Available online at www.sciencedirect.com
SCIENCE DIRECT®
Ocean Modelling 11 (2006) 441–463
Ocean Modelling
www.elsevier.com/locate/ocemod

Parameter sensitivity of three Kalman filter schemes for assimilation of water levels in shelf sea models

J.V.T. Sørensen [a,*], H. Madsen [a], H. Madsen [b]

[a] *DHI Water & Environment, Agern Allé 5, DK-2970 Hørsholm, Denmark*
[b] *Informatics and Mathematical Modelling, Technical University of Denmark, Richard Petersens Plads—Building 321, DK-2800 Kongens Lyngby, Denmark*

Received 19 March 2004; received in revised form 11 March 2005; accepted 15 March 2005
Available online 10 May 2005

Abstract

In applications of data assimilation algorithms, a number of poorly known assimilation parameters usually need to be specified. Hence, the documented success of data assimilation methodologies must rely on a moderate sensitivity to these parameters. This contribution presents a parameter sensitivity study of three well known Kalman filter approaches for the assimilation of water levels in a three dimensional hydrodynamic modelling system. The filters considered are the ensemble Kalman filter (EnKF), the reduced rank square root Kalman filter (RRSQRT) and the steady Kalman filter. A sensitivity analysis of key parameters in the schemes is undertaken for a setup in an idealised bay. The sensitivity of the resulting root mean square error (RMSE) is shown to be low to moderate. Hence the schemes are robust within an acceptable range and their application even with misspecified parameters is to be encouraged in this perspective. However, the predicted uncertainty of the assimilation results are sensitive to the parameters and hence must be applied with care. The sensitivity study further demonstrates the effectiveness of the steady Kalman filter in the given system as well as the great impact of assimilating even very few measurements.

Keywords: Kalman filters; Tide gauges; Parameter sensitivity

* Corresponding author. Tel.: +45 4516 9304.
E-mail addresses: jts@dhi.dk, jts@imm.dtu.dk (J.V.T. Sørensen).

1463-5003/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ocemod.2005.03.002

1. Introduction

Data assimilation methodologies are becoming increasingly applied in the ocean modelling community. As part of this trend, sequential ensemble based techniques have gained significant popularity during the last decade. Like most other statistical approaches to assimilation, such techniques rely on simple assumptions about the sources of the error in the numerical ocean models. The successful application of these assimilation techniques must rely either on extensive calibration of the parameters describing the model errors or on a moderate sensitivity to these parameters. This contribution sets out to examine this sensitivity to assimilation parameters, in order to provide an improved understanding of the nature of assimilation techniques. First, it is instructive to review the foundation of present days efficient sequential assimilation techniques.

The standard approach and hence terminology of sequential estimation techniques is that of the Kalman filter (Kalman, 1960). The original Kalman filter was derived for a linear system with Gaussian error sources. When applied to non-linear and high dimensional systems, the formulation demands vast computational resources and its limitations in terms of Gaussian error assumptions and linearity become clear. Several extensions have been made in an attempt to accommodate for such deficiencies.

Primarily, the problem needs to be solvable on available computational resources. The most widespread techniques for making the problem tractable are ensemble based. Basically these schemes represent the information contained in the error covariance matrix in a reduced space spanned by a small ensemble of states. The ensemble Kalman filter (EnKF) (Evensen, 1994) and the reduced rank square root Kalman filter (RRSQRT) (Verlaan and Heemink, 1997), are examples presented in this paper. Two alternative popular ensemble based approaches are the SEEK filter (Pham et al., 1997) and the SEIK filter (Pham et al., 1998). A recent review of ensemble based Kalman filters is provided in Evensen (2003). Another approach reducing the computational cost uses a simpler description of model dynamics. This can either be done by using a coarser grid for the error covariance modelling in the numerical model (Cohn and Todling, 1996 and Fukumori and Malanotte-Rizzoli, 1995), or by approximating time consuming elements of the numerical model, such as employing cheaper numerical schemes, simpler turbulence closure schemes or assuming geostrophic balance for the error covariance propagation (Dee, 1991).

A significant reduction in computational time can be obtained with the steady Kalman filter, where the model error covariance or the Kalman gain is assumed to be the same at each update time. Fukumori and Malanotte-Rizzoli (1995) derives such a steady gain from limiting theory solving the time invariant Riccati equation. Cañizares et al. (2001) also uses a steady approach, but here the steady gain is calculated as a time average of the EnKF. The steady approach generally reduces computational times with two orders of magnitude compared to the EnKF and is only slightly more computationally demanding than a single execution of a numerical model.

Extensions to the Kalman filter need to accommodate for non-linearities in the model propagation and the measurement equation. Also, bias or coloured noise in the numerical model and the measurements requires attention. Most schemes use a non-linear numerical model for the state propagation, while the forward operator employed for the error covariance propagation ranges from a steady linear operator (Fukumori and Malanotte-Rizzoli, 1995) to a linear expansion in extended Kalman filter applications such as the RRSQRT filter and a full non-linear error propagation in the EnKF.

While the handling and nature of non-linearities in a data assimilating system thus has been widely examined, the importance of using a proper error structure and robustness to error misspecification has gained only sporadic attention. The optimality of the Kalman filter assumes known and unbiased model and measurement errors. However, the estimation of these errors is to some extend subjective and can typically never be estimated from the limited data sets available. Further, structural model errors often lead to biased model states. Dee and da Silva (1998) present a scheme for the simultaneous estimation of the unbiased state and the model bias. Cañizares (1999) and Verlaan (1998) both uses a coloured noise implementation. Sørensen et al. (2004a) investigates the behaviour under misspecification of the model error in the case of a biased forcing. In all cases a clear improvement of the estimate results from correct error structure specification.

In a general data assimilation application the error sources are typically only known to a first or second order approximation and hence misspecification is part of the working conditions. However, good performance is nevertheless demonstrated in schemes, which do not explicitly account for the actual error structure, e.g. Madsen and Cañizares (1999). This must be accredited to a sufficient information content of the measurements and subsequent distribution. Bias is also corrected by a Kalman filter approach assuming no bias, albeit in a suboptimal way (Dee and da Silva, 1998). The specification of error structure and its subsequent propagation only need to provide a good interpolation of the innovation in space and time. Hence, when many data are available, the importance of a proper error model is reduced. In the case of assimilation of tidal gauge data, as considered herein, the measurements are usually sparsely distributed in space. Thus, the error structure provides the mean for updating state elements situated far from points of observation and hence its description becomes more important.

Focusing at the three state-of-the-art assimilation schemes, the EnKF, the RRSQRT filter and the Steady filter, with a coloured noise assumption implemented in a 3D hydrodynamic model, this papers will perform a sensitivity study of the schemes for various parameter settings. Acknowledging that misspecifications are often part of the working conditions such a study provides insight to the effect on performance of uncertain parameters. Hence, calibration can be focused at key parameters and in case of low sensitivity, confidence can be build in the schemes even for moderately misspecified parameters. Any such sensitivity study inherently only applies to the particular system under consideration and its characteristics.

Section 2 will introduce the building blocks of the assimilation approach, which provides the Kalman filter as a special case. The three schemes, which constitute the basis of this study will be described briefly—namely the EnKF, the RRSQRT and the steady Kalman filter. In Section 3 the filter parameters in the schemes are presented and discussed. In Section 4 results are presented for a range of sensitivity twin experiments using an idealised bay test case. Finally, Section 5 summarises and concludes the paper. The notation suggested by Ide et al. (1997) is used throughout.

2. Assimilation approach

It is important to stress that the three assimilation schemes to be presented in this section all rely on approximations, which are inevitably violated to some extend or another. The schemes and their foundation are given a fairly thorough treatment. This is partly done in order to highlight the

approximate nature of the schemes. Further, it simultaneously provides the setting for defining the assimilation parameters and testing the impact of these on the model performance.

The foundation of sequential estimation schemes is a linear model for combining the information contained in a model with measurements in an estimate of state variables. Hence, let $\mathbf{x}^t(t_i) \in \mathbb{R}^n$ be a representation of the true state at time t_i. This could be an array of grid averaged water levels and velocities at all model grid points in the area of interest. It can also contain additional augmented elements from an error model. Let $\mathbf{x}^f(t_i) \in \mathbb{R}^n$ be the model estimate of $\mathbf{x}^t(t_i)$ and $y_i^{\mathbf{o}} \in \mathbb{R}^p$ be a vector of observations at time t_i, which is assumed related to the state vector through the measurement equation,

$$\mathbf{y}_i^{\mathbf{o}} = \mathbf{H}_i \mathbf{x}^t(t_i) + \boldsymbol{\epsilon}_i \tag{1}$$

The operator $\mathbf{H}_i \in \mathbb{R}^{p\times n}$ projects the state space onto the measurement space.

The measurement noise is assumed additive and represented by the random variable, $\boldsymbol{\epsilon}_i \in \mathbb{R}^p$. The relation in (1) is assumed linear.

With the definitions given above and assuming both $\mathbf{x}^f(t_i)$ and $\mathbf{y}_i^{\mathbf{o}}$ to be unbiased, a linear unbiased estimate, $\mathbf{x}^a(t_i)$, of $\mathbf{x}^t(t_i)$ can be obtained as

$$\mathbf{x}^a(t_i) = \mathbf{x}^f(t_i) + \mathbf{K}_i(\mathbf{y}_i^{\mathbf{o}} - \mathbf{H}_i \mathbf{x}^f(t_i)) \tag{2}$$

A good sequential assimilation scheme is characterised by a proper estimation of the elements of the linear operator, $\mathbf{K}_i \in \mathbb{R}^{n\times p}$, which is denoted the Kalman gain. What is meant by proper depends on the application at hand and the properties of the estimate, $\mathbf{x}^a(t_i)$, sought for. Usually assumptions about linearity and unbiasedness are imposed and a least squares approach is taken. This is followed in the next section.

2.1. The BLUE estimator

The linear projection of $\mathbf{x}^t(t_i)$ on $\mathbf{y}_i^{\mathbf{o}}$ provides the best (minimum variance) linear unbiased estimate (BLUE) of $\mathbf{x}^t(t_i)$ (Jazwinski, 1970). The first moment of $\begin{bmatrix}\mathbf{x}^t(t_i)\\ \mathbf{y}_i^o\end{bmatrix}$ is given as $\begin{bmatrix}\mathbf{x}^f(t_i)\\ \mathbf{H}_i\mathbf{x}^f(t_i)\end{bmatrix}$ and its covariance matrix is $\begin{bmatrix}\mathbf{P}_i^f & \mathbf{P}_i^f\mathbf{H}_i^{\mathrm{T}}\\ \mathbf{H}_i\mathbf{P}_i^f & \mathbf{H}_i\mathbf{P}_i^f\mathbf{H}_i^{\mathrm{T}} + \mathbf{R}_i\end{bmatrix}$. Here, $\mathbf{R}_i \in \mathbb{R}^{p\times p}$ is the covariance of the noise process ϵ_i in (1) and $\mathbf{P}_i^f \in \mathbb{R}^{n\times n}$ is the error covariance of the state estimate $\mathbf{x}^f(t_i)$. The linear projection $\mathbb{E}(\mathbf{x}^t(t_i)|\mathbf{y}_i^{\mathbf{o}})$ is given by (2) with $\mathbf{K}_i$ given by

$$\mathbf{K}_i = \mathbf{P}_i^f\mathbf{H}_i^{\mathrm{T}}(\mathbf{H}_i\mathbf{P}_i^f\mathbf{H}_i^{\mathrm{T}} + \mathbf{R}_i)^{-1} \tag{3}$$

The error covariance of the estimated state, $\mathbf{x}^a(t_i)$, is given by

$$\mathbf{P}_i^a = \mathbf{P}_i^f - \mathbf{K}_i\mathbf{H}_i\mathbf{P}_i^f \tag{4}$$

The quest now becomes the estimation of $\mathbf{P}_i^f$ and $\mathbf{R}_i$. The measurement error is usually assumed constant in time and is prescribed according to measurement uncertainty and its representation in the model state. For sequential Kalman filter based algorithms, the discrepancies lies in the approximations made in the estimation of $\mathbf{P}_i^f$ and $\mathbf{x}^f(t_i)$. Three approaches are described in Section 2.2.

It should be noted that the BLUE estimator assumes that $\mathbb{E}(\mathbf{x}^t(t_i)|\mathbf{y}_i^{\mathbf{o}})$ is linear in $\mathbf{y}_i^{\mathbf{o}}$. This is true in the measurement point if the linear relation in (1) is valid. However, the state variables are

generally not linearly related and hence this assumption is not valid. The minimal variance property of the BLUE estimator and hence the Kalman filter only applies to the class of linear functions. Further, the numerical model does in most cases not provide an unbiased estimate of the true state. Since most sequential assimilation schemes employ this estimator, they are also subject to these sources of suboptimality.

2.2. State and error propagation

The basis of the predictions is a numerical hydrodynamic model. In this study the model adapted is MIKE 3, which is developed at DHI Water and Environment, DHI (2001). The numerical model solves for the three-dimensional non-linear oceanic flow. In the present ideal setup, constant density is applied. The code that constitutes a one-time-step-ahead prediction can be regarded as a model propagation operator, M^D_{M3}. With knowledge of the state at time t_{i-1} and the forcing, $\mathbf{u}(t_i)$, it provides the state at time t_i. The state considered consists of velocities and water levels on a specified grid. The forcing is open boundary water levels, sources and sinks, wind velocities and atmospheric pressure. Hence in a standard non-assimilating application of the numerical model, the one-time-step-ahead prediction can be written as,

$$\mathbf{x}^f(t_i) = M^D_{\mathrm{M3}}(\mathbf{x}^f(t_{i-1}), \mathbf{u}(t_i)) \tag{5}$$

In this case the state description and propagation are deterministic.

Acknowledging the approximate nature of a numerical model, a more detailed description must incorporate the model error introduced at each time step and its propagation throughout the system. Thus the propagation operator and the state become stochastic. Let r be the dimension of the space in which model propagation error sources are described. The error introduced by the model propagation, $\boldsymbol{\xi}_i \in \mathbb{R}^r_i$ is evident in the system equation,

$$\mathbf{x}^t(t_i) = M_{\mathrm{M3}}(\mathbf{x}^t(t_{i-1}),\ \mathbf{u}(t_i), \boldsymbol{\xi}_i) \tag{6}$$

The operator M_{M3} is a stochastic extension of M^D_{M3}. Optimally, knowledge of the correct time varying probability density function (pdf) of the noise sequence $\boldsymbol{\xi}_i$ and an initial state could be used to provide an exact stochastic forecast of the state. However, representation and propagation of the full pdf is not a tractable approach. A common approximation is to consider only the first and second order moments of the distribution. In case of Gaussian random fields and a linear model operator this further describes the full probability distribution. In addition, only these two moments are needed for the BLUE estimator in (2) and (3). Another approach is to approximate the propagation of the pdf to a precision, which allows a confident estimation of the first and second order moments. Two schemes for forecasting the first and second order moments in the hydrodynamic model and subsequent update of the state conditioned on observations are presented in Sections 2.2.1 and 2.2.2. A steady Kalman gain reduction of these elaborate techniques is presented in Section 2.2.3. For the two time varying schemes, the initial state error covariance is assumed to be zero. Hence, a short spin-up period is necessary for proper functioning of the time varying filters.

Of prime importance to any assimilation scheme is a proper model and measurement error description. As previously stated, the measurement error is usually based on simple assumptions. The specification of model error is a more difficult task. It is generally assumed proportional to the

model dynamics variability in some way or to originate solely from external forcing fields. The latter approach is applied here. Errors are introduced in the open boundary water level and in the wind velocity. Hence, the numerical model is in itself assumed to be perfect. In coastal and shelf seas this is often a good approximation to actual model inaccuracies, since the flow is strongly driven by these boundary conditions, which in many cases are known to contain rather large errors. A further step is taken by assuming the error to be coloured as described by a first order autoregressive model $M_{\mathrm{AR}(1)}$,

$$\boldsymbol{\xi}_i = M_{\mathrm{AR}(1)}(\boldsymbol{\xi}_{i-1}, \boldsymbol{\eta}_i) = \mathbf{A}\boldsymbol{\xi}_{i-1} + \boldsymbol{\eta}_i \tag{7}$$

The matrix $\mathbf{A} = \mathrm{diag}(\boldsymbol{\alpha})$, where the vector $\boldsymbol{\alpha}$ contains the coefficients of the autoregressive model, assuming the components of $\boldsymbol{\xi}_i$ to be uncorrelated. The noise process $\boldsymbol{\eta}_i$ is assumed Gaussian with zero mean and error covariance matrix, $\mathbf{Q}_i^{\eta} \in \mathbb{R}^{r\times r}$. Hence, $\mathbf{x}^t(t_i)$ is augmented with the boundary and wind error description and an extended operator, $M = (M_{\mathrm{M3}}, M_{\mathrm{AR}(1)})^{\mathrm{T}}$, is introduced. This leads to a system equation with additive noise, which will be used in the remainder of this work,

$$\mathbf{x}^t(t_i) = M(\mathbf{x}^t(t_{i-1}),\ \mathbf{u}(t_i), \boldsymbol{\eta}_i) = M(\mathbf{x}^t(t_{i-1}), \mathbf{u}(t_i), 0) + \begin{bmatrix} 0 \\ \boldsymbol{\eta}_i \end{bmatrix} \tag{8}$$

The error covariance of $\begin{bmatrix} 0 \\ \boldsymbol{\eta}_i \end{bmatrix}$ is $\mathbf{Q}_i = \begin{bmatrix} 0 & 0 \\ 0 & \mathbf{Q}_i^{\eta} \end{bmatrix}$. Running the hydrodynamic model alone with error propagation according to one of the schemes described below, but with no assimilation, will yield a model mean estimate and its error covariance matrix. Hence, the accuracy of model results based on the given assumptions can be addressed and compared to observational evidence.

2.2.1. The ensemble Kalman filter

In the ensemble Kalman filter (EnKF) the propagation of the full pdf is approximated by an ensemble propagation. Both the first and second order moments are calculated as ensemble statistics and used for the update of each ensemble member. The strength of the approach lies in its representation of the full pdf, its handling of non-linearities and its ease of implementation for complex state and error descriptions.

An ensemble of q state realisations is defined at an initial time. In this work, a single model initial state defines all ensembles with zero spread at a pre-initial point of time as the starting conditions of a spin-up period. During this period the forcing errors are propagated throughout the system to provide the initial model error covariance matrix and mean state estimate.

All ensemble members are propagated according to the model operator in (8),

$$\mathbf{x}_j^f(t_i) = M(\mathbf{x}_j^a(t_{i-1}), \mathbf{u}(t_i), \boldsymbol{\eta}_{j,i}), \quad j = 1, \ldots, q \tag{9}$$

The state estimate $\mathbf{x}_j^a(t_{i-1})$ is the update from the previous time step. If no new data were available for update then $\mathbf{x}_j^a(t_{i-1}) = \mathbf{x}_j^f(t_{i-1})$. The model error, $\boldsymbol{\eta}_{j,i}$ is randomly drawn from a predefined Gaussian distribution with zero mean and covariance, $\mathbf{Q}_i$. With each ensemble member propagated by (9), the mean state estimate and model error covariance estimate are provided by the following equations:

$$\hat{\mathbf{x}}^f(t_i) = \frac{1}{q}\sum_{j=1}^{q} \mathbf{x}_j^f(t_i) \tag{10}$$

$$\mathbf{P}_i^f = \mathbf{S}_i^f (\mathbf{S}_i^f)^{\mathrm{T}}, \quad \mathbf{s}_{j,i}^f = \frac{1}{\sqrt{q-1}} (\mathbf{x}_j^f(t_i) - \hat{\mathbf{x}}^f(t_i)) \tag{11}$$

The vector, $\mathbf{s}_{j,i}^{f \in \mathbb{R}^n}$, is the j'th column of $\mathbf{S}_i^f \in \mathbb{R}^{n \times q}$. The update can be performed by (2) and (3), when given the proper interpretation in an ensemble setting. For computational efficiency an algebraically equivalent set of equations are used.

Each ensemble member must be updated rather than the ensemble state estimate, in order to maintain correct statistical properties of the updated ensemble. For the same reason an ensemble of measurements must be generated and used for each ensemble member update accordingly rather than the measurement itself (Burgers et al. (1998)). Hence,

$$\mathbf{y}_{j,i}^o = \mathbf{y}_i^o + \boldsymbol{\epsilon}_{j,i}, \quad j = 1, \ldots, q \tag{12}$$

Randomly generated realisations, $\boldsymbol{\epsilon}_{j,i}$, of $\boldsymbol{\epsilon}_i$ are drawn from a Gaussian distribution with zero mean and error covariance, $\mathbf{R}_i$, and added for each member. The update scheme presented here specifically uses that measurement errors are uncorrelated to assimilate simultaneous measurements sequentially. The updating algorithm for every ensemble member, j, reads, Chui and Chen (1991),

$$x_{j,m}^a(t_i) = x_{j,m-1}^a(t_i) + \mathbf{k}_{i,m}(y_{j,i,m}^o - \mathbf{h}_{i,m} x_{j,m-1}^a(t_i)), \quad m = 1, \ldots, p \tag{13}$$

and $x_{j,0}^a(t_i) = x_j^f(t_i)$. In (13) $y_{j,i,m}^o$ is the m'th element in $y_{j,i}^o$ and $\mathbf{h}_{i,m}$ is the m'th row of $\mathbf{H}_i$. Treating one measurement at a time the Kalman gain is a vector, $\mathbf{k}_{i,m}$, given by,

$$\mathbf{k}_{i,m} = \frac{\mathbf{S}_{i,m-1}^a \mathbf{c}_{i,m}}{\mathbf{c}_{i,m}^{\mathrm{T}} \mathbf{c}_{i,m} + \sigma_{i,m}^2}, \quad \mathbf{c}_{i,m} = (\mathbf{S}_{i,m-1}^a)^{\mathrm{T}} \mathbf{h}_{i,m}^{\mathrm{T}} \tag{14}$$

The m'th diagonal element in $\mathbf{R}_i$ is denoted $\sigma_{i,m}^2$. The matrix $\mathbf{S}_{i,m}^a$ in (14) is calculated as

$$\mathbf{S}_{i,m}^a = [\mathbf{s}_{1,i,m}^a \ldots \mathbf{s}_{q,i,m}^a], \quad \mathbf{s}_{j,i,m}^a = \frac{1}{\sqrt{q-1}} (\mathbf{x}_{j,m}^a(t_i) - \hat{\mathbf{x}}_m^a(t_i)) \tag{15}$$

for $m = 1, \ldots, p$ and $\mathbf{S}_{i,0}^a = \mathbf{S}_i^f$. Now, (13), (14) and (15) provides the update equations of all ensemble members, one measurement at a time.

2.2.2. The reduced rank square root Kalman filter

The reduced rank square root Kalman filter (RRSQRT) is based on the extended Kalman filter formalism, in which the error propagation is calculated using a statistical linearisation of the model propagation operator. It further uses a square root algorithm and a lower rank approximation of the error covariance matrix. Thus, it handles weak non-linearities and it has a concise and smooth representation of the error covariance matrix.

The state propagation is the model forecast of the central estimate,

$$\mathbf{x}^f(t_i) = M(\mathbf{x}^a(t_{i-1}), \mathbf{u}(t_i), 0) \tag{16}$$

The error covariance propagation basically performs the following truncated Taylor Series approximation,

$$\mathbf{P}^f(t_i) = \mathbf{M}_i \mathbf{P}^a(t_{i-1}) \mathbf{M}_i^{\mathrm{T}} + \mathbf{Q}_i, \quad \mathbf{M}_i = \left. \frac{\partial M}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}^f(t_i), \mathbf{u}=\mathbf{u}(t_i), \eta=0} \tag{17}$$

A square root implementation of this propagation and subsequent update has been performed. Denote by $\mathbf{S}^a(t_{i-1})$ the approximation of rank q of the square root of the error covariance matrix $P^a(t_{t-1})$. The propagation of the error covariance matrix approximated according to (17), is then given by,

$$\mathbf{S}^f(t_i) = \left[\mathbf{M}_i\mathbf{S}^a(t_{i-1})\middle|\mathbf{Q}_i^{1/2}\right] \tag{18}$$

$\mathbf{Q}_i^{1/2}$ is the square root approximation of rank r of $\mathbf{Q}_i$. Since the rank of $\mathbf{Q}_i$ is equal to the rank of $\mathbf{Q}_i^{\eta}$, which is r, this corresponds to the square root approximation of $\mathbf{Q}_i^{\eta}$. To calculate the derivatives needed in $\mathbf{M}_i$ a finite difference approximation of M is column-wise adopted as follows,

$$(\mathbf{M}_i\mathbf{S}^a(t_{i-1}))_j = \frac{M(\mathbf{x}^a(t_{i-1}) + \delta s^a_{j,i-1}, u(t_i), 0) - M(\mathbf{x}^a(t_{i-1}), u(t_i), 0)}{\delta} \tag{19}$$

The value of the parameter δ has been discussed in Segers et al. (2000) and is set equal to one according to their recommendation. The propagation step in (18) increases the number of columns in the error covariance matrix from q to $q + r$. Thus a complimentary part of the scheme must provide a mean for reducing the rank of the space similarly. In order to do this a lower rank approximation of $\mathbf{S}^f(t_i)$ in (18) is applied through an eigenvalue decomposition of the matrix $(S^f(t_i))^{\mathrm{T}}\mathbf{S}^f(t_i)$ (Verlaan, 1998). This approach provides efficient calculations, but introduces the need for normalisation in a multivariate setting. The optimal normalisation is application dependent and is here approximated by using normalisation based on energy consideration. Basically, the contribution from potential and kinetic energy to each element of $(\mathbf{S}^f(t_i))^{\mathrm{T}}\mathbf{S}^f(t_i)$ are equal.

Now, having calculated the forecast state and error covariance, the algorithms developed for the EnKF can be followed to provide the update. The state is updated using (13). However, an additional ingredient is needed, namely the update of the error covariance estimate, when the state is updated. This is provided by Cañizares (1999),

$$\mathbf{S}^a_{i,m} = \mathbf{S}^a_{i,m-1} - \frac{\mathbf{k}_{i,m}\mathbf{c}^{\mathrm{T}}_{i,m}}{1 + \sqrt{\frac{\sigma^2_{i,m}}{\mathbf{h}^{\mathrm{T}}_{i,m}\mathbf{h}_{i,m} + \sigma^2_{i,m}}}} \tag{20}$$

The vectors, $\mathbf{k}_{i,m}$ and $\mathbf{h}_{i,m}$ are defined in (14).

2.2.3. The steady Kalman filter

For the Ensemble Kalman filter and the reduced rank square root Kalman filter, the error covariance propagation typically takes the order of 10^2 model executions, Cañizares (1999), which may be too many in operational settings. A well known work around on this is to assume time invariant model and measurement error covariance matrices, $\mathbf{P}_i^f$ and $\mathbf{R}_i$, rendering a time constant Kalman gain, $\mathbf{K}$. However, it can still be difficult to estimate $\mathbf{K}$ without the help from more elaborate methods.

When time invariance is approximately true, both the EnKF and the RRSQRT can provide robust estimates of the gain. Hence a Kalman gain that is still based on model dynamics can be obtained as a time average of the gain from one of these two elaborate methods, Cañizares et al. (2001). The update is still done using (13), but now with a fixed, $\mathbf{k}$, and operating on just a single state forecast.

3. Filter parameters

Now, with the assimilation schemes presented, focus can be moved to the essence of this contribution, the assimilation parameters. In an actual implementation of the filters above, several parameters need to be specified. These mainly relate to the model and measurement error covariance description. This section describes each parameter, while the sensitivity to parameter variations is tested in Section 4.2.

Rank: q. The rank of the model error covariance matrix is essential to the performance of an assimilation scheme. For the ensemble Kalman filter this is equal to the ensemble size, while for the RRSQRT Kalman filter it is the number of leading eigenvalues preserved in the covariance reduction. In either case the rank needs to be large enough to describe the error covariance field with sufficient accuracy, but with the trade-off of increased computational time.

Measurement standard deviation: σ_m. This study only considers uncorrelated water level measurements. Hence, the error specification is simply given by a value of σ_m for each measurement. Both instrumentation error and the error due to lack of representation of state variables need to be taken into account in the specification.

Model error standard deviation: σ_η. The model error specification has a more complex description. The model errors are assumed to be due to an error in the open boundary water level and/or in the wind forcing. The assumed standard deviation, σ_η, of the white noise process, η_i, is naturally a key parameter. An independent set of parameters are specified for each boundary and wind velocities. The relative sizes of model and measurement uncertainty basically determines which source of information that ought to be trusted the most in the state estimation.

Temporal correlation scale: τ. The temporal correlation scale defines the coefficient, α, in the AR(1) process by giving the half time τ of the exponential process.

$$\alpha = 0.5^{(\frac{\Delta t}{\tau})} \tag{21}$$

where Δ_t is the model time step. Note that since the noise enters into an autoregressive process the actual standard deviation, σ_{forcing} of the boundary or wind forcing is given by,

$$\sigma^2_{\text{forcing}} = \frac{\sigma^2_\eta}{1-\alpha^2} \tag{22}$$

A higher temporal correlation allows a more distant effect of an error in the forcing. The formulation ensures that the standard deviation can be specified from (21) and (22) to construct a forcing perturbation that is independent of time step length, Δt, and changes in τ.

If $\alpha = 1$ and σ_n is zero with a suitable initial covariance, then the Kalman filter provides a bias estimate of the forcing terms. A τ close to one and a moderate σ_η approximates this bias estimation, hence enabling the filters to detect a slowly varying bias in the external forcing.

Spatial correlation length: l_c. An exponential correlation model is employed in the definition of the error covariance model, $\mathbf{Q}^\eta$. The spatial correlation length of the model errors plays a key role in defining the correlation structure in the model that ultimately determines the update. A too large spatial correlation scale assumption in the wind velocities can cause an update in data sparse regions based on a measurement, which does not contain any information about this distant area in the real system. On the other hand, a too small spatial correlation scale underestimates the correlation in the model errors and thereby provides a filtered estimate, which is too close to

the model solution. It also increases the effective number of degrees of freedom in the error model, which in turn makes the estimated parameters of the error model more uncertain.

Grid factor: g. The grid factor is introduced as an ad hoc approach to reducing the dimension of the error space. This integer factor simply expresses the number of model grid points in between each error point. The errors are subsequently redistributed using a kriging technique. Such a space reduction is viable because the spatial correlation length often is considerably larger than the grid spacing. However, when the spatial correlation length approaches the distance between error points, this assumption is violated and the spatial correlation length looses its interpretation. Other space reduction techniques, e.g. EOF decomposition, can also be cast in the present framework.

Smoothing factor: s. The steady Kalman filter uses a time average from one of the more elaborate schemes for the generation of the Kalman gain. A greater deal of smoothness can be obtained in the time varying Kalman filter schemes as well through the introduction of an exponential smoothing factor. This number is the proportion of weight given to the Kalman gain calculated at the present time, $\mathbf{K}_i^{\mathrm{KF}}$. The applied gain matrix then becomes

$$\mathbf{K}_i = (1-s)\mathbf{K}_{i-1} + s\mathbf{K}_i^{\mathrm{KF}},\ s \in [0;1] \tag{23}$$

Update interval: d. In practical applications it must be considered whether all data shall be assimilated. Trying to drive the model into an observed regime with effects not represented by the model, can introduce noise and ultimately cause model instabilities. In practice, water levels may typically be provided at half hourly intervals, but the implementation of the data assimilation schemes interpolates the measurements to every model time step and assimilates it. Hence, an update interval is introduced as the final parameter to test the effect of using various subsets of the interpolated water level measurements.

4. Idealised bay experiment

For the purpose of investigating filter performance, when basic assumptions are violated, an idealised, controllable and stable setup was chosen. This choice also facilitates a large number of sensitivity runs to be performed and a comparison to the full true field to be done.

4.1. Setup and basic results

The region under consideration is the hypothetical Ideal Bay situated at 51°N. It is a 200 km by 200 km square bay with an open Northern boundary and simple bathymetry with a maximum depth of 100 m as shown in Fig. 1. The vertical grid spacing is 10 m and the horizontal resolution is 10 km.

Density is constant in this study, which is conducted over a 48 h period. The open Northern boundary is forced with a spatially constant water level signal with a sinusoidal variation in time. The period is 12 h and the amplitude is 1 m. The model is further forced by an artificially generated passing cyclone, which moves 50 km across the bay every 6 h. It has a maximum wind speed of 26 m/s.

The basic solution has a main flow, which is dominated by a Kelvin wave moving cyclonically in the bay. This is superposed with a wind generated flow.

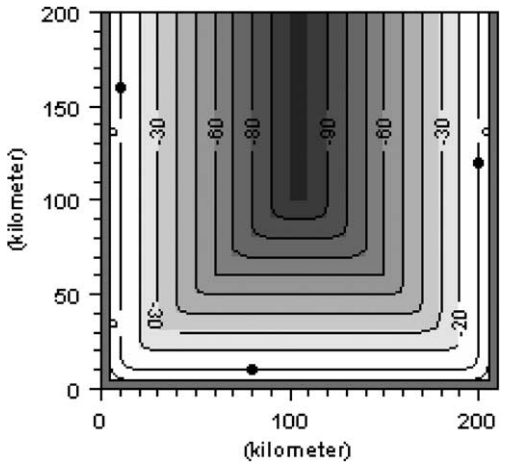

Fig. 1. Bathymetry of Ideal Bay. The three dots indicate the positions, where water level time series were extracted to be used as measurements.

4.2. Parameter sensitivity and robustness assessment

The performance of an assimilation scheme should be examined under ideal conditions as well as under conditions, where assumptions that are typically violated break down. This section does both, but with the focus aimed at the nature of the solution in the latter case. The ideal setup provides knowledge of the exact truth and control of the model errors, hence allowing the choice of assimilation parameters to match or violate the errors and perform a sensitivity study accordingly.

A large number of twin experiments has been performed. In each case the basic run from Section 4.1 is taken as the true state of the system. Water levels are extracted at three locations indicated in Fig. 1 at every 15 min timestep interval. Subsequently uncorrelated Gaussian white noise with a standard deviation of 5 cm is added to each time series in order to represent measurement noise. Only these three time series provide information about the true state in the assimilation procedures. In each perturbed run a different error source is introduced, i.e. the model is run with different open boundaries, wind forcing and/or parameters than the basic run. The ability of the assimilation scheme to correct the flow in the perturbed run to that of the basic flow is examined.

As a measure of the filter performance, the spatially averaged root mean square error ($\overline{\mathrm{RMSE}}$) of water levels, l, calculated over the last 24 h is used,

$$\overline{\mathrm{RMSE}} = \frac{1}{J}\frac{1}{K}\sum_{j=1}^{J}\sum_{k=1}^{K}\sqrt{\frac{1}{I}\sum_{i=1}^{I}\left(l_{\mathrm{true}}(x_j,y_k,t_i) - l_{\mathrm{pert}}(x_j,y_k,t_i)\right)^2} \tag{24}$$

The constants $J, K = (20, 21)$ are the number of grid points in the x and y direction respectively. The constant $I = 96$ is the number of time steps in the period, in which the statistic is calculated. The indices true and pert refer to results from the true and the perturbed experiment respectively. The Kalman filter schemes provide estimates of the state error covariance, $\mathbf{P}_i^a$. Hence, diag($\mathbf{P}_i^a$) is

the model state variance. An indicative diagnostic for the performance of an assimilation scheme is the spatially and temporally averaged standard deviation of water level, Std. Dev., calculated by the scheme.

The error types in the perturbed runs are divided into two groups: Gaussian errors in Section 4.2.1 and typical errors in Section 4.2.2. Gaussian errors refer to error structures that basically fulfill the assumptions of the assimilation schemes if the parameters are chosen correctly. This is where the actual parameter sensitivity study is performed. The typical errors, on the other hand, refer to errors that include other distributions and sources than those assumed in the filters. This latter case is thought to closer resemble a real filter application.

4.2.1. Gaussian errors

The model error assumption lies in two forcing terms: The boundary water level and the wind velocity components. Thus, the investigation and the presentation of the results are divided according to this division. All the results presented in this section attempt to correct the same two perturbed runs, where a random coloured noise realisation has been added to the boundary forcing and wind field respectively. The parameters used to generate these Gaussian perturbations are stated in Tables 1 and 2.

Comparisons between the perturbed runs without data assimilation and the true runs give $\overline{\mathrm{RMSE}}$ values of 0.33 m and 0.58 m for boundary and wind errors, respectively. The maps of the RMSE values for each case is shown in Figs. 2 and 3. Using the correct parameters in the assimilation schemes, but varying the rank of the covariance matrix, yields the $\overline{\mathrm{RMSE}}$ results displayed in Figs. 4 and 5. Since the EnKF uses randomly generated noise, an average over five runs were used to smooth out the worst stochastic variations in the statistics. The assimilation schemes simultaneously provide the standard deviation they use for the update. The spatial averages of these are also included in the figures. A previous study by Madsen and Cañizares (1999) has shown that the RRSQRT filter converges to a good performance at a lower rank than the EnKF, but similar execution times gave similar performance. This is also evident here for the wind error

Table 1
Characteristics of the imposed errors in the open boundary perturbed runs

Boundary spatial correlation scale	100 km
Boundary grid factor	1
Boundary temporal correlation scale	2 h ($\alpha = 0.92$)
Boundary Std. Dev.	0.10 m
Boundary actual Std. Dev.	0.25 m

Table 2
Characteristics of the imposed errors in the wind velocity perturbed runs

Wind spatial correlation scale	300 km
Wind grid factor	3
Wind temporal correlation scale	6 h ($\alpha = 0.97$)
Wind Std. Dev.	3 m/s
Wind actual Std. Dev.	13 m/s

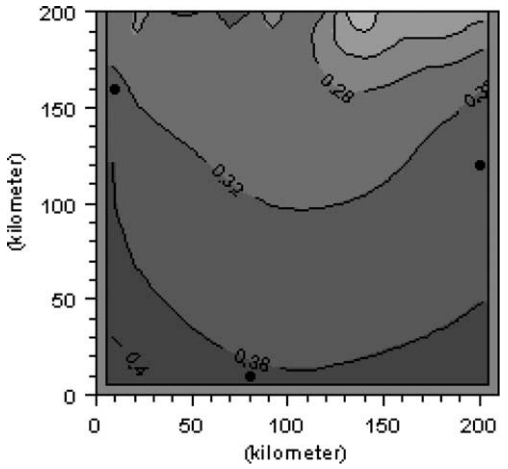

Fig. 2. Spatial distribution of RMSE values between the true run and the false boundary run, where a realisation of the Gaussian process described by Table 1 has been added. The positions of the measurements are indicated by dots.

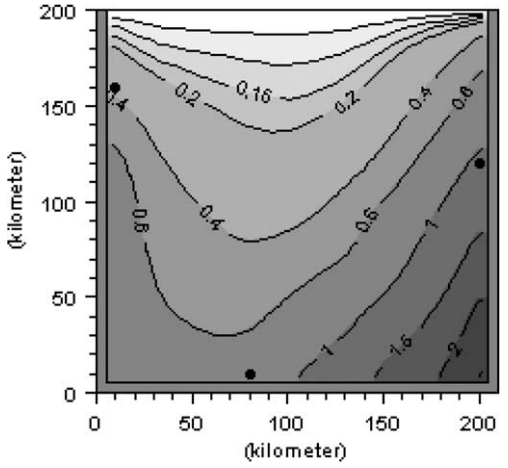

Fig. 3. Spatial distribution of RMSE values between the true run and the false wind run, where a realisation of the Gaussian process described by Table 2 has been added. The positions of the measurements are indicated by dots.

case, but in the boundary error case, low rank EnKF outperforms the RRSQRT. Further investigation shows that this can be accredited to the normalisation procedure required in the eigenvalue decomposition of the RRSQRT scheme and hence a tuning of the normalisation makes the RRSQRT converge faster for the boundary error case as well. In general the figures show good performance with respect to the reduction of overall prediction error as well as the estimation of the prediction error.

Also included in the Figs. 4 and 5 are the results of the steady Kalman filter using a constant Kalman gain obtained as the average gain estimated by the EnKF over the last 24 h. For a low rank of the error covariance the steady filter out-performs the EnKF. This is an example of

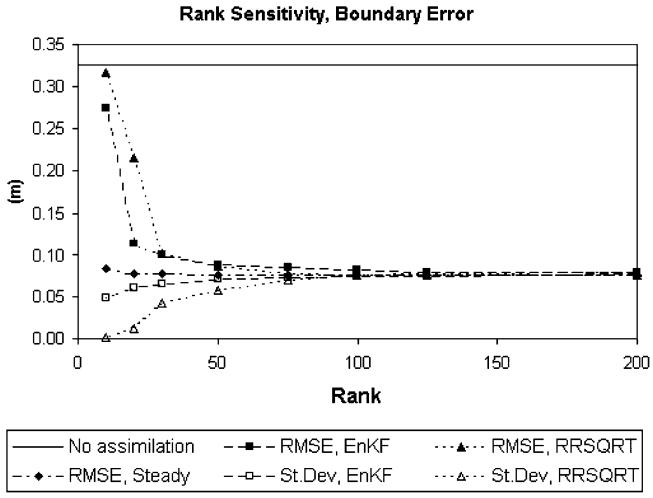

Fig. 4. Sensitivity to rank of error covariance for perturbed boundary runs.

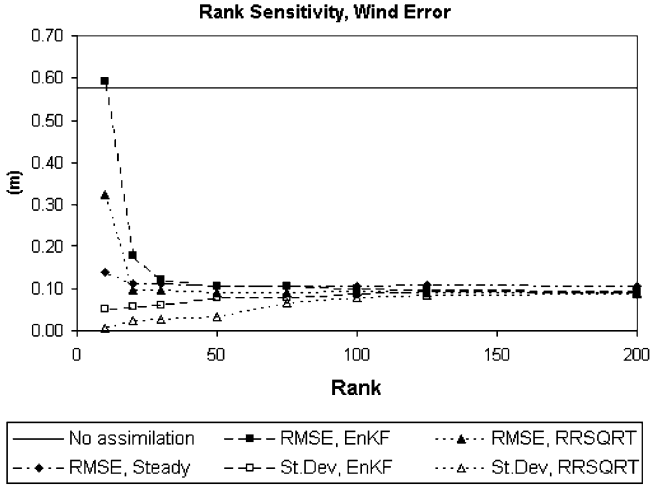

Fig. 5. Sensitivity to rank of error covariance for perturbed wind velocity runs.

bias-variance trade-off. The EnKF attempts to estimate the best unbiased state. However, this is done at the price of a high variance of the estimated parameters in the Kalman gain. The time averaging operation employed for the generation of the time constant gain used in the steady Kalman filter reduces this variance considerably, but deliberately introduces a bias in the estimation by assuming the gain to be time-invariant. For few ensembles, the latter trade-off between bias and variance gives the better performance. However, the Steady filter still requires the execution of one of the more elaborate schemes to construct the steady Kalman gain to be used. Hence, the good performance states that (once computed) it is not worth the effort to calculate the time variability of the error covariance. This has a great impact on filtering efficiency for water level assimilation in

coastal and shelf seas. A more rigorous discussion on this matter is provided in Sørensen et al. (2004b), which takes a regularisation perspective on the estimation of the Kalman gain elements.

As the rank is increased the ensemble based estimate become more and more skillful and the performance gets better. For the boundary error case the estimated gain is actually rather time invariant and hence, even when the EnKF and RRSQRT filter are converged their performance is matched by that of the cheaper Steady filter. The converged gain of the wind error case is more time varying. Thus, in this case the steady Kalman filter performs a little worse than the two time varying filters. Fig. 6 shows six Kalman gain segments corresponding to each of the three measurement positions (10 km, 160 km), (80 km, 10 km) and (200 km, 120 km) in the boundary and wind error case respectively. The structure of the error correlations expressed in the Kalman gain evidently reflects the nature of the imposed error assumption. The gain structures corresponding to assuming a wind error shows a strong correlation to the closed boundaries, where the strongest wind driven signal is present. Similarly the gain structures corresponding to the open boundary water level error assumption reflects the correlation pattern of a cyclonically propagating Kelvin wave. This is evident in the gain segment corresponding to (10 km, 160 km), in which the faster wave speed at the deeper central waters also is seen. The size of the gain is also significantly larger in this point, due to its dynamical proximity to the error source.

In the subsequent experiments, an ensemble size of 100 is used for the EnKF and a rank of 50 is used for the RRSQRT scheme. The steady Kalman filter will be based on the EnKF with 100 ensemble members. For these choices the EnKF gives a $\overline{\text{RMSE}}$ of 0.08 m and 0.10 m and RRSQRT 0.09 m and 0.09 m for boundary and wind error, respectively. Similarly the steady Kalman filter gives a RMSE of 0.08 m and 0.10 m. Maps of the RMSE for the EnKF case are shown in Figs. 7 and 8. Compared to Figs. 2 and 3 these demonstrate the good performance of the filter. The results also show the importance of a good network design for assimilation purposes. For instance, in the boundary error case the Northwestern measurement corrects most of the noise in the Kelvin wave, leaving little error in its cyclonic propagation further South in the bay. Experiments show that leaving out the Southern measurement hardly alters the performance of the schemes.

In all subsequent sensitivity runs, the errors broadly speaking look like Figs. 7 and 8, but with larger amplitudes and some spatial variation. In order to treat the large number of sensitivity runs, only the RMSE average performance is used hereafter.

First the sensitivity to the assumed measurement standard deviation is examined. For each of the three measurements extracted from the basic solution a measurement standard deviation must be specified. The same value is used for all three stations. Table 3 summarises the $\overline{\text{RMSE}}$ results. The most notable result is the robustness of the filters for varying measurement standard deviation. The general picture is a degradation both when the solution is pulled too strongly toward the measurement, where the innovation has an excessive impact on unobserved regions, and when little trust is put in the measurements leading to only minor corrections of the perturbed solution. However, in both directions, extreme and unrealistic values must be assumed to significantly degrade the results. For a measurement standard deviation of 0.005 m the wind error case gives no results. The wind field is updated to unrealistic values because the measurement is trusted so much and the model blows up (indicated by the "–").

The sensitivity to model standard deviation is summarised in Tables 4 and 5 as $\overline{\text{RMSE}}$ values. Here too, a quite robust performance is achieved. In general the behaviour degrades as the values

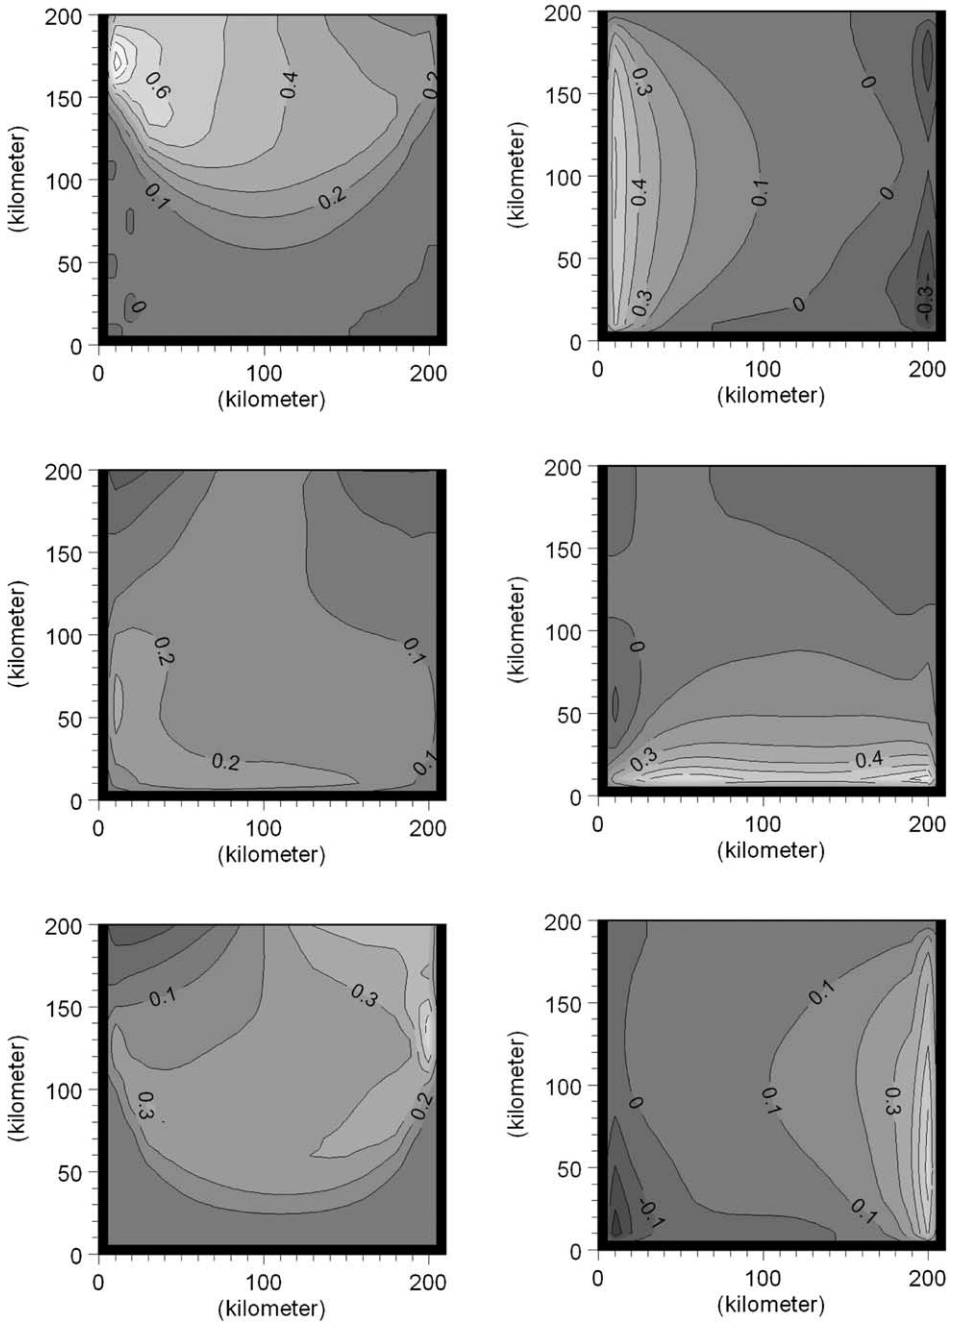

Fig. 6. Kalman gain segments corresponding to each of the three measurement positions (10 km, 160 km) (top), (80 km, 10 km) (center) and (200 km, 120 km) (bottom). The gain from the run with boundary errors are displayed to the left and the gain from the run with wind errors are displayed to the right. The gains are derived as time averages from an Ensemble Kalman filter run with 100 ensembles.

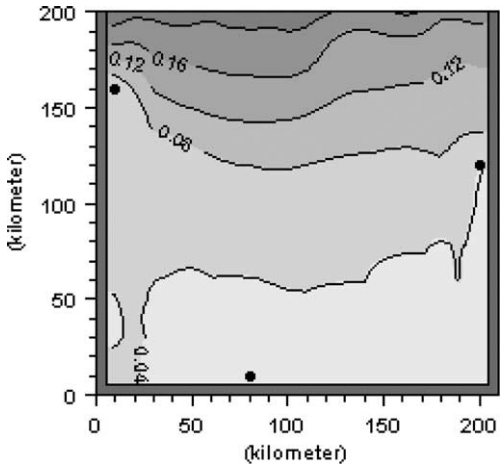

Fig. 7. Spatial distribution of RMSE values between the true run and the 100 EnKF run for the boundary error case. The positions of the measurements are indicated by dots.

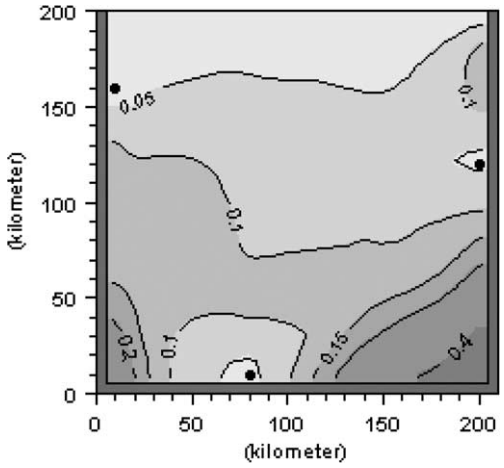

Fig. 8. Spatial distribution of RMSE values between the true run and the 100 EnKF run for the wind error case. The positions of the measurements are indicated by dots.

become too large or too small. An interesting point is that the RRSQRT scheme has its best performance for slightly overestimated model errors. This is explained by the standard deviations estimated by the RRSQRT scheme itself. As shown in Figs. 4 and 5 the RRSQRT in general tends to underestimate the size of the model error and hence it provides a more correct standard deviation estimate with a larger model error assumption. When the wind error standard deviation becomes too high, wind fields, whose solution is not supported by the model can be obtained, hence giving a model blow-up (indicated by the "–").

Now, the sensitivity to the temporal correlation scale is examined. A temporal correlation scale of 0 h gives a white noise assumption and a temporal correlation of ∞ hours gives a random walk

Table 3
Sensitivity to measurement standard deviation for perturbed boundary runs

σ_m	Bound			Wind		
	EnKF	RRSQRT	Steady	EnKF	RRSQRT	Steady
0.005	0.26	0.13	0.14	–	0.14	–
0.01	0.12	0.09	0.10	0.24	0.11	0.11
0.02	0.09	0.08	0.08	0.13	0.10	0.10
0.05	**0.08**	**0.09**	**0.08**	**0.09**	**0.09**	**0.10**
0.15	0.09	0.10	0.09	0.10	0.10	0.11
0.40	0.12	0.12	0.12	0.13	0.12	0.13
1.00	0.18	0.19	0.18	0.19	0.18	0.21

Reference run is marked in bold. All results are $\overline{\text{RMSE}}$ and given in meters (m).

Table 4
Sensitivity to model error standard deviation for perturbed boundary runs

σ_n	EnKF	RRSQRT	Steady
0.001	0.31	0.31	0.31
0.005	0.20	0.19	0.19
0.01	0.13	0.13	0.13
0.05	0.09	0.09	0.08
0.10	**0.08**	**0.09**	**0.08**
0.25	0.10	0.08	0.08
1.50	0.13	0.09	0.10

Reference run is marked in bold. All results are $\overline{\text{RMSE}}$ and given in meters (m).

Table 5
Sensitivity to model error standard deviation for perturbed wind runs

σ_η	EnKF	RRSQRT	Steady
0.05	0.32	0.32	0.36
0.10	0.23	0.22	0.26
0.50	0.12	0.11	0.13
1.00	0.11	0.10	0.11
3.00	**0.09**	**0.09**	**0.10**
5.00	0.11	0.09	0.11
10.00	–	0.10	–

Reference run is marked in bold. All results are $\overline{\text{RMSE}}$ and given in meters (m).

error process. These extremes corresponds to values of σ_{forcing} equal to σ_η and ∞ respectively. Tables 6 and 7 summarise the results. Here too, quite robust performance can be observed. The results are most degraded when there is a small time correlation. The explanation for this is a combination of the smaller resulting standard deviation of the forcing terms and a worse description of the spatial correlations due to the approximate white noise assumption. When coloured noise is adopted, the dynamical propagation transfers the coloured signal into spatial correlations.

Table 6
Sensitivity to model error temporal correlation scale, τ, for perturbed boundary runs

τ (h)	EnKF	RRSQRT	Steady
0.00	0.15	0.15	0.14
0.25	0.12	0.12	0.11
2.00	**0.08**	**0.09**	**0.08**
6.00	0.09	0.09	0.08
24.00	0.10	0.09	0.09

Reference run is marked in bold. All results are $\overline{\mathrm{RMSE}}$ and given in meters (m).

Table 7
Sensitivity to model error temporal correlation scale, τ, for perturbed wind runs

τ (h)	EnKF	RRSQRT	Steady
0.00	0.18	0.15	0.20
0.25	0.16	0.13	0.16
1.00	0.12	0.10	0.12
6.00	**0.09**	**0.09**	**0.10**
24.00	0.13	0.10	0.10

Reference run is marked in bold. All results are $\overline{\mathrm{RMSE}}$ and given in meters (m).

Table 8
Sensitivity to model error spatial correlation scale, l_c, for perturbed boundary runs

l_c (km)	EnKF	RRSQRT	Steady
0	0.11	0.15	0.09
25	0.09	0.08	0.08
100	**0.08**	**0.09**	**0.08**
250	0.08	0.09	0.08
1000	0.09	0.10	0.08

Reference run is marked in bold. All results are $\overline{\mathrm{RMSE}}$ and given in meters (m).

Table 9
Sensitivity to model error spatial correlation scale, l_c, for perturbed wind runs

l_c (km)	EnKF	RRSQRT	Steady
30	0.19	0.18	0.20
100	0.11	0.09	0.13
300	**0.09**	**0.09**	**0.10**
1000	0.08	0.10	0.10

Reference run is marked in bold. All results are $\overline{\mathrm{RMSE}}$ and given in meters (m).

This dependence on a proper spatial correlation is also evident when considering the sensitivity to the spatial correlation scale (Tables 8 and 9). In line with the discussion above about temporal correlations, the case of no spatial correlation has a poorer performance, which shows the

Table 10
Sensitivity to update interval for perturbed runs

d	Bound			Wind		
	EnKF	RRSQRT	Steady	EnKF	RRSQRT	Steady
1	**0.08**	**0.09**	**0.08**	**0.09**	**0.09**	**0.10**
2	0.09	0.10	0.09	0.10	0.09	0.11
4	0.11	0.12	0.11	0.11	0.10	0.12
8	0.16	0.16	0.15	0.17	0.15	0.16
16	0.22	0.21	0.21	0.27	0.25	0.29
32	0.26	0.25	0.26	0.42	0.33	0.38

Reference run is marked in bold. All results are $\overline{\text{RMSE}}$ and given in meters (m).

importance of properly describing the correlations in the state vector. This can be an important factor to control in real setups of the assimilation schemes. The results are in general insensitive to grid factor variations as long as the resulting coarse error grid resolves the assumed spatial error covariance. For the basic spatial correlation lengths of 100km and 300km for the open boundary and wind error respectively, no changes in the RMSE-measure results by changing the grid factor.

The smoothing factor shows hardly any sensitivity at all. This can be accredited to the similar performance of the time varying and the steady schemes. In cases, with more time varying correlation structures, this factor must be expected to play a greater role, providing a smooth transition from the time varying to the steady performance. The smoothing factor opens a possibility of obtaining stable time varying runs of e.g. a low rank EnKF. If the time variation of the error covariance fields are slow, but important to resolve, then this approach might provide an important operational option.

Altering the update interval degrades the result shown in Table 10. This proves a continuous transition to the runs with no assimilation, which corresponds to an update interval equal to ∞. Using less information from the true state gives a lower resemblance with the truth.

All together, in the present test case the assimilation schemes are robust to moderately misspecified parameters for Gaussian error sources that resembles the specified error models. This is encouraging, but does not guarantee good performance for any setup. In particular, care must be taken to ensure a proper model error covariance in sparsely observed systems.

4.2.2. Typical errors

The demonstrated robustness in the Gaussian error case gives some hope that even for more typical error sources, not elaborately taken into account by the schemes, an improved performance can be obtained using good first guess estimates of parameters in the Gaussian framework assumed. This section investigates such behaviour.

Perturbed runs with ten different errors have been conducted. The results of the twin experiments both without and with the three data assimilation schemes are summarised in Table 11. The runs with 'Bound:' and 'Wind:' indicates that only open boundary or wind errors was added and subsequently assumed in the assimilation procedures. The 'Bound:' runs apply one and 3 h phase errors as well as half a meter amplitude error by itself and in combination with the 3 h phase error. The two systematic wind error runs are forced by 20 m/s Westerly winds and a stronger cyclone with a perturbed path referred to as 'False cyclone'. Further, a run was conducted with

Table 11
Sensitivity to typical errors for perturbed runs

Error type	No ass.	EnKF	RRSQRT	Steady
Bound: 1 h phase lag	0.69	0.07	0.07	0.07
Bound: 3 h phase lag†	1.87	0.18	0.18	0.20
Bound: 1.5 × amplitude†	0.62	0.07	0.07	0.07
Bound: 3 h phase lag + 1.5 × amplitude	2.23	0.23	0.22	0.26
Wind: 20 m/s West	0.17	0.07	0.09	0.09
Wind: false cyclone†	0.25	0.12	0.12	0.12
Bed friction 0.5†	0.26	0.06	0.06	0.13
Bathymetry†	0.05	0.04	0.04	0.03
All errors with †'s	2.53	0.51	0.44	0.38
No forcing	1.33	0.19	0.18	0.20

All results are $\overline{\mathrm{RMSE}}$ and given in meters (m).

an erroneous bed friction using a Nikuradse roughness coefficient of 0.5 instead of the true 0.05. The 'Bathymetry' perturbed run refers to a run applying a modified bathymetry with a standard deviation of 1 m compared to the truth. Also, a run, 'All errors with †'s', applying a composite of these errors is included in the study. Finally, a 'No forcing' run was conducted with no wind and open boundary forcing at all, thus giving a solution at rest. In all assimilation runs, the parameters of Tables 1 and 2 were used in the assimilation schemes.

A much improved performance is observed in all cases. In the worst case the residual $\overline{\mathrm{RMSE}}$ is 0.51, which is too large for many applications, but a convincing result considering the size of the errors introduced into the model. Even in the 'No forcing' run, the schemes are actually able to generate a large portion of the signal from the missing forcing terms including the forcing terms themselves through the augmented state vector description. Most of the boundary signal is regenerated by the measurement at (10 km, 160 km) alone. Hence, external phenomena such as surges can be effectively modelled by assimilating measurements close to the open boundaries.

However, the good performance is not matched by a good uncertainty estimate. For instance, both time varying filters estimate an average standard deviation of 0.07 m for the 3h phase lag experiment. The standard deviation estimates for the 'False cyclone' run were 0.06 m and 0.07 m for the EnKF and RRSQRT schemes, respectively, while the numbers for the 'All errors with †'s' are 0.13 m and 0.14 m. This clearly shows the violation of the underlying filter assumptions and the error estimates provided by the filters must be applied with care. At least, data must be retained in an attempt to perform a subsequent validation of the standard deviation estimates.

The biased nature of the errors makes the performance even more dependent on fairly high temporal and spatial correlations than was the case in the previous section. Hence, these are key parameters to consider in the calibration of any data assimilation setup. In general the Steady filter performs well with the 'All errors with †'s' case being the most impressive case. Basically, when all the error assumptions are violated, the elaborate schemes can not be expected to give superior performance. Rather, a certain regularisation of the Kalman gain acknowledges the bias in the estimate hence allowing a reduced variance and in this case, a better performance. However, time dependent schemes are clearly superior for bed friction error by itself. This study demonstrates the success of assimilation schemes despite the unavoidable wrong error assumptions imposed, and also shows how different filters handle different real error sources the best.

5. Summary and conclusions

This paper has presented three known assimilation schemes and described the filter parameters that can typically be varied in an application of the schemes. In a set of experiments in an idealised bay a sensitivity study has been conducted to investigate the filter performance for misspecified error structure in the schemes. The sensitivity to key parameters is vital for the practical use of sequential data assimilation techniques in hydrodynamic modelling. It is demonstrated that the filter performance is robust with respect to low to moderate parameter perturbations in the specification of the noise statistics. For more typical errors such as phase lags, bathymetry, etc., care must be taken to ensure a specification of fairly high temporal and spatial correlations. However, thought should be put into properly setting up every application.

In general the steady filter seems like a good candidate for tide gauge assimilation in coastal areas. This is particularly true if an operational setting is considered. The spatial distribution of the filter performance has further demonstrated that proximity to the stations or dynamically well chosen positions enhances the estimation skill. Hence, the denser and better designed the measurement network is, the better the overall performance. Such a design increases the representative information available. Further, dense networks diminishes the importance of the spatial distribution of the information and thus the correct parameter settings.

References

Burgers, G., van Leeuwen, P.J., Evensen, G., 1998. Analysis scheme in the ensemble Kalman filter. Monthly Weather Review 126, 1719–1724.

Cañizares, R., 1999. On the application of data assimilation in regional coastal models, Ph.D. thesis, Delft University of Technology.

Cañizares, R., Madsen, H., Jensen, H.R., Vested, H.J., 2001. Developments in operational shelf sea modelling in Danish waters. Estuarine, Coastal and Shelf Science 53, 595–605.

Chui, C.K., Chen, G., 1991. Kalman filter with real-time applicationsSpringer Series in Information Sciences, 17. Springer-Verlag.

Cohn, S.E., Todling, R., 1996. Approximate data assimilation schemes for stable and unstable dynamics. Journal of Meteorological Society of Japan 74, 63–75.

Dee, D.P., 1991. Simplification of the Kalman filter for meteorological data assimilation. Quarterly Journal of the Royal Meteorological Society 117, 365–384.

Dee, D.P., da Silva, A.M., 1998. Data assimilation in the presence of forecast bias. Quarterly Journal of the Royal Meteorological Society 124, 269–296.

DHI: 2001. MIKE 3 estuarine and coastal hydrodynamics and oceanography. DHI Water and Environment.

Evensen, G., 1994. Sequential data assimilation with a nonlinear quasigeostrophic model using Monte Carlo methods to forecast error statistics. Journal of Geophysical Research 99 (C5), 10143–10162.

Evensen, G., 2003. The ensemble Kalman filter: Theoretical formulation and practical implementation. Ocean Dynamics 53, 343–367.

Fukumori, I., Malanotte-Rizzoli, P., 1995. An approximate Kalman filter for ocean data assimilation; An example with an idealised Gulf Stream model. Journal of Geophysical Research 100 (C4), 6777–6793.

Ide, K., Courtier, P., Ghil, M., Lorenc, A.C., 1997. Unified notation for data assimilation: Operational, sequential and variational. Journal of Meteorological Society of Japan 75 (1B), 181–189.

Jazwinski, A.H., 1970. Stochastic processes and filtering theoryMathematics in Science and Engineering, 64. Academic Press.

Kalman, R.E., 1960. A new approach to linear filter and prediction theory. Journal of Basic Engineering 82 (D), 35–45.

Madsen, H., Cañizares, R., 1999. Comparison of extended and ensemble Kalman filters for data assimilation in coastal area modelling. International Journal of Numerical Methods in Fluids 31 (6), 961–981.

Pham, D.T., Verron, J., Gourdeau, L., 1998. A singular evolutive Kalman filter for data assimilation in oceanography. Comptes Rendus de l' Academie des Sciences, Paris 326, 255–260.

Pham, D.T., Verron, J., Roubaud, M.C., 1997. Singular evolutive Kalman filter with EOF initialization for data assimilation in oceanography. Journal of Marine Systems 16, 323–340.

Segers, A.J., Heemink, A.W., Verlaan, M., van Loon, M., 2000, Kalman filtering for nonlinear atmospheric chemistry models: second (order) experiences, Technical report, Delft University of Technology.

Sørensen, J.V.T., Madsen, H., Madsen, H., 2004a. Data assimilation in hydrodynamic modelling: On the treatment of nonlinearity and bias. Stochastic Environmental Research and Risk Assessment 18 (4), 228–244.

Sørensen, J.V.T., Madsen, H., Madsen, H., 2004b. Efficient Kalman filter techniques for the assimilation of tide gauge data in three-dimensional modeling of the North sea and Baltic sea system. Journal of Geophysical Research 109 (C3), C03017.

Verlaan, M., 1998, Efficient Kalman filtering algorithms for hydrodynamic models, Ph.D. thesis, Delft University of Technology.

Verlaan, M., Heemink, A.W., 1997. Tidal flow forecasting using reduced rank square root filters. Stochastic Hydrology and Hydraulics 11, 349–368.