimilation with an Ensemble Kalman Filter: Results with Real Observations

P. L. HOUTEKAMER, HERSCHEL L. MITCHELL, GÉRARD PELLERIN, MARK BUEHNER, MARTIN CHARRON, LUBOS SPACEK, AND BJARNE HANSEN

Meteorological Service of Canada, Dorval, Québec, Canada

(Manuscript received 9 October 2003, in final form 24 March 2004)

ABSTRACT

An ensemble Kalman filter (EnKF) has been implemented for atmospheric data assimilation. It assimilates observations from a fairly complete observational network with a forecast model that includes a standard operational set of physical parameterizations. To obtain reasonable results with a limited number of ensemble members, severe horizontal and vertical covariance localizations have been used.

It is observed that the error growth in the data assimilation cycle is mainly due to model error. An isotropic parameterization, similar to the forecast-error parameterization in variational algorithms, is used to represent model error. After some adjustment, it is possible to obtain innovation statistics that agree with the ensemble-based estimate of the innovation amplitudes for winds and temperature. Currently, no model error is added for the humidity variable, and, consequently, the ensemble spread for humidity is too small. After about 5 days of cycling, fairly stable global filter statistics are obtained with no sign of filter divergence.

The quality of the ensemble mean background field, as verified using radiosonde observations, is similar to that obtained using a 3D variational procedure. In part, this is likely due to the form chosen for the parameterized model error. Nevertheless, the degree of similarity is surprising given that the background-error statistics used by the two procedures are rather different, with generally larger background errors being used by the variational scheme.

A set of 5-day integrations has been started from the ensemble of initial conditions provided by the EnKF. For the middle and lower troposphere, the growth rates of the perturbations are somewhat smaller than the growth rate of the actual ensemble mean error. For the upper levels, the perturbation patterns decay for about 3 days as a consequence of diffusive model dynamics. These decaying perturbations tend to severely underestimate the actual error that grows rapidly near the model top.

1. Introduction

The ensemble Kalman filter (EnKF) has been proposed as a method for performing 4D data assimilation (Evensen 1994; Houtekamer and Mitchell 1998, hereafter HM98). The EnKF generates an ensemble of initial states that can in principle be used to initiate an ensemble forecast. A fairly complete recent overview of the work done with the EnKF in the oceanographic and atmospheric sciences can be found in Evensen (2003).

Encouraging results, using an ocean general circulation model and real data, have been obtained by Kepenne and Rienecker (2002). Whitaker et al. (2004) have performed a reanalysis of the atmospheric state using a long series of available surface pressure obser-

ventions. The potential of the EnKF as a basis for numerical weather prediction has been discussed by Lorenc (2003). For the atmospheric applications it is not clear that the EnKF, implemented with a modest ensemble size and with a forecast model having imprecisely known error characteristics (Dee 1995, Orrell et al. 2001), will be competitive with existing data assimilation and ensemble generation methodologies. For instance, at the Canadian Meteorological Centre (CMC), a system simulation approach is used for ensemble prediction (Houtekamer et al. 1996). Currently in the data assimilation component of that system, multiple analysis cycles are run with multiple versions of a spectral forecast model and with an optimal interpolation (OI) data assimilation procedure. In the future, we would like to use the EnKF to perform the multiple analysis cycles.

As part of the historical development of the EnKF, we have performed a series of experiments in increasingly realistic environments ranging from the 3-level quasigeostrophic model used in HM98 to the dry 21-

Corresponding author address: P. L. Houtekamer, Direction de la Recherche en Météorologie, 2121 Route Trans-Canadienne, Dorval, Québec, H9P 1J3, Canada.
E-mail: peter.houtekamer@ec.gc.ca

level primitive equation model used in Mitchell et al. (2002, hereafter MHP). It would appear from these studies that modest ensemble sizes (order 100) are sufficient. However, in these studies we only used simulated observations and there has been a concern, expressed most specifically in MHP, about the perhaps dominant role of poorly understood model imperfections in an operational environment with real observations.

In the present study, the EnKF is implemented with a medium-resolution version of a primitive equation model that includes a complete set of physical parameterizations. A fairly complete set of real observations, including radiance observations from satellites, is used. This allows a first impression of the potential quality of the EnKF in a realistic environment and permits an examination of the error dynamics in a data assimilation cycle and in a subsequent forecast.

In the next section, we describe the experimental environment. In section 3, two simulation experiments are performed to investigate filter behavior with and without a simulated model error having known statistical properties. Subsequently in section 4, using real data, we obtain innovation statistics for the EnKF and compare with innovation statistics of a 3D variational data assimilation (3DVAR) procedure as well as with an ensemble-based prediction of innovation amplitudes. In section 5, the growth rates in an ensemble prediction initiated from the analyses produced by the EnKF are investigated. We summarize our results in section 6.

2. The experimental environment

The experimental environment has been largely inherited from our earlier studies (most recently MHP) with a number of modifications to, in particular, the forecast model and the observational network. These were mostly motivated by a desire to increase compatibility with the tools that are supported at our operational center. This approach has the additional advantage that it allows for direct comparisons with the currently operational deterministic 3DVAR (Gauthier et al. 1999a, hereafter GBF; Chouinard et al. 2001; Chouinard et al. 2002; Sarrazin and Brasnett 2002) and its 4D variational successor, which is under development.

a. The model

At the CMC, different configurations of the Global Environmental Multiscale (GEM) gridpoint model (Côté et al. 1998a; Côté et al. 1998b) are used to produce global as well as regional forecasts. For the EnKF experiments, we use a configuration of the GEM model that is very similar to the version used for operational global deterministic forecasting. A lower 240×120 horizontal resolution is used here with a corresponding increase in the horizontal diffusion with respect to the

operational version. Both versions use the same 28 η levels between the surface and the model top at 10 hPa. We also use the same set of physical parameterizations, including the new subgrid-scale orographic blocking term implemented by Zadra et al. (2003). We modified the treatment of snow over ice to reduce the rate of adjustment to the radiative equilibrium temperature. We also found that we could collocate the model's computational poles with the geographical poles and thus avoid an interpolation step between the forecast model and the data assimilation procedure.

b. The observations

In our first implementation of the EnKF, we try to benefit as much as possible from the existing local infrastructure. We also feel that it is preferable to begin with a relatively simple algorithm and to subsequently add enhancements, guided by the outcome of the early experiments. Therefore, the current treatment of observations in the prototype EnKF is very similar to the corresponding operations for the operational 3DVAR that provides the initial conditions for the high-resolution global forecast.

At our center, data assimilation cycles operate with a 6-h period. Each day a sequence of analyses valid at 0000, 0600, 1200, and 1800 UTC is produced. All available observations are grouped into a time window of 6 h centered on the analysis time. In the data assimilation procedure, it is assumed that all observations are valid exactly at the central analysis time. This assumption is reasonable for, in particular, the radiosonde observations. For other observation types, the data-selection procedures give preference to observations taken close to the central time.

The operational 3DVAR outputs an observation file that contains information on whether an observation has been used in the analysis. From this file we extract all accepted observations. This procedure allows us to benefit from the operational quality control and data-selection procedures, many of which are specific to each particular observation type. For example, radiosonde profiles are subject to hydrostatic, lapse rate, and wind shear checks; aircraft reports are sorted by aircraft identifier and then quality controlled one aircraft at a time; while level-1b microwave radiances from the Advanced Microwave Sounding Unit-A (AMSU-A) instruments are subject to a three-step bias-correction procedure. All observations are also subject to a "background check" that verifies that each observation is reasonably close to the available background field. A data-selection procedure is then applied to reduce the density of aircraft, satellite-wind, and microwave-radiance reports. This thinning procedure reduces the horizontal resolution of aircraft and satellite-wind reports to $\sim 1^\circ$ and $\sim 1.5^\circ$ of latitude, respectively, and microwave-radiance reports to ~ 250 km. Further information about the preprocessing and quality control of aircraft, satellite-wind, and level-1b microwave-radiance reports

in the operational 3DVAR can be found in Chouinard et al. (2001), Sarrazin and Brasnett (2002), and Chouinard et al. (2002), respectively.

By assimilating only those observations that have been accepted by the operational system, the EnKF benefits also from the subsequent variational quality control (Andersson and Järvinen 1999) procedure (hereafter QC-Var) that is currently applied to all observations except the microwave radiances. This procedure verifies that each observation is reasonably consistent with all available information, including nearby observations. In the 3DVAR, the QC-Var smoothly reduces the weight given to outlying observations. It never rejects an observation completely. However in the EnKF, observations mostly accepted by the QC-Var are given full weight, and observations mostly rejected by the QC-Var are completely rejected; there is no middle ground between these two extremes.

Currently, the EnKF does not assimilate observations of surface wind and surface humidity. In an experiment including surface-wind observations, we could not demonstrate a positive impact. We have not yet performed experiments with surface humidity. Nor did we include the humidity from satellite (HUMSAT) data (Garand 1993) that were being used operationally during our experimental period but were later replaced at our center by the direct assimilation of 6.7- μm channel data from the Geostationary Operational Environmental Satellite (GOES) platforms. Finally, since the EnKF runs at lower resolution than the operational system, we remove (i) surface observations that are too far from the EnKF topography, and (ii) upper-air observations that are too close to it.

In Table 1, we list the number of observational items that are actually used by the EnKF on the first day, 19 May 2002, of the assimilation experiments. It may be noted that we assimilate of the order of 10^5 observations per analysis time. With respect to the level-1b microwave radiances, we assimilate channels 3–10 over open ocean and from three to five of these channels over land and ice depending on the height of the topography, as described by Chouinard et al. (2002).

TABLE 1. Number of observations used by the EnKF on the first day of the assimilation experiments. National Oceanic and Atmospheric Administration (NOAA) satellites *NOAA-15* and *NOAA-16* were the two polar-orbiting satellites providing AMSU-A level-1b microwave radiances in May–Jun 2002.

	UTC			
	0000	0600	1200	1800
Radiosonde upper air	40 153	5248	39 941	3718
Radiosonde surface	1086	77	1083	52
Surface	9243	9661	9557	9829
Aircraft	10 905	6781	11 033	11 536
Satellite wind	10 990	10 728	10 844	13 474
<i>NOAA-15</i>	15 942	10 225	15 863	15 131
<i>NOAA-16</i>	17 379	15 808	17 060	17 928
Total	105 698	58 528	105 381	71 668

c. Model error

A proper statistical description of the model error is a crucial component of any implementation of the Kalman filter (Dee 1995). It is not our intention in this paper to develop a realistic description of model error. Rather, we aim to develop an algorithmic framework that will allow us to learn more about the error dynamics in operational data assimilation. This acquired knowledge will hopefully provide guidance for more realistic experiments in the future. Therefore, for our first implementation of the EnKF, we decided to stay close to our center's 3DVAR, which is known to work well for operational atmospheric data assimilation.

To account for model error, we use a model-error covariance matrix from which we sample different random fields for different members of the ensemble (Mitchell and Houtekamer 2000, hereafter MH2000). As suggested in that study, we assume that the model-error covariance, \mathbf{Q} , is of the same functional form as the forecast-error covariance matrix, \mathbf{P}_{3D}^f , used by our center's 3DVAR analysis (GBF), but with a smaller amplitude:

$$\mathbf{Q} = 0.25\mathbf{P}_{3D}^f. \quad (1)$$

We have made a number of simplifications with respect to the complete covariance description that is used in the variational algorithm. Of the independent covariance components for streamfunction, divergence, unbalanced temperature, natural logarithm of specific humidity, and of unbalanced surface pressure, we only implemented the components for streamfunction and unbalanced temperature. For these two components, we neglected the wavenumber dependence of the vertical correlations, the seasonal dependence, and the latitudinal dependence of the variances. To generate an approximately balanced perturbation for u , v , T , and p_s from a streamfunction perturbation field, we now use the same linear regression operator (GBF) as the 3DVAR does.

Our approach of adding a covariance matrix from a 3DVAR algorithm to dynamically evolved covariances is similar to what is done in a hybrid scheme (Hamill and Snyder 2000). However, because we add a sample of model-error fields to a sample of background fields, in our algorithm the rank of the covariance matrix is at most the size of the ensemble. In the classical hybrid scheme, the combined covariance matrix is full rank. The advantage of our approach is that we can use a computationally more efficient direct algorithm for the data assimilation procedure, whereas the hybrid method uses a variational algorithm for each member of the ensemble.

d. The method

The design of an EnKF algorithm that makes optimal use of an ensemble of necessarily limited size is an area

of active research (e.g. Tippett et al. 2003; Hamill and Snyder 2000; Lawson and Hansen 2004). Present-day computers could likely support most of the proposed algorithms. The emphasis of the current paper is on obtaining a “reality check” using real observations. For the data assimilation algorithm, we use a configuration that is very similar to our earlier algorithm (MHP). In particular, we use a pair of ensembles where the Kalman gain used for the assimilation of data into one ensemble is computed from the other ensemble. The algorithm, also known as a double EnKF, is shown schematically in Fig. 1. The forecast model and the analysis both use an η terrain-following vertical coordinate. Note that an interpolation of analysis increments from pressure levels to η levels (Lönnerberg and Shaw 1987) was used in MHP.

Another important change to the EnKF algorithm is the replacement of the satellite-derived thicknesses, which we assimilated previously, by AMSU-A microwave radiances. These are assimilated as described in Houtekamer and Mitchell [2001, hereafter HM2001, Eqs. (1)–(3)]. For the interpolation from the background to a given radiance profile, we first interpolate horizontally to produce a profile of model variables at the given latitude and longitude. This profile is then used to simulate the radiances—that is, brightness temperatures for the channels being assimilated—using a fast radiative transfer model. As was done in the 3DVAR at the time, RTTOV-6 (Saunders et al. 1999, Saunders 2000) was used for this purpose.

The state vector for the analysis consists of the two horizontal wind components, the temperature, and the specific humidity at each of the 28 η levels, as well as the surface pressure and surface skin temperature fields. This skin temperature is required for the assimilation of radiance observations. It is produced by the

forecast model at the end of the 6-h integration and is updated by the analysis, but the updated skin temperature field is not subsequently used by the forecast model for the next 6-h integration.

With the addition of seven more vertical levels and a model top that is now at 10 hPa (as compared to 50 hPa in MHP), it proved beneficial to perform covariance localization not only in the horizontal but also in the vertical (Keppenne and Rienecker 2002; Whitaker et al. 2004). For this covariance localization, we use a fifth-order piecewise rational function [Gaspari and Cohn 1999, Eq. (4.10)] with the natural logarithm of pressure as the vertical coordinate. The localization is such that covariances are forced to zero in two units of $\ln p$. Thus, for example, the covariances associated with a 1000-hPa observation fall to zero at 135 hPa, while those associated with a 10-hPa observation fall to zero at 74 hPa.

To motivate the vertical localization of covariance, we use the final set of background fields of the experiment that is described in detail in section 4. We perform an analysis similar to the one that was presented by HM98 (their Fig. 7) to motivate a horizontal localization. In Fig. 2, we show vertical correlations of temperature with respect to the temperature at model level 23 (approximately at 100 hPa). At each grid point we computed the vertical correlations, and subsequently we computed the global average of these correlations. We notice that for this level the correlations are very narrow. In fact, they are so narrow that at levels 22 and 24 the correlations are already negative. The narrowness of these correlations suggests that we may encounter some difficulties with the assimilation of temperature observations above the tropopause. In principle, the low global mean average correlations could mask locally significant features. To investigate this, we quan-

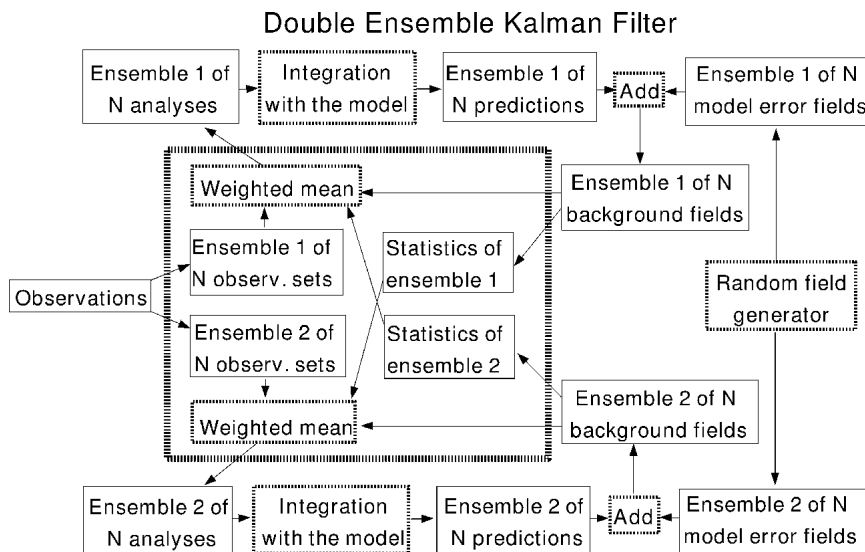


FIG. 1. The procedure used to perform a data assimilation cycle with the double EnKF.

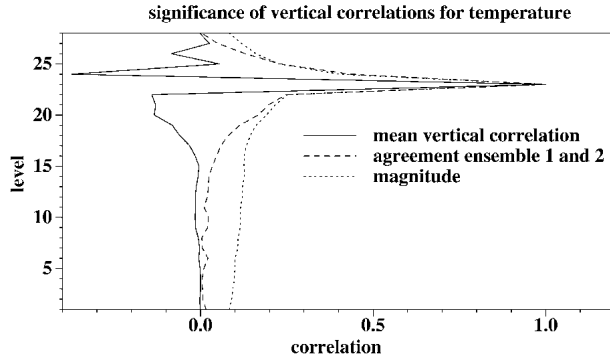


FIG. 2. Vertical correlations in the background fields valid at 1200 UTC 2 Jun 2002 for temperature with respect to the temperature at model level 23. The solid curve shows the average over the globe of the correlations computed over the entire set of twice 64 members. The dashed curve shows the agreement between the two ensembles of the pair (i.e., the square root of corresponding correlations computed from each ensemble of 64 members). The dotted curve is the square root of the mean squared correlations of the entire set.

tify the agreement between the correlation estimates from the two individual 64-member ensembles. At each grid point, the vertical correlations from each of the two ensembles are computed and multiplied together. We then compute the global average of this product and finally take the square root of this average. We obtain a broader curve now, which tends to zero only near level 15 and near the model top. The convergence to zero occurs as the two correlation estimates become uncorrelated values that are distributed about zero. It would appear then that the narrow vertical oscillations in the background temperature field remain well organized over a number of extremes before damping out. The dynamical or algorithmic origin of these upper-level oscillations is not known. We also compute the magnitude of the correlation using all 128 members, by computing the correlation at each grid point, squaring it, taking the global mean, and subsequently taking the root. This value asymptotes to $0.089 \sim 128^{-0.5}$ [HM98, Eq. (12)] as the vertical separation increases and the ensemble loses its information content. Again this value is nearly obtained at level 15 and at the model top. The presence of nonzero spurious correlation estimates suggests that the filtering of ensemble-based correlations of vertically remote variables may be beneficial.

In our algorithm, we apply the horizontal and vertical localizations separately as follows:

$$\mathbf{K}_j = [\rho_V \circ \rho_H \circ (\mathbf{P}_j^f H^T)] [\rho_V \circ \rho_H \circ (H \mathbf{P}_j^f H^T) + \mathbf{R}]^{-1}. \quad (2)$$

Here \mathbf{K}_j is the Kalman gain calculated from ensemble j ; ρ_H and ρ_V are the correlation functions used for horizontal and vertical localization, respectively; \circ denotes the Schur product; and \mathbf{R} is the observation-error covariance matrix. The terms $\mathbf{P}_j^f H^T$ and $H \mathbf{P}_j^f H^T$ can be

calculated from the members of ensemble j as in HM2001 [Eqs. (2) and (3)].

To apply the horizontal and vertical localizations, we require the 3D coordinates (longitude, latitude, and pressure) of each observation and each model variable. The horizontal locations of all observations and variables used here are well defined. The vertical localization is more problematic. The surface pressure is considered to be valid at the surface (where the pressure equals the surface pressure by definition), even though the surface pressure, via the definition of the η coordinate, has an impact on all model variables. The skin temperature is also considered to be valid at the surface. This temperature is important in the assimilation of radiance observations even though these typically correspond to higher model levels. Consequently, the skin temperature will be used properly in the interpolation operator H , but because of the terms ρ_V in the Kalman gain, the radiance innovations may have only a small impact on the skin temperature of the background.

The different microwave-radiance channels are sensitive to different, fairly broad layers of the atmosphere. Each of the eight channels that are assimilated by the EnKF is assigned the (approximate) pressure at which that channel peaks. Thus, for example, channel 3 (the lowest peaking channel) is assigned a pressure of 625 hPa, while channel 10 (the highest peaking channel) is assigned a pressure of 37 hPa. The question of how to localize becomes more complicated for certain other radiance observations that are affected by temperature and humidity values at fairly different vertical locations (e.g., AMSU-B microwave radiances). Such radiances were not being assimilated operationally at the CMC at the time we performed these experiments.

For the experiments described in this paper, the impact of any observation will drop to zero at a horizontal distance of 2800 km and at a vertical distance of two units of $\ln p$. Based on some experiments with different parameters for the localization in which we tried to minimize the radiosonde innovation amplitudes, both values seemed reasonable for use with a pair of 64-member ensembles. At the time of this writing, we do not know if the vertical localization has undesirable side effects. The vertical localization would appear to be beneficial, as evaluated using radiosonde innovation statistics, but its precise formulation is a subject of investigation.

3. Simulation experiments

To validate our EnKF, we perform a simulation experiment in which the statistical descriptions of all sources of error are considered to be known exactly and are made available to the EnKF. Any subsequent discrepancy between the ensemble spread and the ensemble mean error would suggest the presence of an

as-yet-unknown source of error in our implementation. This is similar to previously performed experiments (e.g., MHP, sections 4a and 4b).

To initialize a truth run we take the CMC operational deterministic analysis, valid at 0000 UTC 19 May 2002, and interpolate it to the 240×120 grid used by the model. The truth run, which ends at 1200 UTC 2 June 2002, is performed as a sequence of 6-h integrations.

Using the covariance matrix $\mathbf{P}^f = \mathbf{P}_{3D}^f$, we perturb the truth field valid at 0000 UTC 19 May 2002 to obtain a central background field valid at that same time [HM98, Eq. (7)]. Then, using the same covariance matrix, we obtain a pair of 64-member ensembles centered on that field [HM98, Eq. (8)]. To obtain the first analysis valid at 0000 UTC 19 May 2002, we take the observations that were assimilated operationally and replace them with values interpolated from the truth run and perturbed with a random observational error [HM98, Eq. (6)]. A sequential analysis algorithm (HM2001) is used with batches of 3×200 observations.

Every 6 h, a model-error component with covariance given by (1) is added, as in Fig. 2 of MH2000. For this simulation experiment, the model-error parameters used to increase the spread of the ensemble are the same as those used to describe the difference with the truth run. The parameters of the model-error component used here, in particular the amplitude and the horizontal length scale, have been subject to a certain amount of tuning and represent an estimate of the isotropic model-error component.

Starting at 0000 UTC 19 May 2002, we perform a 6-h data assimilation cycle until 1200 UTC 2 June 2002. We want to see if the ensemble statistics remain representative of the ensemble mean error. In this simulation experiment, the ensemble mean error is computed with respect to the available truth run.

Figure 3 shows some summary statistics for the assimilation cycle. For winds, temperature, and surface pressure, we use an energy norm [MHP, Eq. (9)]. For specific humidity, we select a level ($\eta = 0.631$) with behavior representative of all levels of the lower troposphere. For winds, temperature, and surface pressure, the ensemble spread remains very close to the true ensemble mean error that it simulates. This suggests that the EnKF works well for these variables. However, for specific humidity, the ensemble spread is systematically too small. This may be related to the almost complete absence, from our observational database, of humidity observations in the tropical areas that dominate the humidity content of the atmosphere.

Similar to our earlier results (MHP, section 4b), we observe that the error amplitudes grow mostly because of the model-error component. For winds and temperatures, error amplitudes actually decrease during the 6-h prediction step. There is a subsequent increase due to the addition of a model-error component and, as expected, a decrease due to the assimilation of new observations. For surface pressure, we observe a modest

increase of amplitudes with the model dynamics followed by a significant increase due to model error and a similarly significant decrease due to the data assimilation. For humidity, it is very hard to discern any dynamics from Fig. 3. This is due to the absence of a humidity component in the model-error parameterization and also to the relative lack of observations related to humidity in our observational dataset.

In summary, the filter appears to behave well for winds, temperature, and surface pressure, with no discernible model error being due to deficiencies of the assimilation component of the EnKF. However, the lack of error growth due to model dynamics is in contrast with the classical picture of the analysis cycle being a “breeding ground” for fast growing modes (Toth and Kalnay 1993). This is a serious concern because we would like to use the EnKF to provide the unstable initial conditions for an ensemble prediction system with realistic spread in the medium range.

From these results, it appears that the error dynamics during the assimilation cycle are dominated by the addition, every 6 h, of a model-error term of significant amplitude. It is interesting to investigate the error dynamics in the absence of such a model-error term. We have therefore redone the previous experiment without the regular addition of model error. This simulates how the EnKF would behave if it were used with an atmospheric model that exactly represents the atmospheric dynamics.

The results in Fig. 4 show that error amplitudes, as measured by the ensemble spread, decrease by roughly a factor of 3 for winds, temperature, and surface pressure if a perfect forecast model is used. For humidity, for which we had no model-error term even in Fig. 3, the decrease is less significant. Having no model-error term for any of the model variables, we note that a significant discrepancy between the ensemble spread and the ensemble mean error now develops for all variables. This is similar to what was observed in the perfect-model experiment by MHP (their Fig. 3) and is indicative of the presence of an error or inconsistency in the experimental configuration that we have yet to identify. Possibilities are the presence of inbreeding in a sequential EnKF (HM2001, their Fig. 3) and an unintentional difference between the integration of the truth run and of the individual ensemble members. Finally, we note the extremely modest error growth of both the temperature and humidity due to the model dynamics. This suggests that the lack of perturbation growth, which was observed in Fig. 3, is not simply a consequence of adding parameterized model-error fields that were perhaps fairly unbalanced.

In section 5, we return to the subject of error growth rates when we investigate the dynamics of an ensemble of 5-day forecasts started from initial conditions provided by the EnKF.

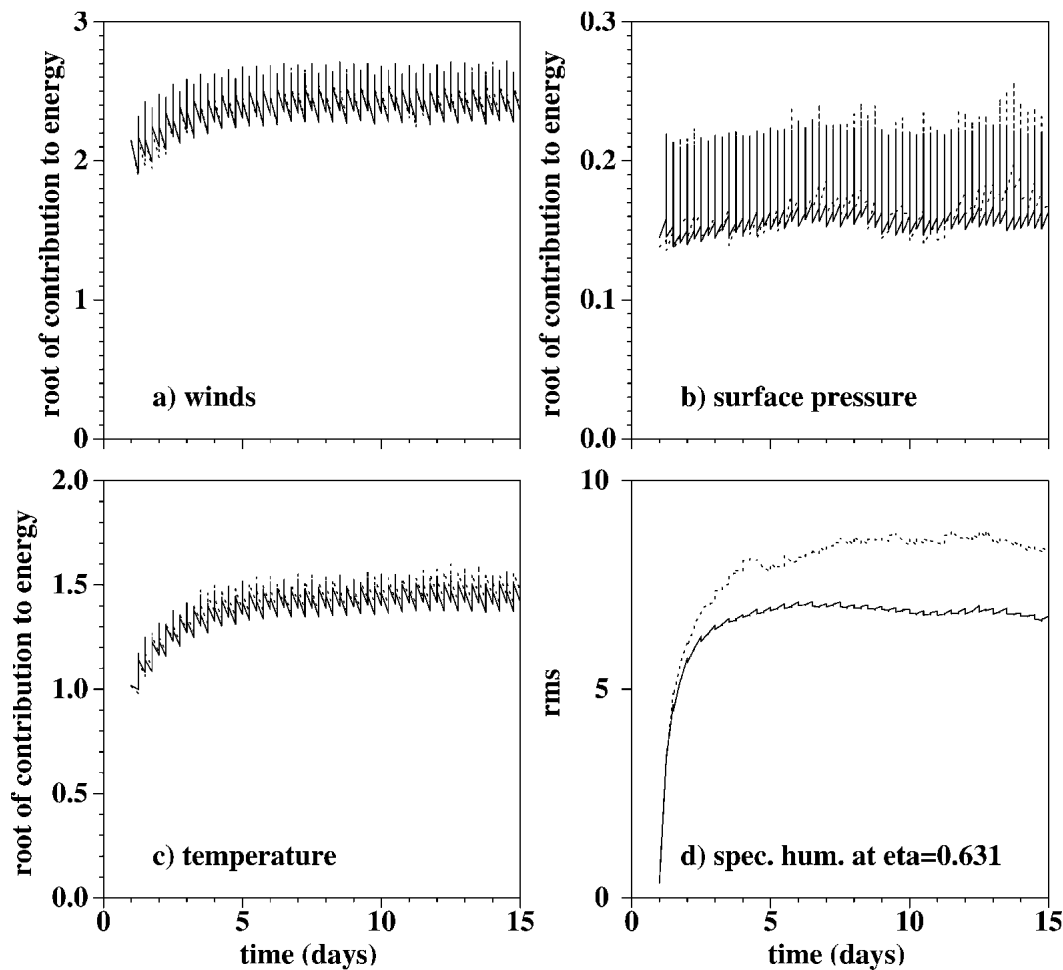


FIG. 3. Analysis, prediction, and forecast error every 6 h for the period between 0000 UTC 19 May 2002 and 0000 UTC 2 Jun 2002 for a simulation experiment with added model error. The solid line is for the rms spread in the ensemble. The dotted line gives the rms error of the ensemble mean with respect to the known truth. All results are for the first ensemble of the pair. The results for (a) winds, (b) surface pressure, and (c) temperature are expressed in terms of a total energy norm. For (d) specific humidity at level 0.631, the global rms value is shown in units of $10^{-4} \text{ kg kg}^{-1}$.

4. An experiment with real data

An experiment with real observations has been performed in order to further evaluate the impact of the different approximations and assumptions that are part of our proposed EnKF configuration. For the proper functioning of the filter, it is important to verify and ensure that the ensemble spread is in general agreement with innovation statistics. We also want to compare the quality of the ensemble mean with the analyses and background fields from a currently operational algorithm.

To assimilate real observations, the EnKF is configured as in the experiment of section 3 that includes model error. Now, however, the truly observed values are assimilated instead of being replaced with values that have first been interpolated from a truth run and then been perturbed with a small random value. The

data assimilation cycle extends from 0000 UTC 19 May 2002 until 1200 UTC 2 June 2002. A set of particularly reliable stations from the global radiosonde network is used for validation purposes. To compare with the radiosonde observations, the ensemble mean and the ensemble spread are interpolated to the 16 standard pressure levels (1000, 925, 850, 700, 500, 400, 300, 250, 200, 150, 100, 70, 50, 30, 20 and 10 hPa). To allow for the spinup of the filter properties (see Fig. 3), the first 5 days of the experiment are discarded so that the validation period runs from 0000 UTC 24 May 2002 until 1200 UTC 2 June 2002. The validations are performed at 0000 and 1200 UTC each day, since it is at these hours that a large number of radiosonde observations are available. The validation statistics are thus averaged over 20 cases, which yield a total of approximately 6500 individual reports.

For the EnKF we performed a number of experi-

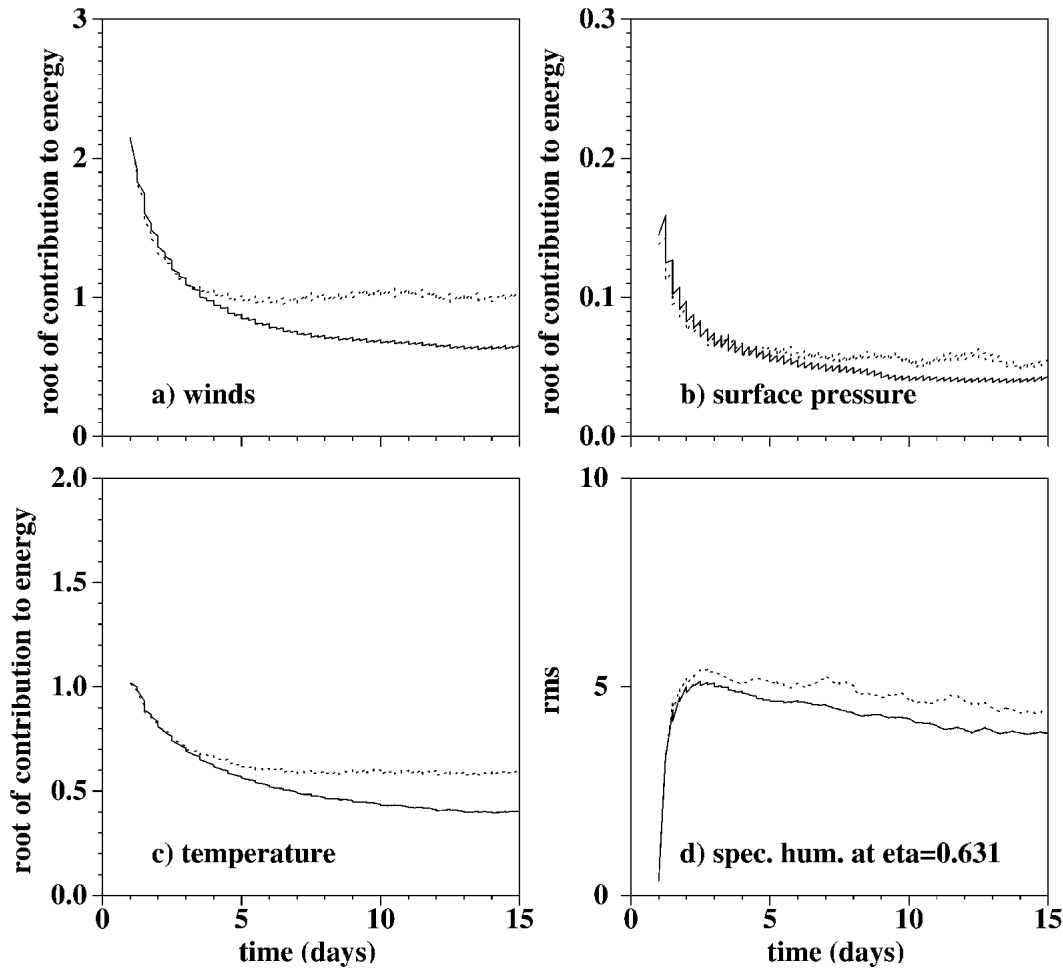


FIG. 4. Analysis and forecast error, every 6 h for the period between 0000 UTC 19 May 2002 and 0000 UTC 2 Jun 2002 for a simulation experiment without model error. (a)–(d) As in Fig. 3.

ments in which we used a separate bias-correction scheme. This algorithm used the zonal-mean analysis increment to update an evolving estimate of the zonal-mean bias. The idea was that the EnKF has been designed to reduce the second moment of the error and one would not expect it to accurately remove the bias that cannot be simulated with the internal dynamics of the model. Equipped with a bias-correction scheme, the EnKF showed significantly less bias, and the second moment was somewhat better as well. However, our subsequent experience was that it was very difficult to isolate the impact of further modifications to the experimental environment. We decided, therefore, to remove the bias-correction scheme and to rely on the success of ongoing research on the forecast model for the future reduction of the bias.

a. Validation of the ensemble spread

A certain amount of prior experimentation has been performed to ensure reasonable behavior of the EnKF.

As mentioned in section 3, the amplitude and the length scale of the model-error component were considered free parameters. The material presented in this subsection is thus to be viewed both as a description of the adjustment procedure and as proof that the filter behaves well.

The basic equation for model-error estimation [Moghaddamjoo and Kiriln 1993; MH2000, Eq. (7)] is

$$\langle v_n v_n^T \rangle = \mathbf{H} \mathbf{P}_n^p \mathbf{H}^T + \mathbf{H} \mathbf{Q}_n \mathbf{H}^T + \mathbf{R}. \quad (3)$$

Here v_n is the innovation vector at time t_n , \mathbf{H} is the forward interpolation from a complete model state to the observations, \mathbf{P}_n^p and \mathbf{Q}_n are the prediction and the model-error covariance matrices, and \mathbf{R} is the observational error covariance. It is possible (Dee 1995; MH2000) to use an adaptive procedure to adjust a small number of model-error parameters using an ensemble-based estimate of $\mathbf{H} \mathbf{P}_n^p \mathbf{H}^T$.

For this study, a nonautomated procedure has been used in which only the diagonal of (3) has been considered:

$$\text{diag}(\langle v_n v_n^T \rangle) = \text{diag}(\mathbf{H}\mathbf{P}_n^f\mathbf{H}^T) + \text{diag}(\mathbf{R}). \quad (4)$$

Here the forecast-error covariance, \mathbf{P}_n^f , is the sum of the prediction error, \mathbf{P}_n^p , and the model-error covariance, \mathbf{Q} . In our experiments, the same model-error covariance, \mathbf{Q} , is used for an entire 14.5-day data assimilation cycle. All terms in (4) can be evaluated from a data assimilation cycle with the EnKF. The innovation is obtained as the difference between observations and the interpolated ensemble mean value [MH2000, Eq. (5)], \mathbf{P}_n^f can be estimated from the ensemble of background fields, and an estimate of \mathbf{R} is available from the data assimilation code. Several data assimilation cycles, with different fixed values for the model-error parameters, were run until a reasonable agreement was obtained between the two sides of (4).

At the end of the 14.5-day cycle, the mean innovation statistics were computed for different variables and levels. It can be seen from Fig. 5 that there is general agreement between the standard deviation of the innovations and the corresponding ensemble estimate. However, there are a few systematic differences. The predicted innovation standard deviation for humidity is generally too small, which is taken to imply that the ensemble spread is too small. This is related to the absence of a humidity component in our model-error description. Some preliminary experiments with parameterized model error for humidity did increase the ensemble spread for humidity but did not reduce the ensemble mean innovation amplitudes for humidity.

For geopotential height, the ensemble spread is generally too small above 300 hPa. Combining this information with the narrow vertical correlations for temperature from Fig. 2 and the satisfying ensemble spread for temperature would suggest that we need broader vertical correlations for the model error at the upper levels. However, it should be noted that the validation of geopotential height is less direct than the validation of temperature. Geopotential height observations have long been used at our center for the deterministic OI analysis (Mitchell et al. 1996). That analysis, which forms the basis of the data assimilation procedure used by the operational ensemble prediction system (Houtekamer et al. 1996), uses geopotential height as an analysis variable. Initially the same was true for the new 3DVAR (Gauthier et al. 1999b). However, the geopotential height variable is no longer used by the 3DVAR (Chouinard et al. 2001), nor is it used by the EnKF. These newer algorithms use temperature and surface pressure as analysis variables for the mass field. A diagnostic program is run to obtain geopotential height fields from the surface pressure and the temperature fields so that they can be compared to geopotential height observations for validation purposes. The narrow vertical structures for temperature in the upper levels (Fig. 2) lead to an inherent uncertainty in the derived values for geopotential. It is not clear how this impacts on the results shown in Fig. 5c.

For winds and temperature, we feel that the agreement is within the uncertainties of the experimental procedure followed here. We note, for instance, that the observational errors used for Fig. 5 are those used in the EnKF for all radiosondes, whereas the validation is restricted to a subset of particularly reliable radiosonde stations. These latter stations, being more reliable, should be assigned a smaller observational error.

Looking at the observational and forecast components in Fig. 5, one notices that they are generally of the same magnitude. This is reassuring because a relatively small ensemble spread could cause the EnKF to give insufficient weight to the observations. That condition could lead to filter divergence. The worst behavior, with a difference in amplitude of a factor of 2, again occurs for humidity and for geopotential height above 300 hPa. It is possible to do the validation separately for small subareas, but such validations are difficult to interpret and they might lead to more complex correlation models for the model error. We do not feel that this is a fruitful area for future EnKF research and we discuss an alternative (perturbing some parameters of the forecast model) in the concluding discussion.

b. Validation of the ensemble mean

Our objective is to implement the EnKF as the data assimilation component of the operational ensemble prediction system at CMC. It would thus replace the currently operational OI system, which runs on a 300×150 horizontal grid. Some comparisons were performed between the innovation statistics for the mean of the ensemble produced by the EnKF and the mean of the ensemble produced by the operational OI system. The statistics were much better for the EnKF, but it was difficult to interpret this result because the operational system uses an ensemble of configurations of a different dynamical model and the OI uses satellite-derived thicknesses instead of directly assimilating the observed radiances. It was impossible to say whether the 4D procedure, the improved treatment of the observations, or the use of the GEM dynamical model was responsible for the observed improvement with respect to the OI system.

We subsequently decided to compare with the 3DVAR that was operational during the spring of 2002. That system used an older version of the GEM dynamical model, a 400×200 horizontal grid, and model computational poles not located at the geographical poles. The observational dataset was almost identical but included observations of surface winds, surface humidity, as well as the HUMSAT observations. The EnKF also used an older description, with larger values, for satellite-wind observational errors, and, finally, observational errors for dewpoint depression observations that did not depend on height. Verification against radiosonde observations of ensemble mean 6-h forecasts from the EnKF and of 6-h forecasts from the 3DVAR gave very similar results, with the two schemes behav-

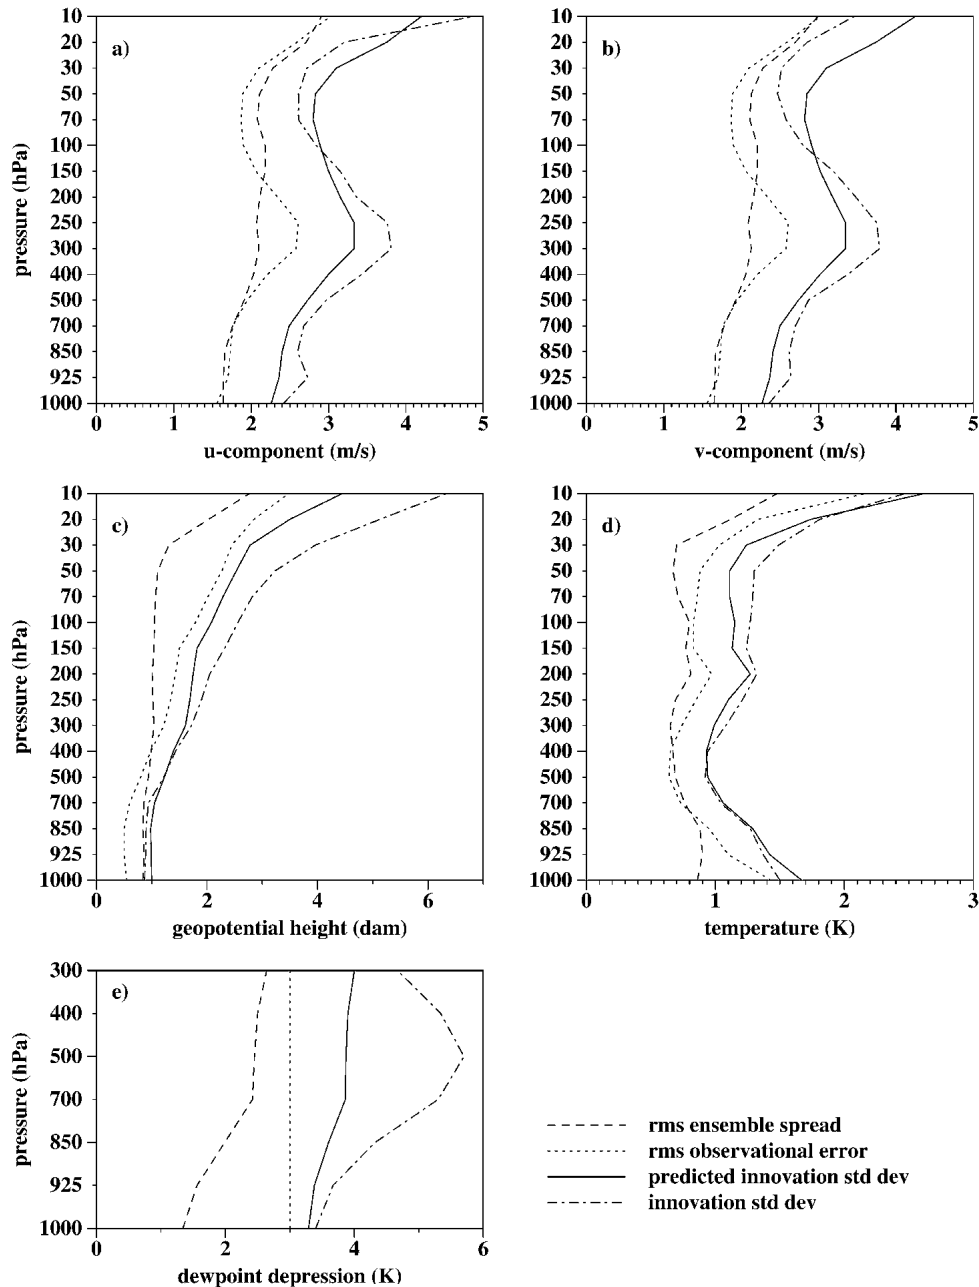


FIG. 5. Comparison of error amplitudes that have been averaged over a 10-day experimental period. Shown are the predicted innovation std dev (solid) that should match the observed innovation std dev (dashed-dotted). The predicted std dev is computed from the rms observational error (dotted) and the rms ensemble spread (dashed).

ing slightly differently in different areas. Interestingly, it was possible to hypothesize that the differences in the model configurations were responsible for a significant part of the observed differences, whereas we had intended to evaluate the impact of changing from a 3D to a 4D procedure.

We therefore decided to rerun the 3DVAR at the same resolution as the EnKF, with exactly the same

dynamical model, and with exactly the same set of observations and observational error statistics. The resulting verifications of 6-h forecasts against radiosonde observations are shown in Fig. 6. Based on these results, it is very difficult to choose one system over the other. The EnKF has smaller standard deviation for humidity but it has a larger bias for that variable. The EnKF also has a larger bias for geopotential height. It would seem

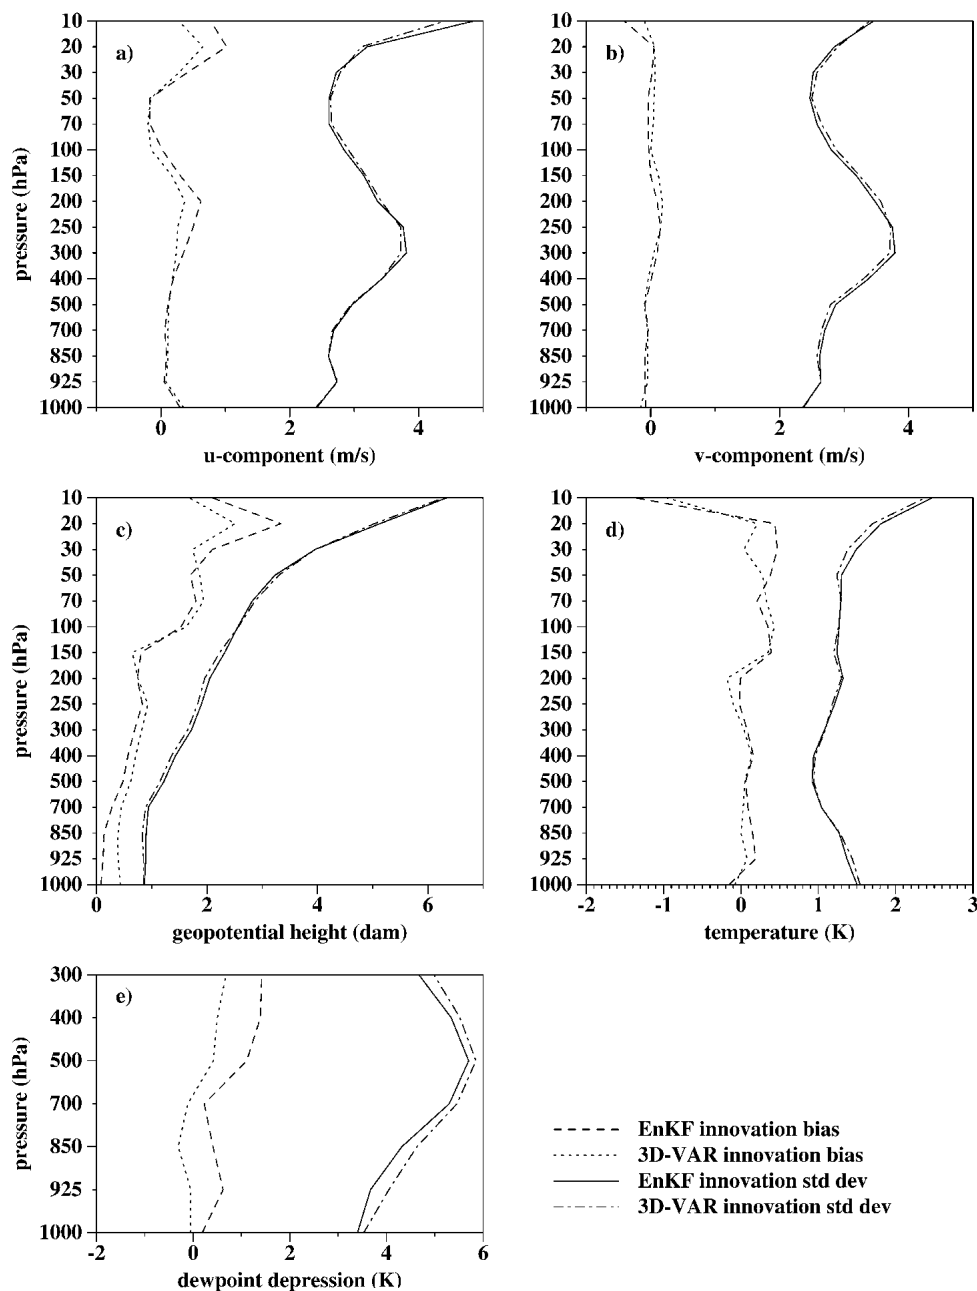


FIG. 6. Comparison of verification scores, obtained over a 10-day experimental period, of the ensemble mean background of the EnKF and the 3DVAR background. The mean value (bias) and std dev of the observed minus interpolated values are shown for the EnKF (dashed and solid, respectively) and the 3DVAR (dotted and dashed-dotted, respectively).

that minor changes in either the 3DVAR or the EnKF could cause either of these systems to improve relative to the other. In any case, a comparison based on just 10 days in one particular season has its limitations.

It is puzzling that these two conceptually different data assimilation systems lead to verifications that are so similar. In fact, verifications over smaller subareas or shorter periods (not shown) generally exhibit the same close agreement. To shed further light on this, we com-

puted the difference between the radiosonde observations and the interpolated analysis values. This measures how closely the analysis draws to the observations it has used. The results are presented in Fig. 7. As each analysis has already used the same observations that it is now being compared with, the values plotted in Fig. 7 are generally smaller than the corresponding innovation amplitudes in Fig. 6. More interestingly, it can be seen that the variational algorithm draws substantially




FIG. 7. Comparison of observed minus interpolated values, averaged over a 10-day experimental period, of the ensemble mean analysis of the EnKF and the 3DVAR analysis. The mean value (bias) and std dev of the observed minus interpolated values are shown for the EnKF (dashed and solid, respectively) and the 3DVAR (dotted and dashed-dotted, respectively).

closer to the wind and humidity observations and also tends to draw closer to the temperature observations. The differences are seen to be largest near the surface and to decrease gradually with decreasing pressure. A different behavior is observed for geopotential height, but, as mentioned previously, geopotential height observations are not assimilated in either the EnKF or the variational algorithm.

The differences in the observed minus analyzed values of the EnKF and the variational algorithm reflect differences in the error statistics of the background field. The EnKF generally has smaller error covariances near the surface. Accounting for the uncertainty in the surface fields (Houtekamer et al. 1996; Keppenne and Rienecker 2002) would likely increase the lower-level ensemble spread in a realistic manner. Because

Fig. 6 suggests that the current implementations of the EnKF and the variational algorithm are roughly of the same quality, and because the two schemes give different weight to observations, it follows that at least one of the two algorithms would benefit from a revision of its background-error statistics. In the variational algorithm one can modify these statistics directly, while in the EnKF changing the model-error description is the most direct way of changing the background-error statistics. The use of a model-error parameterization, based on the 3DVAR forecast-error-covariance matrix, is likely partly responsible for the observed similarity in verification statistics.

5. Growth rates

Some concern was expressed in section 3 about the analysis cycle not being a breeding ground for growing modes in our experiments. This would appear to be in conflict with the following generally accepted belief [but see Orrell et al. (2001) for a different opinion] in the ensemble prediction community: the analysis error contains some patterns that are unstable and that will give rise to rapid error growth. After, say, about 2 days of integration, these growing components will dominate the forecast error. This belief is supported by the additional observation that forecast errors grow quickly in operational forecasts. Therefore, one would expect an ensemble prediction system based on initial conditions that are decaying to be unable to match the observed growth rates without the use of additional measures, such as the addition of a very significant model-error component.

To further investigate the observed (Fig. 3) initial

decay of the perturbations, we performed an ensemble of 5-day integrations starting from initial conditions taken from the data assimilation cycle described in section 4. For simplicity we only used the 64 analyses, valid at 1200 UTC 2 June 2002, that constituted the first of the two ensembles of the pair. We extended the data assimilation cycle for 5 more days so that the ensemble mean analyses, as computed from the first ensemble of the pair, could be used to validate the ensemble of sixty-four 120-h forecasts.

The growth rate of the perturbations, as measured with the energy norm, is displayed in Fig. 8a. The total error energy is seen to decrease for about 24 h, after which time the growth of the growing modes apparently starts to dominate over the decay of the decaying modes. To put these error levels in a context, we display in Fig. 8b the corresponding errors of the ensemble mean validated against the ensemble mean analysis. These latter curves would have zero error at the initial time by construction (not plotted). For short lead times of 6 h, we saw in Fig. 5 that error levels in the ensemble, due in large part to the addition of model error every 6 h, are in approximate agreement with innovation amplitudes. Beyond day 1 the error involved in validating against analyses is less dominant, so we conclude from Fig. 8 that the ensemble spread, which here evolves because of internal dynamics only, grows more slowly than the ensemble mean error. Such behavior was observed in MHP (Fig. 7) and is commonly seen in ensemble prediction systems (e.g., Table 4 of Houtekamer et al. 1996). In an ensemble of medium-range integrations, this difference would be reduced by a simulation of the model-error component.

To permit a better understanding of the dynamical




FIG. 8. The growth rate of (a) perturbations and (b) the ensemble mean error are shown in units of energy for an ensemble of sixty-four 120-h forecasts. The contributions from wind, temperature, and surface pressure are shown by the dashed, dotted, and dashed-dotted curves, respectively.




FIG. 9. The growth rate of (a) perturbations and (b) the ensemble mean error are shown for four different levels in units of energy for an ensemble of sixty-four 120-h forecasts. The combined wind and temperature error-energy contributions at the levels $\eta = 0.0$, $\eta = 0.101$, $\eta = 0.302$, and $\eta = 1.0$ are shown by the solid, dashed, dotted, and dashed-dotted curves, respectively.

behavior of the ensemble of perturbations, Fig. 9a shows the growth of the combined wind and temperature error energy at the model top ($\eta = 0.0$), at two intermediate levels ($\eta = 0.101$ and $\eta = 0.302$), and at the surface ($\eta = 1.0$). At the top of the model, the error energy decays during 84 h, with an overall significant decrease of the error amplitude. Level $\eta = 0.101$ shows a fairly flat growth curve with the minimum being obtained at 60 h. The most rapid perturbation growth is obtained near level $\eta = 0.302$ where the minimum error energy occurs at 6 h. Growth rates drop off toward the surface as shown.

We may again compare these growth rates with the growth of the ensemble mean forecast error as validated against the ensemble mean analysis. Discarding the first 12 h, during which validating against an analysis leads to artificially low error levels, one may note a fairly similar behavior for the levels $\eta = 0.302$ and $\eta = 1.0$. It can be seen, by considering Figs. 9a and 9b, that both the ensemble spread and the actual errors grow smoothly; as expected, the actual error exhibits larger growth rates. For the top levels, the behavior of the spread and the actual error is qualitatively different; the actual errors grow rapidly, whereas the ensemble spread decreases for several days before showing some moderate growth. The observed growth of the actual error likely reflects the limited quality of the forecast model near the model top. The decrease of the ensemble spread likely results from strongly diffusive model dynamics near the model top. The resulting discrepancy between the behavior of the spread and the actual error implies that the simulation of model error near the model top warrants careful attention in EnKF or ensemble prediction applications.

6. Summary and concluding discussion

The EnKF algorithm implemented here has been based on our earlier studies. A configuration consisting of a pair of ensembles (Fig. 1), having a total of 128 members, has been used. Because we aim at an operational application of the EnKF at our center, we have upgraded our environment to be closer to the one used operationally for the deterministic 3DVAR. Thus, we have adopted a forecast model that is very similar to the version used for high-resolution global deterministic forecasting. In particular, the same set of physical parameterizations is employed and the model top has been raised to 10 hPa. The selection and processing of observations is very similar to what is done in the 3DVAR. The EnKF, like most modern data assimilation algorithms, directly assimilates the observed radiances.

In our experiments, we have measured very narrow vertical temperature structures in the upper layers of the model (Fig. 2). This is likely related to our raising of the model top to 10 hPa. However, the precise algorithmic origin of these correlations is not known at the time of this writing. It is clear though that such narrow unresolved structures have a negative impact on the analysis. In this study we use a vertical Schur product to limit the vertical extent over which an observation may have an impact. In our follow-up work, we intend to further investigate the effects of the vertical localization. Perhaps changing either the model or the model-error parameterization would result in smoother vertical structures.

In section 3, we present two simulation experiments in which all observed values are replaced by values in-

terpolated from a truth run (with a subsequent addition of a small random error). In the first of these experiments (Fig. 3), model error is simulated using an isotropic model-error covariance term. We find that error growth in our data assimilation cycle is mainly due to this term. In our experimental setup, this term represents a multitude of terms that have in common that they are not directly simulated in our EnKF implementation. These, so far elusive, terms include imperfections such as (i) errors in the forward interpolation operator; (ii) errors in the specification of the statistics of observations; (iii) errors due to the parameterization of unresolved dynamical and physical processes; and (iv) errors due to imperfectly known surface fields. The current experimental results do not support a view in which rapidly growing baroclinically unstable perturbations emerge, as a simple consequence of model dynamics, in a data assimilation cycle (Toth and Kalnay 1993). This is surprising because baroclinic instability should be well represented in a primitive equation model of modest resolution such as used here. Our results are not in conflict with the singular vector approach (Molteni et al. 1996). The lack of unstable perturbations could result from our way of accounting for model error. It is possible that the use of more realistic error source terms, whose nature remains to be identified, would lead to the highly unstable initial conditions whose presence is postulated by the singular vector method. Our results seem to support the paradigm, proposed by Orrell et al. (2001), in which model error is the main error source for the first few days of a forecast.

For our current experiments, very few humidity observations were available. Consequently, we observe from Fig. 3d that the humidity variable does not systematically benefit from the data assimilation. We expect that this situation will change as our center recently started using AMSU-B radiances for operational data assimilation.

We do not know if a meaningful isotropic parameterization of humidity-related “model error” exists. Consequently, at the present time, we are not adding parameterized model error for the humidity variable. However, we are currently investigating perturbing some parameters of the forecast model in order to simulate forecast-model-related model error. Perturbing parameters related to convection and condensation would likely augment the ensemble spread for humidity.

A second simulation experiment, in which the model is now considered to exactly represent the atmosphere and in which consequently the isotropic model-error term is set to zero, shows (Fig. 4) that in this case the error amplitudes are smaller by roughly a factor of 3 for winds, temperature, and surface pressure. However, this experiment also shows that there is some unknown imperfection or inconsistency in our current implementation of the EnKF. This conclusion is arrived at by

comparing the ensemble spread and the error of the ensemble mean. As in Figs. 3 and 4 of HM98 and Fig. 3 of MHP, this comparison is a useful diagnostic of EnKF performance.

In section 4, we present an experiment in which we assimilated real data. It can be inferred from Fig. 5 that the ensemble spread is in broad agreement with innovation statistics. This is reassuring, because a lack of agreement might eventually lead to filter divergence. We find, however, that the ensemble spread appears to be too small for humidity and for geopotential height above 300 hPa. The lack of spread for humidity is related to the absence of a model-error component for humidity in our experiments. The lack of spread for geopotential height is likely related to the narrow vertical structures already observed in Fig. 2. We note that the EnKF provides a new tool, via the predicted amplitude of the innovations, that can be used to work toward a coherent simulation of all sources of error in a data assimilation cycle. This includes the “model error” component, which has properties that are not well known. It could project, with small amplitude, on highly unstable modes, which would potentially support the breeding and singular vector method for ensemble prediction, or also, with large amplitude, on stable modes.

The quality of the ensemble mean background field is found to be similar to that obtained with a 3DVAR using exactly the same forecast model and the same observational network, as shown by Fig. 6. However, as seen in Fig. 7, the variational algorithm generally draws closer to the observations. This would seem to be due to the generally larger background-error amplitudes used in the variational procedure. Since, as observed, the current background fields have similar quality, this implies that some retuning of the algorithms would likely lead to improved results.

The small difference in quality between the ensemble mean background from the EnKF and the background from the 3DVAR would seem to suggest that the impact of dynamically evolving covariances is fairly small. One may note that a 4D algorithm in principle also allows for an accurate interpolation of model states to the location and time of the observations. The temporal interpolation will likely be implemented in the EnKF as our center also moves to a 4D variational approach and as quality controlled observation sets not centered on synoptic times become available.

We have started investigating whether the initial conditions from the EnKF can be used as the basis for our center’s medium-range ensemble prediction system. From Fig. 8a one observes that the initial conditions provided by the EnKF, when integrated with a unique version of the forecast model, do not immediately show error growth. Even after a few days of integration, the growth rate of the perturbations is below the rate at which the actual ensemble mean error grows (Fig. 8b). From Fig. 9 it is clear that this is, in part, due to an unrealistic lack of perturbation growth near the top of

the model. To couple the EnKF with the existing ensemble prediction system (Pellerin et al. 2003), we might, for an initial implementation, simply interpolate the first 16 initial states provided by the EnKF to the grids used by the 16-member ensemble. The multiple model versions used in the ensemble prediction system sample the model-error component and would help to obtain a larger, more realistic spread in the ensemble. For future implementations, we would of course prefer having a more unified approach in which the short-range ensemble prediction for the EnKF is performed in the same manner as the medium-range ensemble prediction for the ensemble prediction system.

In summary, we note that operationally interesting results can be obtained with an EnKF using an ensemble of moderate size. We, therefore, are continuing to further develop the EnKF so that it may be used as the 4D data assimilation method for the ensemble prediction system at our center.

The results of the comparison with the 3DVAR are perhaps more intriguing than earthshaking. However, because we are not yet simulating the very significant model-error terms, which we are currently only accounting for using an isotropic parameterization, we are not yet in a position to make definite statements about the potential of the EnKF algorithm. It is therefore too early to predict precisely how the fully developed EnKF will compare with fully developed 3D and 4D variational methods.

Acknowledgments. We are grateful to our many colleagues at Direction de la Recherche en Météorologie and the Canadian Meteorological Centre for their help, suggestions, and encouragement. We thank Stéphane Laroche for his careful internal review and the official reviewers for helping us clarify the discussion.

REFERENCES

- Andersson, E., and H. Järvinen, 1999: Variational quality control. *Quart. J. Roy. Meteor. Soc.*, **125**, 697–722.
- Chouinard, C., C. Charette, J. Hallé, P. Gauthier, J. Morneau, and R. Sarrazin, 2001: The Canadian 3D-VAR analysis scheme on model vertical coordinate. Preprints, *14th Conf. on Numerical Weather Prediction*, Fort Lauderdale, FL, Amer. Meteor. Soc., 14–18.
- , J. Hallé, C. Charette, and R. Sarrazin, 2002: Recent improvements in the use of TOVS satellite radiances in the unified 3D-VAR system of the Canadian Meteorological Centre. *Proc. 12th Int. ATOVS Study Conf.*, Lorne, Australia, Bureau of Meteorology Research Centre, 38–44.
- Côté, J., J.-G. Desmarais, S. Gravel, A. Méthot, A. Patoine, M. Roch, and A. Staniforth, 1998a: The operational CMC-MRB Global Environmental Multiscale (GEM) model. Part II: Results. *Mon. Wea. Rev.*, **126**, 1397–1418.
- , S. Gravel, A. Méthot, A. Patoine, M. Roch, and A. Staniforth, 1998b: The operational CMC-MRB Global Environmental Multiscale (GEM) model. Part I: Design considerations and formulation. *Mon. Wea. Rev.*, **126**, 1373–1395.
- Dee, D. P., 1995: On-line estimation of error covariance parameters for atmospheric data assimilation. *Mon. Wea. Rev.*, **123**, 1128–1145.
- 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.
- , 2003: The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.*, **53**, 343–367.
- Garand, L., 1993: A pattern recognition technique for retrieving humidity profiles from Meteosat or GOES imagery. *J. Appl. Meteor.*, **32**, 1592–1607.
- Gaspari, G., and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, **125**, 723–757.
- Gauthier, P., M. Buehner, and L. Fillion, 1999a: Background-error statistics modelling in a 3D variational data assimilation scheme: Estimation and impact on the analyses. *Proc. ECMWF Workshop on Diagnosis of Data Assimilation Systems*, Reading, United Kingdom, ECMWF, 131–145. [Available from ECMWF, Shinfield Park, Reading, Berkshire RG2 9AX, United Kingdom.]
- , C. Charette, L. Fillion, P. Koclas, and S. Laroche, 1999b: Implementation of a 3d variational data assimilation system at the Canadian Meteorological Centre. Part I: The global analysis. *Atmos.–Ocean*, **37**, 103–156.
- Hamill, T. M., and C. Snyder, 2000: A hybrid ensemble Kalman filter–3D-variational analysis scheme. *Mon. Wea. Rev.*, **128**, 2905–2919.
- 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.
- , L. Lefaiivre, J. Derome, H. Ritchie, and H. L. Mitchell, 1996: A system simulation approach to ensemble prediction. *Mon. Wea. Rev.*, **124**, 1225–1242.
- 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.
- Lawson, W. G., and J. A. Hansen, 2004: Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth. *Mon. Wea. Rev.*, **132**, 1966–1981.
- Lönnberg, P., and D. Shaw, Eds., 1987: ECMWF data assimilation scientific documentation. ECMWF, 96 pp. [Available from ECMWF, Shinfield Park, Reading, Berkshire RG2 9AX, United Kingdom.]
- Lorenc, A. C., 2003: The potential of the ensemble Kalman filter for NWP—A comparison with 4D-Var. *Quart. J. Roy. Meteor. Soc.*, **129**, 3183–3203.
- Mitchell, H. L., and P. L. Houtekamer, 2000: An adaptive ensemble Kalman filter. *Mon. Wea. Rev.*, **128**, 416–433.
- , C. Chouinard, C. Charette, R. Hogue, and S. J. Lambert, 1996: Impact of a revised analysis algorithm on an operational data assimilation system. *Mon. Wea. Rev.*, **124**, 1243–1255.
- , P. L. Houtekamer, and G. Pellerin, 2002: Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Mon. Wea. Rev.*, **130**, 2791–2808.
- Moghaddamjoo, A. R., and R. Kiriln, 1993: Robust adaptive Kalman filtering. *Approximate Kalman Filtering*, G. Chen, Ed., World Scientific, 65–85.
- Molteni, F., R. Buizza, T. N. Palmer, and T. Petroliagis, 1996: The ECMWF ensemble prediction system: Methodology and validation. *Quart. J. Roy. Meteor. Soc.*, **122**, 73–119.
- Orrell, D., L. Smith, J. Barkmeijer, and T. N. Palmer, 2001: Model error in weather forecasting. *Nonlinear Processes Geophys.*, **8**, 357–371.
- Pellerin, G., L. Lefaiivre, P. Houtekamer, and C. Girard, 2003: Increasing the horizontal resolution of ensemble forecasts at CMC. *Nonlinear Processes Geophys.*, **10**, 463–468.
- Sarrazin, R., and B. Brasnett, 2002: Modifications in the op-

- erational use of satellite atmospheric motion winds at CMC. *Proc. Sixth Int. Winds Workshop*, Madison, WI, EUMETSAT Publication 35, 147–154.
- Saunders, R. W., 2000: RTTOV-6: Science and validation report. EUMETSAT, 31 pp. [Available from EUMETSAT Satellite Application Facility on NWP, The Met Office, London Road, Bracknell, Berkshire RG12 2SZ, United Kingdom.]
- , M. Matricardi, and P. Brunel, 1999: An improved fast radiative transfer model for assimilation of satellite radiance observations. *Quart. J. Roy. Meteor. Soc.*, **125**, 1407–1425.
- Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, 2003: Ensemble square root filters. *Mon. Wea. Rev.*, **131**, 1485–1490.
- Toth, Z., and E. Kalnay, 1993: Ensemble forecasting at NMC: The generation of perturbations. *Bull. Amer. Meteor. Soc.*, **74**, 2317–2330.
- Whitaker, J. S., G. P. Compo, X. Wei, and T. M. Hamill, 2004: Reanalysis without radiosondes using ensemble data assimilation. *Mon. Wea. Rev.*, **132**, 1190–1200.
- Zadra, A., M. Roch, S. Laroche, and M. Charron, 2003: The sub-grid-scale orographic blocking parameterization of the GEM model. *Atmos.–Ocean*, **41**, 155–170.