Q. J. R. Meteorol. Soc. **00**: 1–19 (2009)
Published online in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/qj.448

Adaptation of a particle filtering method for data assimilation in a 1D numerical model used for fog forecasting

S. Rémy [a], O.Pannekoucke [b], T. Bergot [c] and C.Baehr [d]

[a, b, c] *Météo-France/CNRS CNRM/GAME URA 1357*
[d] *Météo-France/CNRS CNRM/GAME URA 1357*
and
Université de Toulouse Paul Sabatier, Institut de Mathématiques

Abstract: COBEL-ISBA, a boundary layer 1D numerical model, has been developed for the very short term forecast of fog and low clouds. This forecast system assimilates local observations to produce initial profiles of temperature, specific humidity and liquid water content. As fog forecasting is a threshold problem, the model is strongly non linear.
A new assimilation method based on a genetic selection particle filter was tested to produce the initial conditions. The particle filter was adapted for a deterministic forecast and to take into account the time dimension by minimizing the error on a time window. A simplified particle filter was also used to determine the initial conditions in the soil. The filter was tested with two sets of simulated observations. In all cases, the initial conditions produced by this algorithm were of considerably better quality than the ones obtained with a Best Linear Unbiased Estimator (BLUE). The forecast of the control variables and of fog events was also improved. When comparing scores with the ones obtained with an ensemble Kalman filter (EnKF), the particle filter showed better performances for most of the cases. The size of the ensemble impacted the frequency of filter collapse but had a limited influence on the temperature and specific humidity scores.

KEY WORDS data assimilation; 1D model; particle filter; PBL; airports; low visibility conditions,fog

Received ; Revised ; Accepted

1 Introduction

Low visibility conditions often cause problems for many international airports as they may reduce the landing/takeoff traffic, leading to delays or even cancellations of flights. Accurate forecasts of these conditions have become an important issue. Each airport defines a set of visibility and ceiling thresholds below which safety procedures, called Low Visibility Procedures (LVP), are applied. At Paris-Charles De Gaulle airport, the threshold values are set at 600m for visibility and 60m for the ceiling.
Various approaches are employed to forecast low visibility conditions. 1D models are suitable for the nowcasting of radiation fog events for airports located in flat terrain (Bergot and Guédalia (1994a), Bergot and Guédalia (1994b)). They are currently used in real time to forecast fog at local scale in several airports (e.g. Bergot et al. (2005), Clark (2002), Clark (2006), Herzegh et al. (2003)). The 1D boundary layer model COBEL (COde Brouillard à l'Echelle Locale), developped jointly by Météo-France and the Paul Sabatier University was coupled with the land surface scheme ISBA (Interface Sol Biosphère Atmosphère, (Noilhan and Planton (1989), Boone (2000))), as documented in Bergot et al. (2005). This forecasting system has been used to help produce forecast bulletins of LVP conditions at the Paris-Charles de Gaulle airport in France since 2005. These bulletins aim to provide estimated times for the onset and lifting of LVP conditions up to 4 hours in advance.
Fog is a phenomenon that evolves at small spatial and time scales. Modeling the life cycle of fog involves interactions between many parameterizations : turbulence, microphysics, radiative scheme, surface-atmosphere exchanges. This stresses the importance of working with accurate initial conditions : the quality of the COBEL-ISBA forecasts depends much on the initial conditions (Bergot and Guédalia (1994a), Roquelaure and Bergot (2007), Rémy and Bergot (2009a)). As fog modelling involves numerous threshold processes, the model is strongly non-linear. Because they do not require any linear or Gaussian hypothesis, particle filters are an adequate algorithm to produce initial conditions for such a non-linear system. Particle filters (Doucet et al. (2001), Del Moral (2004) and van Leeuwen (2009) among others) are a probabilistic method that aims to estimate the probability density function (pdf) of the first guess given observations through an ensemble of random draws, or particles. The filter consists of two steps: the particles are integrated by the model, and then updated or selected. There exist many kind of particle filters, based on how the updating and/or selection of particles is done at each assimilation step (van Leeuwen (2009),

*Correspondence to: CNRM/GAME, 42 Av Coriolis, F-31057 Toulouse Cedex, France E-mail : samuel.remy@meteo.fr

Baehr and Pannekoucke (2009)). Assimilation schemes that mix both the particle filter and the ensemble Kalman filter (EnKF) has also been developed, with the aim to guide the particles closer to observations.
Our aim is to check if an algorithm based on particle filtering can provide initial conditions for COBEL-ISBA at a reasonable numerical cost. In doing that, we are confronted to the so-called "dimensionality problem" (Snyder et al. (2008)), i.e. that the number of particles needed to adequately represent the prior density could be very large. This problem is highly dependent on the system that is considered and the type of filter that is used. Particle filters with genetic selection for example were shown (Baehr and Pannekoucke (2009)) to be less affected by this problem. A genetic selection particle filter was thus adapted to provide initial conditions for a deterministic run. The computation of the weights was also modified to take into account observations that are available before or shortly after the analysis time.
The framework of this study is outlined in section 2. Two sets of simulated observations were created : one with mostly clear-sky conditions at the initialization, to study the formation of fog, and the other with frequent occurrence of fog and low clouds. Section 3 presents the setup of the particle filter and section 4 shows the results with the two sets of simulated observations, as compared to the operational setup of the assimilation scheme. In section 5, we are going to discuss the impact of the ensemble size on the performance of the particle filter. Finally, section 6 summarizes the results.

2 Framework of the study

2.1 The COBEL-ISBA assimilation-prediction system

2.1.1 The model

The COBEL-ISBA system results from the coupling of the high resolution atmospheric boundary layer 1D model COBEL (Bergot (1993), Bergot and Guédalia (1994a) and Bergot and Guédalia (1994b)) with the 7-layer land-surface scheme ISBA (Noilhan and Planton (1989), Boone (2000)). The atmospheric model possesses a high vertical resolution: 30 levels between 0.5 and 1360 m, with 20 levels below 200 m, to be able to adequately forecast radiative fog events. The physical parameterizations used in COBEL-ISBA consist of:

- a turbulent mixing scheme with a 1.5-order turbulence closure that uses a prognostic turbulent kinetic energy (TKE) equation. The mixing length differs for stable (Estournel (1988)) and for neutral or unstable conditions (Bougeault and Lacarrere (1989)),
- a warm microphysical scheme adapted to fog and low clouds in temperate regions,
- detailed long-wave and short-wave radiation transfer schemes.

COBEL-ISBA is run at one-hourly intervals and provides up to eight hours of LVP forecasts. The inputs to the model are the initial conditions and mesoscale forcings. Mesoscale forcings (i.e. geostrophic wind, horizontal advection and cloud cover above the model column) are given by the Numerical Weather Prediction (NWP) model ALADIN (http://www.cnrm.meteo.fr/aladin).

2.1.2 The operational assimilation scheme

A two-step assimilation scheme using local observations (Bergot et al. (2005)) provides the initial conditions. The observation system used at Paris-Charles de Gaulle airport is designed to provide up-to-date information on the state of the surface boundary layer temperature and moisture. It consists of a weather station which provides 2 m temperature and humidity, visibility and ceiling; a measurement mast that gives temperature and humidity observations at 1, 5, 10 and 30 meters; radiative flux (short-wave and long-wave) observations at 2 and 45 meters; and soil temperature and water content between the surface and -40cm. Observations from the weather stations are available every 6 minutes whereas for other instruments they are available every 15 minutes.
The operational assimilation system uses information from a first guess or background (i.e. a previous 1 hour COBEL-ISBA forecast), local observations, and profiles from the ALADIN NWP model to compute a Best Linear Unbiased Estimator (BLUE) of temperature and specific humidity initial conditions:

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{K}(\mathbf{y}^o - \mathbf{H}\mathbf{x}^b) \quad (1)$$

where

$$\mathbf{K} = \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1} \quad (2)$$

In Eq. 1, $\mathbf{x}^a$ is the analysis, $\mathbf{x}^b$ is the first guess or background, and $\mathbf{y}^o$ are the observations. $\mathbf{K}$ is the Kalman gain that determines how the background is modified to take into account the observations. $\mathbf{B}$ and $\mathbf{R}$ are the error variance/covariance matrices of the background and of the observations respectively, and $\mathbf{H}$ is the forward operator, i.e. the matrix that interpolates information from the model grid to the observation grid. As the dimension of the system is small (30 levels for two control variables), matrices can be explicitly inverted and there is no need for a variational algorithm. In the operational setup, the error statistics are imposed arbitrarily to allow the initial profile to be close to observations near the surface and closer to the ALADIN profiles above.
When a layer of cloud is detected, an additional step uses an algorithm that minimizes the difference between observed and simulated radiative fluxes at the ground and at 45 m to estimate cloud thickness. This algorithm works as follows: the radiation scheme of COBEL is used to compute the modeled radiative fluxes at 2 and 45m, using different initial thicknesses of the fog layer. The best estimate of the initial fog thickness is the one that minimizes the error between modelled and observed radiative fluxes (see Bergot et al. (2005) for more details). The relative humidity profile is then modified within the saturated layer.

DOI: 10.1002/qj

The soil temperature and water content profiles used to initialize ISBA are obtained directly by interpolation of soil measurements.

2.2 Simulated observations

Observing System Simulation Experiments (OSSE) are an adequate tool for studying the accuracy of an assimilation scheme (e.g. Huang et al. (2007)). They consist in generating pseudo-observations by adding perturbations to a reference model run. The pseudo-observations are then assimilated and the initial state and forecast can be compared to the reference run. The advantages of this method are :

- The perfect model hypothesis is true, in agreement with the hypothesis made in the BLUE assimilation algorithm. The errors in the initial conditions originate only in the observations and first guess errors, themselves originating from errors in initial conditions propagated by the previous forecast. The lack of observations for certain parameters (e.g. the thickness or water content of a cloud layer) does not allow the assimilation scheme to entirely correct the errors of the first guess field. The quality of the initial conditions thus depends solely on the observations used and on the assimilation scheme.
- This framework allows observations to be simulated over the whole domain (the boundary layer for this study).
- Lastly, it is possible to create a large variety of observation sets that accommodate our needs for evaluation purposes.

Two sets of simulated observations were made: one for the study of clear-sky nights and of shallow-fog situations (NEAR-FOG), and the other for the study of frequent and deep fogs (FOG) (See Rémy and Bergot (2009a) for more details on how the simulated observations were generated).

2.2.1 The NEAR-FOG situation

Simulated observations corresponding to clear-sky and shallow-fog situations were produced. This observation set will be referred to as NEAR-FOG hereafter. Fifteen days of simulated observations were generated, during which no fog occurred for the first 10 nights. Shallow fog situations developed for the remaining five nights. Their thicknesses did not exceed 10 m. Twenty-one hours of Low Visibility Procedure (LVP) conditions were "observed" for this situation. The skies above the model column were entirely clear, which ensured strong night-time cooling. Figure 1 shows the "true" temperature at 1m and the corresponding liquid water path. Close to ground level, the daily highs lay in the 20-22 ˚C range while the lows were around 8-9 ˚C. Day and night relative humidity varied greatly from 30% to 100%, corresponding to typical conditions observed during autumn and winter over land.

2.2.2 The FOG situation

This situation was designed to study the fog and low cloud life cycle. Fog and low clouds occurred during many nights of the 15-day observation set, hereafter referred to as FOG, because of high moisture combined with strong night-time cooling due to clear skies above the model column. Figure 2 shows the "true" temperature observations at 1m and the "true" liquid water content integrated over the model column. In total, 98 hours of LVP conditions were "observed" in these 15 days, with fog occurring on 11 nights. Stratus also occurred in the upper part of the model column on days 7 and 8. It was not counted as LVP. Various fog situations occurred, from shallow, early-morning fog to fog layers more than 200 m thick.

2.2.3 Reference experiments for NEAR-FOG and FOG

Figure 3 shows the mean Root Mean Square Error (RMSE) and the mean bias of the forecasted temperature and specific humidity versus forecast time and altitude, when the operational setup, as defined in section 2.12, was used. The influence of the observations can be seen in the lower values of RMSE at initialization time below 50m, especially for temperature. For both temperature (figure 3c) and specific humidity (figure 3a), most of the increase of the RMSE occurred during the first two hours of forecast time. For specific humidity, the maximum of RMSE was always at the surface whereas, for temperature, the RMSE no longer showed large differences between the lower and upper part of the domain after 4h of forecast time. The analysis was nearly unbiased for both specific humidity and temperature (figures 3b and d). The specific humidity bias became slightly positive with forecast time, with a maximum close to the ground. A small cold bias also occured for the forecasted temperature (figure 3d) and increased regularly with the forecast time, with maxima close to the ground-level and above the top of the mast (30m).
Figure 4 shows the mean RMSE and bias of temperature and specific humidity when the operational setup was used with the FOG situation. It is interesting to compare it with figure 3. The initial profiles of specific humidity (figure 4a) show a larger RMSE for FOG than for NEAR-FOG over the whole domain. This is mainly due to errors in the initialization of fog and low clouds. The increase of RMSE with forecast time is slower for FOG than for NEAR-FOG and, after two hours of forecast, the values close to the surface are similar for both situations. The RMSE above 100m remain significantly higher for FOG than for NEAR-FOG, for all forecast times. The specific humidity bias (figure 4b) is close to zero for all forecast times below 50m whereas it is negative above that height. For all heights, the specific humidity bias did not vary much with forecast time. The RMSE of forecasted temperature (figure 4c) increases much faster in the lower part of the domain for FOG than for NEAR-FOG (figure 3c) and reaches a maximum of 1K after 7 hours of simulation. A maximum appears between 50 and 150 m of

DOI: 10.1002/qj

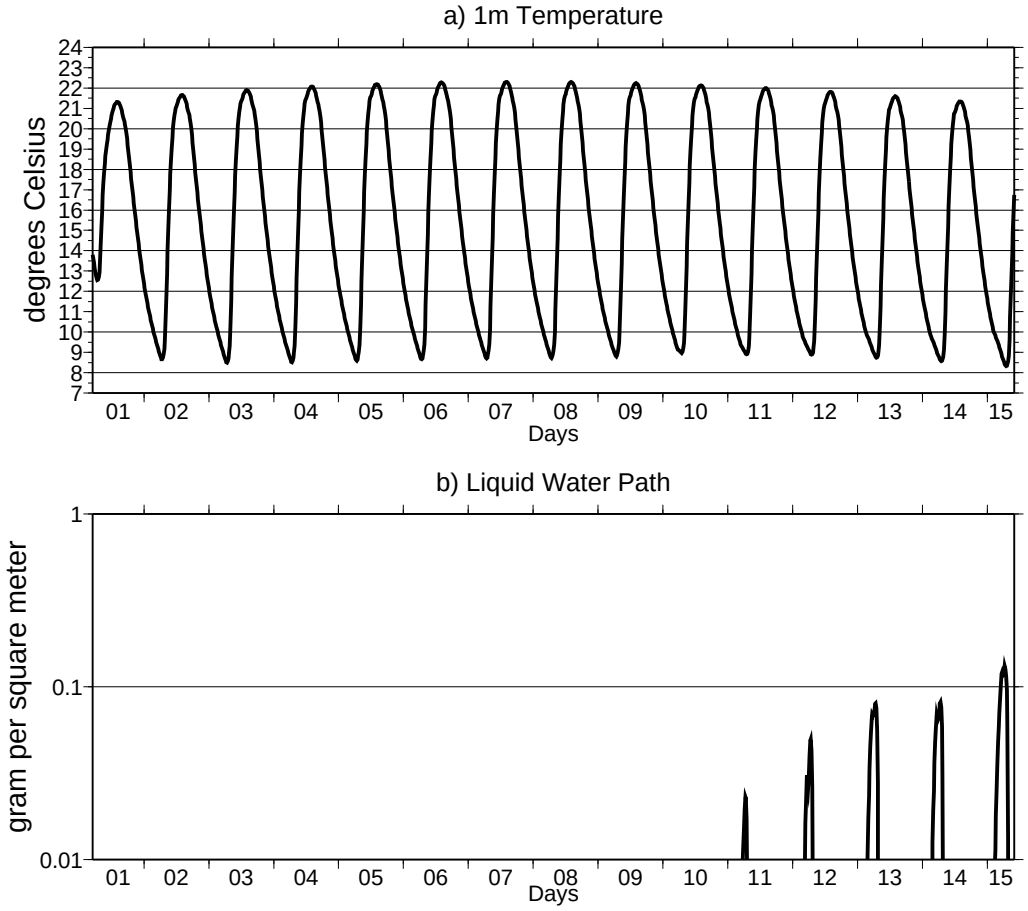

Figure 1. NEAR-FOG : "Truth" for 1m temperature (a) and Liquid water path (b).

altitude, which corresponds to situations where the forecasted height of the fog is different from the simulated observations. The inversion at the top of the fog layer significantly increases the error if the forecasted cloud layer thickness is not the same as the observed one. The temperature bias (figure 4d) also increases with forecast time, with a maximum at the surface.

3 Particle filter-based data assimilation

Particle filters are ensemble-based assimilation algorithm that employ a fully non-linear and non-Gaussian analysis step to estimate the probability distribution function of the model conditioned by the observations. There exist several particle filter algorithms. In this work, a genetic selection particle filter based on the work of Baehr and Pannekoucke (2009) was adapted to a deterministic 1D model. This section presents a background on particle filter, focuses on the genetic selection algorithm, and then show how the particle filter was adapted to supply initial conditions to a deterministic model.

3.1 Fundamentals of particle filtering

Let $(x_k), k \in \mathbb{N}$ be a Markov chain that denotes the model state. $(y_k), k \in \mathbb{N}$ is the sequence of observations. Both are realizations of the random variables $\mathbf{X}_k$ and $\mathbf{Y}_k$ with the probabilities $p(x_k)$ and $p(y_k)$. The aim of filtering algorithms is to estimate the probability $p(x_k|y_k)$. In this work, the hypothesis that a linear relation, denoted by the $\mathbf{H}$ matrix, exists between the observation and the model spaces is made. Non-linear observation operators are possible, but non-necessary in this work. A non-linear dynamical system can be written as:

$$\begin{cases} x_{k+1} = & \mathbf{f}(x_k) + V_k \\ y_k = & \mathbf{H}x_k + W_k \end{cases} \tag{3}$$

$\mathbf{f}$ is the model, $(V_k), k \in \mathbb{N}$ and $(W_k), k \in \mathbb{N}$ are the model and the observation errors respectively; the observation errors are supposed to be independent from each other in time. Particle filters use an ensemble of first guesses $(x_{i,k}), i = 1, .., N$, also called "particles". The subscript k denotes the analysis time iterations, and i the particles. Particle filtering relies on the hypothesis that this ensemble of first guesses is able to approximate the probability $p(x_k)$ through a discrete sum:

$$p(x_k) \quad \overset{N\to\infty}{\sim} \quad \frac{1}{N}\sum_{i=1}^{N} \delta(x_{i,k}) \tag{4}$$

DOI: 10.1002/qj

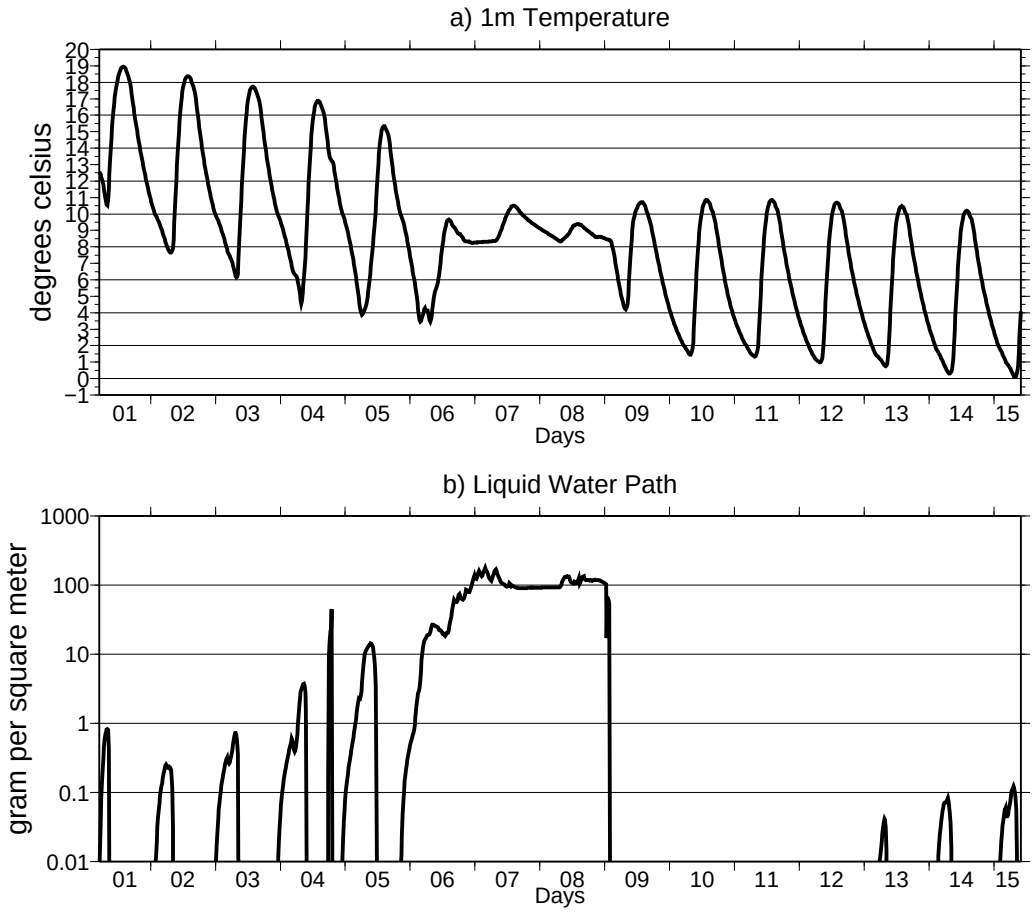

Figure 2. same as figure 1 for FOG.

Then, using the Bayes theorem:

$$p(x_k|y_k) = \frac{p(y_k|x_k)p(x_k)}{\int p(y_k|x_k)p(x_k)dx} \quad (5)$$

$$\overset{N\to\infty}{\sim} \sum_{i=1}^{N} w_i \delta(x_{i,k}) \quad (6)$$

Where $(w_{i,k}), i = 1, .., N$ are the weighting coefficients, given by:

$$w_{i,k} = \frac{p(y_k|x_{i,k})}{\sum_{j=1}^{N} p(y|x_{j,k})} \quad (7)$$

The maximal weight is $w_k^{max} = max_i(w_{i,k})$. The potential fonction is defined for each particle i as follows:

$$G_k(x_i) = p(y_k|x_{i,k}) \quad (8)$$

so that

$$w_{i,k} = \frac{G_k(x_i)}{\sum_{j=1}^{N} G_k(x_j)} \quad (9)$$

3.2 Genetic selection algorithm

Particle filter algorithms differ on whether and how the particles are selected and resampled. The genetic selection algorithm selects the particles that are closer to the observations, i.e. the ones that have larger weights. Resampling is done using only the selected particles. During the selection step, a particle i will be kept with a probability of $G_k(x_i)$ or eliminated with a probability of $1 - G_k(x_i)$. Del Moral (2004) showed that using a multiplicative coefficient ϵ_k, so that a particle i has a $\epsilon_k G_k(x_i)$ probability to be selected and a $1 - \epsilon_k G_k(x_i)$ probability to be eliminated, lowers the error variance of the estimator provided by the particles filter. As in Baehr and Pannekoucke (2009), we chose $\epsilon_k = \frac{1}{max_i G_k(x_i)}$. Once the particles are selected, they are resampled through an importance resampling (IR) algorithm (**?** which uses multinomial draws. This algorithm replicates particles with higher weights. To differentiate them, noise is added to each particle. This noise has to be large enough to differentiate the similar particles that result from the selection step and to range the first guesses probability, but not too large so that the resulting particles have weights that are not too small, i.e. to avoid filter collapse. We chose to add to each particle analysis i a term in the form:

$$\mathbf{x}^a_{i,k} = \mathbf{x}^a_{i,k} + \mathbf{B}^{\frac{1}{2}} \mu_{\mathbf{i,k}} \quad (10)$$

Where $\mu_{\mathbf{i,k}}$ is a vector of random variables drawn from a gaussian distribution $\mathcal{N}(0, 1)$. The noises were thus coherent with the model uncertainty. Bounds were arbitrarily imposed on them so that they were not too large.

DOI: 10.1002/qj

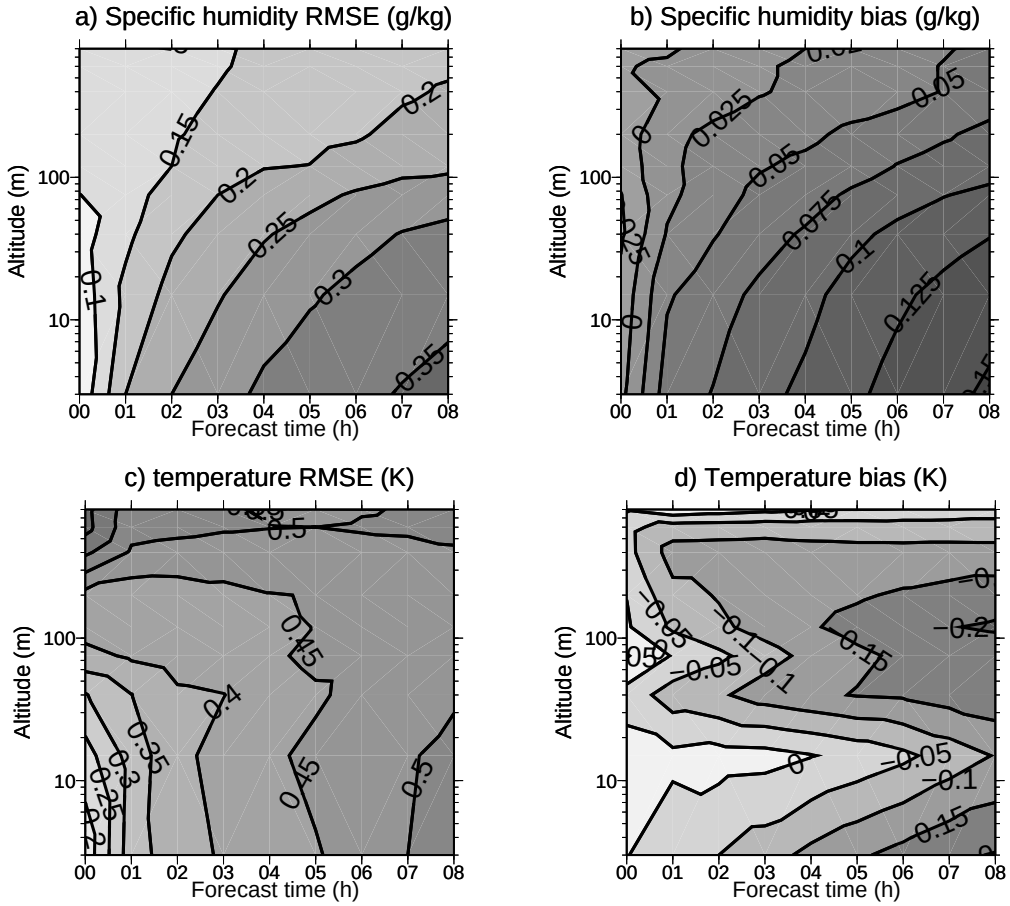

Figure 3. NEAR-FOG : RMSE (left) and bias (right) of temperature (top) and specific humidity (bottom) versus forecast time. Isolines are every 0.05K for temperature bias and RMSE, every 0.05 g/kg for specific humidity RMSE and every 0.025 g/kg for specific humidity bias.

3.3 Dimensionality problem

As particle filters rely entirely on the hypothesis that the background probability $p(x_k)$ can be estimated by a weighted sum of particles, the ensemble has to be large enough to represent accurately enough the probability density function of the first guess. Snyder et al. (2008) showed with the dynamical system proposed by Lorenz (Lorenz (1996)) that the ensemble size needed for a successful implementation of a Sequential Importance Resampling (SIR) particle filter scales exponentially with the problem size. For a 200-dimensional model space, they found at 10^{11} particles were needed to avoid filter collapse or divergence, i.e. a single particle has a weight nearly equal to 1 while all others has very small weights.
The dimensionality problem can be partially reduced depending on what kind of particle filter is used. Baehr and Pannekoucke (2009) showed that a genetic selection algorithm brought convergence of the particle filter with 1000 particles and a 200-dimensional model space, using the same dynamical system as Snyder et al. Furthermore, the frequency of filter divergence also depends on the dynamical system.

3.4 Adaptation of a particle filtering algorithm to a deterministic 1D model

A particle filter with genetic selection was adapted for usage within a deterministic 1D model. The dimensionality problem was partially corrected through the resampling state.

3.4.1 Computation of the weights

As shown by Eq. 7, the weights $w_{i,k}, i = 1, ...N$ are a function of the distance between the particle i and the observations, which is supposed to be known. This function depends much on the law followed by the observation errors, as shown by Del Moral (2004). The hypothesis was made that these errors are Gaussian; as a consequence, the weights are also a Gaussian function of distances. Another advantage of this choice is that the Gaussian function is very discriminative: particles with higher distances will have very small weights. The distance between observations and the particle was taken as the Mahalanobis distance, modified to take into account the background error statistics:

$$p(y_k|x_{i,k}) = p((y_k - \mathbf{H}x_{i,k})^T(\mathbf{R} + \mathbf{HBH}^T)^{-1} (y_k - \mathbf{H}x_{i,k})) \quad (11)$$

DOI: 10.1002/qj

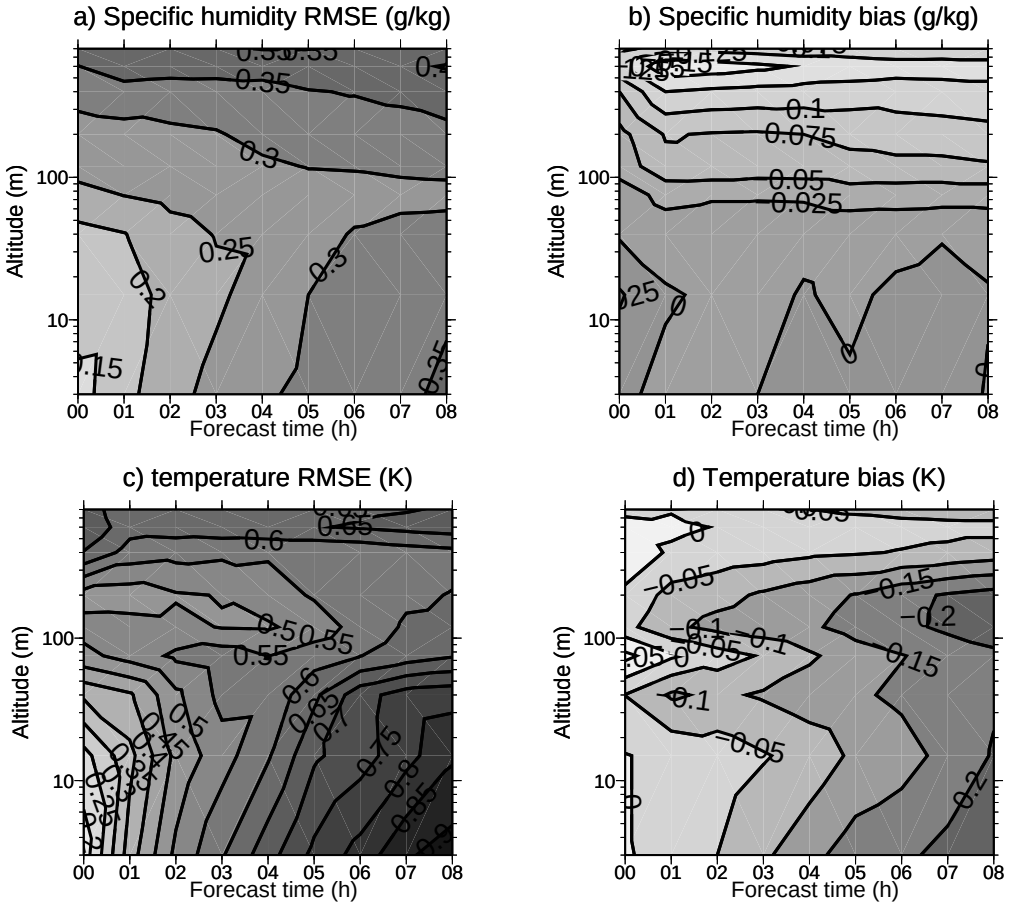

Figure 4. Same as figure 3 for FOG.

The **B** matrix in Eq. 11 was computed directly from the ensemble of first guesses. An issue is the relative importance of temperature and specific humidity in the computation of the modified Mahalanobis distance. As temperature was generally larger than specific humidity in the situations under study, the distance between simulated and observed temperature was often much larger than distance between simulated and observed specific humidity. As a consequence, the overall distance given by Eq. 11 was much more influenced by errors on temperature than on specific humidity. The weights thus depended much more on temperature errors than on specific humidity errors. To solve this problem, the distance between simulated and observed temperature was normalized so that the sum of all temperature distances computed at a given analysis time were made equal to the sum of all distances on specific humidity. As we had no information on the "real" relative impact of temperature and specific humidity on the distance, we chose arbitrarily to equalize their relative influence. The overall distance was then taken as the sum of the specific humidity distance and the normalized temperature distance. It then depends on temperature and specific humidity errors in the same proportion.

In addition, the potential functions were computed at different forecast times of the backgrounds, matching the times when observations were available, so that the observations were assimilated within a time window and not at a single point in time. If a family $(y_k^m), m = 1, ..M$ of observations are available at times m between analysis times $k-1$ and $k+1$, then for each observation y_k^m and each particle $x_{i,m}$ the potential function is computed similarily as with Eq.8:

$$G_k(x_{i,m}) \quad = \quad p(y_k^m|x_{i,m}) \tag{12}$$

The potential function of the particle i over the time window associated with analysis time k is then the product of all potential functions computed at a single time, following Del Moral (2004):

$$\mathcal{G}_k(x_i) \quad = \quad \Pi_{m=1}^{m=M} G_k(x_{i,m}) \tag{13}$$

and

$$\mathcal{W}_{i,k} \quad = \quad \frac{\mathcal{G}_k(x_i)}{\sum_{j=1}^{N} \mathcal{G}_k(x_j)} \tag{14}$$

The maximal weight becomes $\mathcal{W}_k^{max} = max_i(\mathcal{W}_{i,k})$. Figure 5 illustrates this concept. If the weights were computed only as a function of the distance between the particles and observations for single point in time (i.e. analysis time), particle 1 in figure 5 would have had a larger weight than particle 2, as the distance between observations and particle 1 is smaller at analysis time.

DOI: 10.1002/qj

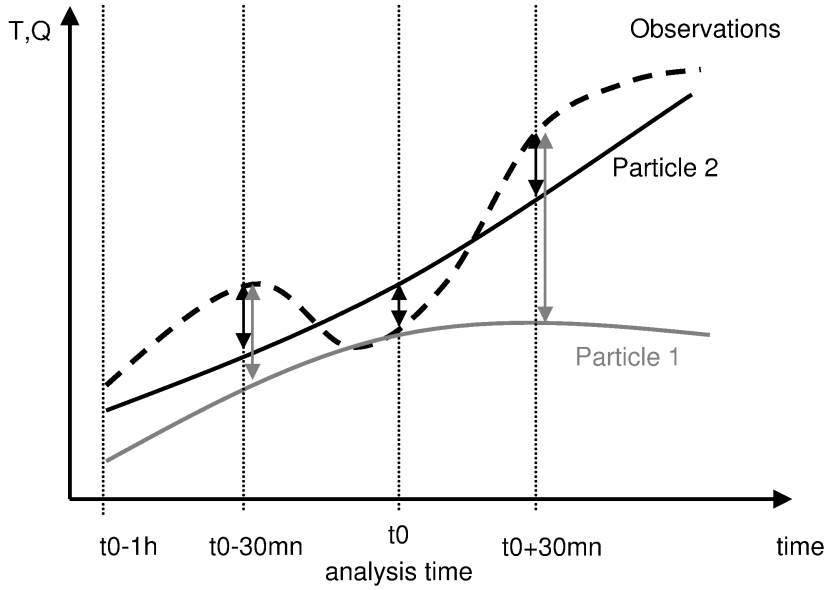

Figure 5. Schematic graph showing two particles and their weights computed with a time window of 30 minutes centered on analysis time.

Particle 1 is however not a good choice, as its trajectory is very different from the sequence of observations, except at analysis time. When the weights computed at forecast times 30 minutes later and 30 minutes earlier than analysis time are taken into account, then particle 2 has a larger overall weight than particle 1, as its trajectory is closer to the sequence of observations.

In our case study, observations from the measurement mast were available every 15 minutes, the weather station provided 2m humidity and temperature every 6 minutes, and ALADIN profiles were available for every hour. The distance between obervations and the background were computed for forecast times varying from 6 minutes to 1h30, i.e. from analysis time minus 54 minutes to analysis time plus 30 minutes. In this way, all available observations were used. This setup thus imposed simulations to be started at least 30 minutes later than analysis time. It is already the case in the operational setup, as the observations that covers the period from analysis time included to 30 minutes later are available around 40 minutes after analysis time.

3.4.2 Determination of the initial conditions

Two possibilities exist for the construction of the initial conditions for the deterministic, non-perturbed run: either take the weighted mean of all particle as the analysis, or the best particle, i.e. the one with the largest weight, as the analysis. Here, we chose the latter option, so that the initial conditions are as close as possible to the observations and that its coherence with the model physics is ensured.

The filter was run with 50 perturbed particles, plus the non-perturbed first guess. Figure 6 shows the frequency of the non-perturbed run being chosen to be the initial conditions versus analysis time. During the nights of NEAR-FOG, the non-perturbed first guess was chosen most of the times; it could be because in a stable atmosphere, the perturbations added to the analysis were better preserved during the simulation than in a neutral or unstable stratification, while in a neutral atmosphere, perturbation to the analysis were quickly smoothed during the simulation. During the day, the perturbed particles were most of the time closer to observations than the non-perturbed first guess. During FOG, the frequent occurrence of fogs changed this pattern; perturbed particles were chosen more often during the nights, because thick fogs or stratus occured, which ma the atmosphere less unstable or neutral. During the day, as for NEAR-FOG, the non-perturbed guess was seldom chosen to be the initial conditions.

The initial temperature and specific humidity provided by this algorithm replaced the ones that were given by the BLUE algorithm. The second step of the assimilation scheme, i.e. the initialization of liquid water content and adjustment of initial specific humidity profiles in case clouds are present at initialization time, was left unchanged.

3.4.3 Frequency of filter collapse

Before anything else, we have to check if the filter does not collapse. Figure 7 presents a frequency histogram of the maximal weight W^{max} for all simulations of the FOG and NEAR-FOG situations. The filter was run with ensemble sizes varying from 50 to 200 members, which is small

DOI: 10.1002/qj

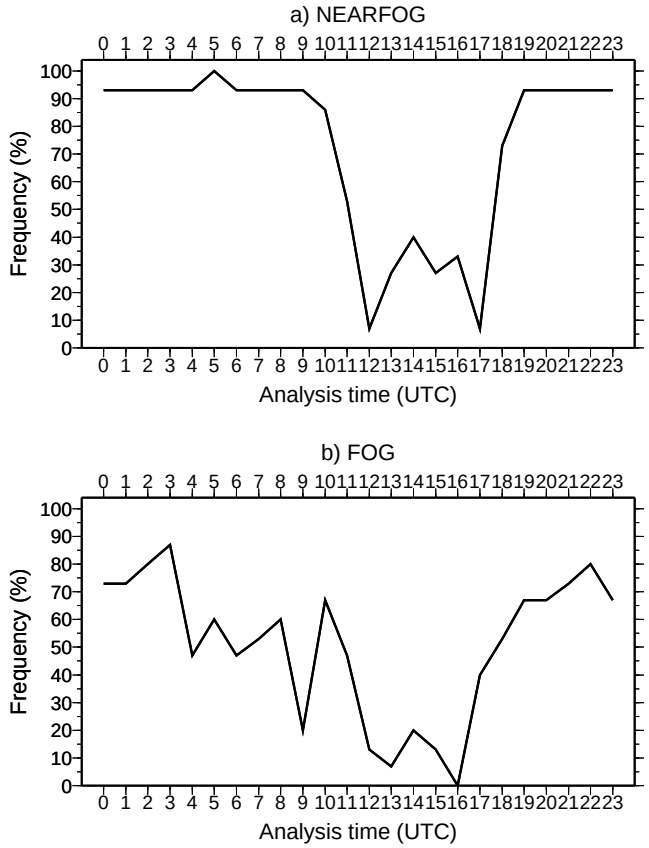

Figure 6. Frequency of the non-perturbed guess being chosen to be the initial conditions versus analysis time, for NEAR-FOG (a) and FOG (b).

compared to ensemble sizes used by Snyder et al. (2008) and Baehr and Pannekoucke (2009). Filter collapse can be diagnosed by diagrams strongly skewed toward the right: when the maximal weight is very close 1. For the FOG situation (figure 7a, c and e), it was not the case. When using 50 members; the filter was already rather convergent: for less than 10% of the simulations, the maximal weight was above 0.95. The diagrams were more and more skewed towards the left with increasing ensemble size, which means that the filter was more and more convergent. This can be explained by the fact that when more particles are available, the best one is likely to stand above the other ones less markedly, in terms of distance to the observations, than when fewer particles are used.

With the NEAR-FOG situation, collapse of the filter was more frequent when using a 50-particles ensemble; the maximal weight was larger than 0.95 for around one analysis in three. The frequency of filter collapse decreased with increasing ensemble sizes. The difference between FOG and NEAR-FOG lay in the occurence of deep fogs during FOG, which provoked a change in the stratification of the atmosphere at night.

Figure 8 shows the maximal weight $\mathcal{W}^{max}$ for every assimilation cycle of PART50 versus lower atmosphere stability, arbitrarily defined here as the gradient of potential temperature in the first 100m of the atmosphere, for NEAR-FOG and FOG. Cases when fog or stratus were present at analysis time are plotted in gray. When fog was present at analysis time, the atmosphere was either very stable, when the fog was shallow, or neutral when it was thicker or in the presence of stratus. There were clearly more situations with weakly stable, neutral or unstable lower atmosphere for FOG, than for NEAR-FOG; and many of these situations were linked to the presence of fog or stratus. For NEAR-FOG, when the atmosphere was stable, the variability of the maximal weight divided by the sum of all weights was less important than with FOG. There appears to be a dependancy between stability and the frequency of filter divergence, with a threshold around -2 to -3K for potential temperature gradient. Below that value, filter divergence was frequent; while it was quite rare when stability was above that value. For FOG, the dependancy is less clear, though overall filter divergence was significantly more frequent for strongly stable atmospheres than with weakly stable, neutral or unstable atmospheres.

This explanation of the different behaviour of the filter depending on the stratification of the atmosphere lies in how the initial perturbations are preserved or smoothed during the simulation. The atmosphere is neutral or weakly unstable during the day and at night if deep fog or stratus occurs. With a neutral or unstable atmosphere, the initial perturbations are quickly smoothed during simulations; the distances of the particles are then rather close and filter divergence is avoided. Stable atmosphere occur during nights with clear-sky or shallow fogs. When

DOI: 10.1002/qj

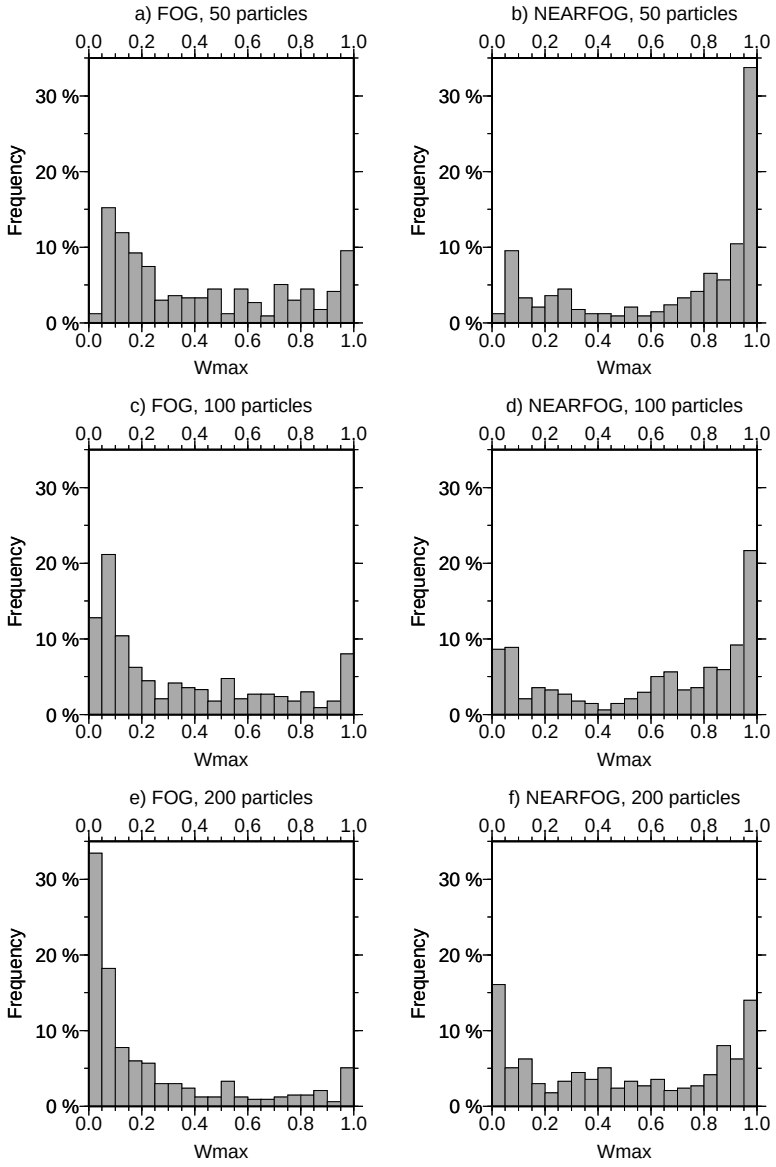

Figure 7. Frequency histogram of $\mathcal{W}^{max}$ for all simulations of the FOG (left) and NEAR-FOG (right) situation; simulations using 50 (a and b), 100 (c and d), 200 (e and f) particles.

the atmosphere is stable, the initial perturbations are not modified much through the simulations and the distances of the particles to the observations are larger; the filter is then likelier to collapse. As shown before, in this case, the non-perturbed guess was often chosen to be the initial conditions.

Filter divergence was linked for most cases to the stratification of the atmosphere. For the same dimension of the model space, it occured less frequently than in Snyder et al. (2008), for ensemble sizes much smaller than were used in their work. The frequent convergence of the modified particle filter was due to the selection stage, which eliminated particles that were distant from the observations. The fact that the noise added to the initial state of the particles was coherent with the model uncertainty and bounded also allowed to run with fewer particles than if they were purely random. The selection step was not included in the kind of particle filter that Snyder et al. used., which could explain the different behaviour of the particle filter. The results in terms of filter divergence frequency versus ensemble size are in the same range as the ones obtained by Baehr and Pannekoucke (2009) for the NEAR-FOG situation. Convergence was more frequent for the FOG situation. perturbations were used.

3.5 Soil data assimilation

In the operational setup, the soil observations are simply interpolated to the ISBA grid to provide initial conditions. During the simulation, COBEL-ISBA adjusts the values of temperature and humidity in the lower levels of COBEL and the upper levels of ISBA through its physical processes, in order to reach some kind of equilibrium that is consistent with its parameterized processes. Figure 9 illustrates this phenomenon; with the operational setup, unrealistic initial values of sensible and latent heat fluxes are quickly adjusted in the first 15 minutes of the simulation.

DOI: 10.1002/qj

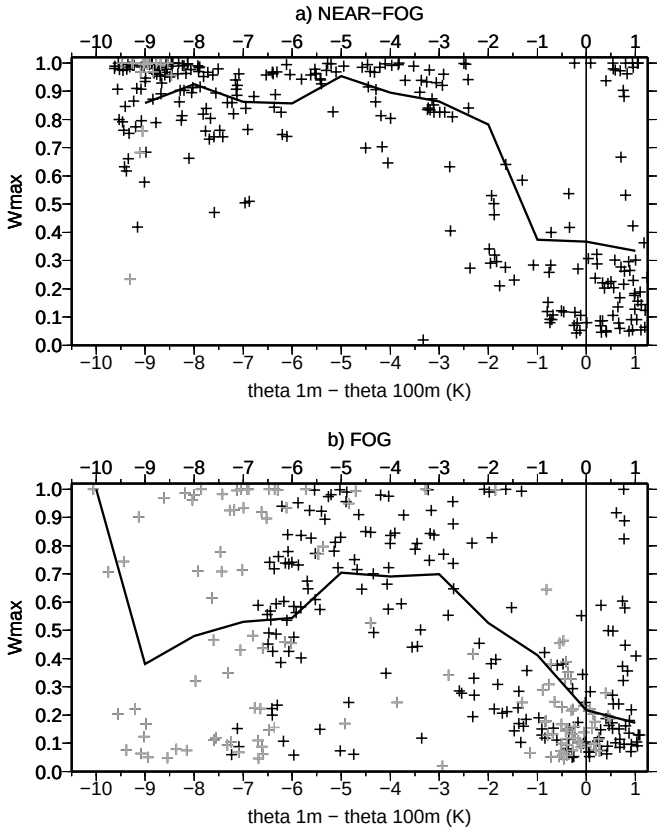

Figure 8. $\mathcal{W}^{max}$ for every assimilation cycles of PART50 with NEAR-FOG(a) and FOG(b) versus potential temperature gradient in the first 100m of the atmosphere. When fog or stratus was present at analysis time, the corresponding cross is plotted in gray. The mean for every stability interval of 1K is plotted by a continuous line.

This adjustment brought a sharp peak in the forecasted 2m temperature and a brutal increase in the forecasted 2m specific humidity. This phenomenon is frequent for simulations with maximal solar radiation and is a source of specific humidity bias.

This problem was especially troublesome within a particle filter, as for many simulations it concerned all particles, perturbed or non-perturbed. When using the adapted particle filter with the original soil initialization, particles all showed the same bias for specific humidity for simulations that started between 10UTC and 15UTC. To prevent this problem, a simplified version of a particle filter was set up to provide the initial conditions for ISBA. The ISBA first guess that was closest to observations of soil temperature and water content was chosen to be the ISBA initial conditions for the non-perturbed run. A random perturbation was added to these profiles to produce the initial conditions of the perturbed particles. That means that the selection step consisted here only to keep the closest particle and to eliminate all others. The distance between observations and the ISBA backgrounds were computed over a time window, as they were for COBEL. The rationale behind this algorithm was to provide ISBA with initial conditions produced by the model itself, as it is the case for COBEL. The adjustment that usually occurs at the beginning of the simulation would then already be taken in account in the initial conditions of both ISBA and COBEL. Figure 9 shows an example of how the problem of the interface between soil and atmosphere was partially solved following the implementation of this algorithm. There was still some adjustment on the sensible heat flux, but the impact on 2m temperature and specific humidity was much smaller than with the operational setup.

4 Results of the filter

The performance of the filter was assessed against the REF experiment to evaluate the improvement or degradation of the new assimilation algorithm as compared to the operational setup. Scores on temperature and specific humidity were computed and, for the FOG situation, the impact of the new assimilation scheme on the quality of the forecast of LVP events was also estimated. The experiments were called PART50, PART100, PART200 depending on the size of the (perturbed) particle ensembles. In this section the results of PART50 are shown; the influence of the ensemble size will be discussed in a specific section.

COBEL-ISBA was designed to forecast radiation fog, which is a phenomenon that occurs in the lower part of the model's domain. As a consequence, when discussing the scores, more emphasis will be put in the first 100m of the domain.

DOI: 10.1002/qj

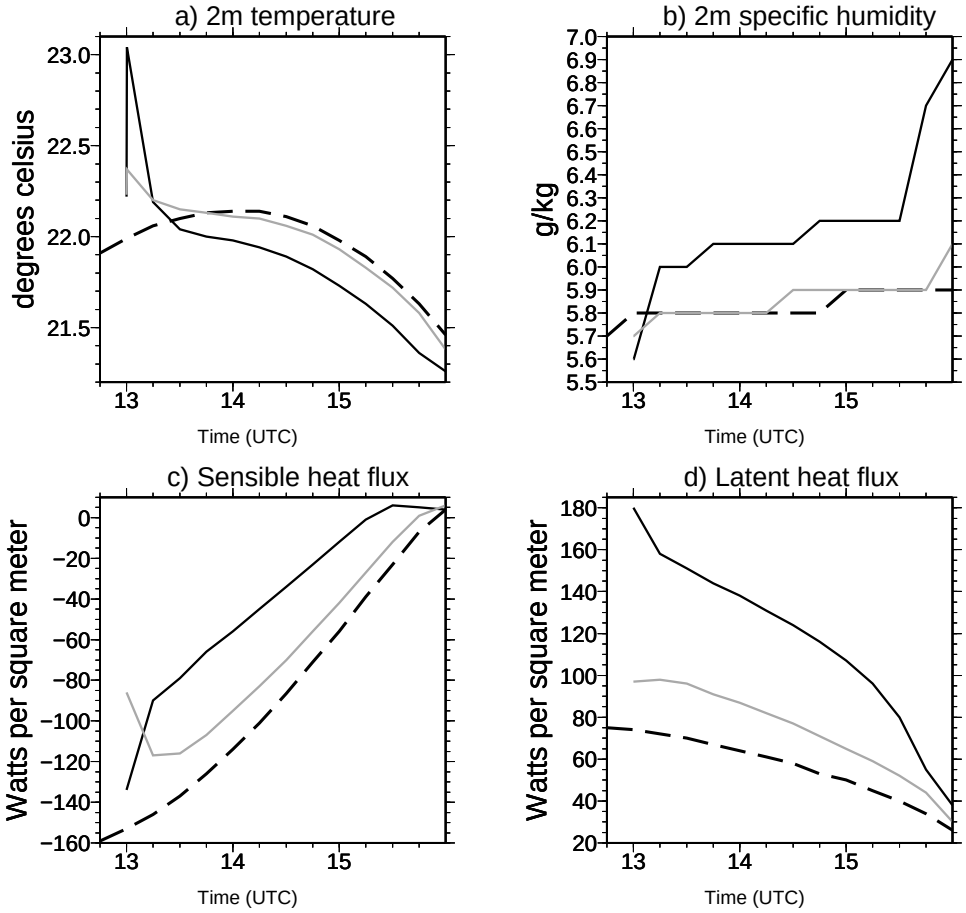

Figure 9. NEAR-FOG: simulation starting at Day 4, 13UTC; 2m temperature (left) and the sum of latent and sensible heat fluxes (right). Observations are plotted by a dashed line, simulation with the operational setup is represented by a black line; with the new assimilation scheme, by a gray line.

4.1 NEAR-FOG situation

Figure 10 shows the Root Mean Square Error (RMSE) of PART50 as a percentage of REF's RMSE for temperature and specific humidity, and also the bias difference between the two experiments for NEAR-FOG. The RMSE of initial temperature was improved by up to 20-40% above 80m and degraded by up to 10% below 20m. For initial specific humidity the RMSE was reduced by 25 to 55% above 100m, and slightly degraded below 20m. PART50 did not improve the initial RMSE in the lowest part of the domain. An explanation for this is that the distance between the particles and observations was minimized over a time window and not just at analysis time; the particle that was selected to be the initial conditions may not be the one closest to the observations at analysis time. Also, the initial conditions of REF were very close to the observations from the mast and the weather station there, since the variances of the measurements from the mast and the weather station used in the BLUE were much smaller than the ones of both the ALADIN profiles and the first guess. The initial temperature bias was slightly degraded by PART50 as compared to REF, that of specific humidity was unchanged below 100m and slightly improved above that.

The usefulness of taking a first guess as the initial conditions and of assimilating data over a time window appeared fully during the forecast. As the initial conditions were coherent with the model's physical processes, the forecast was rapidly of much better quality for PART50 as compared to REF. The improvement reached 35 to 45% for specific humidity and 25 to 30 % for temperature. The bias was also reduced in the lower part of the domain for temperature after 2 hours of forecast and for specific humidity over the whole domain after 1 hour of forecast. This shows that the initialization of ISBA and the interface between COBEL and ISBA worked better with the new algorithm then with the operational setup, as it was shown by previous studies (Rémy and Bergot (2009a)) that a faulty initialization of ISBA is a cause of increasing forecasted bias on temperature and specific humidity.

These results were obtained with a filter that was often diverging during the nights (see figure 7). This was not too detrimental in our case, as the filter was used within a deterministic approach. The most important in this framework is that the filter provides good quality first guesses to be used as initial conditions. The noise added to the initial conditions also contributed to increase the spread of the ensemble even when the filter collapsed.

DOI: 10.1002/qj

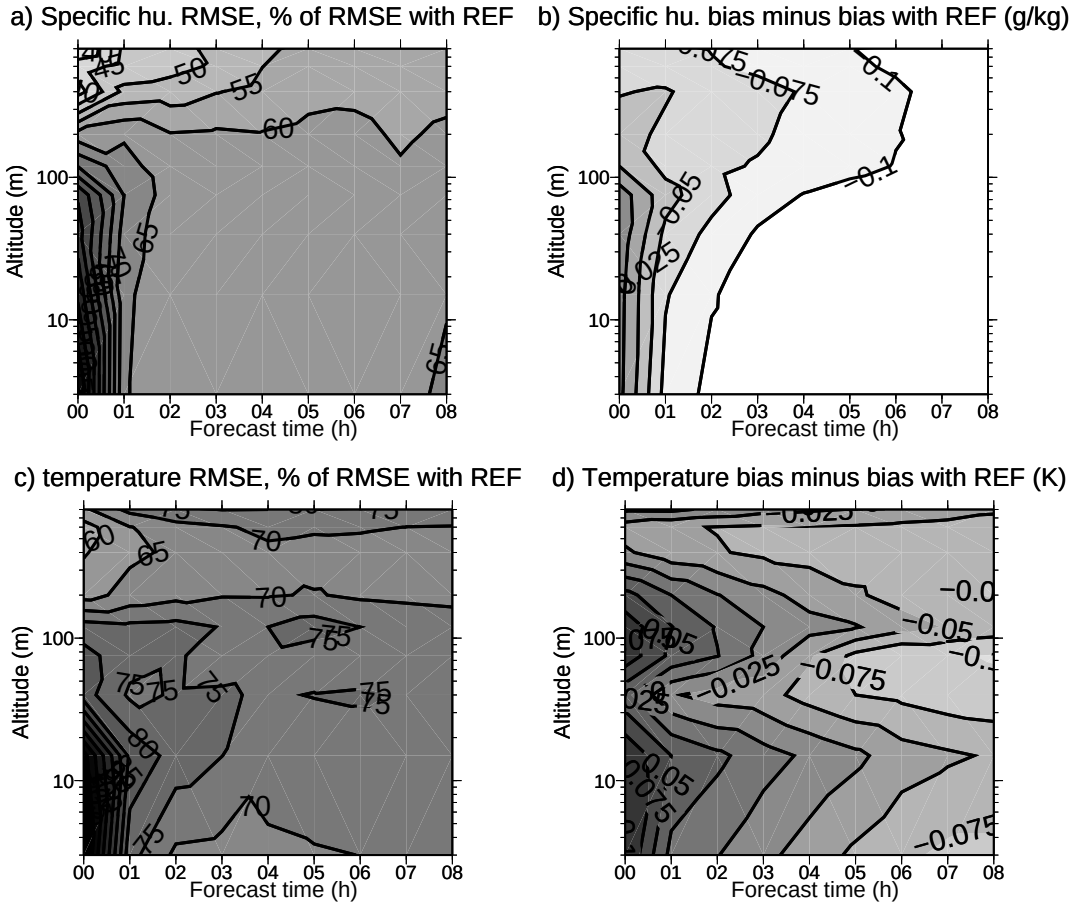

Figure 10. NEAR-FOG: RMSE of PART50 as a percentage of the RMSE of REF (left) and bias of PART50 minus bias of REF (right) versus forecast time, for temperature (top) and specific humidity (bottom).

4.2 FOG situation

4.2.1 Scores on temperature and specific humidity

Figure 11 shows the RMSE of PART50 as a percentage of REF's RMSE for temperature and specific humidity, and also the bias difference between the two experiments for FOG. PART50 improved the initial conditions as compared to REF. For specific humidity the initial RMSE was reduced by 40 to 45 % over the whole domain. As for temperature, the initial RMSE improvement was larger above 50m, with a reduction of 30 to 45% above that height and of 10 to 25% below. For both temperature and specific humidity, the initial bias difference as compared to REF was very small

As for simulations with NEAR-FOG, the temperature RMSE was reduced by larger margin during the simulation than for the initial state. For specific humidity, the improvement is in the same range for the forecast and for the initial state. After one hour of forecast, the RMSE was improved by up to 35-45% for temperature and specific humidity. The bias was slightly degraded in the lower part of the domain for temperature and left unchanged for specific humidity. It was improved in most other part of the domain.

4.2.2 Forecast of LVP events

Figure 12 shows the frequency distribution histogram of the onset and the burnoff time of LVP events, for all simulation times and forecast times, for the FOG situation. Simulations in which fog was already present at initialization time were discarded for the computation of the onset scores. For these simulations, it was meaningless to compare the simulated and observed onset times because the fog events considered had begun before the initialization time. The errors larger than 240 minutes are grouped together in the 240 minutes column. The mean and standard deviation of errors smaller than 240 minutes are also indicated.

The onset time of low visibility conditions was generally forecasted too early for REF: there was small negative bias for this experiment. This bias was corrected and even inverted by PART50, with onset time generally forecasted too late. The errors were generally smaller for PART50 than for REF. The frequency of errors being smaller or equal to 30 minutes was raised from 30% for REF to 45% for PART50 and the standard deviation of the error was smaller. The errors larger than 240 minutes were significantly less frequent. PART50 also improved markedly the prediction of LVP burnoff time as compared to REF. The errors were generally smaller with much fewer errors larger than 240 minutes. The frequency of errors being

DOI: 10.1002/qj

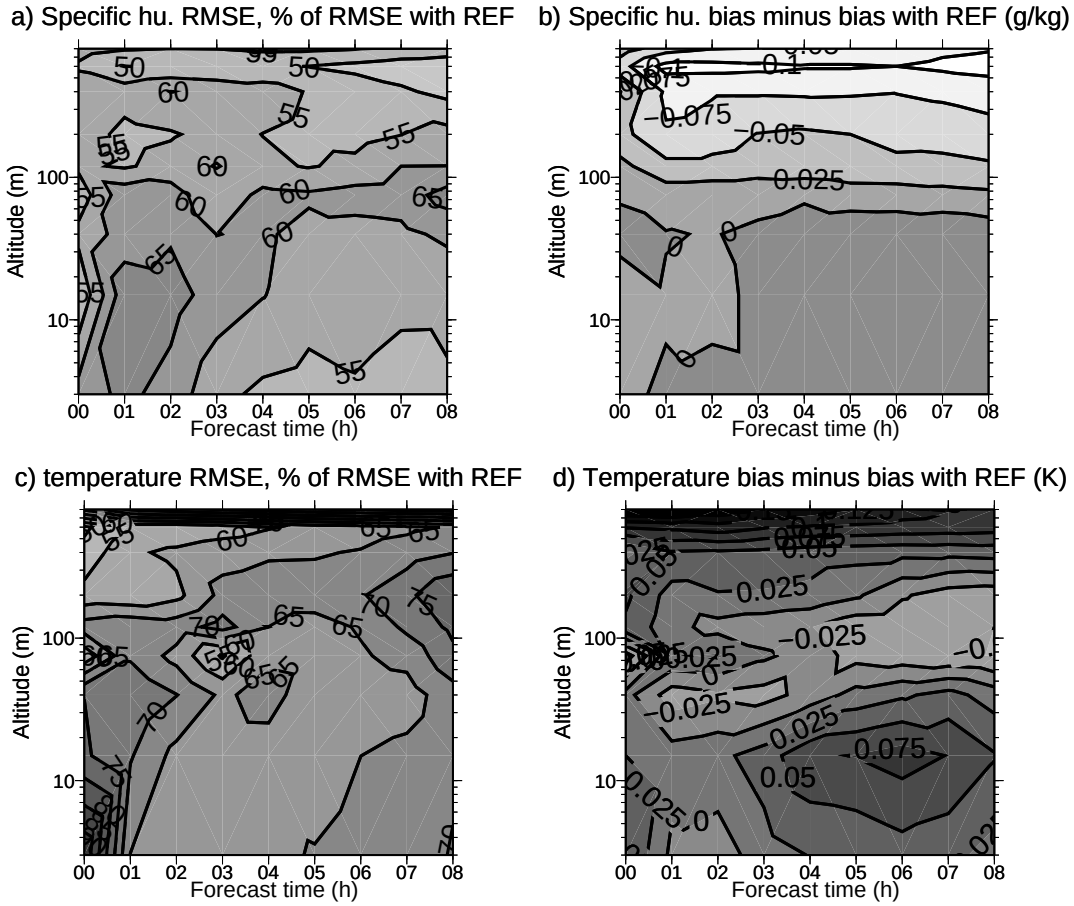

Figure 11. Same as 10 for FOG.

Table I. Hit Ratio (HR) of LVP conditions for various forecast times for the FOG situation and for the REF, PART50 and ENKF32 experiments. EnKF32 values are taken from Rémy and Bergot (2009b).

	1h00	2h00	3h00	4h00	6h00	8h00	all
REF	0.93	0.89	0.89	0.88	0.86	0.84	0.88
PART50	0.93	0.94	0.97	0.98	0.98	0.98	0.97
ENKF32	0.95	0.92	0.93	0.95	0.93	0.93	0.94

Table II. Pseudo False Alarm Ratio (FAR) of LVP conditions for various forecast times for the FOG situation and for the REF and PART50 experiments. EnKF32 values are taken from Rémy and Bergot (2009b).

	1h00	2h00	3h00	4h00	6h00	8h00	all
REF	0.07	0.05	0.07	0.10	0.12	0.18	0.09
PART50	0.01	0.00	0.03	0.01	0.09	0.09	0.04
ENKF32	0.04	0.03	0.02	0.06	0.08	0.15	0.07

smaller or equal to 30 minutes was raised from 40% for REF to 70% for PART50. The negative bias of REF for the forecast of burnoff time was reduced by PART50.

Tables I and II show the Hit Ratio (HR) and pseudo False Alarm Ratio (FAR) of LVP conditions for various forecast times and for the REF and PART50 experiments. In the case of rare event forecasting, such as fog and LVP conditions, the pseudo-FAR is convenient because it removes the impact of the "no-no good forecasts" (no LVP forecast and no LVP observed), which mostly dominate the data sample and hide the true skill of the LVP forecast system. If a is the number of observed and forecasted events, b the number of not observed and forecasted events, and c the number of observed and not forecasted events, HR and pseudo-FAR are then defined as follows:

$$HR = \frac{a}{a+c}; \qquad pseudoFAR = \frac{b}{a+b}$$

Table I shows that the detection of LVP conditions was improved for all forecast times larger than 1 hour, and that the overall hit ratio was significantly higher for PART50 than for REF. The improvement was larger for longer forecast times, corresponding to the largest improvements in temperature and specific humidity RMSE as compared to REF. Also, the hit ratio of LVP conditions did not decrease with time with PART50, while it did with REF. This shows the strong influence of the initial conditions on the forecast when the model error has been removed by using simulated observations. Table II shows that PART50 experienced fewer false alarms than REF. The number of

DOI: 10.1002/qj

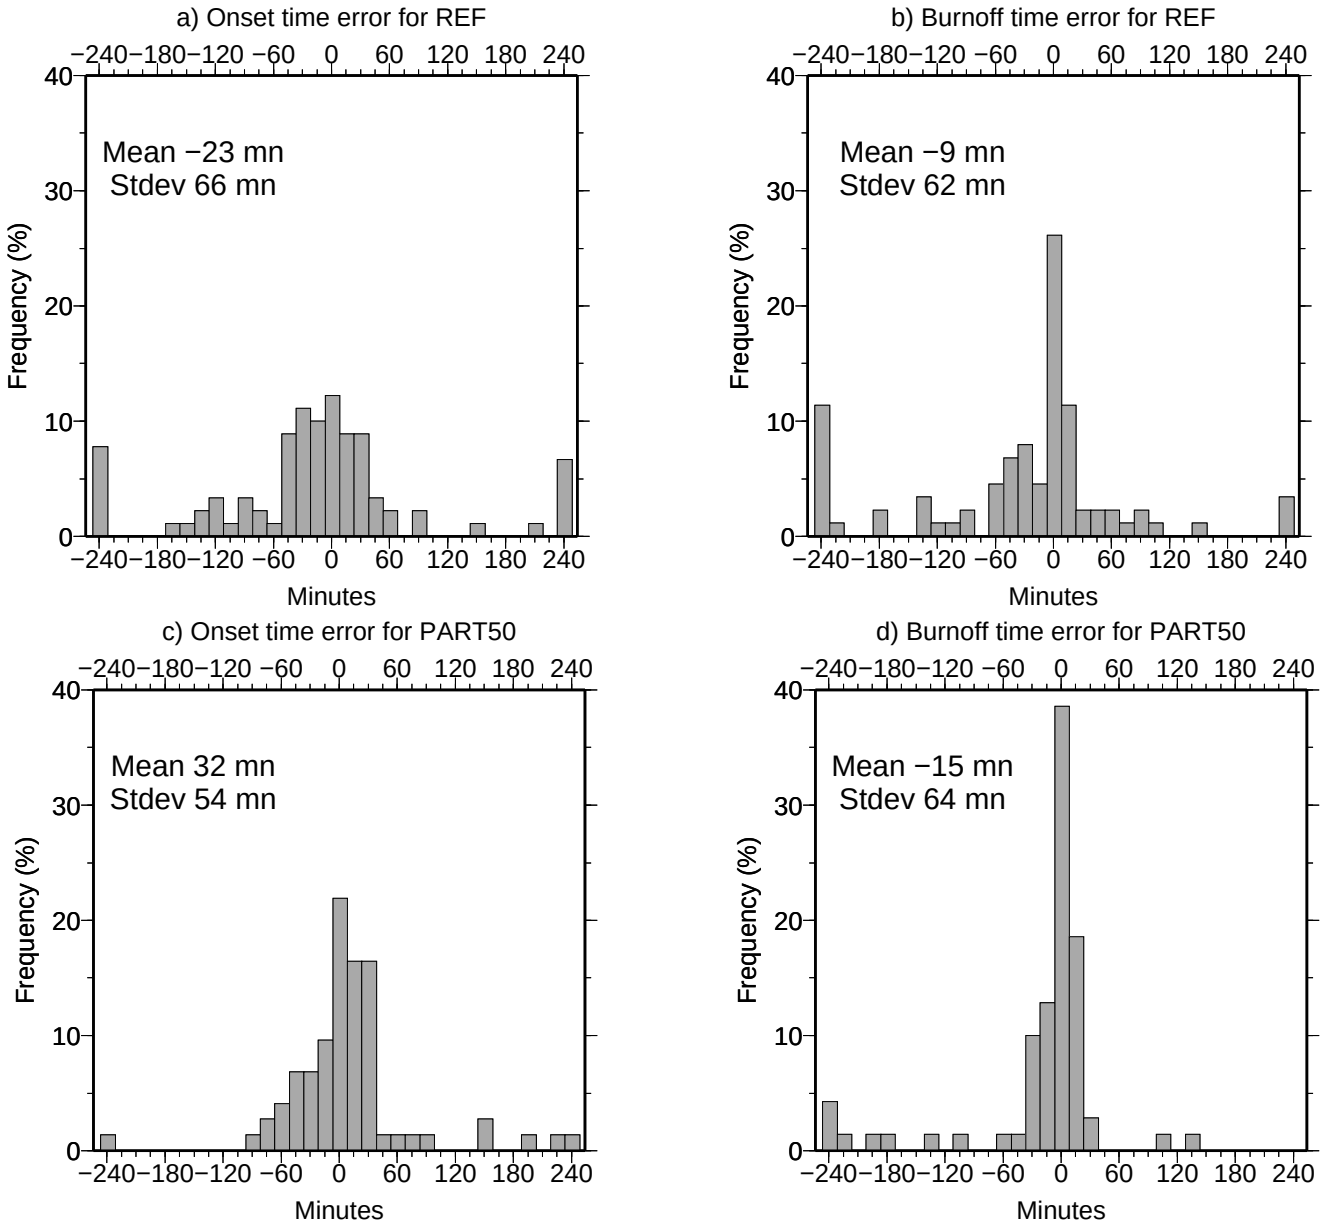

Figure 12. FOG: Frequency distribution histogram of the error on onset time (left, the LVP conditions at initial time are not taken into account) and burnoff time (right) of LVP conditions, in minutes. REF experiment is at the top, PART50 at the bottom. Positive values correspond to a forecast of onset or burnoff that is too late. Errors larger than 240 minutes are grouped in the 240 minutes column. The mean and standard deviation of errors smaller than 240 minutes are indicated.

false alarms did not increase much with forecast time. This is an interesting result since an improvement in both HR and pseudo-FAR is hard to obtain.

4.3 Comparison with an ensemble Kalman filter

The ensemble Kalman filter (Evensen (1994) and Evensen (2003)) is an assimilation scheme that uses an ensemble of first guesses to estimate the background error statistics, which are then used in the BLUE algorithm that computes the initial conditions for the ensemble and the non-perturbed run. This scheme has been implemented in various oceanic and atmospheric models (Houtekamer et al. (2005), Zhang (2005) and Sakov and Oke (2008) among others). A "perturbed observations" version (Burgers et al. (1998)) of the ensemble Kalman filter was run with FOG and NEAR-FOG using ensemble of 8, 16 and 32 members. As the ensembles used were rather small, the covariances were inflated using an adaptive covariance inflation algorithm (Anderson (2007)). The results are described in Rémy and Bergot (2009b). As the ensemble size didn't impact much the quality of initial conditions and forecasts when using the ensemble Kalman filter with simulated observations, it was possible to qualitatively compare the results of the 32 members ensemble Kalman filter (experiment ENKF32) with the ones obtained with PART50.
Figure 13 shows the RMSE of PART50 as a percentage of ENKF32's RMSE for temperature and the bias difference between the two experiments for NEAR-FOG and FOG.

DOI: 10.1002/qj

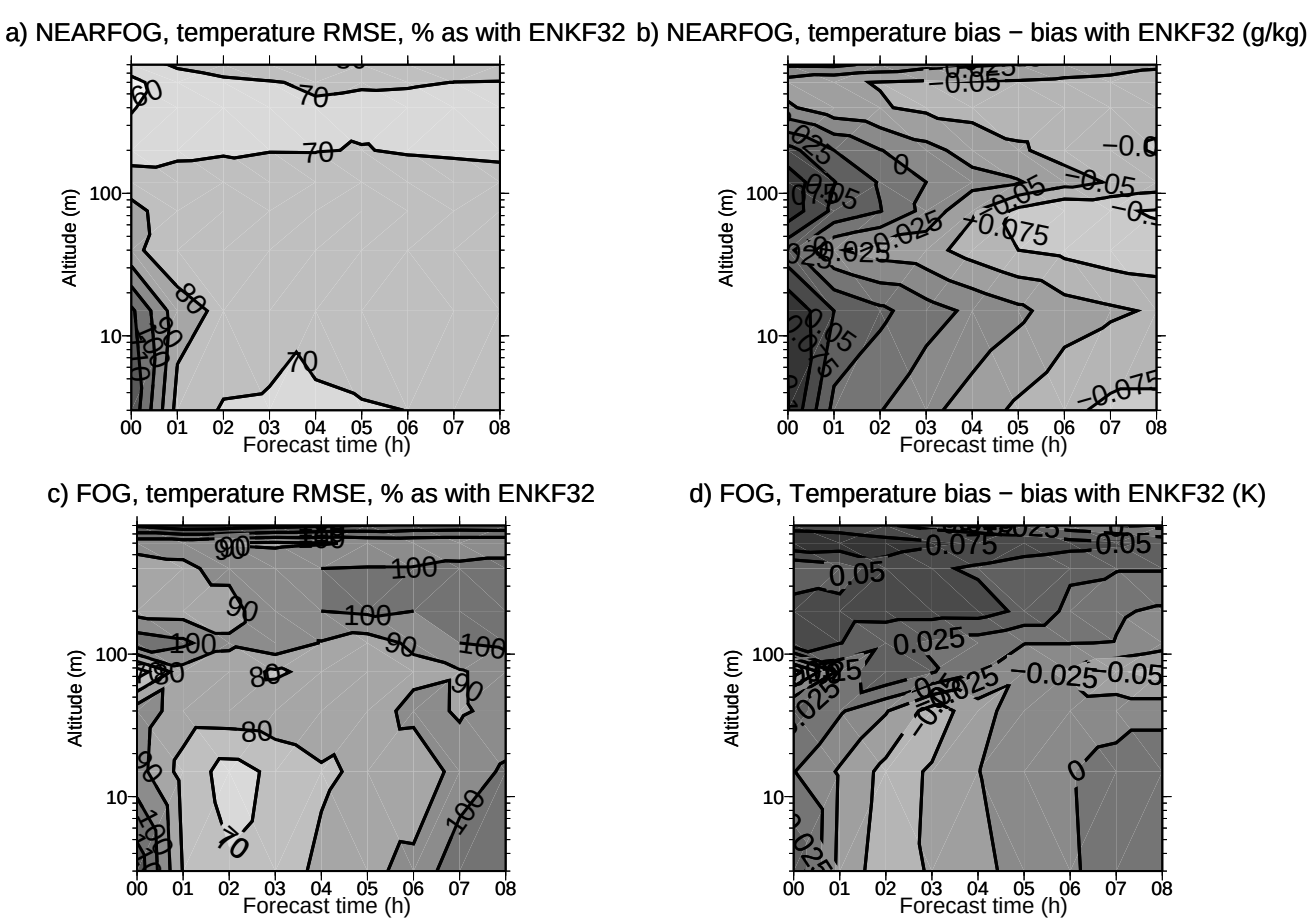

Figure 13. Temperature RMSE of PART50 as a percentage of the RMSE of ENKF32 (left) and bias of PART50 minus bias of ENKF32 (right) versus forecast time, for NEAR-FOG (top) and FOG (bottom). Data for ENKF32 are taken from Rémy and Bergot (2009b).

Specific humidity scores are not shown as they display the same patterns. The RMSE of temperature at initialization time was slightly degraded in the first 20m of the domain and improved elsewhere by PART50 as compared with ENKF32 for NEAR-FOG; while the initial bias was mostly unchanged. Forecasted temperature RMSE was significantly improved for NEAR-FOG, by up to 30%; and the forecasted temperature bias was also smaller for PART50 as compared with ENKF32. The overall improvement of PART50 as compared to ENKF32 increased with forecast time.

For FOG, the initial temperature RMSE was larger for PART50 as compared with ENKF32 below 10 m. Above that height, there was a small improvement of 5 to 10%. The initial bias was similar for both experiments below 100m and slightly larger for PART50 above 100m. During the forecast, PART50 displayed smaller RMSEs than ENKF32 by 10 to 30% below 100m and up to 7 hours of forecast. Elsewhere, the differences between PART50 and ENKF32 were smaller. The forecasted temperature bias was slightly smaller for PART50 as compared to ENKF32 below 80m, and slightly larger above that height.

The HR and pseudo-FAR of LVP conditions were close between the two experiments, as shown by tables I and II. PART50 had an overall HR slightly higher than ENKF32 and a smaller pseudo-FAR. ENKF32 showed a higher detection rate for forecast times of 1 hours while PART50 was better for higher forecast times. For the pseudo-false alarm rate, PART50 and ENKF32 showed scores in the same range, except for a forecast time of 8 hour, for which PART50 was significantly better. PART50 also predicted onset and burnoff time more accurately than ENKF32 (not shown).

5 Impact of the ensemble size

The size of the ensemble influenced the frequency of filter collapse, especially for the NEAR-FOG situation (see figure 7). In this section, the impact on the initial conditions and forecasts is assessed.

Overall, no consistent tendancy can be drawn for the impact of the ensemble size on the RMSE of analyzed and forecasted temperature and specific humidity (not shown). The scores of PART100 and PART200 were slightly better or worse than PART50 depending on the height and the forecast time, but no correlation could be drawn between the quality of these scores and the ensemble size.

The same conclusion holds also for the specific humidity bias (not shown). A consistent impact of the ensemble size on the temperature bias was however found. Figure 14 shows the temperature bias difference between

DOI: 10.1002/qj

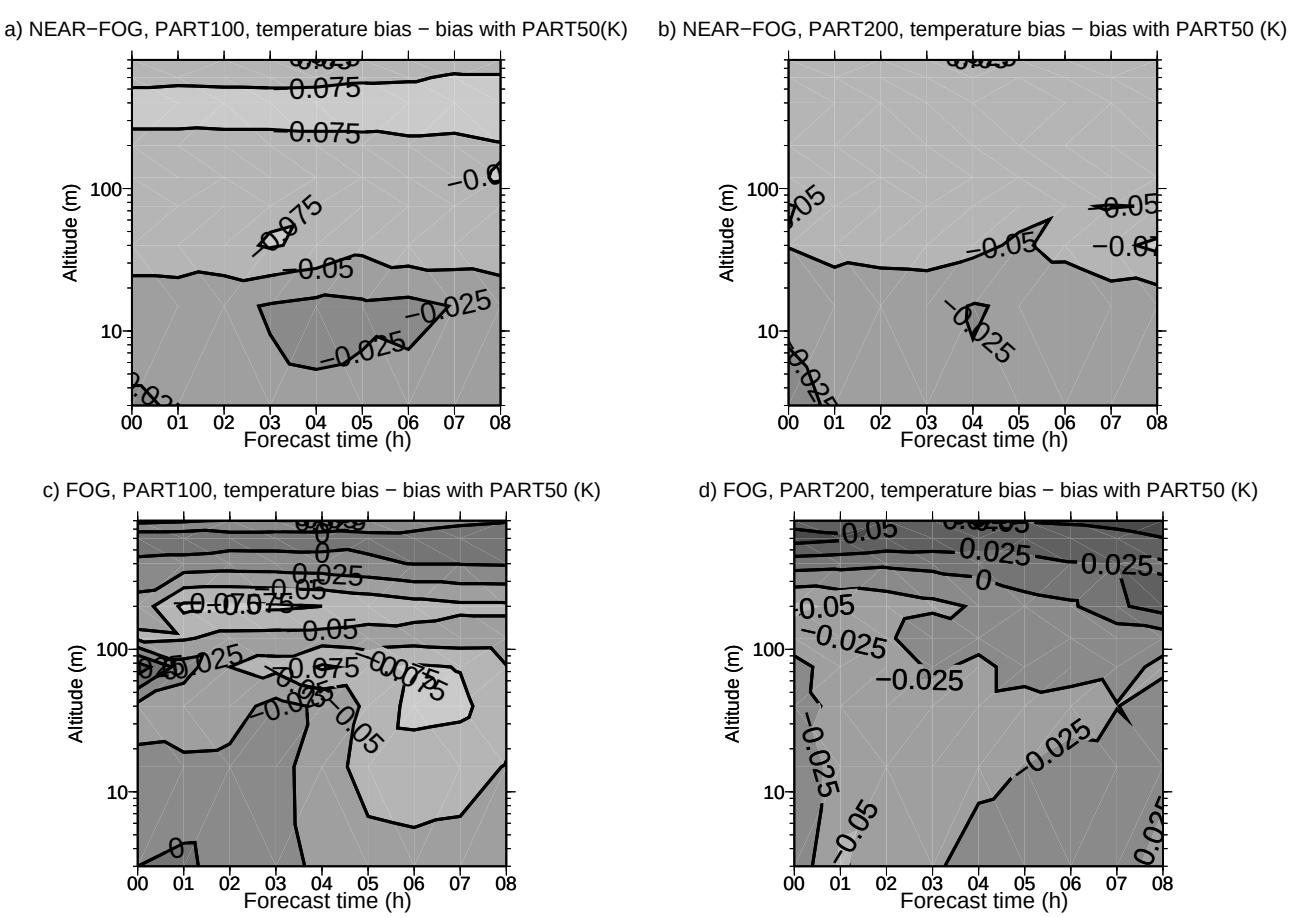

Figure 14. Temperature bias difference between PART100 and PART50 (left) and PART200 and PART50 (right) for NEAR-FOG (top) and FOG (bottom). A negative value indicates that PART100 or PART200 had a smaller temperature bias than PART50.

PART100/200 and PART50, for FOG and NEAR-FOG. It can be seen that for both situations, PART100 and PART200 showed significantly better temperature bias, for the initial profiles as well as for the forecasted ones.

Overall, correlation between the ensemble size and the quality of the initial conditions and forecasts was weak, though the impact of ensemble size on the convergence of the filter was marked. Convergence frequency and the quality of initial conditions appear thus to be decoupled. This could be due to the fact that the filter was used within a deterministic approach: the goal of the filter is to provide an accurate first guess, not to describe fully all the possible states of the background. The noise added to each particle at initial time during the resampling stage explains also this result, as it increased the spread of the ensemble even when the filter diverged. The fact that this noise was bounded also probably helped to keep the best particles, even when the overall number of particles was not very large, close to observations.

This study shows also that for these two particular situations, the filter worked well with 50 particles with simulated observations and that more particles does not bring further information on the probability distribution of the backgrounds. This conclusion has to be confirmed using real observations.

6 Summary and discussion

A challenge of data assimilation is to provide the model with initial conditions that are at the same time close to the true state of the atmosphere, and coherent with the processes that are modelled in the system. Both requirements are harder to reach with strongly non-linear systems such as COBEL-ISBA. An algorithm based on a genetic selection particle filter was developed, with modifications brought to take into account the particle's time trajectories. Experiments using this new assimilation scheme were assessed against experiments using the operational setup of COBEL-ISBA that consists of a BLUE algorithm. This work showed that an algorithm based on particle filtering with genetic selection is able to provide accurate initial conditions to a 1D model using a reasonable numbers of particles, within a simulated observations framework. Filter collapse was less frequent, given the size of the model space and of the ensembles that were used, than with experiments carried out by Snyder et al. (2008), thanks to the genetic selection algorithm and the noise added to each particle at analysis time. The divergence frequency of the filter was shown to depend on the stratification of the atmosphere. Both temperature and specific humidity analysis were improved as compared to the BLUE algorithm used in the operational setup. As the initial conditions were given by the first guess that had the closest

DOI: 10.1002/qj

trajectory to the observations, the initial conditions met the two conditions mentioned above. Thanks to that also, the forecasted temperature and specific humidity were improved by a larger margin than the analyzed ones. The better quality of the initial conditions and forecasts brought better forecasts of LVP events. The final product delivered by COBEL-ISBA, i.e. hours of forecasted occurence or lift-off of LVP conditions, was markedly improved by the new assimilation scheme.

The conclusions on the convergence of the filter and on the adequate number of particles needed to run the model are model-dependant and also situation-dependant: the results vary from one set of simulated observations to another. Nethertheless, they can be helpful for future implementations of assimilation algorithms based on particle filtering, as they provide general insights on the causes and mechanisms of filter divergence.

A next stage will be to test this assimilation scheme with real observations. As the physics underlying the observations and the simulations were the same when using simulated observations, the task of producing initial conditions consistent with the model's physics was simplified. The real atmosphere is non-linear to greater extent than a simulated one; particle filtering is an assimilation algorithm that was designed for non-linear dynamical systems, so it seems fit for the task.

References

Anderson, L., 2007: An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus*, **59**, 210–224.

Baehr, C. and O. Pannekoucke, 2009: Some issues and results on the enkf and particle filters for meteorological models. *Chaos 2009*, **early online release**.

Bergot, T., 1993: Modélisation du brouillard à l'aide d'un modèle 1d forcé par des champs mésoéchelle : application à la prévision. Ph.D. thesis, Université Paul Sabatier, 192 pp., [available at CNRM, Meteo-France, 42 Av. Coriolis, 31057 Toulouse Cedex, France.].

Bergot, T., D. Carrer, J. Noilhan, and P. Bougeault, 2005: Improved site-specific numerical prediction of fog and low clouds: a feasibility study. *Weather and Forecasting*, **20**, 627–646.

Bergot, T. and D. Guédalia, 1994a: Numerical forecasting of radiation fog. part i : Numerical model and sensitivity tests. *Mon. Wea. Rev.*, **122**, 1218–1230.

Bergot, T. and D. Guédalia, 1994b: Numerical forecasting of radiation fog. part ii : A comparison of model simulation with several observed fog events. *Mon. Wea. Rev.*, **122**, 1231–1246.

Boone, A., 2000: Modélisation des processus hydrologiques dans le schéma de surface isba : inclusion d'un réservoir hydrologique, du gel et modélisation de la neige (modeling of hydrological processes in the isba land surface scheme : inclusion of a hydrological reservoir, freezing, and modeling of snow). Ph.D. thesis, Université Paul Sabatier, 207 pp., [available at CNRM, Meteo-France, 42 Av. Coriolis, 31057 Toulouse Cedex, France.].

Bougeault, P. and P. Lacarrere, 1989: Parameterization of orography-induced turbulence in a mesoscale model. *Mon. Wea. Rev.*, **117**, 1872–1890.

Burgers, G., P. V. Leuwen, and G. Evensen, 1998: Analysis scheme in the ensemble kalman filter. *Mon. Wea. Rev.*, **126**, 1719–1724.

Clark, D., 2002: Terminal ceiling and visibility product development for northeast airports. *14th Conf. on Aviation, Range, and Aerospace Meteorology*, AMS, [available at http://www.ll.mit.edu/mission/aviation/publications/publication-files/ms-papers/Clark_2002_ARAM_MS-15290_WW-10474.pdf].

Clark, D., 2006: The 2001 demonstration of automated cloud forecast guidance products for san-francisco international airport. *10th Conf. on Aviation, Range, and Aerospace Meteorology*, AMS, [available at http://jobfunctions.bnet.com/abstract.aspx?docid=321609].

Del Moral, P., 2004: *Feynman-Kac formulae, genealogical and interacting particle systems and applications*. Springer-Verlag.

Doucet, A., N. de Freitas, and N. Gordon, 2001: *Sequential Monte-carlo methods in practice*. Springer-Verlag, 581 pp.

Estournel, C., 1988: Etude de la phase nocturne de la couche limite atmospherique (study of the nocturnal phase of boundary layer). Ph.D. thesis, UniversitÃ© Paul Sabatier, 161 pp., [available at CNRM, Meteo-France, 42 Av. Coriolis, 31057 Toulouse Cedex, France.].

Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using monte-carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, 10 142–10 162.

Evensen, G., 2003: The ensemble kalman filter : theoretical formulation and practical implementation. *Ocean Dynamics*, **53**, 343–367.

Gordon, N., D. Salmond, and A. Smith, 1993: Novel approach to nonlinear/non gaussian bayesian state estimation. *IEE Proceedings*, IEE, Vol. 140, 107–113.

Herzegh, P., S. Benjamin, R. Rasmussen, T. Tsui, G. Wiener, and P. Zwack, 2003: Development of automated analysis and forecast products for adverse ceiling

DOI: 10.1002/qj

and visibility conditions. *19th Internat. Conf. on Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology*, AMS, [available at http://ams.confex.com/ams/pdfpapers/57911.pdf].

Houtekamer, P., H. Mitchell, G. Pellerin, M. Buehner, M. Charron, L. Spacek, , and B. Hansen, 2005: Atmospheric data assimilation with an ensemble kalman filter: results with real observations. *Mon. Wea. Rev.*, **133**, 604–620.

Huang, X., H. Wang, Y. Chen, X. Zhang, S. Tjernkes, and R. Stuhlmann, 2007: An observing system simulation experiment using both mm5 and wrf : experiment configuration and preliminary results. *8^{th} WRF annual meeting*, UCAR, [available at http://www.mmm.ucar.edu/wrf/users/workshops/WS2007/abstracts/p2-2_Huang.pdf].

Lorenz, N., 1996: Predictability: a problem partially solved. *Seminar on predictability*, ECMWF.

Noilhan, J. and S. Planton, 1989: A simple parameterization of land surface processes for meteorological models. *Mon. Wea. Rev.*, **117**, 536–549.

Rémy, S. and T. Bergot, 2009a: Assessing the impact of observations on a local numerical fog prediction system. *Quart. J. Roy. Meteor. Soc.*, **135**, 1248–1265.

Rémy, S. and T. Bergot, 2009b: Ensemble kalman filter data assimilation in a 1d numerical model used for fog forecasting. *Mon. Wea. Rev.*, **in revision**.

Roquelaure, S. and T. Bergot, 2007: Seasonal sensitivity on cobel-isba local forecast system for fog and low clouds. *Pure Appl. Geophys.*, **164**, 1283–1301.

Sakov, P. and P. Oke, 2008: A deterministic formulation of the ensemble kalman filter: an alternative to ensemble square root filters. *Tellus*, **60A**, 361–371.

Snyder, C., T. Bengtsson, P. Bickel, and J. Anderson, 2008: Obstacles to high-dimensional particle filtering. *Mon. Wea. Rev.*, **136**, 4629–4640.

van Leeuwen, P., 2009: Particle filtering in geophysical systems. *Mon. Wea. Rev.*, **early online release**.

Zhang, F., 2005: Dynamics and structure of mesoscale error covariance of a winter cyclone estimated through short-range ensemble forecasts. *Mon. Wea. Rev.*, **133**, 2876–2893.

DOI: 10.1002/qj