Department of Meteorology

Localization in the ensemble Kalman Filter

Ruth Elizabeth Petrie

A dissertation submitted in partial fulfilment of the requirement for the degree of MSc. Atmosphere, Ocean and Climate

August 2008

Abstract

Data assimilation in meteorology seeks to provide a current analysis of the state of the atmosphere to use as initial conditions in a weather forecast. This is achieved by using an estimate of a previous state of the system and merging that with observations of the true state of the system. Ensemble Kalman filtering is one method of data assimilation. Ensemble Kalman filters operate by using an ensemble, or statistical sample, of the state of a system. A known prior state of a system is forecast to an observation time, then observation is assimilated. Observations are assimilated according to a ratio of the errors in the prior state and the observations. An analysis estimate of the system and an analysis estimate of the of the errors associated with the analysis state are also produced.

This project looks at some problems within ensemble Kalman filtering and how they may be overcome. Undersampling is a key issue, this is where the size of the ensemble is so small so as to not be statistically representative of the state of a system. Undersampling can lead to inbreeding, filter divergence and the development of long range spurious correlations. It is possible to implement counter measures. Firstly covariance inflation is used to combat inbreeding and the subsequent filter divergence. Covariance localization is primarily used to remove long range spurious correlations but also has the benefit increasing the *effective* ensemble size.

Specifically this project uses an implementation of the ensemble Transform Kalman filter (ETKF) a deterministic ensemble, with a simple model, to demonstrate the behaviour of the filter when undersampling is present. Covariance inflation was implemented, and was found to increase the accuracy of the analysis state. A new method of covariance localization by Schur product for the ETKF was introduced and implemented. This method was not consistent with the equations of the ETKF. The analysis estimate was detrimentally affected by this technique. By using covariance inflation in conjunction with this localization the performance may be improved. In its current state this implementation does not function as desired.

Acknowledgments

I would very much like to thank my supervisor Dr. Sarah Dance for all her time and effort spent helping me throughout the course of the project. I would also like to thank Dr. Ross Bannister for the assistance that he has given.

Many thanks also to the Royal Meteorological Society for funding me, without which I would not have been able to undertake this MSc. course.

Contents

List of Symbols

Matrices and vectors

$\mathbf{H}$	Linear observation Operator
$\mathbf{I}$	Identity Matrix
$\mathbf{K}$	Kalman gain matrix
$\mathbf{M}$	Linear dynamical model
$\mathbf{P}$	State error covariance matrix
$\mathbf{Q}$	Model error covariance matrix
$\mathbf{R}$	Observation error covariance matrix
$\mathbf{S}$	Ensemble version of innovations covariance matrix
$\mathbf{T}$	Post-multiplier in deterministic formulations of the EnKF
$\mathbf{U}, \mathbf{V}, \mathbf{W}$	Orthogonal matrices
$\mathbf{x}$	State vector
$\mathbf{X}$	State ensemble matrix
$\mathbf{y}$	Observation vector
$\mathbf{Y}$	Observation ensemble matrix
$\mathbf{z}$	Arbitrary vector
$\mathbf{Z}$	Arbitrary matrix
ϵ	Random observation error vector
η	Random model error vector
$\mathbf{\Lambda}$	Diagonal matrix of eigenvalues
$\mathbf{\Sigma}$	Matrix of singular values
ρ	Correlation covariance matrix

Sub and superscripts

a	Analysis value
e	Ensemble value
f	Forecast value
i	Ensemble Member (unless otherwise stated)
j	Time index (unless otherwise stated)
t	True value

Other

t	Time
L	Number of grid points
a	Advection model variable
n	Number of state components
m	Number of observation components
N	Number of ensemble members

Operators

$\langle \mathbf{z} \rangle$	Expectation operator
$\bar{\mathbf{z}}$	Mean value
$'$	Deviation from the mean
$\hat{\mathbf{z}}$	Scaled value
$\mathbf{Z}'$	Perturbation matrix

Abbreviations

KF	Kalman Filter
ENOI	ENsemble based Optimal Interpolation
EnKF	ENsemble Kalman Filter
ETKF	Ensemble Transform Kalman Filter
LEKF	Local Ensemble Kalman Filter
SRF	Square Root Filter
EnSRF	ENsemble Square Root Filter
NWP	Numerical Weather Prediction
SVD	Singular Value Decomposition
RMSE	Root Mean Square Error

List of Figures

Chapter 1

Introduction

1.1 What is Data Assimilation?

Data Assimilation in Meteorology is the process of combining observational data with a prior forecast state of the atmosphere to produce a more accurate estimate of the current state. Before a numerical weather prediction (NWP) model is run it is essential to have as accurate a representation of the current state of the atmosphere as is possible. Observations are made of variables such as pressure, temperature, humidity, wind speed and direction from sources such as radiosondes, areoplanes, satellites, ships and surface instrumentation. In a global model there are of $O(10^7)$ model state variables but only of $O(10^6)$ observations (UK Met Office, 2008). Clearly it is not possible to use the observational data alone to provide the initial conditions for a NWP model as there are not enough observations worldwide. The observations taken are combined with the previous information of the forecast state of the atmosphere to produce an estimate of the current state of the atmosphere to be used as initial conditions for the NWP model. The prior forecast state of the atmosphere is known as the *background state*. After the observations have been assimilated into the background state we obtain the *analysis state* of the system.

There are two main types of assimilation systems: variational assimilation methods and sequential methods. Variational methods are currently the most popular for use in an operational setting. 3D VAR (Lorenc, 1986) and 4D VAR (Le Dimet and Talagrand, 1986) are examples of variational methods. These are iterative methods that seek to minimize a cost function. The cost function is a measure of the difference between the

analysis state and the observations and background state. 3D VAR treats all observations as having been taken at the same time where in 4D VAR the time at which the observations were taken is accounted for.

Sequential data assimilation schemes explicitly solve a series of equations to find the analysis state of a system. Examples of sequential methods are the Kalman Filter (KF) (Kalman and Bucy, 1961) and the various filters that have been derived from the basis of the KF. For example, the ensemble Kalman filter (EnKF) (Evensen, 1994), the ensemble transform Kalman filter (ETKF) (Bishop et al., 2001) and the ensemble square root filter (EnSRF) (Whitaker and Hamill, 2002) are all sequential schemes.

The KF works in two stages. The first stage is solving forecast equations, where the background state of the system is forecast by the model to the time of an observation. The second stage is the analysis stage, the observation is assimilated into the forecast state. This is done according to a ratio of errors in the background state and in the observations. An analysis state of the system is then produced. The formulae for the KF (section 2.1) can be very computationally demanding and this filter is only valid for linear systems. Due to these problems it is not used widely used in operational NWP systems.

The ensemble filters such as the EnKF, the ETKF and the EnSRF are able to deal with non-linear systems. The formulations of the equations in these methods (sections 2.2 and 2.3) are much more computationally efficient and more suited to parallel processing than the KF, making them more viable for use within operational NWP. These methods use the ensemble spread as a statistical sample of forecast uncertainty. Each ensemble member is forecast individually by the model to the time of an observation. The analysis stage uses the ensemble mean and ensemble covariances to assimilate the observation. A new ensemble is then generated by calculating the analysis ensemble perturbations. This will be described in much greater detail in chapter 2.

1.2 Motivation for the project

Ensemble Kalman filtering is hindered by a restriction in the size of the ensemble due to the computational requirements of maintaining a large ensemble. Ensembles that are too small can lead to problems developing within a filter. The specification of the relative weighting to be placed on the background errors is always underestimated (Furrer and

Bengtsson, 2007) and can be exacerbated by using ensembles sizes that are too small. This leads to the filter placing more confidence in the background state of the system and less on the observations. This means that observational data is not influential in adjusting the forecast state adequately when producing the analysis state. If the forecast state is not a good representation of the true state of the system then filter fails to produce a meaningful state estimate. Further the use of a small ensemble can lead to the development of correlations between state components that are at a significant distance from one another where there is no physical relation.

The aim of this project is to study these problems in more detail, to look at methods currently in use to combat them and to attempt to implement these in a simple model using the ETKF.

1.3 Principal goals

To show, in ensemble filters, that

- small ensembles cause errors in the background state to be underestimated
- underestimation of errors in the background state can cause a filter to produce a meaningless estimate
- small ensembles cause spurious correlations develop
- inflating the background error estimates can help the filter produce more meaningful state estimates
- unphysical correlations can be removed
- removing unphysical correlations that develop can help improve the problems of small ensembles
- a new approximation to the removal of unphysical correlations in the ETKF is possible
- the effects of this new approximation

1.4 Outline of Report

In the following chapter the equations of the Kalman filter and the principles behind an ensemble Kalman filter will be introduced. The formulation of the Ensemble transform Kalman filter, used for the experiments of this project, will also be reviewed. Chapter 3 provides a description of the main problems in ensemble Kalman filtering due to small ensembles and two techniques commonly used to overcome them. A review of how other authors have attempted the removal of unphysical correlations is also provided. Chapter 4 gives a description of the model implemented within the ETKF on which the experiments in this project are based. The specific implementation of the ETKF that is used in the experiments is described. One possible new method of removing unphysical correlations, specifically within the ETKF, is formulated. Chapter 5 provides experimental evidence of the problems of small ensembles and the techniques employed to combat them. The results of the new formulation of chapter 4 are analysed. Finally a summary, discussion and details of further work that could be done in this area is presented are chapter 6.

Chapter 2

Kalman Filtering

This chapter consists of detailed review of the formulation of the Kalman Filter (KF). The principles of the ensemble Kalman filter (EnKF) are introduced. The ensemble transform Kalman filter (ETKF), a type of square root filter (SRF), is used in the experiments of this project and consequently will be introduced and described in detail.

2.1 The Kalman Filter

The Kalman Filter (Kalman and Bucy, 1961) is a well established method of sequential method of data assimilation introduced in the 1960's (for a detailed review see Jazwinski (1970)). There are two stages involved in solving the set of equations to reach the analysis state of the system. Firstly the forecast stage; from a given initial estimate, or background state, a forecast state estimate of the system is produced by evolving a given linear model in time. The background state has a known error covariance, it is an important feature of the Kalman filter that these errors are also evolved in time, by the linear model, to produce a forecast error covariance. This is commonly referred to as flow dependence of the background errors. The second stage is the analysis stage. A weighting is produced according to a ratio of the errors between the background state and known errors in the observations. The observations are then assimilated, according to this weighting, that is they are used to adjust the forecast estimate to more accurately reflect the true state of the system. Finally an analysis of the state of the system and an analysis error covariance is produced according to this weighting.

Given a known matrix, $\mathbf{M}_j$, a linear model of some system dynamics at a given time, j,

the true state of the system can be defined as

$$\mathbf{x}^t(t_{j+1}) = \mathbf{M}_j\mathbf{x}^t(t_j) + \eta(t_j). \tag{2.1}$$

- $\mathbf{x}$ is the state vector of the system with dimension n
- $\mathbf{x}^t$ is the true state of the system denoted by the superscript t
- $\mathbf{x}^t(t_j)$ represents the true state of the system at a given time t_j
- $\eta(t_j)$ is the random model error at time t_j

It is assumed that the random model error is unbiased, $\langle\eta(t_j)\rangle = 0$. The angle brackets here denote the expectation of the quantity. The model error covariance is given by $\mathbf{Q}(t_j) = \left\langle\eta(t_j)\eta(t_j)^T\right\rangle$. In general the model error is unknown, it is either set to zero or it can estimated, perhaps itself being modelled (Evensen, 2003). Observations at a given time, $\mathbf{y}(t_j)$ where $\mathbf{y}$ is the observation vector, of dimension m, satisfy

$$\mathbf{y}(t_j) = \mathbf{H}\mathbf{x}^t(t_j) + \epsilon(t_j). \tag{2.2}$$

The linear observation operator, denoted by $\mathbf{H}$, provides a mapping from model space to observation space. This can be done for example, by interpolating from the model grid to the location of an observation. The random observational error is given by $\epsilon(t_j)$ and it is assumed to be unbiased, $\langle\epsilon(t_j)\rangle = 0$. The observations have a known error covariance matrix, $\mathbf{R}$, that is the errors are correlated where $\left\langle\epsilon(t_j)\epsilon(t_j)^T\right\rangle = \mathbf{R}$. It is further assumed that there is no correlation between the random observational error and the random model error at any time.

Superscripts of f, a and t will be used to denote forecast, analysis and true states respectively throughout this report.

During the forecast and analysis stages of the Kalman Filter we wish to find forecasts and analyses that are unbiased, i.e., $\left\langle\mathbf{x}^f - \mathbf{x}^t\right\rangle = 0$, and $\langle\mathbf{x}^a - \mathbf{x}^t\rangle = 0$. The forecast and analysis error covariance matrices, of dimensions $n \times n$, are a measure the size of the correlation of the error between the components of the state. They are given by

$$\mathbf{P}^f = \left\langle(\mathbf{x}^f - \mathbf{x}^t)(\mathbf{x}^f - \mathbf{x}^t)^T\right\rangle, \tag{2.3}$$

and

$$\mathbf{P}^a = \left\langle (\mathbf{x}^a - \mathbf{x}^t)(\mathbf{x}^a - \mathbf{x}^t)^T \right\rangle . \tag{2.4}$$

The forecast stage

The forecast equations of the Kalman filter are given by

$$\mathbf{x}^f(t_j) = \mathbf{M}\mathbf{x}^a(t_{j-1}), \tag{2.5}$$

and

$$\mathbf{P}^f(t_j) = \mathbf{M}\mathbf{P}^a(t_{j-1})\mathbf{M}^T + \mathbf{Q}. \tag{2.6}$$

These forecast equations evolve the previous model analysis state and the previous analysis error covariance in time according to the linear model with the addition of the model error covariance $\mathbf{Q}$, as previously defined. In the case of the first forecast it is necessary to have an initial estimate of the state $\mathbf{x}^a(t_0)$ and an initial estimate of its error covariance $\mathbf{P}^a(t_0)$.

The analysis stage

The analysis update equations begin by calculating

$$\mathbf{K}(t_j) = \mathbf{P}^f(t_j)\mathbf{H}^T(\mathbf{H}\mathbf{P}^f(t_j)\mathbf{H}^T + \mathbf{R})^{-1}. \tag{2.7}$$

The Kalman gain matrix, $\mathbf{K}$, calculates a weighting to give to the observations according to a ratio between the forecast and observational error covariances. This choice of the Kalman gain matrix can be derived from a minimum variance estimate. This means that the Kalman filter provides an optimal solution such that it minimises the cost function

$$J(\mathbf{x}) = \left(\mathbf{x} - \mathbf{x}^f\right)^T \left(\mathbf{P}^f\right)^{-1} \left(\mathbf{x} - \mathbf{x}^f\right) + (\mathbf{y} - \mathbf{H}\mathbf{x})^T (\mathbf{R})^{-1} (\mathbf{y} - \mathbf{H}\mathbf{x}) . \tag{2.8}$$

Using the Kalman Gain matrix the observations at a given time, j, are assimilated to produce an analysis of the state

$$\mathbf{x}^a(t_j) = \mathbf{x}^f(t_j) + \mathbf{K}(t_j)(\mathbf{y}(t_j) - \mathbf{H}\mathbf{x}^f(t_j)). \tag{2.9}$$

The forecast error covariance is updated according to the Kalman gain to provide an analysis error covariance

$$\mathbf{P}^a(t_j) = (\mathbf{I} - \mathbf{K}(t_j)\mathbf{H})\mathbf{P}^f(t_j). \tag{2.10}$$

The Kalman filter is a useful data assimilation technique as the state update equations 2.5 and 2.9 provide unbiased forecasts and analyses. It follows therefore that the forecast and analysis error covariances are exact. The solution is optimal since the Kalman gain is a minimum variance estimate. Importantly the background error covariance is flow dependent.

The Kalman Filter is not currently used in operational NWP as it has a number of drawbacks that make it unsuitable to this environment. The Kalman Filter is only valid for linear systems where NWP models are based on non-linear dynamical models. The Kalman filter is a computationally expensive scheme to implement in NWP models. This is due to the large state spaces of operational NWP models and the consequently large covariance matrices that must be calculated and stored. In a NWP model there are of $O(10^7)$ state variables (UK Met Office, 2008). Error covariances have dimensions $n \times n$, thus the error covariances will be of $O\left(10^7 \times 10^7\right)$. In addition in any assimilation cycle $O(10^6)$ observations must also be assimilated (UK Met Office, 2008). The Kalman gain also requires the inversion of a large matrix, this is computationally expensive and inefficient. It is for these reasons that the Kalman Filter is not used in operational NWP.

2.2 The Ensemble Kalman Filter

The Ensemble Kalman Filter (EnKF) was originally introduced by Evensen (1994). A review article by Evensen (2003) covers many of the subsequent developments. The EnKF allows nonlinear models to be used in a formulation that is based on the ideas of the KF. In the KF a single state estimate is used to produce an analysis. In contrast the key idea in the EnKF is that an ensemble, which is a statistical sample of state estimates, is used to calculate an analysis state for each member. From a forecast ensemble the sample mean and error covariances are calculated. During the analysis stage these are then used to calculate a single Kalman gain that is used to assimilate the observations and produce an analysis state ensemble.

The EnKF is a desirable method for its ability to use nonlinear dynamical models; it does not require tangent linear models to be implemented (Kalnay, 2003). The representation of the state estimate statistically enhances the quality of the forecast, and may provide good initial perturbations for an ensemble based forecast scheme. The EnKF requires that an ensemble of state estimates is maintained. The computational cost of maintaining this ensemble is offset by computational savings within the method. No separate covariance matrix is required to be evolved or updated. In addition the use of a common Kalman gain for each ensemble member reduces the additional computational cost of maintaining an ensemble. The size of the ensemble can also be chosen; by using a smaller ensemble the cost of maintaining the ensemble is reduced. Care must be taken to ensure that the ensemble is not too small so that it remains statistically representative of the model (Kalnay, 2003). Since each ensemble member is evolved in time independently (section 2.2.2) this filter is well suited for parallel processing (Houtekamer and Mitchell, 1998). The EnKF importantly retains the flow dependent nature of the background error covariance matrix of the KF.

There are many difficulties associated with the EnKF. In the EnKF it is essential that the ensemble is statistically representative. Where the ensemble size is too small to be statistically representative of the model the system it is said to be undersampled. Undersampling is a fundamental problem in ensemble Kalman filtering (Anderson, 2001). Undersampling leads to further problems such as inbreeding, filter divergence and the development of long range spurious correlations. These will be discussed in more detail in chapters 3 and 5.

2.2.1 Definition of terms for the EnKF

The terms of equations 2.11 to 2.15 are defined following Evensen (2003), however a different notation has been used. Define an initial ensemble of size N, then we have the state vector for each ensemble member, i, $\mathbf{x}_i$ for $i = 1, 2, ..., N$. The state vector for each ensemble member has dimension n. The ensemble mean is defined as

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{x}_i. \tag{2.11}$$

The ensemble error covariance $\mathbf{P}_e$ is written as

$$\mathbf{P}_e = \frac{1}{N-1}\sum_{i=1}^{N}\left(\mathbf{x}_i - \bar{\mathbf{x}}\right)\left(\mathbf{x}_i - \bar{\mathbf{x}}\right)^T. \tag{2.12}$$

Here there is a division over the ensemble by *N-1* and not N this ensures that $\mathbf{P}_e$ is an unbiased estimate of the covariance $\mathbf{P}_e$ (Barlow, 1989). The ensemble error covariance will be of dimension $n \times n$. Define a state ensemble matrix as

$$\mathbf{X} = \frac{1}{\sqrt{N-1}}\left(\mathbf{x}_1 \;\; \mathbf{x}_2 \;\; ...\mathbf{x}_N\right), \tag{2.13}$$

where each column is the state estimate for an individual ensemble member. The ensemble matrix of state estimates, $\mathbf{X}$, is of dimensions $n \times N$. An ensemble perturbation matrix can then be written as

$$\mathbf{X}' = \frac{1}{\sqrt{N-1}}\left(\mathbf{x}_1 - \bar{\mathbf{x}} \;\; \mathbf{x}_2 - \bar{\mathbf{x}} \;\; ... \;\; \mathbf{x}_N - \bar{\mathbf{x}}\right). \tag{2.14}$$

This allows us to write the ensemble covariance matrix as

$$\mathbf{P}_e = \mathbf{X}'\mathbf{X}'^T. \tag{2.15}$$

The matrix $\mathbf{X}'$ is a square root of $\mathbf{P}_e$. The definition of a matrix square root of $\mathbf{A}$ given by $\mathbf{A} = \check{\mathbf{A}}\check{\mathbf{A}}^T$, such that $\check{\mathbf{A}}$ is a square root of $\mathbf{A}$, is often used in engineering and meteorology (Tippett et al., 1999). This square root is not the unique, symmetric, square root; this is now shown. If $\mathbf{Z}$ is an $n \times n$ orthogonal matrix then

$$\mathbf{A} = \check{\mathbf{A}}\check{\mathbf{A}}^T = (\check{\mathbf{A}}\mathbf{Z})(\check{\mathbf{A}}\mathbf{Z})^T \tag{2.16}$$

since

$$(\check{\mathbf{A}}\mathbf{Z})(\check{\mathbf{A}}\mathbf{Z})^T = (\check{\mathbf{A}}\mathbf{Z})(\mathbf{Z}^T\check{\mathbf{A}}^T) \tag{2.17}$$

Now if $\mathbf{Z}$ is orthogonal then $\mathbf{Z}\mathbf{Z}^T = \mathbf{I}$ and so

$$(\check{\mathbf{A}}\mathbf{Z})(\check{\mathbf{A}}\mathbf{Z})^T = \check{\mathbf{A}}\check{\mathbf{A}}^T = \mathbf{A} \tag{2.18}$$

2.2.2 The forecast stage

The forecast equation of the EnKF and its various formulations, is similar to the KF forecast equation (equation 2.5). It is expressed as

$$\mathbf{x}_i^f = \mathbf{m_j}\left(\mathbf{x}_i^a(t_{j-1})\right) + \eta_i(t_{j-1}). \tag{2.19}$$

The model $\mathbf{m}_j$ is a nonlinear model. The convention of lower case letters to denote nonlinear operators and upper case letters to denote linear operators is used throughout. The term $\eta_i(t_{j-1})$ is pseudo random model error from a distribution of mean zero and covariance $\mathbf{Q}$. An initial analysis ensemble is required $\mathbf{X}^a(t_0)$ from this an initial error covariance $\mathbf{P}_0^a$ can be determined according to equation 2.4. There is no longer a forecast of the previous error covariance as now this is will not be explicitly calculated but will be incorporated into the analysis equations according to equation 2.15. Each ensemble member is forecast independently by the full nonlinear dynamical model and this formulation is therefore well suited to parallel processing. From this ensemble forecast the ensemble mean $\overline{\mathbf{x}^f}$ and the forecast error covariance of the ensemble $\mathbf{P}_e^f$ can be calculated.

The analysis stage will only be considered at one time and so from here the time index j will be dropped.

2.2.3 General formulation of the EnKF analysis equations for a square root filter

Filters that use perturbed observation ensembles, such as the EnKF, are known as stochastic filters those that do not use perturbed observations are known as deterministic filters (Hamill, 2006). Square root filters (SRF) are deterministic filters. The general formulation of a square root filter is now presented and is based on the formulation in Tippett et al. (1999). Firstly a forecast observation ensemble is defined

$$\mathbf{y}_i^f = \mathbf{H}(\mathbf{x}_i^f). \tag{2.20}$$

Where each observational ensemble member $i = 1, ..., N$ is denoted by $\mathbf{y}_i^f$. The mean of the observation ensemble is given by

$$\overline{\mathbf{y}^f} = \overline{\mathbf{H}\mathbf{x}^f}. \tag{2.21}$$

The ensemble observation perturbations are given by

$$\begin{aligned} \mathbf{y}_i' &= \mathbf{H}(\mathbf{x}_i) - \overline{\mathbf{H}(\mathbf{x})} && (2.22)\\ &= \mathbf{H}(\mathbf{x}_i) - \mathbf{H}(\bar{\mathbf{x}}) && (2.23)\\ &= \mathbf{H}\left(\mathbf{x}_i - \bar{\mathbf{x}}\right) && (2.24) \end{aligned}$$

This could be extended to cover a situation where a nonlinear observation operator is required (Livings et al., 2008). An observation ensemble perturbation matrix $\mathbf{Y}'$, with each column of this matrix belonging to an individual ensemble member $\mathbf{y}_i'$ for $i = 1, ..., N$, is defined as

$$\mathbf{Y}' = \mathbf{H}\mathbf{X}'. \tag{2.25}$$

The analysis stage begins by calculating the ensemble Kalman gain which is analogous to the Kalman gain given the KF (equation 2.7) given by

$$\mathbf{K}_e = \mathbf{P}_e^f \mathbf{H}^T \left(\mathbf{H}\mathbf{P}_e^f \mathbf{H}^T + \mathbf{R}\right)^{-1}. \tag{2.26}$$

Using the definitions of the ensemble perturbation matrices (equations 2.15 and 2.25) we can write the ensemble Kalman gain as

$$\mathbf{K}_e = \mathbf{X}'^f (\mathbf{Y}'^f)^T \left(\mathbf{Y}'^f (\mathbf{Y}'^f)^T + \mathbf{R})^T\right)^{-1}. \tag{2.27}$$

The term $(\mathbf{Y}'^f (\mathbf{Y}'^f)^T + \mathbf{R})^T$ appears often in the analysis equations and so this is defined as

$$\mathbf{S} = \mathbf{Y}'^f (\mathbf{Y}'^f)^T + \mathbf{R}, \tag{2.28}$$

such that

$$\mathbf{K}_e = \mathbf{X}'^f (\mathbf{Y}'^f)^T (\mathbf{S})^{-1}. \tag{2.29}$$

This Kalman gain appears different to the Kalman gain of the KF. As the number of ensemble members increases towards infinity this gain matrix converges on that of

equation 2.7, in the case of a linear model. It is used in the calculation of the analysis ensemble mean given by

$$\begin{aligned} \overline{\mathbf{x}^a} &= \overline{\mathbf{x}^f} + \mathbf{K}_e \left(\mathbf{y} - \overline{\mathbf{y}^f} \right) && (2.30) \\ &= \overline{\mathbf{x}^f} + \mathbf{K}_e \left(\mathbf{y} - \overline{\mathbf{H}\mathbf{x}^f} \right) && (2.31) \end{aligned}$$

We can write an ensemble equation equivalent to equation 2.10 as

$$\begin{aligned} \mathbf{X}'^a (\mathbf{X}'^a)^T &= \mathbf{P}_e^a \\ &= (\mathbf{I} - \mathbf{K}_e \mathbf{H}) \, \mathbf{P}_e^f \\ &= \left(\mathbf{I} - \mathbf{X}'^f (\mathbf{Y}'^f)^T \mathbf{S}^{-1} \mathbf{H} \right) \mathbf{X}'^f (\mathbf{X}'^f)^T \\ &= \mathbf{X}'^f \left(\mathbf{I} - (\mathbf{Y}'^f)^T \mathbf{S}^{-1} \mathbf{Y}'^f \right) (\mathbf{X}'^f)^T. \end{aligned} \quad (2.32)$$

In practice 2.32 is not directly calculated. The ensemble perturbations are instead updated as

$$\mathbf{X}'^a = \mathbf{X}'^f \check{\mathbf{T}}. \quad (2.33)$$

Where $\check{\mathbf{T}}$ is an $N \times N$ matrix that is a square root of $\mathbf{I} - (\mathbf{Y}'^f)^T \mathbf{S}^{-1} \mathbf{Y}'^f$ such that

$$\check{\mathbf{T}} \check{\mathbf{T}}^T = \mathbf{I} - (\mathbf{Y}'^f)^T \mathbf{S}^{-1} \mathbf{Y}'^f. \quad (2.34)$$

Square root formulations of the ensemble Kalman filter are deterministic. They do not use perturbed observations which introduce additional noise into the problem, particularly when the number of ensemble members is small (Hamill, 2006). SRF's do not explicitly calculate the forecast error covariance matrix, only the ensemble state perturbations; this is of benefit as it is efficient and accurate computationally (Hamill, 2006). SRF's still encounter similar problems to those of the EnKF such as undersampling, inbreeding, filter divergence and development of long range spurious correlations. These are the subject of much discussion in chapters 3 and 5.

2.3 The Ensemble Transform Kalman Filter

The Ensemble Transform Kalman Filter (ETKF) was introduced by Bishop et al. (2001), it was modified by Wang et al. (2004) to be unbiased, and it is a type of square root filter. This section presents the ideas of the ETKF. This formulation makes use of the identity

$$\mathbf{I} - (\mathbf{Y}'^{f})^T\mathbf{S}^{-1}\mathbf{Y}'^{f} = \left(\mathbf{I} - (\mathbf{Y}'^{f})^T\mathbf{R}^{-1}\mathbf{Y}'^{f}\right)^{-1}. \tag{2.35}$$

This can be verified if we multiply $\mathbf{I} - (\mathbf{Y}'^{f})^T\mathbf{S}^{-1}\mathbf{Y}'^{f}$ by $\mathbf{I} - (\mathbf{Y}'^{f})^T\mathbf{R}^{-1}\mathbf{Y}'^{f}$ using the definition of $\mathbf{S}$ given in equation 2.28. Since $\mathbf{R}$ will ordinarily be diagonal the calculation of $\mathbf{R}^{-1}$ is simple compared to that of $\mathbf{S}^{-1}$ which will not necessarily have a simple structure. The eigenvalue decomposition

$$(\mathbf{Y}'^{f})^T\mathbf{R}^{-1}\mathbf{Y}'^{f} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T, \tag{2.36}$$

is calculated where the matrix $\mathbf{U}$ is orthogonal and $\mathbf{\Lambda}$ is diagonal. It follows that

$$\mathbf{I} - (\mathbf{Y}'^{f})^T\mathbf{S}^{-1}\mathbf{Y}'^{f} = \mathbf{U}(\mathbf{I} + \mathbf{\Lambda})^{-1}\mathbf{U}^T. \tag{2.37}$$

Hence the square root matrix, $\check{\mathbf{T}}$, of equation 2.33 that we seek is

$$\check{\mathbf{T}} = \mathbf{U}(\mathbf{I} + \mathbf{\Lambda})^{-\frac{1}{2}}\mathbf{U}^T. \tag{2.38}$$

This is the formulation of the $\check{\mathbf{T}}$ matrix that gives us the ensemble transform Kalman filter. The matrix $(\mathbf{I} + \mathbf{\Lambda})$ is diagonal making the calculation of $(\mathbf{I} + \mathbf{\Lambda})^{-\frac{1}{2}}$ straightforward. The final term $\mathbf{U}^T$ is required to ensure the filter is unbiased (Livings et al., 2008).

The ETKF has the benefits of being a deterministic filter. It is a beneficial algorithm for implementation as it is able to rapidly calculate the forecast error covariance (Bishop et al., 2001). It is computationally faster than the ensemble sqaure root filter (EnSRF) of Whitaker and Hamill (2002), (Hamill, 2006). Another benefit of the ETKF is the ability, given the location of an observation and its error statistics, to identify an optimal observation site. This is explained in detail in Bishop et al. (2001), but is not the subject of this project.

The ETKF is subject to the same problems as all ensemble Kalman filters, i.e undersampling, inbreeding, filter divergence and development of long range spurious correlations (see chapter 3). A major drawback in the ETKF is the problems encountered when covariance localization (section 3.2.2), the removal of long range spurious correlations (3.1.4), is attempted. This is related to the forecast error covariance not being computed directly, but represented in the equations in terms of the ensemble state perturbations.

The formulation of the Kalman filter and the ensemble transform Kalman filter have been explained in detail. In the following chapter problems inherent in ensemble Kalman filtering will be introduced as will methods that are available to negate them.

Chapter 3

Issues concerning small ensembles

In operational NWP the size of the ensemble used is restricted by computational expense. A restriction of the number of ensemble members can cause problems. This chapter discusses these problems and looks at some methods employed to overcome them.

3.1 Problems caused by small numbers of ensemble members

3.1.1 Undersampling

The success of the EnKF is highly dependent on the size of the ensemble employed. This issue has been investigated widely for example in Houtekamer and Mitchell (1998). The success of an ensemble filter will be dependent upon the ensemble being statistically representative. The ensemble must span the model sub-space adequately (Oke et al., 2007).

Current NWP models have state spaces of $O(10^7)$ (UK Met Office, 2008) and thus can require a large ensemble to adequately represent the statistics. This can come at a large computational cost. Typically an ensemble filter uses a smaller number of ensemble members than the size of the state. If it is so small that the ensemble is not statistically representative the system is said to be undersampled.

Undersampling introduces three major problems in ensemble filtering; these are inbreeding, filter divergence and the development of long range spurious correlations. These will be discussed in the subsections that follow.

3.1.2 Inbreeding

Inbreeding is a term used by Houtekamer and Mitchell (1998) to describe a particular phenomenon in ensemble filtering that can arise due to undersampling. In general this term is used to describe a situation where the analysis error covariances are systematically underestimated after each of the observation assimilations. The analysis error covariance $\mathbf{P}^a$ should always be less than that of the forecast error covariance (Furrer and Bengtsson, 2007), since it is defined as

$$\mathbf{P}^a(t_j) = (\mathbf{I} - \mathbf{K}(t_j)\mathbf{H})\mathbf{P}^f(t_j). \tag{3.1}$$

The forecast error covariance $\mathbf{P}^f = \left\langle (\mathbf{x}^f - \mathbf{x}^t)(\mathbf{x}^f - \mathbf{x}^t)^T \right\rangle$ is a measure of uncertainty in the forecast estimate state of the system. The calculation of the ensemble Kalman gain for a square root filter was defined in section 2.2.3 as

$$\mathbf{K}_e = \mathbf{P}_e^f \mathbf{H}^T \left(\mathbf{H}\mathbf{P}_e^f \mathbf{H}^T + \mathbf{R} \right)^{-1}. \tag{3.2}$$

The ensemble Kalman gain uses a ratio of the error covariance of the forecast background state and the error covariance of the observations to calculate how much emphasis or weight should be placed on the background state and how much weighting should be given to the observations. The forecast state estimate is adjusted by the observations in accordance with the ratio of background and observation covariance matrices in the Kalman gain. Therefore if the either the forecast background error or observational error is incorrectly specified then the adjustment of the forecast state will be incorrect.

Inbreeding in small ensembles that do not adequately span the model subspace can occur due to sampling errors in the covariance estimate (Lorenc, 2003). The smaller the ensemble is, the greater the degree of undersampling is present and the greater the chance is of underestimated forecast error covariances, or inbreeding (Ehrendorfer, 2007). In ensemble filters where each member of the ensemble is updated, at the analysis stage, by the same observations and the same Kalman gain matrix there is a tendency for the filter to underestimate the analysis covariance (Whitaker and Hamill, 2002). Ensemble

Kalman filters that use perturbed observations, such as the EnKF of Evensen (1994), have additional sampling errors in the estimation of the observation error covariances, this in turn makes it more likely that inbreeding will occur (Whitaker and Hamill, 2002). Square root filters do not use perturbed observations and so this source of inbreeding is negated. Inbreeding can also occur if the model error term $\mathbf{Q}$ is omitted as this can also cause the error covariances to be too small (Lorenc, 2003).

Inbreeding is a potential source of filter divergence (Hamill et al., 2001) and the development of spurious long range correlations. Undersampling produces a reduced rank representation of the background error covariance matrix (Hamill et al., 2001). In cases where the undersampling is severe there is a tendency for the variances and physical covariances to be underestimated.

3.1.3 Filter divergence

Filter divergence occurs when an incorrectly specified analysis state is unable to be adjusted by observation assimilations to more accurately represent the true state of the system. One cause of filter divergence is inbreeding (Hamill et al., 2001). If the covariances in the forecast estimate become too large then there is little certainty in the forecast estimate state of the system. Less weighting is given to the forecast state of the system and more to the observations. Conversely, if the covariances become too small then there is high degree of certainty in the forecast estimate state of the system. Accordingly in the Kalman gain more weighting is placed on the forecast state of the system and less on the observations. As the size of the analysis covariances decreases, i.e the greater the degree of inbreeding, the more relative weighting a filter assigns to the background state and the more the observational data is ignored during the assimilation. The standard deviation of the analysis state ensemble is smaller than that of the forecast state ensemble; the ensemble members are converging. Due to the small error covariances, or inbreeding, it becomes impossible for the filter to then adjust an incorrectly specified forecast state to accurately represent the true state and filter divergence has occurred.

3.1.4 Spurious correlations

Spurious correlations are correlations in the forecast error covariance between state components that are not physically related and they are normally at a significant distance from each other. Where all the observations have an impact on each state variable large long range spurious correlations may develop (Anderson, 2001). The consequence of these is that a state variable may be incorrectly impacted by an observation that is physically remote. As the size of the ensemble and the true correlation between state components decreases, so the error in the covariance error estimate, relative to the true correlation, greatly increases (Hamill et al., 2001). In the physical world it is expected that at a distance from a given observation point the true correlation will decrease. In a NWP the size of the error relative to the true correlation at grid points remote from the observation point in the forecast error covariance matrix will therefore be expected to increase (Hamill et al., 2001). It was demonstrated by Hamill et al. (2001) that the analysis estimate provided by the EnKF was less accurate when the error in the covariance estimate, known as noise, is greater than the true correlation, known as the the signal. Since the correlations at a distance are expected to be small and the relative error increases with distance, then it is expected that state components distant from the observations have a greater noise to signal ratio. These are long range spurious correlations and they degrade the quality of the analysis estimate. Further Hamill et al. (2001) show that the noise to signal ratio is a function of ensemble size. Larger ensembles which more accurately reflect the statistics have less associated noise. Thus the problem of spurious correlations is associated with undersampling. Lorenc (2003) shows that when an ensemble is generated by random sampling from a probability distribution function (pdf), the forecast error covariance will have an error proportional to $1/N$, where N is the size of the ensemble.

3.2 Methods to combat problems of undersampling

Various methods have been implemented in an attempt to negate these problems. Methods known as covariance inflation and covariance localization (Hamill et al., 2001) are the methods implemented in the model described in section 4 and results of these will be demonstrated in section 5.3. Consequently they are described in detail in the following subsections.

These are not the only methods that have been researched for possible implementation. Alternative methods of covariance localization include a double EnKF (Houtekamer and Mitchell, 1998) and the local ensemble Kalman filter (LEKF) (Ott et al. (2004), Szunyogh et al. (2005)). The double EnKF employs parallel ensemble data assimilation cycles. Error covariances are estimated by one ensemble and this is then used to update the other. This can help prevent the cycle of systematic inbreeding. The LEKF is a square root filter where the analysis equations are solved exactly in a local subspace spanned by all the ensemble members using all the available observations within the predefined localized space. This filter is explored in more detail in Kalnay et al. (2006).

3.2.1 Covariance inflation

Covariance inflation is one method of correcting an underestimation in the forecast error covariance matrix. Introduced by Anderson and Anderson (1999) the aim is to increase the forecast error covariances by inflating, for each ensemble member, the deviation of the background error from the ensemble mean by a percentage. Prior to a new observation being assimilated in any new cycle the background forecast deviations from the mean are increased by an inflation factor, r,

$$\mathbf{x}_i^f \leftarrow r\left(\mathbf{x}_i^f - \bar{\mathbf{x}}^f\right) + \bar{\mathbf{x}}^f, \tag{3.3}$$

where $\leftarrow$ represents the replacement of the previous value.

The inflation factor, r, is normally chosen to be slightly greater than 1.0. The specification of an optimal inflation factor may vary according to the size of the ensemble (Hamill et al., 2001). It was found in the experiments by Hamill et al. (2001) that a 1% inflation factor was nearly optimal for all numbers of ensemble members. However in the experiments of Whitaker and Hamill (2002) the optimal values of the inflation factors were 7% for the EnKF and 3% for EnSRF. The size of the inflation factor chosen will depend on various factors. It will be dependent on the dynamical model implemented. If the dynamical model has significant error growth or decay the size of the inflation factor required may be affected. The size of the inflation factor may also be dependent on the type of ensemble filter used, and further the covariance filtering length scale (see 3.2.2) chosen as in Whitaker and Hamill (2002). Thus the choice of the inflation factor for a given experiment is somewhat heuristic in nature. This method of overcoming the problem of inbreeding is commonly used, for example in Hamill et al. (2001); Whitaker

and Hamill (2002); Anderson (2001) and Oke et al. (2007). Where there are local linear physical balances in the system dynamics a large inflations may lead to this balance being disrupted (Anderson, 2001). Inflation factors do not help to correct the problem of long range spurious correlations, for this a more sophisticated approach is required.

3.2.2 Covariance localization

Covariance localization (Houtekamer and Mitchell, 2001; Hamill et al., 2001; Whitaker and Hamill, 2002) is a process of cutting off longer range correlations in the error covariances at a specified distance. It is a method of improving the estimate of the forecast error covariance. It is ordinarily achieved by applying a Schur product (Schur, 1911), also known as the Hadamard product (Horn, 1990), to the forecast error covariance matrix.

A Schur product involves an elementwise product of matrices written as $A \circ B$, where A and B have the same dimensions. If i is the row index and j is the column index the Schur product is calculated as

$$(A \circ B)_{ij} = A_{ij}B_{ij} \qquad \text{(no summation)}. \tag{3.4}$$

Properties of the Schur product

For two real matrices A and B that are of the same dimensions.

(a) If A and B are positive semidefinite, then so is $A \circ B$, Horn (1990),

(b) If B is positive definite and if A is positive semidefinite with all its main diagonal entries positive, then $A \circ B$ is positive definite, Horn (1990),

(c) $A \circ (BC) \neq (A \circ B)C$

(d) $A(B \circ C) \neq (A \circ B)C$

For the proofs of the properties (a) and (b) see Horn (1990). Points (c) and (d) will be needed for approximations in the inclusion of the Schur product in the EnKF (section 3.3) and ETKF (section 4.4). These are now proved by writing the identities in a matrix index notation, where i is the row index, j is the column index and k is an index to be

summed over.

$$\begin{aligned}
[A \circ B]_{ij} &= A_{ij}B_{ij} \\
[(A \circ B)C]_{ij} &= \sum_{k=1}^{N} A_{ik}B_{ik}C_{kj} \\
[BC]_{ij} &= \sum_{k=1}^{N} B_{ik}C_{kj} \\
[A \circ (BC)]_{ij} &= \sum_{k=1}^{N} A_{ij}B_{ik}C_{kj}
\end{aligned}$$

Thus it is concluded that $A \circ (BC) \neq (A \circ B)C$. Also

$$[A(B \circ C)]_{ij} = \sum_{k=1}^{N} A_{ik}B_{kj}C_{kj}, \tag{3.5}$$

so it is now concluded that $A(B \circ C) \neq (A \circ B)C$.

General implementation of covariance localization by a Schur product

To achieve covariance localization by Schur product a function, ρ, is normally defined to be a correlation function with local support. Local support is a term meaning that the function is only non zero in a small (local) region and is zero elsewhere. The correlation function is commonly taken to be the compactly supported 5^{th} order piecewise rational function as defined in Gaspari and Cohn (1999), such that

$$\rho = \begin{cases} -\frac{1}{4}\left(|z|/c\right)^5 + \frac{1}{2}\left(|z|/c\right)^4 + \frac{5}{8}\left(|z|/c\right)^3 - \frac{5}{3}\left(|z|/c\right)^2 + 1, & 0 \leq |z| \leq c, \\ \frac{1}{12}\left(|z|/c\right)^5 - \frac{1}{2}\left(|z|/c\right)^4 + \frac{5}{8}\left(|z|/c\right)^3 + \frac{5}{3}\left(|z|/c\right)^2 & \\ -5\left(|z|/c\right) + 4 - \frac{2}{3}\left(c/|z|\right), & c \leq |z| \leq 2c, \\ 0, & 2c \leq |z|. \end{cases} \tag{3.6}$$

Here z is the Euclidean distance between either the grid points in physical space or the distance between a grid point and the observation location; this is dependent on the implementation (see section 3.3). A length scale c is defined such that beyond this the

correlation reduces from 1.0 and at a distance of more than twice the correlation length scale the correlation reduces to zero. The length scale is generally set to be $c = \sqrt{10/3}l$ where l is any chosen cutoff length scale. The factor $\sqrt{10/3}$ is included to tune the correlation function to be optimal (Lorenc 2003) such that the final localised global average error variance is closet to that of the true probability distribution.

To achieve covariance localization a Schur product is taken between the forecast background error covariance matrix, $\mathbf{P}^f$, calculated from the ensemble, and a correlation function with local support, ρ.

In the EnKF the ensemble Kalman gain is given by

$$\mathbf{K}_e = \mathbf{P}_e^f \mathbf{H}^T \left(\mathbf{H}\mathbf{P}_e^f \mathbf{H}^T + \mathbf{R}\right)^{-1}. \tag{3.7}$$

The forecast error covariance appears twice within this equation and strictly speaking the Schur product should be taken with each of these occurrences such that we have

$$\mathbf{K}_e = \left(\rho \circ \mathbf{P}_e^f\right) \mathbf{H}^T \left[\mathbf{H}\left(\rho \circ \mathbf{P}_e^f\right) \mathbf{H}^T + \mathbf{R}\right]^{-1}. \tag{3.8}$$

Since ρ is a covariance matrix and $\mathbf{P}^f$ is a covariance matrix then it can be proved that $\rho \circ \mathbf{P}^f$ is also a covariance matrix (Horn, 1990).

This product has the effect of filtering out long range spurious correlations, or noise. This can been seen schematically.

Figure 3.1 shows representations of covariance matrices. Each pixel in the grid represents a covariance such that pixel ij is a measure of covariance between the i^{th} and j^{th} state components. The pixels are coloured according to the sizes of the covariances in the matrix. Figure 3.1(a) shows an example of a correlation matrix, ρ (equation 3.6). In addition to being defined as the Gaspari and Cohn (1999) compactly supported correlation function periodic boundary conditions have been accounted for, by defining the last 7 state components to be the same as the first. The number of state components to have covariances at the boundary was heuristically chosen. The boundary conditions were included to satisfy the conditions of the experiment in chapter 4. These can be seen in the top right and bottom left corners. The cutoff length scale was chosen, again heuristically, to be $\mathrm{l} = 5$. The forecast error covariances were produced using, $N = 8$, ensemble members for a state space of size, $n = 100$. The full details of the experimental desgin are described in chapter 4. Figure 3.1(b) is an example of a forecast error covariance at

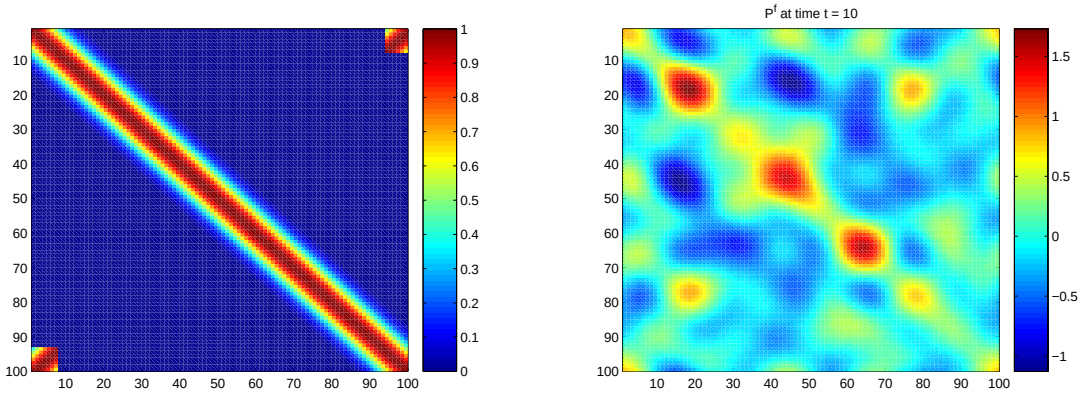

(a) Covariance correlation function ρ

(b) Sample $\mathbf{P}^f$

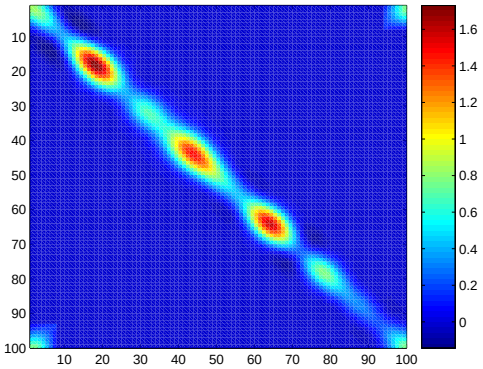

(c) $\rho \circ \mathbf{P}^f$

Figure 3.1: Illustration of a correlation function, ρ as defined in equation 3.6, a forecast error covariance, $\mathbf{P}^f$ as defined in equation 2.15, and the Schur product composition of these matrices.

time t = 10s defined in equation 2.15. Figure 3.1(c) shows a Schur product composition of this correlation covariance matrix and the forecast error covariance matrix. Clearly the correlations that existed at a distance have been removed while the correlations with local state components, and those at the boundaries, have been maintained.

The distance at which correlations in the error covariances are cutoff, i.e. reduced to zero, can be defined in two ways. Firstly as in Hamill et al. (2001) as the Euclidean distance between a grid point and the observation location and secondly as in Oke et al. (2007) where the distance is defined as the Euclidean distance between grid points in physical space. The filtering length scale is of primary importance. It is essential that while spurious correlations are removed by the correlation function that correctly specified physical correlations are not excessively damped but maintained. If the filter length

scale is too long, so as to correctly capture all the dynamical correlations, then many of the spurious correlations may not be removed; the covariance may be "noisy causing an overly adjusted variance deficient ensemble" (Ehrendorfer, 2007). If the length scale is too short then important dynamical correlations may be lost. The process of defining the correct length scale for a given system is at present a heuristic process.

The application of a Schur product to the forecast error covariance primarily achieves the removal of long range spurious correlations. An important benefit of the Schur product is that it increases the rank of the covariance matrix to which it is applied (Hamill et al., 2001; Oke et al., 2007). The rank of the forecast error covariance matrix is at most $N-1$ where N is the number of ensemble members, i.e. it is rank deficient. The increase in number and magnitude of the rank of the forecast error covariance by including a Schur product is shown in figure 3.2. The correlation matrix ρ and error covariance $\mathbf{P}^f$ used to produce figure 3.2 correspond to those of figure 3.1. In the list of properties of the Schur product, point (b) of the Schur product is very important. It implies that providing the correlation matrix is positive semidefinite with all main diagonal entries positive then the forecast error covariance, and subsequently the the analysis error covariance will become strictly positive definite.

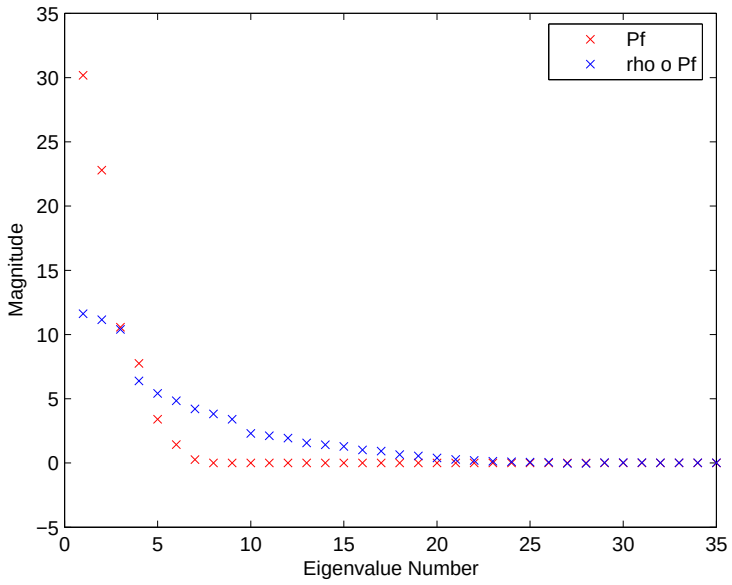

Figure 3.2: Eigenvalue spectrum before and after a Schur product is applied to a forecast error covariance matrix, $t = 10s$.

The x-axis is truncated at 35 as all values beyond this are zero and it allows the details

of the eigenvalues to be seen more clearly. Increasing the rank of the forecast error covariance means that the number of degrees of freedom of available for assimilating the observations also increases thus the analysis state should better represent the observations (Ehrendorfer, 2007). This is extremely useful when a system is undersampled as the *effective* size of the ensemble is increased (Oke et al., 2007). The use of the correlation function, where there is an introduction of many zeros into the error covariance matrices also has the advantage of computational savings in the calculations of the analysis error covariance (Lorenc, 2003).

If the covariance localization is used in the calculation of the Kalman gain in a SRF to update the ensemble deviations then the relation $\mathbf{P}^a = (\mathbf{I} - \mathbf{KHP})^f$ will only be satisfied if the correlation covariance is a Heavyside step function, i.e. it is only either zero or one (Whitaker and Hamill, 2002). The common correlation function in use, equation 3.6 of Gaspari and Cohn (1999) is not a step function and therefore this relationship is not held strictly. Information on physical dynamical balances is held within the error covariance matrices; the modification of this matrix by inclusion of a Schur product is known to disrupt these balances (Oke et al., 2007).

3.3 Implementations of the Schur product

The way in which the Schur product is implemented is far from standardized. This section looks at some methods of implementation. All select the correlation function to be the compactly supported 5^{th} order piecewise rational function as defined in Gaspari and Cohn (1999) (equation 3.6), except Oke et al. (2007) where the exact specification of the correlation function is not stated.

One approximation is not to apply the Schur product to both occurrences of the $\mathbf{P}^f$ matrix but only one of them (equation 3.8). Hamill et al. (2001) use a stochastic ensemble Kalman Filter, i.e. one which uses perturbed observations, and the Schur product is applied only on $\mathbf{P}_e^f\mathbf{H}^T$ and not to the term $\mathbf{HP}_e^f\mathbf{H}^T$. They state that there is the use of only a simple linear observation operator, $\mathbf{H}$, and that this constitutes as only a small approximation. Hamill et al. (2001) define the distance cutoff of the correlation function as the Euclidean distance between a grid point and the observation location.

Houtekamer and Mitchell (2001) also using a stochastic ensemble Kalman Filter reformulate the analysis equations to include the Schur product in their filter. The reformulation

shown here follows Houtekamer and Mitchell (2001). The Kalman Gain is reformulated from equation 3.8. An approximation is made such that $\left(\rho \circ \mathbf{P}^f\right)\mathbf{H}^T = \rho \circ \left(\mathbf{P}^f\mathbf{H}^T\right)$. Although this is not true in the general sense (see section 3.2.2) this approximation is made here since the structure of $\mathbf{H}$ is generally diagonal. "The forward interpolation of $\mathbf{H}$ involves operations on the columns of grid points and/or interpolations or finite differences on the horizontal grid, where ρ is a relatively broad function" (Houtekamer and Mitchell, 2001). This is the justification for changing the order in which the Schur product is applied and approximating the Kalman Gain matrix as

$$\mathbf{K}_e = \left[\rho \circ \left(\mathbf{P}_j^f\mathbf{H}^T\right)\right]\left[\rho \circ \left(\mathbf{H}\mathbf{P}_j^f\mathbf{H}^T\right) + \mathbf{R}\right]^{-1}. \tag{3.9}$$

The reformulated equations are

$$\left[\rho \circ \left(\mathbf{H}\mathbf{P}_i^f\mathbf{H}^T + \mathbf{R}\right)\right] d_i = \mathbf{y}_{ij} - \mathbf{H}x_i^f, \quad i = 1, ..., N. \tag{3.10}$$

and

$$\mathbf{x}_i^a = \mathbf{x}_i^f + \left[\rho \circ \left(\mathbf{P}_j^f\mathbf{H}^T\right)\right] d_i, \quad i = 1, ..., N. \tag{3.11}$$

Since we seek to find $d_i = \mathbf{y}_{ij} - \mathbf{H}x_i^f\left[\rho \circ \left(\mathbf{H}\mathbf{P}_i^f\mathbf{H}^T + \mathbf{R}\right)\right]^{-1}$. Multiplying both sides of the equation by $\left[\rho \circ \left(\mathbf{H}\mathbf{P}_i^f\mathbf{H}^T + \mathbf{R}\right)\right]$ avoids the calculation of the matrix inverse. Thus the term $\rho \circ \left(\mathbf{H}\mathbf{P}_i^f\mathbf{H}^T + \mathbf{R}\right)$, is not inverted directly but the equation (3.10) is solved for d_i. Houtekamer and Mitchell (2001) use a perturbed observation filter and so the additional index j on the observations is necessary to include the one perturbation per ensemble member. The analysis state, including the Schur product, is calculated according to equation 3.11.

Oke et al. (2007) use both an ensemble Kalman Filter and ensemble based Optimal Interpolation (EnOI) assimilation schemes. In the experiments run by Oke et al. (2007) the Kalman gain is written as

$$\mathbf{K}_e = \left[\rho \circ \left(\mathbf{P}^f\mathbf{H}^T\right)\right]\left[\left(\rho \circ \left(\mathbf{H}\mathbf{P}^f\right)\mathbf{H}^T\right) + \mathbf{R}\right]^{-1}. \tag{3.12}$$

The Schur product is applied to both occurrences of the forecast error covariance. The correlation function used by Oke et al. (2007) is defined as being dependent only on the Euclidean distance between the ith and jth state element in physical space. The forecast error covariance is reduced at distances from a given state element in accordance with a

Gaussian function with an e-folding length scale.

Whitaker and Hamill (2002) use an ensemble square root filter. In this EnSRF the ensemble mean and perturbations are calculated independently according to

$$\overline{\mathbf{x}^a} = \overline{\mathbf{x}^f} + \mathbf{K}(\overline{\mathbf{y}'} - \mathbf{H}\overline{\mathbf{x}^f}), \tag{3.13}$$

$$\mathbf{x}'^a = \mathbf{x}'^f + \tilde{\mathbf{K}}(\mathbf{y}' - \mathbf{H}\mathbf{x}'^f), \tag{3.14}$$

where $\tilde{\mathbf{K}} = \alpha\mathbf{K}$.

$$\mathbf{K} = \mathbf{P}^f\mathbf{H}^T(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}, \tag{3.15}$$

and

$$\alpha = \left(1 + \sqrt{\frac{\mathbf{R}}{\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R}}}\right)^{-1}. \tag{3.16}$$

Note that Whitaker and Hamill (2002) use a serial processing of the observations, which implies that $\mathbf{HPH}^T$ is a scalar, thus allowing the division in equation 3.16. It is not explicitly stated by Whitaker and Hamill (2002) where the covariance is applied. In this method the forecast error covariance only appears in the term for α. Whitaker and Hamill (2002) make it clear that the assumption of a correlation is not included in the derivation of 3.16. They also state that inclusion of such a correlation matrix, where it is not a step function, means that the equations of the filter will no longer be satisfied.

In the ETKF of Bishop et al. (2001) and the implementation used in this project of Livings (2005) (see section 4.3) the forecast error covariance $\mathbf{P}^f$ does not appear explicitly. Rather it is implied through the appearance of the ensemble state forecast perturbation matrix $\mathbf{X}'^f$ (equation 3.8). This is in contrast to the SRF used by Whitaker and Hamill (2002), where the forecast error covariance appears as $\mathbf{P}^f$. It is therefore more difficult to implement a Schur product covariance localization in the ETKF and it is difficult to see how this may be achieved without losing some of the numerical benefits that are due to doing the calculations only in terms of square roots.

The problems of undersampling, inbreeding, filter divergence and development of long range spurious correlations have been introduced. Methods known as covariance inflation and covariance localization were explained theoretically and the general formulation of the implementation of these were given. A review of the way in which other authors have implemented covariance localization by Schur product was presented. In the next chapter a description of the methodology of the model to be implemented in this project is given.

The specific implementation of the ETKF used is described and a new approximation of covariance localization by Schur product in the ETKF is made.

Chapter 4

Experiment Design

This chapter describes the model that was implemented within the ETKF as basis for the experiments of chapter 5. Details of a new approximation of covariance localization by Schur product for the ETKF is also introduced.

4.1 Advection Model

The model that was implemented in the ETKF, to experiment with localization techniques, is based on the one-dimensional linear advection model used by (Oke et al., 2007) solving the equation

$$\frac{\partial a}{\partial t} + u\frac{\partial a}{\partial x} = 0, \qquad 0 \leq x \leq L, \quad t \geq 0. \tag{4.1}$$

Here a is an arbitrary model variable, which could be, for example, representative of a concentration of a gas. The advection speed is given by u and the size of the domain is given by L. The solution to the advection equation satisfies periodic boundary conditions such that $a(0,t) = a(L,t)$. The general solution of this partial differential equation is determined from the initial conditions, $a(x,0) = f(x)$ and it is given by

$$a(x,t) = f(x - ut). \tag{4.2}$$

where $f(x)$ is an arbitrary function, such that it satisfies the periodic boundary conditions. These will be satisfied if $f(x)$ is a periodic function. The variable a is set as a linear combination of five sine waves. It is then used as the initial conditions to solve equation 4.1 analytically, where

$$a_i^0 = \sum_{k=0}^{5} A_k \sin\left(\frac{2\pi k}{L} i + \phi_k\right). \tag{4.3}$$

For this function, in general a_i^n, is the value of a at the ith grid point, at the nth time step. The amplitude of the wave, A_k and the phase, ϕ_k are randomly generated uniformly distributed samples with $0 < A_k < 1$ and $0 < \phi_k < 2\pi$.

4.2 Experiment Setup

In these experiments the domain size is L = 100, the advection speed is chosen to be $u = 1ms^{-1}$, which allows the function to be shifted by 1 grid point at every time step if we also choose $\delta x = 1m$, where δx is the grid spacing, and $\delta t = 1s$. The observation window is 121 seconds.

To solve the advection equation the general solution of equation 4.2 is used instead of integrating equation 4.1 directly . Then the analytic solution is advanced in time by a shifting algorithm (Oke et al., 2007).

$$a_i^{n+1} = a_{i-1}^n \qquad i = 1, ..., L, \tag{4.4}$$

Numerical rounding errors will be introduced each time the solution is shifted via this algorithm. To satisfy the boundary conditions we have $a_1^n = a_L^{n+1}$ (Oke et al., 2007).

To be able to analyse the results from the experimental runs an 'identical twin' experiment is run. Identical twin experiments calculate an exact solution and a forecast and analysis state for each time step. The exact solution is referred to as the true solution or truth function. The truth function is generated by solving the advection equation 4.1, analytically using the initial condition (equation 4.3), at each time step and for each grid point. This solution will serve as the exact, true solution, of the advection equation; though it will be subject to small numerical rounding errors these will not be considered. The forecast stage of the filter takes each ensemble member and applies the

shifting algorithm (equation 4.4) to advect the solution in time and space. The analysis state produced by the filter can then be compared to the true solution in time and space to assess how well the filter is performing under different conditions.

To ensure that the model that was implemented worked correctly within the ETKF the true solution at the initial time was processed using the shifting algorithm (equation 4.4) of section 4.1. The shifting algorithm advects the individual ensemble members in space at subsequent time step and it is essential to ensure that the model has been implemented correctly before implementing any experimental tests. Figure 4.1 shows the true solution being advected by the shifting algorithm.

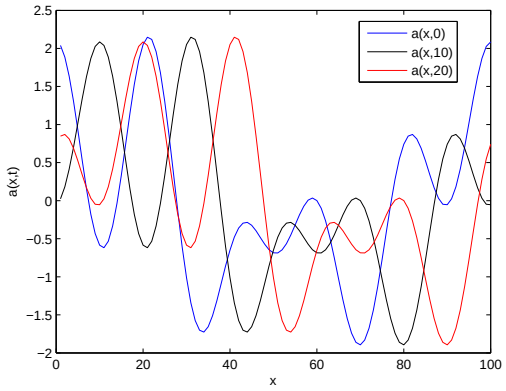

Figure 4.1: True solution advecting along the domain from 0 to 20 seconds. The blue line shows the initial truth function at 0 seconds, the black line shows the truth function at 10 seconds and the red line shows the truth function at 20 seconds.

The true solution can be clearly seen advecting across the domain, with no change in amplitude or phase and it is concluded that the algorithm is performing correctly.

The solution of the truth function is recorded for each grid point at every time. Observations are then made of these true solutions to be used in the assimilations in the filter. Observations in space are taken every five grid points across the whole domain. Observations in time are taken every ten time steps, starting at t = 10 s, no random observational error is included, i.e. in equation 2.2, $\epsilon = 0$. However there is an observational covariance matrix, $\mathbf{R}$, as described in the formulation of the equations of the ETKF (section 2.3).

It is assumed that the model is perfect, i.e. there is no model error term, $\eta = 0$ in equation 2.19. The model is one of simple linear advection; it is not therefore expected that there would be any growth in error of the forecast analysis between the observa-

tion assimilations. Error growth is associated with non linearity and chaotic dynamics (Ehrendorfer, 2007), and this model has neither. Consequently there would not be any expected amplification of the analysis ensemble spread between observations due to evolution of the model alone.

An initial ensemble of eight members is generated. Each ensemble member is based on the linear combination of sine waves as defined in equation 4.3, however, each individual member of the ensemble is randomly allocated a different amplitude and phase. Each ensemble perturbation will have a variance that is slightly greater or less than one. The initial ensemble distribution is shown in figure 4.2.

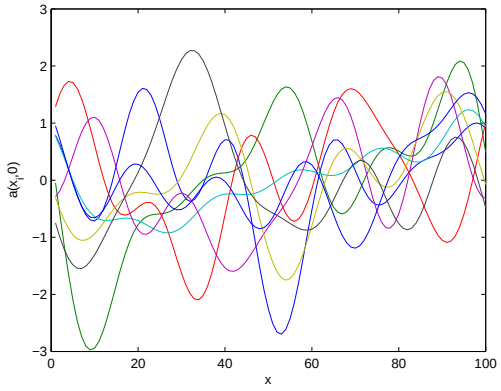

Figure 4.2: Solution over the whole domain of all the initial ensemble members. Each colour represents a different ensemble member.

Consistency between the experiments is essential if the results are to be accurately compared. The amplitude and phase of the initial conditions to generate the true solution and the amplitude and phase of each ensemble member is generated by a random number generator. To ensure consistency it is necessary to ensure that the same random phase and amplitude, in the initial conditions, used for generating the true solution are allocated in every model run. Similarly in every model run each individual ensemble member has the same random amplitude and phase allocated. This is achieved by ensuring that in every model run the random number generator always generated the same sequence of numbers.

After these initial parameters are specified the ETKF is then run, this involves the calculation of the forecast and analysis states. These are calculated as described by the method in section 4.3.

Various post processing procedures are carried out. The forecast estimate is merged with analysis and the forecast and analysis error covariances are calculated according to their definitions $\mathbf{P} = \mathbf{X}'\mathbf{X}'^T$. Analysis tools such as the mean, standard deviation and root mean square error (RMSE) are also calculated.

The further data and the analysis tools generated will be used in chapter 5.

4.3 Implementation of the ETKF

The implementation of the ETKF that will be used throughout this project is as in Livings (2005). This code was developed and validated by Livings (2005) for a different model. The following formulation of the ETKF differs in parts to that in section 2.3, however they are equivalent in analysis. This implementation is the more efficient and accurate formulation of the ETKF as it reduces the number of multiplicative operations therefore reducing rounding errors. The following description of the implementation follows that of Livings (2005). A scaled forecast observation ensemble perturbation matrix is introduced.

$$\hat{\mathbf{Y}}^f = \mathbf{R}^{-\frac{1}{2}}\hat{\mathbf{Y}}'^f, \tag{4.5}$$

and it follows from this that

$$(\mathbf{Y}'^f)^T\mathbf{R}^{-1}\mathbf{Y}'^f = (\hat{\mathbf{Y}}^f)^T\hat{\mathbf{Y}}^f. \tag{4.6}$$

The effect of scaling the observations by the square root of the observation error covariance $\mathbf{R}^{-\frac{1}{2}}$ is to normalise the observations of physical quantities with differing standard deviations in error. This has the effect of reducing numerical rounding errors. (Livings, 2005)

The singular value decomposition (SVD) (Golub and Van Loan, 1996) of the scaled forecast observation perturbation matrix $(\mathbf{Y}'f)^T$ is found according to

$$(\mathbf{Y}'^f)^T = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T. \tag{4.7}$$

The matrices $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices. $\mathbf{U}$ has dimensions $N \times N$, $\mathbf{\Sigma}$ has dimensions $N \times m$ and $\mathbf{V}$ has dimensions $m \times m$. The eigenvalues of $(\mathbf{Y}'^f)^T$ are given by $\mathbf{\Lambda} = \mathbf{\Sigma}\mathbf{\Sigma}^T$. This SVD is computed so that the calculation of $(\hat{\mathbf{Y}}^f)^T\hat{\mathbf{Y}}^f$ is avoided and

the matrices used in the analysis equation are simpler and so more suited to efficient calculation.

The ensemble perturbation update is given by

$$\mathbf{X}'^{a} = \mathbf{X}'^{f}\mathbf{T} = \mathbf{X}'^{f}\mathbf{U}(\mathbf{I}+\mathbf{\Lambda})^{-\frac{1}{2}}\mathbf{U}^{T}. \tag{4.8}$$

This differs to the implementation given in Livings (2005); the term $\mathbf{U}^T$ has been included to obtain an unbiased version as in Livings et al. (2008).

The ensemble Kalman Gain can now be written as

$$\begin{aligned}\mathbf{K}_e &= \mathbf{X}'^{f}(\mathbf{Y}'^{f})^{T}\left(\mathbf{Y}'^{f}(\mathbf{Y}'^{f})^{T}+\mathbf{R}\right)^{-1} \\ &= \mathbf{X}'^{f}(\hat{\mathbf{Y}}'^{f})^{T}\left(\hat{\mathbf{Y}}'^{f}(\hat{\mathbf{Y}}'^{f})^{T}+\mathbf{I}\right)^{-\frac{1}{2}} \\ &= \mathbf{X}'^{f}\mathbf{U}\mathbf{\Sigma}\left(\mathbf{\Sigma}^{\mathbf{T}}\mathbf{\Sigma}+\mathbf{I}\right)^{-1}\mathbf{V}^{T}\mathbf{R}^{-\frac{1}{2}}.\end{aligned} \tag{4.9}$$

For computational efficiency the product

$$\mathbf{z} = \mathbf{\Sigma}(\mathbf{\Sigma}^{T}\mathbf{\Sigma}+\mathbf{I})^{-1}\mathbf{V}^{T}\mathbf{R}^{-\frac{1}{2}}(\mathbf{y}-\overline{\mathbf{y}^{f}}) \tag{4.10}$$

is built from right to left, this avoids storing large matrix and involves the calculation of matrix-vector products rather than matrix-matrix products which are obviously more computationally efficient calculations. Then the ensemble mean is updated using

$$\overline{\mathbf{x}^{a}} = \overline{\mathbf{x}^{f}} + \mathbf{X}'^{f}\mathbf{U}\mathbf{z}. \tag{4.11}$$

4.4 Localization within the ETKF

The inclusion of a Schur product within an ETKF is highly problematic. This section investigates these problems and a possible new way of approximating the inclusion of a Schur product within the ETKF.

4.4.1 Problem of localization in the ETKF

The Schur product as a method of covariance localization has been implemented by various authors (see section 3.3). The Schur product with the forecast error covariance is traditionally applied in the Kalman gain previously written as

$$\mathbf{K}_e = \left(\rho \circ \mathbf{P}_e^f\right) \mathbf{H}^T \left[\mathbf{H}\left(\rho \circ \mathbf{P}_e^f\right) \mathbf{H}^T + \mathbf{R}\right]^{-1}. \tag{4.12}$$

The nature of the implementation of the ETKF is to split the forecast error covariances using $\mathbf{P}_e = \mathbf{X}'\mathbf{X}'^T$ as in section 2.2. In the implementation used within these experiments (Livings, 2005) the analysis update equation is

$$\overline{\mathbf{x}^a} = \overline{\mathbf{x}^f} + \mathbf{X}'^f \mathbf{U} \Sigma \left(\mathbf{\Sigma}^T \mathbf{\Sigma} + I\right)^{-1} \mathbf{V}^T \mathbf{R}^{\frac{1}{2}} \left(\mathbf{y} - \overline{\mathbf{y}^f}\right). \tag{4.13}$$

Where $\mathbf{U}$ is an orthogonal matrix of eigenvalues from the singular value decomposition of the scaled observation forecast perturbations (see section 4.3).

Now in this expression we have no explicit reference to the forecast error covariance matrix $\mathbf{P}^f$ although this is implied through $\mathbf{X}'^f$. This causes problems when we wish to introduce a Schur product which acts on the forecast error covariances and not on the perturbation matrix. The correlation function is defined in physical space and has the same dimensions as the error covariances, $n \times n$, where n is the size of the state. The dimensions of the perturbation matrix, $\mathbf{X}'$, are $n \times N$ where N is the size of the ensemble. The application of the Schur product requires that the matrices are of the same dimensions. Lorenc (2003) and Hamill (2006) state that it is not possible to include a Schur product in the ETKF as currently defined.

It is not possible to apply the correlation function with local support to the full forecast error covariance in the ETKF as it is currently defined. The forecast error covariance could be reconstructed according to $\mathbf{P}^f = \mathbf{X}'\mathbf{X}'^T$. This could be done just prior to an observation assimilation then the operation $\rho \circ \mathbf{P}^f$ could be applied in full without problem, although this would mean that the method would not strictly be a square root filter. However, the remainder of the implemented algorithm of section 4.3 requires the use of the perturbation matrix $\mathbf{X}'$ and not $\mathbf{P}^f$. To maintain the ETKF equations as they stand the perturbation matrix would need to be calculated from the forecast error covariance, i.e. the process $\mathbf{P}^f = \mathbf{X}'\mathbf{X}'^T$ would have to be reversed. It is possible to find the square root of $\mathbf{P}^f$ as it is a square matrix, $(n \times n)$, however the dimensions of

$\mathbf{X}'$ are $n \times N$. It is not clear how it would be possible to deconstruct the forecast error covariance back into a matrix of perturbations. Therefore it is not possible to implement the localization in this way.

4.4.2 New approximation of localization in the ETKF

Any square matrix that is positive semi-definite can be written as $A = \check{A}\check{A}^T$ (section 2.2.1) where $\check{A}$ is a square root of the matrix A. Applying this to the correlation matrix we can write $\rho = \check{\rho}\check{\rho}^T$. This is not the unique symmetric square root of ρ. The way in which the square root is defined will have an impact on the implementation. If the unique symmetric square root is found the matrix becomes sparse so the non unique, non symmetric definition of the square root was preferred. The square root of the matrix ρ is calculated by a singular value decomposition (SVD) (Golub and Van Loan, 1996), such that

$$\begin{aligned} \rho &= \mathbf{W}\mathbf{\Psi}\mathbf{Y} \\ \check{\rho} &= \mathbf{W}\mathbf{\Psi}^{-\frac{1}{2}} \end{aligned}$$

The matrices $\mathbf{W}$ and $\mathbf{Y}$ are orthogonal matrices, $\mathbf{\Psi}$ is a diagonal matrix and all are of dimension $n \times n$. The application of the Schur product between the correlation function with local support, ρ, and the forecast error covariance $\mathbf{P}^f$ can be written

$$\rho \circ \mathbf{P}^f = \left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right). \tag{4.14}$$

Initially it was investigated if this could be written as

$$\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right) = \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T, \tag{4.15}$$

however this is not the case and this can be proved by using the summation notation for matrices as previously used in section 3.2.2. If

$$\left[\check{\rho}\check{\rho}^T\right]_{ij} = \sum_{k=1}^{N} \check{\rho}_{ik}\check{\rho}_{jk}, \tag{4.16}$$

and

$$\left[\mathbf{X}'\mathbf{X}'^{T}\right]_{ij} = \sum_{k=1}^{N} \mathbf{X}'_{ik}\mathbf{X}'_{jk}, \tag{4.17}$$

then

$$\left[\left(\check{\rho}\check{\rho}^{T} \circ \mathbf{X}'\mathbf{X}'^{T}\right)\right]_{ij} = \sum_{k=1}^{N} \check{\rho}_{ik}\check{\rho}_{jk} \sum_{m=1}^{N} \mathbf{X}'_{im}\mathbf{X}'_{jm}. \tag{4.18}$$

Now,

$$\left[\check{\rho} \circ \mathbf{X}'\right]_{ij} = \check{\rho}_{ij}\mathbf{X}_{ij} \tag{4.19}$$

and so

$$\left[\left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^{T}\right]_{ij} = \sum_{k=1}^{N} \check{\rho}_{ik}\mathbf{X}'_{ik}\check{\rho}_{jk}\mathbf{X}'_{jk}. \tag{4.20}$$

So it can be concluded that $\left(\check{\rho}\rho^{T}\right) \circ \left(\mathbf{X}'(\mathbf{X}')^{T}\right) \neq \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^{T}$.

Horn (1990) defines a factorization of the Schur product as

$$A \circ B = \left[AA^{T} \circ BB^{T}\right]^{\frac{1}{2}} C, \tag{4.21}$$

where C is a contraction matrix with a norm of less than or equal to one. This also implies that the factorization attempted in equation 4.15 is not possible in the way hoped.

To attempt to include the Schur product in the ETKF the relation of equation 4.15 was introduced in to the derivation of the equations of the ETKF (section 2.3) and the Livings (2005) implementation. It was used as an approximation rather than an exact equivalence as this is known to be untrue in general. The derivation of the equations of section 2.3 and the Livings (2005) implementation will now be followed but with the inclusion of a Schur product. Starting from the definition of the ensemble Kalman gain

$$\begin{aligned}
\mathbf{K}_e &= \mathbf{P}_e^f\mathbf{H}^T\left(\mathbf{H}\mathbf{P}_e^f\mathbf{H}^T + \mathbf{R}\right)^{-1} \\
&= \left(\rho \circ \mathbf{P}_e^f\right)\mathbf{H}^T\left(\mathbf{H}\mathbf{P}_e^f\mathbf{H}^T + \mathbf{R}\right)^{-1} \\
&= \left(\check{\rho}\check{\rho}^{T} \circ \mathbf{X}'\mathbf{X}'^{T}\right)\mathbf{H}^T\left(\mathbf{H}\mathbf{P}_e^f\mathbf{H}^T + \mathbf{R}\right)^{-1}
\end{aligned}$$

Now the approximation $\left(\check{\rho}\check{\rho}^{T}\right) \circ \left(\mathbf{X}'(\mathbf{X}')^{T}\right) \approx \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^{T}$ is made. The correlation covariance matrix $\check{\rho}$ is of dimensions $n \times n$ and the matrix $\mathbf{X}'$ is of dimensions $n \times N$ and therefore the Schur product as stated in the approximation is not possible in the

current formation. To overcome this the perturbation matrix $\mathbf{X}'$ is extended to correct the dimensions; $(n-N)$ columns of zeros are added to this matrix so that the dimensions of this extended perturbation matrix becomes $n \times n$. This new matrix is denoted by $\tilde{\mathbf{X}}' = (\mathbf{X} \quad 0)$. The forecast error covariance is preserved such that $\mathbf{P}^f = \tilde{\mathbf{X}}'\tilde{\mathbf{X}}'^T$. The choice of adding the zero columns at the end was arbitrary. They could have been distributed throughout the columns of $\mathbf{X}'$ and the forecast error covariance would still have been preserved. The effect of adding this column of zeros is simply to add several copies of the ensemble mean to the ensemble. Now it is possible to partition the correlation matrix $\check{\rho}$ into two sets of columns. The first N columns become $\check{\rho}_1$ such that it is of dimensions $n \times N$. The second set of columns is set to zero, of dimensions $n \times n - N$, so that $\check{\rho}$ is partitioned $(\check{\rho}_1 \quad 0)$. The Schur product of the square root of the correlation function and the ensemble perturbation matrix is now given by

$$\check{\rho} \circ \mathbf{X}' = (\check{\rho}_1 \circ \mathbf{X}' \quad 0) \tag{4.22}$$

The columns of zeros are then dropped leaving a matrix of the correct dimensions $n \times N$ This approximation may destroy some of the details of the perturbation matrix. This will be explored in chapter 5.

The Kalman gain is now approximated as

$$\mathbf{K}_e = (\check{\rho} \circ \mathbf{X}') (\check{\rho} \circ \mathbf{X}')^T \mathbf{H}^T \left(\mathbf{H}\mathbf{P}_e^f\mathbf{H}^T + \mathbf{R}\right)^{-1} \tag{4.23}$$

Equation 4.23 has required the use of the identity 4.15, which is implemented as an approximation as it has been proved not to be an exact equivalence. A further approximation is now required that $(\check{\rho} \circ \mathbf{X}')^T \mathbf{H}^T = \mathbf{Y}'$. Following the general formulation of a SRF (Tippett et al., 1999) the Kalman gain is written

$$\mathbf{K}_e = (\check{\rho} \circ \mathbf{X}') \mathbf{Y}'^T \left(\mathbf{Y}'\mathbf{Y}'^T + \mathbf{R}\right)^{-1}. \tag{4.24}$$

Substituting $\mathbf{Y}'\mathbf{Y}'^T + \mathbf{R} = \mathbf{S}$ as in section 2.2.3 yields

$$\mathbf{K}_e = (\check{\rho} \circ \mathbf{X}') \mathbf{Y}'^T (\mathbf{S})^{-1}. \tag{4.25}$$

Now by writing the definition of the analysis error covariance as $\mathbf{P}_e^a = \mathbf{X}'^a (\mathbf{X}'^a)^T$ and since $\mathbf{P}_e^a = (\mathbf{I} - \mathbf{K}_e\mathbf{H}) \mathbf{P}_e^f$ an expression for the analysis perturbation can be found as

follows; firstly substituting in the redefined Kalman gain yields

$$\begin{aligned}\mathbf{X}'^a(\mathbf{X}'^a)^T &= \left(\mathbf{I} - \left(\check{\rho} \circ \mathbf{X}'^f\right)(\mathbf{Y}'^f)^T\mathbf{S}^{-1}\mathbf{H}\right)\mathbf{X}'^f(\mathbf{X}'^f)^T, \\ &\approx \left(\check{\rho} \circ \mathbf{X}'^f\right)\left(\mathbf{I} - (\mathbf{Y}'^f)^T\mathbf{S}^{-1}\mathbf{Y}'^f\right)(\mathbf{X}'^f)^T. \end{aligned} \tag{4.26}$$

This requires yet another approximation such that we consider $\left(\check{\rho} \circ \mathbf{X}'^f\right)$ equivalent to $\mathbf{X}'^f$ to allow this to be taken as a common factor, finally this yields

$$\mathbf{X}'^a = \left(\check{\rho} \circ \mathbf{X}'^f\right)\check{\mathbf{T}}. \tag{4.27}$$

This is similar in form to the equation of section 2.2.3, and the matrix $\check{\mathbf{T}}$ is identical, thus the equations of section 2.3 are followed exactly, to define the square root matrix $\check{\mathbf{T}}$ to be used in the implementation. Clearly the inclusion of the Schur product requires many approximations to be consistent with the equations of the ETKF.

For the Livings (2005) implementation the Kalman gain (equation 4.9) becomes

$$\mathbf{K}_e = \left(\check{\rho} \circ \mathbf{X}'^f\right)\mathbf{U}\left(\mathbf{\Sigma}\left(\mathbf{\Sigma}\mathbf{\Sigma}^T + \mathbf{I}\right)^{-1}\mathbf{V}^T\mathbf{R}^{-\frac{1}{2}}\right). \tag{4.28}$$

The analysis equation (4.13) is then given by

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^f + \left(\check{\rho} \circ \mathbf{X}'^f\right)\mathbf{U}\Sigma\left(\mathbf{\Sigma}^T\mathbf{\Sigma} + I\right)^{-1}\mathbf{V}^T\mathbf{R}^{\frac{1}{2}}\left(\mathbf{y} - \overline{\mathbf{y}}^f\right). \tag{4.29}$$

In effect this implies that the localization is only applied to the term $\mathbf{P}^f\mathbf{H}^T$. Hamill et al. (2001) also makes the approximation of only applying the localization to $\mathbf{P}^f\mathbf{H}^T$.

This implementation of the Schur product in the ETKF is merely an approximation and it certainly does not satisfy the equations of the ETKF. How good these approximations are will be explored in chapter 5.

In this chapter a detailed description of the model, the formulation of the particular implementation of the ETKF used, the problems inherent in implementing covariance localization by Schur product in the ETKF and an approximate implementation of Schur product for the ETKF were discussed. The next chapter looks at the general behaviour of the ETKF and demonstrates the problems of ensemble Kalman filtering. Analysis of the effect of including covariance inflation and the approximated covariance local-

ization by Schur product are made. The assumptions made in the derivation of this new approximation will also be examined. How good the approximation is can then be determined.

Chapter 5

Results

In this chapter results from model experiments without localization are presented to highlight the problems of undersampling, inbreeding, filter divergence and the development of spurious correlations. Results from model experiments incorporating localization techniques will also be presented and compared.

5.1 Performance of the filter without localization

The model was run to determine the performance of the ETKF without any localization techniques applied. This is a model of linear advection (sections 4.1 and 4.2), no error growth between observations is expected as error growth is associated with non-linearity and chaotic dynamics (Hamill et al., 2001; Ehrendorfer, 2007). Initially the experiments were run with an ensemble of eight members, generated as described in section 4.2. Graphs of the results were plotted to make analyses of the performance. All plots used in this chapter are comparable at a given time. All parameters, for example the numbers of observations in space and time, other than those specifically mentioned have been held constant.

5.1.1 General behaviour of the filter

The first plots that will be examined are those showing the evolution in time of the solution for individual components of the state vector.

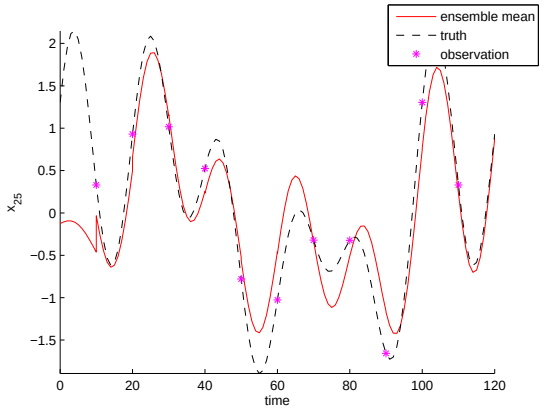

Figure 5.1: Evolution in time of the 25th state component, an observed component

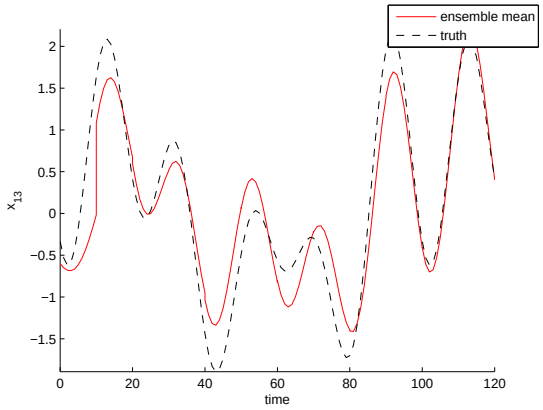

Figure 5.2: Evolution in time of the 13th state component, an unobserved component.

Figures 5.1 and 5.2 show state components, x_{25} and x_{13}, evolving in time over the whole assimilation window. These were chosen as x_{25} is an observed component and x_{13} is an unobserved component (observations are of every fifth component, section 4.2). Shown in these figures is the ensemble mean in red and the true solution in black. The observations are shown as pink stars, with an observation variance of 1.0. In both these figures we can see a smooth evolution over the first 10 seconds up to the time of the first observation, at 10 seconds. In both we can a distinct jump in the solution which moves the ensemble mean from its forecast trajectory to an analysis position closer to the observation. At the time of the second observation, 20 seconds, there is a significantly smaller alteration of the forecast trajectory towards the observation at the analysis solution. Although small there is still a small effect of assimilating this second observation. At the third observation, 30 seconds, the forecast for the 25th grid point is very close to the observation so we would

not expect much effect from the assimilation of this observation. As the model evolves in time it is clear to see that the observations have a diminishing effect on the adjustment of the forecast state estimate. By looking at the 10th state component, but only over the first half of the assimilation window, the diminishing effect of the observations can be seen (figure 5.3(a)).

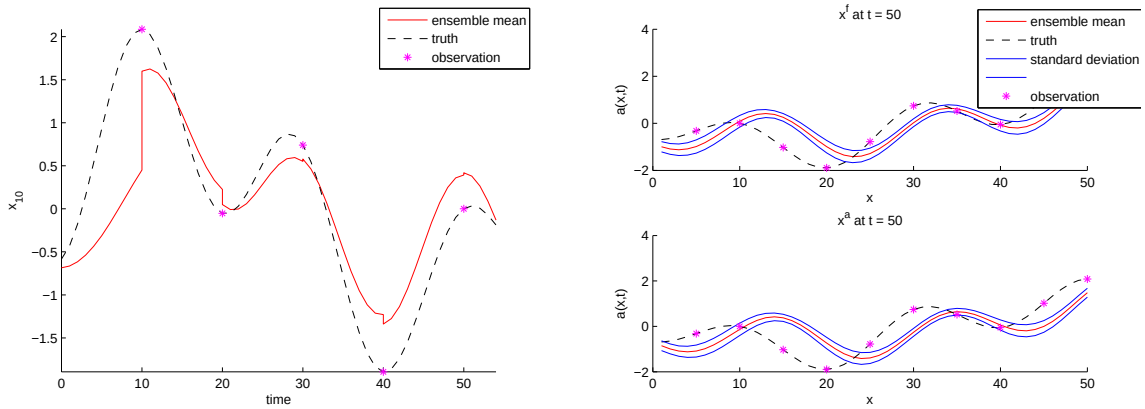

(a) Evolution in time for the first 60 seconds of the 10th state component

(b) Solution over the whole domain just before the fifth assimilation

Figure 5.3: Comparison of the fifth assimilation.

Again the large jump in the trajectory at the time of the first assimilation is clear in figure 5.3(a). At the second assimilation there is a clear adjustment of the forecast estimate, but at the third assimilation there is a significantly reduced effect of the observation. At the fourth assimilation, 40 seconds, the observation is at a great distance from the forecast solution if the assimilation was working well it would be expected that the forecast solution would be more significantly influenced by the observation at this time. At the fifth assimilation again the observation is not close to the forecast estimate, in this case the forecast has been adjusted away from the solution. This could be due to the multivariate nature of the analysis, i.e. a nearby observation is affecting this component or it could be due to the reduced number of degrees of freedom. Figure 5.3(b) is a plot of the forecast and anlysis means over the whole domain at the 5^{th} observation assimilation. Also shown is ensemble standard deviation, the true solution and the observations. By close examination of figure 5.3(b) at the 10^{th} state component it can be seen that the ensemble mean lies above the observation before the assimilation and below it after the assimilation. The solution **may** have been affected by the assimilation of the observation at the 15^{th} state component, which it has moved towards.

Now we can look at the behaviour of the solution over the whole domain at different

times. Figure 5.4 shows the ensemble mean, in red, and the ensemble mean plus and minus the standard deviation of the ensemble, in blue. These are plotted over the whole domain at the initial time $t = 0s$. Plotted alongside these is the true solution at this time, the black dashed line. The initial ensemble mean does not accurately represent the true solution. When the true solution lies between the upper and lower bounds of the standard deviation of the ensemble it will be referred to as *encapsulating* the true solution. At time $t = 0s$ the ensemble standard deviation is wide. For some state components, even including the upper and lower bounds of the standard deviation, the ensemble estimates far under or overestimate the true solution. If an event is taken from a Gaussian distribution then it is expect that 68.27 % of the samples will lie within one standard deviation of the mean (Barlow, 1989). Therefore it is not expected that the initial sample should represent the true solution at the initial time across the whole domain.

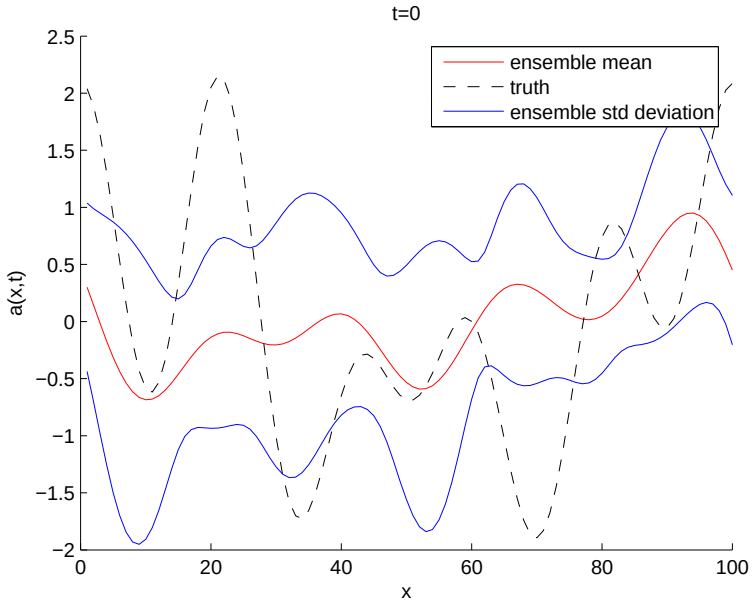

Figure 5.4: Initial solution over the whole domain.

Figure 5.5 shows a similar type of plot of the ensemble mean and standard deviations. This figure shows the forecast solution just before the assimilation in the top plot and the analysis solution just after the assimilation in the lower plot.

Before the assimilation the ensemble mean does not accurately represent the true solution and the ensemble standard deviation over the domain encapsulates a good proportion of the true solution. After the first assimilation the ensemble mean has been adjusted

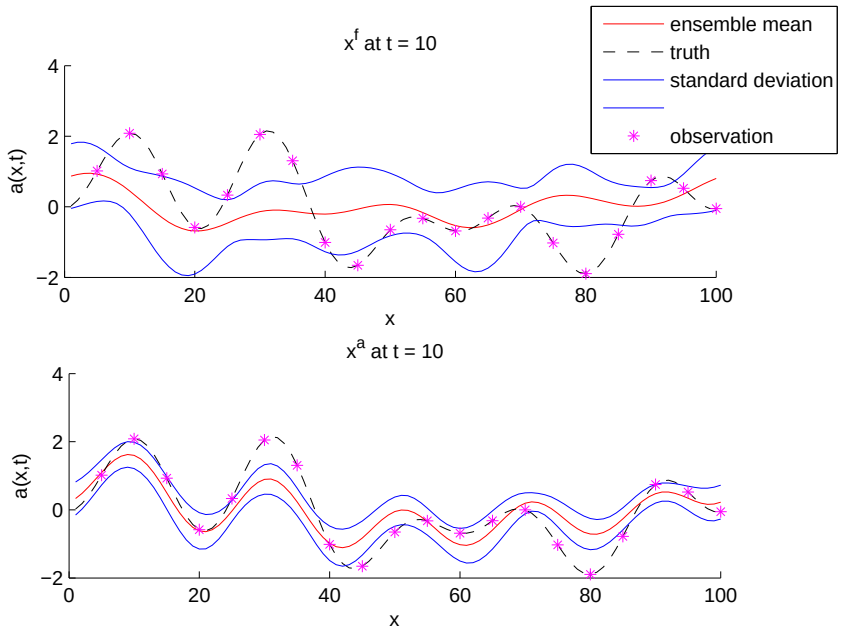

Figure 5.5: Forecast and analysis estimates, over the whole domain, for the first assimilaton.

significantly and is now a much more accurate representation of the truth function. Importantly it can be seen that the standard deviation of the ensemble members has become significantly smaller than before the assimilation. This general behaviour of the filter is consistent with that expected in an ensemble Kalman filter, (Kalnay, 2003; Evensen, 2003; Hamill et al., 2001)

5.1.2 Inbreeding

Inbreeding (section 3.1.2) is when the analysis error covariances are systematically underestimated after each time of an observation assimilation. The structure of the $\mathbf{P}^f$ matrix can be explored by plotting a representation of this matrix. A $n \times n$ grid is constructed according to equation 2.15; each pixel in the grid represents an element of the covariance matrix such that pixel ij is a covariance between the i^{th} and j^{th} state component. Each pixel is coloured according to a scale to give a representation of the size of the correlation in the covariance matrix. The covariance matrices have been plotted such that the orientation of the axis reflects conventional matrix representations.

Figure 5.6(a) shows a representation of the P^f just before the first assimilation and

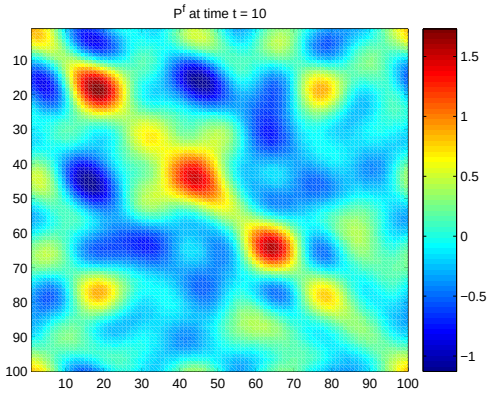

(a) P^f just before the first assimilation t = 10s

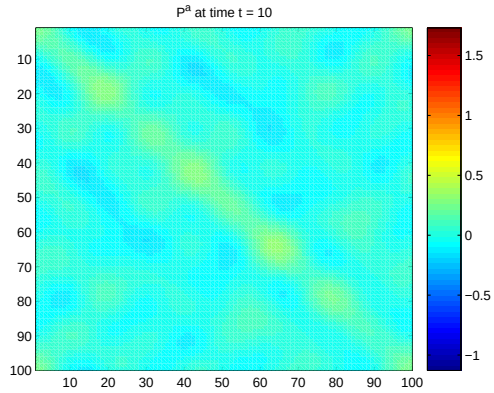

(b) P^a just after the first assimilation

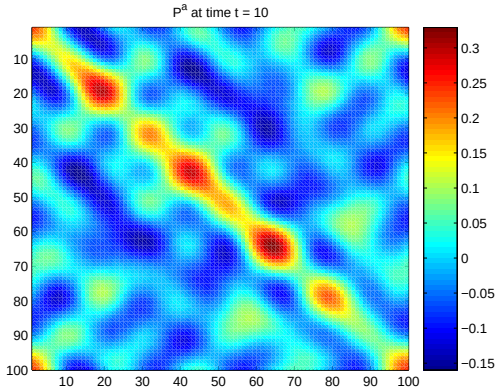

(c) $\mathbf{P}^a$ at the first assimilation with a reduced colour scale

Figure 5.6: Error covariances over the first assimilation t = 10s

figure 5.6(b) is a representation of the P^a matrix just after the first assimilation. In figure 5.6(c) the colour scale has been reduced to see more clearly the structure of the $\mathbf{P}^a$ matrix after the first assimilation.

There is a clear leading diagonal structure before the assimilation, where values on the leading diagonal are variances. There is a diagonal band around the leading diagonal that show the correlations between a point and its neighbouring points. This diagonal structure remains after the assimilation process although the sizes have been significantly decreased. There are also correlations in the bottom left and top right corners, these are due to the periodic boundary conditions. The final grid point in the domain is equal to the first; so these correlations are physical and are present both before and after the assimilation. Due to the periodic nature of the sine wave combination some correlations at a distance from a given state component are expected. The size of the covariances

just before the first assimilation are between -1 and 1.5 where just after the assimilation they have decreased significantly. It can be seen in figure 5.6(c) that the sizes of the covariances after the assimilation have decreased now ranging from -0.04 to 0.1. The size of the forecast error covariances have decreased by an order of magnitude after the first assimilation process. It would be expected that the error covariances decrease after the assimilation process since "$\mathbf{P}^a$ is negatively biased" (Furrer and Bengtsson, 2007). The problem of inbreeding is expected and has been demonstrated. This behaviour of a reduction in the size of the forecast error covariances at the assimilation, when the forecast error covariance becomes the analysis error covariance, is repeated at each assimilation although the reduction in the size is not as severe as at the first assimilation.

5.1.3 Filter Divergence

Inbreeding is a source of filter divergence (Hamill et al., 2001). As the size of the forecast error covariance decreases there will be an associated decrease in the size of the adjustment of the forecast state estimate.

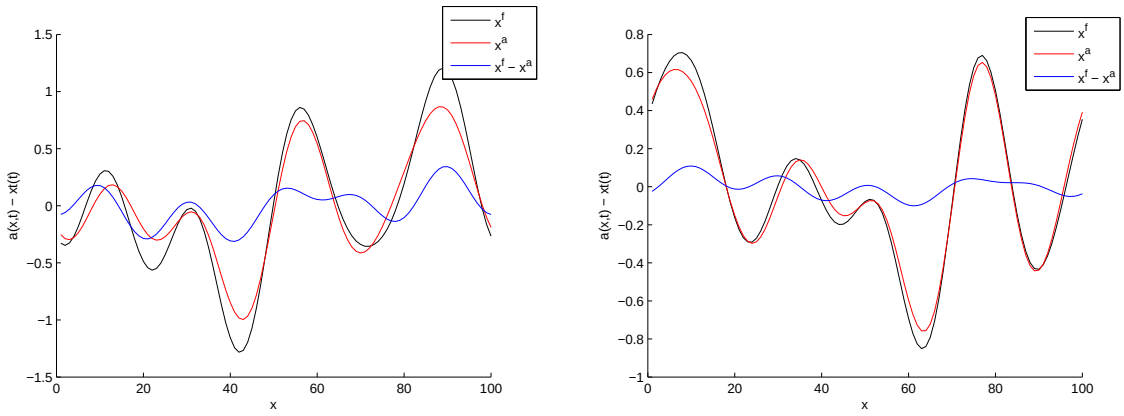

(a) Difference between forecast and analysis states at the second assimilation

(b) Difference between forecast and analysis states at the fourth assimilation

Figure 5.7: Comparison of the size of the difference between the forecast and analysis ensemble means.

Figure 5.7 shows the ensemble forecast mean, relative to the truth, in black; the ensemble analysis mean, relative to the truth is in red and the difference between them is in blue. The difference between these is a measure of the impact of the assimilation. The greater the difference between the forecast and analysis the more influence the assimilation of an observation has had. Figure 5.7(a) is for the second assimilation, 20 seconds and figure

5.7(b) is at the fourth assimilation, 40 seconds. The difference between the forecast and the analysis estimates at the fourth assimilation is significantly less than at the second assimilation. Indicating that the state analysis and the state forecast are very similar. The observations are becoming less effective in adjusting the forecast solution to more accurately represent the truth. This is filter divergence, due to inbreeding. The filter has a diminishing capacity to adjust the forecast estimate more accurately reflect the true solution, due to an underestimation in the forecast error covariance. This is consistent with the results of Hamill et al. (2001) who states that during subsequent assimilation cycles a "variance-deficient ensemble further underestimates the background error statistics, disregarding even more the influence of new observations". Further they say that this can worsen with time and in the presence of chaotic dynamical systems, error growth/decay between assimilations, can lead to a useless forecast.

Filter divergence is associated with a convergence in the spread of the ensemble members (see section 3.1.3). Figure 5.8 demonstrates this.

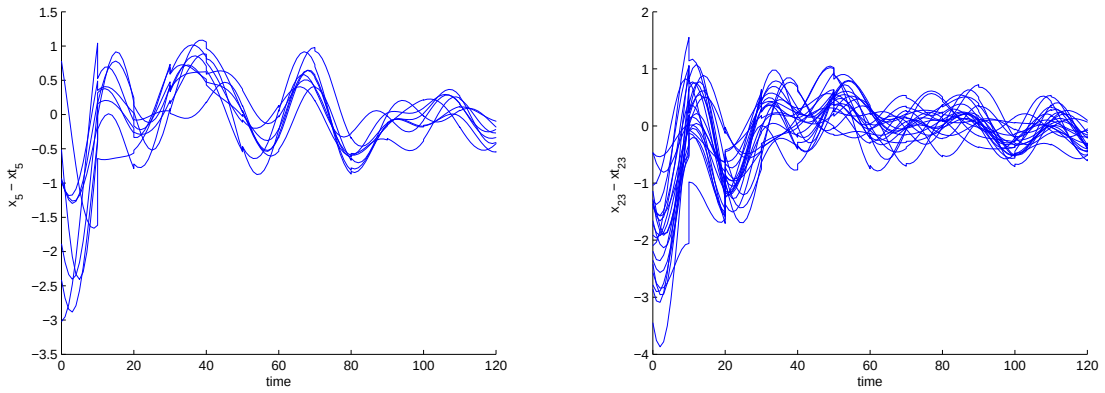

(a) Evolution in time of eight ensemble members for state component x_5 (an observed component) relative to the true solution

(b) Evolution in time of twenty ensemble members for state component x_{23} (an unobserved component) relative to the true solution

Figure 5.8: Ensemble convergence.

Figure 5.8(a) shows how the ensemble spread of eight ensemble members decreases in time. Each blue line represents an individual ensemble member. Observations are made at every five grid points. This figure shows the evolution of the fifth state component, a component that has an observation. The forecast and analysis estimates are plotted relative to the truth, such that a perfect estimate would have a value of zero. Over time it can be seen that the ensemble spread is decreasing, the ensemble members are converging. They are converging on an estimate that is close to zero. Figure 5.8(b) shows a similar

type of plot except the evolution is of the 23^{rd} state component for twenty ensemble members. There is no observation of this component however the same basic behaviour of convergence of the ensemble members can be seen. This larger ensemble converges more quickly. This is desirable if the analysis state estimate is a good representation of the truth but if the analysis state is a poor estimate of the truth then this is an undesirable feature as the filter will have diverged.

These results showing filter divergence due to inbreeding are consistent with those found by others for example Hamill et al. (2001); Anderson (2001); Whitaker and Hamill (2002); Houtekamer and Mitchell (2001)

5.1.4 Spurious Correlations

In this model setup the model is advected by one grid point at each timestep. Now $dt = 1s$ and $u = 1ms^{-1}$, and so it follows that $dx = 1m$. Observations are taken at every 10 timesteps, in this time the solution will have been advected along by 10 grid points. This means that errors will also have been advected by 10 grid points between observations. This gives a lower estimate of a length scale for physical correlations which can be used a cutoff length to filter noise. The initial conditions for each ensemble member of this model are given by a sum of sine functions (section 4.1). Since sine function are periodic some correlations at a distance will be physical. The period of the function is 100 seconds, thus we would expect correlations at a distance of 50 grid points. This gives an upper bound to the cutoff length for physical correlations. This determination of cutoff length is heuristic and is used only as guide. It could be approximated that correlations that develop at a distance greater than 20 grid points are spurious. There are many such correlations that have developed as can be seen in the representations of the $\mathbf{P}^f$ and $\mathbf{P}^a$ matrices in figures 5.6.

5.1.5 Ensemble size, undersampling and standard deviation

This model is based on simple linear advection and no errors are associated within the model except those of rounding errors. This implies that after the first observation has been assimilated, where the ensemble mean becomes a reasonable representation of the true solution, it is not expected that the ensemble mean will deviate from this good representation during the period between the observation assimilations. This is due to

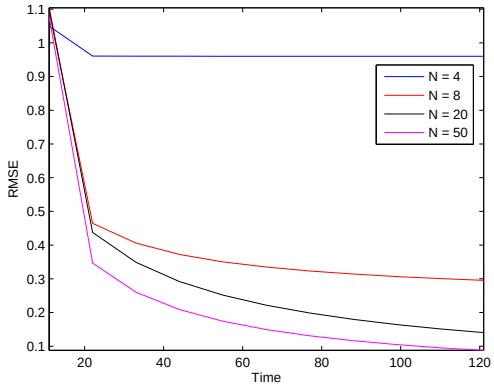

Figure 5.9: The RMSE of different numbers of ensemble members as a function of time.

this model having no chaotic nonlinear dynamics. This can also be seen in the smooth evolutions of the forecast solutions between observations in the figures of section 5.1.1. A reasonable representation of the true solution has an impact on the analysis of the effectiveness of localization in the ETKF.

The size of the ensemble in ensemble Kalman filtering is of critical importance and is essential that the ensemble size is sufficient to adequately represent the model state. The accuracy of the forecast estimates can be analysed by calculating the root mean square error (RMSE) between the true solution and the forecast or analysis estimates. The definition of the root mean square is given in Barlow (1989). This can be easily adapted to give the RMSE; it is calculated from

$$x_{rmse} = \sqrt{\frac{1}{n}\sum \left(\mathbf{x}^t - \mathbf{x}^f\right)^2}. \tag{5.1}$$

Figure 5.9 shows that the RMSE decreases significantly with the size of the ensemble. The largest increase in accuracy is between four and eight ensemble members. Where four ensemble members has a RMSE of approximately one for the duration of the assimilation window, the RMSE of the forecast with eight ensemble members decreases to around 0.3 by the end of the assimilation window. There is a sharp decrease in the RMSE over the first 21 seconds for ensembles of eight, twenty and fifty members due to the increase in the accuracy of the estimate by the assimilation of the observations. It can be seen that when the number of ensemble members is small compared to the size of the state that the forecasts and analyses are less accurate. Thus it can be concluded that undersampling

can have a dramatic effect on the quality of the analysis provided, this is consistent with results of Hamill et al. (2001).

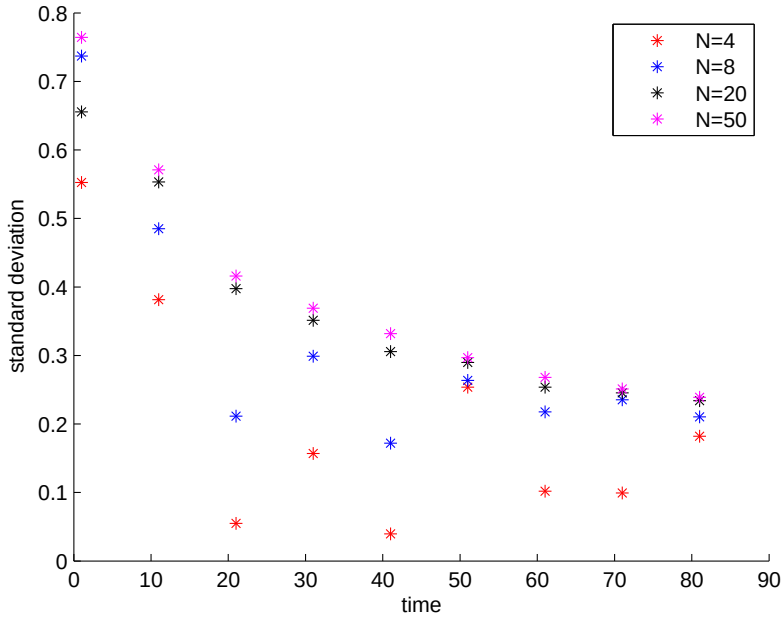

Figure 5.10: Standard deviations of analysis ensemble spread.

Figure 5.10 shows the magnitude of the analysis standard deviations of the ensemble spread during the first 90 seconds of the assimilation for different numbers of ensemble members. Over time the general trend of the standard deviations is to decrease for all different sizes of ensemble members. There is however anomalous behaviour when the ensemble size is reduced to four or eight ensemble members. At each alternating assimilation there is a pattern of a decrease in standard deviation followed by an increase. The state space is of size 100, however due to the ensemble being generated by a linear combination of five sine waves, each with a random phase and amplitude, the size of the model subspace is 10 (Oke et al., 2007). For the following experiments we can consider ensembles of less than 10 to be undersampled.

5.2 Undersampled behaviour

The behaviour of the model when there are only 4 ensemble members and the state is significantly undersampled is now examined.

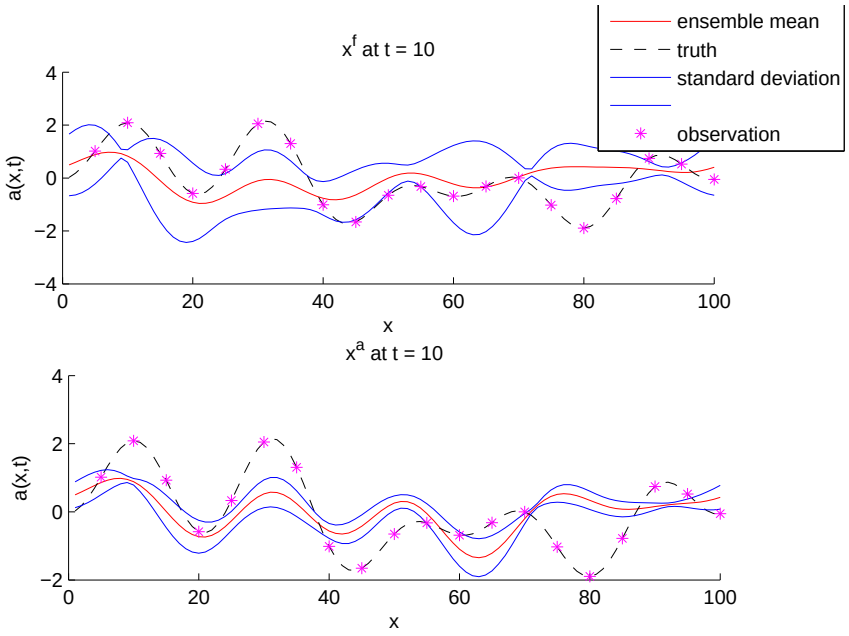

Figure 5.11: Forecast and analysis estimates, over the whole domain, for four ensemble members at the first assimilation.

Figure 5.11 shows the forecast (top figure) and analysis (bottom figure) estimates over the whole domain for the first assimilation when there are only four ensemble members. There is great difference in the analysis estimate of four ensemble members from when eight members were used (figure 5.5). As before initially the ensemble mean does not accurately represent the true solution, however the standard deviation of the ensemble members does not do as good a job of encapsulating the truth function as when there are eight ensemble members, as would be expected. After the assimilation it is obvious that that the incorporation of the observation has been significantly less effective as the ensemble mean does not accurately represent the true solution. The spread of the ensemble members has again become significantly reduced.

Figure 5.12 shows the forecast error covariances over the first assimilation when there are only four ensemble members; these plots are created in the same way as those of figure 5.6. The general structure of the forecast and analysis error covariances is the same as for eight ensemble members. The same leading diagonal structure can be seen, as can the periodic boundary correlations and the spurious correlations at a distance. The behaviour of the error covariances over the assimilation is to decrease with the assimilation as expected Furrer and Bengtsson (2007). It is in size that the covariances differ from those with eight ensemble members. Just before the first assimilation for four

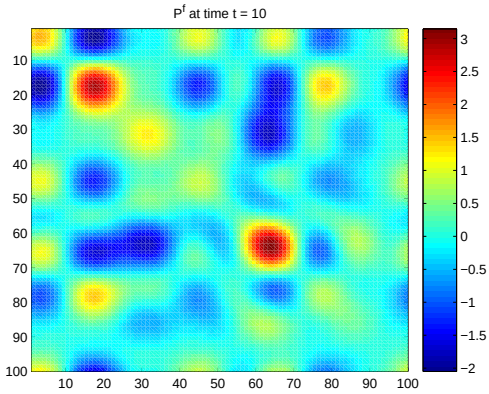

(a) P^f just before the first assimilation

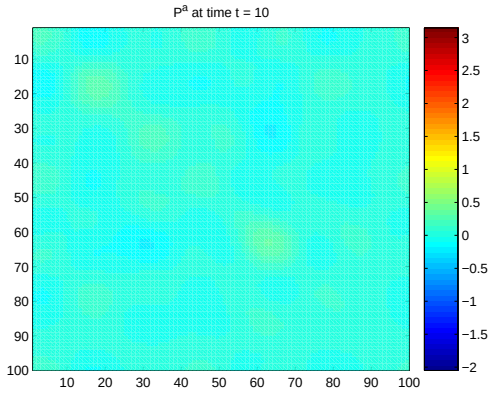

(b) P^a just after the first assimilation

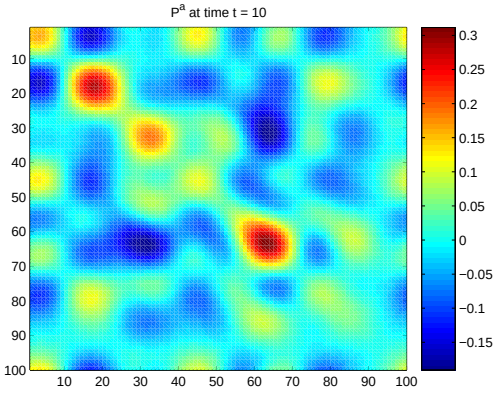

(c) P^a just after the first assimilation with a reduced colour scale

Figure 5.12: Error covariances over the first assimilation t = 10s for four ensemble members.

ensemble members the covariances range from -2 to 3 compared to -1 and 1.5 when there are eight ensemble members. Just after the first assimilation they range from -0.15 to 0.3 compared to -0.04 to 0.1 for eight ensemble members. Although initially the covariances are large enough to place enough confidence in the observations at the first assimilation, the scheme does not do a good job of correcting the forecast, perhaps this is due to the significant undersampling or there not being enough observations which are every five grid points. For four ensemble members the ensemble mean is not a good representation of the true state after the first assimilation and the forecast error covariances have been reduced. In this situation we have significant undersampling and inbreeding so the scheme places much more confidence in the a-priori state and the observations are not assimilated effectively leading to filter divergence.

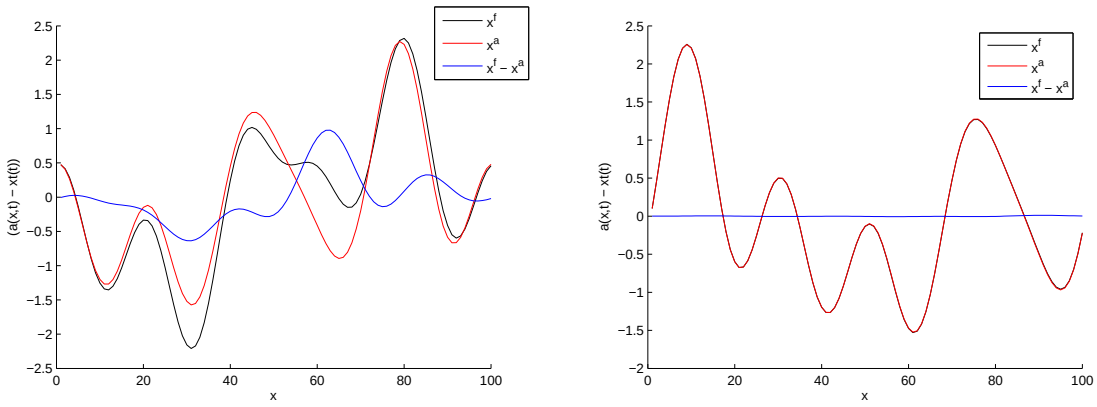

(a) Difference between forecast and analysis states at the first assimilation for four ensemble members

(b) Difference between forecast and analysis states at the fourth assimilation for four ensemble members

Figure 5.13: Comparison of the size of the difference between the forecast and analysis ensemble means for four ensemble members.

Figure 5.13 is a similar plot to that of figure 5.7, where the blue line is a measure of the adjustment of the forecast estimate, by the observation assimilation, to the analysis estimate. At the first assimilation, figure 5.13(a), the adjustment of the forecast estimate to the analysis estimate is clear to see. By the fourth assimilation, figure 5.13(b), the adjustment is almost negligible. The filter divergence in this situation will be more problematic as the ensemble mean is not a good representation of the truth function. Figure 5.13 shows that filter divergence is associated with the undersampling.

5.3 Assimilation runs with localization

5.3.1 Inclusion of the inflation factor

An inflation factor was included in the ETKF (section 4.3) according to the equation $\mathbf{x}_i^f \leftarrow r\mathbf{X}'^f + \bar{\mathbf{x}}^f$ as in section 3.2.1. The root mean square error(RMSE) (equation 5.1) between the the truth function and the forecast and analysis means were calculated. The RMSE was plotted as a function of time, figure 5.14, to analyse the improvements made by the inclusion of the inflation factor.

Table 5.1 shows an analysis of these figures. The percentage improvements were calculated between the RMSE for r = 0 at 121 seconds. Where there are only four ensemble

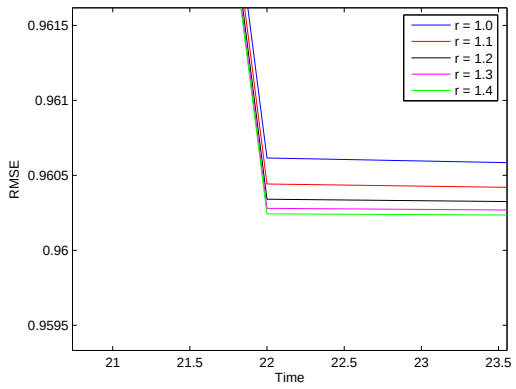

(a) RMSE for varying inflation factor with 4 ensemble members

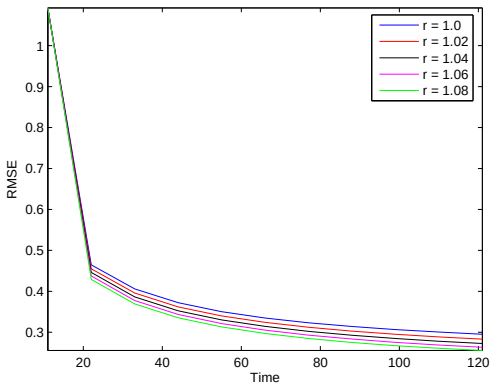

(b) RMSE for varying inflation factor with 8 ensemble members

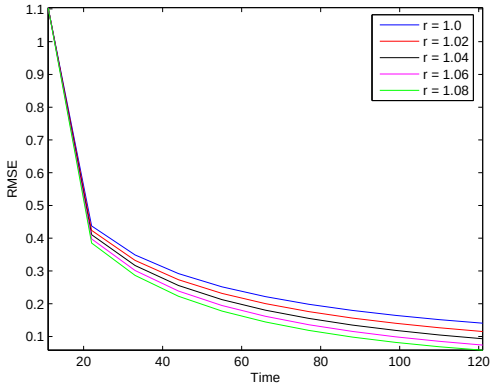

(c) RMSE for varying inflation factor with 20 ensemble members

Figure 5.14: RMSE as a function of time for varying inflation factors and different numbers of ensemble members.

members the RMSE was calculated at 23.5 seconds as it has become approximately static.

Number of ensemble members	**r**	**% improvement**
4	1.5	0.04
8	1.08	12
20	1.08	57

Table 5.1: Table to show the % improvement of the RMSE for different numbers of ensemble members by the inclusion of an inflation factor

The data from table 5.1 and figure 5.14(a) shows that the inclusion of an inflation factor for an ensemble of four members makes little to no improvement in the quality of the

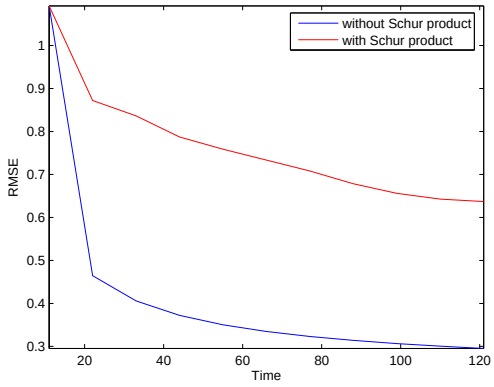

Figure 5.15: Comparison of the RMSE with and without the implementation of a Schur product.

analysis mean. This can be seen by the scale of the RMSE on the y-axis of figure 5.14(a), the very large inflation factors used and the very small improvement of only 0.04% after the second assimilation. The larger the ensemble, the greater the effect of the inflation factor. For twenty ensemble members an inflation factor of $r = 1.08$ has a 57% improvement in the RMSE. It is expected that the greater the size of the ensemble the greater the impact of the inflation factor (Hamill et al., 2001). If the ensemble spread is very small, as is consistent with a small number of ensemble members, then it will require a large multiplicative factor to significantly increase the spread. For a larger number of ensemble members the ensemble spread will be much greater and the multiplication by a factor will have a more significant impact. Hamill et al. (2001) find that inflation factors of only 1% or 2 % are adequate to improve their analysis for 100 ensemble member. For a 400 member ensemble they find that only a 0.25% - 1% is required. This is consistent with the results shown in table 5.1. As the number of ensemble members increases so the greater the influence of the inflation factor.

5.3.2 Inclusion of a Schur product

A Schur product was implemented as described in section 4.4.2 in an attempt to achieve covariance localization within the ETKF.

The overall performance of the filter can be examined using the RMSE.

Figure 5.15 shows the RMSE of the solution, in time, for eight ensemble members, with

and without a Schur product being applied. Where the Schur product is applied the cutoff length scale used is of $l = 20$. It is clear that the inclusion of a Schur product in the manner described in section 4.4.2 has significantly impacted the effectiveness of the ETKF to estimate the state of the system. Unfortunately this means that the approximation described is a bad one and could not be recommended for use in this current state.

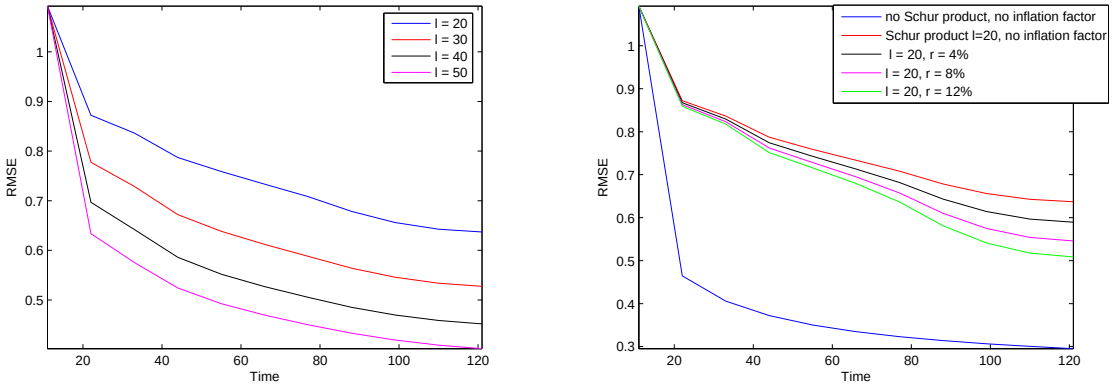

(a) Comparison of the RMSE with the implementation of a Schur product and varying the cutoff length

(b) Comparison of the RMSE with the implementation of a Schur product and varying the inflation factor

Figure 5.16: Performance of the ETKF when varying the cutoff length and the inflation factor.

If the distance of the cutoff length in the Schur product is altered there is an alteration in the RMSE. Figure 5.16(a) shows the RMSE for various cutoff lengths, where an ensemble of eight members is used. As the cutoff length becomes larger the RMSE decreases towards that of the RMSE where there was no Schur product included. The effect of the Schur product is reduced as the cutoff length increases, so the RMSE in this instance is reducing.

The effectiveness of the Schur product can be tuned by increasing the inflation factor. Figure 5.16(b) shows that when a Schur product is applied in conjunction with an inflation factor the state estimate improves as can be seen by the drop in the RMSE. This implies that it may be possible to include a Schur product in this way and tune the inflation factor to improve the quality of the analysis.

The effect on the error covariances can be examined. Figure 5.17 shows the forecast and analysis error covariances over the first assimilation. Many of the spurious correlations have not been removed (figure 5.17(c)). This figure also shows that the maximum size of

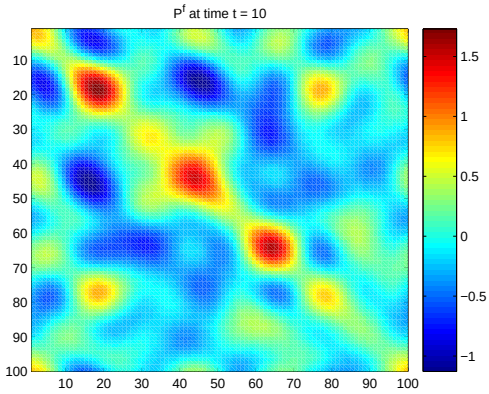

(a) P^f just before the first assimilation t = 10s

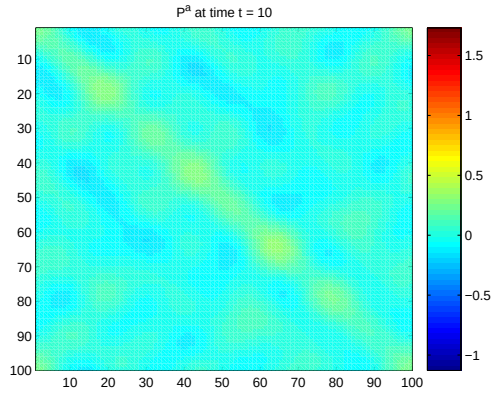

(b) P^a just after the first assimilation

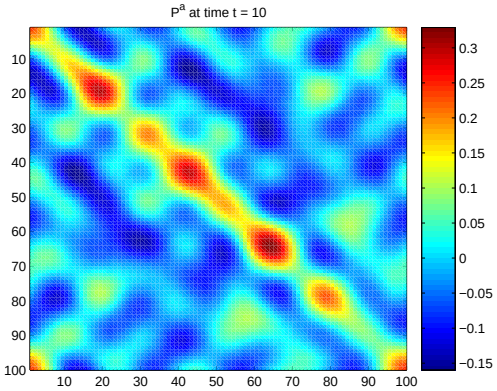

(c) $\mathbf{P}^a$ at the first assimilation with a reduced colour scale

Figure 5.17: Error covariances over the first assimilation t = 10s with the inclusion of a Schur product.

the correlations has been affected. The range of covariances after the Schur product is from -0.15 to 0.3. In the idealized Schur product implementation (figure 3.1) there is no reduction in the maximum size of the correlations. This is indicating that physical variances have been reduced by this Schur product approximation. Figure 5.17 is indicating that the Schur product in this form is not functioning as desired.

A crucial benefit of localization is the increase in rank of the forecast error covariance, as this increases the effective size of the ensemble (Oke et al., 2007; Ehrendorfer, 2007), (see section 3.2.2). The effect of this implementation of the Schur product on the eigenvalues of the forecast error covariance can be plotted. To produce figure 5.18 the ETKF for eight ensembles was run and the eigenvalues of the forecast error covariance were plotted for $t = 10s$. The approximated Schur product was introduced into the ETKF, with a cutoff

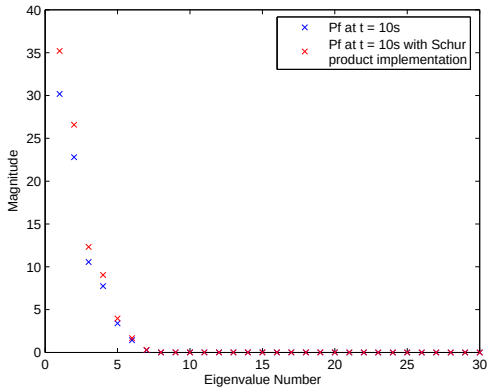

Figure 5.18: Effect on the eigenvalues of the approximated ETKF implementation of the Schur product.

length of $l = 20$, and an inflation factor of $r = 1.08$ was used. The x-axis is truncated at 30 as all values beyond this are zero and it allows the details of the eigenvalues to be seen more clearly. The inclusion of the Schur product has not increased the number of eigenvalues in the way expected by the general implementation of the Schur product (section 3.2.2). The magnitude of the eigenvalues has increased however this is likely to be due to the inclusion of the inflation factor. Thus a major benefit of increasing the rank of the forecast error covariance is not experienced with this approximated version of the Schur product for the ETKF.

The approximation introduced to include the Schur product in the manner described in section 4.4, $\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right) = \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T$, can be examined by plotting these matrices.

The perturbation matrix used to produce the figures of 5.19 was at time $t = 10s$, for eight ensemble members. It is important to note that the additional approximation of extending the columns of the perturbation matrix as in section 4.4.2 was made in producing figures 5.19(a) and 5.19(b). It is clear to see from figure 5.19 that this approximation is not a very good one. It was shown analytically in section 4.4.2 that this was merely an approximation. However figure 5.19 gives a more qualatative result. Clearly the size of the differences around the leading diagonal are substantial in comparison to the values in this area without the inclusion of a Schur product. This indicates that the approximation $\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right) = \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T$ is a poor one.

It is concluded that this approximation of covariance localization by Schur product for

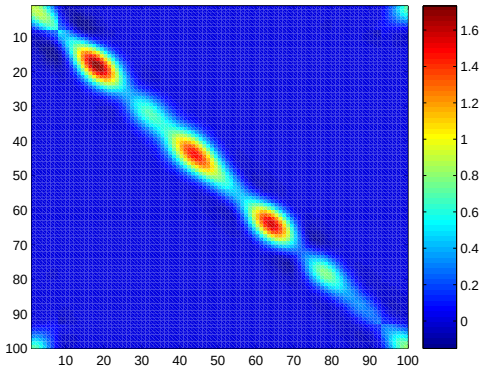

(a) $\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right)$

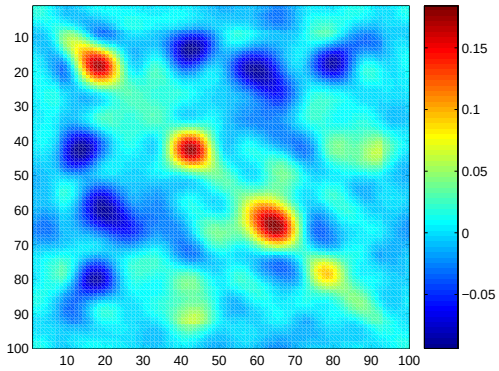

(b) $\left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T$

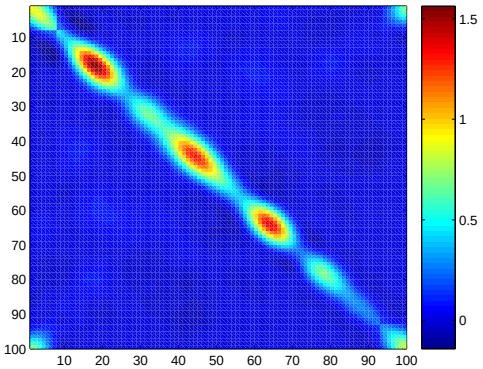

(c) $\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right) - \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T$

Figure 5.19: Approximation of $\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right) = \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T$.

the ETKF is not a good one. It suffers many problems; the RMSE is worsened by its inclusion, spurious corrleations are not effectively removed, the rank of the forecast error covariance is not increased and as in any application of Schur product to the forecast error covariance dynamical balances may be disrupted (Oke et al., 2007).

In this chapter the problems of undersampling, inbreeding, filter divergence and the development of long range spurious correlations have been demonstrated for ETKF. The performance of the new approximation of including a Schur product for the ETKF, section 4.4.2, was analysed. It was found that this implementation was a poor approximation. The following chapter provides a conclusion to this report.

Chapter 6

Conclusion

6.1 Summary and Discussion

The motivation for this project was to examine the effects of undersampling in ensemble Kalman filtering. Undersampling is when the number of ensemble members is small compared to the size of the state. Where the ensemble, which provides a statistical sample of the estimates, is so small that the ensemble does not statistically represent the state of the system subsequent problems may arise. In this project the specific implementation of Livings (2005) of the ensemble transform Kalman filter (ETKF), introduced by (Bishop et al., 2001), was used to investigate this problem.

A detailed description of the formulation of the Kalman filter (KF) (Kalman and Bucy, 1961) and the ETKF was presented (chapter 2). The ETKF is deterministic ensemble filter, it does not use perturbed observations which introduce noise into the system. As it is an ensemble filter it is able to deal with nonlinear models. It is a beneficial form of ensemble Kalman filer due its efficiency and speed computationally (Hamill, 2006).

The three problems caused by undersampling investigated in this project were inbreeding, filter divergence and the development of long range spurious correlations. Inbreeding is the successive reduction in the estimate of the size of the forecast error covariance, $\mathbf{P}^f$ after each observation assimilation (Houtekamer and Mitchell, 1998). This is anticipated as the analysis error covariance is always less than the forecast error covariance (Furrer and Bengtsson, 2007). Filter divergence occurs when the analysis estimate of the system is incorrectly specified, perhaps as a consequence of undersampling, and this analysis is

unable to be corrected by the observations due to inbreeding (Hamill et al., 2001). The Kalman gain of the ETKF (equation 2.26) determines how much weighting should be given to the observations as a ratio of the errors associated with the observations and the a-priori, or background, state. Underestimation of the forecast error covariance leads to less relative weighting being placed on the observations and more to the a-priori state and an incorrect convergence of the ensemble members, consistent with Hamill et al. (2001) (figure 5.8). As a result the observations have a diminishing effect on adjusting the forecast state to a more accurate analysis state. Long range spurious correlations are correlations between state components that are not physically related and typically they are at a significant distance from each other. These are represented in the forecast error covariance matrix. The development of these spurious correlations is as a result of undersampling (Hamill et al., 2001).

Two methods that are currently used to overcome these problems are covariance inflation and covariance localization. Covariance inflation (Anderson, 2001) applies an inflation factor to the forecast error covariance in an attempt to negate the problem of inbreeding and subsequent filter divergence. This is done by inflating the deviations of the background error from the ensemble mean, for each ensemble member, as in equation 3.3. Covariance localization is the removal of the long range spurious correlations (Hamill et al., 2001). This is typically achieved by taking a Schur product, elementwise multiplication, of the forecast error covariance and a correlation function such that it removes spurious correlations. The correlation function generally used is the 5^{th} order compactly supported piecewise function defined by Gaspari and Cohn (1999) (equation 3.6).

To demonstrate these problems a simple linear advection model (equation 4.1) was implemented within the ETKF of Livings (2005).

It was found that the ETKF exhibited all the problems expected due to undersampling. Inbreeding due to undersampling was shown in figures 5.6 and 5.12. The size of the covariances decreased from the forecast error covariance to the analysis error covariance. This is as expected since the forecast analysis covariance is "negatively biased" (Furrer and Bengtsson, 2007). The inbreeding led to filter divergence and was demonstrated in figures 5.7 and 5.13. The reduction of the size of the forecast error covariance after the assimilation and the reduction of the effectiveness of the assimilation of observations is consistent with results of Hamill et al. (2001) and Furrer and Bengtsson (2007). The development of long range spurious correlations due to undersampling can be seen in the representations of the error covariances of figures 5.6 and 5.12. This behaviour is also

consistent with Hamill et al. (2001).

Covariance inflation was successfully implemented in the ETKF according to equation 3.3. This can be seen in figure 5.14. The inclusion of the inflation factor reduced the RMSE of the analysis state estimate compared to the value of true solution; thus aiding the ETKF in producing more meaningful results. This is again consistent with results of other authors i.e. Anderson (2001); Hamill et al. (2001) and Whitaker and Hamill (2002).

Covariance localization by use of a Schur product is shown schematically in figure 3.1. The application of the Schur product to a forecast error covariance removes the spurious correlations (Hamill et al., 2001). The distance at which the correlations are cutoff is specified within the definition of the correlation function, ρ, (equation 3.6). The choice of this cutoff length scale is heuristic and can depend on the physical nature of the model implemented. The implementation of a Schur product is also seen to increase the rank of the forecast error covariance, figure 3.2 (Oke et al., 2007). This is an important benefit of the Schur product method of covariance localization as it increases the "effective size" of the ensemble and as such could be a useful tool in undersampled systems. The same result is demonstrated by Oke et al. (2007).

The implementation in the ETKF of covariance localization by use of a Schur product was problematic. The nature of the ETKF, which is a square root filter, is not use the forecast error covariance matrix explicitly (see section 2.2.3). The forecast error covariance is implied through the appearance of the ensemble state perturbation matrix. According to Lorenc (2003) and Hamill (2006) it is not possible to achieve covariance localization by a Schur product in the ETKF as it is currently defined.

An approximation of the inclusion of the Schur product in the ETKF was sought. Approximations within the equations of the ETKF were made to attempt implementation. It was approximated that

$$\left(\check{\rho}\check{\rho}^T\right) \circ \left(\mathbf{X}'(\mathbf{X}')^T\right) = \left(\check{\rho} \circ \mathbf{X}'\right)\left(\check{\rho} \circ \mathbf{X}'\right)^T, \tag{6.1}$$

although this is not true in the general sense (section 4.4.2 and figure 5.19). Modifications were then necessary to ensure that the matrices of the ensemble state perturbations and the correlation function square root were of the same dimensions. The Schur product

was applied in the Livings (2005) implementation in the Kalman gain as

$$\mathbf{K}_e = \left(\check{\rho} \circ \mathbf{X}'^f\right) \mathbf{U} \left(\mathbf{\Sigma} \left(\mathbf{\Sigma}\mathbf{\Sigma}^T + \mathbf{I}\right)^{-1} \mathbf{V}^T \mathbf{R}^{-\frac{1}{2}}\right). \tag{6.2}$$

This in effect only applies the localization to the term $\mathbf{P}^f\mathbf{H}^T$; Hamill et al. (2001) also only apply the Schur product to this term. A further approximation of $\left(\check{\rho} \circ \mathbf{X}'^f\right)^T \mathbf{H}^T = \mathbf{Y}'$ was also required. This implementation of the Schur product is now consistent with square root formulation. Although now the analysis statistics are no longer strictly consistent with the desired properties of an ensemble Kalman filter (Livings et al. (2008)). For example the relation $\mathbf{P}^a = (\mathbf{I} - \mathbf{KH})\mathbf{P}^f$ is no longer strictly held.

This method of implementing the Schur product increases the RMSE of the analysis estimate compared to the true solution; figure 5.15, this is highly undesirable. However it may be possible to improve the performance of the covariance localization by the simultaneous use of an inflation factor, figure 5.16(b). An important benefit of this implementation is the increased rank of the forecast error covariance. Figure 5.18 shows that there is no increase in the rank of the forecast error covariance when this approximated Schur product is implemented within the ETKF. This is a major disadvantage of this approximation as it does not allow the "effective" ensemble size to be increased (see section 3.2.2). The correlation length scale chosen will also affect the the overall success of this implementation (figure 5.16(a)). If the length is too short important dynamical correlations maybe lost, if it is too long then spurious correlations may not be removed (Ehrendorfer, 2007). Spurious correlations that the Schur product was included to overcome have not been removed in the way desired (figure 5.17). This new approximation is highly heuristic and requires many crucial approximations. An effective inclusion of the Schur product in the ETKF has not been found within the scope of this project.

6.2 Further work

Use of a more complex model

This new method of implementing the Schur product could be tested with a more complex dynamical model. A more complex dynamical model would allow error growth in the solution between the observation times (Ehrendorfer, 2007). The effect on the RMSE between the analysis state and the true solution with the Schur product could then be

analysed. Implementing the Schur product together with an inflation factor may also be possible in a more complex model. This may have the effect of increasing the accuracy of the analysis estimate and spurious correlations being removed.

The effect of this new implementation on dynamical balance

It is known that the inclusion of a Schur product detrimentally influences the dynamical balances with the model (Oke et al., 2007). The effect of this implementation on the dynamical balance should be investigated. This could be done using the same method as Oke et al. (2007). A second model variable that is related to the first could be introduced. The effect of the localization on the this model variable could then be examined. This is essential if a more complex dynamical model were to be implemented.

Bias introduced by the new implementation of the Schur product

Bias in a filter is when the ensemble statistics are no longer centred on the ensemble mean (Livings (2005) and Livings et al. (2008)). This would be seen in the matrix of the analysis state where the columns no longer sum to zero. The effect of the new implementation of the Schur product on the introduction of bias into the filter should also be investigated.

Using a different method to filter the spurious correlations

A different method of covariance localization within the ETKF could be explored. At present the Schur product method is far from ideal. An alternative method of removing the spurious correlations, such as truncated Fourier series, could be investigated.

Inclusion of model error

The model error term $\mathbf{Q}$ is set to zero, that is we assume that the model is perfect. This is widely made assumption in the testing of new techniques. The inclusion of model error is an area of ongoing research. It is important to examine the effects of including this term, particularly for if a nonlinear, and more chaotic model was implemented.

Bibliography

Anderson, J. L., 2001: An Ensemble Adjustment Kalman Filter for data assimiliation. *Mon. Wea. Rev.*, **129**, 2884–2903.

Anderson, J. L. and S. L. Anderson, 1999: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev.*, **126**, 2741–2758.

Barlow, R. J., 1989: *A guide to the Use of Statistical lMethods in the Physical Sciences.* John Wiley and Sons.

Bishop, C. H., B. J. Etherton, and S. J. Manjundar, 2001: Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Wea. Rev.*, **129**, 420–436.

Ehrendorfer, M., 2007: A review of issues in ensemble-based Kalman filtering. *Meteorol. Z.*, **16**, 795–818.

Evensen, G., 1994: Sequential data assimilation with a nonlinear quasigeostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99(C5)**, 10 143–10 162.

Evensen, G., 2003: The Ensemble Kalman Filter: Theoretical formulation and practical implementation. *Ocean Dynamics.*, **53**, 343–367.

Furrer, R. and T. Bengtsson, 2007: Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *Journal of Multivariate Analysis*, **98**, 227–255.

Gaspari, G. and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Q. J. R. Meteorol. Soc.*, **125**, 723–757.

Golub, G. H. and C. F. Van Loan, 1996: *Matrix computations*. The Johns Hopkins University Press, 3rd edition.

Hamill, T., J. S. Whitaker, and C. Snyder, 2001: Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.*, **129**, 2776–2790.

Hamill, T. M., 2006: *Ensemble forecasting and data assimilation: two problems with the same solution - In: Predictability of Weather and Climate, T. Palmer, R. Hagedorn (Eds.)*. Cambridge University Press, Cambridge, 124-156 pp.

Horn, R., 1990: The Hadamard product. *Matrix theory and applications*, C. R. Johnson, Ed., American Mathematical Society, Proceedings of Symposia in Applied Mathematics, Vol. 40, 87–169.

Houtekamer, P. L. and H. L. Mitchell, 1998: Data assimilation using an Ensemble Kalman Filter technique. *Mon. Wea. Rev.*, **126**, 796–811.

Houtekamer, P. L. and H. L. Mitchell, 2001: A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, **129**, 123–137.

Jazwinski, A. H., 1970: Stochastic Processes and Filtering Theory. *Academic Press*, **57**, 32–45.

Kalman, R. and K. Bucy, 1961: New results in linear prediction filtering theory. *Trans. AMSE J. Basic Eng.*, **83D**, 95–108.

Kalnay, E., 2003: *Atmospheric modelling: Data assimilation and predictability*. Cambridge University Press, Cambridge.

Kalnay, E., B. R. Hunt, E. Ott, and I. Szunyogh, 2006: *Ensemble forecasting and data assimilation: two problems with the same solution - In: Predictability of Weather and Climate, T. Palmer, R. Hagedorn (Eds.)*. Cambridge University Press, Cambridge, 157-176 pp.

Le Dimet, F.-X. and O. Talagrand, 1986: Variational algorithms for analysis and assimilation of meteorological observations: Theoretical aspects. *Tellus*, **38A**, 97–110.

Livings, D. M., 2005: Aspects of the Ensemble Kalman Filter. *Reading University Masters Thesis*.

Livings, D. M., S. L. Dance, and N. K. Nichols, 2008: Unbiased ensemble square root filters. *Physica D.*, **237/8**, 1021–1028, http://dx.doi.org/10.1016/j.physd.2008.01.005.

Lorenc, A. C., 1986: Analysis methods for numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, **122**, 1177–1194.

Lorenc, A. C., 2003: The potential of the ensemble Kalman filter for NWP - a comparison with 4D-VAR. *Q. J. R. Meteorol. Soc.*, **129**, 3183–3203.

Oke, P. R., P. Sakov, and S. P. Corney, 2007: Impacts of localisation in the EnKF and EnOI: experiments with a small model. *Ocean Dynamics.*, **57**, 32–45.

Ott, E., et al., 2004: A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, **56A**, 415–428.

Schur, I., 1911: Bemerkungen zur theorie der beschrnkten bilinearformen mit unendlich vielen vernderlichen. *J. reine angew. Math.*, **140**, 1–28.

Szunyogh, I., E. J. Kostelich, G. Gyarmati, D. J. Patil, B. R. Hunt, E. Kalnay, E. Ott, and J. A. Yorke, 2005: Assessing a local ensemble Kalman filter: perfect model experiments with the National Centers for Environmental Prediction global model. *Tellus*, **57A**, 528–545.

Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamil, and J. S. Whitaker, 1999: Ensemble square root filters. *Mon. Wea. Rev.*, **131**, 1485–1490.

UK Met Office, 2008: Observations. *http://www.metoffice.gov.uk/research/nwp/observations/*, last accessed 2008.

Wang, X., C. H. Bishop, and S. J. Julier, 2004: Which is better, an ensemble of positive-negative pairs or a centered spherical simplex ensemble. *Mon. Wea. Rev.*, **132**, 1590–1605.

Whitaker, J. S. and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924.