University of Reading

Observation impact, domain length and parameter estimation in data assimilation for flood forecasting

Article

Accepted Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Cooper, E. S., Dance, S. L. ORCID: https://orcid.org/0000-0003-1690-3338, Garcia-Pintado, J., Nichols, N. K. ORCID: https://orcid.org/0000-0003-1133-5220 and Smith, P. J. ORCID: https://orcid.org/0000-0003-4570-4127 (2018) Observation impact, domain length and parameter estimation in data assimilation for flood forecasting. Environmental Modelling and Software, 104. pp. 199-214. ISSN 1364-8152 doi: https://doi.org/10.1016/j.envsoft.2018.03.013 Available at https://centaur.reading.ac.uk/76153/

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

To link to this article DOI: http://dx.doi.org/10.1016/j.envsoft.2018.03.013

Publisher: Elsevier

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the End User Agreement.

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Observation impact, domain length and parameter estimation in data assimilation for flood forecasting.

E.S. Cooper[1], S.L. Dance[1,2], J. Garcia-Pintado[3], N.K. Nichols[1,2], P.J. Smith[2]

[1]Department of Meteorology, University of Reading, [2]Department of Mathematics and Statistics, University of Reading, [3]MARUM Center for Marine environmental Sciences and Department of Geosciences, University of Bremen.

Abstract

Accurate inundation forecasting provides vital information about the behaviour of fluvial flood water. Using data assimilation with an Ensemble Transform Kalman Filter we combine forecasts from a numerical hydrodynamic model with synthetic observations of water levels. We show that reinitialising the model with corrected water levels can cause an initialization shock and demonstrate a simple novel solution. In agreement with others, we find that although assimilation can accurately correct water levels at observation times, the corrected forecast quickly relaxes to the open loop forecast. Our new work shows that the time taken for the forecast to relax to the open loop case depends on domain length; observation impact is longer-lived in a longer domain. We demonstrate that jointly correcting the channel friction parameter as well as water levels greatly improves the forecast. We also show that updating the value of the channel friction parameter can compensate for bias in inflow.

Keywords Data assimilation, inundation forecasting, fluvial flooding, observation impact, joint state-parameter estimation, ensemble Kalman filter.

Highlights

- Data assimilation is applied to simulated flood forecasts and SAR-like observations
- Reinitialisation shock due to water level correction is removed using a novel method
- Observation impact is linked to domain length when updating only water levels
- Updating the channel friction parameter leads to marked improvement in forecast skill
- Updating the channel friction parameter can compensate for biased inflow

Software Availability

The inundation simulations in this work were generated using Clawpack 5.2.2, a collection of FORTRAN and python code available from http://www.clawpack.org/. Details of the amended Clawpack source code as used in this work are freely available on request from the corresponding author, as is the python code used to perform data assimilation on the inundation simulation output. Please contact e.s.cooper@pgr.reading.ac.uk for details.

1 Introduction

Data assimilation can improve the accuracy of predictions from flood inundation models by combining forecasts from the model with observations of the system, taking into account uncertainty in both the model predictions and the observations. In this study we use a sequential data assimilation method comprising a forecast-update dynamic feedback loop. During each forecast step, the numerical model runs an inundation simulation. When an observation (or set of observations) is available the simulation is interrupted and the update step is performed; updating combines observational data and model predictions to give a better estimate of the state. The next forecast step then starts, with the adjusted water levels as the initial condition. An update is carried out each time a new observation or set of observations is available.

There are a number of numerical inundation models that can predict the behaviour of flood water given information about the topography of the domain and the amount of water flowing into the area, e.g. HEC-RAS, Telemac, LISFLOOD-FP (HEC-RAS Development Team; Hervouet, 2000; Neal et al., 2012). In a real flood situation, topographical information is often available in the form of a digital terrain model (DTM) and inflow estimates may come from an upstream gauge, or as output from a hydrological model. Observations of the flood may be available from a variety of different sources. These include river depth and flow rate measurements from gauges, and authors have used these data in assimilation schemes , e.g. Mure-Ravaud et al. (2016). However, many catchments are ungauged and the number of gauges worldwide is in decline (Vrsmarty et al., 2001). Observations of flood extent can be obtained from aerial photos, although the cloudy conditions associated with heavy rainfall often limit the usefulness of this information source. Recently, much attention has been paid to the use of synthetic aperture radar (SAR) satellite images in delineating flood extent, since such observing systems have all weather and day and night capability. Water depth information can then be retrieved from SAR satellite images using a high quality digital terrain model (DTM) as described in Mason et al. (2012) and Brown et al. (2016). Such techniques for extracting information from SAR images are well established, e.g.Thornhill et al. (2012), Mason et al. (2010), Scott et al. (2008) and Scott and Mason (2007).

Various authors e.g. Lai and Monnier (2009), Matgen et al. (2007) and Schumann et al. (2009) have used data assimilation techniques to highlight the fact that although observations from SAR can cover a large spatial area, the usefulness of the information they contain is limited in time. Assimilating data from one or more river gauges can help to mitigate this, as shown by Lai and Monnier (2009) and Hostache et al. (2010), but we consider here the situation in which only time-sparse satellite derived water level data is available for assimilation. This leads to a situation in which data assimilation can provide a good analysis - i.e. can correct water levels very well at the time of observations, but the model forecast then moves quickly away from the true water levels during the subsequent forecast step. This short lived improvement in the water levels has been shown in studies such as Andreadis et al. (2007), Neal et al. (2009) and Garcia-Pintado et al. (2013), which use ensemble Kalman Filter data assimilation algorithms, as well as in Matgen et al. (2010), where a particle filter approach is taken. This result indicates that water levels in a river flood situation are not strongly sensitive to

initial conditions. In fact, water levels are likely to be more dependent on inflow and model parameter values, and updating one or a combination of these is therefore necessary.

In order to address the short-lived nature of the forecast improvement, authors such as Andreadis et al. (2007), Matgen et al. (2010), Giustarini et al. (2011), Garcia-Pintado et al. (2015), Garcia-Pintado et al. (2013) and Mason et al. (2015) have carried out data assimilation including on-line correction of inflow along with water levels. Inflow correction is shown in all of these studies to give much better forecast accuracy over time than correcting water levels alone. Less attention has been paid to the effect of errors in model parameters in sequential data assimilation, despite the fact that several studies, including Andreadis and Schumann (2014) and the comprehensive review paper by Grimaldi et al. (2016), indicate that model parameters are likely to have an important influence on the behaviour of the flow.

One study in which parameter effects are investigated is Garcia-Pintado et al. (2015), in which water levels, inflows and several model parameter values were updated simultaneously using an ensemble Kalman filter technique. The study used LISFLOOD-FP to model the flooding of the river Severn and tributaries near Tewksbury, UK, in 2014, assimilating real SAR-derived water level observations. A large improvement in forecast skill was seen when inflow was corrected along with water levels, leading to good agreement between the forecast and independently measured gauge data. In this case, channel friction parameter estimation alongside estimation of water levels, inflows and other parameters was not found to improve the forecast significantly, despite the fact that water behaviour is strongly influenced by this parameter. The question of whether the retrieved friction parameter value was correct was left open as the true value for the system was not known.

In this study we address open questions about the role of the channel friction parameter in data assimilation for inundation modelling. We use a similar data assimilation technique to that in Garcia-Pintado et al. (2015) in twin experiments with an idealised topography and an unbiased inflow. This allows us to separate out and further investigate the effect of channel friction retrieval on the forecast. We find that, in contrast with Garcia-Pintado et al. (2015), online estimation of the channel friction parameter along with water levels leads to a large improvement in the forecast skill in our experiments. The twin experiments also show that our data assimilation method is capable of finding an accurate value for the channel friction parameter, even when water depth observations are only available on the flood plain during a flood.

We also investigate the effect of domain length on forecast skill, showing that because the assimilation is able to correct water levels in areas where there are no observations, the time taken for corrected water levels to decay back to the open loop (no assimilation) case is longer for a physically longer domain. Further, we demonstrate that when reinitialising the numerical model after an assimilation, an initialisation shock can occur. We demonstrate an efficient and effective technique for removing this shock, leading to more accurate forecasts in the hours immediately following an assimilation.

This paper is organised as follows: In section 2 the numerical inundation model is described, the data assimilation method is outlined and our novel re-initialisation method is demonstrated. In section 3 the experimental configurations for various simulations are described. Section 4 shows the effect of including online

channel friction parameter estimation along with water level estimation, and compares results from different length domains. Section 5 draws conclusions about the effects of domain length and channel friction parameter estimation.

2 Methodology

In this section we describe the methods used in this study. In section 2.1 the numerical inundation model is outlined. Section 2.2 contains information about the data assimilation method used. In section 2.3 we discuss the impact of assuming the water has only hydrostatic momentum at the start of a forecast step and describe our approach to dealing with problems caused by this assumption.

2.1 Numerical inundation model

In this study we use a numerical flood model we have developed using Clawpack (Clawpack Development Team, 2014; Mandli et al., 2016; LeVeque, 2002), an open source collection of FORTRAN and python code that can be used to solve a wide variety of conservation laws. Clawpack uses finite volume methods and sophisticated Riemann solvers to treat systems of partial differential equations; in this work the equations of interest are the 2D shallow water equations that describe how river and flood water will move in space and time. The model splits the domain of interest into N cells and calculates the water depth in each cell. The code is capable of dealing with shocks in the solution, such as bores that may occur following a sudden increase of inflow into a particular river stretch. Clawpack deals effectively with the wet-dry interfaces which are present in an inundation event, and preserves depth non-negativity (George, 2008).

The shallow water equations for two spatial dimensions, x and y, can be written as (e.g.LeVeque (2002))

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{F(q)}}{\partial x} + \frac{\partial \mathbf{G(q)}}{\partial y} = \mathbf{R}(\mathbf{q}), \tag{1}$$

where **R(q)** is a source term and **q** is a vector of conserved quantities

$$\mathbf{q} = \begin{bmatrix} h \\ hu \\ hv \end{bmatrix}, \tag{2}$$

h represents depth of the fluid, and u and v represent velocity in the x and y directions respectively.

In equation (1), **F(q)** and **G(q)** represent fluxes of the conserved quantities in the x and y directions respectively. For the shallow water equations these are

$$\mathbf{F(q)} = \begin{bmatrix} hu \\ hu^2 + \frac{1}{2}gh^2 \\ huv \end{bmatrix} \quad \text{and} \quad \mathbf{G(q)} = \begin{bmatrix} hv \\ huv \\ hv^2 + \frac{1}{2}gh^2 \end{bmatrix}, \tag{3}$$

where g is acceleration due to gravity.

The effect of friction is modelled as a source term in Clawpack, since the friction force acts to reduce the momentum of the water. The magnitude of the momentum reduction is strongly dependent on a Manning's friction coefficient, n, and the flow of water is very sensitive to the value of this parameter. We have also added an inflow source term to the Clawpack code to model water arriving in the river stretch of interest, as further described in Appendix A and Cooper et al. (2013). The time step for the hydrodynamic model is variable, and automatically adjusted in the code to preserve numerical stability.

Correct specification of the solution at the boundaries of the computational domain is vital for the stability of any numerical scheme. To achieve this, Clawpack adds a user-specified number of 'ghost' cells (2 by default) next to each cell at a domain boundary. The domain is effectively extended in all directions by the addition of these ghost cells and the behaviour of the solution at the boundaries then depends strongly on the values of calculated model quantities in the ghost cells. We use non-reflecting outflow (extrapolating) boundary conditions in which values of $\mathbf{q}$ are extrapolated from the cell next to the boundary into the ghost cells at each time step. This is called a zero order extrapolation in LeVeque (2002).

Another important factor is the representation of the domain topography at and across the domain boundaries. The default in the code is to copy the value of the domain elevation at the boundary into the ghost cells. This represents a situation where there is no slope in bathymetry or topography across any boundaries. This is not suitable for the downstream boundary in our experiments, at which the majority of the water leaves the domain. A more physically realistic situation for the downstream boundary is to extrapolate the slope of the domain into the ghost cells at the boundary and changes have been made to the code to accommodate this.

2.2 Data assimilation

2.2.1 State estimation

In data assimilation, a state vector is used to represent the state of a physical system. In this work the state vector, $\mathbf{x} \in \mathbb{R}^N$, comprises water depths in each of N computational cells. Sequential data assimilation algorithms comprise two steps: a forecast (or prediction) step and an update (or analysis) step. In the prediction step, an estimate of the state, $\mathbf{x}$ is evolved forward in time using the forecast model

$$\mathbf{x}(t_{k+1}) = \mathfrak{M}(\mathbf{x}(t_k)), \tag{4}$$

where $\mathfrak{M}$ is the forecast model, in this case the non-linear numerical shallow water equation model described in section 2.1. In the update step the forecast is updated to take account of observations of the state. We assume that the observations can be described by

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \epsilon, \tag{5}$$

where $\mathbf{y} \in \mathbb{R}^p$ is a vector of observations and $\mathbf{x}$ is the true state of the system. Since the observations may be indirect and not located at model cell centres, an observation operator, $\mathbf{H} : \mathbb{R}^N \rightarrow \mathbb{R}^p$ is required, which maps the state vector into observation space. For this work, $\mathbf{H}$ is assumed to be a linear operator. The observation

noise, ϵ is assumed to be unbiased stochastic noise with covariance $\mathbf{R} \in \mathbb{R}^{p\times p}$. The ensemble Kalman filter used here is based on the Kalman filter. In the Kalman filter, whenever observations are available the state and error covariance matrix are updated Kalman (1960) according to

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{K}(\mathbf{y}_{obs} - \mathbf{H}\mathbf{x}^f) \tag{6}$$

and

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}^f, \tag{7}$$

where forecast and analysis quantities are denoted by f and a superscripts respectively, $\mathbf{I} \in \mathbb{R}^{N\times N}$ is the identity matrix and $\mathbf{P} \in \mathbb{R}^{N\times N}$ is the state error covariance matrix. The matrix $\mathbf{K} \in \mathbb{R}^{N\times p}$ is the Kalman gain, given by

$$\mathbf{K} = \mathbf{P}^f\mathbf{H}^T(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}, \tag{8}$$

and $\mathbf{R} \in \mathbb{R}^{p\times p}$ is the observation error covariance matrix.

In the ensemble Kalman filter (Evensen, 1994), an ensemble of state vectors is used to represent a statistical sample of the forecast or analysis uncertainty. Each ensemble member represents one possible realisation of the true state of the system, given uncertainties in initial conditions and/or model parameters. For an ensemble comprising M state vectors, $\mathbf{x}_i$, $(i = 1, 2...M)$, a mean state at any time can be defined as

$$\overline{\mathbf{x}} = \frac{1}{M}\sum_{i=1}^{M}\mathbf{x}_i. \tag{9}$$

The mean of the ensemble, $\overline{\mathbf{x}}$, represents an estimate of the true state of the system.

For any ensemble, an ensemble perturbation matrix $\mathbf{X} \in \mathbb{R}^{N\times M}$ can be defined as

$$\mathbf{X} = \frac{1}{\sqrt{M-1}}(\mathbf{x}_1 - \overline{\mathbf{x}}\ \ \mathbf{x}_2 - \overline{\mathbf{x}}\ \ .....\ \ \mathbf{x}_M - \overline{\mathbf{x}}). \tag{10}$$

The ensemble error covariance matrix, $\mathbf{P} \in \mathbb{R}^{N\times N}$ can then be calculated from

$$\mathbf{P} = \mathbf{X}(\mathbf{X})^T. \tag{11}$$

The forecast step for an ensemble system requires each state vector in the ensemble to be evolved by the forecast model according to equation (4). In the update step the forecast ensemble is combined with observations of the state to produce a 'corrected' ensemble of state vectors called the analysis ensemble. The analysis ensemble is then used as a set of initial conditions for the next forecast step. This forecast-update cycle can be repeated many times and an analysis ensemble calculated whenever observations of the system are available. The ensemble update equations separately update the ensemble mean and the ensemble perturbations according to

$$\overline{\mathbf{x}^a} = \overline{\mathbf{x}^f} + \mathbf{K}(\mathbf{y} - \overline{\mathbf{H}\mathbf{x}^f}) \tag{12}$$

and

$$\mathbf{X}^a = \mathbf{X}^f \mathbf{T}. \tag{13}$$

The vector $\overline{\mathbf{x}^a}$ is the analysis state (the mean of the analysis ensemble), $\overline{\mathbf{x}^f}$ is the mean of the forecast ensemble and $\mathbf{K} \in \mathbb{R}^{N \times p}$ is an ensemble version of the Kalman gain (as shown in equation (16)). The matrix $\mathbf{T} \in \mathbb{R}^{M \times M}$ updates the perturbations such that the state error covariance calculated by using $\mathbf{X}^a$ in equation (11) matches that given by the Kalman covariance update, equation (7) (Kalman, 1960). There is not a unique solution for $\mathbf{T}$; here we use an unbiased, symmetric square root formulation known as an ensemble transform Kalman Filter (ETKF), following the approach of Livings et al. (2008), Livings (2005) and Garcia-Pintado et al. (2013). In this approach we define a forecast observation ensemble comprising M forecast observation vectors, $\mathbf{y}_i^f$, $(i = 1, 2...M)$ such that

$$\mathbf{y}_i^f = \mathbf{H}(\mathbf{x}_i^f). \tag{14}$$

The forecast observation ensemble has a mean, and a perturbation matrix $\mathbf{Y}^f$, defined in the same way as for the state ensemble matrix.

We define a matrix $\mathbf{D}$ as

$$\mathbf{D} = \mathbf{Y}^f (\mathbf{Y}^f)^T + \mathbf{R}; \tag{15}$$

the Kalman gain $\mathbf{K}$ can then be written in terms of the forecast perturbation matrices $\mathbf{X}^f$ and $\mathbf{Y}^f$,

$$\mathbf{K} = \mathbf{X}^f (\mathbf{Y}^f)^T \mathbf{D}^{-1}. \tag{16}$$

Substituting equation (16) for $\mathbf{K}$ on the right hand side of equation (7), and using equations (13) and (11) on the left hand side shows that the matrix $\mathbf{T}$ in equation (13) then needs to satisfy

$$\mathbf{T}(\mathbf{T})^T = \mathbf{I} - (\mathbf{Y}^f)^T \mathbf{D}^{-1} (\mathbf{Y}^f). \tag{17}$$

Using the Sherman-Woodbury-Morisson identity for the right hand side of equation (17), as in equation (15) of Tippett et al. (2003), this becomes

$$\mathbf{T}(\mathbf{T})^T = (\mathbf{I} + (\mathbf{Y}^f)^T \mathbf{R}^{-1} \mathbf{Y}^f)^{-1}. \tag{18}$$

A scaled forecast observation ensemble perturbation matrix can then be introduced,

$$\widehat{\mathbf{Y}^f} = \mathbf{R}^{-\frac{1}{2}} \mathbf{Y}^f. \tag{19}$$

Performing a singular value decomposition (Golub and Van Loan, 1996) of $(\widehat{\mathbf{Y}^f})^T$ gives a factorisation such that

$$(\widehat{\mathbf{Y}^f})^T = \mathbf{U} \Sigma \mathbf{V}^T, \tag{20}$$

where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices with dimensions (M by M) and (p by p) respectively. The columns of $\mathbf{U}$ and $\mathbf{V}$ are the left and right singular vectors of $(\widehat{\mathbf{Y}}^f)^T$ respectively, and the diagonal elements of the (M by p) matrix Σ are the singular values of $(\widehat{\mathbf{Y}}^f)^T$. Combining equations (18), (19) and (20) gives

$$\mathbf{T}(\mathbf{T})^T = \mathbf{U}(\mathbf{I} + \Sigma \Sigma^T)^{-1} \mathbf{U}^T, \tag{21}$$

and a solution for $\mathbf{T}$ is therefore

$$\mathbf{T} = \mathbf{U}(\mathbf{I} + \Sigma\Sigma^T)^{-\frac{1}{2}}\mathbf{U}^T. \tag{22}$$

This is the solution used in this work.

2.2.2 Joint state-parameter estimation

Section 2.2.1 describes how the ETKF can be used to update the water levels in a computational domain, given observational data. It is also possible to update values of uncertain forecast model parameters as part of the same process. This is achieved by state augmentation, in which parameters are appended to the state vector (Smith et al., 2013; Navon, 1998; Evensen et al., 1998; Smith et al., 2009, 2011). The augmented state vector, $\mathbf{x}_{aug}$, is then given by

$$\mathbf{x}_{aug} = \begin{bmatrix} \mathbf{x} \\ \mathbf{b} \end{bmatrix}, \tag{23}$$

where $\mathbf{b} \in \mathbb{R}^m$ is a vector of m parameters and $\mathbf{x}_{aug} \in \mathbb{R}^{N+m}$. In this work, we are interested in updating just one parameter, the Manning's friction coefficient in the river channel, n_{ch}. This means that $\mathbf{b}$ is scalar in this case.

We assume that the value of n_{ch} does not change with time over the course of a particular flood. This means that the value of n_{ch} is constant during the forecast step and only updates at assimilation times. The forecast equation for the augmented state vector is then given by

$$\mathbf{x}_{aug}(t_{k+1}) = \begin{bmatrix} \mathfrak{M}(\mathbf{x}(t_k)) \\ \mathbf{b}(t_k) \end{bmatrix}, \tag{24}$$

where $\mathfrak{M}$ is the forecast model as in equation (4).

The ETKF update equations (12) and (13) can be applied to the augmented state vector in the same way as described in section 2.2.1. The ensemble assimilation scheme then takes into account covariances between errors in the state vectors and the parameter(s). These covariances act to correct the parameter value according to information from observations as part of the same process that corrects water levels in the domain.

Estimating parameter values in this way has a number of advantages over a more traditional offline calibration approach. Firstly, the updating of the parameter values is performed with information from current observations. Calibrating parameters with data from previous events risks using out of date information which does not take into account changes to the river bed due to, for example, erosion or sediment transport. Calculating parameters using data assimilation also allows the value to change on shorter timescales during a flood event as the value is assumed to be constant during forecast steps, but updates each time new observational information is available. Additionally, off-line calibration and tuning of parameters can be computationally costly and needs to be performed ahead of a flooding event which may occur with little warning.

In our synthetic experiments we have assumed that the channel friction parameter, n_{ch} is the same value for the whole channel. In a real setting, it would be necessary to take an approach like that in Garcia-Pintado

et al. (2015), where different values of n_{ch} were assumed for different types of channel such as large rivers and small streams.

2.3 Hydrostatic initialisation shock

The ETKF is a sequential method and this means that each time observations are available an assimilation can be carried out. In order to perform an assimilation, the forecast model is interrupted. The ensemble states (water levels) are updated by the ETKF and each ensemble member then restarts running in Clawpack, reinitialised with the new water levels. In some approaches to data assimilation for inundation modelling, e.g. Lai and Monnier (2009); Hostache et al. (2010); Ricci et al. (2011), the state vector contains information about water flow or discharge rates as well as water levels. The flow rates are therefore updated along with water levels as part of the assimilation process. In contrast, in this study, as in Garcia-Pintado et al. (2015), the state vector contains only water depth information (plus parameters when considering the joint state-parameter estimation problem).

In Garcia-Pintado et al. (2015), a hydrostatic assumption was made for reinitialisation, i.e. the water was assumed to have zero momentum at each forecast restart time. The effect of this assumption in our domain can be seen by interrupting a simulation and restarting without performing any data assimilation, i.e. reinitialising with identical water levels as before, but with zero momentum everywhere. Comparing the root mean square error (RMSE) between the water depths predicted by a reinitialised version and continuous version of the simulation then gives a measure of how the hydrostatic assumption at restart affects the forecast. The RMSEs are measured over the whole domain and defined as

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(h_i^r - h_i^c)^2}, \tag{25}$$

where h_i^r is the forecast water depth after reinitilaisation in the ith cell and h_i^c is forecast water depth in the same cell without reinitialisation. The number of cells in the domain is N, as before.

The open circles in figure 1 show RMSEs between the reinitialised forecast and the continuous forecast. Figure 1 shows that the consequence of using a hydrostatic assumption is that the error between the continuous and restarted cases is large at times less than approximately four hours in this system. This means that forecasting the behaviour of flood water at these times is problematic. The error becomes negligible by approximately four hours after the assimilation time.

In order to correct for this without adding flow information to the state vector, we assume that the water in each cell has the same velocity (u and v in the x and y directions respectively) after the assimilation as it did before the assimilation. This gives a state vector of lower dimension (approximately one third as many entries) than the approach of Lai and Monnier (2009), Hostache et al. (2010) and Ricci et al. (2011), thus reducing computational expense, while avoiding problems caused by assuming zero flow rates at the start of each forecast step as in Garcia-Pintado et al. (2015). Reinitialising with the same flow rate (hu and hv) values in each cell

would also be possible, but since the water levels will have changed in some cells due to the assimilation, this is likely to lead to unrealistic behaviour. Reinitialising the water with the pre-analysis velocity values removes the large shock shown in figure 1; the resulting RMSE values are shown in figure 1 with filled circles. The very small RMSE values shown by the filled circles show that the method is effective in removing the initialisation shock.

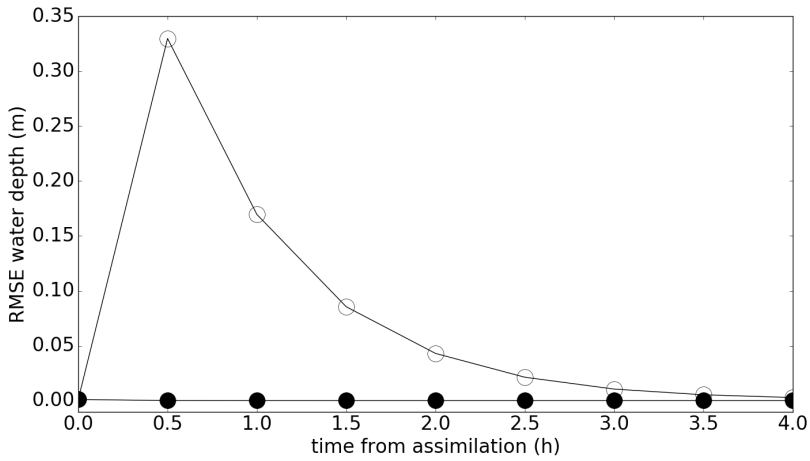

Figure 1: RMSE in water depth over the domain. Open circles show the RMSE between the continuous truth and restarted truth for a hydrostatic assumption. Filled circles show the RMSE when reinitialising with forecast velocities.

We performed a simple comparison of the values of hu and hv obtained using our approach ('simple calculation value') with those calculated during an assimilation in which hu and hv were included in the state vector ('analysis value'). We compared values at each assimilation time in an identical twin experiment in which we update both the water levels and the channel friction parameter. Figure 2 shows some typical results (from the assimilation at 28h in the SPL experiment as described in section 3.2).

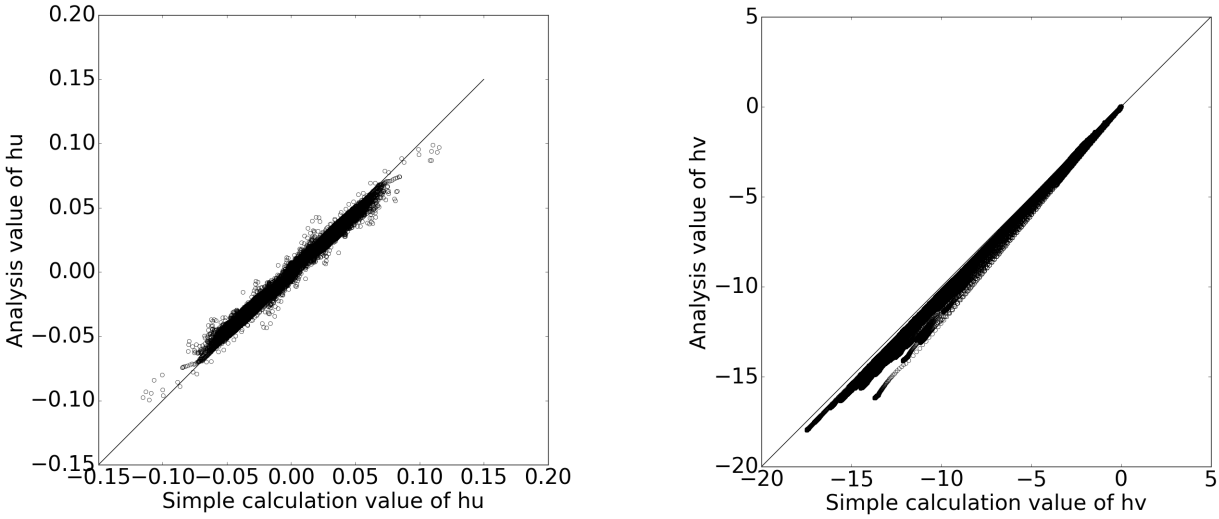

(a) (b)

Figure 2: Comparison of (a)hu and (b)hv values at one assimilation time.

Figure 2a shows the values of hu obtained using the simple method described here plotted against the values of hu calculated by including hu and hv in the assimilation state vector; figure 2b shows the same results for values of hv. The values are close to the identity line and therefore in broad agreement. We found that the agreement between the methods was better at later assimilation times, when the forecast error is low. This is because both methods provide values close to the pre-assimilation (forecast) values when the adjustment by the assimilation is relatively small. The values of n_{ch} obtained when including hu and hv in the state vector were almost identical (less than 0.001% difference) to the values obtained when the state vector comprises only water levels. We observed no instabilities in the solution at initialisation times using our technique.

3 Experimental design

3.1 Model domain

All of the experiments referred to in this paper are carried out in domains with a simplified river valley-like topography. We use two domains in this work, the 'long' domain describes an area of 20 km by 250m and is shown in figure 3. Note that the axes are not to the same scale. The 'regular' domain is 5km by 250m, and is identical to the long domain for $15 \leq y \leq 20km$ and $0 \leq x \leq 250m$. The domains are gently sloping symmetrical valleys with a 50m wide central river channel as shown in figure 3. The grid cell size for the computation is 10m by 10m in all cases. The river channel is defined to be the central 5 grid cells in the x direction for all values of y; the rest of the domain is defined as the flood plain. The cross section for both domains is the same. The domain has an upstream-downstream slope of 0.08% and the slope of the floodplain towards the river is 0.8%.

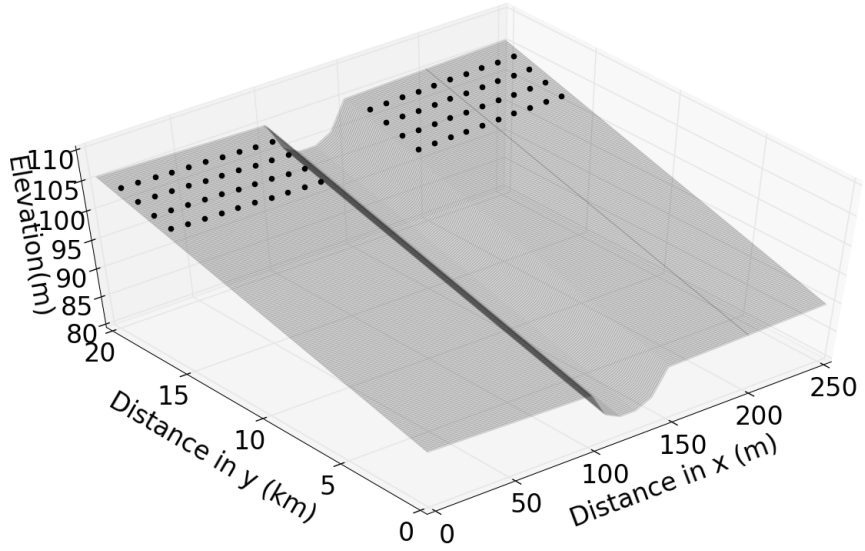

Figure 3: Elevation in metres of the long test domain used for the assimilation experiments. The regular domain is defined by $15 \leq y \leq 20$km and $0 \leq x \leq 250$m . The black circles indicate the positions of water elevation observations used in the data assimilation.

3.1.1 Identical Twin Experiments

In this study, we use an ETKF in identical twin experiments. Identical twin experiments are commonly carried out in order to test a data assimilation system as well as to generate information about the model to which data assimilation is applied, e.g. Evensen (1994). In such experiments, a numerical model is used to generate a 'truth' output for a set of known initial conditions and model parameters. The truth run for these experiments is a continuous run of our inundation model for 112 hours using a time varying inflow shown by the solid black circles in figure 4. For the first four hours, the inflow is set to be constant in order to fill the river channel with water in a spin-up period. The inflow from t = 4 hours onwards is based on some real hydrograph data from a flood of the river Severn near Tewksbury in the summer of 2007. The inflow comprises hourly values and linear interpolation is used to give flow rates between the hourly points. The truth run uses a value for the channel Manning's friction parameter of $n_{ch} = 0.04$, which is the value given for a natural stream by Maidment and Mays (1988). The Manning's friction parameter on the flood plain, n_{fp}, is likely to be higher due to vegetation and here we use a value of $n_{fp} = 0.05$.

The inundation model is also used to generate a 100 member ensemble of flood realisations. This ensemble represents a forecast of the true flood given uncertainty in upstream inflow and the channel friction parameter values. The number of ensemble members is relatively small compared to the dimension of the state vector, which contains 125,000 water depths for the short domain and 500,000 for the long domain. However, none of the problems which indicate undersampling of a system (spurious correlations or ensemble collapse as out-

lined by Petrie and Dance (2010)) are seen, suggesting that this number of ensemble members is sufficient to represent the system in this case. It was therefore not necessary to apply any localisation or inflation in these experiments.

Each member of the ensemble is driven by a different inflow and has a different friction parameter. The ensemble inflows are generated by adding time correlated random errors to the 'true' inflow; the ensemble inflows and true inflow are shown in figure 4. The variance for the inflow distribution is a proportion of the inflow, since the error in measured or predicted flow is likely to be flow-dependent as in Garcia-Pintado et al. (2015). The standard deviation for the generated inflow distribution is $0.15\times$ inflow, which is the same as the value used by Garcia-Pintado et al. (2015) and fits within the range of errors in measured flow rates (4% to 43%) reported in Di Baldassarre and Montanari (2009). No bias was applied to the inflow ensemble so that the mean inflow is very similar to the true inflow. An unbiased inflow removes the effect of an incorrectly specified inflow, since this has already been studied elsewhere. Choosing to use an unbiased inflow allows the effect of the incorrect channel friction parameter value to be studied in isolation; we briefly address the effect of biased inflow at the end of section 4.2.

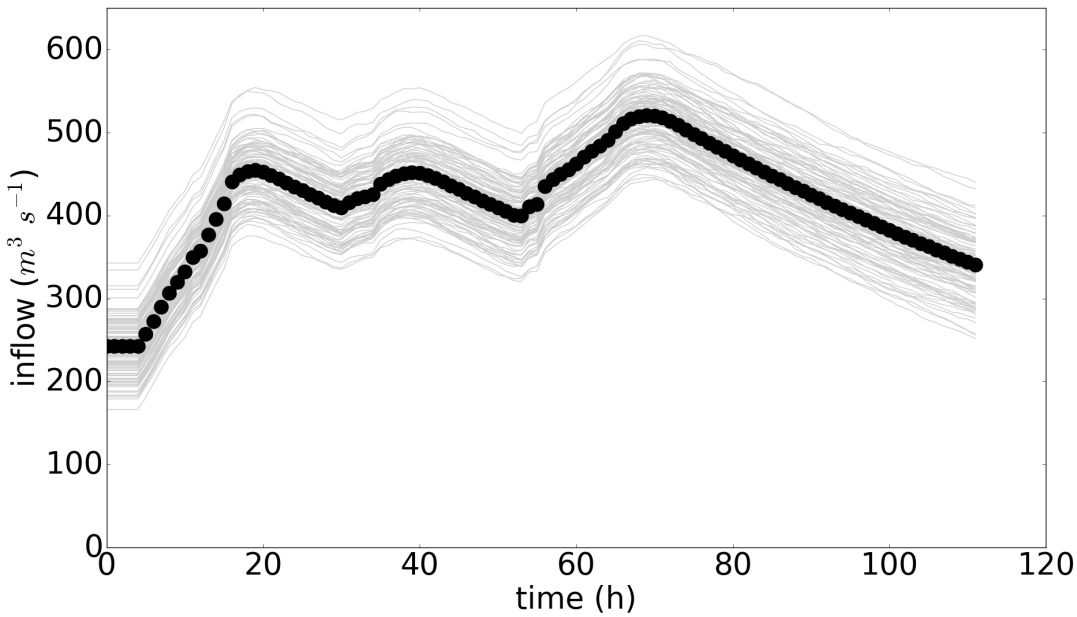

Figure 4: Inflow ensemble with time. Circles show the true inflow values and the grey lines show the ensemble inflows.

Ensemble channel friction parameters were generated by selecting from a Gaussian distribution centered on a 'wrong' initial value to reflect the fact that this parameter varies between catchments and will not generally be known before the start of a flood event. The channel friction parameter is also not likely to be directly measurable as it relates to the specific way in which Clawpack models friction. Different numerical models with different friction parameterisations have been shown to generate different optimal friction parameters for the same data for this reason (Horritt and Bates, 2002). Centering the channel friction ensemble on an

incorrect value also enables us to test whether the data assimilation scheme can retrieve the correct value from an incorrect starting point. The channel friction parameter for each ensemble member was selected from a Gaussian distribution, $\mathcal{N}(0.05, 0.01)$. The true value of $n_{ch} = 0.04$ falls within one standard deviation of the mean of this distribution. The value for the friction parameter on the flood plain was set to the true value for all ensemble members, i.e. $n_{fp} = 0.05$.

Data assimilation using an ETKF is carried out on the forecast ensemble, using synthetic observations generated from the truth as described in the next section. Since the ETKF is a sequential method the ensemble is evolved forward in time with Clawpack between observation times, and an assimilation is carried out every 12 hours. We also consider the 'open loop' case in which a forecast ensemble runs with no assimilation.

3.2 Experimental configurations

Various data assimilation experiments are carried out, each for a total period of 112 hours, including a 4 hour spin-up period with constant inflow to allow the river channel to fill up (as shown in figure 4). Synthetic observations of water depth are taken from the truth at 12 hourly intervals and assimilated with an ETKF. Twelve hourly assimilation intervals were chosen as this represents the smallest likely return time for SAR-equipped satellites at present.

Since we are running identical twin experiments, we know the true water levels everywhere in the domain, as well as the water levels forecast by the ensemble. It is theoretically possible to use SAR derived flood extent observations along with a high quality digital terrain model (DTM) to derive water levels in all parts of a real domain. One approach in a synthetic experiment is therefore to directly use water level observations in all parts of the domain as a proxy for SAR derived information. Such an approach is used by, for example, Lai and Monnier (2009). In reality, SAR images can reliably provide information about water elevation only at a few points along the flood extent, as demonstrated in Mason et al. (2012).

In this paper, we do not use all the available water levels. Instead we use synthetic observations of water levels taken directly from the truth run in the positions shown by black circles in figure 3. This approach replicates a situation in which four reliable flood delineation positions are available from a SAR image at $y = 16, 17, 18$ and 19km. We assume that we have a SAR image covering the domain from $y = 15$ to 20km, and that water elevation at each of the four flood edge positions can be obtained from a DTM; this water level can then be extrapolated perpendicular to the channel to give water elevation in each floodplain cell in the cross sections where we have delineation observations (i.e. $y = 16, 17, 18, 19$km). Although extrapolating water elevation across a cross section in this way would also give information about the water elevation in the channel, we exclude observations in the channel in this experiment. This is because topography information in the channel is likely to be much less accurate than that for the floodplain, making water depths less certain. Noise was added to the synthetic observations to represent uncertainty due to instrument error, flood extent determination and typical DTM errors. The noise added to the observations is Gaussian with a standard deviation of 0.25m; this is the same value as used in experiments with real data in Garcia-Pintado et al. (2015).

We perform a series of experiments using both the regular and long domain. Further details of the individual experiments are as follows:

- **Case SOR**: State-only estimation in the regular domain, with the ensemble of channel friction parameters distributed about an incorrect 'first guess' as described in section 3.1.1.
- **Case SOL**: State-only estimation in the long domain; all other details as for case SOR.
- **Case SPR**: State and channel friction parameter estimation in the regular domain with the initial channel friction parameter distributed about an incorrect 'first guess' as described in section 3.1.1 .
- **Case SPL**: State and channel friction parameter estimation in the long domain; all other details as for case SPR.
- **Case OR**: Open loop ensemble forecast in the regular domain; this is a free running ensemble forecast with the same initial conditions as cases SOR and SPR but without assimilation of observations.
- **Case OL**: Open loop ensemble forecast for the long domain.

The positions of the observations are the same for both the long and regular domains. This corresponds to observations spread throughout the whole of the regular domain, and observations only in the upstream part of the long domain. The observation errors are the same for the four different configurations that use observations.

4 Results and discussion of assimilation

4.1 State only estimation (SOR and SOL)

Results from state-only assimilation experiments are shown in figure 5. The graphs shows the RMSE over the whole domain between the forecast ensemble mean water depths and the true water depths at three hourly intervals from the time of the first assimilation at 16h. Here, RMSE is defined as

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\overline{(h_i^f - h_i^t)^2}}, \tag{26}$$

where h_i^f and h_i^t are the forecast ensemble mean and true water depth in cell i respectively. This definition means that the error is averaged over a larger area for the longer domain than for the regular domain. The values of RMSEs, though broadly similar, are therefore not directly comparable between domains.

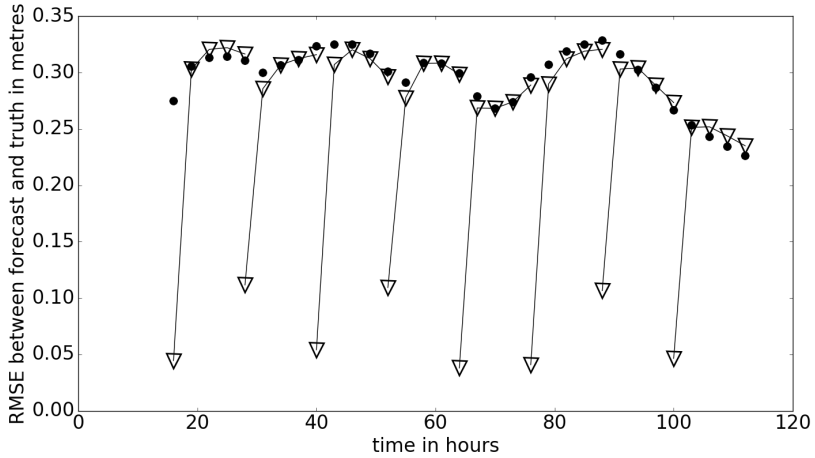

(a) RMSE between the ensemble forecast mean and the true water depth ($\overline{\mathbf{x}^f} - \mathbf{x}^t$) in the regular domain. Triangles show results for state only estimation (SOR), with assimilation carried out at t=16h, 28h, 40h, 52h, 64h, 76h, 88h and 100h. Circles show the open loop (OR) case for the regular domain (no assimilation).

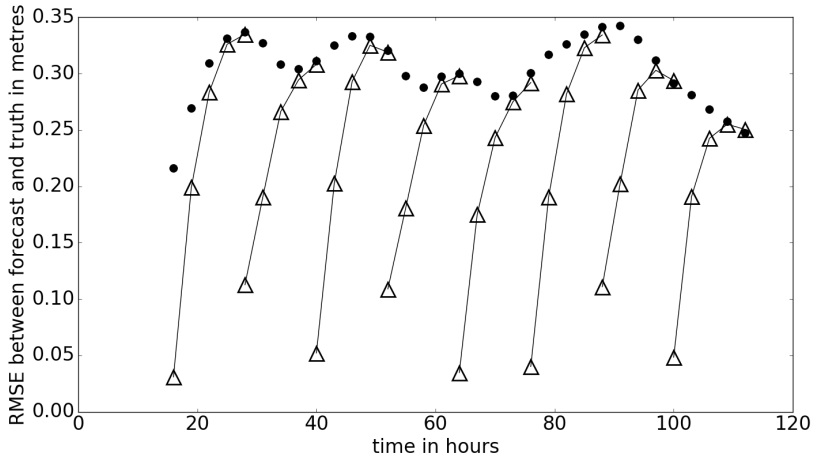

(b) RMSE between the ensemble forecast mean and the true water depth ($\overline{\mathbf{x}^f} - \mathbf{x}^t$) in the long domain. Triangles show results for state only estimation (SOL), with assimilation carried out at t=16h, 28h, 40h, 52h, 64h, 76h, 88h and 100h. Black circles show the open loop case (OL) for the long domain (no assimilation).

Figure 5: RMSEs for state-only estimation in the regular and long domains.

Figure 5a shows the RMSE between the forecast ensemble mean and true water depths with time for the regular domain (SOR). Figure 5b shows the same results for the longer domain (SOL). In both domains the ETKF produces a good analysis. The difference between the analysis ensemble mean water levels and the true water levels is very small at the time of each observation, and a large improvement is seen compared to the open loop forecast. However, for the regular domain in particular, the forecast skill is quickly lost during each subsequent forecast step and the RMSE quickly relaxes towards the open loop case. Comparing 5a and 5b shows that the RMSE at each analysis time is broadly similar for the two different domains. The results also

show that the forecast skill persists for longer in the long domain; the forecast takes longer to relax to the open loop case in the longer domain than in the regular domain. This means that the same observations are having a longer-lived impact on the forecast when a longer stretch of river and floodplain is considered.

In order to further understand why the observation impact is longer-lived in the long domain, the evolution of the error during the forecast step can be investigated. Figure 6 shows the long domain in plan view with the error between the forecast ensemble mean and the true water levels in each cell. The errors are shown at several times during the forecast after assimilation at t = 52h and before assimilation at t = 64h. In this particular forecast period the inflow is increasing steadily, but similar patterns are seen for forecast periods in which the inflow is varying in other ways.

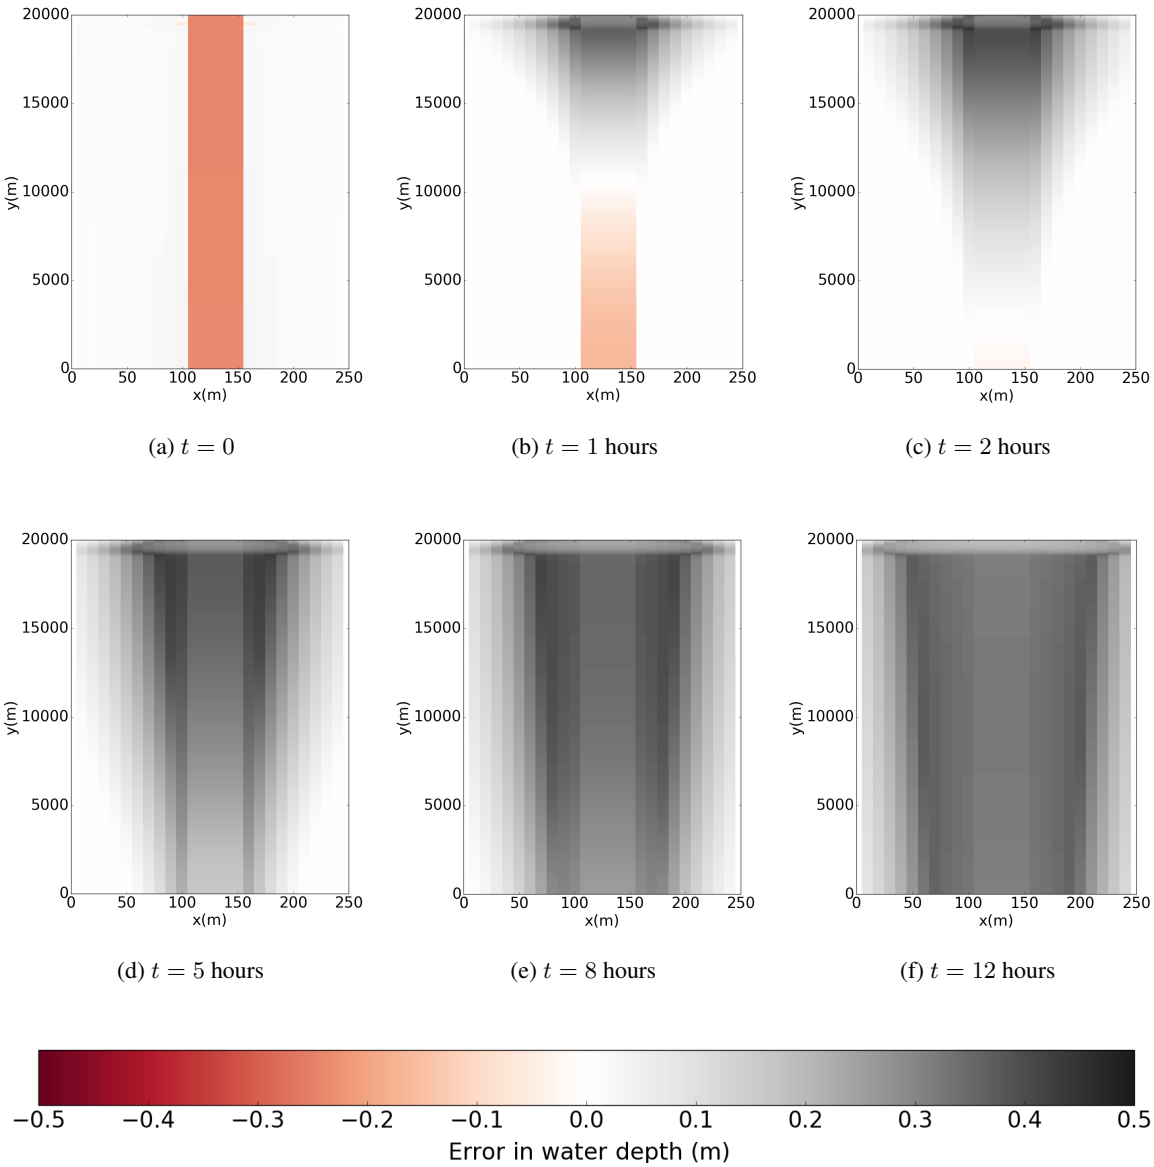

(a) $t = 0$ (b) $t = 1$ hours (c) $t = 2$ hours

(d) $t = 5$ hours (e) $t = 8$ hours (f) $t = 12$ hours

Figure 6: Forecast (ensemble mean) water depth minus true water depth in the long domain, shown in plan view for case SOL. Times are measured from assimilation at t = 52h. For reference, the true water depth on the floodplain varied between 0 and 0.4m during this forecast period. Water depth in the centre of channel varied between 5.9m and 6.4m.

Figure 6a shows the difference between the forecast ensemble mean water depths and the true water depths at the observation/assimilation time. The error between the forecast ensemble mean and true water depths is small at all points in the domain at this time. The difference between the forecast ensemble mean and the truth is highest in the channel; this reflects the fact that there are no observations of channel water depth used in the assimilation. Figure 6b shows the error between the forecast ensemble mean and the truth 1 hour after the assimilation. The error at the downstream end of the domain remains small, while a large, positive error can be seen in the upstream part. A positive error here means that the forecast is overestimating the water depth.

In figure 6c, the errors at 2 hours after assimilation are shown. The physical area in which the forecast is overestimating water depths has grown in the hour between 6b and 6c, and penetrated further downstream into the domain. This pattern of error growth from upstream to downstream continues over time in figures 6d (5 hours after assimilation), 6e (8 hours after assimilation) and 6f (12 hours after assimilation).

The observed pattern of error growth explains why the observations have impact in the longer domain for a longer time. The regular domain is defined by $15 \leq y \leq 20$km and $0 \leq x \leq 250$m (as shown in figure 3). The errors in water depth in this part of the long domain become very large within a few hours of the assimilation, as can be seen in figure 6c. This means that the error in the whole of the regular domain becomes large very quickly. In contrast, the RMSE for the long domain includes the downstream area where the water levels are still corrected from the assimilation; the RMSE is therefore lower for a longer time as the error takes a longer time to reach the downstream part of the longer domain.

During the forecast step, figure 6 shows that the error in the forecast resulting from incorrect specification of the n_{ch} parameter starts at the upstream end of the domain, and propagates downstream with time. This pattern of error growth is the same as that which would be expected from a bias in inflow, as noted in e.g. Andreadis et al. (2007). A biased inflow acting on corrected water levels will clearly degrade the forecast water levels close to the upstream boundary first and this error will propagate downstream with the flow. The results shown in figure 6 therefore indicate that errors due to incorrect inflow specification and those due to incorrect specification of the channel friction parameter may be difficult to separate out in a real flood event.

The low RMSEs between the analysis and the truth in the long domain highlight the fact that the ETKF is able to correct the water levels in areas for which there are no observations. The state error covariance matrix generated by the ensemble perturbations is such that information from the observations is spread throughout the domain, enabling corrections to be made to the state at the downstream part of the domain when observations are available only at the upstream end. This is further demonstrated in figure 7, which shows the difference between the forecast ensemble mean and the truth pre-assimilation (figure 7a) and post-assimilation (figure 7b), plotted in the long domain at t = 52h. Figure 7c shows the increments applied in the long domain as a result of observations in the upstream part of the domain only. The figure clearly shows that adjustments are made to water levels in the whole domain.

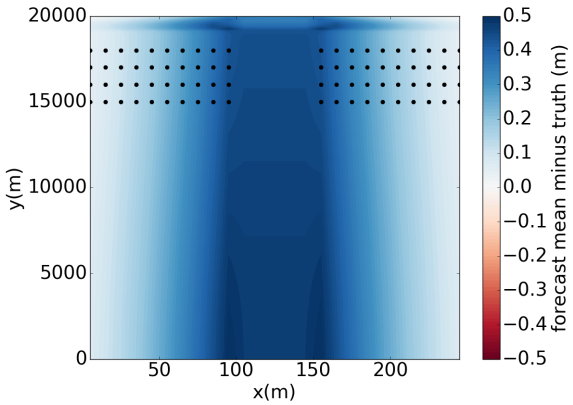

(a) Difference between the forecast ensemble mean and the truth ($\overline{\mathbf{x}^f} - \mathbf{x}^t$) plotted in the long domain (pre assimilation) at t = 52 hours.

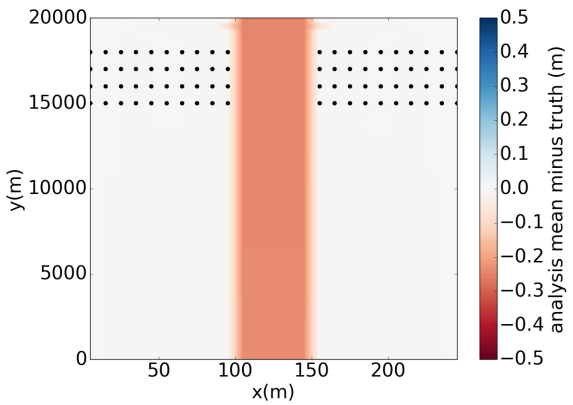

(b) Difference between the analysis ensemble mean and the truth ($\overline{\mathbf{x}^a} - \mathbf{x}^t$) plotted in the domain (after assimilation) at t = 52 hours.

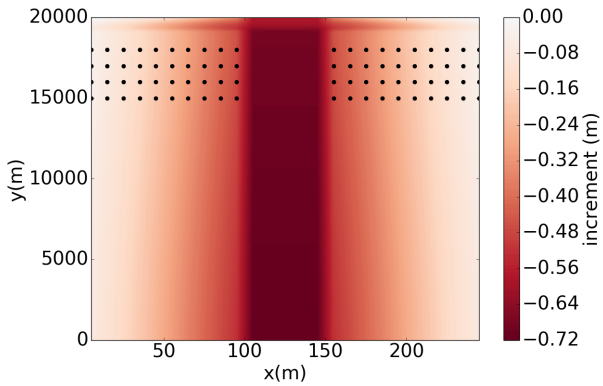

(c) Increments applied to the forecast ensemble mean for the assimilation at t = 52 hours

Figure 7: Difference between the forecast ensemble mean and the truth at t = 52 hours, (a) pre assimilation and (b) post assimilation; (c) shows the increment applied to the forecast to compute the analysis ($\overline{\mathbf{x}^a} - \overline{\mathbf{x}^f}$). In each plot the black circles show the position of the observations.

For a reliable ensemble, the RMSE should be close to the spread of the ensemble, where the spread is defined as the square root of the average ensemble variance (see e.g. Fortin et al. (2014)). Reliability plots (of RMSE vs. spread) should therefore produce points which lie close to the identity line. Such plots can be used to diagnose ensemble collapse, where the spread of an ensemble becomes unrealistically small.

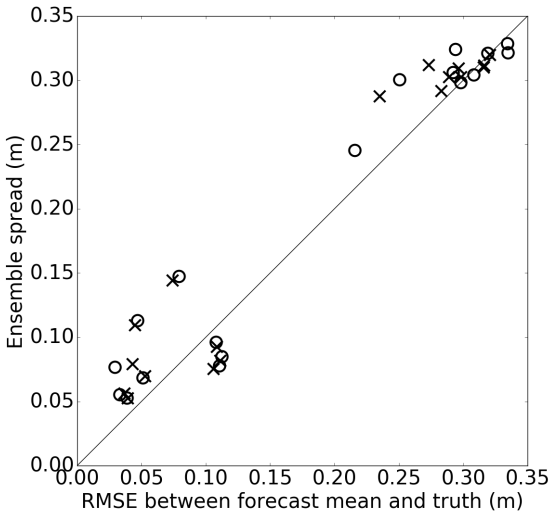

Figure 8: Reliability plot for the SOR and SOL experiments, showing forecast and analysis RMSE vs ensemble spread at each assimilation time. Circles are for the long domain, crosses for the regular domain. For an ideal ensemble, all points would lie on the identity line.

Figure 8 is a reliability plot for the SOR and SOL experiments. The points all lie close to the identity line, indicating that the ensemble spread is adequate to capture the uncertainty in the forecast. There is no indication of ensemble collapse. The points form two clusters, with large error, large spread values before an assimilation and low error, small spread values for analysis ensembles.

4.2 State and parameter estimation (SPR and SPL)

Considering a longer stretch of river in our idealised domain, as in section 4.1, shows an improvement in the forecast skill of the ETKF, in that the RMSE increases more slowly. By design, much of the error between the truth and the forecast comes from the incorrect channel friction parameter in these experiments. In this section we show results from using data assimilation to jointly estimate the state and the channel friction parameter in both domains.

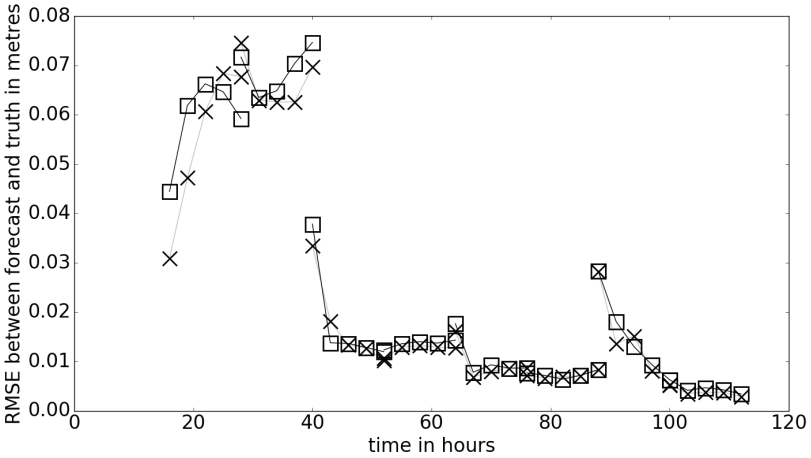

Figure 9: RMSE between the ensemble mean and the truth over the whole domain for joint state-parameter estimation; crosses denote RMSE for the long domain (SPL) and squares denote RMSE for the regular domain (SPR).

Figure 9 shows the RMSEs with time for the state-parameter estimation experiments in the regular and long domain. Comparison of figures 5 and 9 demonstrates a very clear improvement in the forecast for joint state-parameter estimation over the state only case (note the different scales on the y axes). Joint state-parameter estimation markedly increases the observation impact for exactly the same observations. The extra computional cost of estimating the friction parameter along with the state is extremely small, as it adds only a single extra component to the state vector. Joint state-parameter estimation is therefore a very efficient way of producing a much better forecast in this situation.

The small RMSE between the truth and the forecast for joint state-parameter estimation is of the order of the observation error. There is no longer any significant difference between the results for the different length domains and this implies that the error growth must occur at similar rates along the length of the long domain when the friction parameter is estimated.

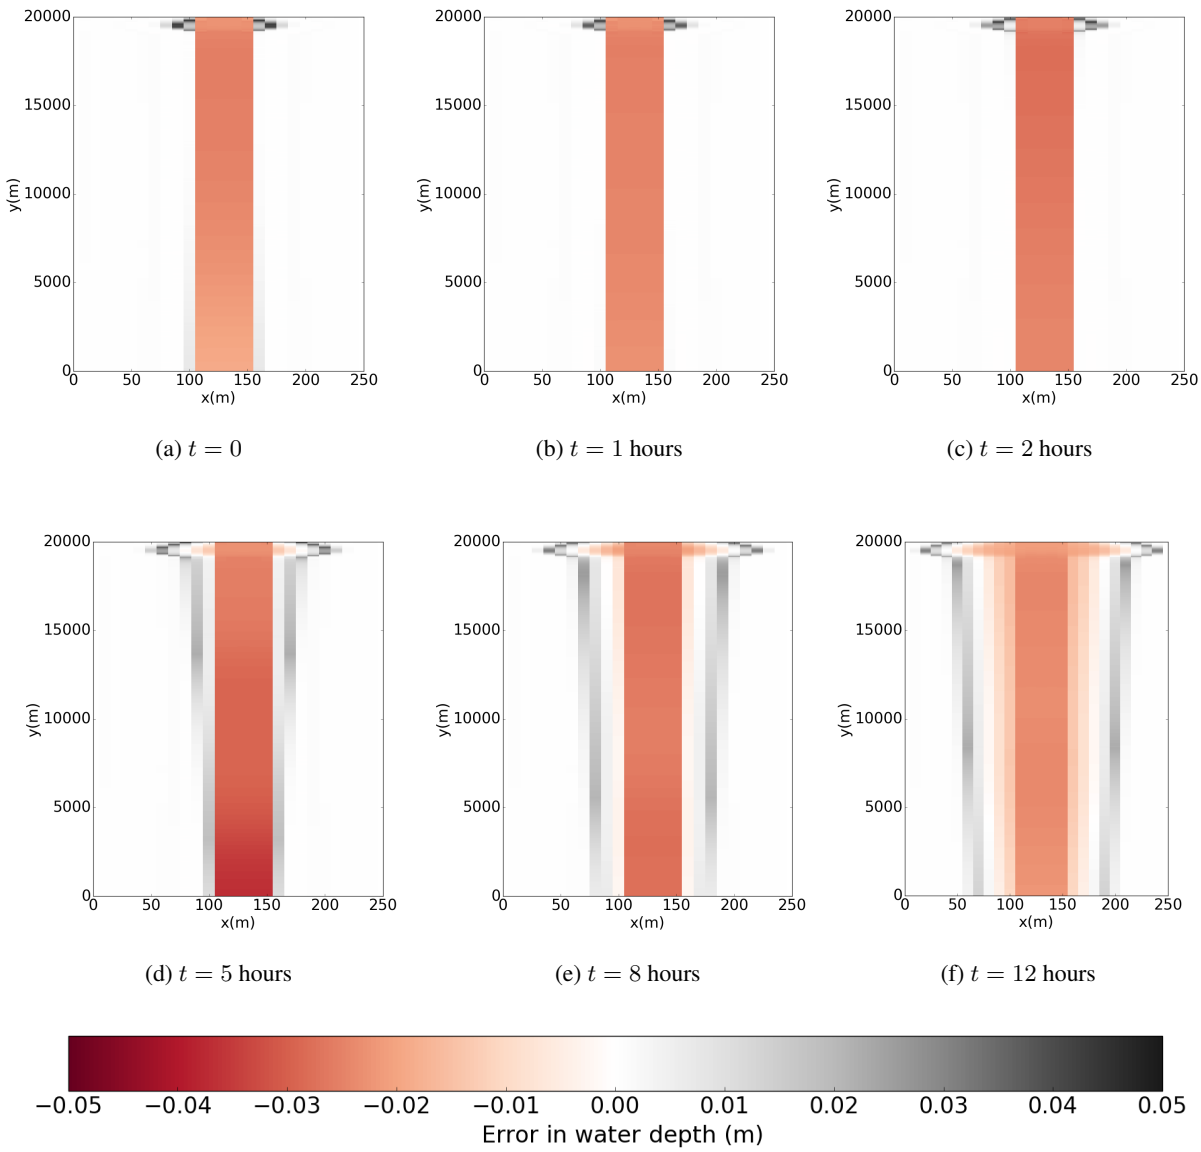

(a) $t = 0$ (b) $t = 1$ hours (c) $t = 2$ hours

(d) $t = 5$ hours (e) $t = 8$ hours (f) $t = 12$ hours

Figure 10: Forecast (ensemble mean) water depth minus true water depth in the long domain, shown in plan view for case SPL. Times are measured from assimilation at t=52h. For reference, the true water depth on the floodplain varied between 0 and 0.4m during this forecast period. Water depth in the centre of channel varied between 5.9m and 6.4m.

Figure 10 shows the the difference between the forecast ensemble mean and the true water levels, plotted in plan view in the long domain, at several times during the forecast step starting at t = 52 hours for the SPL experiment. Note that the scale in figure 10 is ten times smaller than in figure 6. Figure 10a shows the difference between the forecast ensemble mean water depths and the true water depths at the observation/assimilation time. As in the SOL experiment, the error between the forecast ensemble mean and true water depths is small at all points in the domain at this time. Figure 10b shows the error between the forecast and the truth 1 hour after the assimilation; figures 10c (2 hours after assimilation), 10d (5 hours after assimilation), 10e (8 hours after

assimilation) and 10f (12 hours after assimilation) show how the error evolves. Unlike in the SOL experiment, the error does not propagate from upstream to downstream; instead, the error grows at a similar rate along the length of the domain. This further suggests that the upsteam-downstream error growth seen in the SOL experiments is due to incorrect friction parameter specification.

Figure 10 shows that the mean forecast overestimates water depth on the floodplain but underestimates water depth in the channel. This reflects the fact that although each ensemble member predicts a physically realistic water level that is flat in cross section, the ensemble mean forecast is not necessarily flat and therefore not physically realistic. This is because the forecast mean water depth in each cell is the average value predicted by the ensemble members. An example of this situation is shown in figure 11.

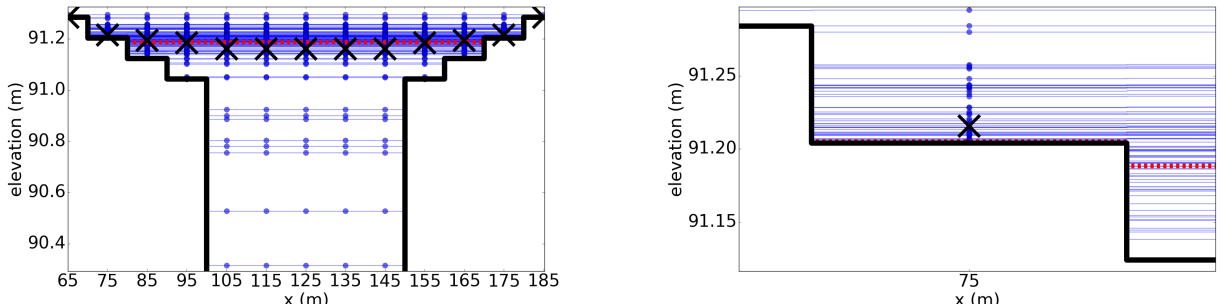

(a) (b)

Figure 11: Cross section of the domain showing bathymetry as a solid black line. The true water level is shown as a dashed red line, water levels predicted by each ensemble member are shown as blue cirlces. The mean forecast in each model cell is shown as a cross. Figure 11a shows the central part of the domain from $65 \leq x \geq 185$m. Figure 11b shows the forecast water levels and resulting forecast mean in the cell centred at 75m in greater detail.

Figure 11 shows the true flood level at one cross section of the domain and the water level predicted by each ensemble member. The crosses show the mean ensemble water level in each cell in the cross section. In the channel (e.g. at $x = 125$m), figure 11a shows the ensemble predictions are such that the ensemble mean is slightly lower than the true water level. Beyond the edge of the true flood on the flood plain, the true water depth relative to the topography is zero, and most ensemble members also predict zero water depth. However, as shown in figure 11b for the cell at $x = 75$m, there are cells beyond the flood edge in which some ensemble members predict non-zero water depth and the ensemble mean is therefore a very small positive water depth; this simply follows from equation 9. It is therefore possible for the ensemble mean to predict water levels deeper than the truth in cells beyond the true flood edge where there are a number of ensemble members which predict non-zero depth, even when the water level is under-predicted in the channel. It should be noted that the errors are very small.

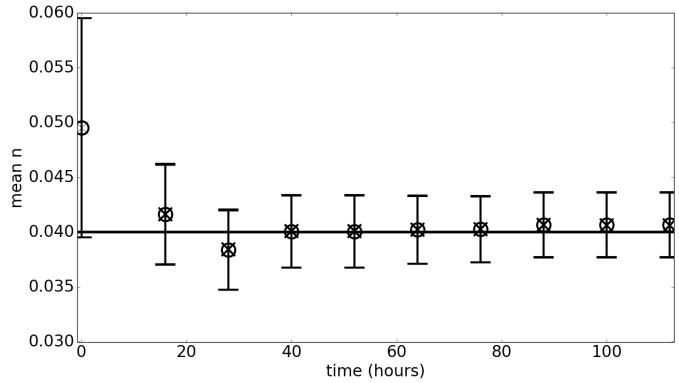

Figure 12: Calculated (analysis) mean channel friction parameter values at each assimilation time. The true value is shown by the horizontal line. Open circles show the values for the long domain, crosses for the regular domain. The error bars show one standard deviation of the analysis parameter distribution for the long domain; values for the regular domain are very similar.

Figure 12 shows the analysis ensemble mean value of the channel friction parameter at each assimilation time. The true value of the channel friction parameter, n_{ch}, indicated by the solid horizontal line, and the incorrect initial mean value is shown at time zero. The error bars show one standard deviation of the analysis n_{ch} distribution. The results show that the joint-state parameter data estimation produces a good estimate of the value of n_{ch}, and that the ensemble mean values are almost identical for both the regular and long domains. It is notable that the convergence of the estimated channel friction parameter value to the truth is achieved with water depth observations only taken on the floodplain.

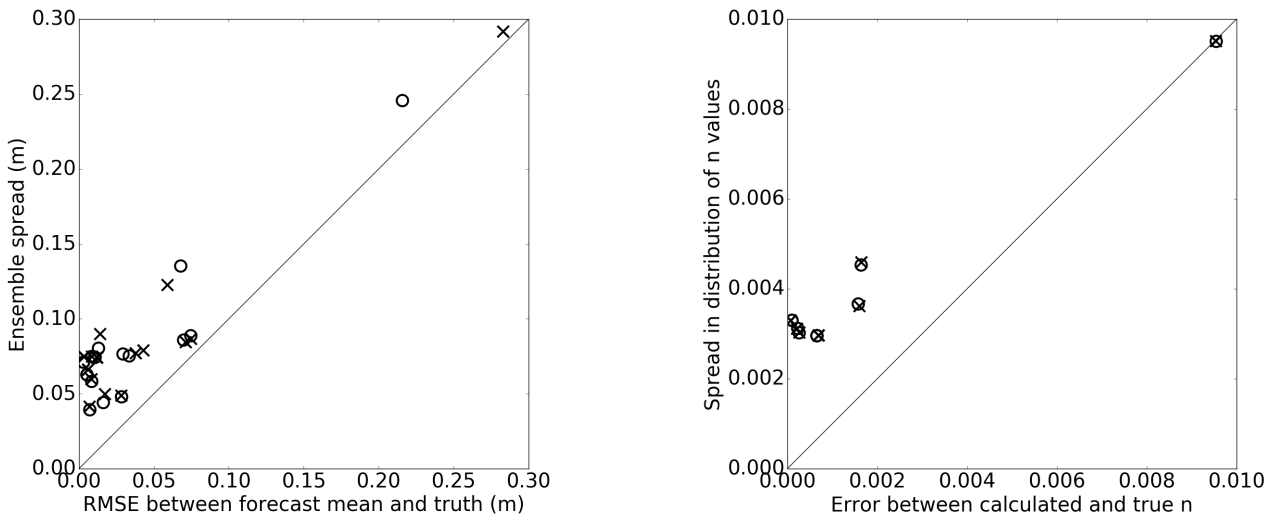

(a) (b)

Figure 13: Reliability plot for the SPR and SPL experiments. Circles are for the long domain, crosses for the regular domain. Plots show error vs spread for (a) the ensemble of water levels and (b) the n_{ch} distribution.

Reliability plots for the ensemble of state vectors are given in figure 13a and for the calculated parameter distributions in figure 13b. The plots show no evidence of ensemble collapse, which can be a problem in ensemble data assimilation schemes. In fact the spread of the state and parameter distributions remains broader than the size of the error in our experiments, which minimises the risk of overfitting the value of n_{ch} at a particular time in the simulation. If the spread of either the parameter or the state ensembles became too small it would be necessary to use inflation techniques (see e.g. Anderson (2007)).

These results show a clear advantage in jointly estimating the channel friction parameter alongside the model state; this contrasts with the findings of Garcia-Pintado et al. (2015), where no clear improvement in forecast water levels was seen. There are a number of possible reasons for this, one being that in Garcia-Pintado et al. (2015) convergence of the channel friction parameter value to a steady value was slow compared to the timescales of the flood event. The difference between our results and those in Garcia-Pintado et al. (2015) may also be related to the fact that there are more sources of uncertainty in Garcia-Pintado et al. (2015), which used real data and real topography, rather than the idealised situation in this study. One significant source of time-varying error not accounted for in the setup used by Garcia-Pintado et al. (2015) and not present in our experiments is lateral inflows (see Bermudez et al. (2017)). It may also be that the initial parameter value used in Garcia-Pintado et al. (2015) was already close to the true value, such that the error in the parameter was not a large source of error, whereas our initial guess was incorrect by design. A better initial guess would also explain the smaller changes to n_{ch} produced by the data assimilation in Garcia-Pintado et al. (2015). Alternatively, it may be that the presence of the initialisation shock in Garcia-Pintado et al. (2015) prevented convergence to a more accurate channel friction parameter.

Another possible reason for the contrasting importance to the forecast of updating the channel friction parameter may be that in Garcia-Pintado et al. (2015) a bias correction was also made to the inflow as part of the data assimilation. It may be that inflow correction was also indirectly including some of the effects of the incorrect n_{ch} parameter. This seems likely, since the upstream-downstream error propagation pattern seen in figure 6 would also be expected for an incorrectly specified inflow. Thus it may be that no added benefit was gained from including n_{ch} estimation in addition to inflow bias estimation because correcting error in inflow bias also compensates for any error in n_{ch}. To test the interdependence of inflow error and error in n_{ch} we conducted some state-parameter estimation experiments exactly as for the SPL experiments but with biased inflows. No correction was made to the biased inflow, but the value of n_{ch} was updated at each assimilation time.

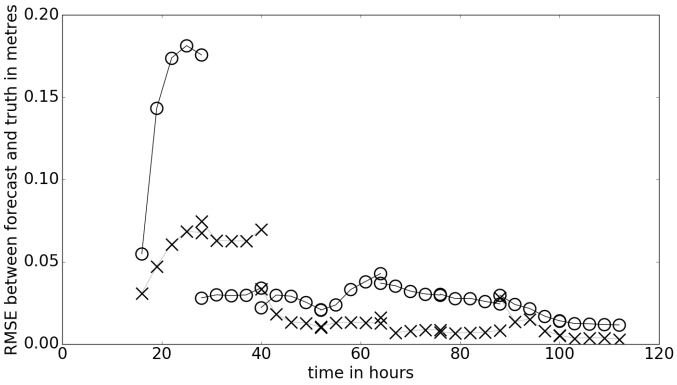

Figure 14: RMSEs between the mean forecast and the truth for positive inflow bias in the long domain shown as open circles. The corresponding RMSEs for the unbiased inflow are shown as crosses for comparison.

Figure 14 shows water level RMSEs with time for the long domain with biased and unbiased inflows, and demonstrates that correction of the channel friction parameter allows the forecast to predict accurate water levels, even with a biased inflow. The correction to the value of n_{ch} is therefore compensating for inflow bias.

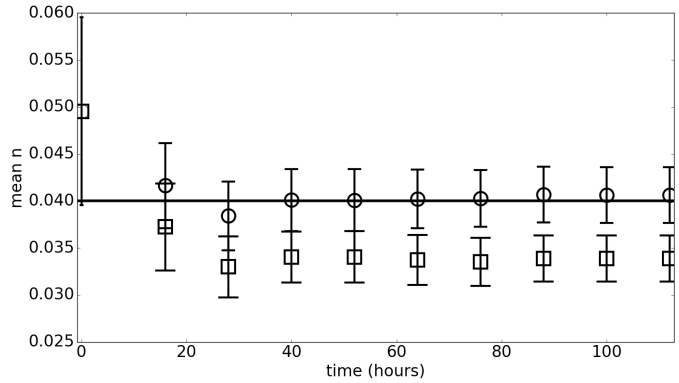

Figure 15: Calculated values of mean n_{ch} for biased inflow in the long domain shown as squares. Circles show the corresponding values for unbiased inflow. The horizontal line shows the true value and the error bars show one standard deviation of the calculated n_{ch} distribution.

The updated values of n_{ch} with time for biased inflow are shown in figure 15. When the inflow is biased, the value of n_{ch} does not converge to the true value, but instead to a value which allows it to compensate for inflow bias and minimise errors in water levels in the domain. This clearly indicates an interdependence between errors in channel friction parameter value and inflow.

Whatever the reason for the differences to the results in Garcia-Pintado et al. (2015), the work presented here shows that our inundation forecast model is sensitive to the channel friction parameter, and that the correct value of the parameter can be retrieved from observations of water depth by using joint state-parameter data assimilation during a flood event with unbiased inflow information.

5 Conclusions

In this study, we have investigated the effect of domain length and channel friction parameter estimation in data assimilation for flood inundation forecasting. We have also demonstrated that assuming water has zero momentum at the start of each forecast step can cause an initialisation shock. The period of time for which this shock then causes problems for the forecast is likely to be domain dependent; in this study we found that the impact of the shock disappeared within a period of approximately 4 hours. We developed a novel method to reinitialise water velocities in each model grid cell with pre-assimilation values, and showed that this approach eliminated the initialisation shock.

In agreement with Andreadis et al. (2007), Neal et al. (2009), Garcia-Pintado et al. (2013) and Matgen et al. (2010), we found that assimilating SAR-like water levels in a state-only data assimilation system gives a time limited improvement in the forecast skill, since such improvement can only persist for as long as the information is relevant in the domain. We have shown that considering a longer domain extends the time over which observations have an impact on the forecast, even when no extra observations are used. This is because

the ETKF is able to correct water levels downstream of the observations due to strong covariances between the errors in water levels in different parts of the domain. In a domain with more realistic topography, the correlations between the errors in water depth in different parts of the domain may not be as strong, and are likely to be more complex. However, work by authors such as Garcia-Pintado et al. (2015) indicates that water level errors in large, real domains are correlated, as depths can be corrected at considerable distances from SAR-derived observations. Additionally, understanding the effects of domain length and friction parameter estimation without the extra complications of topographical features is important in understanding the fundamental sensitivities of such systems.

We have shown that in the forecast period following an assimilation, the difference between the forecast and the truth when the channel friction parameter is incorrectly specified and not updated grows faster in the upstream part of the domain, and then propagates downstream. This error propagation is the reason that the forecast ensemble retains skill for a longer time in the long domain; the errors grow more slowly in the downstream areas which are not considered in the regular domain.

Jointly estimating the channel friction parameter along with the water levels is shown to produce a significantly better forecast for the same observations at very little extra cost. This was not seen in Garcia-Pintado et al. (2015), in which inflow and parameters were estimated simultaneously. Correcting the channel friction parameter also eliminates the differences in forecast error growth for the two different domain lengths, as errors grow at similar rates in the upstream and downstream parts of the domain. We have shown that it is possible to estimate a good value for the channel friction parameter, even when water level observations are available only on the floodplain.

In summary, we have shown that in the case where there is no inflow bias but channel friction is incorrectly specified, assimilating SAR-like water levels from floodplains provides a time limited improvement in the forecast when only water levels are corrected. The time over which the forecast is improved depends on the length of the domain of interest. Authors such as Andreadis et al. (2007); Matgen et al. (2010); Giustarini et al. (2011); Garcia-Pintado et al. (2015, 2013) have shown that inflow correction can lead to a marked improvement in forecast skill. We have shown that jointly estimating the channel friction parameter along with the water levels also provides a clear improvement in the forecast at all times and can retrieve an accurate channel friction parameter value. Our results suggest that it may be difficult to separate out errors due to incorrect specification of inflow and incorrect specification of channel friction when carrying out ensemble data assimilation for inundation modelling. This is because the character of the errors in the forecast resulting from these two sources of uncertainty are similar; this interdependence explains our finding that updating the value of the channel friction parameter can compensate for a bias in inflow. Further study is required to see how well each of these conclusions are applicable to more complex and realistic topography, and for real satellite derived observations. In this way, the work here may serve to enhance operational flood forecasting potential.

Acknowledgements

The authors gratefully acknowledge the NERC studentship supporting Elizabeth Cooper, and CASE sponsorship from the Satellite Applications Catapult. This work was also supported in part by NERC grants NE/K00896X/1 and NE/K008900/1 as well as EPSRC grant EP/P002331/1 and the NERC National Centre for Earth Observation (NCEO). The software used in this article can be obtained from the corresponding author on request, subject to licensing conditions.

Appendix A Inflow and friction source terms

In order to model fluvial flooding in a stretch of river, the amount of water flowing into that stretch needs to be modelled. Here we describe the implemetation of a new inflow source term in Clawpack. Clawpack is able to solve systems of partial differential equations with or without source terms and is designed such that the user can introduce new code to describe additional source terms. Our new inflow source term has been implemented in the code in a similar way to the pre-existing friction source term. The code for the friction source term is outlined in section A.1 and the new inflow source term is described in section A.2.

A.1 Pre-existing friction source term

Friction between the fluid and the channel in which it is flowing acts as a momentum source in the shallow water equations. This is represented in Clawpack with a source term of the form

$$\mathbf{R(q)}_{friction} = \begin{bmatrix} 0 \\ -\gamma(hu) \\ -\gamma(hv) \end{bmatrix}, \tag{A.1}$$

where γ is given by

$$\gamma = \frac{gn^2\sqrt{(hu)^2 + (hv)^2}}{h^{\frac{7}{3}}}, \tag{A.2}$$

h is water depth and u and v are velocities in the x and y directions. Acceleration due to gravity is denoted g and n is Manning's friction coefficient. This coefficient describes the roughness of the channel bed in which the water is flowing and in practical applications its value is usually determined empirically. The value of n is specified by the user in the Clawpack code, and can vary over the domain if specified in the simulation setup. The units of n are $sm^{-1/3}$.

In Clawpack, inhomogeneous sets of equations are solved using the method of fractional stepping described in LeVeque (2002) p.380-395. This method splits the equation into two simpler problems; one homogeneous conservation law and one inhomogeneous partial differential equation - which can be solved independently over the same time step. The solutions are then combined in an alternating fashion to give a solution to the whole problem. For a friction source term, the set of problems to be solved are the homogeneous system

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{F(q)}}{\partial x} + \frac{\partial \mathbf{G(q)}}{\partial y} = 0, \tag{A.3}$$

and the source term equation

$$\frac{\partial \mathbf{q}}{\partial t} = \mathbf{R(q)}_{friction}, \tag{A.4}$$

with $\mathbf{R(q)}_{friction}$ as in A.1 and $\mathbf{q}$ a vector of conserved quantities (see equation (2)). For each time step in the code, equation (A.3) is advanced from t_i to t_{i+1} to give intermediate values of h_*, hu_* and hv_* in $\mathbf{q}_*$. The values in $\mathbf{q}_*$ are then used to solve equation (A.4) over the same time step. This introduces a 'splitting error' into the solution of order Δt = t_{i+1}-t_i, making the whole method only first order accurate. A more accurate splitting method ('Strang splitting') is available for implementation in the code, but the first order accuracy has been found to be more stable and sufficient in practice for a wide range of applications (LeVeque, 2002, 1997).

A.2 Novel inflow source term

For inundation simulations, water entering the domain of interest can be modelled as a source term. We have added a novel inflow source term to the Clawpack code to model river-like flow. In operational situations, information regarding this source term may be available from an upstream gauge as a mass flow rate, Q, measured in m^3s^{-1}. In an ungauged catchment, the same information could be generated using a rainfall run-off model. Here we use hourly values of Q based on gauge data; linear interpolation is carried out between the hourly values. A water mass flow rate can be turned into a source term, S expressed in ms^{-1} (c.f. the term 'Sce' in Hervouet (2007) p.31, which has the same units and can include rainfall, infiltration etc) as long as the area of the domain or 'footprint' over which the water is added is known. For water added over an area A, $S = Q/A$. The equation for the inflow source term is then given by

$$\frac{\partial \mathbf{q}}{\partial t} = \mathbf{R}(\mathbf{q})_{inflow}, \tag{A.5}$$

where

$$\mathbf{R}(\mathbf{q})_{inflow} = \begin{bmatrix} S \\ 0 \\ 0 \end{bmatrix}. \tag{A.6}$$

Equations A.5 and A.6 show that for each time step Δt, the change in h due to the incoming water will depend on the value of the inflow source, S, over the same Δt. The extra water arriving in the domain creates extra water height, and is assumed here to arrive without any momentum; the water is subject only to hydrostatic momentum effects. This inflow source term has been implemented in the code in the following way

- determine in which grid cells the source term will be applied. This is reasonably arbitrary but must be such that the solution remains stable (we used trial and error in this experiment);

- calculate the total area that the inflow cells cover in the domain, A;
- for a given mass flow rate Q, calculate S for each value of Q by dividing by A;
- at the relevant grid points extract depth , h_*, as calculated from equation (A.3);
- calculate the change to h_* due to incoming water from a discretisation of equation (A.6) using a Crank-Nicholson scheme (Crank and Nicolson, 1996)

$$h = h_* + \Delta t \frac{S(t) + S(t + \Delta t)}{2}; \tag{A.7}$$

- use the new value of h from equation (A.7) to solve for the next time step.

A.3 Combining friction and inflow source terms

The source terms described in this section are applied in a sequential manner in the code. For each time step, the inflow source term calculates the new water depths in the relevant parts of the domain and then the friction source term is applied to the new water depths.

References

J. L. Anderson. An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus A: Dynamic Meteorology and Oceanography*, 59(2):210–224, 2007. doi: 10.1111/j.1600-0870.2006.00216.x.

K. M. Andreadis and G. J.-P. Schumann. Estimating the impact of satellite observations on the predictability of large-scale hydraulic models. *Advances in Water Resources*, 73:44 – 54, 2014. ISSN 0309-1708. doi: http://dx.doi.org/10.1016/j.advwatres.2014.06.006.

K. M. Andreadis, E. A. Clark, D. P. Lettenmaier, and D. E. Alsdorf. Prospects for river discharge and depth estimation through assimilation of swath-altimetry into a raster-based hydrodynamics model. *Geophysical Research Letters*, 34(10), 2007. ISSN 1944-8007. doi: 10.1029/2007GL029721. L10403.

M. Bermudez, J. C. Neal, P. D. Bates, G. Coxon, J. E. Freer, L. Cea, and J. Puertas. Quantifying local rainfall dynamics and uncertain boundary conditions into a nested regional-local flood modeling system. *Water Resources Research*, 53(4):2770–2785, 2017. ISSN 1944-7973. doi: 10.1002/2016WR019903.

K. M. Brown, C. H. Hambidge, and J. M. Brownett. Progress in operational flood mapping using satellite synthetic aperture radar (SAR) and airborne light detection and ranging (LiDAR) data. *Progress in Physical Geography*, 40(2):196–214, 2016. doi: 10.1177/0309133316633570.

Clawpack Development Team. Clawpack software, 2014. Version 5.2.2. URL http://www.clawpack.org.

E. Cooper, S. Dance, J. Garcia-Pintado, N. Nichols, and P. Smith. Improving inundation forecasting using data assimilation. *Mathematics Report Series, University of Reading*, 2013. URL `http://www.reading.ac.uk/web/FILES/maths/mathsreport.pdf`.

J. Crank and P. Nicolson. A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type. *Advances in Computational Mathematics*, 6(1):207–226, 1996. ISSN 1572-9044. doi: 10.1007/BF02127704.

G. Di Baldassarre and A. Montanari. Uncertainty in river discharge observations: a quantitative analysis. *Hydrology and Earth System Sciences*, 13(6):913–921, 2009. doi: 10.5194/hess-13-913-2009.

G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical research*, 99(C5):10142–10162, 1994.

G. Evensen, D. P. Dee, and J. Schröter. *Parameter Estimation in Dynamical Models*, pages 373–398. Springer Netherlands, Dordrecht, 1998. ISBN 978-94-011-5096-5. doi: 10.1007/978-94-011-5096-5_16.

V. Fortin, M. Abaza, F. Anctil, and R. Turcotte. Why should ensemble spread match the rmse of the ensemble mean? *Journal of Hydrometeorology*, 15(4):1708–1713, 2014. doi: 10.1175/JHM-D-14-0008.1.

J. Garcia-Pintado, J. C. Neal, D. C. Mason, S. L. Dance, and P. D. Bates. Scheduling satellite-based SAR acquisition for sequential assimilation of water level observations into flood modelling. *Journal of Hydrology*, 495:252 – 266, 2013. ISSN 0022-1694. doi: http://dx.doi.org/10.1016/j.jhydrol.2013.03.050.

J. Garcia-Pintado, D. Mason, S. L. Dance, H. Cloke, J. C. Neal, J. Freer, and P. D. Bates. Satellite-supported flood forecasting in river networks: a real case study. *Journal of Hydrology*, 523:706–724, 2015.

D. L. George. Augmented Riemann solvers for the shallow water equations over variable topography with steady states and inundation. *Journal of Computational Physics*, 227(6):3089–3113, 2008.

L. Giustarini, P. Matgen, R. Hostache, M. Montanari, D. Plaza, V. R. N. Pauwels, G. J. M. De Lannoy, R. De Keyser, L. Pfister, L. Hoffmann, and H. H. G. Savenije. Assimilating sar-derived water level data into a hydraulic model: a case study. *Hydrology and Earth System Sciences*, 15(7):2349–2365, 2011. doi: 10.5194/hess-15-2349-2011.

G. H. Golub and C. F. Van Loan. Matrix computations. 1996. *Johns Hopkins University, Press, Baltimore, MD, USA*, pages 374–426, 1996.

S. Grimaldi, Y. Li, V. R. N. Pauwels, and J. P. Walker. Remote sensing-derived water extent and level to constrain hydraulic flood forecasting models: Opportunities and challenges. *Surveys in Geophysics*, 37(5): 977–1034, 2016. ISSN 1573-0956. doi: 10.1007/s10712-016-9378-y.

HEC-RAS Development Team. Hec-ras software. http://www.hec.usace.army.mil/software/hec-ras/.

J. M. Hervouet. Telemac modelling system: an overview. *Hydrological Processes*, 14(13):2209–2210, 2000. ISSN 1099-1085. doi: 10.1002/1099-1085(200009)14:13 <2209::AID-HYP23 >3.0.CO;2-6.

J.-M. Hervouet. *Hydrodynamics of Free Surface Flows*. John Wiley and Sons Ltd, 2007.

M. Horritt and P. Bates. Evaluation of 1D and 2D numerical models for predicting river flood inundation. *Journal of Hydrology*, 268(14):87 – 99, 2002. ISSN 0022-1694. doi: http://dx.doi.org/10.1016/S0022-1694(02)00121-X.

R. Hostache, X. Lai, J. Monnier, and C. Puech. Assimilation of spatially distributed water levels into a shallow-water flood model. Part II: Use of a remote sensing image of mosel river. *Journal of Hydrology*, 390(34):257 – 268, 2010. ISSN 0022-1694. doi: http://dx.doi.org/10.1016/j.jhydrol.2010.07.003.

R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.

X. Lai and J. Monnier. Assimilation of spatially distributed water levels into a shallow-water flood model. Part I: Mathematical method and test case. *Journal of Hydrology*, 377(12):1 – 11, 2009. ISSN 0022-1694. doi: http://dx.doi.org/10.1016/j.jhydrol.2009.07.058.

R. J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, 2002.

R. J. LeVeque. Wave propagation algorithms for multidimensional hyperbolic systems. *Journal of Computational Physics*, 131(2):327–353, 1997.

D. Livings. Aspects of the Kalman filter. *MSc thesis, Unversity of Reading. http://www.reading.ac.uk/web/FILES/maths/Livings.pdf*, 2005.

D. M. Livings, S. L. Dance, and N. K. Nichols. Unbiased ensemble square root filters. *Physica D: Nonlinear Phenomena*, 237(8):1021 – 1028, 2008. ISSN 0167-2789. doi: http://dx.doi.org/10.1016/j.physd.2008.01.005.

D. Maidment and L. Mays. *Applied Hydrology*. McGraw-Hill series in water resources and environmental engineering. Tata McGraw-Hill Education, 1988. ISBN 9780070702424.

K. T. Mandli, A. J. Ahmadia, M. Berger, D. Calhoun, D. L. George, Y. Hadjimichael, D. I. Ketcheson, G. I. Lemoine, and R. J. LeVeque. Clawpack: building an open source ecosystem for solving hyperbolic pdes. *PeerJ Computer Science*, 2:e68, 2016. doi: 10.7717/peerj-cs.68.

D. Mason, T. Scott, and S. Dance. Remote sensing of intertidal morphological change in Morecambe Bay, U.K., between 1991 and 2007. *Estuarine, Coastal and Shelf Science*, 87(3):487 – 496, 2010. ISSN 0272-7714. doi: http://dx.doi.org/10.1016/j.ecss.2010.01.015.

D. Mason, G.-P. Schumann, J. Neal, J. Garcia-Pintado, and P. Bates. Automatic near real-time selection of flood water levels from high resolution synthetic aperture radar images for assimilation into hydraulic models: A case study. *Remote Sensing of Environment*, 124:705 – 716, 2012. ISSN 0034-4257. doi: http://dx.doi.org/10.1016/j.rse.2012.06.017.

D. Mason, J. Garcia-Pintado, H. Cloke, and S. Dance. The potential of flood forecasting using a variable-resolution global digital terrain model and flood extents from synthetic aperture radar images. *Frontiers in Earth Science*, 3:43, 2015. ISSN 2296-6463. doi: 10.3389/feart.2015.00043.

P. Matgen, G. Schumann, F. Pappenberger, and L. Pfister. Sequential assimilation of remotely sensed water stages in flood inundation models. In *IAHS-AISH Symposium on remote Sensing for Environmental Monitoring and Change Detection*, number 316, pages 78–88, 2007.

P. Matgen, M. Montanari, R. Hostache, L. Pfister, L. Hoffmann, D. Plaza, V. R. N. Pauwels, G. J. M. De Lannoy, R. De Keyser, and H. H. G. Savenije. Towards the sequential assimilation of sar-derived water stages into hydraulic models using the particle filter: proof of concept. *Hydrology and Earth System Sciences*, 14(9): 1773–1785, 2010. doi: 10.5194/hess-14-1773-2010.

M. Mure-Ravaud, G. Binet, M. Bracq, J.-J. Perarnaud, A. Fradin, and X. Litrico. A web based tool for operational real-time flood forecasting using data assimilation to update hydraulic states. *Environmental Modelling & Software*, 84:35 – 49, 2016. ISSN 1364-8152. doi: http://dx.doi.org/10.1016/j.envsoft.2016.06.002.

I. Navon. Practical and theoretical aspects of adjoint parameter estimation and identifiability in meteorology and oceanography. *Dynamics of Atmospheres and Oceans*, 27(1):55 – 79, 1998. ISSN 0377-0265. doi: http://dx.doi.org/10.1016/S0377-0265(97)00032-8.

J. Neal, G. Schumann, P. Bates, W. Buytaert, P. Matgen, and F. Pappenberger. A data assimilation approach to discharge estimation from space. *Hydrological Processes*, 23(25):3641–3649, 2009. ISSN 1099-1085. doi: 10.1002/hyp.7518.

J. Neal, G. Schumann, and P. Bates. A subgrid channel model for simulating river hydraulics and floodplain inundation over large and data sparse areas. *Water Resources Research*, 48(11):n/a–n/a, 2012. ISSN 1944-7973. doi: 10.1029/2012WR012514. W11506.

R. E. Petrie and S. L. Dance. Ensemble-based data assimilation and the localisation problem. *Weather*, 65(3): 65–69, 2010. ISSN 1477-8696. doi: 10.1002/wea.505.

S. Ricci, A. Piacentini, O. Thual, E. L. Pape, and G. Jonville. Correction of upstream flow and hydraulic state with data assimilation in the context of flood forecasting. *Hydrology and Earth System Sciences*, 15(11): 3555–3575, 2011.

G. Schumann, P. D. Bates, M. S. Horritt, P. Matgen, and F. Pappenberger. Progress in integration of remote sensing derived flood extent and stage data and hydraulic models. *Reviews of Geophysics*, 47(4), 2009. ISSN 1944-9208. doi: 10.1029/2008RG000274. RG4001.

T. Scott and D. Mason. Data assimilation for a coastal area morphodynamic model: Morecambe Bay. *Coastal Engineering*, 54(2):91 – 109, 2007. ISSN 0378-3839. doi: http://dx.doi.org/10.1016/j.coastaleng.2006.08.008.

T. R. Scott, P. J. Smith, S. L. Dance, D. C. Mason, M. J. Baines, N. K. Nichols, K. J. Horsburgh, P. K. Sweby, and A. S. Lawless. Data assimilation for morphodynamic prediction and predictability. In *Coastal Engineering, Proceedings of the 31st International Conference*, pages 2386–2398. 2008. doi: 10.1142/9789814277426_0197.

P. J. Smith, S. L. Dance, M. J. Baines, N. K. Nichols, and T. R. Scott. Variational data assimilation for parameter estimation: application to a simple morphodynamic model. *Ocean Dynamics*, 59(5):697, 2009. ISSN 1616-7228. doi: 10.1007/s10236-009-0205-6.

P. J. Smith, S. L. Dance, and N. K. Nichols. A hybrid data assimilation scheme for model parameter estimation: application to morphodynamic modelling. *Computers & Fluids*, 46(1):436–441, July 2011. 10th ICFD Conference Series on Numerical Methods for Fluid Dynamics (ICFD 2010).

P. J. Smith, G. D. Thornhill, S. L. Dance, A. S. Lawless, D. C. Mason, and N. K. Nichols. Data assimilation for state and parameter estimation: application to morphodynamic modelling. *Quarterly Journal of the Royal Meteorological Society*, 139(671 Part B):314–327, January 2013.

G. Thornhill, D. Mason, S. Dance, A. Lawless, N. Nichols, and H. Forbes. Integration of a 3D variational data assimilation scheme with a coastal area morphodynamic model of Morecambe Bay. *Coastal Engineering*, 69:82 – 96, 2012. ISSN 0378-3839. doi: http://dx.doi.org/10.1016/j.coastaleng.2012.05.010.

M. K. Tippett, J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker. Ensemble square root filters. *Monthly Weather Review*, 131(7):1485–1490, 2003. doi: 10.1175/1520-0493(2003)131 <1485:ESRF >2.0.CO;2.

C. Vrsmarty, A. Askew, W. Grabs, R. Barry, C. Birkett, P. Dll, B. Goodison, A. Hall, R. Jenne, L. Kitaev, J. Landwehr, M. Keeler, G. Leavesley, J. Schaake, K. Strzepek, S. Sundarvel, K. Takeuchi, and F. Webster. Global water data: A newly endangered species. *Eos*, 82(5):54+56+58, 2001. doi: 10.1029/01EO00031.