

UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

School of Geography

**Flood forecasting and adaptive sampling with spatially
distributed dynamic depth sensors**

By

Jeffrey Neal

Thesis for the degree of Doctor of Philosophy

July 2007

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS
SCHOOL OF GEOGRAPHY

Doctor of Philosophy

Flood forecasting and adaptive sampling with spatially distributed dynamic depth sensors

By Jeffrey Neal

The movement of computational power and communications capabilities onto networks of sensors in the environment through the concept of pervasive or ubiquitous computing has initiated opportunities for the delivery of ground-based data in real-time and the development of adaptive monitoring systems. Measurements of water level taken by a network of wireless sensors called 'FloodNet' were assimilated into a one-dimensional hydrodynamic model using an ensemble Kalman filter, to create a forecasting model. The ensemble Kalman filter led to an increase in forecast accuracy of between 50% and 70% depending on location for forecast lead times of less than 4 hours.

This research then focused on methods for targeting measurements in real-time, such that the power limited but flexible resources deployed by the FloodNet project could be used optimally. Two targeting methods were developed. The first targeted measurements systematically over space and time until the forecasting model predicted that the probability of the water level exceeding a pre-defined threshold was less than 5%. The second method targeted measurements based on the expected decrease in forecasted water level error variance at a validation time and location, quickly calculated for various sets of measurements by an ensemble transform Kalman filter. Estimates of forecast error covariance from the ensemble Kalman filter and ensemble transform Kalman filter were significantly correlated, with correlations ranging between 0.979 and 0.292. Targeting measurements based on the decrease in forecast error variance was found to be more efficient than the systematic sampling method. The ensemble transform Kalman filter based targeting method was also used to estimate the 'signal variance' of theoretical measurements at any computational node in the hydrodynamic model. Furthermore, time series data, different sensors types and measurements of floodplain stage could all be taken into account either as part of the targeting process or prior to measurement targeting.

Table of contents

1 Introduction.....	1
1.1 Thesis aims	3
2 Literature review.....	6
2.1 Distributed flood models	6
2.1.1 Data sources and issues arising therein.....	8
2.1.2 Parameter estimation, calibration and validation.....	11
2.2 Data assimilation and flood forecasting.....	15
2.2.1 Flood forecasting	18
2.2.2 State updating with Kalman filters	19
2.3 Adaptive sampling	20
2.3.1 Adaptive sampling with Kalman filters	24
2.3.2 The ensemble transform Kalman filter	25
2.4 Summary.....	27
3 Methods	28
3.1 Hydraulic modelling	28
3.1.1 Channel model	28
3.1.2 Floodplain flows	30
3.2 Model calibration: GLUE	34
3.2.1 The Likelihood Measure	36
3.3 Data assimilation with the Kalman filter	38
3.3.1 The system model	39
3.3.2 The measurement model.....	40
3.3.3 The predictor step	41
3.3.4 The Analysis step.....	42
3.3.5 Parameter estimation.....	45
3.4 Dealing with non-linearity	46
3.5 Ensemble Methods: The ensemble Kalman filter	47
3.5.1 State propagation	48

3.5.2 Measurement model.....	50
3.5.3 Analysis step.....	51
3.5.4 Parameter estimation.....	52
3.5.5 Practical implementation	55
3.6 Adaptive sampling: The ensemble transform Kalman filter.....	57
3.7 Summary	61
4 Flood inundation model updating using an Ensemble Kalman filter and spatially distributed measurements.	62
4.1 Introduction.....	62
4.2 Methods	63
4.2.1 Hydraulic modelling	63
4.2.2 Forecasting: The Ensemble Kalman Filter	64
4.3 Application Case study	67
4.3.1 Study area and data.....	67
4.3.2 Model calibration.....	76
4.3.3 Application of EnKF to a 1D hydraulic model of the River Crouch.....	82
4.3.4 Estimation of autoregressive model coefficients.....	84
4.4 Results and discussion	86
4.4.1 Increasing forecast accuracy using the EnKF.....	86
4.4.2 Forecast accuracy in relation to sampling design.....	91
4.4.3 Estimating forecast uncertainty using the EnKF	92
4.5 Conclusions.....	97
4.6 Summary.....	99
5 Evaluating the utility of the ensemble transform Kalman filter for adaptive sampling when updating a hydrodynamic model.....	100
5.1 Introduction.....	100
5.2 Implementation of the ETKF on the RC model.....	103
5.2.1 Simulation model and updating procedure	105
5.2.2 Parameterisation issues	108
5.3 Results: A comparison between EnKF and ETKF forecasts of error covariance.	110

5.3.1	Analysed variances	111
5.3.2	Forecast variances: Single sensor example.....	113
5.3.3	Forecast variances: multiple sensor example.....	132
5.4	Discussion.....	136
5.5	Conclusions.....	139
5.6	Summary	141
6	Adaptive space-time sampling with wireless sensor nodes for flood forecasting .	142
6.1	Introduction.....	142
6.1.1	The FloodNet approach	142
6.2	Study site and model setup	145
6.2.1	Sensor node hardware.....	147
6.2.2	Measurement data	148
6.2.3	Forecasting model.....	149
6.3	Adaptive sampling methods and results	150
6.3.1	Method 1	152
6.3.2	Method 2.....	158
6.4	Discussion.....	168
6.5	Conclusions.....	172
6.6	Summary	174
7	Optimal sampling design with an ensemble transform Kalman filter	175
7.1	Introduction.....	175
7.2	Signal variance for in-channel sensors	177
7.2.1	Flow dependent error propagation and signal variance	190
7.2.2	Signal variance after one measurement	197
7.2.3	Signal variance and time series data.....	203
7.3	Signal variance at the study site.....	213
7.4	The influence of measurement variance on signal variance	219
7.5	Discussion.....	212
7.6	Conclusions.....	226
7.7	Summary	228

8 Discussion	229
8.1 Updating hydrodynamic model state	229
8.1.1 Error sources and parameter uncertainty	232
8.2 Adaptive sampling	233
8.2.1 Adaptive sampling with the ETKF	235
8.3 Measurement data issues	236
9 Conclusions	239
10 References	243
Appendix A: Simple KF example: bucket water levels	260
Appendix B: Ensemble transform Kalman filter code	261

List of figures

Fig. 3.1: Simplified diagram of the forces acting on the water mass in an open channel, where G is the weight of the water mass. The bed slope $\tan(\omega)$ can replace the direction of flow component $\sin(\omega)$ when the angle is small.	30
Fig. 3.2: DEM and example reservoir used to parameterise ISIS floodplain reservoir nodes.	31
Fig. 3.3: Cross-section between floodplain reservoirs, used to parameterise ISIS floodplain sections. The same principle applies to ISIS spill sections between the channel nodes and floodplain reservoirs.	33
Fig.3.4: Schematic diagram of the KF. The predictor step advances the model state and state error covariance to the next time step using the system model. The corrector step then updates these forecasts with measurements of the state variables. Analysed state and state error variance estimates are then returned to the system model as updated initial conditions.	44
Fig. 4.1: Map of the River Crouch and River Roach, Essex, UK. Scale 1:150,000. To the south is the Thames Estuary, with the North Sea to the east. The box indicates the field site where FloodNet sensors have been deployed.	68
Fig. 4.2: Topographic map of study site derived from LiDAR data overlain by sensor locations and model topology.	70
Fig. 4.3: The River Crouch at sensor 10 looking west. Sensor 1 is located on the far bank of the river in the centre of the image. Photo was taken approaching high tide on the 26 th March 2004.	70
Fig.4.4: The River Crouch near sensor 1 and Hullbridge, looking west. Sensor 1 is located next to the boat. Photo was taken approaching low tide on the 26 th March 2004.	71
Fig. 4.5: Area of floodplain near to sensor 4 looking north west at low and mid tide.	72
Fig. 4.6: Area of floodplain looking south towards sensors 4 and 5. Photo was taken at high tide on the 16 th March 2004.	72
Fig. 4.7: LiDAR DEM of Brandy Hole, with air photo overlay. Looking west over Brandy hole towards Hullbridge.	73
Fig. 4.8: Diagram showing model topology at Brandy Hole. The areas that contribute to each reservoir are also shown as polygons in Fig 4.2. The area shown is 2.5 km across.	74
Fig. 4.9a: Likelihood histograms for 18 consecutive tidal events between the 21 st and 30 th May 2004. Plots show the effect on likelihood of varying channel roughness, using measurements from sensor 1.	79
Fig. 4.9b: Likelihood histograms for 18 consecutive tidal events between the 21 st and 30 th May 2004. Plots show the effect on likelihood of varying floodplain roughness, using measurements from sensor 3.	80
Fig. 4.9c: Likelihood histograms for 18 consecutive tidal events between the 21 st and 30 th May 2004. Plots show the effect on likelihood of varying boundary condition magnitude, using measurements from sensor 1.	81
Fig. 4.10: Plot of RC model stage predictions against stage measurements at	87

four sensor locations: sensor 1 * , sensor 2 O , sensor 3 Δ , sensor 4 + .	
Fig. 4.11: Plot of residuals between RC model predictions and measured stage against predicted stage at sensor 1. Dots are joined by event.	87
Fig. 4.12: Plot of updated RC model stage forecasts against measured stage at sensor 1 after one time step Δ and after four time steps * . Data are shown from 11 tidal cycles.	89
Fig. 4.13: Plot of updated RC model stage forecasts against measured stage at sensor 1 * and sensor 3 O , after 4 time steps.	90
Fig. 4.14: Example forecast at sensor 1 during tidal cycle 6. Forecast stage after 1, 2 and 4 filter time steps with measurement data overlain.	90
Fig. 4.15: Difference between squared errors in forecast from run A and run B, for sensor 3 (upper plot) and sensor 1 (lower plot). A positive value indicates an improved forecast of stage from run A over run B.	92
Fig. 4.16: Post analysis ensemble state standard deviation, with the number of measurement time steps / filter time steps shown on x-axis. The y-axis shows analysed state standard deviation in metres.	94
Fig. 4.17: Scatter plots of forecast error at sensor 1 against sensor 3 and difference in forecast error against stage at sensor 1.	97
Fig. 5.1: Adaptive sampling timings.	103
Fig. 5.2: Map of sensor locations and RC model topology around the validation location at sensor 2.	107
Fig. 5.3: ETKF analysed error variance plotted against the EnKF analysed error variance, for 39 tidal events at the validation location, after six target times.	112
Fig. 5.4: ETKF analysed signal variance plotted against the EnKF analysed error, for 39 tidal events at the validation location, after six target times.	112
Fig. 5.5: Time series plots showing forecast stage at sensor 2 from target times 1 to 6 (dark lines). Observations were available from sensor 2 at each target time. The background forecast is indicated by the light lines, whilst measurements at five minute intervals from sensor 2 are shown as stars.	114
Fig. 5.6: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times from the assimilation of the routine measurement against background error variance.	115
Fig. 5.7: Signal variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement.	117
Fig. 5.8: Forecast error variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement.	118
Fig. 5.9: Plot of model background states (stage and flow) and their variance for channel cross sections, with signal variance at validation time and location shown for theoretical measurements taken at any RC model node along the river at target time one. Measurement variance was 0.005 m.	119

Fig. 5.10: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times after assimilating the routine measurement against background error variance. Measurements were simulated from the ensemble mean.	121
Fig. 5.11: Forecast error variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Measurements were simulated from the ensemble mean.	123
Fig. 5.12: Signal variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Measurements were simulated from the ensemble mean.	124
Fig. 5.13: Plot of the signal variance at the validation time against change in ensemble mean (after assimilating the routine measurement). There are 39 dots in total each of which is a separate tidal event. Crosses indicate the corresponding results where the measurement was generated from the ensemble mean.	125
Fig. 5.14: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times from the assimilation of a measurement at the first target time after assimilating the routine measurement.	126
Fig 5.15: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times after assimilating a measurement at the first target time against the background error variance after the routine measurement. Measurements were simulated from the ensemble mean.	127
Fig 5.16: Forecast error variance at seven target times and the validation time after the assimilation of a measurement at the first target time and routine measurement time. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement.	128
Fig 5.17: Forecast error variance at seven target times and the validation time after the assimilation of a measurement at the first target time. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Measurements were simulated from the ensemble mean.	129
Fig 5.18: ETKF predictions of error variance (lower plots) and EnKF predictions of error variance (upper plots) at three times after assimilating two or three measurements at target time 6 against the background error variance after the routine measurement.	133
Fig 5.19: ETKF predictions of error variance (lower plots) and EnKF predictions of error variance (upper plots) at three times after assimilating two or three measurements at target time 6 against the background error variance after the routine measurement.	134
Fig. 6.1: Conceptual diagram of the two FloodNet control loops. An inner loop comprising of a peer-to-peer wireless network and an outer loop	143

moving measurements from the sensor network to a database and forecasting model, and then returning sampling requests to the sensor network.

- Fig. 6.2: Schematic diagram of four different updating or data assimilation methods: 1. Input updating. 2. State updating. 3. Parameter updating. 4. Error prediction. Also shown is the adaptive sampling concept, where grey boxes define FloodNet specific control loops. 145
- Fig. 6.3: Location of sensor nodes and RC model topology. 147
- Fig. 6.4: Adaptive sampling timings. 151
- Fig. 6.5: Scatter plot of 39 high tides at the location of sensor 8 as predicted by the RC model without data assimilation (Dots) and measured by sensor 8 (crosses). 153
- Fig 6.6: Results of adaptive sampling Method 1. 154
- Fig. 6.7: Event 12 showing evidence of a phase difference between the background forecast mean and the observed tide along with major model instability. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event. 156
- Fig. 6.8: Tide 35 where the background simulation was similar to the actual tide. Includes the model forecast at stopping point. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event. 157
- Fig 6.9: Tide 18 where the background simulation was greater in magnitude than the actual tide. Includes model forecast at stopping point. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event. 157
- Fig. 6.10: Tide 5 where the background simulation was less than the actual tide. Includes the model forecast at stopping point and shows evidence of model instability. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event. 158
- Fig. 6.11: Bar chart of signal variance for Q possible measurements spread over eight sensor locations and five target times. 161
- Fig. 6.12: Bar chart of signal variance for Q possible measurements spread over eight sensor locations and five target times, given a measurement from sensor 2 at target time 5. 162
- Figure 6.13: Results of adaptive sampling Method 2. The plot shows forecast error variance statistics prior to data assimilation, after the routine measurement and after assimilating a further one to six measurements chosen by the ETKF. 167
- Fig. 6.14: Results of adaptive sampling Method 2 with median results only. 167
- Fig 7.1.a-i: Plots of model states (stage and flow) and their uncertainties (variance) for channel cross-sections, with log signal variance at 178-

validation location for validation and target times given a measurement with variance 0.005. Note: all axes are constant between plots except log signal variance. a) Plots of states and signal 4 hours prior to validation time. b) Plots of states and signal 3 hours 40 minutes prior to validation time. c) Plots of states and signal 3 hours 20 minutes prior to validation time. d) Plots of states and signal 3 hours prior to validation time. e) Plots of states and signal 2 hours 40 minutes prior to validation time. f) Plots of states and signal 2 hours 20 minutes prior to validation time. g) Plots of states and signal 2 hours prior to validation time. h) Plots of states and signal 1 hour 40 minutes prior to validation time. i) Plots of states and signal 1 hour 20 minutes prior to validation time.	182
Fig. 7.2: Top: Map of RC model channel nodes (black dots). The study site is at Hullbridge. Bottom: Long section of river thalweg and bank heights.	183
Fig. 7.3: Plot of variance in stage at the target location and time against the signal variance at the validation location and target time. Results are shown for 39 consecutive tidal events at each channel node in the RC model.	187
Fig. 7.4: Plot of stage variance at the target location against distance from the validation location. Results are shown for 39 consecutive tidal events at each channel node in the RC model.	187
Fig. 7.5: Plot of target time covariance between target location and validation location against target time signal variance. Results are shown for 39 consecutive tidal events at each channel node in the RC model.	188
Fig. 7.6: Plot of target time signal variance against target time stage variance at the target location and target time error covariance between the target location and validation location. Results are shown for 39 consecutive tidal events at each channel node in the RC model.	188
Fig. 7.7: Plot of target time signal variance against target time stage variance at the target location and target time error covariance between the target location and validation location, from an alternative angle. Results are shown for 39 consecutive tidal events at each channel node in the RC model.	189
Fig. 7.8: Plot of error covariance between target location and validation location against distance from the validation location. Results are shown for 39 consecutive tidal events at each channel node in the RC model.	189
Fig 7.9.a-i: Plots of model states (stage and flow) and their uncertainties (variance) for channel cross sections, with measurement with variance 0.005. Validation location was 1.9 km downstream. a) Plots of states and signal 4 hours prior to validation time. b) Plots of states and signal 3 hours 40 minutes prior to validation time. c) Plots of states and signal 3 hours 20 minutes prior to validation time. d) Plots of states and signal 3 hours prior to validation time. e) Plots of states and signal 2 hours 40 minutes prior to validation time. f) Plots of states and signal 2 hours 20 minutes prior to validation time. g) Plots of states and signal 2 hours prior to validation time. h) Plots of states and signal 1 hour 40 minutes prior to validation time. i) Plots of states and	192- 196

signal 1 hour 20 minutes prior to validation time.

- Fig 7.10.a-h: Plots of model states (stage and flow) and their uncertainties (variance) for channel cross sections, with measurement with variance 0.005. After assimilating one measurement at one of three locations 4 hours prior to the validation time with the ETKF. a) Plots of states and signal 3 hours 40 minutes prior to validation time. b) Plots of states and signal 3 hours 20 minutes prior to validation time. c) Plots of states and signal 3 hours prior to validation time. d) Plots of states and signal 2 hours 40 minutes prior to validation time. e) Plots of states and signal 2 hours 20 minutes prior to validation time. f) Plots of states and signal 2 hours minutes prior to validation time.g) Plots of states and signal 1 hour 40 minutes prior to validation time. h) Plots of states and signal 1 hour 20 minutes prior to validation time. 198-202
- Fig. 7.11.a-h: Results of assimilating a time series of synthetic measurements at 3 locations within the RC model domain. The top plot in each figure displays mean stage estimates from the background ensemble. Estimated stage variance after assimilating a measurement 4 hours prior to the validation time is displayed for each of the three example sensors. Signal variance for the target time and locations was plotted for each example sensor. Flow, flow variance and stage were not updated during the ETKF analysis. Note: all axes are constant between plots except signal variance. The lower three plots on each figure show the target and validation time expected signal variance given either the upstream, validation or downstream sensor. a) Plot of signal variance 3 hours 40 minutes prior to the validation time. b) Plot of signal variance 3 hours 20 minutes prior to the validation time. c) Plot of signal variance 3 hours prior to the validation time. d) Plot of signal variance 2 hours 40 minutes prior to the validation time. e) Plot of signal variance 2 hours 20 minutes prior to the validation time. f) Plot of signal variance 2 hours prior to the validation time. g) Plot of signal variance 1 hour 40 minutes prior to the validation time. h) Plot of signal variance 1 hour 20 minutes prior to the validation time. 205-209
- Fig. 7.12: Forecast state at sensor 2 at six target times (dark lines). Light lines show background simulations made before data assimilation, whilst black stars are measurements at five minute intervals from sensor 2. 211
- Fig. 7.13: Histogram of background validation time simulated state (top left), followed by histograms of validation time forecast state after assimilating 1 through to 6 measurements from sensor 2 with the EnKF approach. The bottom right plot shows the mean EnKF simulations of state (dark line) and state error variance (lighter line). Also shown are the corresponding ETKF predictions of state error variance. 212
- Fig. 7.14.a-f: Signal variance at the validation time and sensor 8 (RC model node F17) during tide 26. Signal variance is shown for theoretical sensors at 22 locations on the floodplain and 12 locations in the channel for six target times. The background image is a DEM derived from LiDAR data. a) Log signal variance 3 hours 40 minutes prior to the validation time. b) Log signal variance 3 hours 20 minutes prior to the validation time. c) Signal variance 3 hours prior to the validation 224-216

time. d) Signal variance 2 hours 40 minutes prior to the validation time. e) Signal variance 2 hours 20 minutes prior to the validation time. f) Signal variance 2 hours prior to the validation time.	
Fig. 7.15: Signal variance (coloured dots) at the validation time and upstream validation location during tide 26. Signal variance is shown for theoretical sensors at 22 locations on the floodplain and 12 locations in the channel for six target times.	218
Fig. 7.16: Plots showing signal variance estimates when assimilating one measurement during any of the 39 consecutive tidal events or at any channel node in the RC model assuming a measurement variance of 1 m on the x -axis against the signal variance assuming the measurement error to be either 0.5, 0.1, 0.05, 0.01, 0.005, 0.001, 0.0005, 0.0001.	220
Fig. 7.17.a-b: a) Plot of measurement variance against mean signal variance at the target time and validation time at the validation location. The mean is taken from 39 consecutive tides and all channel nodes in the RC model. b) Plot of measurement variance against mean signal variance at the target time and validation time at the validation location (log scales).	221- 222

List of tables

Table 4.1: Summary of Flow duration for EA gauge on the River Crouch at Wickford.	68
Table 4.2 Summary of forecasting data set. All samples collected at 15 minute intervals.	76
Table 4.3: Initial sampling range for input variables calibrated using GLUE.	77
Table 4.4: Table of RMSE between forecasted and measured stage averaged over the forecast period.	88
Table 4.5: r values for forecasted uncertainty with different sampling strategies denoted by: (A) 15 minute sampling interval from sensors 1 to 4, (B) 15 minute sampling interval from sensor 1, (C) hourly samples from sensors 1 to 4 (D), hourly samples from sensor 1. h indicates that the r value was calculated using hourly measurements only, whilst the 'adjusted' results use a linear regression model to remove systematic errors.	95
Table 5.1: Description of adaptive sampling notation.	102
Table 5.2: Selected RC model parameters.	110
Table 5.3: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at seven times between a measurement time t_r and a validation time t_v . From a	115

background forecast t_i an EnKF was used to assimilate one measurement t_r . Forecast error variance was then calculated by propagating an ensemble of RC model simulations to t_v . The ETKF used the same background forecast and the same measurement error to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.

Table 5.4: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at seven times between a measurement time t_m and a validation time t_v . From a background forecast an EnKF was used to assimilate one measurement. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C. 122

Table 5.5: This table contains the correlations between ETKF and EnKF estimates of forecast error variance at seven times between a measurement time t_r and a validation time t_v . The background forecasts were produced assimilating one measurement with the EnKF at t_r , then running an ensemble of RC model simulations to t_v . From the background forecast an EnKF was used to assimilate one measurement at the first target time. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. Correlation A lists the correlation between EnKF and ETKF error variance estimates. Correlation B lists the correlation between the background forecast error variances at each target time and the validation time. Correlation C shows the correlation between signal variance estimates. 131

Table 5.6: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at two times between a measurement time t_m and a validation time t_v . From a background forecast an EnKF was used to assimilate two or three measurements. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same 135

measurement error covariance to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.

Table 5.7: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at two times between a measurement time t_m and a validation time t_v . From a background forecast an EnKF was used to assimilate two or three measurements. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error covariance to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.	136
Table 6.1: Sensor nodes and measurement data.	149
Table 6.2: Selected measurements from tide 21 chosen by adaptive sampling Method 2 in terms of signal variance.	164
Table 6.3: Selected measurements from tide 21 chosen by adaptive sampling Method 2 in terms of least signal variance.	164
Table 7.1: Table describing example sensor locations. Distances upstream of the validation location are denoted by negative numbers.	191

Acknowledgement

The author is very grateful to Professor Peter Atkinson and Dr Craig Hutton for their insightful comments and discussions, enthusiasm for the subject and patience as supervisors.

This research was part funded by the UK Department of Trade and Industries next wave technologies and markets program through the FloodNet project. Data used in this case study was kindly provided by FloodNet, Halcrow Ltd, The Environment Agency of England and Wales, United Kingdom Hydrographic Office and Associated British Ports Marine Environment Research. The author is also grateful to Professor Dave DeRoure (ECS, Southampton), MAC Ltd, IBM and Mervin Smith.

Finally, I would like to thank all the other geography postgraduates for making the Southampton PhD experience so enjoyable and my parents whose unwavering support has been exceptional.

1 Introduction

Timely, reliable and accurate flood warnings have the potential to mitigate some of the negative impacts of flooding (Carsell *et al.*, 2004). Generating these warnings in real-time requires integrated systems that detect an impending event and disseminate alerts to stakeholders in and around the area at risk with sufficient lead times (Khatibi *et al.*, 2003). Observations such as rainfall or river flow taken during the detection step may be used directly as the basis for issuing a warning. Alternatively, a flood forecasting system may be put in place prior to the warning step to make predictions of river state variables such as stage and flow, usually with the aim of increasing the forecast accuracy (Sprokkereef, 2001; Aubert, 2003; Madsen and Skotner, 2005) or extending the lead time of the forecasts to give advanced warnings (De Roo *et al.*, 2003; Werner *et al.*, 2005). These systems often include a numerical model to simulate the expected flows over the forecast lead time and a method for considering current observations.

Data assimilation is the process of incorporating observations into model simulations. McLaughlin (2002) distinguishes between three types of data assimilation problem; interpolation, smoothing and filtering. These are discussed further in chapter 2. There are several data assimilation methods that can be applied to filtering problems, of which the method used in chapters 4, 5, 6 and 7 of this thesis for updating model states with measurement data is just one example.

The concept behind state updating is that updating the current states of a numerical model to form a closer match with real-time measurements will lead to a reduction in forecasted state error compared to that obtained with the uncorrected model (Aubert *et al.*, 2003). This error reduction occurs because the errors in a short forecast are different from those over the longer time periods that are relevant to the simulation model structure and calibration (Young, 2002). For linear systems where there is uncertainty in both the simulation model and measurements, optimal state updating can be achieved by minimising the errors between model state predictions and measurements using a Kalman filter (Kalman, 1960; Refsgaard, 1997; Schreider *et*

al., 2001; Young, 2002). For models in which state changes exhibit non-linear behaviour approximations to this approach such as the extended Kalman filter (Maybeck, 1979) and ensemble Kalman filter EnKF (Evensen, 1994) have been developed and applied to river and coastal flow models (Madsen and Cañizares, 1999; Shiiba *et al.*, 2000; Hartnack and Madsen, 2001; Madsen and Skotner, 2005; Andreadis *et al.*, 2007). Ensemble based methods where model error statistics are derived from multiple state simulations are usually easier to implement than the extended Kalman filter, especially as the degree of process non-linearity increases (Evensen, 1994; Madsen and Cañizares, 1999). Nevertheless, characterising and quantifying the sources of uncertainty in the simulation model and generating ensembles with appropriate error statistics remains a major problem (Butts *et al.*, 2005).

The accuracy of a forecast depends on the accuracy of its simulation model and the efficiency of the data assimilation algorithm (Refsgaard, 1997). Where the simulation model accuracy will depend on the errors introduced as a result of uncertainty in model structure, boundary conditions and parameters, whilst the efficiency of the data assimilation algorithm, as a means of reducing state uncertainty, is important because it will influence the reduction in state uncertainty brought about by assimilating measurements, which themselves will contain errors. These measurements will be collected according to some sampling framework, which may not be optimised to provide data with as much information content as possible. For spatial sampling problems geostatistical methods can be used to optimise sampling design based on an underlying assumption of the degree of spatial dependence in state values. However, in sequential data assimilation, where time forms an important factor in state dynamics, methods of real-time sampling design based on the Kalman filter have been developed. In the field of meteorological modelling Monte Carlo adaptive sampling methods such as the ensemble transform Kalman filter (ETKF) (Bishop *et al.*, 2001) have been developed and applied to address a similar set of issues that would be expected in the field of flood forecasting. Namely, how to approximate an optimal spatio-temporal sampling scheme with a given set of resources when the underlying process is known to be non-linear? To the author's knowledge the application of these adaptive sampling

methods (which are introduced in chapter 2) in a flood forecasting context has not been previously attempted.

An additional motivation behind researching optimal real-time sampling design relates to recent developments in environmental wireless sensor networks (De Roure, 2005). These pervasive sensor networks can be reconfigured in real-time, thus facilitating adaptive sampling, usually with the aim of saving power or reducing data transmission. The FloodNet project, which provided the measurement data used in this study, created an experimental wireless sensor network specifically designed to monitor river flows and levels. It required a means of prioritising which measurements to collect and transmit from the field in real-time in order to save battery power.

1.1 Thesis aims

Flood forecasting systems have the potential to improve the management of flood events and thereby mitigate some of their negative impacts. Whilst it is recognised that flood forecasting systems require procedures to disseminate and react to forecasts, research in this thesis was limited to flood detection and modelling. In recent years, stochastic predictions of state variables and the delivery of probabilistic rather than deterministic predictions have been sought by environmental managers interested in hydrological forecasting (Krzysztofowicz, 2001). When flood forecasting, the principal advantage of a probabilistic approach is the ability to calculate flood risk based on the probability of a particular magnitude of flood event and the expected losses from such an event. This allows the value of risk avoidance and risk reduction strategies to be assessed against their cost. For example, not building a house on a floodplain is an example of risk avoidance, whilst installing a forecasting system is a means of risk reduction, both of which may be unnecessary if the risk is low. The desire for stochastic rather than deterministic predictions was a core aim of this thesis which applied to all subsequent aims. The overall aims of the thesis were as follows:

1. Build a forecasting model based on:
 - a. A simulation model capable of simulating states at unmonitored locations.
 - b. A data assimilation algorithm that can utilise real-time measurement data that are spatiotemporally discontinuous and corrupted by errors.
2. Assess the ability of the forecasting model to simulate state error propagation
3. Develop a method to quickly estimate forecast error variance simulated by the forecasting model.
4. Develop and evaluate an adaptive sampling framework that could be applied to a network of wireless sensor nodes with limited power.
5. Develop and test a methodology to design optimal sampling frameworks for flood forecasting.

These general aims were addressed by four research chapters each of which had its own set of specific aims.

In chapter 4 'Flood inundation model updating using an Ensemble Kalman filter and spatially distributed measurements' the principal aim was to utilise a state updating method and a hydraulic model to estimate stage and stage uncertainty, over a 4 hour forecast period, at specific validation locations. Integral to this aim was the accuracy of state forecasts, and uncertainty estimates, and the robustness of the model when subjected to different spatio-temporal sampling regimes.

In chapter 5 'Evaluating the utility of the ensemble transform Kalman filter for adaptive sampling when updating a hydrodynamic model' the aim was to evaluate a method, known as the ensemble transform Kalman filter (ETKF) (Bishop *et al.*, 2001), to quickly estimate the error variance of a stage forecast given one or more possible stage measurements from the sensor network.

In chapter 6 'Adaptive space-time sampling with wireless sensor nodes for flood forecasting' the aim was to develop and evaluate adaptive sampling techniques that facilitate a reduction in the need to transmit real-time data from a network of sensor nodes in the field (called FloodNet) to the forecasting model introduced in chapter 4.

In chapter 7 'Optimal sampling design with an ensemble transform Kalman filter' the aim was to test the response of the ETKF to several test scenarios, useful for the design of a sensor network deployment. These tests included an investigation of the spatio-temporal changes in signal variance from theoretical sensors both within the main river channel and on the floodplain, as well as measurements with different error variances.

These research chapters were supported by a literature review in chapter 2 and a methods review in chapter 3.

2 Literature review

This chapter reviews some of the research conducted on two of the key components of a flood forecasting system, namely the simulation of system dynamics (section 2.1) and data assimilation methods for updating simulations with real-time measurements (section 2.2). Some of the recent research on adaptive sampling techniques is also reviewed in section 2.3. Each of these sections represent major areas of research, although not always centred on the sphere of flood forecasting. Therefore, the section on simulation models is limited to distributed hydrodynamic models, whilst the section on data assimilation is limited to sequential data assimilation algorithms based on the Kalman filter. In turn, this limits the section on adaptive sampling to research that is compatible with the previous sections.

2.1 Distributed flood models

Water flow is a spatially three-dimensional temporally varying process. The dynamics or hydrodynamics of this process are well known and can be described by the three-dimensional Navier-Stokes equations (Hervouet and Van Haren, 1996). Solving these equations is difficult and computationally demanding, even for small real world problems. Therefore, it is common for hydrodynamic flow models to be based on simplifications of these equations when simulating within channel flows and flood inundation. Despite three-dimensional processes having been found to be important at the interface between channel and floodplain flows (Knight and Shiono, 1996), floodplain flows are generally thought of as two-dimensional processes, whilst one-dimensional models are often regarded as being adequate for simulating channel flow. Simplified process representations include, in decreasing complexity, codes based on: two-dimensional finite element (Bates *et al.*, 1992), simple finite volume (Horritt, 2004), two-dimensional diffusive wave (Bates and DeRoo, 2000; Yu and Lane, 2006), one-dimensional dynamic wave (Moussa and Bocquillon, 1996) and kinematic wave models (Bates and DeRoo, 2000; Singh, 2001).

Two-dimensional finite element solutions have been shown to be capable of representing significant floodplain flow complexity (Bates *et al.*, 1992, 1997, 1998; Nicholas and Mitchell, 2003). Although, Horritt and Bates (2001) found that a kinematic wave model of channel flows coupled to a simple storage cell raster model of floodplain flows (LISFLOOD-FP, Bates and DeRoo, 2000) performed as well as a more complex two-dimensional finite element method, given the available calibration and validation data. The principal barrier to a more precise assessment of the validity of the flow dynamics and an assessment of the accuracy of state simulations was the lack of distributed calibration and validation data with sufficient resolution and temporal coverage. Horritt and Bates (2002) compared LISFLOOD-FP with a one-dimensional model (HEC-RAS) and a two-dimensional finite element model (TELEMAC-2D), finding that the HEC-RAS and TELEMAC-2D models made equally accurate predictions of flood inundation area given calibration data in the form of flood wave travel time or inundated area. However, the raster cell model (LISFLOOD-FP) was only able to make comparable predictions of inundated area when calibrated against inundated area data. The calibration of the raster cell model for flood inundation impaired its ability to predict flood wave travel times because the optimal parameterisation of the model for inundation area and flood wave travel time were found to be in different areas of the parameter space.

One-dimensional channel flow models are popular because they are computationally simpler and require fewer parameters than their two-dimensional counterparts (Horritt and Bates, 2001). The one-dimensional Saint-Venant equations can be used to simulate unsteady open channel flow in the along-channel directions only, thus omitting lateral and vertical flow. Solutions to the full Saint-Venant equations are known as dynamic wave models. They account for local acceleration, convective acceleration, pressure, gravity, friction, eddy loss, wind shear and lateral inflows (Wilson, 2004). Simplifications to the dynamic wave model include diffusive wave models that neglect local and convective acceleration terms (Chow, 1988) and kinematic wave models that neglect local acceleration, convective acceleration and pressure terms (Singh, 2001). Neither of these simplifications is suitable in situations where significant backwatering is expected (e.g. where there are interactions between tidal and fluvial flows).

Codes based on dynamic wave models such as ISIS Flow (HR Wallingford) and MIKE 11 (Danish Hydraulic Institute) can be applied where waves propagate both downstream due to catchment runoff and upstream due to tidal forcing. When the domain of interest includes areas of floodplain it is common for one-dimensional models to be coupled to a separate model of floodplain flow that calculates flow based on the water surface elevation and Manning's equation. This can be done on either a regular grid (e.g. LISFLOOD-FP) or for irregular reservoirs (e.g. ISIS Flow). More details are given in section 3.1.2.

Further to the choice of process representation, and thereby model code, the spatial and temporal domain over which state is simulated must be discretised due to finite computational resources. Thus, for a given code the amount of computation and number of parameters required by the model will increase as the spatial resolution becomes finer. That is, it becomes more difficult in terms of computer time and memory to implement models as the size of the domain increases, the discretisation of the domain becomes finer and the process representation becomes more complex. Unfortunately, the spatial resolution of the model elements has been found to affect inundation extent estimates due to the topographic smoothing and loss of small scale features that accompany coarser resolutions (Yu and Lane, 2006), especially in urban areas where the presence of small linear features such as hedges and curb stones can have significant effects. This means that parameter values are often not directly transferable between model codes, model structures and model applications.

2.1.1 Data sources and issues arising therein.

Some of the most important parameters of hydraulic models such as the wetted perimeter, channel width, slope etc. are those which relate to the topography of the model domain. Digital elevation models (DEM) have become the principal methods of representing these parameters. Taking the measurements needed to generate a DEM of a natural river and its floodplain is not trivial, since thousands of elevation observations are usually required to provide the necessary spatial detail, especially since topographic variation as small as 10 cm can have a significant effect on

inundation extent (Anderson *et al.*, 1996). One approach applicable to small domains or providing cross sections for one-dimensional models is to survey the area of interest with a total station (Nicholas and Mitchell, 2003) or differential global positioning system. These methods provide point elevations that can be targeted by the surveyor to breaks in slope so that important topographic features are retained at minimal sampling expense. However, ground surveys are impractical over large areas (such as floodplains) and for measuring more than a few thousand points because the surveyor must physically move between sample points.

Data for generating DEM's of floodplain topography or areas over a few hundred square metres can come from a number of sources including contour maps, interferometric synthetic aperture radar (InSAR) and airborne laser altimetry (LiDAR). Wilson and Atkinson (2005) found that contour data lead to an overly smoothed representation of floodplain topography, which omitted many features that might act as important controls on floodplain flow, whilst currently available InSAR data were found to give an overly rough representation of floodplain topography. Both data were found to demonstrate significant problems when used to estimate flood risk using a raster cell inundation model. In contrast LiDAR can provide point measurements of surface elevation over large areas, often at a finer spatial resolution than the grid, mesh or node spacing of the hydrodynamic model (Marks and Bates, 2000; French, 2003). Prior to the availability of LiDAR data, providing suitable topographic parameters was a major barrier to the development of two-dimensional hydrodynamic models of out-of-bank flows at the reach scale (Bates *et al.*, 1998).

One difficulty with LiDAR data is that the elevation points collected relate to the surface off which the laser pulse returned. Therefore, raw LiDAR data can only be used to generate digital surface models (DSM), which include vegetation, trees and buildings, along with the ground level points needed to generate a DEM. A number of image segmentation algorithms have been developed to separate the underlying topography from the features present in the raw data (Cobby *et al.*, 2001). The motivations behind removing these features include:

1. They are often much smaller than the resolution at which the model simulation is conducted making it impossible to account for all of them as surface features.
2. They may be porous rather than solid surfaces that must be overtopped for flow to occur, especially in the case of vegetation.

The DSM features removed when generating a DEM are often either assumed to have little influence on flow and ignored or treated at a sub-grid scale using roughness or porosity terms. Current LiDAR technology has been found to predict the elevation of short vegetation (such as grass) with a root mean squared error of 0.14 m and the elevation of the underlying surface to around 0.17 m (Cobby *et al.*, 2001). Cobby *et al.* (2001) also found that the accuracy of ground elevation estimates decreases over dense and tall vegetation. Recently, a method for fusing Ordnance Survey map data with LiDAR data from urban areas has been developed (Mason *et al.*, 2007) allowing the location of features such as building edges, roads, hedges and pavements to be defined more accurately.

Removing features such as trees and buildings from raw LiDAR data or any DSM may result in gaps in the spatial coverage. It may be necessary to interpolate between these gaps using methods such as inverse distance weighting or kriging (Lloyd and Atkinson, 2006) before parameterisation of the hydrodynamic model. If the simulation model operates over a regular grid it will also be necessary to interpolate from the LiDAR point data to a regular grid (Lloyd and Atkinson, 2002). For hydraulic models with flexible grids, methods have been developed to generate model discretisations from raw topographic data that take account of local topographic variation (Bates *et al.*, 2003), vegetation cover (Cobby *et al.*, 2003) and urban features (Mason *et al.*, 2007).

Measurements of state variables such as stage and flow are often required for the purpose of calibration and validation (explained in more detail in the next section). Distributed flood inundation models are difficult to calibrate and validate because it is difficult to take spatial and temporal observations of flood extent, stage and flow. Several studies such as Horritt and Bates (2001; 2002) have cited a lack of suitable

validation and calibration data as a barrier to further model development. However, progress is being made. Recently, multiple images of flood inundation at different times have become available for a single flood event (Bates *et al.*, 2006). These have allowed dynamic changes in inundation extent to be validated for the first time, although, more data is still required. Furthermore, the use of *in situ* or ground-based measurements is becoming more feasible with the advent of simpler and cheaper sensors, gauges, and loggers (Troch *et al.*, 2003). The type of sensor nodes introduced in section 6.3.1 is an example.

Hunter *et al.* (2005) investigated the utility of five data sources for calibrating flood inundation models derived from; two gauges measuring stage internal to the model domain, the discharge at the downstream boundary of the model, a shoreline estimate from a satellite radar image, a shoreline estimate from an air photo and a ground survey of maximum free surface elevation conducted after the event. Internal stage values and the shoreline estimates were found to offer considerable potential for reducing uncertainty in parameter specification, whilst the discharge at the downstream boundary was less effective.

2.1.2 Parameter estimation, calibration and validation

For most applications the quickest and easiest method of parameter estimation is to guess parameters based on the modeller's experience. An educated guess can usually be made based on prior experience from other locations or applications of the simulation model that appear to have similar properties. In some cases, conceptual models may have been developed to aid this process. Estimation of the roughness coefficient in Manning's equation based on lookup tables is one example where conceptual models and user experience are commonly used. Many textbooks on river hydraulics or geomorphology will quote roughness values suitable for parameterisation of different geomorphologic features such as gravel bedded rivers, open grasslands or floodplain forests, to list just a few.

If parameters relate to something physically meaningful and observable it may be possible to determine values by taking measurements in the laboratory or field. In

general, measurements collected in the field tend to be made on small supports, which are sparsely spread over the domain of interest, where the support is defined the area, shape and orientation of space represented by a sample (Curran and Atkinson, 1998). As a consequence parameters are often estimated from measurements that are inherently uncertain with disparities existing between the scales of measurement, the scale of the process and the scale of the model (Atkinson and Tate, 2000). Where measurements are collected sparsely in space, parameters may need to be interpolated or scaled to meet model requirements, introducing additional uncertainties.

For models that are distributed over large spatial domains remote sensing provides an attractive means of estimating model parameter values, mainly due to its low cost and continuous spatial coverage relative to ground survey. A by-product of the LiDAR image segmentation algorithms introduced in the previous section is a set of point data representing above surface features. Mason *et al.* (2003) used this information to classify the floodplain into regions of differing vegetation height, which were subsequently used to specify spatially distributed friction coefficients. Wilson and Atkinson (2007) used Landsat TM imagery to classify different vegetation types from which spatially distributed friction was estimated. Although, distributed friction has a small effect on inundation depths and inundation extent for a study site on the river Nene, UK.

One problem with measuring or guessing parameter values is that parameters found to be appropriate at one location or for a particular model structure are by no means guaranteed to be appropriate in another context, when subjected to validation. For example, it is now widely established that calibrated bed roughness is scale dependent (Lane, 2005) both in terms of the model architecture and the volume of flow (Romanowicz and Beven, 2003). This means that there may be one or more region(s) of parameter space that are more accurate simulators of state measurements than the physically meaningful value of the parameters, which themselves may have been imprecisely defined and subject to errors. These regions of parameter space occur because the model process representation and structure are a simplification of an unknown reality that has only been partially sampled by the measurement data,

which will also contain errors. In fact all measurements, process representations and state estimates will contain errors as they are all simplified representations of reality.

The process of identifying parameter values that lead to the most accurate simulations of observed conditions is known as model calibration. When model parameters are calibrated to some measurement data they are generally referred to as effective parameters rather than physically meaningful, because they may no longer pertain to the properties which they represent. Cunge (2003) argued that the calibration of model parameters to a particular event destroys the predictive capabilities of the model for other magnitude events. For this reason, Connell *et al.* (2001) applied published parameter values for two ‘un-calibrated’ distributed flood inundation models, despite both models underestimating flood level and extent. However, it may be possible to avoid reducing a models predictive capability should enough data be available on events of differing magnitude to conduct a split-sample calibration and validation approach, where the data is divided into two sets of independent measurements. The first measurement set is then used to calibrate the model, whilst the second is reserved to validate the simulations made by the calibrated model.

Here, it will be argued that there is not a single set of ‘correct’ parameter values, but that many possible combinations of parameter values will be reasonable simulators of the system, given some validation data, thereby endorsing the belief that distributed hydrodynamic models are equifinal. The concept of model equifinality (Beven and Binley, 1992) rejects the notion of an optimal model structure, set of parameters and forcing terms. Instead, many different combinations of model structures, parameters, and forcing terms within a pre-defined space will perform equally or almost as well as simulators of state, when subjected to measurement data. On this basis, it can be argued that the idea of an optimal parameter set is generally ill-founded and should be rejected, if uncertainty in a model is to be recognised explicitly. Allowing multiple parameter sets and/or forcing terms to be considered as simulators of the system within a specific model structure, it follows that multiple state predictions will be made. The challenge is, therefore, to represent the uncertainty in model forcing terms and parameters such that they feed through to

create a range of predictions that can be considered to represent prediction uncertainty.

Several hydrological studies in applications such as rainfall-runoff modelling (Freer *et al.*, 1996; Lamb *et al.*, 1998; Cameron *et al.*, 1999; Beven and Freer, 2001), two-dimensional river flow modelling (Hankin *et al.*, 2001), soil saturation (Franks *et al.*, 1998), one-dimensional hydrodynamic flow (Pappenberger *et al.*, 2005a; Werner, 2004), flood inundation prediction (Aronica *et al.*, 1998, 2002; Bates *et al.*, 2005; Romanowicz *et al.*, 1996), and atmospheric pollutant deposition (Zak and Beven, 1999; Page *et al.*, 2004a, 2004b) have found that for a given model structure many different parameters and forcing term sets perform almost as well as simulators of the system, given the available validation data. To a lesser extent equifinality in model codes and structures (Butts *et al.*, 2004; Georgakakos *et al.*, 2004) has also been investigated.

Hankin *et al.* (2001) point out that for distributed models of channel and floodplain flow, equifinality is not surprising as there are typically many degrees of freedom in the specification of parameters and forcing terms, within even a single model structure. Often there are many more parameters than can be supported by the available calibration data. Thus, distributed models are generally overparameterised (Beven, 1989) in that the parameters of each model element, cell or node (depending on the model type) cannot be treated in isolation from its spatial neighbours. To address this problem parameters tend to be lumped together either spatially or according to some system of classification. For example, estimates of Manning's roughness coefficient are frequently calibrated on a global basis or by partitioning the model into a few parameter regions based on land cover (Wilson and Atkinson, 2004) or by separation of channel and floodplain (Nicholas and Mitchell, 2003). Furthermore, it is often necessary to fix some parameters or link parameters using fixed ratios to prevent the problem getting out of hand (Refdgaard, 1997), even though this is tantamount to saying there is no uncertainty in that particular parameter. The selection of 'fixed' parameters is usually a combination of how well the modeller thinks a particular parameter can be estimated from observations and the sensitivity of model output to changes in that parameter. In addition to

parameterisation problems errors can be expected to propagate from incomplete and biased model structure, random and systematic errors in model forcing terms, and random and systematic errors in validation and calibration data (Butts *et al.*, 2004).

2.2 Data assimilation and flood forecasting

Data assimilation is the process of optimising the extraction of information from observations for combination with numerical models (Babovic *et al.*, 2001).

Depending on the problem to be solved, data assimilation methods can be sorted into three categories: interpolation, smoothing and filtering (McLaughlin, 2002).

McLaughlin (2002) defines each category as follows; interpolation characterises a time-invariant system at a single measurement time, smoothing characterises a time-dependent system through a time interval, and filtering characterises a time-dependent system at the most recent measurement time. This thesis addresses the problem of forecasting river level (stage) and flow from the most recent measurement time. To this end, data assimilation will be considered predominantly from a filtering viewpoint because it provides the most efficient means of assimilating real-time data into simulation models.

When filters are combined with simulations from numerical models the implementation of the filtering process is commonly referred to by the term model updating. Thus, numerical model simulations, such as those that predict river state variables such as flow and stage, can be updated using measurement data by a filter. A forecast is only made when the updated model is used to simulate to a time ahead of the most recent measurement. It is usually assumed that the accuracy of forecasts produced by such systems depends on the accuracy of the simulation model and the efficiency of the updating routine (Refsgaard, 1997). Updating methods can be split into four categories depending on the variables being modified (Refsgaard, 1997). These are:

1. Input updating
2. Parameter updating

3. Error prediction or correction
4. State updating.

Each of these model updating categories will now be introduced in more detail.

Input updating

Input updating methods involve generating new estimates of input data and then running them through the simulation model. Input updating for river flow models may be as simple as using the most recently available measurement data from upstream of the forecast location and using the time lag between the hydrographs at the upstream and forecast location to determine the forecast lead time. However, forecast lead times can be significantly increased by taking updated outputs from hydrological and meteorological models as input data (Knebl *et al.*, 2005).

Numerical models that might act as inputs to hydraulic models have received research attention, including forecasting models of rainfall using weather radar (Moore *et al.*, 2004), soil moisture and catchment properties (Walker *et al.*, 2003; Reichle *et al.*, 2002), catchment runoff (Aubert *et al.*, 2003; Schreider *et al.*, 2001; Lee and Singh, 1998; Todini, 1994), snowmelt, river flow (Shiiba *et al.*, 2000; Hartnack and Madsen, 2001) and tidal surge (Cañizares *et al.*, 1998). Research has recently been undertaken into model coupling and the propagation of uncertainty between models (Pappenberger *et al.*, 2005b).

Parameter updating

Parameter updating methods seek to update the parameters of simulation models based on observed errors. For the purpose of river flow forecasting these methods are less common than the other three types, but, have been utilised by simulation models based around transfer functions (Young, 1984). An example can be found in Lees *et al.* (1994) where a single input-single output transfer function was used to predict downstream water level based on upstream observations along the River Nith at Dumfries. A gain parameter was included in the transfer function which was then estimated in real-time by comparison with downstream observations using a

stochastic time variable parameter estimation algorithm (Young, 1984). The adaptive gain approach allowed seasonal changes in the catchment characteristics and river dynamics to be accounted for by the transfer function, resulting in an increase in forecast accuracy compared to the equivalent non-adaptive model. There may be similar benefits in updating certain distributed hydrodynamic models parameters.

Error prediction

Error prediction methods analyse past errors between observations and model predictions to predict future model errors (Dobson and Davies, 1990). One example is the autoregressive moving average (ARMA) technique used as part of the River Flow Forecasting System (RFFS) (Moore *et al.*, 1990). More details on ARMA type models can be found in Bras and Rodrigues-Iturbe (1985). It could be argued that error prediction is conceptually unappealing in the sense that it does not make changes to model state that may have a significant influence on post update state dynamics (Moore *et al.*, 1990). However, error prediction is usually easy to implement and does not seek to intervene in any way with the operation of the simulation model, unlike the other three methods. Cunge (2003) argues that this is beneficial because it limits the scope for corrections or adjustments to the model itself that do not have a physical basis.

State updating

State updating methods can be used to adjust internal model state based on observed errors and can be coupled with input and parameter updating algorithms (Refsgaard, 1997). The principal aim is to estimate state as accurately and precisely as possible at the initialisation of the forecast. The correction of physically meaningful state estimates is appealing because in theory this has the potential to provide a more accurate starting point from which to simulate future model state than a non-updated model (Aubert, 2003). More accurate definitions of initial states leads to more accurate forecasts because the errors in a short forecasts are different from those over the longer time periods that were relevant to the simulation model structure and calibration (Young, 2002). Bearing in mind the above point made by Cunge (2003) regarding the benefits of error prediction, state updating assumes that the dynamics

of the simulation model will not be significantly affected, in an adverse manner, by the update process. This includes numerical instability which may result from the specification of physically unrealistic initial conditions and an assumption that the model parameters are suitably based on the physical processes that their location in parameter space does not change markedly during the short time period after the update.

2.2.1 Flood forecasting

Systems designed for processing real-time data to produce flood forecasts tend to be individually tailored to meet various requirements such as the spatial scale over which they are required to operate, the desired accuracy of forecasts and the length of the lead time (Werner *et al.*, 2005). Pre-warning systems such as the European Flood Forecasting System (De Roo *et al.*, 2003), have the ability to provide lead times of 5-10 days through the utilisation of medium range numerical weather prediction models. However, at shorter lead times or for longer rivers, operational flood forecasting systems begin to take account of measured hydrological and meteorological data to enable a more accurate forecast than can be achieved with the pre-warning systems. Examples include the River Flow Forecasting System (Moore *et al.*, 2004), and the Hydrologic Engineering Center's Hydrologic Modelling System/ River Analysis System (HEC-HMS/RAS) (Knebl *et al.*, 2005). In large river basins a combination of rainfall-runoff models, routing models and gauge data will generally give adequate lead times, such that flood warning can be issued. For example, the FloRIJN system on the Rhine has a 3-4 day lead time between flows observed upstream and their downstream impact (Sprokkereef, 2001). In this thesis, only the distributed hydrodynamic channel and floodplain flow component of basin hydrology will be investigated further such that a single process model could be used when running the forecasting model, avoiding the need to run and couple other model types.

The use of dynamic physically-based spatially distributed models for flood inundation and flow prediction has led to interest in the application of these models in a flood-forecasting context. The detail with which processes are represented

within a domain of interest both in terms of spatio-temporal resolution and code complexity is often driven by the application. For flood-forecasting applications the speed of model execution is critical if forecast with sufficient lead-times are to be issued (Cunge, 1980). This inevitably leads to a preference for the simplest model capable of simulating state variables to a desired accuracy. For operational systems one-dimensional flood models are often used (if at all) in preference to more complex process representations because they require less input, validation and calibration data and are computationally less intensive. Many distributed flood models were developed primarily with the aim of producing a single deterministic prediction of flood inundation given a set of driving forces and optimal or best guess parameters that were either estimated through calibration or determined from field measurement. Whilst this type of model can be applied easily in a flood-forecasting context, through the input of the most recently available forcing data (see example in Knebl *et al.* 2005), conveying a realistic estimate of forecast uncertainty and taking advantage of real-time data for state correction is difficult.

2.2.2 State updating with Kalman filters

For linear models, a Kalman filter (Kalman, 1960) provides the optimal solution to the state correction problem and may be extended to give an approximate solution for non-linear models (Refsgaard, 1997). The extended Kalman filter (EKF) (Maybeck, 1979) is one such method for non-linear systems. Aubert *et al.* (2003) used an EKF to assimilate soil moisture measurements into a conceptual rainfall-runoff model, finding that the method could increase the accuracy of streamflow forecasts, especially during flood events. However, the high computational cost incurred when obtaining covariance estimates is a significant drawback of this method, which several studies have sought to address using reduced rank square root filters (Verlaan and Heemink, 1997; Madsen and Cañizares, 1999; Cañizares *et al.*, 1998). This filtering method was applied by Shiiba *et al.* (2000) to a one-dimensional hydraulic flood routing model to forecast stage and discharge. The model was found to make precise estimates of channel state without an excessive increase in computational time relative to the deterministic model.

In an ocean modelling context, Evensen (1992) found that unbounded error growth as a result of approximations in the EKF error covariance propagation equations caused instability and closure problems, even though the model under investigation was only weakly non-linear. Evensen (1994), therefore, proposed a Monte Carlo based approach to approximate the KF, which has since become known as the ensemble Kalman filter (EnKF). Hartnack and Madsen (2001) used an EnKF to remove various errors in boundary conditions from a synthetic one-dimensional hydraulic model (MIKE 11) with simulated input data sets. Madsen and Skotner (2005) continued this theme by using a fixed user-specified Kalman gain (weighting given to errors between observations and model states) for measurement locations, as a computationally cost-effective alternative to propagating an ensemble of model runs. The method was applied to an operational flood forecasting model of Metro Manila, the Philippines. Another computationally cost-effective method would be to compute the Kalman gain from an EnKF off-line and use a steady Kalman filter implementation (Cañizares *et al.*, 2001). Recently, more efficient square root implementations of the EnKF that do not require measurements to be perturbed have become available (Tippett *et al.*, 2003; Evensen, 2004). Andreadis *et al.* (2007) used the square root implementation of the EnKF from Evensen (2004) to recover depth and discharge estimates for a reach of the Ohio River by assimilating synthetic satellite altimetry measurements into an ensemble of corrupted simulations from a LISFLOOD-FP based model. The synthetic measurements were sampled from an arbitrary 'truth' simulation using the Jet Propulsion Laboratories Instrument Simulator. Andreadis *et al.* (2007) concluded that estimates of discharge from altimetry measurements would be best provided through a data assimilation strategy built around a river hydrodynamic model; although no alternative methods were presented.

2.3 Adaptive sampling

There is always a cost of some form associated with the collection, storage and processing of data. Data on the environment form a representation of reality through the sampling framework used in the collection of data (Atkinson and Tate, 2000). It

is therefore important that this sampling framework maximises the delivery of information given the resources available for data collection. The sampling framework within a domain of interest has several parameters that when changed can influence information content of data. These include: the density of samples, the sampling scheme, the sample size and the support of each sample. The sampling scheme, density and sample size can be split into the spatial sampling framework and the temporal sampling framework. Snepvangers *et al.* (2003) states that the spatial and temporal domains are fundamentally different and that time should not be thought of as an additional dimension to the spatial sampling problem. This is because observations over space represent a multidimensional state that shows a degree of autocorrelation depending on the scale of spatial variation. Observations over time are one-dimensional and describe the evolution of state, which is dependent on the previous state and the factors driving the state change.

For many applications (including flood forecasting) the resources available for data collection will be limited. When considering the deployment of resources to provide measurements to an updating routine there are numerous decisions to make that relate to the sampling framework. For example:

1. The number of sensors to deploy needs to be decided as this is likely to determine the sample size and density. This number is likely to be constrained by the cost of the sensors and their intrusiveness, which in turn is influenced by the type of sensor used. The precision and accuracy of the sensor measurements should also be taken into account. For example, installing a weir in a river channel may help make more accurate estimates of the stage-discharge relationship at a monitoring site than deploying a pressure transducer in the existing channel but will be more expensive and invasive.
2. The location of sensors and rate at which they sample will determine the spatio-temporal pattern of the data, which should ideally be informed by the spatial and temporal variability of the state being observed. For example, the sampling rate required to reconstruct flow variation past a point on a lowland river, where flow changes gradually over time, would be different to that for an upland river where

flow is known to change rapidly. When deciding the spatial location of samples, spreading them evenly throughout the domain being monitored may be a good first approximation of an efficient sampling strategy. However, the presence of floodplain storage areas and tributaries at certain locations may cause uncertainty to propagate in a manner that is not easily approximated by distance. Furthermore, it may not be possible to deploy sensors in certain locations.

3. The optimal location of sensors and the required sampling rate may not be constant over time. For example, as a flood wave travels along a river the optimal location for collecting data on the flood peak will change. Furthermore, errors in measurements may change with the state being measured.

In the spatial domain geostatistical methods have been utilised for around 25 years to make decisions about spatial sampling strategies (Burgess and Webster, 1980; Heuvelink and Webster, 2001). These strategies usually have the objective of finding the smallest sample size that constrains the uncertainty in the phenomena under observation within acceptable limits. Assuming that there are no abrupt changes in spatial structure over the domain to be sampled, the most efficient sampling scheme in a two-dimensional space will be a systematic one based on an equilateral triangular grid (Burgess *et al.*, 1981). Similarly, when sampling over time the most efficient method where no prior information is available will be a systematic one with equal sampling interval.

However, in a non-linear dynamic system where state and state uncertainty varies sampling with an even spatio-temporal coverage may not be the most efficient framework, particularly if the covariance between points in space and time varies with state. Furthermore, extremes of state (that occur infrequently and space and time) may be of more interest, especially if the state extreme poses a risk to some asset. For example, above a particular volume of flow a river may burst its banks and flood the area adjacent to it. As these volumes of flow occur infrequently in space and time data collected at certain times and locations will be more informative regarding the magnitude of the flow than others. Furthermore, it may be important to

have an accurate estimate of the flow magnitude at locations where valuable assets reside, but not at locations where overbank flows will cause little or no damage.

Robinson and Glenn (1999) define adaptive sampling as measurements whose distribution in time and space are adjusted. By taking this very general definition almost any sampling scheme that changes over space or time when historical or prior information is taken into account could be called adaptive. In a forecasting context the objective of the sampling scheme should be such that data is collected in a way that facilitates future forecasts meeting some objective. Ideally this should be done in an optimal manner, given the *a priori* information available in the current or most recent forecast. If the optimal sampling scheme that achieves this objective changes over time given the most recent *a priori* forecast then sampling will need to be adaptive.

Adaptive sampling strategies may include both a routine component and an adaptive component (Bishop *et al.*, 2001). The routine component comprises measurement devices that are fixed spatially and temporally, such as satellite remote sensing platforms and traditional data loggers. The adaptive component comprises devices that can vary their sampling such as; pervasive sensor networks (DeRoure, 2005), robotic vehicles (Curtin *et al.*, 1993; Rahimi, 2003), people, airborne remote sensing platforms etc. The Long-term Ecosystem Observatory (LEO) (Glenn *et al.*, 2000) is an example of an observational system where adaptive sensor networks have been deployed 9 km off the coast of Tuckerton, New Jersey, USA. The network allows a variety of sensors to be controlled from remote locations via electrical or fibre-optic cables on the sea floor. Another oceanographic application of adaptive sampling was the Haro Strait experiment, where a fixed network of buoys was deployed along with two autonomous underwater vehicles (AUV) (Curtin *et al.*, 1993) to monitor the boundary between fresh and saline water in the Haro Strait. The AUV were controlled by a feedback loop from an onshore data-analysis effort (Nadis, 1997).

One approach to adaptive sampling is to target measurements to minimise state error variance. In geostatistics this can be done using kriging, by targeting measurement

locations with the greatest kriging variance. For spatio-temporal problems a Kalman filter based approach can be used.

2.3.1 Adaptive sampling with Kalman filters

As mentioned previously, the KF is a recursive optimal means of estimating the state of a linear process by minimising the mean squared error of the state estimate (Kalman, 1960), which can be extended to give approximate solutions for non-linear models or applied to error statistics generated from state ensembles (Evensen, 1994). KF's provide a framework for propagating different sources of uncertainty through hydrological models (Butts *et al.*, 2005), whilst incorporating uncertainty in measured properties. Thus, KF-based approaches to state updating have several properties which make them attractive tools for supporting adaptive space-time sampling. Four properties are introduced here:

1. Estimates of state error covariance can be used to identify locations in the model's state space with the greatest uncertainty at a particular point in time. Observations can then be targeted to these locations.
2. The effects of reducing state uncertainty through the assimilation of observations can be propagated to future 'validation times', which may be of more interest than the present time. Furthermore, any location or region in model state space may be used as a 'validation location', regardless of where observations are taken.
3. The filter can estimate state uncertainty for a sequence of 'what if' scenarios, whereby a given observation may be considered in the light of other observations, allowing efficient comparisons between different sampling routines. Sampling components that are not adaptive can also be incorporated into estimates of error covariance before the adaptive component is designed using a technique known as serial assimilation (Bishop *et al.*, 2001).
4. The contribution of a particular observation and its uncertainty in reducing state uncertainty can be quantified in terms of its effect on error covariance (the reduction

is hereafter termed signal variance), allowing rigorous comparisons to be made between sensor deployment options. Comparisons can also be made between deploying competing instrument types provided that the error statistics of each instrument are known.

For a linear process the reduction in forecast state error variance or ‘signal variance’ (Majumdar *et al.*, 2002) brought about by assimilating a possible measurement set can be calculated prior to the measurements being taken, provided that the covariance between measurement errors is known. To calculate signal variance the likely magnitude of the state measurements, which are needed to issue an updated forecast, need not be known. Instead it is only necessary to define possible times when and locations where measurements could be taken and the error covariance between measurements. These properties are usually available prior to measuring, unlike the magnitude of the measurements which by definition are unknown until they are taken. Repeatedly calculating signal variance for different measurements or measurement sets provides a quantitative framework for targeting measurements.

The linear KF may not be suited to applications involving non-linear models. An obvious alternative would be to use an ensemble based approach. Bishop *et al.* (2001) developed a version of the KF known as the ensemble transform Kalman filter (ETKF) to assess the ability of a fixed observational network and adaptive airborne reconnaissance missions to reduce forecast variance in a non-linear atmospheric winter storm model.

2.3.2 The ensemble transform Kalman filter

Along with other Kalman filters the ETKF estimates the reduction in error variance as a result of observations (Bishop *et al.*, 2001). Since its conception by Bishop *et al.* (2001) the ETKF has been used as a tool for targeting observations in several meteorological and oceanographic applications. The underlying objective is usually to predict the reduction in state error variance at a pre-defined validation location or region that would result from collecting data at a target location some time before the validation time. Bishop *et al.* (2001) and Majumdar *et al.* (2002) refer to the

validation time as a verification time. However, as this thesis is primarily concerned with validating forecasts against observations rather than verify against other forecasts it seems counterintuitive to create a verification time and location.

Majumdar *et al.* (2002) used an ETKF for planning flight tracks, along which GPS dropwindsondes were deployed to monitor winter storm systems off the US western seaboard. The ETKF was able to estimate the signal variance via a statistical rescaling of the National Centres for Environmental Predictions (NCEP) medium-range forecasting model (MRF). Signal variance was deduced by running parallel implementations of the MRF model with and without the targeted data. The statistical rescaling was necessary because the covariance estimates used by the MRF and ETKF were different.

ETKF based adaptive sampling techniques have been employed in an oceanographic context to adjust the behaviour of autonomous ocean vehicles, collecting data on cold water up welling around Monterey Bay, California, USA, for the purpose of increasing the forecast accuracy of models of ocean circulation (<http://www.mbari.org/mb2006/>). An overview of oceanographic data assimilation and uncertainty propagation focusing on examples from Monterey Bay was presented by Lermusiaux *et al.* (2006).

2.4 Summary

This literature review focused on three aspects of the distributed flood forecasting problem:

1. It provides a narrative on recent developments in the field of flood inundation modelling and the parameterisation of these distributed models.
2. It reviews recent innovations in the data assimilation literature with a specific focus on a family of sequential methods based on the Kalman filter.
3. It introduces adaptive sampling, with a focus on examples of the Kalman filter being used as an aid to sampling design in the fields of oceanography and meteorology.

The review indicates that there are few instances of Kalman-type data assimilation algorithms being applied to distributed inundation models for the purpose of estimating terrestrial water storage, flood forecasting and adaptive sampling. Especially for situations where out-of-bank flows are simulated or measured. Thus, the potential of Kalman-type data assimilation methods to increase the accuracy of simulations and the efficiency of data collection campaigns is unknown and for the most part can only be inferred from analogous situations in other research fields. However, before these methods can be evaluated in a state estimation, forecasting and sampling context that is specific to a flood inundation application, a description of potential methods and their practical implementation is required. Therefore, the next chapter provided additional technical detail on a chosen subset of simulation, filtering and sampling methods.

3 Methods

In this chapter, background information is presented on the methods subsequently applied to the problem of flood forecasting and adaptive sampling on the River Crouch. The chapter is divided into five sections that describe the background theory to the methods applied in chapters 4, 5, 6, and 7. The structure of this chapter is intended to be similar to the literature review. Thus, the chapter begins with simulation modelling and parameterisation followed by model updating, it then ends with sampling design methods.

3.1 Hydraulic modelling

Chapter 2.1 provided a short review of the different types of hydraulic models and some example applications. In this thesis, only one model code, ISIS (HR Wallingford), will be used to simulate water flows at the study site. The theoretical background behind this code will now be introduced as two coupled components.

1. An in-channel flow model
2. An out-of-channel or floodplain flow model

3.1.1 Channel model

Open-channel flows are often considered only in the dominant stream-wise flow direction as, for example, in the popular one-dimensional models based around the St Venant equations (Cunge *et al.*, 1980). The ISIS code used in this study estimates open-channel flow by solving numerically a dynamic wave model that comprises a pair of one-dimensional non-linear partial differential equations. The first of these is the conservation of mass or continuity equation (3.1), which controls the balance between discharge from a cross-section and the area of flow at the cross-section:

$$\frac{\partial Q}{\partial L} + \frac{\partial A}{\partial t} = q \quad (3.1)$$

where Q is the within-channel volumetric flow rate or discharge, L is the distance to the next downstream cross-section, A is the cross-sectional area of flow, q is the lateral flow into or out of the channel and t is the time step. The second equation is the conservation of momentum or dynamic equation (3.2), which seeks a balance between gravity, inertia, diffusion and friction:

$$\frac{\partial Q}{\partial t} + \frac{\partial}{\partial L} \left(\frac{\beta Q^2}{A} \right) + gA \frac{\partial h}{\partial L} - gAS_f = 0 \quad (3.2)$$

where β is the momentum correction coefficient, g represents acceleration due to gravity and h is the water surface elevation above datum or stage. The frictional slope S_f is found by:

$$S_f = \left(\frac{Qn}{AR^{2/3}} \right)^2 \quad (3.3)$$

where n is Manning's roughness coefficient and R is the hydraulic radius. The hydraulic radius being equal to the cross sectional area of flow A divided by the wetted perimeter P . A number of assumptions are made when applying these equations (Chow *et al.*, 1988):

1. Within-channel flows are predominantly one-dimensional, in that single values of discharge and stage are sufficient to describe state at a particular cross-section.
2. The bed slope is fixed and small enough to allow the small angular approximation to be used. This approximation assumes that the component of the water mass between cross-sections that generates the force that drives flow indicated by $\sin(\omega)$ in Fig. 3.1 can be approximated by $\tan(\omega)$.
3. The channel is approximately straight and not subject to significant secondary flows.
4. The density of water plus solutes and sediment remains constant.

5. The effects of channel roughness, turbulence, lateral and vertical flows are accounted for by the channel conveyance through an estimate of Manning's coefficient and the momentum correction coefficient.

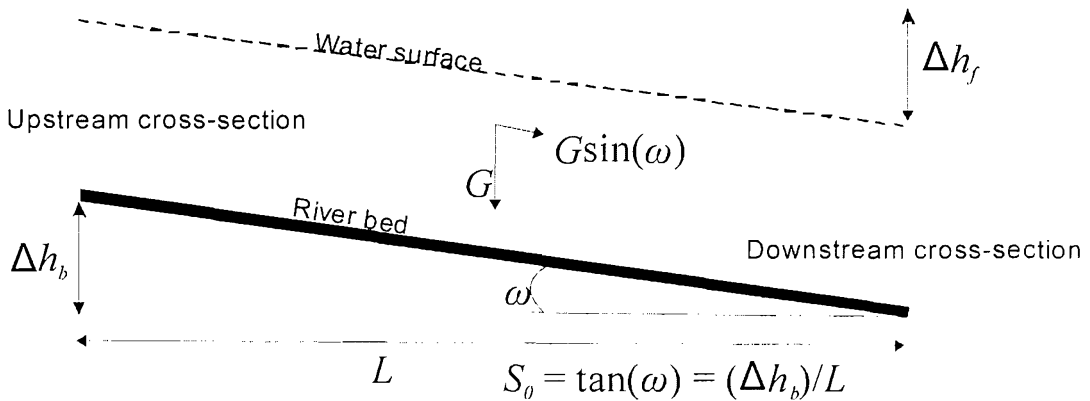


Fig. 3.1: Simplified diagram of the forces acting on the water mass in an open channel, where G is the weight of the water mass. The bed slope $\tan(\omega)$ can replace the direction of flow component $\sin(\omega)$ when the angle is small.

3.1.2 Floodplain flows

ISIS represents out-of-bank flows using a network of storage cells (reservoirs) connected by floodplain sections which control the flow between adjacent reservoirs. Reservoirs are coupled to channel nodes via spill units. It is assumed that inertial effects across reservoirs are negligible and that the water level within the reservoir is flat. Conservation of mass is ensured during simulations for each reservoir by the continuity equation:

$$q_{net} - A(h) \frac{\partial h}{\partial t} = 0 \quad (3.4)$$

Where q_{net} is net flow from all floodplain sections or spill units connected to the reservoir, h is water surface elevation or stage and A is the surface area of the reservoir. Since the surface area of the reservoir may change with stage, surface area is treated as a function of stage generated from topographic data. One method of parameterising this relationship is illustrated by Fig. 3.2 where a LiDAR derived

DEM has been used to estimate the area of a flat water surface (planar area) at different elevations. This stage-area relationship is then plotted on the reservoir geometry graph. Between approximately 1 m and 2.4 m above datum water surface area increases slowly with water level elevation as only the small channels in the reservoir fill. Between 2.4 m and 3 m above datum there is a rapid increase in water surface area with water surface elevation as the salting is inundated. Above 3 m the reservoir is almost fully inundated causing surface area to rise slowly with elevation.

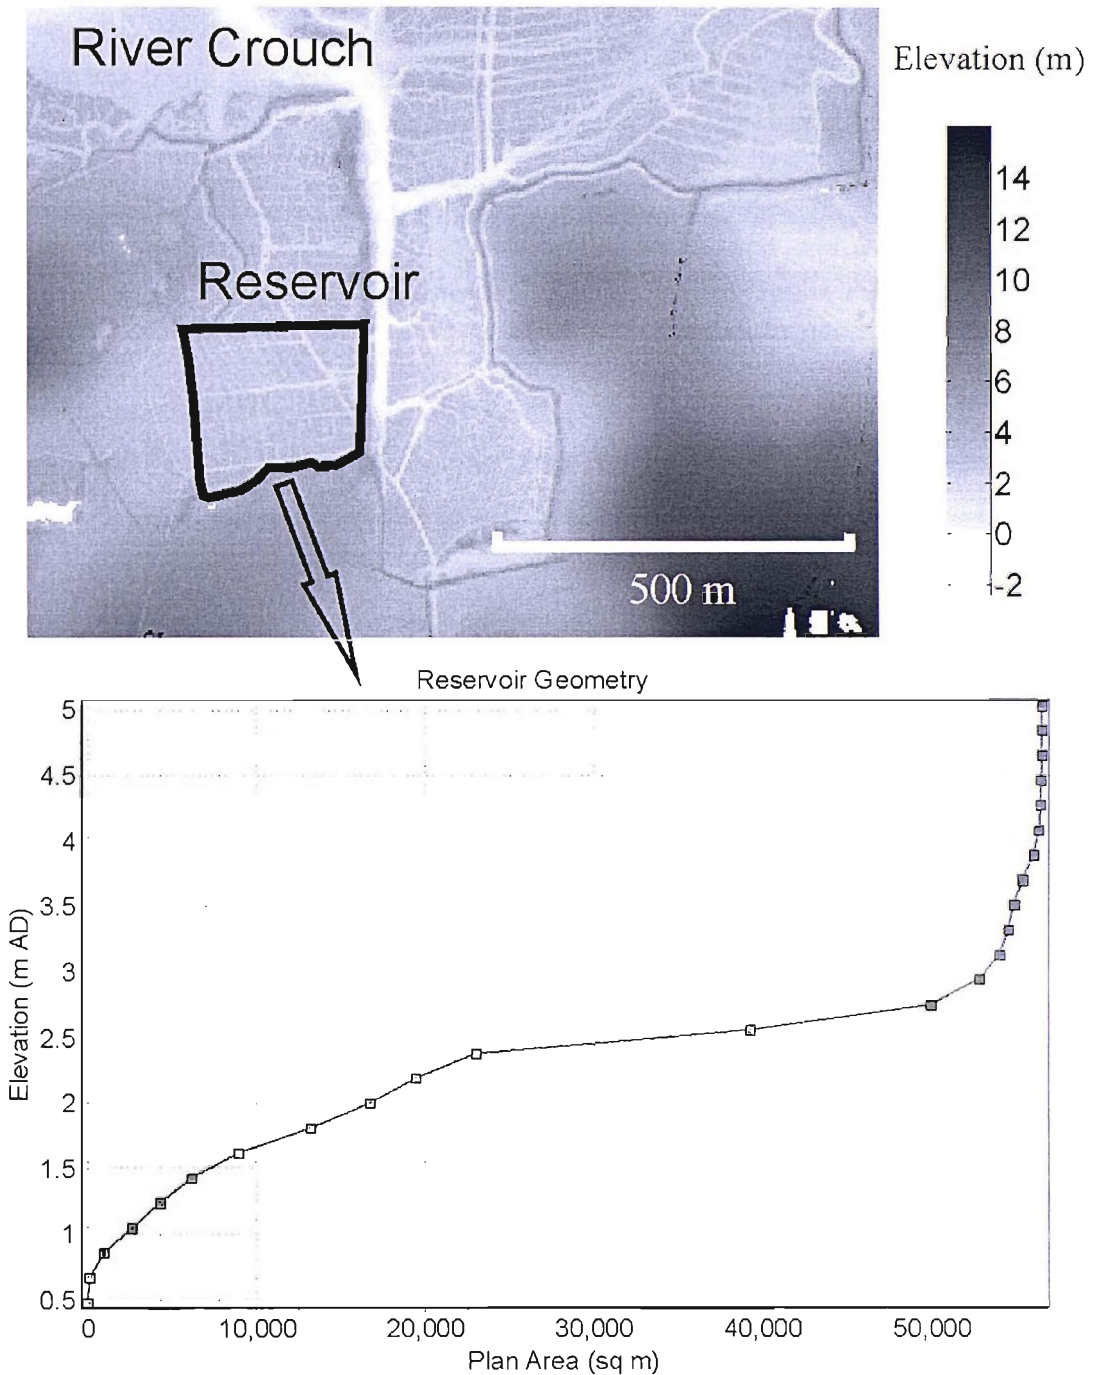


Fig. 3.2: DEM and example reservoir used to parameterise ISIS floodplain reservoir nodes.

As stated previously q_{net} is the combined effect of flow to and from the reservoir either from the channel nodes via spills or between reservoirs via floodplain sections. Flow between reservoir units is controlled by Manning's frictional flow equation:

$$q = \frac{A^{5/3}}{nP^{2/3}} s^{1/2} \quad (3.5)$$

where n is Manning's friction coefficient, s is water surface slope between reservoir centroids, A is the flow area taken from a cross-section of elevation between the two reservoirs and P is the wetted perimeter of the flow. Flow is calculated for each break in slope on the elevation cross-section known as a segment. Fig. 3.3 shows the cross-section from the top edge of the reservoir in Fig. 3.2. Eq. 3.5 can be applied to each of the 49 segments in this cross-section by first calculating the flow area, wetted perimeter and water surface slope as follows:

Area

$$A = (d_2 c_1 + \max(md_2 c_1 d_1 c_2)) b / d, \quad (3.6)$$

Wetter perimeter P

$$P = ((g_{l1} - g_{l2})^2 + b^2)^{1/2}, \quad (3.7)$$

Slope

$$s = (c_1 - c_2) / d, \quad (3.8)$$

where c_1 is the water depth above section in the upstream reservoir, c_2 is the water depth above section of the downstream reservoir, g_{l1} is the ground level at the left hand end of segment, g_{l2} is the ground level at the right end of segment, d_1 is the distance from the centre of the upstream reservoir to the centre of the floodplain section, d_2 is the distance from the centre of the downstream reservoir to the centre of

the floodplain section, m is the modular limit, b is the width of the segment and d is the distance between the upstream and downstream reservoirs. Flows between reservoirs and the channel are controlled by weir-type equations. The equation used varies between free flow, drowned flow and zero flow depending on the difference in stage between the river and reservoir and the modular limit. These equations are described in detail in the ISIS user manual (HR Wallingford, 2007).

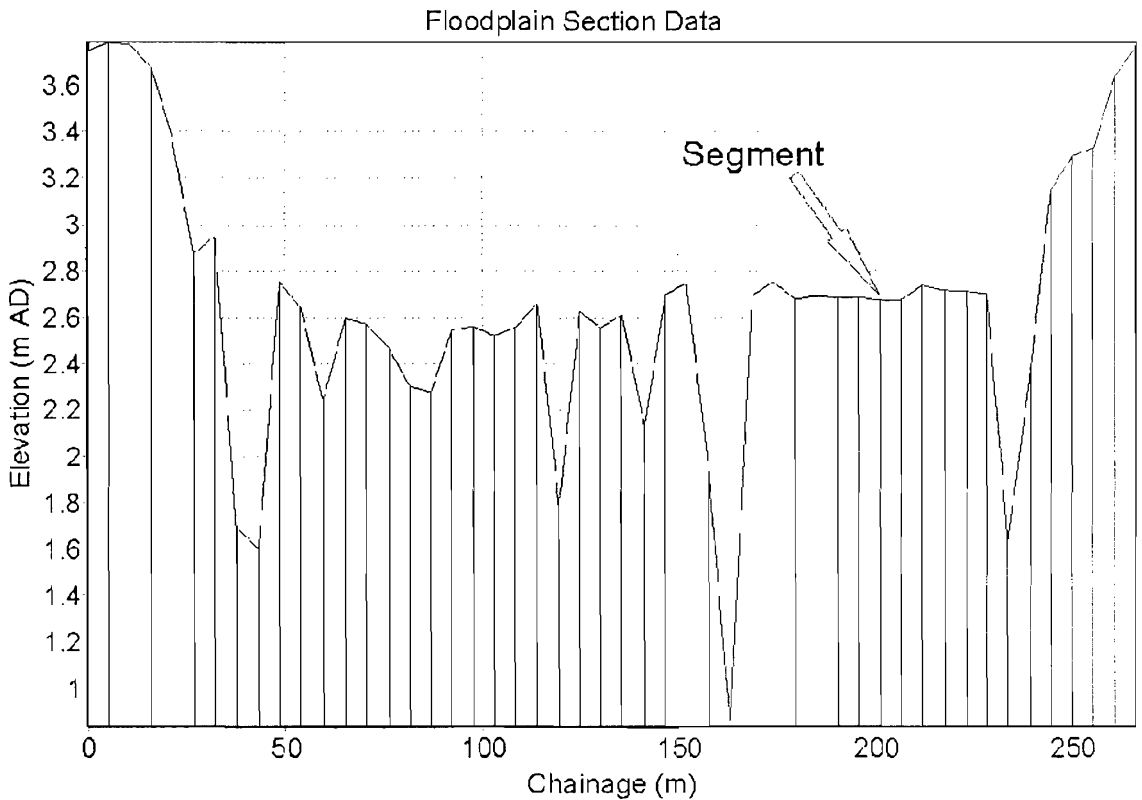


Fig. 3.3: Cross-section between floodplain reservoirs, used to parameterise ISIS floodplain sections. The same principle applies to ISIS spill sections between the channel nodes and floodplain reservoirs.

3.2 Model calibration: GLUE

The Generalised Likelihood Uncertainty Estimator (GLUE) procedure (Beven and Binley, 1992) is a procedure for estimating uncertainty in sets of model parameters. It has been applied to physically-based distributed models for the purpose of estimating: effective roughness parameters of one-dimensional hydrodynamic models (Pappenberger *et al.*, 2005), flood inundation probability (Aronica *et al.*, 1998b, 2002; Romanowicz and Beven, 1998; Romanowicz *et al.*, 1996) and two-dimensional river flow estimation (Hankin *et al.*, 2001). GLUE is not restricted to parameter space as it can also be used to assess uncertainties in boundary conditions (Page *et al.*, 2004; Zack and Beven 1999), initial conditions and model structure (Hankin *et al.*, 2001; Butts *et al.*, 2004). Although this brief introduction will concentrate on parameter uncertainty, the principles described will in most cases be equally applicable to model structure and boundary conditions.

GLUE requires a large number of model simulations for which state estimates can be extracted and compared with observations. Each model run has a parameter set drawn randomly from a pre-defined parameter space, where the parameter space must be large enough to encompass all model parameter sets which lead to ‘behavioural’ simulations. The term behavioural is used to describe all parameter sets which lead to a reasonable simulation of observed state. Sets of parameters are considered, rather than the effect of varying a single parameter, meaning that interactions between parameters are explicitly considered within this process (Beven and Binley, 1992). Here, changes to parameter sets refer to changing the numeric values of parameters as opposed to selecting different parameter combinations, in the sense that a model may attempt to represent reality using a number of different parameter combinations.

For each parameter set likelihood can be calculated based on the fit between the simulation produced by that parameter set and the observations (section 3.2.1). The likelihood values can then be used to estimate the uncertainty in the parameters which make up the parameter set. The GLUE procedure upholds a concept known as

equifinality because no single parameter set is assumed to be optimal within the given model architecture (Beven, 2000). Although one parameter set may achieve the largest likelihood based on some objective function and validation data, other parameter sets that achieve user defined reasonable or behavioural predictions are also considered. Parameter sets that give predictions which do not meet the user's definition of reasonable are termed 'non-behavioural' and removed from further analysis. Whilst this method allows the user to accept any number of parameter sets as behavioural, the weighting given to each parameter set by the pre-defined likelihood function will give less weight to parameter sets that do not recreate observations accurately.

GLUE is an extension of Bayesian inference that is less formal but easier to implement (Beven *et al.*, 2000). A disadvantage of the method, when compared with a classical Bayesian inference approach, is that it assumes all errors (between observed and simulated states) are due to uncertainty in the parameters being estimated, thereby not accounting for other error sources such as measurement errors and input errors. As a consequence estimates of parameter values will be biased because no account will be taken of the variability introduced by the other error sources (Mantovan and Todini, 2006).

There are five requirements to conduct a GLUE procedure outlined by Beven and Binley (1992):

1. A definition of a likelihood measure (referred to more generally as an objective function).
2. An initial range of parameters from which to take parameter sets and a method to select parameter sets, based on a uniform distribution of an unduly wide parameter range, in the absence of information to the contrary.
3. A procedure for using likelihood weights in uncertainty estimation.
4. A procedure for updating likelihood weights recursively as new data become available.
5. A procedure for evaluating the value of additional data.

These steps are discussed in detail by a number of authors (such as: Beven and Binley, 1992; Beven 2000; Hunter *et al.* 2005). Only the choice of likelihood measure will be introduced here because the GLUE method is being used as a parameter estimation tool in support of other thesis aims, rather than an integral component of the thesis aims in its own right.

3.2.1 The Likelihood Measure

The likelihood of a particular set of parameters being a simulator of a system within the constraints of the model structure can be tested by comparing simulated state with observation. Likelihood is a function of how well the model conforms to the observed behaviour of the system. A set of parameters with a likelihood of zero indicates that the set does not represent the behaviour of the system given the available observations. A likelihood of 1 would indicate that the parameter set was a perfect representation of the system, given the available observations. Often the user will define a lower limit of likelihood at some value above zero, below which parameter sets are considered not to represent the behaviour of the system and are reset to a likelihood of zero and termed ‘non-behavioural’. Once ‘non-behavioural’ parameters are removed the likelihood values can be scaled to give a likelihood index between zero and one.

The likelihood measure can be anything that evaluates the goodness-of-fit between the model predictions and observations. The choice of likelihood measure is subject to the judgment of the modeller and will usually depend on the overall objectives of the modelling process. The simplest likelihood measure would evaluate the difference between the predicted and observed state estimates for a given point in time:

$$\ell = \frac{1}{|e|} \tag{3.9}$$

where e is the residual of the observation from the prediction. Another simple measure might be the sum of squared error or the error variance (Beven, 2000):

$$\sigma_e^2 = \frac{1}{i-1} \sum_{t=1}^i (y_t - x_{t,k}) \quad (3.10)$$

Where y_t is an observation at time t , $x_{t,k}$ is the simulated state given parameter set k and i is the number of time steps. A transformation is required to calculate likelihood. Nash and Sutcliffe (1970) suggest the following formulation, which they term a measure of model efficiency:

$$\ell = \left(1 - \frac{\sigma_e^2}{\sigma_o^2} \right); \quad \sigma_e^2 < \sigma_o^2 \quad (3.11)$$

where σ_o^2 is the variance of the observations. This likelihood measure constrains behavioural parameters to those with a residual variance within the variance seen in the observations. The likelihood will be zero when the variance of the residuals equals the variance in the observations (the model is no better than the mean of the observations) and one when all the residuals are zero (Beven and Binley, 1992).

Another likelihood function might be an inverse scaled sum of squares of the residuals (Box and Taio, 1992):

$$\ell = (\sigma_e^2)^{-n} \quad (3.12)$$

where n is a user-defined parameter that scales the sum of the squares of the residuals. If $n = 0$ all simulations will have equal likelihood and when $n = \infty$ the most accurate parameter set will have a likelihood of 1 whilst all others will have a likelihood of 0. The selection of n will affect the shape of the likelihood distribution with larger n values accentuating the weight given to more likely parameter sets. The scaling function n can also be applied to the Nash-Sutcliffe method as discussed in Franks *et al.* (1998) and Freer *et al.* (1996):

$$\ell = \left(1 - \frac{\sigma_e^2}{\sigma_o^2}\right)^{-n}; \sigma_e^2 < \sigma_o^2 \quad (3.13)$$

There are many other possible likelihood functions including Bayesian likelihood functions (Romanowicz, 1996) and pseudo maximum likelihood functions (Van Straten, 1983). There are also functions where the likelihood value is one if the prediction falls within measurement uncertainty, but, less than one either side of this range (Page *et al.*, 2004). Likelihood measures have also been developed to use binary pattern flood inundation data collected synoptically by remote sensors. In these situations simulated flood inundation can be treated as a binary output (e.g. wet or dry) and compared with the observed inundation extent using a contingency table (Aronica *et al.*, 2002; Hunter *et al.*, 2005).

3.3 Data assimilation with the Kalman filter

Kalman filter type algorithms can be traced from the 1960s with the development of the discrete linear Kalman filter (KF) (Kalman, 1960). The KF provides an optimal recursive solution to the discrete-data linear filtering problem, where a process governed by a linear stochastic differential equation controls the system state (Welch and Bishop, 2002). Put simply, the KF algorithm can be used to recursively assimilate observations of a system's state into a model of that system's state. The KF relies on input variables generated by two models: a system model and a measurement model. The system model simulates changes in model state over time and the growth of model uncertainty, whilst the measurement model accounts for measurement uncertainty and maps measurements onto the system model states. These two models will now be introduced along with the assumptions that must be satisfied for the KF to be considered optimal, followed by a description of the KF itself.

3.3.1 The system model

The state of a system of interest at a specific time t can be represented by the state vector \mathbf{x}_t , which is made up of n predicted states within the model domain.

$$\mathbf{x}_t = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}_t \quad (3.14)$$

Uncertainty in these state estimates can be included through the addition of a stochastic component to the state vector, such that the true state of the system is a realisation of the state vector plus an unknown process error $\boldsymbol{\varepsilon}_t$.

$$\mathbf{x}_t = \mathbf{x}_t^{\text{true}} + \boldsymbol{\varepsilon}_t \quad (3.15)$$

Here, it will be assumed that this process error term is a random variable which is Gaussian with zero mean, such that the probability distribution of $\boldsymbol{\varepsilon}_t$ is:

$$p(\boldsymbol{\varepsilon}_t) \sim N(0, \mathbf{Q}_t) \quad (3.16)$$

Where \mathbf{Q}_t is the n -by- n process noise covariance matrix describing the covariance between the systems state errors at time step t . The diagonal terms of the matrix \mathbf{Q} are the variances of each state estimate.

The first step of the KF is to represent the propagation of the state vector from time $t-1$ to time t as a linear stochastic difference equation.

$$\mathbf{x}_t = \mathbf{M}\mathbf{x}_{t-1} + \mathbf{B}\mathbf{u}_{t-1} + \boldsymbol{\varepsilon}_{t-1} \quad (3.17)$$

Where \mathbf{M} is a n -by- n state transition matrix that describes how the state of the system changes linearly from the previous time interval $t-1$ to the next t . \mathbf{x}_{t-1} is the n -by-1 state vector from the previous time step and \mathbf{x}_t is the n -by-1 predicted state vector for

the next point in time t . \mathbf{u}_t is a forcing term or boundary condition, and \mathbf{B} is an n -by- l matrix that relates this boundary condition to the state vector, where l is the number of boundary conditions. The vector $\boldsymbol{\varepsilon}_{t-1}$ is the stochastic component representing system uncertainty or process noise described above.

3.3.2 The measurement model

The measurement model describes the relationship between system state and available measurement data. Assuming there are no errors in the measurements and that the measurements and system states relate linearly to each other, the vector of measurements $\mathbf{y}_t^{\text{true}}$ will relate to the true system state $\mathbf{x}_t^{\text{true}}$ via the m -by- n measurement operator \mathbf{H} :

$$\mathbf{y}_t^{\text{true}} = \mathbf{H}\mathbf{x}_t^{\text{true}} \quad (3.18)$$

Where there are m measurements $\mathbf{y}_t^{\text{true}}$ is an m -by-1 vector. Measurement uncertainty can be accounted for by including a stochastic term in this equation, such that $\mathbf{x}_t^{\text{true}}$ represents the true state of the system that is unknown due to the effect of measurement noise $\boldsymbol{\eta}_t$:

$$\mathbf{y}_t = \mathbf{H}\mathbf{x}_t^{\text{true}} + \boldsymbol{\eta}_t. \quad (3.19)$$

Here $\boldsymbol{\eta}_t$ is assumed to be a Gaussian white noise process with zero mean that is uncorrelated with the process noise $\boldsymbol{\varepsilon}_t$. As with the system model the probability distribution of the random variable $\boldsymbol{\eta}_t$ is

$$p(\boldsymbol{\eta}_t) \sim N(0, \mathbf{R}_t) \quad (3.20)$$

where \mathbf{R}_t is an m -by- m measurement noise covariance matrix describing the covariance between the measurement errors at time step t . For simplicity, it is usually assumed that there is no correlation between measurement errors. Under these

circumstances \mathbf{R} is a matrix of zeros, except for the variance values for each measurement making up its diagonal terms. As $\boldsymbol{\eta}_t$ is assumed to be independent of $\boldsymbol{\varepsilon}_t$, \mathbf{R} is independent of \mathbf{Q} and \mathbf{B} .

The system and measurement models from equations 3.17 and 3.19 respectively, can be used to form the algorithm for the KF. The KF algorithm is based on two sets of equations usually known as the predictor and analysis equations (Welch and Bishop 2001). These equations are sometimes referred to as the time and measurement update equations or prediction and correction schemes, depending on the author. The two-step cycle of the filter uses the predictor equations to estimate the *a priori* state of the system and its covariance for a future point in time using the system model. This state estimate is then adjusted using the corrector equations when measurement data become available, resulting in *a posteriori* state and covariance estimates. Here superscript *a* signifies *a posteriori* state at the previous time step and superscript *f* is forecasted state.

3.3.3 The predictor step

The system model simulates the state vector and its uncertainty at a future point in time based on the system state and state uncertainty at the previous time step. This is achieved by advancing a deterministic estimate of system state (Eq. 3.21) and adding noise to the estimate of the uncertainty in state (Eq. 3.22). The resulting state vector and error covariance estimates are *a priori* as indicated by the superscript *f* symbol because they represent our knowledge of state before measurement assimilation:

$$\begin{aligned}\mathbf{x}_t^f &= \mathbf{M}\mathbf{x}_{t-1}^a + \mathbf{B}\mathbf{u}_{t-1} \\ \mathbf{P}_t^f &= \mathbf{M}\mathbf{P}_{t-1}^a\mathbf{M}^T + \mathbf{Q}_{t-1}\end{aligned}\tag{3.21-22}$$

Where \mathbf{P}_t^f is the error covariance matrix associated with \mathbf{x}_t^f and \mathbf{P}_{t-1}^a is the error covariance matrix associated with \mathbf{x}_{t-1}^a . The *n-by-n* process noise covariance matrix $\mathbf{Q}_t = \text{cov}\{\boldsymbol{\varepsilon}_t\}$ is taken from Eq. 3.16 and *T* indicates the transpose. The simulation of state is deterministic as it is based on the change from the previous state described by

the state transition matrix plus any input driving forces. Since $\boldsymbol{\varepsilon}_t$ is assumed to have a mean of zero it does not occur here.

When initialising the KF it is necessary to make an initial estimate of \mathbf{x}_0 and \mathbf{P}_0 . \mathbf{x}_0 can be anything but \mathbf{P}_0 must be positive and non-zero if observations are to be assimilated at t_0 . If \mathbf{P}_0 is zero then there is no uncertainty in the system model and therefore no reason to assimilate noisy measurement data. Furthermore, if \mathbf{P}_0 is small the filter will initially adjust very slowly to new measurement data, whereas if \mathbf{P}_0 is large the filter may initially adjust very rapidly to new measurements leading to initial instabilities. Nevertheless, Cahill *et al.* (1999) found the KF to be quite robust when given reasonable initial conditions, finding them to have little effect on the final result given enough time steps and measurement data.

3.3.4 The Analysis step

The predictor step simulates the state of the system based on a linear process model. It also passes an estimate of uncertainty in that state estimate. The objective of the analysis step is therefore to gain a more accurate understanding of the current state of the system and to reduce the uncertainty in that estimate through the assimilation of measurement data that are themselves uncertain. For linear systems where there is uncertainty in both the simulation model and measurements, optimal state updating can be done by minimising the errors between model state predictions and measurements using a KF algorithm (Kalman, 1960; Refsgaard, 1997; Schreider *et al.*, 2001; Young, 2002). This is achieved by first calculating the Kalman gain matrix \mathbf{K} , which is then used as a weighting matrix for assimilating measurement innovations $\mathbf{y} - \mathbf{H}\mathbf{x}^f$ into the state vector \mathbf{x} to give an updated estimate of system state that is conditioned in an optimal manner on both modelled and measured state, such that:

$$\begin{aligned} \mathbf{K} &= \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} \\ \mathbf{x}^a &= \mathbf{x}^f + \mathbf{K} (\mathbf{y} - \mathbf{H} \mathbf{x}^f) \end{aligned} \quad (3.23-24)$$

Where \mathbf{H} is a linear observation operator that maps measurements onto states. Here, the reference to time t is dropped as all state vectors and covariance matrices accrue simultaneously in time. The difference between state and measurement vector ($\mathbf{y} - \mathbf{H}\mathbf{x}^f$) is the innovation vector \mathbf{d} (Welch and Bishop 2002). To clarify this process a hypothetical example for a single point based on a bucket filling with rainwater is presented in Appendix A. The sequence of equations 3.23 to 3.24 shows how the Kalman gain \mathbf{K} is used to weight the difference between the *a priori* state vector \mathbf{x}^f and the measurement vector \mathbf{y} when updating to the posterior state vector. The Kalman gain can also be used to update state error covariance such that,

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}^f \quad (3.25)$$

where \mathbf{I} is an $n \times n$ identity matrix. See Appendix A for a single point example.

The KF is the best unbiased linear predictor of state because it estimates the posterior state vector by adjusting the *a priori* state vector depending on the weighted difference between the *a priori* state vector and the measurement model, in a way that minimises *a posteriori* state error variance (Maybeck, 1979). The KF uses all the state information available, in the form of the system model, measurement model and initial conditions. This means that any measurement data, however corrupted by noise (that is assumed to be Gaussian in this case), will be beneficial to the estimation of state because noisy measurement data will be given less weight by the Kalman gain matrix when updating the state vector. Similarly, even inaccurate models can be corrected by measurement data, although the quality of forecast may deteriorate rapidly in such circumstances. In systems that require complex non-linear models to represent physical processes the Kalman filter has been used as a means of correcting much simpler linear models. Dee (1991), for example, uses a simplification of model dynamics for data assimilation into an atmospheric flow model.

Feedback from the KF analysis is accomplished by using the *a posteriori* state vector as the initialisation for the next model time step. This process is often described as recursive or a predictor-corrector loop because predictions are made into the future

and then that prediction is adjusted in real-time, before predicting again at the next time step, as illustrated by Fig. 3.4.

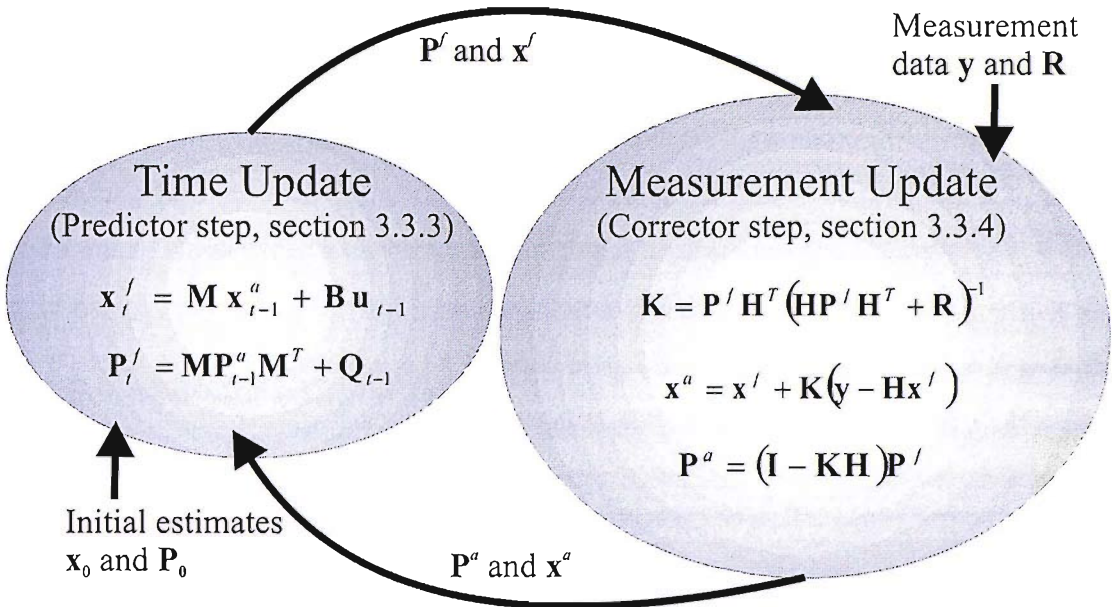


Fig.3.4: Schematic diagram of the KF. The predictor step advances the model state and state error covariance to the next time step using the system model. The corrector step then updates these forecasts with measurements of the state variables. Analysed state and state error variance estimates are then returned to the system model as updated initial conditions.

The KF recursive nature is advantageous as there is no need to store or conduct computations on historical data. In fact, the analysed state variable \mathbf{x}^a and its error covariance matrix \mathbf{P}^a consider all previous states, with previous state estimates becoming less influential on the current state estimate as they move further back in time. This is advantageous for large system models and situations where there are many measurement times.

For completeness, it is worth listing the assumptions that are made by the KF:

1. The state transition matrix, measurement operator and relationship between the model forcing terms and state are all linear.
2. The process noise and measurement noise are independent of one another and can be represented as Gaussian white noise processes with zero mean.

When these assumptions can be satisfied, the filter is the best linear unbiased predictor of state because it minimises the estimated error covariance of the posterior state based on the process and measurement error covariance prior to filtering (Welch and Bishop, 2002). When the above assumptions are not valid the KF will still provide the best linear estimate of state, but not the best estimate of state. It may be the case that the filter can still provide reasonable estimates of state and state error covariance. However, it may be necessary to use a version of the filter which attempts to account for the nonlinear behaviour of the system.

3.3.5 Parameter estimation

In the previous sections, the Kalman filter was expressed as a means of updating internal model state through the assimilation of measurement data. For most models, including those of environmental systems, the values of model parameters are almost guaranteed to be uncertain to some degree. The use of Kalman type filters for the estimation and adjustment of parameters is an appealing extension, especially if the scheme permits the adjustment of parameters in real-time. Here, only one method is presented based on augmentation of the state vector to include a parameter vector. Alternative methods such as one based on maximum likelihood and another based on a dual Kalman filter have been applied in a hydrological context by Bras and Rodrigues-Iturbe (1985) and Moradkhani *et al.* (2004), respectively.

3.3.5.1 Augmented state vectors for parameter estimation

In this section, a method of joint parameter-state estimation or state augmentation will be described, where errors in parameter and boundary condition estimates are assumed to be correlated with both time and errors in state estimation. Recalling the system model from section 2.3.1 (Eq. 3.17) the model operator is re-expressed as being dependent on the unknown parameters θ to give:

$$\mathbf{x}_t = \mathbf{M}(\theta_t)\mathbf{x}_{t-1} + \mathbf{B}\mathbf{u}_t + \boldsymbol{\varepsilon}_{t-1}. \quad (3.26)$$

Therefore, the propagation of the state vector \mathbf{x} is described by the linear model \mathbf{M} which is a function of the parameters $\boldsymbol{\theta}$, and boundary conditions \mathbf{u} . The parameter vector to be estimated can now be augmented with the state vector and updated as part of the KF analysis, such that the state vector becomes (Bras and Rodriguez-Iturbe, 1985):

$$\tilde{\mathbf{x}}_t = \begin{bmatrix} \mathbf{x}_t \\ \boldsymbol{\theta}_t \end{bmatrix} \quad (3.27)$$

Measurements can now be assimilated using the KF analysis algorithm described previously. If the parameters are known to be a constant then their values at the next time step can be identical to the analysed estimate. However, in situations where the parameters might change over time it may be desirable to treat them as random variables. One method of doing so, presented by Bras and Rodriguez-Iturbe (1985), assumes that the parameter values follow a random walk, such that:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\varepsilon}_{\boldsymbol{\theta},t-1} \quad (3.28)$$

where $\boldsymbol{\varepsilon}_{\boldsymbol{\theta}}$ is a zero mean random variable. This allows the parameter to be continuously re-estimated over time as new measurement become available. However, since the parameters change over time the state transition matrix becomes nonlinear. To solve this problem a method for dealing with nonlinear state transitions and error propagation is required. The next section of this review will move onto methods for dealing with nonlinear systems, with the subject of parameter estimation revisited for an ensemble-based method in section 3.5.4.

3.4 Dealing with non-linearity

The KF is the best linear unbiased predictor of system state in terms of error variance, where the state can be described by a linear stochastic difference equation. For many environmental applications, including river flow forecasting, the representation of process with a linear model has been found to be inappropriate

(Evensen, 1994). When the system is nonlinear, forecast state becomes some function f of model state at the previous time step \mathbf{x}_{i-1}^a (analysed in the case of Eq. 3.29), the model forcing terms \mathbf{u}_i and any model errors $\boldsymbol{\varepsilon}_i$:

$$\mathbf{x}_i^f = f(\mathbf{x}_{i-1}^a, \mathbf{u}_i, \boldsymbol{\varepsilon}_i), \quad (3.29)$$

rather than a linear state transition matrix (Drécourt, 2003). The relationship between measurements and the system state may also be nonlinear such that the measurement model becomes:

$$\mathbf{y}_i = H(\mathbf{x}_i^f, \boldsymbol{\eta}_i) \quad (3.30)$$

There are various approximations to the Kalman filter that support the use of nonlinear system and measurement models including the extended Kalman filter (EKF) (Maybeck, 1979) and the unscented Kalman filter (UKF) (Drécourt, 2003). However, for complicated highly nonlinear systems with a large number of states typical of many geoscience applications, including flood forecasting, determining the probability distribution of forecast state can be computationally expensive and mathematically difficult with these methods (Evensen 1992; 1994, Drécourt, 2003). An alternative approach is to use a scheme based on Monte Carlo simulation, where the probability distribution of forecast state is estimated from an ensemble of model simulations.

3.5 Ensemble Methods: The ensemble Kalman filter

The EnKF was developed by Evensen (1994) as an alternative to the EKF, for dealing with non-linear data assimilation problems. The need for the EnKF arises from several implementation problems associated with the EKF. Firstly, the state transition matrix must be re-defined based on the trajectory of the model at the point of data assimilation. This does not affect the state estimation, but will affect error covariance propagation and becomes more of a problem as the non-linearity between time step increases. Evensen (1992) found that unbounded error growth as a result of

approximations in the error covariance propagation equation caused instability and closure problems, even though the model under investigation was only weakly non-linear. For the EKF to be used in this situation Evensen concluded that an extensive coverage of measurement data or higher order closure equations would be necessary to dampen out instability. Such measures are often impractical in an operational sense, especially as the inclusion of higher order error statistics would result in additional computational expense. As noted previously the computational expense of the EKF is a significant drawback of the method. Evensen (1994) therefore proposed a Monte Carlo based approach to approximate the Kalman filter error covariance matrix \mathbf{P} , which has since become known as the EnKF.

3.5.1 State propagation

The EnKF represents the probability density of the system state estimate by Monte Carlo sampling of the state variable. An ensemble of state vectors \mathbf{X} replaces the model state vector used by the KF such that:

$$\mathbf{X} = \begin{bmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,K} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,K} \\ \vdots & \vdots & \ddots & \vdots \\ x_{N,1} & x_{N,2} & \cdots & x_{N,K} \end{bmatrix} \quad (3.31)$$

Where k denotes an ensemble member and there are K ensemble members. The ‘predictor step’ propagates each of the ensemble state estimates forward in time according to the non-linear simulation model M , the state vectors from the previous time step $\mathbf{x}_{k,t-1}^a$, the models boundary conditions (forcing terms) and any model parameters. Stochastic terms should be included for all sources of input and parameter uncertainty such that the ensemble spread is appropriate for the accuracy of the simulation model. In practice, it is not possible to represent or define all sources of model uncertainty, so only the dominant sources of uncertainty are usually considered. These can usually be selected by a combination of sensitivity tests and uncertainty estimation procedures conducted off-line or by adopting a previously

applied approach from the literature. In the example below the system model has one source of uncertainty at the model boundary condition:

$$\mathbf{x}_{k,t}^f = M(\mathbf{x}_{k,t-1}^a, u_t + \varepsilon_{k,t}) \quad (3.32)$$

Further stochastic terms can be added to represent uncertainty in model parameters should this be desired. The mean of the state vector ensemble will be equivalent to a deterministic model run and, therefore, represents the most accurate estimate of state when the ensemble distribution of states is Gaussian:

$$\overline{\mathbf{x}}_t^f = \frac{1}{K} \sum_{k=1}^K \mathbf{x}_{k,t}^f \quad (3.33)$$

where the overline denotes an average over the ensemble. The error covariance matrices can be represented by the difference between the ensemble members and the true state of the system:

$$\mathbf{P}^f = \overline{(\mathbf{x}_{k,t}^f - \mathbf{x}_t^{true})(\mathbf{x}_{k,t}^f - \mathbf{x}_t^{true})^T} \quad (3.34)$$

becoming Eq. 3.35 after the assimilation of measurement data.

$$\mathbf{P}^a = \overline{(\mathbf{x}_{k,t}^a - \mathbf{x}_t^{true})(\mathbf{x}_{k,t}^a - \mathbf{x}_t^{true})^T} \quad (3.35)$$

The true state \mathbf{x}^{true} is of course unknown, preventing the calculation of true error covariance matrices. However, it is possible to estimate the error covariance matrix for each ensemble by calculating the ensemble covariance matrices around the ensemble mean (Evensen, 2004)

$$\mathbf{P}_{k,t}^f \cong \tilde{\mathbf{P}}_{k,t}^f = \overline{(\mathbf{x}_{k,t}^f - \overline{\mathbf{x}}_t^f)(\mathbf{x}_{k,t}^f - \overline{\mathbf{x}}_t^f)^T} \quad (3.36)$$

and

$$\mathbf{P}_{k,t}^a \cong \tilde{\mathbf{P}}_{k,t}^a = \overline{(\mathbf{x}_{k,t}^a - \overline{\mathbf{x}}_t^a)(\mathbf{x}_{k,t}^a - \overline{\mathbf{x}}_t^a)^T}. \quad (3.37)$$

The spread of the ensemble around the ensemble mean is, therefore, used to approximate the error covariance. It should of course be noted that even if state uncertainty is Gaussian this approximation will contain errors proportional to one over the square root of the number of ensembles, meaning that the larger the number of ensemble members the more accurate the approximation of \mathbf{P} . Since there can be an infinite number of ensembles there will also exist an infinite number of ensemble covariance matrices. When the number of ensembles reaches infinity the distribution of states over the ensemble will describe the true probability density function of the system's state variables.

3.5.2 Measurement model

The formulation of the measurement equation for the EnKF was described by Burgers *et al.* (1998), clarifying earlier work by Evensen (1994), by stating the need to treat measurements as random variables to prevent a fall in the estimate of forecast error covariance caused by neglecting measurement uncertainties in the forecast ensemble of state. To treat measurements as random variables it is necessary to create an ensemble of measurements based on the addition of noise to the measurement data before assimilation with the ensemble of system model states. The magnitude of this noise is dependent on the uncertainty in the measurement data, which may result from instrument errors, scaling issues etc. It is assumed here that the uncertainty in measurements is Gaussian and that errors are unbiased and not correlated with each other. Errors in measured and simulated state are also assumed to be uncorrelated. The ensemble of measurements can be generated by:

$$\mathbf{y}_{k,t} = \mathbf{y}_t + \boldsymbol{\eta}_{k,t} \quad (3.38)$$

Where $\boldsymbol{\eta}$ is a Gaussian white noise process representing the uncertainty in measurement, \mathbf{y} is the measurement vector, t is the time and k counts from 1 to the

number of ensembles K . If the relationship between measurements and model state is linear the EnKF measurement model can be defined in a similar manner to the KF measurement model in Eq. 3.19 as:

$$\mathbf{y}_{k,t} = \mathbf{H}\mathbf{x}_t^{true} + \boldsymbol{\eta}_{k,t} \quad (3.39)$$

The covariance of the measured data \mathbf{R} may also be calculated from the measurement ensemble to be used in the analysis step. Although the analysis step presented here will use \mathbf{R} , calculating this matrix is not an efficient means of implementing the filter in practice as shown in section 3.5.5.

3.5.3 Analysis step

In the EnKF analysis step a perturbed measurement vector from the measurement model is assimilated with a member of the forecast state ensemble to yield analysed state, such that:

$$\mathbf{x}_{k,t}^a = \mathbf{x}_{k,t}^f + \mathbf{P}^f \mathbf{H}^T (\mathbf{H}\mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} (\mathbf{y}_{k,t} - \mathbf{H}\mathbf{x}_{k,t}^f) \quad (3.40)$$

This equation is similar to the KF analysis equation in that $\mathbf{P}^f \mathbf{H}^T (\mathbf{H}\mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}$ is the Kalman gain and $\mathbf{y}_{k,t} - \mathbf{H}\mathbf{x}_{k,t}^f$ is an innovation vector. The difference with the KF is that:

1. The analysis is repeated for K ensemble members rather than a single deterministic state vector.
2. Error covariance is approximated from the state ensemble as explained in section 3.5.1.

An efficient matrix-based solution to this analysis developed by Evensen (2003) is presented in section 3.5.5 of this thesis.

By using an ensemble of model simulations to propagate state errors there is no need to calculate a linear state transition matrix based on the trajectory of the model at time t . Hence the problems introduced by the EKF at this stage are avoided. Furthermore, the computational burden of propagating the error covariance is now linked to the size of the ensemble and the speed of the simulation model rather than the size of the covariance matrix, although calculating the Kalman gain can still be expensive when the number of measurements and states is large.

The EnKF assumes the distribution of ensembles to be Gaussian for the purpose of calculating the mean and covariance statistics to be used in the data assimilation or corrector stage. If the state ensemble is non-Gaussian this distribution will be preserved to some extent in the updated state ensemble. However, the error statistics used to calculate the Kalman gain weighting matrix assume normally distributed error statistics for both the measurement and system models. In the context of flood inundation modelling ensemble distributions are unlikely to be Gaussian. Nevertheless, if the errors introduced by assuming a Gaussian distribution are small relative to other error sources this assumption is appealing because it makes filtering easier. Methods for considering non-Gaussian error statistics include kernel based approximation of the probability distribution (Anderson and Anderson, 1999) and the so called particle filter.

3.5.4 Parameter estimation

In section 3.3.5 a state augmentation based method of using the KF to estimate an unknown parameter was presented. The incorporation of a dynamic parameter caused the state transition to become nonlinear, thereby presenting a problem that was not suited to the KF analysis. The EnKF provides a framework for approximating state and state error covariance propagation of a nonlinear simulation model. Therefore, we now revisit the parameter estimation problem with the EnKF instead of the KF for a boundary condition error estimation problem that is analogous to a parameter estimation problem. Recalling the system model in section 3.5.1, correlation would be expected between the state errors and the errors in the boundary conditions (and/or model parameters).

$$\mathbf{x}_{k,t}^f = M(\mathbf{x}_{k,t-1}^a, u_t + \varepsilon_{k,t}) \quad (3.41)$$

In fact for a 1D hydrodynamic river flow model with a stage-time boundary condition the correction made to state at the model boundary must be similar to that made to the boundary condition error at that time. Failure to adjust the boundary condition error in a consistent manner with state at the boundary of the model domain will result in a discontinuity between model state and boundary conditions, which has the potential to cause significant model instability. Therefore, the boundary condition error must be augmented with the state vector (Richel *et al.*, 2002) such that:

$$\tilde{\mathbf{x}}_{k,t}^f = \begin{bmatrix} \mathbf{x}_{k,t}^f \\ \varepsilon_{k,t} \end{bmatrix} \quad (3.42)$$

where $\tilde{\mathbf{x}}_{k,t}^f$ is an augmented state vector containing both forecast model states and boundary condition errors. This augmented state vector can be updated using the analysis equation in section 3.5.3, provided the measurement operator \mathbf{H} is adjusted to take account of the larger state vector by including a matrix of zeros, to indicate that measurements do not map onto the boundary condition error component of the state vector $\tilde{\mathbf{x}}_{k,t}^f$:

$$\tilde{\mathbf{H}} = [\mathbf{H} : \mathbf{0}] \quad (3.43)$$

The filter will require an initial distribution of errors ε_0 such that there is initial uncertainty in the boundary condition error. If the boundary error is believed to be constant the filter can be run as is without any error dynamics. Under these circumstances, the overall variance of the boundary error ensemble will be reduced each time the analysis step is run, along with the model state \mathbf{x} . Therefore, the boundary error distributions produced by the filter will converge towards a single value. If the boundary error is not constant over time the filter will eventually diverge or move away from the truth (Jazwinski, 1970) as the distribution of errors becomes

too narrow. A consequence of divergence is that additional observations have little impact on boundary errors once divergence has occurred (Anderson and Anderson, 1999). To avoid the analysis phase progressively reducing the distribution of errors, they must be treated as dynamic variables (Aksoy *et al.*, 2004). For the boundary errors to be dynamic a technique must be chosen to inflate the posterior error variance to counter the effect of the EnKF analysis.

The dynamics of the boundary errors (or parameters) can be governed by almost any process that is an appropriate representation of the error dynamics. This process will vary depending on the application and may not be fully understood. Here an example will be presented where the boundary condition error is represented using a first order autoregressive process, which is often adequate for short forecasts (Bras and Rodriguez, 1985):

$$\varepsilon_{k,t} = \rho\varepsilon_{k,t-1} + \sigma_\varepsilon(1 - \rho^2)^{1/2}W_{k,t} \quad (3.44)$$

where ρ is a lag-one auto-correlation coefficient the boundary errors at $t-1$ and t , σ is the variance and W is a zero mean, variance one, normally distributed random variable, thus if $\varepsilon_{k,t-1}$ is the boundary error then $\varepsilon_{k,t}$ is the boundary error at the next time step. The time step must be defined prior to calculating ρ and W . Madsen and Canizares (1999) present a similar approach for a distributed boundary condition, where spatial correlation existed between the distributed errors. This needed to be specified through a covariance model such that a random field could be generated (Evensen, 2003).

3.5.4.1 Dual Kalman filters for parameter estimation

In dual estimation, parameter evolution needs to be defined. This is usually done with a random walk model (state augmentation) although Moradkhani *et al.* (2004) suggest the use of kernel smoothing. For dual estimation the EnKF system model is described by:

$$\begin{aligned}\mathbf{x}_{k,t+1}^f &= M(\mathbf{x}_{k,t}^a, \mathbf{u}_{k,t}, \boldsymbol{\theta}_{k,t+1}^f) \\ \hat{\mathbf{y}}_{k,t+1}' &= \mathbf{H}(\mathbf{x}_{k,t+1}^f, \boldsymbol{\theta}_{k,t+1}^f)\end{aligned}\quad (3.45-46)$$

Parameters are updated by assimilating the perturbed observations

$$\boldsymbol{\theta}_{k,t+1}^a = \boldsymbol{\theta}_{k,t+1}^f + \mathbf{K}_{t+1}^\theta (\mathbf{y}_{t+1} + \boldsymbol{\varepsilon}_{k,t+1} - \hat{\mathbf{y}}_{k,t+1}') \quad (3.47)$$

where \mathbf{K}_{t+1}^θ is the Kalman gain for the parameters. The updated parameters are then used to regenerate the system model with updated parameters after which the state vector is updated using the standard EnKF prediction and analysis schemes.

Using the conceptual hydrologic model HyMOD as the system model Moradkhani *et al.* (2005) applied this method for parameter estimation and one-day ahead streamflow forecasting on the Leaf River (USA). The EnKF was found to estimate similar parameter values to a conventional batch simulation approach but had the capability to update in real time and did not require the storage of past model results.

3.5.5 Practical implementation

The computational scheme outlined here for implementing the EnKF is based on that of Evensen (2003). In the following examples the time notation has been dropped to simplify the notation. Firstly, the ensemble of state vectors must be placed in the state ensemble matrix \mathbf{X} , the superscript f has been dropped to simplify the notation:

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K] \quad (3.48)$$

The state ensemble matrix \mathbf{X} has the dimensions N -by- K , where N is the size of the state vector and K is the number of ensemble members. The ensemble mean can be calculated by:

$$\bar{\mathbf{X}} = \mathbf{X} \mathbf{1}_K \quad (3.49)$$

Where $\mathbf{1}_K$ is a K -by- K matrix where each element is equal to $1/K$. This allows for an ensemble perturbation matrix to be calculated:

$$\mathbf{X}' = \mathbf{X} - \bar{\mathbf{X}} \quad (3.50)$$

An estimate of the covariance matrix \mathbf{P}' can then be made for the ensemble perturbation matrix:

$$\tilde{\mathbf{P}} = \frac{\mathbf{X}'(\mathbf{X}')^T}{N-1} \quad (3.51)$$

However, it will be shown later that the covariance matrix $\tilde{\mathbf{P}}$ need not be calculated, thus saving on computational time. Measurements are treated in the same way as the state ensemble:

$$\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_K] \quad (3.52)$$

where \mathbf{y}_k is a vector of measurements perturbed by $\boldsymbol{\eta}_k$ and \mathbf{Y} is an M -by- K matrix, where M is the number of measurements and K is the ensemble size. An ensemble perturbation matrix \mathbf{Y}' can be calculated using the same method used to calculate the state perturbation matrix or defined as:

$$\mathbf{Y}' = [\boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \dots, \boldsymbol{\eta}_K]. \quad (3.53)$$

Measurement error covariance can be calculated using Eq. 3.51 by substituting \mathbf{X}' with \mathbf{Y}' . The KF analysis Eq. 3.23-24 can be adapted to run on the approximate ensemble covariance matrices:

$$\mathbf{X}^a = \mathbf{X} + \tilde{\mathbf{P}}\mathbf{H}^T(\mathbf{H}\tilde{\mathbf{P}}\mathbf{H}^T + \tilde{\mathbf{R}})^{-1}(\mathbf{Y} - \mathbf{H}\mathbf{X}) \quad (3.54)$$

where $\tilde{\mathbf{P}}\mathbf{H}^T(\mathbf{H}\tilde{\mathbf{P}}\mathbf{H}^T + \tilde{\mathbf{R}})^{-1}$ is the Kalman gain matrix \mathbf{K} . \mathbf{H} is the m -by- n linear measurement operator that map's measurements onto the state vector. The most accurate estimate of the state vector can be found by the mean of the ensemble \mathbf{X} . The computation of \mathbf{P} and \mathbf{R} can be avoided by:

$$\mathbf{X}^a = \mathbf{X} + \mathbf{X}'\mathbf{X}'^T\mathbf{H}^T(\mathbf{H}\mathbf{X}'\mathbf{X}'^T\mathbf{H}^T + \mathbf{Y}'\mathbf{Y}'^T)^{-1}(\mathbf{Y} - \mathbf{H}\mathbf{X}) \quad (3.55)$$

A scheme for solving $(\mathbf{H}\mathbf{X}'\mathbf{X}'^T\mathbf{H}^T + \mathbf{Y}'\mathbf{Y}'^T)^{-1}$ using eigenvalue decomposition and an alternative analysis scheme for situations where m is very large was presented by Evensen (2003). Further work has been done to develop analysis schemes that do not require measurements to be perturbed (Whitaker and Hamill, 2002; Tippett *et al.*, 2003; Evensen, 2004). By removing the noise introduced by the perturbed measurement data, Anderson (2001) found the filter was more accurate relative to the standard EnKF analysis scheme presented here for a global atmospheric model. The increase in accuracy was expected to be most noticeable for small ensembles, although it was noted that the improvement may not occur for all cases.

3.6 Adaptive sampling: The ensemble transform Kalman filter

Ensemble based Kalman filters (Evensen, 1994) assume state error covariance \mathbf{P} can be approximated from an ensemble of perturbations around the ensemble mean that have been scaled by the square root of $K-1$, where K is the number of simulations in the ensemble. When the perturbations in question have been generated by an ensemble of model simulations, the error covariance of these background ensembles can be approximated by:

$$\mathbf{P}(t|\mathbf{H}_i) = \mathbf{Z}(t|\mathbf{H}_i)\mathbf{Z}^T(t|\mathbf{H}_i), \quad (3.56)$$

Where superscript T indicates a matrix transpose, t is time and \mathbf{Z} is a matrix of state perturbations with K columns scaled by the square root of $K-1$ (K being the number of ensemble members) such that:

$$\mathbf{z}(t, k | \mathbf{H}_i) = [\mathbf{x}(t, k | \mathbf{H}_i) - \bar{\mathbf{x}}(t | \mathbf{H}_i)] / \sqrt{K - 1}, \quad (3.57)$$

where \mathbf{x} is the state vector of an ensemble member k at time t , the over-line denotes the ensemble mean and the linear observation operator \mathbf{H}_i indicates that the perturbations are conditioned on any measurements that may have been assimilated up to time t_i . State error covariance can be updated at the measurement time t_m using the standard KF update

$$\mathbf{P}(t_m | \mathbf{H}_m) = \mathbf{P}(t_m | \mathbf{H}_i) - \mathbf{P}(t_m | \mathbf{H}_i) \mathbf{H}_m^T (\mathbf{H}_m \mathbf{P}(t_m | \mathbf{H}_i) \mathbf{H}_m^T + \mathbf{R}_m)^{-1} \mathbf{H}_m \mathbf{P}(t_m | \mathbf{H}_i) \quad (3.58)$$

where \mathbf{R} is measurement error covariance. The ETKF analysis is different in that a linear transform \mathbf{T} of the ensemble perturbations is calculated such that analysed error variance from Eq. 3.59 is consistent with that from Eq. 3.58

$$\mathbf{P}(t_m | \mathbf{H}_m) = \mathbf{Z}(t_m | \mathbf{H}_i) \mathbf{T}_m \mathbf{T}_m^T \mathbf{Z}^T(t_m | \mathbf{H}_i), \quad (3.59)$$

A detailed derivation of the ensemble transform \mathbf{T} was presented by Bishop *et al.* (2001). Therefore, the following section will present only the information required for its practical implementation. Since this article focuses on adaptive sampling an additional complication will be added to the notation at this point to deal with the possibility of there being Q potential measurement sets rather than one. Thus, in accordance with Table 5.1 the q^{th} measurement set has the operator \mathbf{H}^q and error covariance \mathbf{R}^q . It will be assumed that the measurement error covariance matrix of the q^{th} potential sample set is known and can be mapped onto state with the linear observation operator \mathbf{H}^q . A further assumption is that errors in these targeted measurements are uncorrelated with errors in any previously assimilated measurements. Under these assumptions Wang *et al.* (2004) defined the transform associated with the q^{th} sample set as:

$$\mathbf{T}^q = \mathbf{C}^q (\mathbf{\Gamma}^q + \mathbf{I})^{-1/2} \mathbf{C}^{qT}, \quad (3.60)$$

Where \mathbf{I} is a K by K identity matrix, \mathbf{C}^q is a K by K matrix of eigenvectors and $\mathbf{\Gamma}$ is a K by K diagonal matrix of eigenvalues of the matrix $\widehat{\mathbf{Z}}^T \widehat{\mathbf{Z}}$ where

$$\widehat{\mathbf{Z}} = \left(\mathbf{R}_m^q\right)^{-\frac{1}{2}} \mathbf{H}_m^q \mathbf{Z}(t_m | \mathbf{H}_i). \quad (3.61)$$

One useful aspect of the ETKF is that transformed ensemble perturbations $\mathbf{Z}(t_m | \mathbf{H}_m^q) = \mathbf{Z}(t_m | \mathbf{H}_i) \mathbf{\Gamma}^q$ represent ensemble perturbations given the measurements used to calculate the transform. This allows routine measurements to be assimilated once before multiple adaptive sampling scenarios are considered (Bishop *et al.*, 2006). Furthermore, what-if scenarios can be built up where the decision to make a measurement is based on the previous decisions to make other measurements. This ‘serial’ approach (Bishop *et al.*, 2001) vastly reduces the size of the adaptive sampling problem. For example, for a situation where there are 10 possible sampling locations, but the means to take only three samples there are 120 possible sampling designs. However, if the location of each sample is decided in series only 27 possible sample locations need to be evaluated. Although, this approach is not optimal in terms of data redundancy, it provides a practical means of conducting a procedure that would otherwise be computationally excessive in situations where there are many possible measurement locations.

At this point, a KF would propagate the analysed covariance matrix through time to the next time of interest, which in this example is the validation time v , using linear dynamic operator \mathbf{M} and adding model error covariance \mathbf{Q} such that:

$$\mathbf{P}(t_v | H_m^q) = \mathbf{M}(t_v, t_m) \mathbf{P}(t_m | H_m^q) \mathbf{M}^T(t_v, t_m) + \mathbf{Q}(t_v, t_m), \quad (3.62)$$

whereas an EnKF based approach would propagate the ensemble to the validation time using the simulation model. The ETKF implementation in Majumdar *et al.* (2002) avoids this computationally expensive operation by substituting the perturbations at the validation time $\mathbf{Z}(t_v | \mathbf{H}_i)$ for the linear dynamic operator giving:

$$\mathbf{P}(t_v | \mathbf{H}_m^q) = \mathbf{Z}(t_v | \mathbf{H}_i) \mathbf{\Gamma}_m \mathbf{T}_m^T \mathbf{Z}^T(t_v | \mathbf{H}_i). \quad (3.63)$$

The validation time can now be any time at which background ensemble perturbations are available. Note that, this is only valid if the dynamics of the ensemble perturbations about the ensemble mean $\bar{\mathbf{x}}(t | \mathbf{H}_i)$ are similar to those that would have been obtained from $\bar{\mathbf{x}}(t | \mathbf{H}_m)$.

In situations where the objective of the adaptive sampling problem is limited to choosing between feasible sample sets, it is only necessary to calculate the reduction in error covariance expected from each of the Q sample sets. This reduction in error covariance is known as ‘signal covariance’ \mathbf{S} , with diagonal terms giving the signal variance. The motivation behind calculating signal covariance instead of error covariance is that it is theoretically easier to calculate because it is independent of model errors \mathbf{P}^f (Majumdar *et al.*, 2002). Using the eigenvalues and eigenvectors calculated previously signal variance is defined by Bishop *et al.* (2001; 2006) as:

$$\mathbf{S} = (\mathbf{Z}^f \mathbf{C}^q) \mathbf{\Gamma}^q (\mathbf{\Gamma}^q + \mathbf{I})^{-1} (\mathbf{Z}^f \mathbf{C}^q)^T, \quad (3.64)$$

Analysed error covariance is, therefore, the forecast covariance minus the signal covariance

$$\mathbf{P}^a = \mathbf{P}^f - \mathbf{S}. \quad (3.65)$$

3.7 Summary

This methodology chapter in conjunction with chapter 2 has reviewed selected data assimilation methods, which in fields such as oceanography and meteorology have been shown to provide solutions to aim similar to those in section 1.1. Specifically, data assimilation and sampling design methods have been presented that can, in theory, be used in conjunction with nonlinear process models and noisy measurements. Over the next four chapters these methods will now be tested through a series of experiments using the study site introduced in the next chapter.

4 Flood inundation model updating using an Ensemble Kalman filter and spatially distributed measurements.

4.1 Introduction

The Environment Agency of England and Wales (EA) aims to deliver accurate, reliable and timely flood forecasts, such that action can be taken to alleviate risks to person and property (Environment Agency, 2003). This objective involves the development of systems for real-time flood forecasting which estimate the probability of a threshold level or flow being exceeded in the near future. Making such a forecast requires some form of real-time enabled model and updating procedure that estimates these states along with their associated uncertainty (Young, 2002). Forecasting state uncertainty requires the identification and estimation of the sources of model and observational uncertainties, and robust methods for dealing with them. An understanding of the spatio-temporal distribution and propagation of these sources of uncertainty and the sensitivity of forecasts to them within the model domain are also required.

Many physically-based distributed flood models were developed primarily with the aim of producing a single deterministic prediction of flood inundation given a set of driving forces and optimal or best guess parameters. The aim of this chapter is to utilise a state updating method and a hydraulic model to estimate stage and stage uncertainty, over a 4 hour forecast period, at specific validation locations. Integral to this aim was the accuracy of state forecasts, and uncertainty estimates, and the robustness of the model when subjected to different spatio-temporal sampling regimes. In turn, the accuracy of state forecasts was dependent on the structure, boundary conditions and parameterisation of the simulation model and the identification of sources of uncertainty within both the system model and observational data. The principal reason for choosing a 4 hour forecast period was a lack of measurement data at low flows for the study site (section 4.4.1), which was

chosen because it provided an opportunity to repeatedly test the assimilation of both floodplain and channel stage data. Although, longer lead times would have been preferable this is more than the 2 hour minimum required by the Environment Agency of England and Wales (EA) (Defra, 2004).

The hyperparameters of the updating procedure were set up for a particular data set. These parameters controlled the source and propagation of errors within the simulation model and, therefore, the weight given to observations by the updating procedure. The accuracy of the forecasts relative to the simulation model alone and the validity of the uncertainty estimates were then assessed under these ideal conditions. The effect on forecast accuracy of reducing the number of data points available to the filter was then assessed, along with the objective of maintaining realistic uncertainty estimates when subjected to different spatio-temporal sampling frameworks.

4.2 Methods

4.2.1 Hydraulic modelling

The ISIS code (HR Wallingford) was used in this study to estimate within channel-flow by solving numerically a dynamic wave model that comprises a pair of one-dimensional non-linear partial differential equations, based around the St Venant equations (Cunge *et al.*, 1980). The first of these is the conservation of mass or continuity equation (4.1), which controls the balance between discharge from a channel cross-section and the area of flow at the cross-section:

$$\frac{\partial Q}{\partial d} + \frac{\partial A}{\partial t} = q \quad (4.1)$$

where Q is the within-channel volumetric flow rate or discharge, d is the location coordinate, A is the cross-sectional area of flow, q is the lateral flow into or out of the channel and t is time. The second equation is the conservation of momentum or

dynamic equation (4.2), which seeks a balance between gravity, inertia, diffusion and friction:

$$\frac{\partial Q}{\partial t} + \frac{\partial}{\partial d} \left(\frac{\beta Q^2}{A} \right) + gA \frac{\partial h}{\partial d} - gAS_f = 0 \quad (4.2)$$

where β is the momentum correction coefficient, g represents acceleration due to gravity and h is the water surface elevation above datum or stage. The frictional slope S_f is found by:

$$S_f = \left(\frac{Qn}{AR^{2/3}} \right)^2 \quad (4.3)$$

where R is the hydraulic radius and n is Manning's roughness coefficient.

Out-of-bank flows were represented using storage cells (reservoirs) connected by floodplain sections where flow was controlled by Manning's equation. These reservoirs were connected to the channel via spill units based on weir type equations. It was assumed that dynamic effects across reservoirs were negligible.

4.2.2 Forecasting: The Ensemble Kalman Filter

The discrete dynamic deterministic hydraulic model or simulation model outlined above can be described by:

$$\mathbf{x}_k^f = M(\mathbf{x}_{k-1}^a, \mathbf{u}_k, \theta) \quad (4.4)$$

where \mathbf{u} represents model forcing terms such as the upstream and downstream boundary conditions, θ represents the model parameters such as bed roughness and M represents the nonlinear model operator which, in this case, represents advancing the hydraulic model state from time step $k-1$ to k . Superscript a denotes *a posteriori* state whilst, superscript f is forecasted state (*a priori*). \mathbf{x} is the model state vector,

which is made up of all the state variables in the model domain that are to be updated by the filter.

Uncertainty in the state estimate is generated by adding errors to the model forcing terms and parameters or even to the state vector itself depending on the source of the uncertainty. Error covariance matrices \mathbf{P} for each ensemble member can be estimated from the difference between the ensemble member and the true state. The true state \mathbf{x}^{true} is unknown, preventing the calculation of true error covariance matrices. However, error covariance \mathbf{P} can be estimated by calculating the ensemble covariance matrices $\tilde{\mathbf{P}}$ around the ensemble mean (Evensen, 1994):

$$\mathbf{P}_k^f \cong \tilde{\mathbf{P}}_k^f = \overline{(\mathbf{x}_{i,k}^f - \overline{\mathbf{x}}_k^f)(\mathbf{x}_{i,k}^f - \overline{\mathbf{x}}_k^f)^T} \quad (4.5)$$

where i counts from one to the number of ensembles. Measurements can be treated in a similar manner. The EnKF measurement equations outlined by Burgers *et al.* (1998) treat measurements as random variables through the generation of a measurement ensemble. The measurement model can be defined as:

$$\mathbf{y}_{i,k} = \mathbf{H}_k \mathbf{x}_k^{\text{true}} + \eta_{i,k} \quad (4.6)$$

where \mathbf{y} is the measurement vector and η is a Gaussian white noise vector representing the uncertainty in the measurements. \mathbf{H} is the measurement operator that maps measurements \mathbf{y} onto the state vector \mathbf{x} . The measurement error covariance matrix \mathbf{R} must also be defined as with the state covariance matrix, such that $\mathbf{R} = \text{cov}\{\eta_{i,k}\}$. Errors in measurement and state estimates are assumed to be uncorrelated. Recently, ensemble Kalman filters that do not require perturbed measurements have become available but are not considered here. The interested reader is referred to Madsen and Cañizares (1999), Whitaker and Hamill (2002) and Evensen (2004).

4.2.2.1 Analysis step

In the analysis step each randomly perturbed measurement vector is assimilated with a member of the prior system model state vector ensemble to generate an updated estimate of model state. The algorithm used in this study can be found in Evensen (2003). At this point it should be noted that the matrices \mathbf{P} (Eq. 4.5) and \mathbf{R} are not explicitly calculated for computational efficiency. However, the EnKF state update is easier to keep track of when described as follows, with reference to time removed.

Firstly, the gain matrix \mathbf{K} (Eq. 4.7) is calculated. This forms a weighting matrix which is used to assimilate members of the measurement ensemble into the state vector ensemble.

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} \quad (4.7)$$

Secondly, each member of the measurement ensemble is assimilated with a member of the state ensemble, resulting in an ensemble of analysed state estimates.

$$\mathbf{x}_i^a = \mathbf{x}_i^f + \mathbf{K}(\mathbf{y}_i - \mathbf{H}\mathbf{x}_i^f) \quad (4.8)$$

The analysed state error covariance \mathbf{P}^a of the state estimate can be calculated using equation 4.5 by substituting analysed state vectors for forecast state vectors. The sequence of equations 4.7 to 4.8 shows how the Kalman gain \mathbf{K} is used to weight the difference between the *a priori* state vector and the measurement vector when updating to the posterior state vector. The KF is optimal in the sense that it estimates the posterior state vector by adjusting the *a priori* state vector depending on the weighted difference between the *a priori* state vector and the measurement model, in a way that minimises *a posteriori* state variance. However, the KF is only optimal in a minimum *a posteriori* variance sense in the linear case. Furthermore, the EnKF is suboptimal in the sense that covariance is estimated from the ensemble. Feedback from the analysis equations is accomplished by using the *a posteriori* state vector as the initialisation for the next model time step. The recursive nature of the KF is advantageous as there is no need to store or conduct computations on historical data.

4.3 Application Case study

4.3.1 Study area and data

The field site used in this study lies 20 km inland along the River Crouch (Essex, UK). The river flows for approximately 30 km west to east from Basildon (British grid reference TQ 730 895) into the North Sea (British grid reference TR 037 961) (see Fig. 4.1) The underlying geology of the region is that of soft sedimentary rock laid down during the Tertiary Period, overlain by deposits from Pleistocene glacial periods. The catchment land use is predominantly agricultural with a number of large towns including Basildon, Wickford, South Woodham Ferrers and Burnham-on-Crouch, along with numerous small towns and villages. The flow duration summary for the river provided by the EA notes that a steady base flow is maintained above the tidal limit by discharge from sewage treatment works at Basildon. The Crouch basin covers an area of 71.8 km² with an average annual rainfall of 572 mm per year (Environment Agency). It is relatively flat and low lying to the extent that the maximum altitude within the catchment is 118 m and tidal conditions extend some 20 km inland, two thirds of the total length of the river. Non-tidal flows on the upper third of the river are monitored by an EA gauge at Wickford (British grid reference TQ 748 934). A summary of the flow duration observed at this gauge is shown in Table 4.1. Saline conditions at the study site range from near fully marine (~35 ppm) to partial influence (~10 ppm). A tidal site was chosen to ensure repeated observations of flood inundation, whilst the long and narrow estuary (see Fig. 4.1) allowed flow between the study site and the reaches tidal boundary to be treated as a one-dimensional problem.

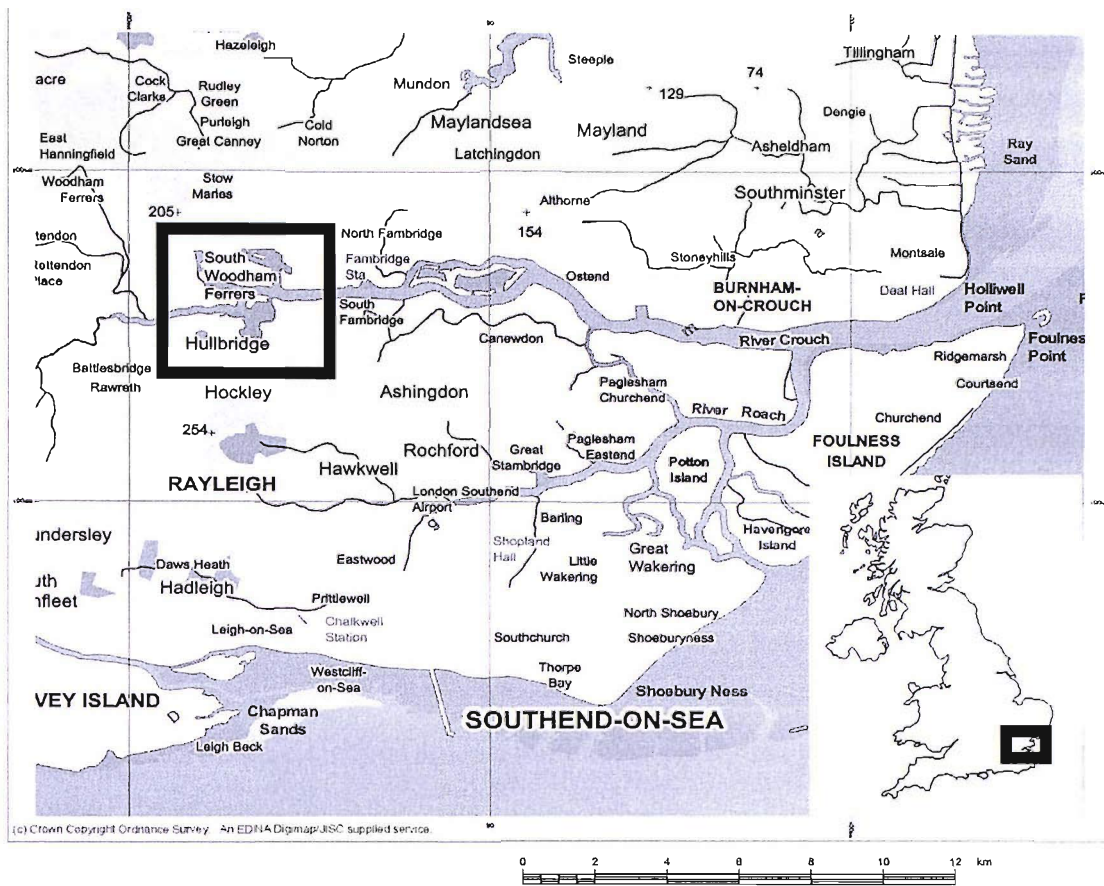


Fig. 4.1: Map of the River Crouch and River Roach, Essex, UK. Scale 1:150,000. To the south is the Thames estuary, with the North Sea to the East. The box indicates the field site where FloodNet sensors have been deployed.

Table 4.1: Summary of Flow duration for EA gauge on the River Crouch at Wickford.

Grid Reference	51 (TQ) 748 934
Mean Flow	$0.35 \text{ m}^3 \text{ s}^{-1}$
95% exceedance (Q95)	$0.046 \text{ m}^3 \text{ s}^{-1}$
10% exceedance (Q10)	$0.776 \text{ m}^3 \text{ s}^{-1}$

The lower 2/3 of the river is characterized by estuarial silt and mud of less than 0.2 mm. Where sea walls or embankments are not present the river is flanked by areas of tidal salting. The tidal range at Burnham-on-Crouch varies between 3.2 meters at neap tide and 5.0 meters at spring tide. Flow velocities are approximately 0.8 m/s during mean tidal conditions. A dredging experiment conducted by Winterwerp *et al.*

(2002) at Burnham-on-Crouch concluded that the channel bed sediment consisted of 40% Clay, 59% silt and 1% sand. The grain sizes D_{50} and D_{90} were estimated to be 4 μm and 40 μm respectively. To the east of Burnham-on-Crouch approximately 5 km from the sea the River Roach flows into the Crouch (Fig. 4.1).

An ISIS 1D hydraulic model was created of a 21 km stretch of the river from Battlesbridge to Burnham-on-Crouch (RC model). Channel cross-sections at the monitoring site (Brandy Hole, Fig. 4.2) were spaced at approximately 200-250 m intervals and were created using bathymetry data provided by Associated British Ports marine environmental research, whilst cross-sections upstream and downstream of the site were provided by Halcrow Ltd, based on bathymetry data. Cross-sections upstream of the study site were separated by 75-300 m, whilst downstream cross-sections were spaced at 350-1500 m intervals. Downstream boundary conditions were set using tidal predictions at 10-minute intervals for the port of Burnham-on-Crouch. These were supplied by The United Kingdom Hydrographic Office and are based on adjusted harmonic constants calculated for the standard port of Walton-on-the-Naze (Essex). Upstream boundary conditions were set at $0.35 \text{ m}^3 \text{ s}^{-1}$ based on the mean flow recorded by the EA gauge at Wickford, approximately 2 km upstream of Battlesbridge. The principal reason for setting this flow constant was that no flow data were available from this gauge. However, running identical model simulations with the Q90 and Q10 flows derived from the gauges flow duration curve resulted in a change in stage at the study site of $<0.01 \text{ m}$. Therefore, the stage and flow at the study site are tidally dominated.

The study site at Brandy Hole consisted of 2.5 km of estuary and approximately 1 km^2 of tidal salting situated behind a derelict sea wall (see Fig. 4.2). At its mouth the river is approximately 700 m wide and up to 14 m deep at high tide. By the eastern extent of the Brandy Hole site the channel width has decreased to around 350 m at high tide with a depth of up to 10 m. At the western extent of the field site channel width is up to 190 m with a depth of up to 5 m (see Fig. 4.3). At low tide the channel becomes un-navigable at the study site (see Fig. 4.4).

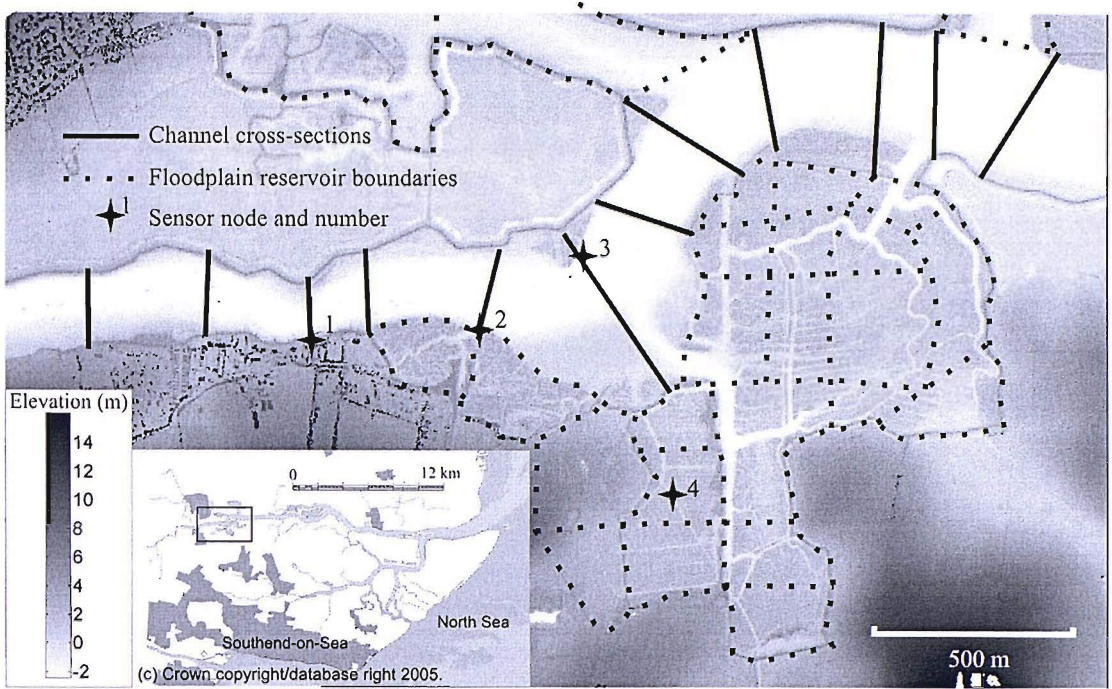


Fig. 4.2: Topographic map of study site derived from LiDAR data overlain by sensor locations and model topology.




Fig. 4.3: The River Crouch at sensor 3 looking west. Sensor 1 is located on the far bank of the river in the centre of the image. Photo was taken approaching high tide on the 26th March 2004.



Fig.4.4: The River Crouch near sensor 1 and Hullbridge, looking west. Sensor 1 is located next to the boat. Photo was taken approaching low tide on the 26th March 2004.

The tidal saltings at Brandy Hole are characterised by numerous small channels typically with steep banks up to a meter high. Some of these channels were artificially cut when the sea wall was first breached, whilst others have formed without human intervention. The presence of artificial channels leads to the striped drainage pattern that can be seen on some areas of the salting particularly around sensor 4 (Fig. 4.2). The eastern and southern limits of the salting are set by an embankment. The river flows across the northern side of the salting whilst its western extent is limited by a gentle slope towards higher ground that makes the salting edge difficult to delineate. At low tide the salting was observed to almost completely drain (Fig. 4.5) whilst becoming fully inundated on a spring tide (Fig. 4.6).




Fig. 4.5: Area of floodplain near to sensor 4 looking north west at low and mid tide.



Fig. 4.6: Area of floodplain looking south towards sensors 4 and 5. Photo was taken at high tide on the 16th March 2004.

The 2.5 km stretch estuary at Brandy Hole was represented using 13 cross-sections, with the salting or floodplain characterised by 22 reservoir units connected to each other by floodplain sections and coupled to the channel. Reservoir units were defined by extracting stage area relationships for polygons overlaying a digital elevation model (DEM). The DEM (Fig. 4.7) was derived from Light detection and ranging (LiDAR) data (Markes and Bates.2000), flown at low tide on 26th December 2003. The user defined reservoir polygons were intended to be as square and similar in size

as possible, with boundaries between polygons drawn at right angles to the perceived flow direction based on user interpretation of the DEM (Fig. 4.2) and an air photo. When placing floodplain sections between two reservoirs it is necessary to choose a cross-section that is representative of the average ground conditions between the two reservoir centroids. Although every effort was made to achieve this, it is in essence a subjective process and there may be no single cross-section that best represents the ground surface at all flows and stages. Spill units were more easily defined and for the most part follow the derelict sea wall that runs parallel to the channel or the highest ground between channel and floodplain. However, where the sea wall has been removed or seriously degraded the boundary between floodplain and channel is less clear. The embankment to the south and east of the floodplain formed the edge of the model domain although the ground on the landward side of the embankment is below high water the embankment was never overtopped during the study period. The model topology can be seen in Fig. 4.8.



Fig. 4.7: LiDAR DEM of Brandy Hole, with air photo overlay. Looking west over Brandy hole towards Hullbridge.

The study site was monitored using four pressure transducers measuring water level in the channel and on the floodplain (Fig 4.2). The measurements taken by these sensors are subject to uncertainties that will now be summarised. The combined effect of non-linearity, hysteresis and repeatability on the accuracy of measurement is $\pm 0.1\%$ for water depths above 1 m of H_2O . A further $\pm 0.1\%$ per annum may be added to account for long-term stability, whilst temperature effects result in a further $\pm 0.3\%$ in depths above 3.5 m and $\pm 0.6\%$ below 3.5 m (Druck, 2004).

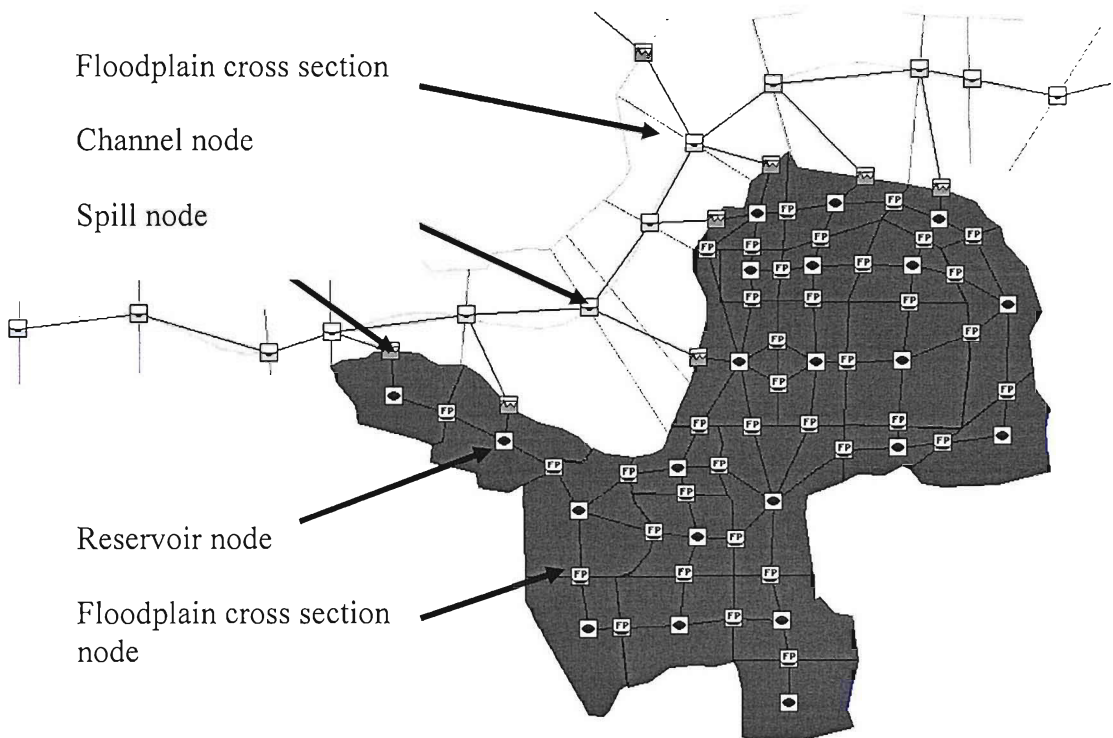


Fig. 4.8: Diagram showing model topology at Brandy Hole. The areas that contribute to each reservoir are also shown as polygons in Fig 4.2. The area shown is 2.5 km across.

The Druck sensors self calibrate using atmospheric pressure as a baseline. Changes in atmospheric pressure are small relative to water pressure and are not listed as a cause of inaccuracy in the published specifications table (Druck, 2004). Differences in the density of fresh and saline waters will affect sensor readings. The manufacture specifies a density of 1.00 g/cm^3 for fresh water and 1.025 g/cm^3 for saline water (Druck 2004b). Estuaries present a problem in this case as their brackish waters vary in density over time. The sensors at Brandy Hole have been calibrated for fresh

water. EA data from Hullbridge (TQ 80982 95565) collected between 1997 and 2000 indicate that salinity varied between 2.3 g/l and 33.4 g/l, where a fully marine environment would have a salinity of around 35 g/l. As a consequence dense saline water could result in overestimates of stage by up to 2.5%, based on these figures. Sensor locations were measured using differential GPS with an estimated accuracy in the vertical plane in this case of ± 0.02 m (Bowdidge, 2004).

Each sensor measurement comprised the mean and variance of 32 samples taken over a 30 second period. Each sensor used IEEE 802.11 wireless Ethernet to transmit its data to a database. Although sensors could transmit data via each other to the database, the limited range of the wireless technology prevented any one sensor being placed more than ~ 700 m from another. This and a desire to have at least one sensor on the floodplain were the principal constraints on the location of each sensor in the network. However, there were a number of technical and computer science issues that put additional constraints to the locations of the 4 sensors:

1. Sensor 1 required mains power. Hence its location on a pier (see Fig. 4.1).
2. Sensor 1 was the only node capable of sending measurements off-site, meaning all other sensors required some means of transmitting their measurements to this node.
3. At least 1 sensor should be unable to directly communicate with sensor 1 without routing its measurement data via another sensor. Sensor 4 met this aim.
4. There must be more than one route that a sensor unable to directly communicate with sensor 1 could transmit its data to sensor 1. Sensors 2 and 3 served this purpose by both providing a possible means of sensor 4 routing its measurement data to sensor 1.

Data acquired over 18 consecutive tidal events from 21st May to 30th May 2004 were used to calibrate the River Crouch (RC) model (calibration data), with data from 11 tidal events from the 1st June to 6th June 2004 used as forecasting test data (forecasting data). A summary of the forecasting data set is provided in Table 4.2. The number of data produced by each sensor is related to its elevation and spatial

location, which resulted in sensor 1 delivering more samples than the other sensor. Stage measurements from sensors 1, 2 and 4 were typically similar, with differences between sensors measurement in the range of 0.01 to 0.03 m. There was little structure to these differences suggesting they were most likely the result of local effects such as wind, waves and measurement errors, rather than a combination of tidal forcing and estuary morphology. At sensor 3 stage was observed to be 0.05 to 0.1 m below that at sensor 1 during the incoming tide and at high tide. On the outgoing tide the difference was typically less than 0.01 m and could be both positive and negative, giving the appearance of a time lag between the two locations on the incoming tide and a flat water surface on the outgoing tide. This effect was observed repeatedly from tidal event to event.

Table 4.2 Summary of forecasting data set. All samples collected at 15 minute intervals.

	Measurements available	Sensor elevation	Location of sensor
Sensor 1	198	0.66 m	Channel
Sensor 2	112	1.59 m	Channel
Sensor 3	137	1.07 m	Floodplain
Sensor 4	89	1.77 m	Channel

4.3.2 Model calibration

The generalised likelihood uncertainty estimator (GLUE) procedure (Beven and Binley, 1992) has been applied to physically-based distributed models for the purpose of estimating effective roughness parameters (Pappenberger *et al.*, 2005a) and flood inundation probability (Romanowicz and Beven, 1998). A detailed description of the procedure will not be presented here as the method used has been published elsewhere (Beven and Binley, 1992).

A comparison of model predictions of stage with observations of stage at sensor 1 was used to evaluate the likelihood of a particular set of parameters being an acceptable representation of the system within the constraints of the model structure. A measure of model efficiency was used:

$$\ell_i = \left(1 - \frac{\sigma_e^2}{\sigma_o^2} \right) \quad (4.9)$$

Where σ_o^2 is the variance of the residuals between observations and the downstream tidal boundary and σ_e^2 is the variance in the residuals between observations y and model state predictions x given variable set Φ_i . As a consequence, the likelihood measure was a measure of the hydraulic model performance above that of the tidal harmonic. Known limitations associated with this likelihood measure include an inability to consider autocorrelation between measurement errors or variations in measurement variance. The main reason for using equation 4.9 is its widespread use in other hydrological applications. However, a heteroscedastic maximum likelihood estimator measure could have been used to consider autocorrelation in measurement variance. The interested reader is referred to Hunter *et al.* (2005).

The GLUE procedure was used to estimate uncertainty in three variables from the RC model. The three variables of interest were a global channel roughness parameter, a global floodplain roughness parameter and the magnitude of the downstream tidal boundary. The initial uniform range or prior distribution of each variable to be calibrated is shown in Table 4.3. Uniform priors were used to reflect a lack of knowledge regarding the distribution of variable uncertainty prior to the calibration procedure being run.

Table 4.3: Initial sampling range for input variables calibrated using GLUE.

Variable	Range sampled (uniform distributions)
Channel roughness	0.01 – 0.04 (Manning's n)
Floodplain roughness	0.02 – 0.1 (Manning's n)
Boundary error	-0.5 – 0.5 (m)

All other model parameters and boundary conditions were held constant, based on either published or default values. An identical initial stage and flow condition for each model simulation based on a steady state simulation at low flow was created 24

hours before the start of the calibration period. The intervening period, comprising of approximately 2 tidal cycles was necessary to remove the effect of the constant initial conditions on simulation likelihood. The GLUE calibration procedure consisted of 1000 model runs over the calibration data period. State predictions were extracted from each model run for comparison with observations. The parameter likelihood distributions were found to vary between incoming and outgoing tides. As such calibrating model parameters to data on the incoming tide was detrimental to prediction accuracy on the outgoing tide when compared with fitting to the complete data set. Adjusting the timings of the simulations and re-calculating the likelihood distributions did not align the distributions. Therefore, as the objective of this particular model was to produce short-range forecasts the model calibration considered observations from the incoming tide only rather than the complete time series. The calibration results can be seen in Fig. 4.9a–9c.

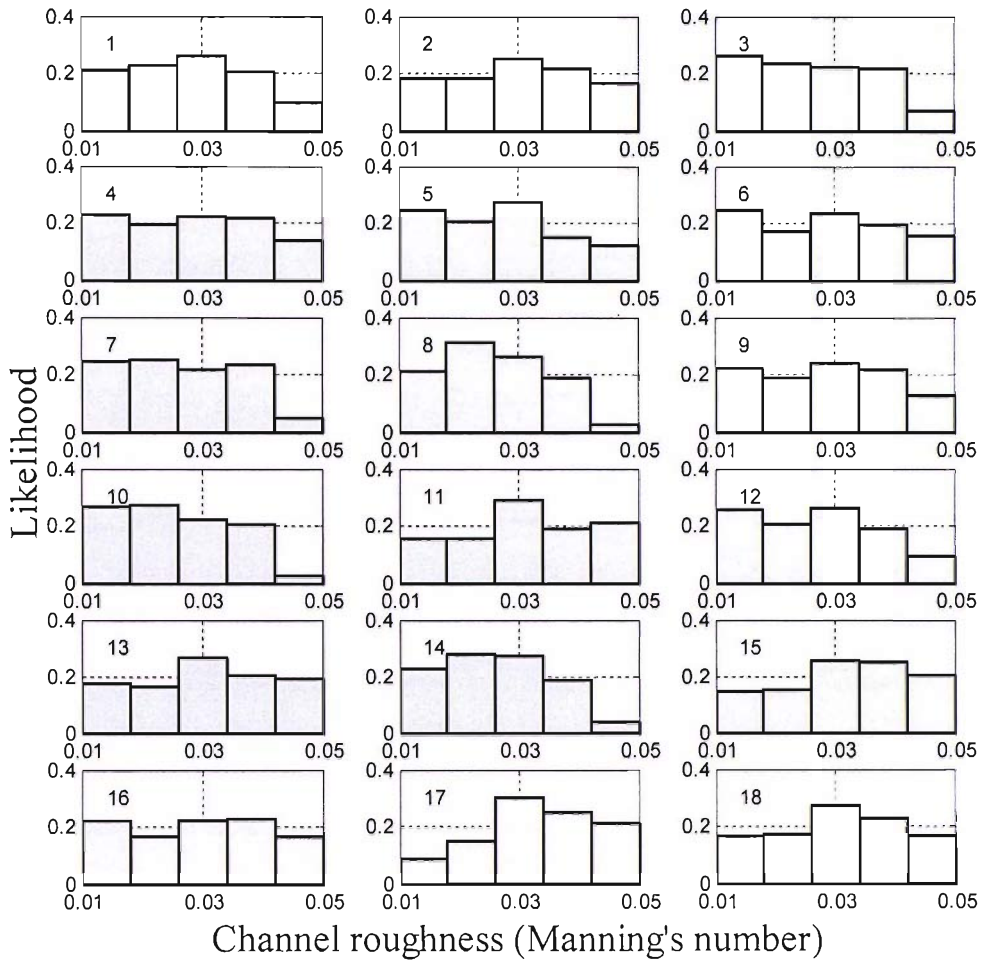


Fig. 4.9a: Likelihood histograms for 18 consecutive tidal events between the 21st and 30th May 2004. Plots show the effect on likelihood of varying channel roughness, using measurements from sensor 1.

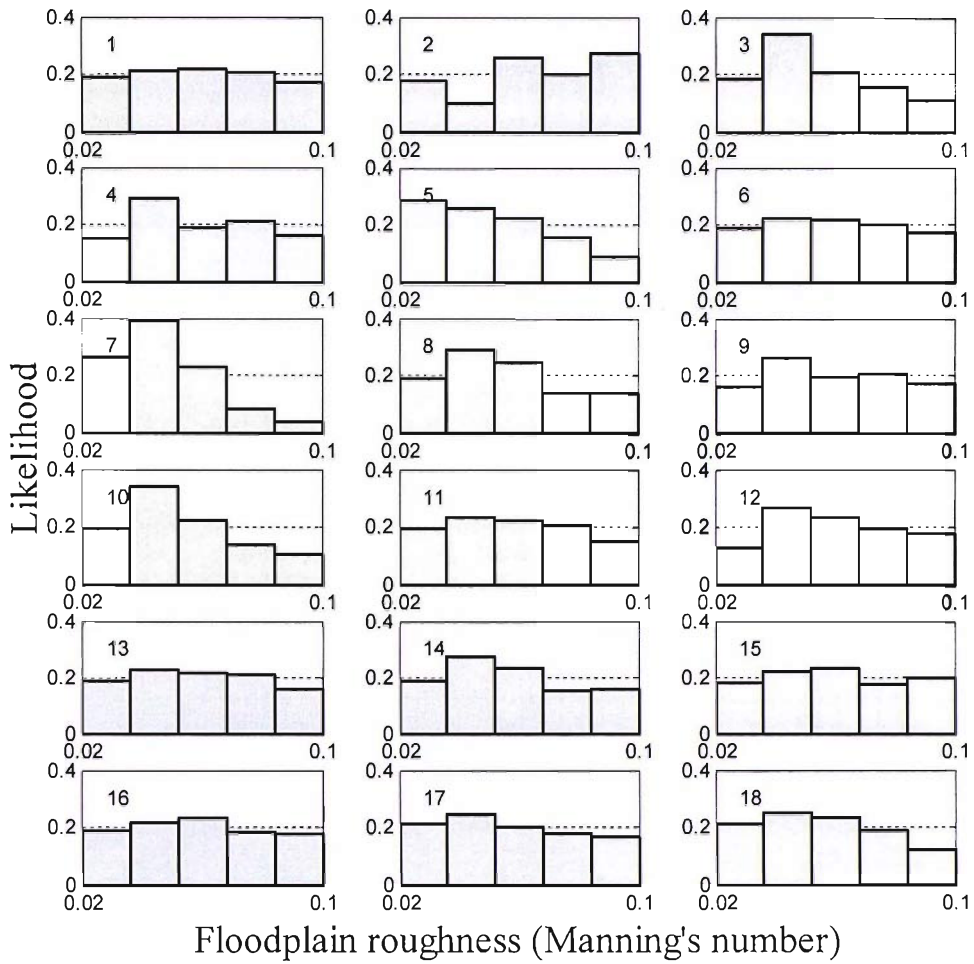


Fig. 4.9b: Likelihood histograms for 18 consecutive tidal events between the 21st and 30th May 2004. Plots show the effect on likelihood of varying floodplain roughness, using measurements from sensor 3.

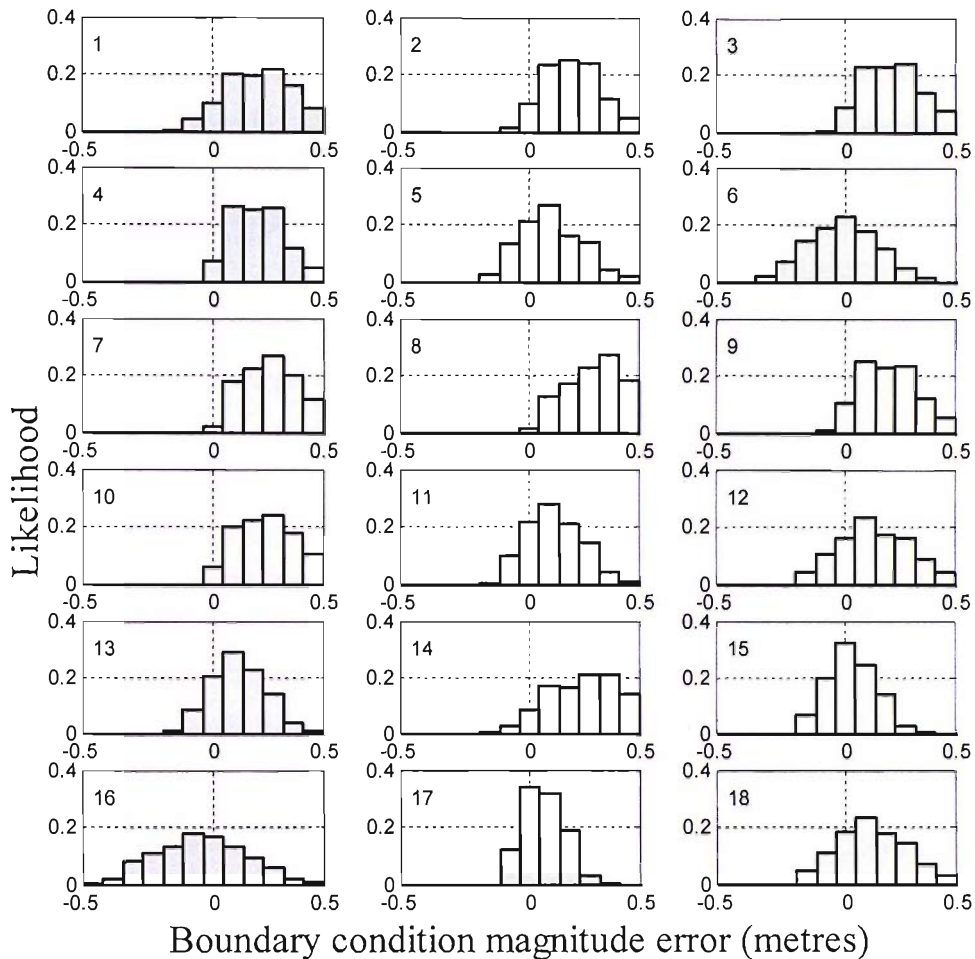


Fig. 4.9c: Likelihood histograms for 18 consecutive tidal events between the 21st and 30th May 2004. Plots show the effect on likelihood of varying boundary condition magnitude, using measurements from sensor 1.

Fig. 4.9a shows that the best estimate of the global channel roughness coefficient remained reasonably stable over the 18 tidal events at 0.024, although any physically meaningful value of roughness within the prior range of 0.01 and 0.04 could produce reasonable results. Larger values of roughness tended to require an increase in boundary condition magnitude relative to smaller values of roughness to maintain behavioural likelihood values. Fig. 4.9b indicates that the value of floodplain roughness, which influences flow between reservoir units on the floodplain, had a limited effect on simulation likelihood (at sensor 3) relative to the other variables during most events. Although simulations with larger roughness values tended to have lower likelihoods during some events any of the values sampled could produce

a behavioural simulation. Fig. 4.9c shows that the best estimate of boundary condition magnitude error η varied from event to event around a mean of +0.12 m. It is not clear from this limited time series of results whether the RC model would have a tendency to under-predict stage at sensor 1 given unaltered boundary conditions because the prevailing atmospheric conditions over the 18 tidal cycles or measurement errors could have produced this effect. However, the results do show that uncertainty in boundary condition magnitude over any one event was generally less than the variation in boundary condition magnitude between events. Consequently, errors in simulations were most likely to have propagated from the model boundary because simulation predictions are less sensitive to channel roughness.

4.3.3 Application of EnKF to a 1D hydraulic model of the River Crouch

An EnKF was applied to the RC model for the purpose of updating model state and boundary conditions. The state vector was made up of 161 stage and flow values from each channel, reservoir and floodplain section in the model. Unique stage and flow initial conditions were created for each ensemble member by running simulations with perturbed boundary conditions from 21st May to 30th May 2004. Initial conditions in the channel did not persist from one tide to the next, although water depths of a few centimetres were predicted to reside at some floodplain locations between tides.

Sensors were located either along cross-sections in the channel or close to reservoir centroids on the floodplain (Fig. 4.2). This allowed observations to be mapped directly onto the corresponding model predictions in the state vector. The water surface was assumed to be flat across channel and floodplain sections. Measurements of flow were not taken. However, in channel flow needed to be updated, along with stage, to prevent state instabilities when the RC model was initialised with updated state and boundary conditions. Flow can be estimated using stage-discharge relationship for each cross-section built up for several steady flow simulations.

However, such a method is not appropriate here where the tide creates strong backwatering effects that lead to a varying relationship between stage and discharge over time. Therefore, flow values were augmented with stage values into the state vector and updated as part of the EnKF analysis phase (Hartnack and Madsen, 2001) such that:

$$\mathbf{x}_{i,k} = \begin{bmatrix} \mathbf{h}_{i,k} \\ \mathbf{Q}_{i,k} \end{bmatrix} \quad (4.10)$$

where \mathbf{h} is a vector containing the stage at each cross-section and \mathbf{Q} is the associated flow (note \mathbf{Q} is a vector not matrix despite capitalisation). Consequently, the EnKF was permitted to alter the volume and spatial distribution of water within the RC model domain based on observations of stage within the channel and on the floodplain.

It was assumed that the downstream tidal boundary u_k was the only source of uncertainty, although this setup could be extended to include other inputs or any model parameter such as channel roughness. Under this assumption the model was formulated as a stochastic nonlinear dynamic system with errors $\varepsilon_{i,k-1}$ in forcing terms:

$$\mathbf{x}_{i,k}^f = M(\mathbf{x}_{i,k-1}^a, u_{k-1} + \varepsilon_{i,k-1}) \quad (4.11)$$

Errors in forcing terms (the downstream tidal boundary) were assumed to be correlated in time and adequately approximated by a first-order autoregressive model, due to the short periods of time between state updates:

$$\varepsilon_k = \alpha \varepsilon_{k-1} + \delta_k \quad (4.12)$$

where α is a lag-one autocorrelation coefficient and δ is a zero-mean Gaussian white noise component with variance σ . Errors in the forcing term were updated as part of the data assimilation scheme by augmenting them with the state vector (Madsen and

Cañizares, 1999). Consequently, the state vector now comprises of estimates of stage, flow and forcing term errors. However, the variance and autocorrelation coefficient of the forcing term error model were not known *a priori* and needed to be estimated.

4.3.4 Estimation of autoregressive model coefficients

The coefficient of the autoregressive model had an impact on the spread of the state ensemble at future points in time and therefore the weighting given to measurements at future time steps. If the autocorrelation is too small the ensemble will spread rapidly as the model is advanced to future points in time. This is likely to result in an over-estimate of model uncertainty, degrading the value of the forecast. It will also result in more weight being given to measurements at future time steps, which may lead to rapid adjustments in model forecasts based on uncertain measurements and in some cases increase the likelihood of model instabilities. An autocorrelation coefficient that is too large may result in an ensemble spread which is too narrow at future points in time. This has the effect of giving less weight to measurements at future time steps and could result in divergence, where the prior distribution is so narrow that further observations become essentially irrelevant (Anderson and Anderson, 1999).

The method used to optimise the coefficients of the autoregressive model is presented below. The objective was to find a decay coefficient that led to an ensemble forecast distribution that was statistically indistinguishable from the truth as approximated by future observations, over many filter time steps. This method was applied by Anderson and Anderson (1999) and Anderson (2001) for tuning a filter operating on the Lorenz-63 weather forecasting model and is presented below.

Firstly, the time averaged root mean squared error between the ensemble forecast mean and measurement data $RMSE_{em}$ was calculated by:

$$RMSE_{em} = \sqrt{\frac{1}{n} \sum_{j=1}^n (y_j - \bar{x}_j)^2} \quad (4.13)$$

where y_j is a measurement, n represents the number of measurements at times j and \bar{x}_j is the ensemble forecast mean.

Secondly, the time averaged root mean squared error was calculated for each ensemble member $RMSE_e$. This was achieved by substituting the predictions from each ensemble member x_{ij} for the ensemble mean in the above equation. Thirdly, the ratio (Ra) between $RMSE_{en}$ and $RMSE_e$ was calculated. Given an ensemble of size N the expected value of this ratio (where the truth represented by the time series of measurements is statistically indistinguishable from the ensemble mean) should be (Anderson, 2001):

$$E(Ra) = \sqrt{(N+1)/2N} \quad (4.14)$$

Finally, the ratio of Ra and $E(Ra)$ can be calculated.

$$r = \frac{Ra}{E(Ra)} \quad (4.15)$$

If r is greater than 1 the spread of the ensemble forecast is on average too small to represent subsequent measurement. If r is less than 1 the spread of the forecast ensemble is too large hence, the closer the value of r to 1 the better the estimate of forecast uncertainty according to this measure. Since all uncertainty in the RC model was assumed to propagate from the downstream boundary condition the amount of noise introduced at this location will be related to the value of r . Good values for the autocorrelation coefficient a and variance of the noise component will result in a value of r closest to 1.

An r value of 0.995 at sensor 1 was achieved based on data from all sensors over the forecasting test period, using a correlation coefficient of 0.995 and variance of 0.035 as parameters of the downstream boundary condition AR(1) model, with 10 minute

time steps. These coefficients represent the best value of r achieved by changing manually the autoregressive model coefficients and running the filter.

4.4 Results and discussion

4.4.1 Increasing forecast accuracy using the EnKF

A simulation by the RC model with boundary conditions set using deterministic tide predictions was produced for the forecasting test period between the 1st and 6th June 2004. A scatter plot of predicted against observed stage at the four monitored locations can be seen in Fig. 4.10. Fig. 4.10 shows that the RC model had a tendency to over-predict stage during the simulation period by up to 0.4 m. This is most likely due to the variations in boundary condition magnitude identified by the calibration procedure in Fig. 4.9c. Fig. 4.11 plots the residuals from Fig. 4.10 (predicted minus observed stage) grouped by event against predicted stage at sensor 1. The plot shows some evidence of hysteresis indicating a lag between predicted and measured stage, with residuals tending towards greater positive errors at predicted stages of approximately 2 m. The residuals highlight some possible limitations associated with the parameters, boundary conditions and structure of the RC model. However, it is possible that this effect may be due to the sensor calibration being less effective in shallow water. Nevertheless, the pattern of the residuals indicates the dominance of inter event prediction errors over intra event prediction errors. There is a temptation to conclude that the RC model systematically over-predicts stage. However, this was not the case for the calibration data set (where it generally under predicted) and may not have occurred if a longer time series of data had been available. Furthermore, action to correct this globally over the test data would be to the detriment of the events that are predicted well by the RC model. The RMSEs of the simulation model predictions are shown in Table 4.4. The degree to which these can be accounted for by the proposed updating scheme will now be investigated.

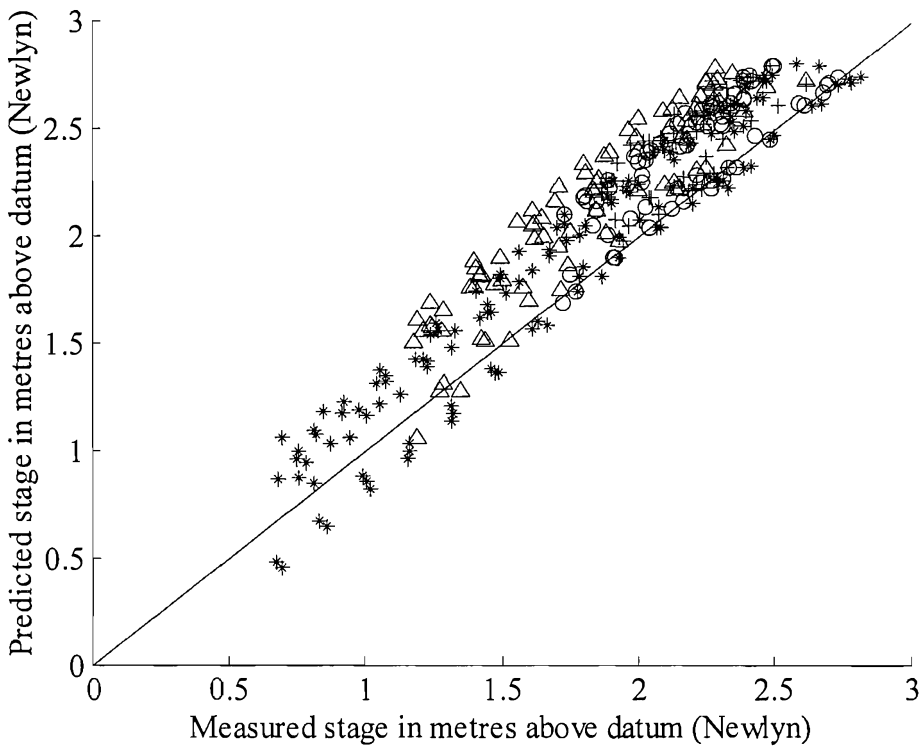


Fig. 4.10: Plot of RC model stage predictions against stage measurements at four sensor locations: sensor 1 *, sensor 2 O, sensor 3 Δ, sensor 4 +.

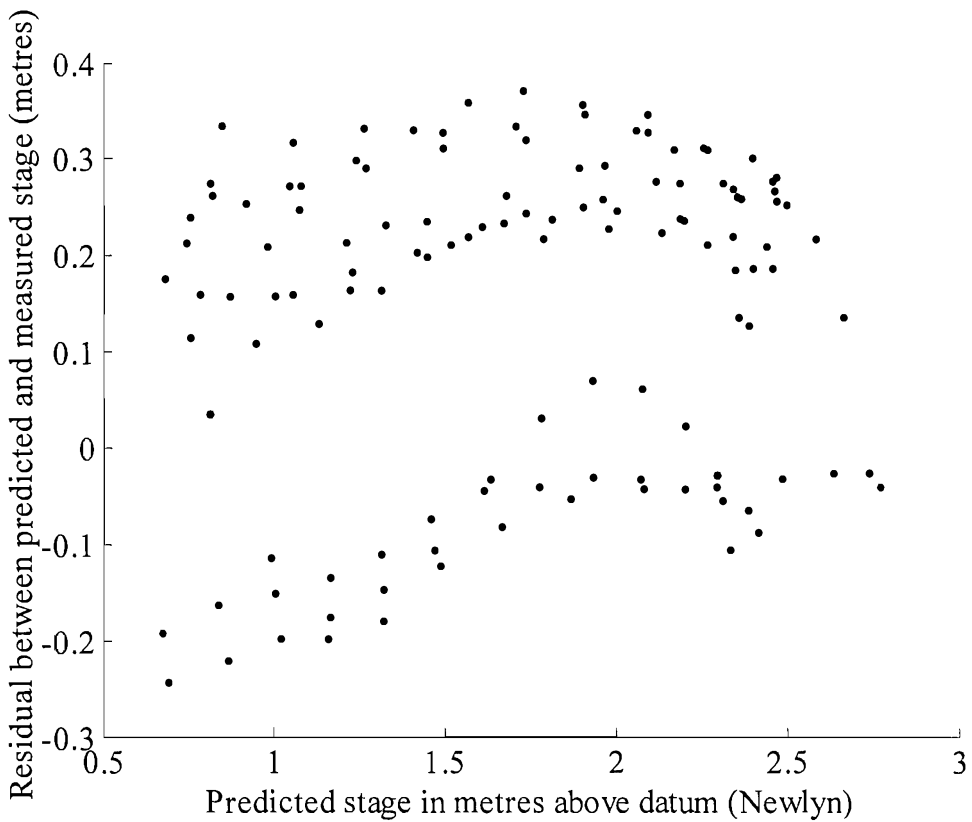


Fig. 4.11: Plot of residuals between RC model predictions and measured stage against predicted stage at sensor 1. Dots are joined by event.

Table 4.4: Table of RMSE in metres between forecasted and measured stage averaged over the forecast period.

Sensor	No Update	1 update time	2 update times	4 update times
1 (channel)	0.211	0.116	0.071	0.064
2 (channel)	0.250	0.231	0.178	0.165
3 (floodplain)	0.393	0.267	0.214	0.201
4 (channel)	0.289	0.243	0.190	0.177

An objective of the updating procedure was to facilitate a forecast of stage that was more accurate than the prediction by the simulation model alone. For the remainder of sections 4.5 and 4.6 a forecast by the RC model covers the period of time between the most recent state update (filter time step) and the expected time of high tide. The ensemble mean was assumed to be the most accurate estimate of state made by the model. Therefore, statistics regarding the accuracy of the forecast refer to the forecast mean, unless stated otherwise. Fig. 4.12 is a scatterplot of forecasted versus observed stage at sensor 1 after one and four filter time steps each of which was 15 minutes apart. Errors between forecast and measured stage in Fig. 4.12 are smaller than those between predicted and measured stage in Fig. 4.10. The magnitude of these errors also decreases as more filter time steps elapsed on an event as summarised by the RMSE errors in Table 4.4. The largest reductions in forecast RMSE occur as a result of the first two filter time steps (falling from 0.22 m to 0.07 m at sensor 1) indicating a rapid response to measurement data. Further decreases in RMSE were much smaller in magnitude and are to some extent likely to be the result of shorter forecast lead times. RMSE was smallest at sensor 1, especially for forecasts. However, this is not unexpected because most of the measurements early on in each tidal cycle were collected by sensor 1, during which time the ensemble was usually less constrained relative to later in an event.

Fig. 4.13 compares forecasted stage estimates at sensors 1 and 3 with measurements after four filter time steps. At both locations the magnitude of errors in stage forecasts decreased relative to the RC model alone in Fig. 4.10. However, the errors between the updated RC model and measurements from sensor 3 were systematically

larger than for sensor 1. This is because over-fitting of state to the more numerous measurements from sensor 1 caused systematic state errors at sensor 3 to remain uncorrected by the filter. The source of the systematic error may be due to an unrepresentative model structure or parameters. However, the consistency of the error over a range of stages and tidal events suggests measurement bias may be influential in producing this error. Nevertheless, the RMSE between forecasts and measurements after four filter time steps (Table 4.4) decreased by almost 70% at sensor 1 and 50% at sensor 3, when compared to the simulation model alone. Therefore, the EnKF based updating procedure facilitated the delivery of more accurate forecasts of river stage by the simulation model, although it was only possible to test over forecast lead times of less than 4 hours with the available data. Example forecasts are shown in Fig. 4.14.

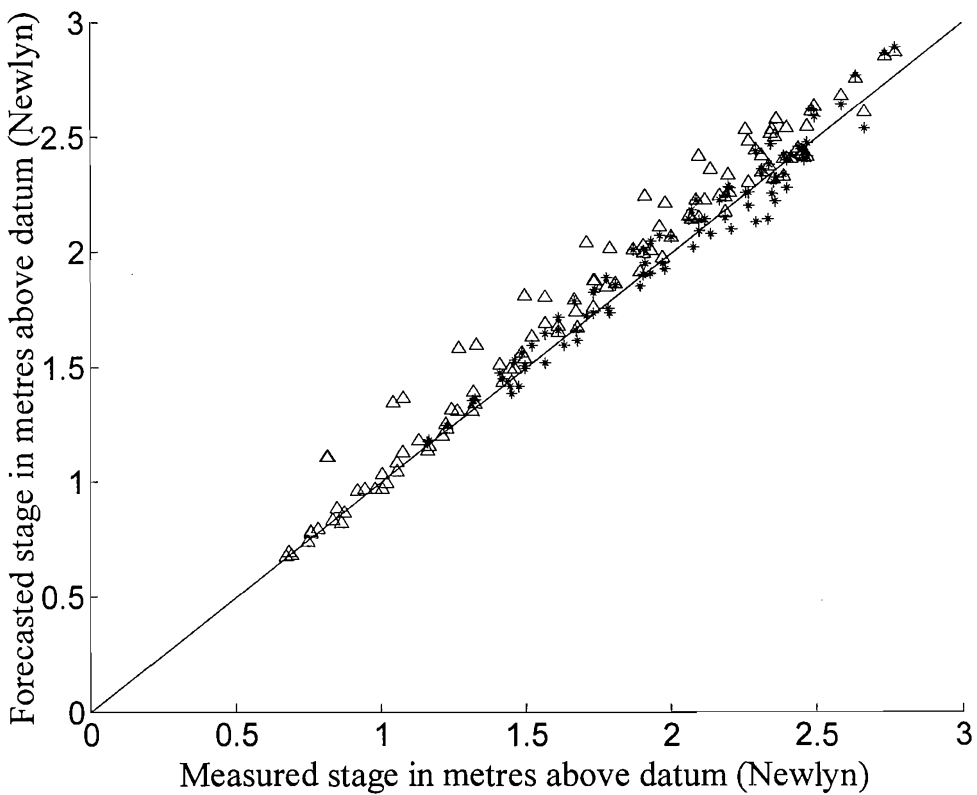


Fig. 4.12: Plot of updated RC model stage forecasts against measured stage at sensor 1 after one time step Δ and after four time steps $*$. Data are shown from 11 tidal cycles.

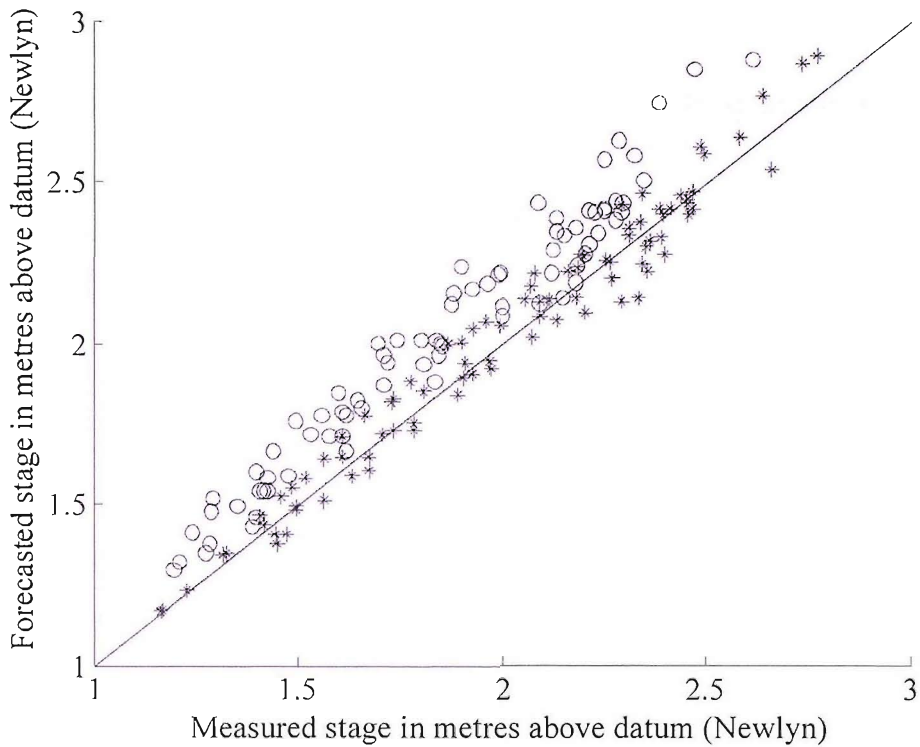


Fig. 4.13: Plot of updated RC model stage forecasts against measured stage at sensor 1 * and sensor 3 O, after 4 time steps.

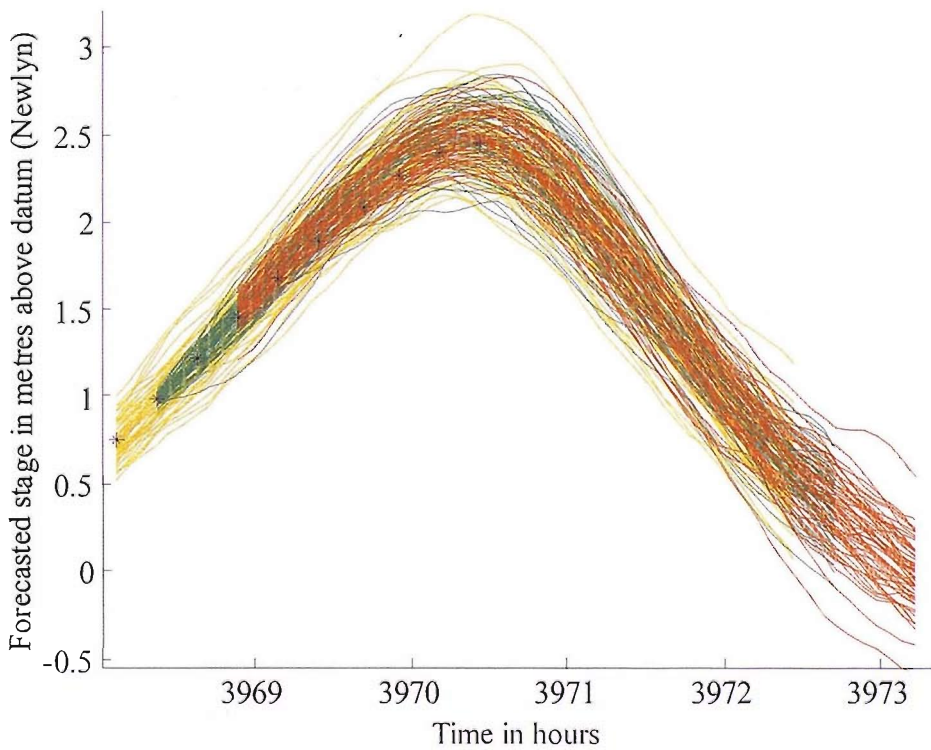


Fig. 4.14: Example forecast at sensor 1 during tidal cycle 6. Forecast stage after 1 (yellow lines), 2 (green lines) and 4 (red lines) filter time steps with measurement data overlain. Stage measurements are shown as stars.

4.4.2 Forecast accuracy in relation to sampling design

Four sampling strategies were simulated from the forecasting data archive. These were:

Run A: All data, Collect measurements every 15 minutes from sensors 1 to 4

Run B: Collect measurements every 15 minutes from sensor 1

Run C: Collect measurements every hour for sensors 1 to 4

Run D: Collect measurements every hour from sensor 1

The accuracy of stage forecasts when all spatial measurement data were included (run A) increased at sensor 3 relative to using only sensor 1 as shown by the upper plot in Fig. 4.15, where the squared errors in run A were subtracted from the squared errors of run B. No discernible increase in forecast stage accuracy was observed at sensor 1 (Fig. 4.15. lower plot). A similar effect was seen when comparing runs C and D. In some cases, the inclusion of spatially distributed measurements lead to an increase in errors at sensor 1. Once again this is the product of the systematic errors between locations 1 and 3. The outliers in the bottom left corner of Fig. 4.15 indicate a substantially less accurate forecast by run A, which occurred due to a combination of the state ensemble being less constrained early on in an event, large state error covariance between channel and floodplain and the measurement uncertainty at sensor 3 being small relative to that at sensor 1. Therefore, the inclusion of spatially distributed measurements led to an increase in forecast accuracy at the locations included away from sensor 1, at least for a short period of time. The inclusion of measurements from all four sensors had less effect overall on forecast accuracy at sensor 1 and was detrimental to forecast accuracy under some conditions (assuming unbiased measurements at sensor 1) again due to systematic errors and the large state error covariance at this site.

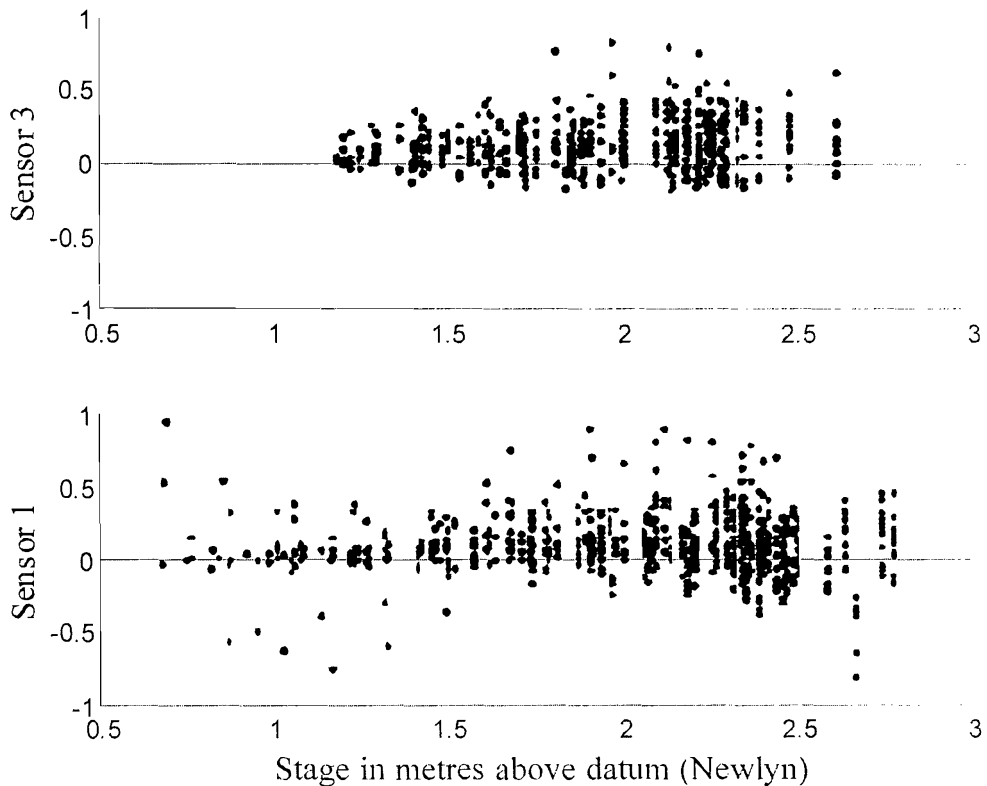


Fig. 4.15: Difference between squared errors in forecast from run A and run B, for sensor 3 (upper plot) and sensor 1 (lower plot). A positive value indicates an improved forecast of stage from run A over run B.

4.4.3 Estimating forecast uncertainty using the EnKF

The accuracy of the forecast mean is an important consideration of any forecasting model. However, to estimate flood risk (in this case derived from an ensemble of peak stage values) over the forecast period a realistic measure of uncertainty in peak stage is required (as represented by the variation within the ensemble). For any location within the model domain this is dependent on the error covariance estimates made by the simulation and measurement models and the error covariance propagated by the system model ensemble.

Fig. 4.16 shows the decrease in the standard deviation of analysed state at sensors 1 and 3 after each filter time step during an event. For runs A and B, with a 15 minute sampling interval, state uncertainty decreases rapidly over the first three to four filter

time steps before stabilising at around 0.05 m. This is seen at both sensors 1 and 3, although the first three filter time steps at sensor 3 are affected by predictions of zero flow at this location by some ensemble members, resulting in very small estimates of standard deviation. In this case study, almost all ensemble members have predicted flow at sensor 3 before measurements from this location are used for updating. However, had data not been available in the channel or the spread of the ensemble been such that a number of ensembles predicted zero flow on the floodplain, it is possible that there would be an overly low covariance between floodplain and channel node, potentially, allowing the floodplain to be updated whilst leaving the channel at a lower stage. Runs C and D exhibit a relatively small fall in state standard deviation as a result of the 1 hour sampling strategies compared to runs A and B. Therefore, the reduction in analysed state variance brought about by using four sensors rather than just one was small relative to the effect of altering the temporal sampling rate from 1 hour to 15 minute intervals. Although, the covariance between measurement locations was large due to the sensors being only around 500 m apart, meaning that there was a substantial amount of data redundancy spatially at this site. Fig. 4.16 also shows that run A produced fewer extremely small state standard deviation estimates compared to run B. Therefore, the inclusion of spatially distributed measurements tended to prevent the creation of the very constrained ensembles in run B, which resulted from the occasional very small estimate of measurement uncertainty by sensor 1. This effect illustrates a problem with using dynamic estimates of measurement uncertainty based on repeat sampling over a short period, in favour of a stationary estimate of measurement uncertainty calculated offline.

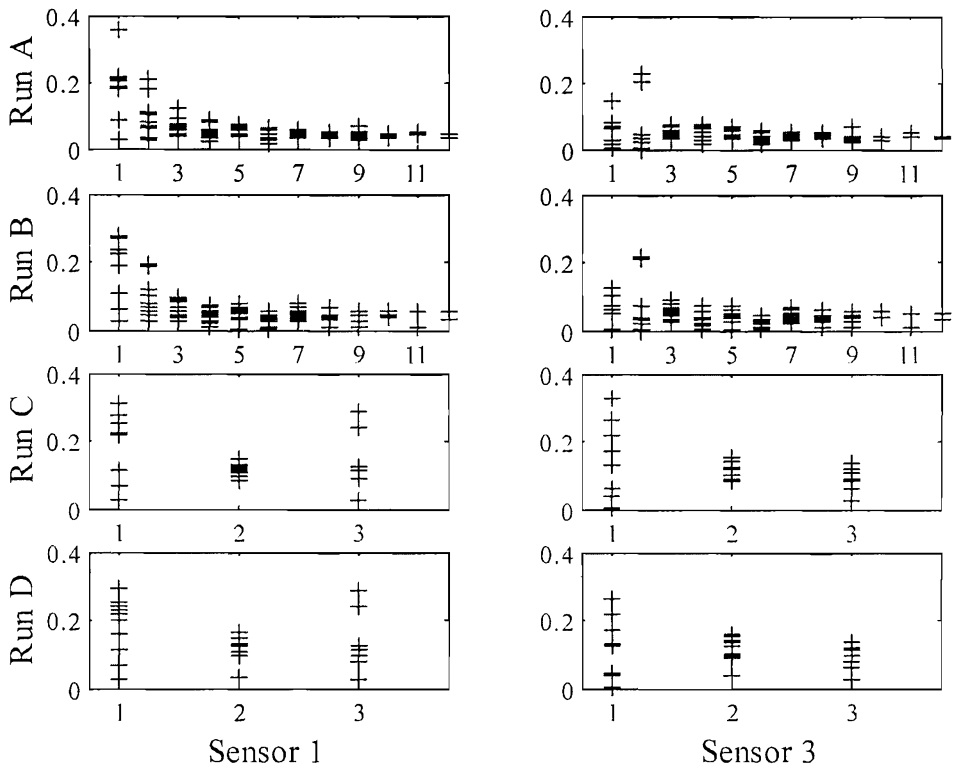


Fig. 4.16: Post analysis ensemble state standard deviation, with the number of measurement time steps / filter time steps shown on x -axis. The y -axis shows analysed state standard deviation in metres.

Table 4.5 analyses the suitability of the uncertainty estimates at each sensor location (r value) over the forecasting test period using the RMSE based method presented in section 3.4. Forecasted uncertainty at sensor 1 using a 15 minute sampling interval was the most representative of forecast accuracy, achieving an r value of 1.02. This result was expected as the boundary condition model was calibrated using this sensor and sampling setup. During run A forecast uncertainty was underestimated at sensors 2, 3 and 4, indicated by r values of 1.25, 1.51 and 1.52 respectively, mainly due to the systematic errors identified in Fig. 4.10 and 4.13. Therefore, the presence of systematic errors at some locations in either the RC model or measurement data adversely affected the filter's ability to adequately forecast uncertainty away from sensor locations. The average magnitude of forecasted stage errors at each sensor location increased with distance from sensor 1, although more sensors with greater distances between them will be required to assess to what extent this occurred

because of coincidental sensor bias or some form of model structural or parameterisation problem.

Table 4.5: r values for forecasted uncertainty with different sampling strategies denoted by: (A) 15 minute sampling interval from sensors 1 to 4, (B) 15 minute sampling interval from sensor 1, (C) hourly samples from sensors 1 to 4 (D), hourly samples from sensor 1. h indicates that the r value was calculated using hourly measurements only, whilst the ‘adjusted’ results use a linear regression model to remove systematic errors.

Sensor	Run A	Run Ah	Run B	Run Bh	Run Ch	Run Dh
1	1.03	1.05	1.04	1.04	0.91	0.94
2	1.25	1.26	1.29	1.27	1.34	1.34
3	1.51	1.46	1.53	1.51	1.59	1.51
4	1.52	1.58	1.54	1.49	1.65	1.66
3 adjusted	-	-	1.18	1.18	-	1.12

The residual errors between measurements and RC model predictions at sensors 1 and 3 are shown in Fig. 4.17 (left). A linear regression of these variables gives a gradient of 1.05 and offset of -0.141 m, with an R^2 of 0.90. If the gradient is set as 1 the offset changes to -0.147. The inverse of this offset was applied to sensor 3 measurements resulting in an r value of 1.18 at sensor 4 for run B. Therefore, accounting for the systematic error using a linear model improved the representativeness of the forecast ensemble. Fig 4.17 (right) plots the difference between residual errors at sensors 1 and 3 against measured stage at sensor 1 and indicates a slight linear relationship ($R^2 = 0.24$) between the difference in errors and stage, suggesting a similar stage-error relationship to that shown in Fig 4.10, indicating errors in the RC model or measurements are slightly stage dependent. Nevertheless, model errors are still greater at sensor 3 and therefore, an additional source of uncertainty is required between sensor 1 (channel) and sensor 3 (floodplain) for this particular model. Alternatively, it may be desirable to reduce this uncertainty by using a different model for the floodplain flows, especially as how best to characterise this uncertainty (given that the calibration procedure showed the roughness parameter had very little effect on flow) is not clear and likely to depend

on floodplain discretisation. Considering the coarse spatial discretisation of the floodplain and one-dimensional flow equations used by the RC model, a two-dimensional flow model may be better able to represent the characteristics of the measurement data, although at an increased computational cost.

The one hour sampling setup was found to be more accurate at sensor 1 than the uncertainty estimates suggest as indicated by r values of 0.91 and 0.94 for runs C and D, suggesting the filter gave insufficient weight to measurement data in this example. The tendency to over-estimate uncertainty at sensor 1 when using the coarser temporal sampling strategies of runs C and D suggests that a more representative propagation of uncertainty over time may be required, especially in the case of run D where the model was updated with data from the validation location only, thus, to some extent negating difficulties with systematic differences with other sensors. To clarify that this effect did not occur because the regularity of forecasts varied between runs [A, B], and [C, D], r values were calculated for runs A and B using hourly forecasts only (Table 4.5) and show only small differences from the 15 minute forecasts at sensor 1, relative to those between runs [A, B] and [C, D].

The use of the downstream boundary condition and an AR(1) model as a means of introducing simulation errors into the ensemble was in this case unable to generate representative stage uncertainty when exposed to different sampling regimes, although the uncertainties allocated to measurement data may also have a role here. Furthermore, uncertainty due to internal model errors, particularly between channel and floodplain were under-estimated with this setup. More measurements from other locations on the floodplain are required if the relative influences of model error and measurement error on predicted state uncertainty are to be identified. Given that the GLUE procedure conducted on the calibration data set identified uncertainty in the value of channel roughness the inclusion of this source of uncertainty in future versions of the EnKF-RC model seems a logical next step. The inclusion of additive or multiplicative noise into state estimates as filter hyperparameters in order to represent structural uncertainties may also be worth pursuing. The AR(1) model at the boundary condition may be more appropriate when not used to compensate for the under-representation of model structural and parameter errors. However,

substituting the tidal harmonic with an ensemble of tide and surge predictions would be a useful experiment for generating longer forecast lead times and accounting for structural changes to the downstream boundary other than stage magnitude.




Fig. 4.17: Scatter plots of forecast error at sensor 1 against sensor 3 and difference in forecast error against stage at sensor 1.

4.5 Conclusions

The RC simulation model was applied in a forecasting context by updating state predictions with an EnKF. Forecasts made after updating model state and boundary conditions were more accurate in terms of RMSE than pre-update simulations by 0.147 m to 0.046 m depending on measurement location, forecast lead time and number of measurements assimilated (Table 4.4). The dominant factor preventing further increases in analysed state accuracy was the presence of spatially-distributed systematic errors either in the RC model or measurements.

The RC model state ensemble converged after 2-3 measurements from sensor 1 early on in each event, meaning that the more numerate measurements from this sensor and any accompanying systematic errors tended to dominate updated state predictions not just locally but globally over the study site.

The filter over-predicted state uncertainty, as indicated r values of 0.91 and 0.94, when the sampling rate was reduced to hourly, indicating that measurements were not given enough weight relative to model predictions. Estimation of state uncertainty was also hampered by spatially-distributed systematic errors. Although the EnKF-RC model formulated here could be calibrated to make representative estimates of uncertainty at a single location, the extrapolation of these uncertainties to other locations was not adequate. When the systematic error at sensor 4 was removed state uncertainty was still under-estimated at this location, as shown by an r value of 1.18. Therefore, additional uncertainty sources within the simulation model domain were required along with an improved method for determining filter hyperparameters. Furthermore, a two-dimensional model of the floodplain flows may be able to replicate more accurately the observed differences between channel and floodplain stage measurements. More spatially distributed sensor data are required to establish to what extent the systematic errors observed are due to model or sensor bias. Ideally further tests of this approach would involve a complete temporal coverage of measurements from at least one location and an increased distance between sensors.

4.6 Summary

This chapter examined critically the application to a site along the River Crouch, Essex of a river-flow forecasting approach based on a one-dimensional hydraulic flow simulation model updated using real-time data within an ensemble Kalman filtering framework. Given a specified validation location and forecast period the objective of the forecasting model was to estimate water level more accurately with updating than without. The method used to estimate both model state and state uncertainty was evaluated in terms of its forecast accuracy and representation of forecast uncertainty. The ensemble Kalman filter lead to an increase in forecast accuracy of between 50% and 70% depending on location. The hyperparameters of the filter could be calibrated to make estimates of forecast uncertainty at a specific location, where the most data were available. However, the presence of systematic errors in the simulation model and especially measurement data meant that uncertainty estimates were inaccurate at other locations. Although, the major source of uncertainty in this model came from the boundary condition, additional uncertainty within the model domain was required, particularly between channel and floodplain. Changing the temporal sampling rate and spatial density of samples had little effect on the accuracy of forecasts at this site. However, uncertainty was underestimated when the temporal sampling rate was decreased, indicating that the relative uncertainties prescribed to the simulation model and measurement model were inadequate.

5 Evaluating the utility of the ensemble transform Kalman filter for adaptive sampling when updating a hydrodynamic model

5.1 Introduction

Numerical models that simulate river states variables such as flow and stage can be used for the purpose of forecasting by incorporating real-time observations using a process often referred to as model updating. Assuming that the accuracy of forecasts produced by such a system depends on the accuracy of the simulation model, the efficiency of the updating routine and the measurement data available (Refsgaard, 1997). The sampling framework which determines the supply of measurements to the updating routine will also affect forecast accuracy. For example, when updating model states collecting too few data can lead to consequent uncertainties in the initial state, which when propagated through to a forecast time, can lead to unacceptably uncertain state forecasts. Collecting many data may provide a more precise (in terms of variance) definition of initial state, but conversely may be inefficient due to data redundancy. Bearing in mind that there is generally a cost associated with measuring, storing and utilising data, optimising data acquisition becomes important, particularly where model outputs may be utilised in a real-time decision making process. This chapter addresses the issues associated with the application of this process to flood modelling, with a particular interest in the application of adaptive data sampling.

Data on river flow, level and flood inundation extent can be collected by *in situ* sensors connected to data loggers or, for inundation extent, by remote sensors (Horritt, 2001). Where real-time data are desired gauge data can be transferred off-site using telemetry (Knott, 1999). Systems designed for processing real-time data to produce flood forecasts tend to be individually tailored to meet various requirements such as the spatial scale over which they are required to operate, the desired accuracy of forecasts and the length of the lead time (Werner *et al.*, 2005). Pre-warning systems such as the European Flood Forecasting System (De Roo *et al.*, 2003), have

the ability to provide lead times of 5-10 days through the utilisation of medium range numerical weather prediction models. At the forecast lead times associated with pre-warning systems the assimilation of river flow data is generally unnecessary since any effect on state variance is likely to be washed out of the model domain before the 5-10 day lead time.

At the river basin scale, operational flood forecasting systems begin to take account of measured hydrological and meteorological data to enable a more accurate forecast than can be achieved with the pre-warning systems. Examples include the River Flow Forecasting System (Moore *et al.*, 2004), and the Hydrologic Engineering Center's Hydrologic Modelling System/ River Analysis System (HEC-HMS/RAS) (Knebl *et al.*, 2005). In large river basins a combination of rainfall-runoff models, routing models and gauge data will generally give adequate lead times, such that flood warnings can be issued. For example, the FloRIJN system on the Rhine has a 3-4 day lead time between flows observed upstream and their downstream impact (Sprokkereef, 2001).

The motivation behind the study presented here was to utilise data collected by a network of wireless sensor nodes, with limited power resources for transmitting data in real-time (De Roure, 2005), to forecast water level (stage) and flow in a short lag time response basin. The deployed sensor nodes carried out point measurements of stage at specific locations along the main channel of the river and on the floodplain, using pressure transducer based instrumentation.

Each sensor node on the network was 'intelligent' in that it could be reconfigured, perform automated decisions based on protocols and carry out basic data processing in real-time. Due to a research interest in power conservation the nodes were configured to transmit real-time data at a low background reporting rate, but could be reconfigured to increase reporting rates at specific data-critical time periods.

Adapting the framework under which data are transmitted as an event of significance unfolds is necessary if these limited but flexible resources are to be used efficiently. Key to this adaptive sampling was a method of prioritising at what time and,

critically, from which sensors to transmit data to a forecasting model, a process otherwise known as targeting. Describing adaptive sampling in a data assimilation context led Bishop *et al.* (2001) to suggest an extension to the unified data assimilation notation proposed by Ide *et al.* (1997). The motivation behind the extension to this notation was that it is necessary to distinguish between routine and adaptive components of the observation network, as well as defining the times at which sampling decisions are made and the validation times and locations for which the sampling is being optimised. Table 5.1 describes the suggested notation as described by Bishop *et al.* (2001), whilst Fig. 5.1 illustrates this notation for a simple scenario. This notation is adopted in this thesis.

Table 5.1: Description of adaptive sampling notation.

Notation	Description
$()^r$	Superscript r denotes routine measurements. These measurements are those expected regardless of any possible adaptive components to the sampling framework. Decisions regarding the deployment of adaptive resources should take into account the likely effect of these measurements on error covariance.
$()^q$	Superscript q denotes adaptive measurements. There may be many Q possible combinations of measurements (sometimes referred to as measurement networks) that could form the adaptive component. It is usually assumed that q includes routine measurements collected at the same time.
$()_m$	Subscript m denotes the ‘target time’ at which the adaptive observations are to be collected.
$()_d$	Subscript d denotes the time at which a decision regarding which of the Q possible measurement combinations will be collected must be made such that time is available to collect the data.
$()_v$	Subscript v denotes the validation time and location time and sometimes location for which the observation network has been optimised. This time can be the same as the target time, but is more likely to be some point in time after the target time.




Fig. 5.1: Adaptive sampling timings.

This chapter evaluates a method, known as the ensemble transform Kalman filter (ETKF) (Bishop *et al.*, 2001), to estimate quickly the error variance of a stage forecast given one or more possible stage measurements from the sensor network. The ETKF is part of a family of sequential data assimilation algorithms based around the Kalman filter and was introduced in section 3.6. Given the information from the ETKF (i.e., given that the ETKF works satisfactorily), the expected best measurement times and locations in terms of forecast stage error variance can be targeted. This chapter does not execute the actual targeting, which is left for the next chapter, but rather focuses on a thorough evaluation, using archive data from the FloodNet sensor network (section 6.1.1), of the ETKF method on which the targeting will be based.

5.2 Implementation of the ETKF on the RC model.

In this thesis, a numerical model was used to simulate stage dynamics. The model was a one-dimensional hydrodynamic model of a 22 km reach of the River Crouch (Essex, UK), referred to hereafter as the RC model, and is known to be process nonlinear. Therefore, a Monte Carlo based approximation of the KF called the ensemble Kalman filter (EnKF) (Evensen, 1994), was used to perform the state updates. The RC model and the setup of the EnKF were described in more detail by chapter 4.4. However, it is worth noting that propagating the ensemble of RC model simulations after each measurement time to the next and issuing ensemble forecasts to a validation time were computationally expensive relative to implementing the EnKF update (see Chapter 4). Together, the RC model and EnKF form what is described as the RC forecasting model in the remainder of this paper.

The forecasting model as presented in chapter 4 is unsuitable as an adaptive sampling tool because the EnKF required the measurement magnitudes to update the state ensemble. If measurements with appropriate covariance are simulated, the RC forecasting model could be used to choose between different measurements sets. However, repeatedly propagating the ensemble to estimate the effect of different measurement sets on forecast error variance would be computationally unworkable. Therefore, the principal reason for using an ETKF is that it can rapidly obtain signal variance estimates at future times using a linear transformation of ensemble state perturbations (Majumdar *et al.*, 2002). Providing it is acceptable to assume that the propagation of the ensemble perturbations is approximately linear, which is presently untested in a flood forecasting and ETKF context, the same transformation can be applied to the ensemble perturbations at the validation times. The advantage of this is that once the perturbations have been generated (by the RC model in this case) they can be used repeatedly for different sample sets, at no additional computational cost except that incurred by calculating the transform and transforming the perturbations (Majumdar *et al.*, 2002). The ETKF estimates the effect of potential measurements on the ensemble of state perturbations, rather than the states themselves as would be the case with an EnKF. Therefore, it does not update mean state or require the likely magnitude of the measurement to be known (Wang and Bishop, 2003).

To the authors' knowledge, the ETKF has not been applied previously to a hydrodynamic model. The principal aim of this chapter was, therefore, to verify the signal variance forecasts made using the ETKF by comparison with stage variance estimates from the RC forecasting model. The results of this comparison are presented in section 5.4, followed by a discussion and a conclusion. The ability to use the ETKF to support adaptive sampling, e.g., in future studies, depends critically on the benchmarking of the ETKF against the EnKF presented in this chapter.

For simplicity, the method for estimating forecast error variance prior to data collection will be known as the ETKF approach, the state updating algorithm will be known as the EnKF approach, whilst the simulation model on which both the ETKF and EnKF depend will be referred to as the RC model. The next section describes the setup of the simulation model which generates the ensemble of flow and stage (these

include background forecasts for the ETKF), the EnKF based model updating method and the measurement data.

5.2.1 Simulation model and updating procedure

A hydrodynamic model of the River Crouch (RC model) was built on the ISIS flow program (HR Wallingford) and was used to estimate within channel-flow by solving numerically a dynamic wave model that comprises a pair of one-dimensional non-linear partial differential equations, based around the St Venant equations (Cunge *et al.*, 1980). As the simulation model used was presented in section 4.4 the following section acts as a reminder of the model function and parameterisation rather than a comprehensive description, although an adjustment to the treatment of roughness uncertainty is made. The RC model domain extended 21 km along the river from Battlesbridge to Burnham-on-Crouch. Out-of-bank flows at the study site were simulated using Manning's equation and represented using a network of 22 interlinked storage cells (reservoirs). The RC model was used to simulate an ensemble of river flows based on ensembles of initial conditions, tidal boundary conditions and, in a change from section 4.3, the global roughness parameter.

Forecasting was implemented by updating the model's initial and boundary conditions with measured data, using an EnKF (Evensen, 2003). This algorithm updated estimates of initial stage and flow in a near optimal manner (in terms of minimising error variance). The updated or analysed initial and boundary conditions were then used to initialise further model simulations. The RC model can be described by the following state-space formulation.

$$\mathbf{x}_k^f(t_{i+1}) = M(\mathbf{x}_k^a(t_i), \theta_k, u_i + \varepsilon_k(t_i)) \quad (5.1)$$

Where, a state vector \mathbf{x} is made up of the state variables stage and flow, M is the non-linear model operator which, in this case, represents advancing the hydraulic model state from t_i to t_{i+1} , θ represents the global channel and floodplain roughness parameters, sampled from a pre-defined distribution, superscript a denotes analysed state (*a posteriori*) whilst superscript f is forecasted state (*a priori*) and u represents

the downstream boundary conditions which have been corrupted by errors ε . These errors were assumed to be correlated in time and were approximated by a first-order autoregressive model (Bras and Rodrigues-Iturbe, 1985; Madsen and Cañizares, 1999):

$$\varepsilon_k(t_{i+1}) = \rho\varepsilon_k(t_i) + \sigma_\varepsilon(1 - \rho^2)^{1/2}W_k(t_i) \quad (5.2)$$

where ρ is a lag-one autocorrelation coefficient, σ is the variance of the error process and W is a zero-mean Gaussian white noise random variable with variance of 1. The variance σ and autocorrelation coefficient ρ of this model are estimated in section 5.2.2. The downstream boundary condition u was set using tidal predictions at 10-minute intervals for the port of Burnham-on-Crouch. These were supplied by The United Kingdom Hydrographic Office and are based on adjusted harmonic constants calculated for the standard port of Walton-on-the-Naze (Essex). A flow-time hydrograph was used at the upstream boundary and was set at a constant $0.35 \text{ m}^3\text{s}^{-1}$ based on the mean flow recorded by the Environment Agency gauge at Wickford, approximately 2 km upstream of Battlesbridge.

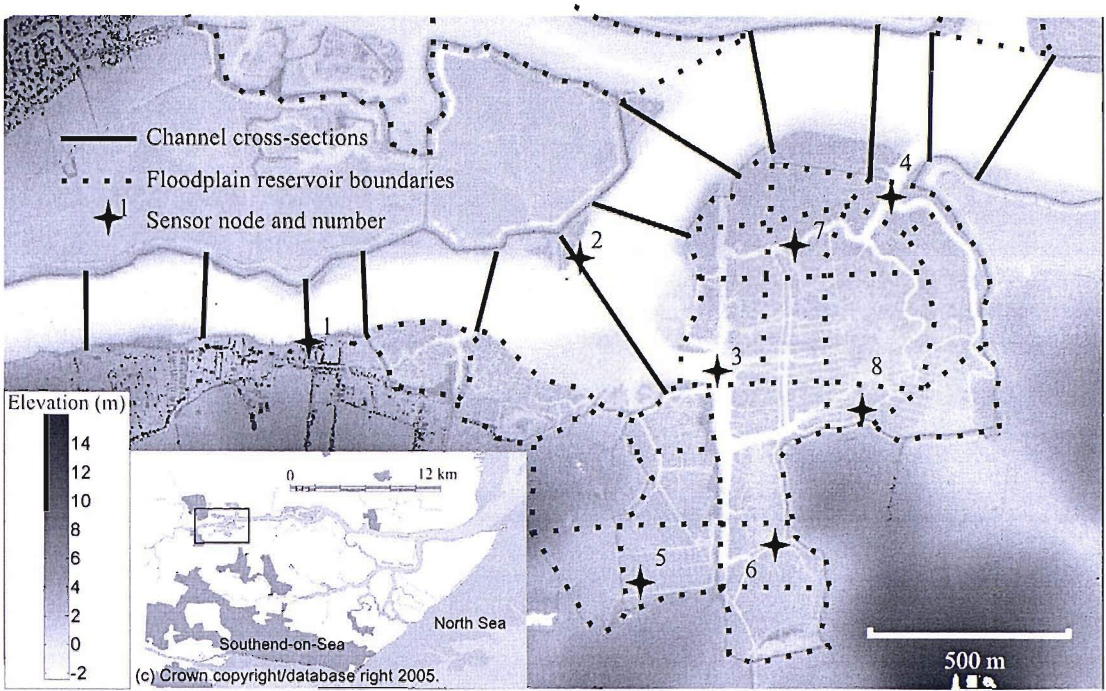


Fig. 5.2: Map of sensor locations and RC model topology around the validation location at sensor 2.

An archived measurement data set from eight sensors (Fig. 5.2) was used in this study. Measurements were available at 5 minute intervals from the 4th to 26th November 2005. This will be referred to as the test period and was chosen because it represented the only period of time during which eight sensors were operational at this site. As the principal purpose of this test was to compare the two error variance estimation methods all measurements were assigned a variance of 0.005 m to simplify the procedure. This information was used to generate Gaussian measurement ensembles, which were mapped onto the RC model state vector using the linear measurement operator \mathbf{H} . Four of the eight sensor nodes were located in the channel with the rest spread over the salting. The varying heights of the sensors meant that most were dry at some point over the tidal cycle. The measurement model was defined as (Burgers *et al.*, 1998):

$$\mathbf{y}_{i,k} = \mathbf{H}_i \mathbf{x}_k^{true}(t_i) + \eta_{i,k} \quad (5.3)$$

where \mathbf{y} is the measurement vector and η is a Gaussian white noise vector representing the errors in the measurements. \mathbf{H} is the measurement operator that

maps measurements \mathbf{y} onto the state vector \mathbf{x} and is linear in this example. Errors in measurements and state estimates were assumed to be uncorrelated, whilst each measurement was assumed to map directly onto a point on the model state vector.

The EnKF analysis step updated each model state ensemble member to give updated initial conditions for the RC model using (Evensen, 2003):

$$\mathbf{x}_k^a(t_i) = \mathbf{x}_k^f(t_i) + \mathbf{P}^f(t_i) \mathbf{H}_i^T (\mathbf{H}_i \mathbf{P}^f(t_i) \mathbf{H}_i^T + \mathbf{R}_i)^{-1} (\mathbf{y}_{k,i} - \mathbf{H}_i \mathbf{x}_k^f(t_i)) \quad (5.4)$$

where \mathbf{P}^f is the forecast state error covariance estimated from the forecast state ensemble (Evensen, 1994). The RC model and forecasting approach were described in more detail in section 4.3.

5.2.2 Parameterisation issues

A generalised likelihood uncertainty estimation (GLUE) procedure (Beven and Binley, 1992; Romanowicz and Beven, 1998; Pappenberger *et al.*, 2005) was used to estimate channel and floodplain roughness uncertainty. The procedure assumes equifinality in that many combinations of parameters and boundary conditions may result in acceptable simulated results, given the data available to validate each simulation. The procedure comprised 1500 simulations by the RC model from the 4th to 26th November 2005. The three variables selected for calibrations were channel roughness, floodplain roughness and the magnitude of errors at the tidal boundary condition. The Nash-Sutcliffe efficiency measure (Nash & Sutcliffe, 1970) was used to calculate the likelihood of each parameter set:

$$\ell = \left(1 - \frac{\sigma_e^2}{\sigma_o^2} \right) \quad (5.5)$$

where σ_o^2 is the variance of the observations and σ_e^2 is the variance in the residuals between observations \mathbf{y} and model state predictions \mathbf{x} . The method was identical to that of section 4.3.2 although the measurement data were different.

The purpose of the calibration was to estimate channel roughness uncertainty such that an ensemble of channel roughness parameters could be generated. The resulting roughness mean and standard deviation are shown in Table 5.2.

The parameters of the tidal boundary autoregressive model were estimated from the errors e between measurements and a best estimate simulation by the RC. The simulation used the most likely roughness parameters found during the GLUE procedure and unaltered tidal boundary conditions. The variance and autocorrelation coefficients were estimated from sample data using measurements from sensor 2 (Bras and Rodrigues-Iturbe, 1985):

$$\hat{\sigma}^2 \cong s_x^2 = \frac{1}{N} \sum_{i=1}^N (e_i - \bar{e})^2 \quad (5.6)$$

and

$$\hat{\rho} \cong r = \frac{\frac{1}{N} \sum_{i=1}^{N-1} (e_i - \bar{e})(e_{i+1} - \bar{e})}{s_x^2} \quad (5.7)$$

where the over-line denotes the mean of the errors between simulated and measured stage. Sensor 2 was chosen because it provided the most continuous record of stage over the study period. It is noted that the statistics calculated here may differ from those that would have been calculated had measurements of stage been available at the model boundary. The resulting coefficients and model parameters are shown in Table 5.2. Note that these calibration methods were chosen for their simplicity and ease of implementation, and are designed to give reasonable parameters to support the subsequent tests rather than provide an example of best practice. An example of a parameter estimation algorithm that is integrated with an EnKF can be found in Vrugt *et al.* (2005).

Table 5.2: Selected RC model parameters

Mean channel roughness $\bar{\theta}$	0.030
Channel roughness standard deviation σ_{θ}	0.004
AR model correlation $\hat{\rho}$	0.987
AR model variance $\hat{\sigma}_{\epsilon}^2$	0.322

5.3 Results: A comparison between EnKF and ETKF forecasts of error covariance.

The suitability of the ETKF as a means of targeting measurements was assumed to be dependent on its ability to estimate analysed error variance and signal variance, which resulted from assimilating the same measurements with the EnKF based forecasting model, as well as the ETKF's ability to estimate forecast error variance and signal variance at future validation or target times. This leads to two questions, which form the focus of the present investigation:

1. To what extent can the ETKF estimate the analysed error variance and signal variance that results from implementing the EnKF analysis (Eq. 5.4) at the measurement time.
2. To what extent can the ETKF estimate the forecast error variance and signal variance that results from implementing the EnKF at the measurement time and propagating the updated initial conditions to future measurement times and the validation time with the RC model.

Section 5.3.1 compares analysed error variance estimates from the ETKF and EnKF methods. The comparison of forecast error variances has been split into those obtained using a single measurement location (section 5.3.2) and those obtained using multiple measurement locations (section 5.3.3). In both cases, two parallel implementations of the EnKF were run: one using real measurements and the other

using synthetic measurements generated from the ensemble state mean, but with the same error statistics.

5.3.1 Analysed variances

A comparison of the EnKF and ETKF analysis algorithms was conducted by assimilating measurements from sensor 2 (BNG easting 582758) into a background ensemble of simulations produced by the RC model. Measurements were assimilated during 39 tides at eight times equally spaced between 3 hours 40 minutes and 1 hour 20 minutes prior to high tide. The estimates of analysed error variance and signal variance were all conditioned on single measurements, and were independent of previous measurements in this sense. Fig. 5.3 is a scatter plot of analysed error variance estimated by the EnKF (Eq. 5.3-4) against analysed error variance estimated by the ETKF (Eq. 3.59-61), whilst Fig. 5.4 shows the same but for signal variance.

There was a correlation of 1.000 between the analysed signal variance estimates in Fig. 5.3, and a correlation of 1.000 between estimates of error variance in Fig. 5.4. These correlations between the EnKF and ETKF were expected as the two algorithms use the same forecast error covariance, making them equally sensitive to measurements, unlike the example in Majumdar *et al.* (2002). The same correlations were found when using four sensors instead of one. This indicates that at the measurement time the ETKF analysis is constant with the KF analysis when given the same error statistics.




Fig. 5.3: ETKF analysed error variance plotted against the EnKF analysed error variance, for 39 tidal events at the validation location, after six target times.




Fig. 5.4: ETKF analysed signal variance plotted against the EnKF analysed error, for 39 tidal events at the validation location, after six target times.

5.3.2 Forecast variances: Single sensor example

Having established that the ETKF estimates of analysed error variance at the routine measurement time were identical to those of the EnKF, the focus was shifted to its ability to predict forecast error variance at future times. Ideally, the estimates of error variance and signal variance from the EnKF-based forecasting model and the ETKF would be equivalent. However, as the processes underpinning the RC model are known to exhibit non-linear behaviour there is unlikely to be an exact correspondence.

A direct comparison between the ETKF and EnKF was made for a scenario where a time series of measurements from sensor 2 (BNG easting 582758), with variance 0.005, were assimilated at 20 minute intervals from 3 hours 40 to 1 hour 20 minutes prior to the validation time. The time when the first of these measurements was taken will be referred to as the routine measurement time t_r (Fig. 5.1). The experiment was set up in five steps and conducted repeatedly over 39 tidal cycles.

Step 1: The ETKF was run on the background forecasts to estimate signal variance (Eq. 3.65) and analysed error covariance (Eq. 3.59-61) at the routine measurement time t_r .

Step 2: The transform \mathbf{T} from step 1 was used to update the background ensemble state perturbations at the validation and target times (Eq. 3.64). These updated ensemble perturbations were then used to estimate forecast error variance, expected from further observations at the next target times (Eq 3.59-61).

Step 3: The EnKF was used to update the ensemble state based on a real measurement with the same variance as the synthetic measurement used in steps 1 and 2 (Eq. 5.4). The updated state was then propagated by the RC model (Eq. 5.1-2) to the target and validation times.

Step 4: Ensemble state perturbations generated after the measurement had been assimilated in step 3 were then calculated. These were then used as a background

forecast by the ETKF to estimate the signal variance from further measurements at the target times (Eq. 3.65).

Step 5: Steps 1 to 4 were repeated for a further measurement collected by sensor 10 at the first target time 20 minutes after the measurement time.

Ensembles of stage simulations produced by the RC model after steps 3-5 are shown for an example tide in Fig. 5.5. The background forecast ensemble made at time t_i is also shown as lighter lines behind the updated forecast. Fig. 5.6 plots the ETKF (red) and EnKF (black) predictions of error variance at the validation time and seven target times at 20 minute intervals from the routine measurement time, against the background error variance at those times. Background error variance refers to the variance before the measurement was assimilated. Thus, each simulation was dependent on the RC model and the k^{th} roughness parameter and boundary condition. All RC model nodes are included as single points on the plots.




Fig. 5.5: Time series plots showing forecast stage at sensor 2 from target times 1 to 6 (dark lines). Observations were available from sensor 2 at each target time. The background forecast is indicated by the light lines, whilst measurements at five minute intervals from sensor 2 are shown as stars.

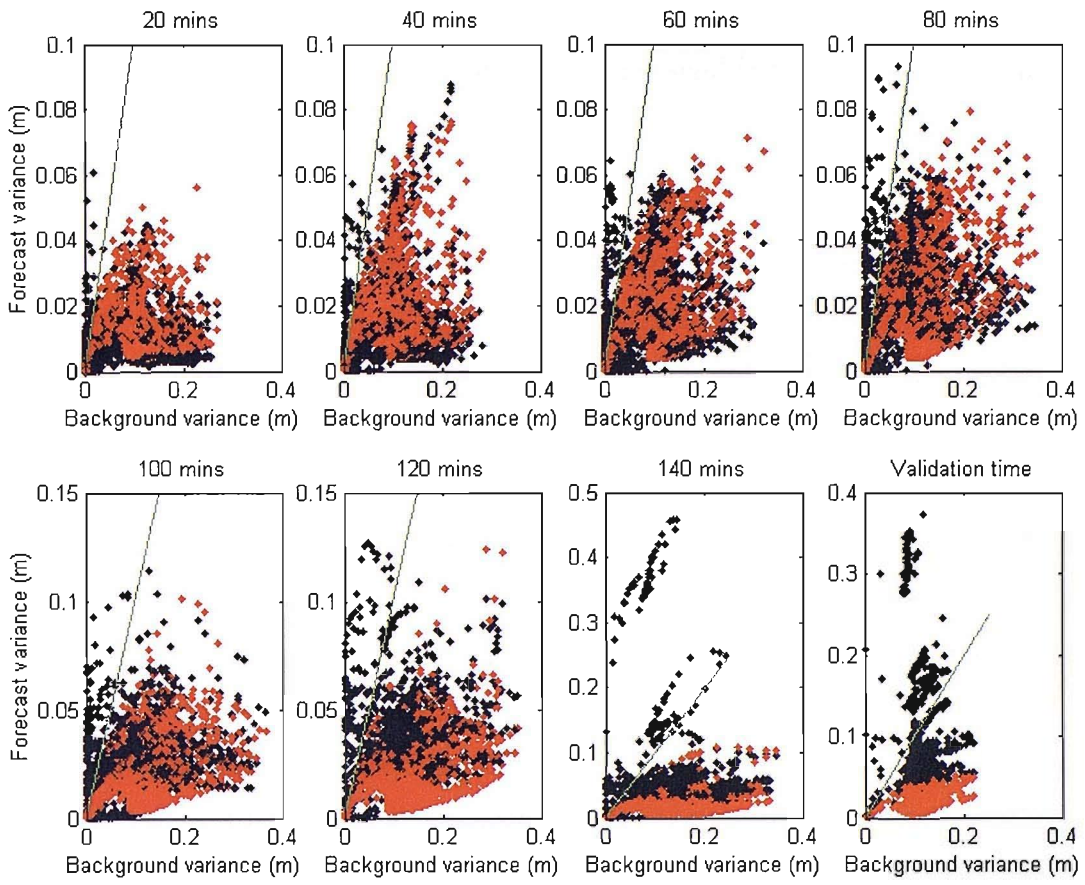


Fig. 5.6: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times from the assimilation of the routine measurement against background error variance. The green line is a representation of perfect agreement between the EnKF and ETKF forecast variances.

A continuous line demarks $y = x$. For any point above this line error variance increased as a result of assimilating the routine measurement, conversely error variance decreased for any point below the line. For target times close to the routine measurement time the scatter of points predicted by the EnKF and ETKF were similar, both in terms of the extent and clustering of points. As the time between the routine measurement time and the target time increased the ETKF approach began to under predict the variance of the EnKF approach. The correlation between EnKF and ETKF error variance estimates decreased as the time from the routine measurement increased (Table 5.3, correlation A). Table 5.3, correlation B, shows the correlation between the background forecast error variance at each target time and the validation

time. The correlation increased as the lag between times decreased. The decrease in correlation coefficient A (Table 5.3) from 0.796 to 0.292 suggests that the dynamics of the updated ensemble perturbations about the ensemble mean $\bar{\mathbf{x}}(t|\mathbf{H}_r)$ are different to those of the background forecast $\bar{\mathbf{x}}(t|\mathbf{H}_i)$.

When signal variance was compared the two approaches were more alike. The correlations between ETKF and EnKF signal variance (Table 5.3, correlation C) remain above 0.95 for two hours after the routine measurement, but drop to 0.679, 20 minutes later and 0.445 by the validation time.

Table 5.3: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at seven times between a measurement time t_r and a validation time t_v . From a background forecast t_i an EnKF was used to assimilate one measurement t_r . Forecast error variance was then calculated by propagating an ensemble of RC model simulations to t_v . The ETKF used the same background forecast and the same measurement error to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.

Time from t_r	20	40	60	80	100	120	140	220
Time to t_v	200	180	160	140	120	100	80	0
Correlation A	0.796	0.870	0.750	0.597	0.568	0.550	0.292	0.397
p_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Correlation B	0.005	0.004	0.036	0.051	0.053	0.462	0.832	1
p_value	0.000	0.001	0.000	0.000	0.000	0.000	0.000	-
Correlation C	0.997	0.993	0.983	0.969	0.968	0.959	0.679	0.445
p_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

A box plot of the signal variance at the target and validation times (Fig. 5.7) shows a tendency for the ETKF to over predict the EnKF signal variance, especially at the validation time where the upper quartile of the EnKF distribution is less than the

lower quartile of the ETKF distribution. There were some noticeable EnKF outliers at the last target time and validation time caused by model instability during two of the 39 tides. These increases in variance show up as clusters of points above the continuous line in Fig. 5.6 and correspond to tides 13 and 40. The ETKF rarely predicted that the error variance would increase after assimilating the routine measurement and never to the extent of the EnKF approach. Relative to the ETKF, the EnKF based approach predicted more occasions where error variance increased after assimilating the routine measurement (Fig. 5.8). The most noticeable instances of increased error variance occur for very low values of background error variance. Conversely, the EnKF was also observed to predict greater decreases in error variance than the ETKF, when the background variance was low.




Fig. 5.7: Signal variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Boxes show the upper quartile, median and lower quartile of forecast signal variance; whiskers extend over the range of the data with crosses indicating outliers.




Fig. 5.8: Forecast error variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Boxes show the upper quartile, median and lower quartile of forecast error variance; whiskers extend over the range of the data with crosses indicating outliers.

A long section of the river, plotting the background forecast of stage and error variance during a tidal event (Fig. 5.9), shows that the lowest values of signal variance occurred upstream of the incoming tide. To a certain extent this is a trait of this river because the fluvial component of the flow was small relative to the tidal component. However, it does confirm that the difference in error variance between the EnKF and ETKF method when the background error variance was low was an effect seen at the edge of the incoming tidal wave. In this case, the edge included areas of the floodplain not already inundated. Therefore, the location of the front edge of the incoming tide was more sensitive (in terms of variance) to changes in the ensemble mean than the main body of the estuary, where it seems more appropriate

to assume that the dynamics of the ensemble perturbations are similar after assimilating the routine measurement.

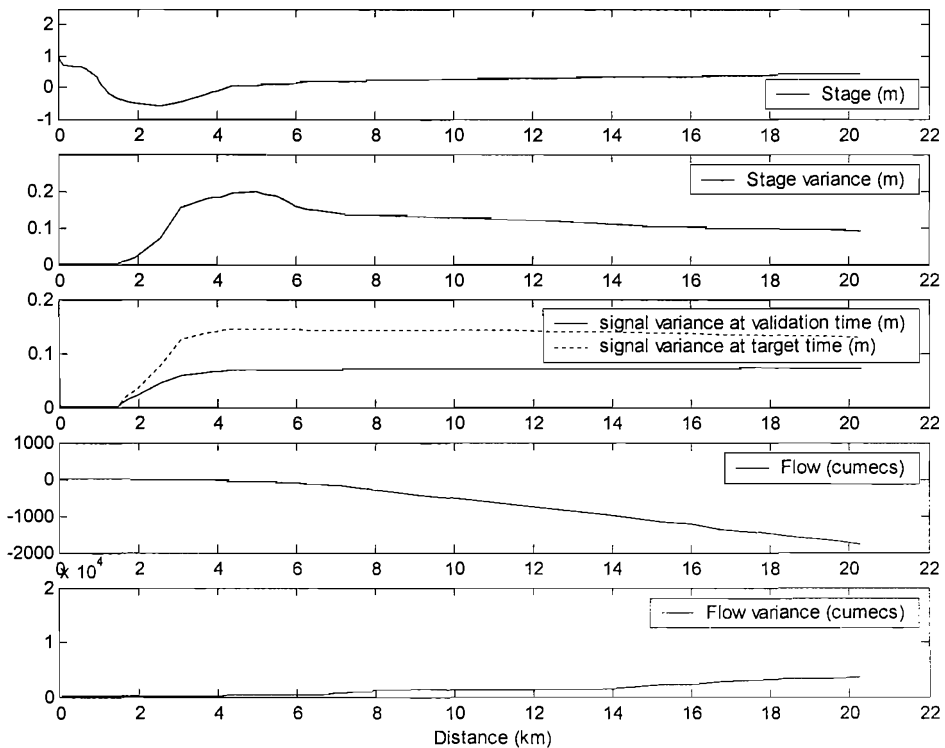


Fig. 5.9: Plot of model background states (stage and flow) and their variance for channel cross sections, with signal variance at validation time and location shown for theoretical measurements taken at any RC model node along the river at target time one. Measurement variance was 0.005 m.

The propagation of ensemble perturbations was sensitive to changes in the ensemble mean, particularly around the interface of tidal and fluvial flows. An experiment to compare the ETKF to the EnKF estimates of forecast error variance when the ensemble mean remained approximately constant was conducted. In this experiment, synthetic measurements were generated from the ensemble mean rather than measurement data. Errors in the synthetic measurements were assumed to be independent of each other and normally distributed around the mean, with a variance equal to that used by the ETKF. Except for the synthetic measurement ensembles, steps 1 to 5 of the experimental setup shown above were conducted to give the following results.

Fig. 5.10 plots the ETKF (red) and EnKF (black) predictions of error variance at the validation time and seven target times at 20 minute intervals from the routine measurement time, against the background error variance at those times. The ETKF predictions of error variance were identical to those shown in Fig. 5.6 because the background forecasts were identical for the two experiments. When measurements were generated from the ensemble mean the EnKF made different predictions of error variance to the real measurement case. There were fewer black dots above the green line, indicating that the updating process did not increase error variance as often when the mean remained constant. In contrast to the real measurement example the ETKF tended to slightly over predict error variance at target times one to five when the background variance was below 0.05. Correlations between EnKF and ETKF predictions of forecast error variance (Table 5.4, correlation A) ranged from 0.858 at the first target time to 0.437 at the validation time, with the largest fall occurring between target times two and four. Correlations A were greater than those in Table 5.3 where real measurements were used. The correlations between EnKF and ETKF estimates of signal variance were 0.670 at the validation time compared with 0.445 when the EnKF ensemble mean was allowed to change.

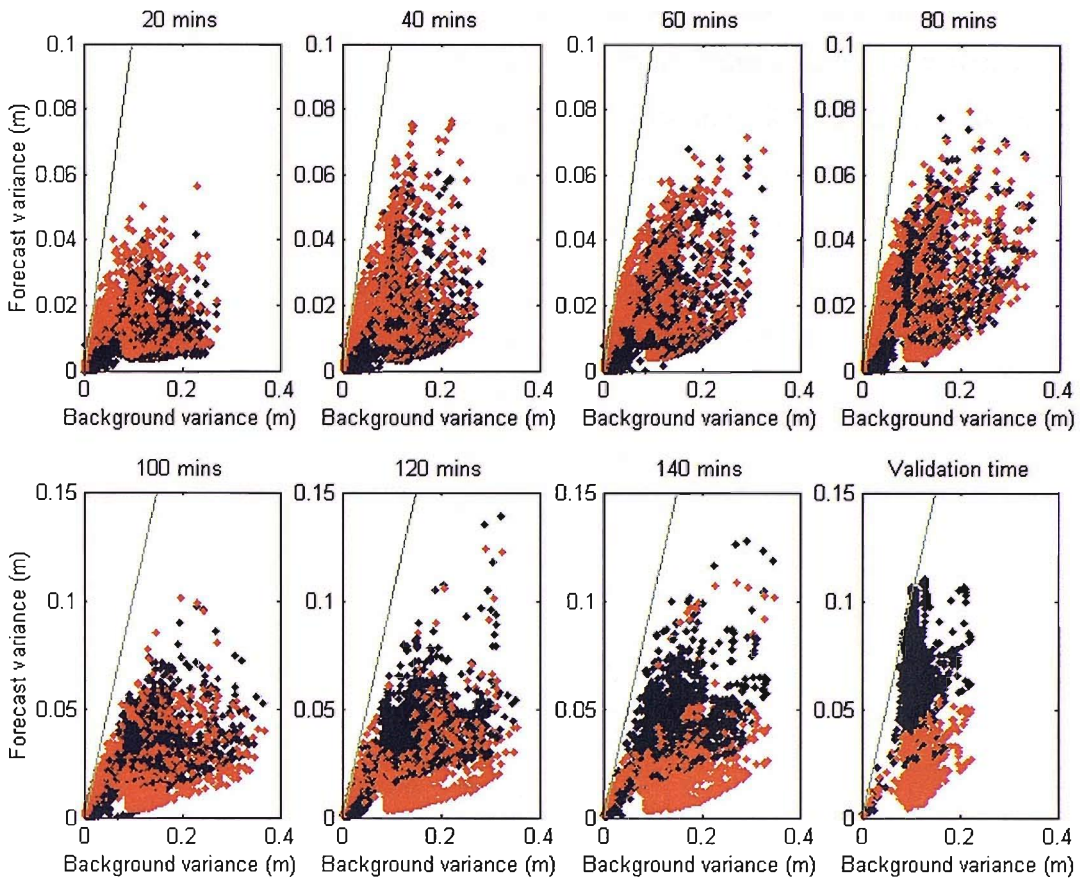


Fig. 5.10: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times after assimilating the routine measurement against background error variance. Measurements were simulated from the ensemble mean.

Table 5.4: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at seven times between a measurement time t_m and a validation time t_v . From a background forecast an EnKF was used to assimilate one measurement. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.

Time from t_r	20	40	60	80	100	120	140	220
Time to t_v	200	180	160	140	120	100	80	0
Correlation A	0.858	0.887	0.787	0.686	0.612	0.620	0.638	0.431
p_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Correlation B	0.005	0.004	0.036	0.051	0.053	0.462	0.832	1
p_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-
Correlation C	0.998	0.994	0.987	0.975	0.964	0.963	0.956	0.670
P_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Fig. 5.11 is a box plot of the forecast error variance from Fig. 5.10. The ETKF under predicted error variance and over predicted signal variance (Fig. 5.12), especially as the time lag from the routine measurement increases (this was as also seen when using real measurements in Fig. 5.7 and 5.8). However, the two clusters of outliers at target time seven (En 7) and the validation time (En v) in Fig. 5.7 and 5.8, which occurred when the simulation model became unstable, were not outliers in Fig. 5.11 or 5.12.




Fig. 5.11: Forecast error variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Measurements were simulated from the ensemble mean. Boxes show the upper quartile, median and lower quartile of forecast error variance; whiskers extend over the range of the data with crosses indicating outliers.




Fig. 5.12: Signal variance at seven target times and the validation time after the assimilation of the routine measurement. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Measurements were simulated from the ensemble mean. Boxes show the upper quartile, median and lower quartile of forecast signal variance; whiskers extend over the range of the data with crosses indicating outliers.

The outliers in Fig. 5.7 and Fig. 5.8 caused by model instability were not evident when the ensemble mean remained constant, suggesting that changing the ensemble mean was a potential cause of instability in RC model simulations. Whether or not the model instabilities apparent in Fig. 5.7 and Fig. 5.8 are linked to large changes in the ensemble mean can be crudely inferred by comparing the signal variance at the validation time with the change in ensemble mean at the routine measurement time. This is shown by the dots on Fig. 5.13, where mean signal variance was negative after assimilating the routine measurement during four tides, all of which coincided with falls in the ensemble mean greater than 0.3 m. Furthermore, the greatest decrease in ensemble mean (0.85 m) corresponded with the greatest increase in mean error variance (0.22 m). The same affect was not observed for increases in the

ensemble mean although these changes were of smaller magnitude than those which caused an increase in error variance. Bearing in mind that a decrease in ensemble mean should result in a smaller mean variance over the whole model domain as the tide advances up the river more slowly than expected, the increase in variance is likely to be due to model instability. Generating synthetic measurements from the ensemble mean always resulted in a positive signal variance, although there was one outlier close to zero and three cases where the ensemble mean increased by up to 0.2 m.




Fig. 5.13: Plot of the signal variance at the validation time against change in ensemble mean (after assimilating the routine measurement). There are 39 dots in total each of which is a separate tidal event. Crosses indicate the corresponding results where the measurement was generated from the ensemble mean.

After assimilating the routine measurement with the EnKF the RC model generated a new background forecast, from the analysed initial conditions $\mathbf{X}(t_r | \mathbf{H}_r)$. The EnKF and ETKF were both used to assimilate a measurement at the first target time from the same location as the routine measurement. Tides 13 and 40 were removed from

the analysis because the RC model had become unstable during these events. Comparisons of the EnKF and ETKF predictions of error variance at the validation time and seven target times are shown in Fig. 5.14, whilst the corresponding results using the ensemble mean as the measurement are shown in Fig. 5.15.

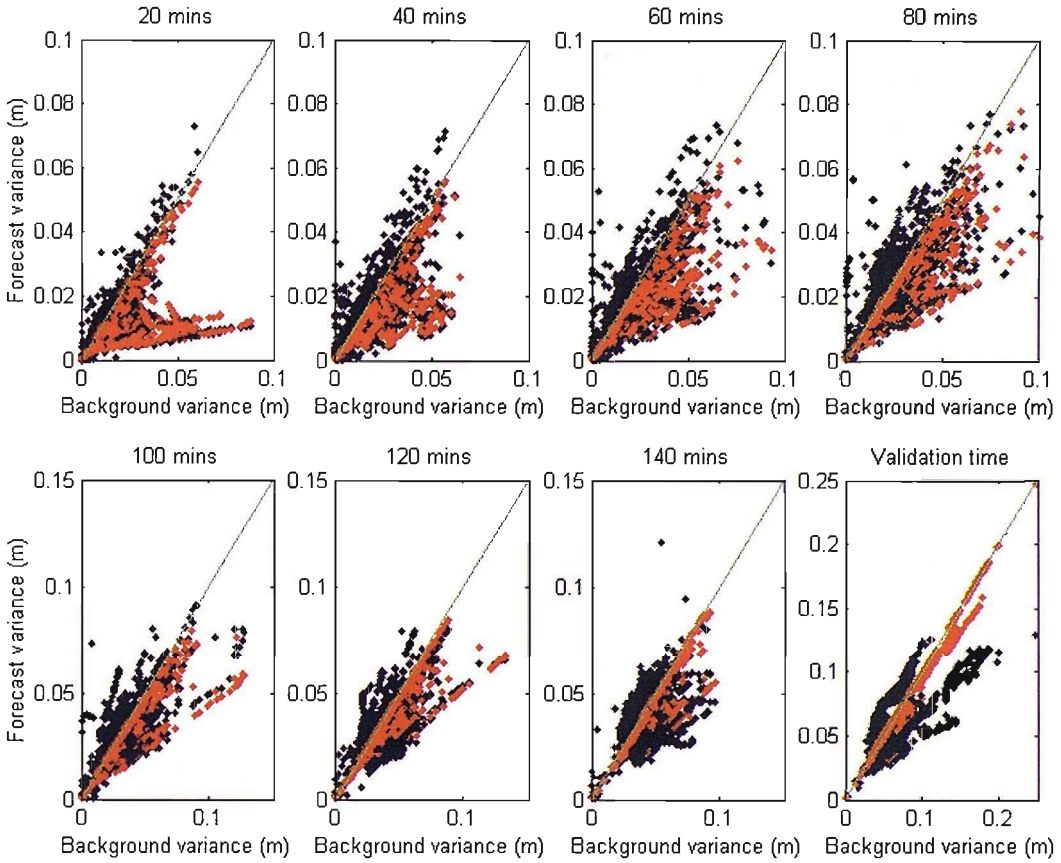


Fig. 5.14: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times from the assimilation of a measurement at the first target time after assimilating the routine measurement.

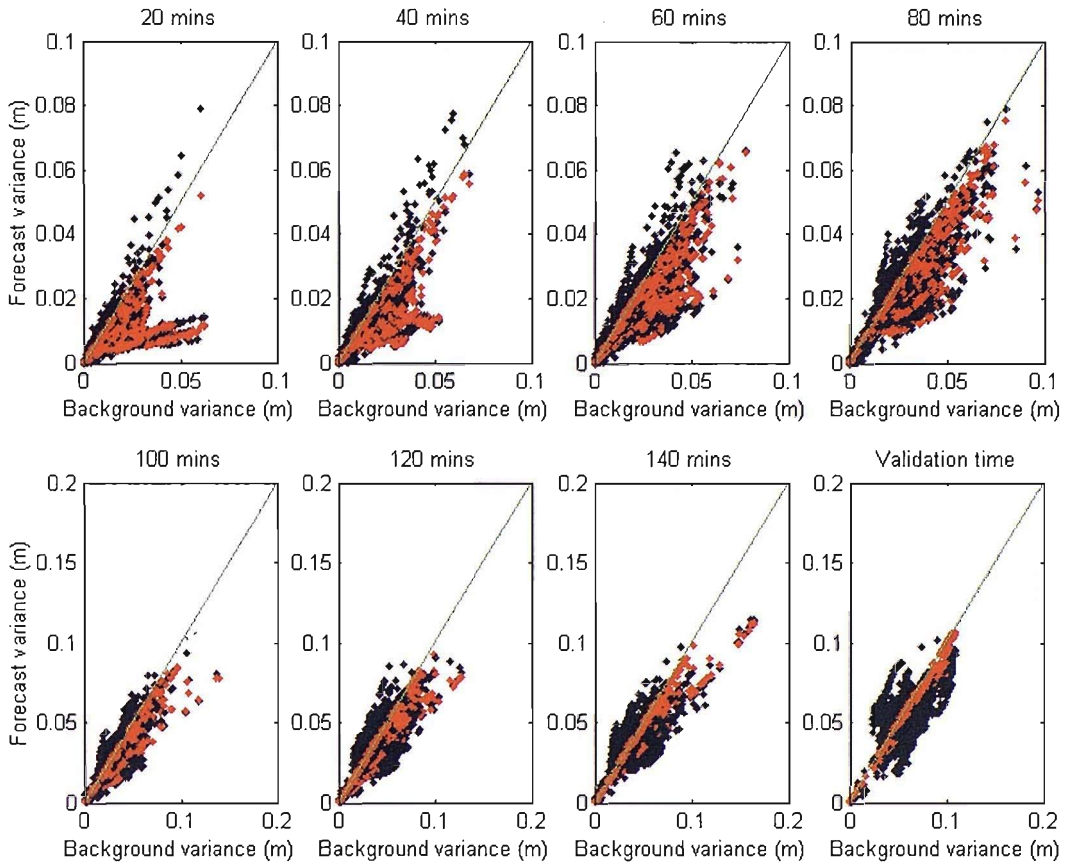


Fig 5.15: ETKF predictions of error variance (red) and EnKF predictions of error variance (black) at varying times after assimilating a measurement at the first target time against the background error variance after the routine measurement. Measurements were simulated from the ensemble mean.

Increases in error variance for small values of background variance, attributed to changes in the ensemble mean when assimilating the routine measurement, were not evident in Fig. 5.14. The scatter plots of forecast error variance in Fig. 5.14 are similar to those obtained when using the ensemble mean as the measurement (Fig. 5.15). The box plots of forecast error variance supported this, as the median and spread of error variance that resulted from assimilating the actual measurement was similar to that of the synthetic measurement (Fig. 5.16 and 5.17 respectively). These results indicate that the ensemble mean changed less when assimilating the measurement at the first target time relative to the routine measurement time. This was expected due to the short time between the measurements and reduction in forecast variance at the target time as a result of the routine measurement.




Fig 5.16: Forecast error variance at seven target times and the validation time after the assimilation of a measurement at the first target time and routine measurement time. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Boxes show the upper quartile, median and lower quartile of forecast error variance; whiskers extend over the range of the data with crosses indicating outliers.




Fig 5.17: Forecast error variance at seven target times and the validation time after the assimilation of a measurement at the first target time. Results are shown using the EnKF approach (En) and the ETKF approach (ET). Target times are at 20 minute intervals from the routine measurement. Measurements were simulated from the ensemble mean. Boxes show the upper quartile, median and lower quartile of forecast error variance; whiskers extend over the range of the data with crosses indicating outliers.

Fig. 5.14 shows that the ETKF tended to under predict error variance when compared to the EnKF based approach, except at the validation time. Unlike when the routine measurement was assimilated, the under prediction was greater at lead times closer to the measurement time. This was supported by the box plots of forecast error variance in Fig 5.16, where the EnKF approach resulted in a convex shape to the growth in median error variance over time, whereas the ETKF predicted a concave shape in median error variance growth between target times two and seven. Therefore, over the 37 tidal cycles shown here; errors initially grew quicker in the EnKF based approach than the ETKF approach before becoming more alike. This result suggests that the growth in error is slightly different after data assimilation or that the

assumptions made by the ETKF are not ideal. Model instabilities caused by large changes in the ensemble mean relative to the ensemble spread have previously been a cause of error growth in this model. These cannot be ruled out as a cause of the convex shape of the error growth when using the EnKF approach.

Despite the above problems the ETKF was able to make predictions of error variance and signal variance that were significantly correlated with the EnKF (Table 5.5). Forecast error variance was reduced by around one order of magnitude due to the assimilation of the routine measurement, with the magnitude of the fall depending on the time lag from the routine measurement. The lower error variance relative to measurement variance of this background forecast made it easier for the ETKF to predict error variance relative to before the routine measurement was assimilated. Hence, the correlations between ETKF and EnKF predictions of error variance tended to be greater in Table 5.5 than Table 5.3. Conversely, the correlations between signal variances were not as large in Table 5.5.

Table 5.5: This table contains the correlations between ETKF and EnKF estimates of forecast error variance at seven times between a measurement time t_r and a validation time t_v . The background forecasts were produced assimilating one measurement with the EnKF at t_r , then running an ensemble of RC model simulations to t_v . From the background forecast an EnKF was used to assimilate one measurement at the first target time. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. Correlation A lists the correlation between EnKF and ETKF error variance estimates. Correlation B lists the correlation between the background forecast error variances at each target time and the validation time. Correlation C shows the correlation between signal variance estimates.

Time from t_m	20	40	60	80	100	120	140	200
Time to t_v	180	160	140	120	100	80	60	0
Correlation A	0.929	0.896	0.856	0.805	0.792	0.739	0.565	0.682
p_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Correlation B	-0.019	0.033	0.103	0.164	0.198	0.269	0.428	1
p_value	0.138	0.011	0.000	0.000	0.000	0.000	0.000	-
Correlation C	0.968	0.892	0.798	0.712	0.602	0.555	0.379	0.211
p_value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

The test carried out here was limited in that measurements from only one location were assimilated. A more extensive test would include measurements from locations in the model domain which were not strongly correlated with each other. The close proximity of the sensors to one another in this case study meant that the measurements from them were highly correlated, which made it difficult to run data sampling scenarios which were spatially more complex with real data. It is possible to simulate measurements at other locations in the RC model domain and run them through the ETKF. However, error variance predictions by the RC model were sensitive to changes in the ensemble mean, especially at the flood edge. Therefore, it was not possible to assess how well the ETKF would approximate the behaviour of the EnKF with measurements from other locations because it is likely to depend on

the magnitude of the change in the ensemble mean. Nevertheless, the ETKF and EnKF can be compared against one another when given data from the floodplain sensor nodes as well as the channel sensor nodes.

5.3.3 Forecast variances: multiple sensor example

This experiment used sensors two, four and six, each of which was around 700 m apart. A routine measurement from sensor two was assimilated at the routine measurement time by the EnKF to give an updated set of initial conditions. From these initial conditions a background forecast with error variance identical to that in Fig. 5.6 was created. As in the previous section events 13 and 40 were discarded due to RC model instability. The two sensors on the floodplain were at higher elevations than sensor two in the channel, meaning that data were unavailable at these locations until the water levels were higher. Therefore, this experiment assimilated data from target time five, one hundred minutes after the routine measurement time rather than target time one as used previously. Despite choosing a target time closer to high tide measurements from sensor six were not always available hence causing the measurement model to vary between using two and three measurements.

The top three plots in Fig. 5.18 show forecast error variance estimates at the validation time, and target times six and seven obtained using the EnKF approach to assimilate measurements at target time five. The lower three plots show forecast error variance using the ETKF approach to assimilate the same data. Both approaches predict that error variance will increase between target time five and the validation time. Complementary to the single sensor example, after the routine measurement the ETKF and EnKF error variances were similar. However, the ETKF predicted a greater range of variances than the EnKF at the validation time. The correlation between estimates of error variance was small at target times six (0.509) and seven (0.307) but increased to 0.516 at the validation time (Table 5.6). Estimates of signal variance had correlations between 0.811 and 0.318 at target time six and the validation time, respectively (Table 5.6). The most obvious difference between the EnKF and ETKF scatter plots of forecast error variance (Fig. 5.18) was the greater scatter of the EnKF error variances at background error variances below 0.1 m. As

shown by Fig. 5.19, this scatter was less evident when measurements were generated from the ensemble mean, an observation supported by the correlation of 0.979 and 0.660 at target time six and the validation time respectively (Table 5.7). The model instability at target time seven during one event, seen as a cluster of points above the continuous line reduced the correlation at this target time to 0.554.

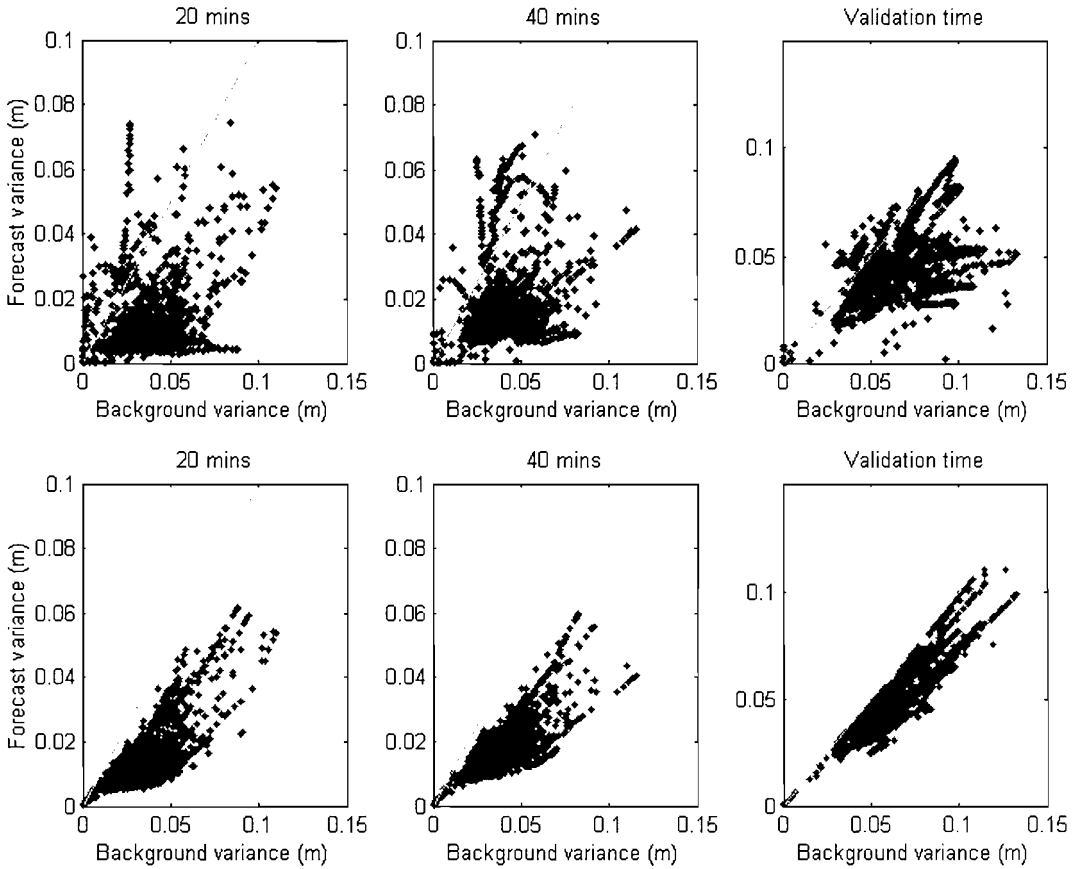


Fig 5.18: ETKF predictions of error variance (lower plots) and EnKF predictions of error variance (upper plots) at three times after assimilating two or three measurements at target time 6 against the background error variance after the routine measurement.

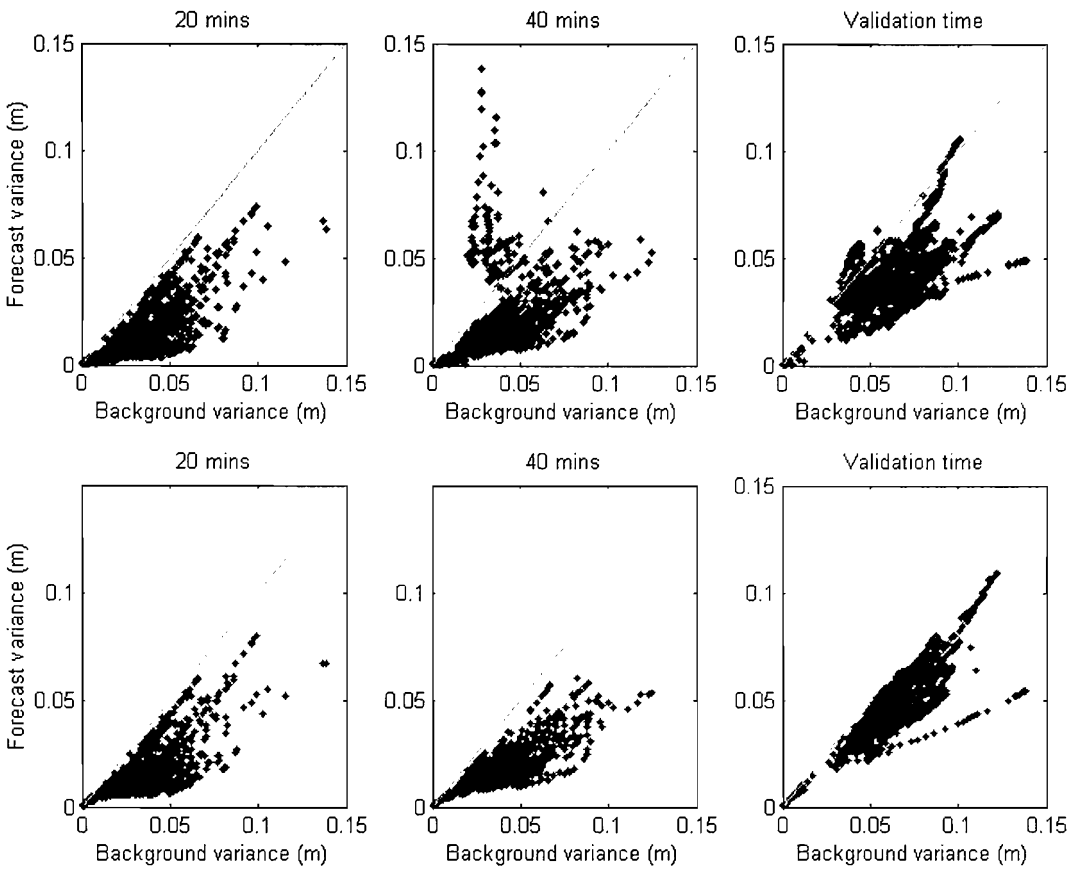


Fig 5.19: ETKF predictions of error variance (lower plots) and EnKF predictions of error variance (upper plots) at three times after assimilating two or three measurements at target time 6 against the background error variance after the routine measurement.

Table 5.6: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at two times between a measurement time t_m and a validation time t_v . From a background forecast an EnKF was used to assimilate two or three measurements. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error covariance to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.

Time from t_m	20	40	100
Time to t_v	100	80	0
Correlation A	0.509	0.307	0.565
p_value	0.001	0.000	0.000
Correlation B	0.089	0.229	1
p_value	0.000	0.000	-
Correlation C	0.811	0.710	0.318
p_value	0.000	0.000	0.000

Table 5.7: This table contains the correlations between ETKF and EnKF (Correlation A) estimates of forecast error variance at two times between a measurement time t_m and a validation time t_v . From a background forecast an EnKF was used to assimilate two or three measurements. The forecast error variance was then calculated by propagating an ensemble of RC model simulations to the validation time. The ETKF used the same background forecast and the same measurement error covariance to calculate an ensemble transform. This transform was then applied to the background forecast at the selected times up to the validation time. The correlation between background error variance at the target times and the validation time are shown as correlation B. Correlations between the ETKF and EnKF estimates of signal variance are shown as correlation C.

Time from t_m	20	40	100
Time to t_v	100	80	0
Correlation A	0.979	0.554	0.660
p_value	0.000	0.000	0.000
Correlation B	0.035	0.098	1
p_value	0.009	0.000	-
Correlation C	0.988	0.717	0.489
p_value	0.000	0.000	0.000

5.4 Discussion

The ETKF allowed model state error covariance estimates to be obtained rapidly from an ensemble of background forecasts at both the time of the observation and future validation times. This approach made it computationally feasible to test many observation sets or sequences because the ensemble of model states needs to be propagated only twice: once to issue the background forecast and then again after the data collected have been assimilated to issue a forecast. Forecast error variance was shown to change when the ensemble mean changed as a result of assimilating actual measurements particularly at the flood edge, as this was the most sensitive region of the model domain to changes in the ensemble mean. This is not ideal given that the

location of the shoreline is likely to be a primary concern to the user, along with the flood depth duration and arrival time at validation locations.

The ETKF was dependent on the RC process model providing state error statistics. Like most model updating algorithms including the EnKF used here, the ETKF variance forecasts were strongly coupled to the simulation model. In this paper, we have not compared the forecasted errors with observed errors but have seen evidence of how instabilities in the RC model have an adverse effect of the correlation between error variance estimates made by the ETKF and EnKF based approaches. This makes it difficult to assess the significance of the assumptions made by the ETKF such as linear error propagation, over the EnKF based approach.

The introduction of unsuitable initial stage and flow conditions into the RC model by the EnKF was the principal factor that caused RC model instability. Incidentally, there was no evidence of model instability when stopping and restarting simulations with unaltered states. It could be argued that the weak point in the comparison made here was the stability of the RC model when given updated initial and boundary conditions, implying that the ETKF would be a more accurate predictor of signal variance had the simulation model on which the two data assimilation methods rely been more robust. Given the instability of the forecasting model used here constraints need to be developed on the updating process to ensure the stability of model simulations after the update. Especially, as Fig. 5.13 showed evidence that the likelihood of model failure increased with the difference between simulated and measured stage.

As a means of targeting data the ETKF has several properties which make it a useful tool for targeting measurements. It provides a rigorous framework for assessing the contribution of measurement data to forecast error variance based on the Kalman filter, hence it is able to incorporate uncertainty from both the measurement and simulation models. Furthermore, as the error statistics can be derived from distributed process models it is possible to consider measurements from previously unmonitored locations based within the model domain. The link to the process model should in theory allow the extrapolation of results to previously unmonitored

conditions, such as those occurring during extreme events, with some degree of confidence.

The instability of the RC simulation model when updated would be a significant barrier to its use in an operational context. However, the information on spatio-temporal signal variance generated by the ETKF such as that in Fig. 5.9 may be useful when planning the deployment of a sensor network like the one used here. Although the RC model was calibrated using data collected by the sensor network, the ETKF could be applied to simulation models of previously unmonitored reaches or to climate based what-if scenarios.

The ETKF method can potentially provide a platform for assimilating data from a number of diverse sources; not just the point measurements of stage used here. For example, targeting techniques similar to the one developed here may be applicable when collecting data on flood inundation extent, either through remote (Horritt, 2001) or ground based survey (Nicholas & Mitchell, 2003). Ground based surveys, in particular, are generally conducted on an opportunistic basis where a method for estimating the likely benefit of the sampling effort could be useful, not least because a lot of effort is required to collect even a small data set. The collection of flood inundation images could be timed to coincide with time periods where the process model was most sensitive to inundation extent. The ETKF is relatively easy to implement in a situation where an ensemble model is present since the model statistics necessary to run the algorithm can all be generated from the state ensembles. Therefore, in circumstances where ensemble simulation models have been developed previously much of the work necessary to apply the ETKF has been done already.

This study has not attempted to compare the performance of the ETKF against any other method of adaptive sampling. Principally this was because the stability of the RC model was a barrier to a more detailed analysis of the efficiency and accuracy of the ETKF approach. Thus, other adaptive sampling methods may prove to be more appropriate in this context. Another, possible avenue of research is that the models

used to process data may also be adapted where misfits occur between model estimates and data (Lermusiaux *et al.*, 2004).

5.5 Conclusions

The RC model was used to simulate 60 background stage predictions for 40 consecutive tides. Errors were introduced to each simulation by varying the roughness parameter of the model and introducing coloured noise at the model boundary. The first tide was used as a warm-up period for the model and discarded from further analysis.

An EnKF assimilated measurements of stage for the purpose of updating RC model state variables (stage and flow) and boundary conditions. These updated state variables and boundary conditions were used as initial conditions for further simulations (or forecasts) of state variables, principally stage. The ability of the ETKF to estimate forecast error variance and signal variance *a posteriori* of these measurements, before they were assimilated was assessed by comparison with simulations by the RC model which were updated with an EnKF. At the measurement time signal and error variance estimates by both methods were similar.

Two experiments were conducted. The first used a series of measurements from a single location and the second measurements from 2-3 locations. For both experiments estimates of forecast error covariance from the two filters were significantly correlated. Correlations ranged between 0.979 and 0.292. Correlations tended to be lower as time after the measurement time increased. The ETKF underestimated forecast error variance and conversely overestimated signal variance relative to the EnKF approach, when assimilating the routine measurement. Correlations between the variance estimates of the two approaches were reduced during some tides due to simulation instability, the likelihood of which increased the greater the change in ensemble mean.

Error variance was most sensitive to changes in the ensemble mean at the boundary between tidal and fluvial flows and in shallow regions towards the edge of the flood envelope on the salting. Therefore, when the EnKF changed the mean of subsequent simulated ensembles the ETKF was a less accurate predictor of error variance at these locations, relative to locations that had been inundated for some time.

The ETKF was a useful tool for quickly estimating the effect on state variance of assimilating measurement data into the hydrodynamic model used here. It, thus, provides a means of quantifying the ‘usefulness’ (in terms of error or signal variance) of possible sampling schemes. Since the Kalman filter provides a framework for considering errors resulting from uncertainty in both the simulation model and measurement data, it seems natural to base any adaptive sampling framework on similar principles, not least because measurements could be collected by different sensor platforms with different error statistics.

5.6 Summary

This chapter compares two Monte Carlo sequential data assimilation methods based on the Kalman filter, for estimating the effect of measurements on simulations of state error variance made by a one-dimensional hydrodynamic model. The first method used an ensemble Kalman filter to update state estimates, which were then used as initial conditions for further simulations. The second method used an ensemble transform Kalman filter to quickly estimate the effect of measurement error covariance on forecast error covariance without the need to re-run the simulation model. Estimates of forecast error covariance from the two filters were significantly correlated, with correlations ranging between 0.979 and 0.292. The motivation behind the study was to assess the ability of the ETKF to target possible measurements, as part of an adaptive sampling framework, before they are assimilated by an EnKF based forecasting model on the River Crouch, Essex, UK. The ETKF was found to be a useful tool for quickly estimating the error covariance expected after assimilating measurements into the hydrodynamic model. It, thus, provided a means of quantifying the ‘usefulness’ (in terms of error variance) of possible sampling schemes. Since the Kalman filter provides a framework for considering errors resulting from uncertainty in both the simulation model and measurement data, it seems natural to base any adaptive sampling framework on similar principles, not least because potential measurements could be collected by multiple sensor platforms with different error statistics.

6 Adaptive space-time sampling with wireless sensor nodes for flood forecasting

6.1 Introduction

Several studies such as Sprokkereef (2001), Madsen and Skotner (2005) and chapter 4 of this thesis have reported increases in forecast accuracy when implementing state updating techniques on hydrodynamic models with stage measurements. This chapter is concerned with the supply of measurement data to a state updating forecasting model based on an EnKF and a hydrodynamic simulation model of the River Crouch, UK (chapter 4). Specifically, adaptive sampling techniques that facilitate a reduction in the need to transmit real-time data from a network of sensor nodes in the field to the forecasting model are developed and evaluated. The motivation behind this aim and some objectives are outlined in section 6.1.1, followed in section 6.2 by a description of the study site, sensor nodes, measurement data and forecasting model. Section 6.3 proposes two adaptive sampling methods and tests them with the forecasting system outlined in section 6.4, followed by a discussion and some conclusions.

6.1.1 The FloodNet approach

The use of *in situ* or ground-based measurements to provide inputs to flood forecasting systems has become more feasible with the advent of simpler and cheaper sensors, gauges, and loggers (Troch *et al.*, 2003). The FloodNet project, which advocated the movement of computational power and communications capabilities onto networks of sensors in the environment through the concept of pervasive or ubiquitous computing, has initiated opportunities for the delivery of ground-based data in real-time and the development of adaptive monitoring systems. The FloodNet network is adaptive in the sense that individual sensors on the network are able to change their behaviour depending on the prevailing environmental and

infrastructural circumstances, whilst considering the status of other sensors on the network as well as data requests from external sources (DeRoure, 2005).

The FloodNet sensor network presents an opportunity to develop a monitoring system (Fig. 6.1) where the collection and communication of real-time data is influenced by requests from an offsite flood forecasting model, although simulations from a local model on the sensor nodes could in theory also be used by such a network instead or in addition to the offsite forecasting model (Hughes *et al.*, 2006). One objective of the approach adopted in the project was to reduce the power drain on the sensor nodes by communicating data in real-time only when required by a forecasting model, hence reducing the need to generate or store as much power at each individual sensor, extending the lifespan of sensors for a given power resource and reducing maintenance requirements.

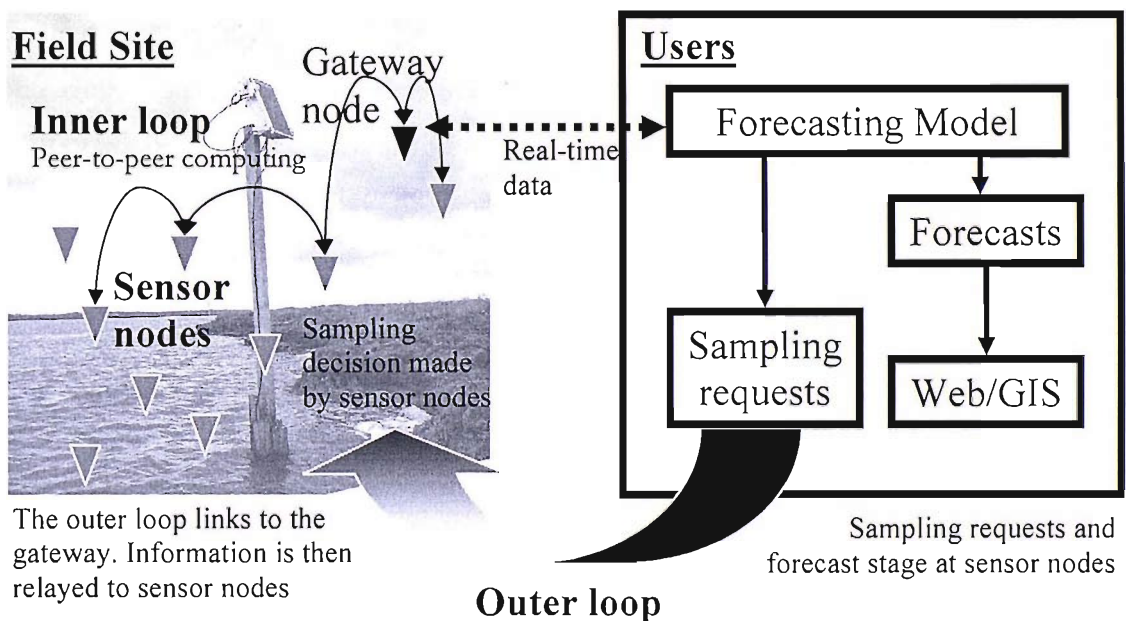


Fig. 6.1: Conceptual diagram of the two FloodNet control loops. An inner loop comprising of a peer-to-peer wireless network and an outer loop moving measurements from the sensor network to a database and forecasting model, and then returning sampling requests to the sensor network.

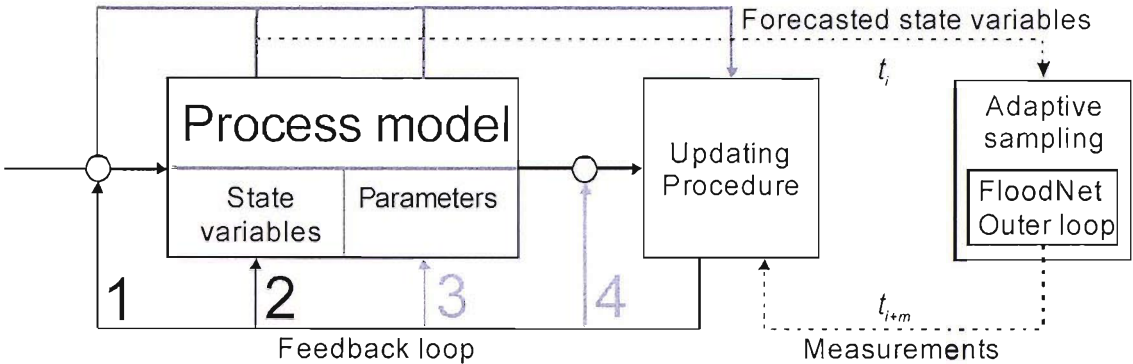
The functionality of the FloodNet system was based on two control loops (Fig. 6.1). The ‘inner loop’ operated on the wireless sensor nodes themselves and could be autonomous of any off-site intervention as required. Each sensor node on the

network had a set sampling rate, requiring samples to be taken at 5 minute intervals. However, an individual node was able to decide its reporting rate (regularity of data transmission) based on a 'demand value'. This could be derived locally based on a local metric such as the residual error between expected and observed water level or it could be prescribed by an external source. The adaptive behaviour was enabled through the FloodNet Adaptive Routing protocol (FAR) (Zhou and De Roure, 2007), which sets out rules for the operation of the network. When battery power was limited the protocol included this factor when deciding the reporting rate and the route data take through the network to the gateway node. The gateway node being a special node linked to a General Packet Radio Service (GPRS) modem that relays data to an off-site database. The 'outer loop' facilitated feedback from external users of the data by allowing 'demand values' to be prescribed for each sensor on the sensor network.

In the present application, the forecasting model was a principal component of the outer loop. Therefore, data leaving the study site were archived and then assimilated with predictions made by the simulation model. Hence, demand values returned to the 'inner loop' should reflect the data requirements of the forecasting model. The novelty of this approach lies in the feedback of information from the forecasting model to the sensor nodes in the field as the driver of the network's adaptive behaviour.

So far the topology of the FloodNet system has been considered from the point of view of the sensor network. However, it is also useful to consider the operation of the FloodNet system as an extension of a more general model updating problem. Using the schematic diagram from Refsgaard (1997) and WMO (1992) as a base to describe the process of forecasting and model updating, the effect of adaptive sampling on model updating is described in Fig. 6.2. This illustrates how outputs from the process model, which may have been updated previously, are the basis for determining demand. When transferred to the inner loop these demand values influence the future measurements that are provided to the updating procedure, completing the cycle. Two methods for generating demand values based on ensemble predictions are

discussed later in this paper, after a description of the study site, process model and updating procedure on which they depend.



Forecasting and simulation after Refgaard (1997) and WMO (1992)

Fig. 6.2: Schematic diagram of four different updating or data assimilation methods: 1. Input updating. 2. State updating. 3. Parameter updating. 4. Error prediction. Also shown is the adaptive sampling concept, where grey boxes define FloodNet specific control loops.

6.2 Study site and model setup

The field site used in this study lies along the River Crouch in Essex, UK. The river flows for approximately 30 km west to east from Basildon (BNG: TQ 730 895) into the North Sea at BNG: TR 037 961 (Fig. 6.3). The underlying geology of the region is that of soft sedimentary rock laid down during the Tertiary Period, overlain by deposits from Pleistocene glacial periods. The catchment land cover is predominantly agricultural with several large towns including Basildon, Wickford, South Woodham Ferrers and Burnham-on-Crouch, along with numerous small towns and villages. The Crouch basin covers an area of 71.8 km² with an average annual rainfall of 572 mm per year (Environment Agency). It is flat and low lying to the extent that the maximum height above datum of the catchment is 118 m and tidal conditions extend some 20 km inland, two thirds of the total length of the river. Non-tidal flows on the upper third of the river are monitored by a gauge at Wickford (BNG: TQ 748 934). The contribution of flow to the estuary from the upstream non-tidal river is small

relative to the tidal flows. The lower two thirds of the river are characterized by estuarine silt and mud of less than 0.2 mm. Where embankments are not present the river is flanked by areas of tidal salting. The tidal range at Burnham-on-Crouch varies between 3.2 m at neap tide and 5.0 m at spring tide.

The field site itself was situated in the Brandy Hole area of the river. It consisted of a 2.5 km stretch of estuary and approximately 1 km² of tidal salting that sits behind a derelict sea wall (Fig. 6.3). At its mouth the river is approximately 700 m wide and up to 14 m deep at high tide. By the eastern (downstream) extent of the field site the channel width decreases to around 350 m at high tide with a depth of up to 10 m. At the western (upstream) extent of the field site the channel width is up to 190 m with a depth of up to 5 m. The tidal salting is characterised by numerous small channels typically with steep banks around one or two metres high. Some of these channels were cut artificially when the sea wall was first breached, whilst others have formed over time. The presence of artificial channels leads to the striped drainage pattern that can be seen on some areas of the salting particularly around sensor 5. The eastern and southern limits of the salting are set by an embankment. The river flows across the northern side of the salting whilst its western extent is limited by a gentle slope towards higher ground that makes the salting edge difficult to delineate. At low tide the salting was observed to drain almost completely whilst becoming fully inundated on a spring tide.

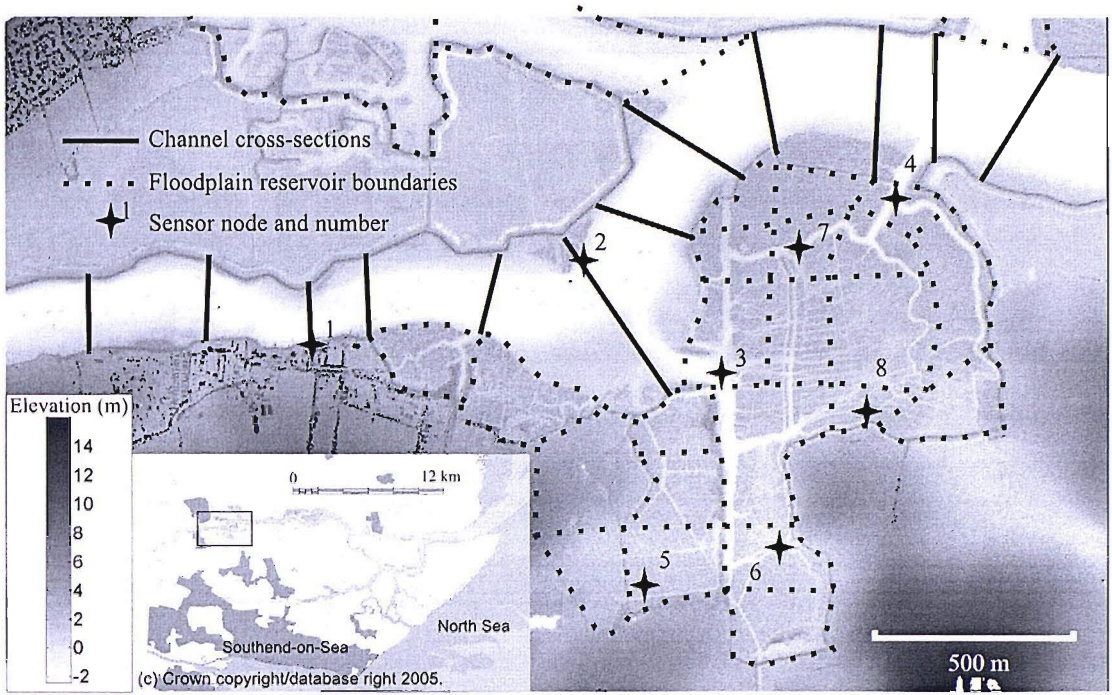


Fig. 6.3: Location of sensor nodes and RC model topology.

6.2.1 Sensor node hardware

The study site was monitored using eight Druck 1830 series pressure transducers (GE Druck) measuring water level in the channel and on the floodplain. Each pressure transducer formed part of a FloodNet sensor node, which included a BitsyX single board computer (Applied Data Systems), an Intel PXA255 RISC microprocessor to provide field processing capabilities and a solar panel as a power source. The sensor nodes were linked via a wireless IEEE 802.11b computer network to a GPRS modem on a gateway node, enabling real-time data transmission to the user from the site. The nodes were able to communicate over a range of around 600 to 800 m. A combination of this and the presence of obstructions such as sea walls, trees and buildings at the Brandy Hole field site, meant that most nodes did not have a direct link with the gateway node. Consequently, data from sensors on the periphery of the network had to route their data via other nodes to reach the gateway.

Taking measurements and storing them locally on a sensor requires very little power (around 60 mW) relative to activating the single board computer and transmitting data (around 1200 mW). Power can be saved by collecting and sending data less often. A FloodNet sensor can store up to 12 measurements (a measurement includes

the sensor measurement, a measurement variance estimate, the battery power and the time) before it needs to transmit the measurements, overwrite previous measurements or move the measurement data onto the single board computer. This functionality allows a short archive of data to be collected regardless of the need for the data in real-time. It also allows nodes that become isolated for any reason to continue collecting data until they re-establish connectivity with the gateway or have their data downloaded manually. Therefore, it is power efficient to send data in larger batches rather than as a continuous stream. More information on the FloodNet sensor nodes can be found in Kaun (2005).

6.2.2 Measurement data

The archive of data used in this paper was collected between 4th and 26th November 2005 by eight sensors. Selected attributes of the eight sensors deployed at the field site are summarised in Table 6.1. No sensor was inundated all of the time. Sensor 2 was located at the lowest elevation and collected the most complete time series of data typically collecting measurements in excess of 4 hours prior to high tide. The other sensors were at higher locations and collected less data. A measurement is defined as the mean of 32 repeat samples from which an estimate of measurement variance was also made. These variance values could have been used, as demonstrated in chapter 4. However, to simplify the implementation the average of these variance values (0.005 m) was taken and applied to all measurements.

Table 6.1: Sensor nodes and measurement data.

Sensor	Elevation (m)	No measurements	Easting (m)	Northing (m)	Location
1	0.634 m	2439	582311.7	195783.1	Channel
2	-0.917 m	4143	582980.9	195958.4	Channel
3	-0.278 m	3424	583356.3	195660.0	Floodplain*
4	0.397 m	2688	583861.1	196189.6	Floodplain*
5	1.340 m	800**	583142.3	195096.2	Floodplain
6	1.518 m	1345	583559.8	195270.1	Floodplain
7	1.244 m	1781	583554.7	196016.5	Floodplain
8	1.820 m	1000	583859.8	195531.7	Floodplain

* Sensor node was <100m from channel.

** Data missing due to sensor node failure

6.2.3 Forecasting model

Measurements from the FloodNet sensor network were assimilated with state simulations from a one dimensional hydrodynamic model (RC model), using an ensemble Kalman filter (EnKF). The RC simulation model was based on the ISIS flow code (HR Wallingford). The model simulates stage and flow over 161 spatially distributed nodes of which measurements were available at up to eight. The model can be described in its state space format as:

$$\mathbf{x}(t, k) = M(\mathbf{x}(t-1, k), \theta(k), u(t-1) + \varepsilon(t-1, k)), \quad (6.1)$$

where \mathbf{x} is the N -by-one model state vector (stage and flow at each node and the downstream boundary condition error), M is the model operator, θ is the global roughness parameter of all model nodes, u is the downstream stage-time boundary condition of the model and ε is an error in the downstream boundary condition generated by a first order autoregressive model. Superscript a refers to ‘analysed’ state estimates made after the assimilation of measurements, whilst f is forecasted state made by the simulation model, but based on previously assimilated measurements. Subscript k indicates that the variable is part of an ensemble, whilst t

is the time step. In this paper, the ensemble size K was always 60. More details on the structure and parameterisation of the RC model were presented in chapter 5.

State was updated by an EnKF with the general form

$$\mathbf{x}^a(k|\mathbf{H}) = \mathbf{x}^f(k) + \mathbf{P}^f \mathbf{H}^\top (\mathbf{H} \mathbf{P}^f \mathbf{H}^\top + \mathbf{R})^{-1} (\mathbf{y}(k) - \mathbf{H} \mathbf{x}^f(k)), \quad (6.2)$$

where \mathbf{y} is a vector of perturbed measurement data, \mathbf{R} is the measurement error covariance and \mathbf{H} is a linear measurement operator which maps the measurement vector \mathbf{y} onto the state vector \mathbf{x} . A detailed description and derivation of the EnKF algorithm can be found in Evensen (2003). The relationship between measurements and the true state of the system can be expressed by the measurement model

$$\mathbf{y}(k) = \mathbf{H} \mathbf{x}^{true}(k) + \boldsymbol{\eta}(k), \quad (6.3)$$

where $\boldsymbol{\eta}$ is a vector of uncorrelated errors sampled randomly from a distribution with zero mean and variance 0.005, which when added to the measurement data is used to generate the measurement ensemble. \mathbf{P}^f is the forecast state error covariance estimated from the forecast state ensemble (Evensen, 1994) such that:

$$\mathbf{P}^f \cong \tilde{\mathbf{P}}^f = \overline{(\mathbf{x}^f(k) - \overline{\mathbf{x}^f}) (\mathbf{x}^f(k) - \overline{\mathbf{x}^f})^\top}, \quad (6.4)$$

where the over-line denotes the ensemble mean. Modelled state error covariance \mathbf{P} was assumed to be independent of the measurement error covariance \mathbf{R} .

6.3 Adaptive sampling methods and results

The RC simulation model produced ensemble state estimates between a time t_i and a validation time t_v at high tide. For the purpose of this experiment t_i was always four hours prior to the validation time. This maximum forecast lead time would ideally be greater. However, the location of the FloodNet sensors limited the lead time on

measurements to four hours. The measurement data from the eight FloodNet sensors were collected at five minute intervals and could have been transmitted at this rate for short periods should this have been necessary. However, two practical considerations prevented this entire data set from being considered. Firstly, each snapshot of RC model state required 16Kb of memory, which, for the 60 model ensembles used here, summed to approximately 1Mb. Archiving these snapshots for the 40 tidal cycles would have required almost 2Gb of memory before any analysis had been conducted. The second reason for not considering more target times was related to the setup of the hydrodynamic model. Put simply, the parameter files that control the operation of the ISIS flow model permit the user to specify only ten state snapshot times during any particular simulation. These snapshot files were needed to provide the state vector of each ensemble member \mathbf{x}_k to the EnKF. One snapshot file was needed at the validation time and one to record the initial conditions. Therefore, it was convenient to have eight target times between the first background forecast and the validation time. Target times were set at 20 minute intervals from t_i , which left the last target time at 1 hour 20 minutes prior to the validation time, as illustrated by the adaptive sampling timings in Fig. 6.4.




Fig. 6.4: Adaptive sampling timings.

Measurements taken by the FloodNet sensor nodes had a single purpose, which was to populate the EnKF measurement model (Eq. 6.3) such that the simulation model state could be updated (Eq. 6.2) and an ensemble forecast issued. The theoretical starting point for the adaptive sampling problem was that more data would give a more accurate representation of current reality and that this more accurate representation of initial reality will then lead to an increase in forecast accuracy. Ideally, there would be no uncertainty in initial conditions; i.e., an infinite number of measurements with no uncertainty had been collected over the model domain. This is

not possible, especially in this case where there are only eight sensors with limited power. Therefore, the adaptive sampling problem was defined as comprising of two competing objectives:

1. To supply measurements such that a sufficiently accurate forecast can be issued as early as possible or before a given time.
2. To make as few real-time data transmissions as possible in order to save power on the sensor node.

Two adaptive sampling scenarios called Method 1 and Method 2 were implemented in this study and are described below. Broadly speaking, the first is designed to complement objective 1 and the second is designed to complement objective 2.

6.3.1 Method 1

In Method 1 data are requested from the FloodNet sensor nodes until the probability of the water level at a validation location exceeding a threshold over a pre-determined forecast period is less than 5%. The ensemble of model simulations was used to calculate the probability p of the water level h exceeding a threshold level a at a validation location j at a validation time t_v , It was assumed that the validation location could be directly mapped onto the model state vector. Therefore, the forecasted probability was conditional on the parameters, initial conditions and boundary conditions of the forecasting model as well as the measurement data used, leading to:

$$p(h(j, t_v) > a(j)) = \frac{1}{K} I(x(k, j, t_v) > a(j)) \quad (6.5)$$

where I represents an indicator function. A default reporting rate of once every 20 minutes was set for each sensor. For this test, sensor 8 (see Fig. 6.3) was selected as the validation location because of its location away from the main channel at a high point on the salting (Table 6.1). The threshold a was set to 3.2 m. After each new measurement set was received the RC model ensemble was updated and simulations

were run to the validation time. The probability of the threshold a being exceeded at a validation location (at sensor 8) was then estimated. If this probability was above 0.05 the reporting rate would remain at the default level. However, if this was not the case the reporting rate could be reset to something less frequent when the next report was sent. In this example, no more reports were made between the decision being made and the validation time.

The threshold a from Eq. 6.5 was set at 3.2 m above ordnance datum (Newlyn) for a validation location at sensor 8. The measured maximum stage at this location for each tide is shown in Fig. 6.5, along with the mean of the background simulations made by the RC model without data assimilation. Only events 21 and 23 were observed to cross the threshold, whilst events 20, 21 and 23 were predicted to exceed the threshold by the background simulations. Observed maximum stage was both greater than and less than the simulated values. None of the tides shown in this series breached the sea defences along the river and, therefore, did not represent conditions where flows and stores of water might occur in areas outside the model domain.




Fig. 6.5: Scatter plot of 39 high tides at the location of sensor 8 as predicted by the RC model without data assimilation (Dots) and measured by sensor 8 (crosses).

The results of adaptive sampling Method 1 will now be summarised. Fig. 6.6 shows the forecasted probability of exceeding threshold a , the number of reports needed to reach the stopping criteria and forecast peak stage as boxplots. For 27 of the 39 tides one report at the first target time was required to meet the stopping criteria. Of the remaining 12 tides five met the stopping criteria after the first target time but before the final target time, whilst seven required all nine target times of which two were observed to exceed the stage threshold. The probability of flooding was one during tide 21 indicating that all the 60 forecast state estimates exceeded threshold a and 0.67 for the other event observed to exceed the threshold (tide 23). The two tides observed to exceed the threshold had the greatest probability estimates. Other events with a significant probability of exceeding the threshold after nine target times were 2, 16, 17, 24 and 25 with probability values of 0.23, 0.52, 0.07, 0.13 and 0.08 respectively. For 22 tides no simulations exceeded the threshold; therefore, returning a probability of zero due to the finite number of ensembles.




Fig 6.6: Results of adaptive sampling Method 1. Box plots show the upper quartile, median and lower quartile of simulated high tide; whiskers extend over the range of the data with crosses indicating outliers. Circles indicate observed stage peaks.

If measurements had been reported at each of the nine target times over the 39 tides there would have been 351 reports from the sensor network. In this example, only 120 were required by the adaptive sampling method, leaving 231 reporting times that were not required. However, as the instruction to stop reporting was retrospective of the measurement being assimilated an additional report (at the next target time) was required for each tide where all nine target times were not used. Therefore, an additional 32 reports were needed, giving a total of 152 reports, or a saving of 199 reports. These results are likely to change markedly for different periods of time and different thresholds.

The boxes at the bottom of Fig. 6.6 show the upper quartile, median and lower quartile of the simulations of peak tide; whiskers extend over the range of the data with crosses indicating outliers. The observed stage peaks are shown as circles. Ideally, observed peak stage should almost always fall within the range of box plot whiskers and occur within the boxes around 50% of the time. However, there were eight tides (numbers 11, 13, 14, 15, 16, 17 and 40) during which the observed stage was below the box plot distribution. This represents a significant number of events where the uncertainty in the forecasts has been under-estimated the bias being induced by the updating procedure or the simulation model becoming unstable. Evidence for the latter two will now be presented in the form of example events.

For tides 11 and 13 there was a phase difference between the background predictions and observed stage. This phase difference manifested itself as a large stage error in the subsequent forecast as shown by event 13 in Fig. 6.7. The subsequent simulations were less accurate because the updating procedure assumes that all observed errors at the boundary condition constitute a magnitude error. The correction to the downstream boundary conditions drives the forecast above observations rather than adjusting the timing of the event. This trait represents a significant limitation of the error model at the downstream boundary of the RC model.




Fig. 6.7: Event 12 showing evidence of a phase difference between the background forecast mean and the observed tide along with major model instability. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event.

Event 13 shows model instability shortly after the update time, which has caused the mean state estimates to fall rapidly, before the driving forces at the model boundary take over to wash out the effect of the instability. Tides 14-17 were all neap tides of a lesser magnitude than expected. It could be argued that the filter was reacting sluggishly to the measured data and, thus, the forecasts were inaccurate. However, the error model at the model boundary is the principal cause of the error rather than the filter. Example tides where the simulation model and observations were in phase and free of major instabilities are shown for a tide similar to the background forecast (Fig. 6.8, event 35), a tide of lower than expected magnitude (Fig. 6.9 event 18) and a tide of greater than expected magnitude (Fig. 6.10 event 5). These forecasts were made from the stopping target times.




Fig. 6.8: Tide 35 where the background simulation was similar to the actual tide. Includes the model forecast at stopping point. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event.




Fig 6.9: Tide 18 where the background simulation was greater in magnitude than the actual tide. Includes model forecast at stopping point. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event.




Fig. 6.10: Tide 5 where the background simulation was less than the actual tide. Includes the model forecast at stopping point and shows evidence of model instability. The dots are observed stage; the bold line is the RC model prediction of stage without data assimilation; the thin lines are RC model forecasts with data assimilation at the stopping time for this event.

6.3.2 Method 2

This method aims to inform the reporting pattern of each sensor using an ensemble transform Kalman filter (ETKF) (Bishop *et al.*, 2001). This method was shown to make reasonable estimates of analysed and forecast error variance under certain conditions for the RC model in chapter 5. Ensembles of RC model simulations were available from 4 hours prior to the validation time t_i to the validation time t_v (high tide). The EnKF was used to assimilate a routine measurement from sensor 2, 20 minutes after t_i . From these updated initial and boundary conditions a ‘background’ forecast made up of 60 RC model simulations was issued. From these background simulations an N -by- K state matrix \mathbf{X} was extracted at five target times equally

spaced at 20 minute intervals between 3 hours 20 minutes and 2 hours prior to the validation time. A state ensemble was also extracted at the validation time (Fig. 6.4 timings plot). Perturbations about the mean state were calculated for each state ensemble member and scaled by the square root of the ensemble size minus one, to give \mathbf{z} :

$$\mathbf{z}(t, k | \mathbf{H}^r) = [\mathbf{x}(t, k | \mathbf{H}^r) - \bar{\mathbf{x}}(t | \mathbf{H}^r)] / \sqrt{K - 1}, \quad (6.6)$$

where t is the time and superscript r indicates that the forecasted state perturbations were dependent on the routine measurement 3 hours 40 minutes prior to the validation time. By repeating Eq. 6.6 for each member of the ensemble the columns of the state perturbation matrix $\mathbf{Z}(t | \mathbf{H}^r)$ can be populated. This was done for each of the target times and the validation time. State error covariance at any of the validation or target times was approximated by:

$$\mathbf{P}(t | \mathbf{H}^r) = \mathbf{Z}(t | \mathbf{H}^r) \mathbf{Z}^T(t | \mathbf{H}^r). \quad (6.7)$$

As shown by Figs 6.3 and 6.4 and Table 6.1 there were five target times and eight sensor locations at which measurements could be taken. Each possible measurement was considered independently giving up to 40 or Q possible measurements. For the q^{th} measurement the signal covariance matrix \mathbf{S} was calculated using the method described by Majumdar *et al.* (2002) whereby:

$$\mathbf{S}(t_{i+M} | \mathbf{H}^q) = (\mathbf{Z}(t_{i+M} | \mathbf{H}^r) \mathbf{C}^q) \mathbf{\Gamma}^q (\mathbf{\Gamma}^q + \mathbf{I})^{-1} (\mathbf{Z}(t_{i+M} | \mathbf{H}^r) \mathbf{C}^q)^T, \quad (6.8)$$

where \mathbf{I} is a K -by- K identity matrix, \mathbf{C}^q is a K -by- K matrix of eigenvectors and $\mathbf{\Gamma}$ is a K -by- K diagonal matrix of eigenvalues of the matrix $\hat{\mathbf{Z}}^T \hat{\mathbf{Z}}$ (Livings, 2005) where:

$$\hat{\mathbf{Z}} = (\mathbf{R}^q)^{-\frac{1}{2}} \mathbf{H}^q \mathbf{Z}(t_{i+M} | \mathbf{H}^r), \quad (6.9)$$

and subscript m indicates the target time of the q^{th} measurement. An approximation of the signal covariance matrix at the validation time given the q^{th} measurements $\mathbf{S}(t_v|\mathbf{H}^q)$ was then provided by substituting the ensemble perturbations at the target time in Eq. 6.8 with those at the validation time (Bishop *et al.*, 2006):

$$\mathbf{S}(t_v|\mathbf{H}^q) = (\mathbf{Z}(t_v|\mathbf{H}^r)\mathbf{C}^q)\mathbf{\Gamma}^q(\mathbf{\Gamma}^q + \mathbf{I})^{-1}(\mathbf{Z}(t_v|\mathbf{H}^r)\mathbf{C}^q)^T. \quad (10)$$

From this matrix the signal variance at the validation time and location was extracted as a measure of the information content of the q^{th} measurement, given the routine measurement. This allowed a table of the expected signal at each sensor location and target time to be built up. This is illustrated for tide 21 by Fig. 6.11, where signal variance at the validation time and location is shown for eight sensor locations and five target times. The signal tended to increase the closer the target time to the validation time, with the exception of sensor 6. The signals obtained from the four sensors nearest the channel (sensors 1, 2, 3 and 4) were similar. Sensor 7 was forecasted to be dry at each target time and was, therefore, expected to yield no signal. Sensor 5 (the furthest sensor from the channel) provided a weaker signal than the other sensors, whilst sensor 6 initially provided a relatively strong signal until target times 4 and 5. The measurement with the greatest signal was that from sensor 2 at target time 5. However, it is worth noting that all the sensors close to the floodplain produced similar signals at this target time. Therefore, it may be equally acceptable to choose one of the other sensors should there be circumstances which make using sensor 2 less appealing. For example, sensor 2 reported the routine measurement, whilst the other sensors were dormant. Should sensors 1, 3 or 9 have more power as a result it may be desirable to use these sensor instead of sensor 2.




Fig. 6.11: Bar chart of signal variance for Q possible measurements spread over eight sensor locations and five target times.

Upon choosing the measurement from sensor 2 at target time 5 the next step was to update the ensemble perturbations at each target time and the validation time such that the next measurement could be chosen. For the chosen measurement location and time a transform matrix \mathbf{T}^q was calculated such that:

$$\mathbf{P}(t_{i+M} | \mathbf{H}^q) = \mathbf{Z}(t_{i+M} | \mathbf{H}^r) \mathbf{\Gamma}^q \mathbf{T}^{qT} \mathbf{Z}^T(t_{i+M} | \mathbf{H}^r), \quad (6.11)$$

where $\mathbf{P}(t_{i+M} | \mathbf{H}^q)$ is the analysed error covariance at the target time associated with having assimilated the routine measurement and the q^{th} expected measurement.

Wang and Bishop (2003) showed that \mathbf{T} could be calculated by:

$$\mathbf{T}^q = \mathbf{C}^q (\mathbf{\Gamma}^q + \mathbf{I})^{-1/2} \mathbf{C}^{qT}. \quad (6.12)$$

The ensemble state perturbations were updated at the other four target times and the validation time by substituting the ensemble perturbations at the chosen measurement time with those at the other times in Eq. 6.11. Conveniently, the state perturbations after assimilating the chosen measurement $\mathbf{Z}(t|\mathbf{H}^q)$ are the product of $\mathbf{Z}(t|\mathbf{H}^r)\mathbf{\Gamma}^q$. The next measurement, of a possible $Q-1$ measurements, was chosen by repeating Eq's 6.8 to 6.12 using the updated state perturbations instead of the routine state perturbations. For the previous example tide, the signal variances at each sensor location and target time after assimilating a potential measurement from sensor 2 at target time 5 are shown in Fig 6.12. Relative to the signal variances in Fig. 6.11 the signal variances at target time 5 are reduced six fold for the channel sensor nodes, but less so at earlier target times and for the floodplain sensor nodes. The greatest signal variances now tended to occur 60 minutes earlier at target time 2, with sensor six providing the single largest signal variance at this time.




Fig. 6.12: Bar chart of signal variance for Q possible measurements spread over eight sensor locations and five target times, given a measurement from sensor 2 at target time 5.

In theory, Method 2 has targeted measurements within the constraints of the sensor network topology in terms of the reduction in peak tide variance or signal variance. From here onwards, sets of measurements that are believed to lead to the greatest reduction in state variance at high tide will be called ‘best’ measurement sets. It should be equally possible to target the measurements that are believed to lead to the smallest reduction in state variance at high tide. These measurement sets will be referred to as the ‘worst’ measurement sets. For tide 21 this was done by selecting the measurement with the least signal (sensor 5 at target time 3) from Fig. 6.11. Sensor measurements of no water (zero depth) were removed from the list of possible sensor measurements prior to the implementation of Method 2. This was necessary because the sensors were not always at the lowest point of the channel cross section or floodplain reservoir, meaning that the RC model nodes could be wet without a measurement being taken. Therefore, the absence of a measured water depth did not confirm that the area represented by the model node was dry; only that the water level was below the sensor. In a real time context, prior knowledge of when sensors are dry can only be estimated from the state ensemble. So there is a risk of targeting measurements at dry sensor nodes, especially if the validation time and location is such that the method prefers measurements towards the flood edge.

Method 2 was used to select the six best and worst measurements in terms of peak signal variance for 39 tides. For tide 21 the six best measurements in order of preference are shown in Table 6.2, with the six worst shown in Table 6.3. The first and second columns of Tables 6.2 and 6.3 list the sensor node and target time from which to take a measurement. Also shown are the ETKF estimates of signal variance and error variance *a posteriori* to the measurement being assimilated. The final column contains the measurement itself. The first sensor selected by the best algorithm was a channel node at the closest target time to the validation time. The next sensor selected (sensor 6) was on the floodplain towards the validation location around the time this location would have been inundated for the first time. The remaining measurements are near the main channel at either early or late target times. Signal variance drops by about half between the first and second targeted measurement and then by a further two thirds by measurement three. Incidentally,

the signal variance of the routine measurement at sensor 2 was 0.0508 m, almost ten times the first targeted measurement, whilst being further in time from the validation time than any of the potential targeted measurements.

Table 6.2: Selected measurements from tide 21 chosen by adaptive sampling Method 2 in terms of signal variance.

Sensor	Target time	Signal variance (m)	Error variance (m)	Measurement (m)
-	prior	-	0.0951	-
2	routine	0.0508	0.0443	0.892
2	5	0.0052	0.0390	2.385
6	2	0.0024	0.0366	2.005
3	5	0.0008	0.0358	2.375
2	2	0.0005	0.0353	1.958
4	1	0.0005	0.0349	1.552
3	2	0.0004	0.0345	1.988

Table 6.3: Selected measurements from tide 21 chosen by adaptive sampling Method 2 in terms of least signal variance.

Sensor	Target time	Signal variance (m)	Error variance (m)	Measurement (m)
-	prior	-	0.0951	-
2	routine	0.0508	0.0443	0.892
5	3	<0.0000	0.0442	1.248
9	1	0.0001	0.0441	1.552
6	2	0.0008	0.0434	0.848
5	5	0.0008	0.0426	2.411
5	4	0.0004	0.0422	1.622
5	5	0.0008	0.0413	2.022

Selecting the measurements that were expected to yield the least signal variance at the validation time resulted in a preference for floodplain sensors at target times away from the validation time. Table 6.3 shows that sensor 5 at target time 3 was selected as the least useful measurement in terms of signal variance. Referring to the sensor locations map in Fig. 6.3 this sensor was farthest from the channel and

farthest from the validation location of all the floodplain sensors. Since most of the errors in this model propagate from the RC model's tidal boundary it seems reasonable to assume that error propagation is the principal driver when selecting sensor locations rather than the proximity to the validation location. Conceptually, this makes sense as the tide is a wave of low frequency relative to the sampling interval, such that, sampling the peak of the tide gives more information on its amplitude at a distant point in space and time, than sampling the validation location prior to the arrival of the wave peak.

When targeting measurements to obtain the greatest signal variance (Table 6.2), signal variance decreased rapidly as the number of measurements increased. This was not the case when selecting the worst signal, where signal variance fluctuated as more measurements were selected due to the combined effect of selecting the worst measurements (hence leaving those with greater signals) and decreasing error covariance.

Testing the ETKF estimates of signal variance against those from the EnKF approach was computationally difficult, as it involved running the RC model repeatedly from the target times to the validation time. For each tide, the EnKF and simulation model was run from t_i to t_v for one measurement, then again for two measurements, and again for measurements three through six. To mitigate this problem the test was run only on tides 18 to 40, the reasons being that tides 18 and 19 were the first not to require one or all of the available target times when applying Method 1 and tide 21 was the largest tide in the series. Since the ETKF required only a few minutes on a 2.4 GHz PC to conduct the same analysis, it was run for all 39 events, with events 2 through 17 discarded from subsequent analysis so that a direct comparison could be made with the EnKF results. For both the EnKF and ETKF based approaches the variance forecasts for each tide were calculated given varying numbers of measurements. From tides 18 to 40 the median, greatest and smallest estimates of forecast error variance made by the two approaches were extracted. The median was used because the distribution was skewed towards larger variances. A comparison between the two approaches is made in Fig. 6.13, with forecast error variance at the validation time and location on the y-axis, and the total number of targeted

measurements assimilated on the x -axis. Also shown are the prior error variance and the forecast error variance at the validation time and location, after assimilating the routine measurement. The bold lines on Fig. 6.13 indicate the median forecast error variance, with the thinner lines indicating the largest and smallest variances. The continuous lines are taken from the EnKF-based forecast error variance, whilst the broken lines are variances estimated by the ETKF. Results are shown for the best measurement case (darker lines) and the worst measurement case (lighter lines) described above. There were, therefore, four sets of results:

1. EnKF with best measurement set.
2. EnKF with worst measurement set.
3. ETKF with best measurement set.
4. ETKF with worst measurement set.

The plot shows that the greatest fall in forecast error variance usually occurred as a result of assimilating the routine measurement, followed by the first targeted measurement in the best measurements case, but the third targeted measurement in the worst measurement case. The worst measurement scenario tended to be distributed over greater values of error variance than the best measurement scenario. The upper and lower extent of the variances as denoted by the thinner lines shows a closer agreement between the lower extents than the upper extents. Therefore, the ETKF was more accurate when predicting the lower extent of the variance range than the upper extent, where anomalies are present after targeted measurements 1 and 5.




Figure 6.13: Results of adaptive sampling Method 2. The plot shows forecast error variance prior to data assimilation, after the routine measurement and after assimilating a further one to six measurements chosen by the ETKF. Continuous lines summarise the median (thick line) or range (thin lines) of forecast error variance from the EnKF, whilst broken lines denote equivalent estimates by the ETKF. Black lines indicate that the ‘best’ measurement set was assimilated; whilst red lines indicate that the ‘worst’ measurement set was assimilated.

The ETKF always predicted that assimilating a measurement would result in a decrease in forecast error variance, although, these decreases were small for the first worst measurement and later best measurements. This is easier to see in the close-up of the median error variance estimates from the four sets of model runs shown in Fig 6.14. The median of the ETKF predictions were similar to the EnKF for the first and second targeted measurements, with the worst measurement set having the greater median error variance. When more than two targeted measurements were assimilated the ETKF variance forecast was less than the EnKF variance forecast. In fact, the median error variance increased as more measurements were assimilated. This increase in EnKF error variance was due to the likelihood of model instability in RC

simulations increasing as more measurement data were assimilated by the EnKF (see chapter 5).




Fig. 6.14: Results of adaptive sampling Method 2 with median results only.

6.4 Discussion

As with many sampling design problems, adaptive sampling is a trade off between a desire to collect as much information as possible through observation and the limitation of only being able or willing to make a finite number of observations. The river monitoring problem outlined here is both spatial and temporal in that state changes over both of these dimensions. Crucially, these changes are not entirely random but are correlated in a manner that can be partially predicted. Through an understanding of the processes and driving forces that govern these changes it was possible to simulate state with the RC model. By incorporating uncertainty into the forcing terms and parameters of the process representation ensembles of equally likely states were simulated, allowing state to be treated as a random variable rather

than a deterministic prediction. In this application, state error covariance was approximated from state ensembles generated by the RC model, and used as part of a near-optimal sequential data assimilation process based on an EnKF. The outcome of this process was an updated state ensemble conditioned on both the model predictions and measurements of state, which was used to provide the initial conditions for a state forecast. Method 1 was based upon these forecasts since it demanded data from sensor nodes until the probability of state exceeding a threshold was less than 5%.

Method 1 was inefficient in that the measurement data collected were not prioritised based on any measure of their usefulness, but were collected simply because the forecasting model predicted a significant probability of flooding. This inefficiency contradicted objective two of the adaptive sampling problem outlined in section 3. The ensemble of model simulations needed to be propagated from the time of the targeted measurement to the validation time each time a decision about further sampling was made. This was computationally expensive since the simulations were by far the most computationally intensive component of the forecasting model. To implement Method 2 the simulation model had to be run twice between the routine measurement time and the validation time, once to generate the background forecasts and once to assimilate the measurements and issue a forecast. Therefore, when Method 1 required more than one target time it was computationally more intensive than Method 2. Furthermore, the computation required to implement Method 2 is known *a priori* whereas the computation required for Method 1 is not. A further problem with Method 1 is that it is not suitable as a means of deciding when and where to collect measurements as it only provides a stopping criterion, rather than a framework for targeting observations.

The information content or usefulness of the measurements fed into the EnKF was directly related to the reduction in error (co)variance brought about by assimilating them. It follows that the greater the reduction in error variance the more useful a given number of measurements and the more efficient the sampling design, complementing adaptive sampling objective two. Method 2 complements objective two because it provided a framework for targeting measurements based on signal

variance. This allowed it to be more efficient than Method 1 because measurements were targeted based on their signals, rather than simply measuring at the next opportunity. However, the probability of flooding could only be estimated if it was assumed that the ensemble mean was not changed by the subsequent measurement data. Herein lies the principal problem with Method 2 in that it does not provide a direct solution to objective one of the adaptive sampling problem. For example, although signal variance can be estimated the change in the ensemble mean is unknown and, thus, the probability of state exceeding a threshold is also unknown. Therefore, Method 2 would appear best suited to a scenario where only a finite number of resources are available with which to take measurements, rather than the more flexible situation found here because it cannot provide a probability value until after the measurements have been collected. However, for the River Crouch example presented here the increase in signal variance decreased rapidly as more measurements were assimilated (Fig. 6.13). The ETKF predicted that 95% of the signal variance from the routine measurement plus six targeted measurements was available from two targeted measurements and that 99% was available from three targeted measurements. In fact, this estimate is probably conservative because the routine measurement would not have been selected as the optimal first measurement due to its timing. A further point, which is possibly specific to this model, was that state instabilities in the simulation model increased as more measurement data were assimilated, making it counterproductive to assimilate more than two-to-three targeted measurements in practice. This may be solved through improvements to the simulation model or additional rules concerning the specification of initial conditions prior to issuing forecasts.

Method 2 raises a number of issues for the design and operation of the sensor network at this site:

1. Most of the sensors at this site add little to signal variance since model errors develop and propagate over much greater distances than those of the deployment.
2. Taking the maximum 12 measurements in an hour and then transmitting a real-time measurement once an hour is likely to provide almost as much

signal variance as transmitting every 20 minutes. Thus, for an adaptive sampling procedure to be required the sensors would need to be capable of storing more measurement data but need not transmit data as often.

3. In theory Method 2 could be used to plan a redeployment of the sensor nodes, such that, the sensor network layout is optimised for signal variance, potentially providing a framework for mitigating the situation outlined in point 1, where most of the sensors add little to the signal variance.

The final point in the above list is as yet untested. It is not known to what extent ETKF will be able to distinguish between potential sensor locations, especially as the distance between a sensor and the validation location increases. For example, will spurious covariance relationships between possible sensor locations and other locations in the model state space lead to a physically unrealistic sampling framework being proposed? This issue will be investigated further in the next chapter.

One thing to consider about both adaptive sampling methods is that they collect measurements only on certain aspects of the state dynamics as determined by the simulation model. The preference for certain measurements may lead to bias if used as part of a parameter estimation procedure because the model parameters would be tuned to the targeted measurement(s). Thus, forecasting algorithms with online parameter estimation may not be suitable in situations where adaptive sampling has been used to collect measurements, instead of a sampling framework designed to gather calibration data. This is especially true in the case of the ETKF where the targeting of measurements relates directly to the parameterisation of the simulation model and sources of uncertainty used to generate the state ensemble.

Whilst the adaptive sampling methods presented here both meet the original aim of providing a framework for transmitting less real-time data from the FloodNet sensor network there are several outstanding issues raised by the experiments conducted. For the forecasting model, simulation instability had a detrimental effect on model forecasts not accounted for by the ETKF. An improved definition of model parameters and sources of uncertainty was also required, especially as they may be

different to those of relevance in a model calibration situation. From an adaptive sampling viewpoint the test needs to be conducted at a location where there are significant differences between the information provided by different sensors and a genuine need for multiple sensors. A larger river catchment with multiple regions of rainfall-based flow generation or significant tidal and precipitation induced flows would provide a more robust test. This is particularly important for Method 2 where the assumptions made by the ETKF over those of the EnKF were not adequately tested by the topology of the sensor network. For example, how accurate would the ETKF estimates of signal variance be in a situation where the validation location was displaced spatially from possible target locations, and there was a two day lead time between a flood peak arriving at each?

Not all FloodNet sensor nodes had a direct link to the gateway node, with which to transmit their data offsite. These nodes required intermediate nodes through which data could be routed to the gateway node. Method 2 did not consider this routing component when targeting measurement from sensor nodes on the periphery of the sensor network relative to the gateway node. Since a node would need to be on and using power in order to act as a relay between a targeted node and the gateway node it could send its own measurements at little extra power cost. This property of the FloodNet network needs to be incorporated into Method 2 to more accurately reflect the cost of taking a measurement. Furthermore, there is the potential to speed up sampling decisions by making them on the FloodNet ‘inner loop’ in a more bottom up approach than those advocated here.

6.5 Conclusions

For the RC forecasting model used here measurements taken systematically over space and time were shown to be less efficient than those targeted by an ETKF at reducing forecast state errors. Therefore, when the propagation of state errors can be simulated and the objective of the exercise is to forecast certain extreme conditions, rather than characterise the spatio-temporal state dynamics as a whole, adaptive

sampling can be a more efficient means of sampling than a simple systematic method.

Measurements taken by the FloodNet sensor network were highly spatially correlated due to the relatively small distances between sensor nodes. This meant that assimilating measurements from multiple sensors had a limited effect on error variance and provided a limited test of the ETKF used by Method 2. The forecasting model was prone to state instability, especially as more measurements were assimilated. This issue requires further investigation and prevented a more detailed analysis of the difference between ETKF and EnKF derived estimates of error covariance.

6.6 Summary

The movement of computational power and communications capabilities onto networks of sensors in the environment through the concept of pervasive or ubiquitous computing has initiated opportunities for the delivery of ground-based data in real-time and the development of adaptive monitoring systems. Measurements of water level taken by a network of wireless sensors called 'FloodNet' were assimilated into a one-dimensional hydrodynamic model using an ensemble Kalman filter, to create a forecasting model. This research focused on methods for targeting measurements in real-time to be assimilated by the forecasting model, such that these power limited but flexible resources could be used optimally. Two targeting methods were developed. The first targeted measurements systematically over space and time until the forecasting model predicted that the probability of the water level exceeding a pre-defined threshold was less than 5%. The second method targeted measurements based on the expected decrease in forecasted water level error variance at a validation time and location, quickly calculated for various sets of measurements by an ensemble transform Kalman filter. Targeting measurements based on the decrease in forecast error variance was found to be more efficient than the systematic sampling method.

7 Optimal sampling design with an ensemble transform Kalman filter

7.1 Introduction

The principal objective of chapter 6 was to adapt the sampling framework of the FloodNet sensor network in real-time to enable a power saving on the sensor nodes. A method was developed that could target stage measurements to a set of discrete times and sensors based on signal variance at a validation time and location. Signal variance was calculated from an ensemble of background simulations by the RC model using an ETKF and estimates of measurement variance. In the discussion, it was suggested that the ETKF may be used to target measurements to any computational node in the RC model domain (e.g. any point at which stage was simulated) and to assess the use of sensors with different error statistics (e.g. different measurement variance). Targeting measurements to previously unmonitored locations is appealing in the FloodNet case study because it provides a means of re-deploying the sensor nodes to optimal locations along the river. More generally, when this capability is combined with the capability to consider different sensor error statistics, the ETKF provides a quantitative means of assessing the benefits (in terms of signal variance) of sampling with almost *any* sensor type at *any* location.

To the author's knowledge the ETKF has not been applied previously to the problem of designing an optimal sampling framework for a river flood forecasting model. Therefore, the response of the algorithm to a number of test scenarios, useful for the design of a sensor network deployment, was investigated. Specifically, the spatio-temporal changes in signal variance from theoretical sensors within the main river channel were calculated for the validation location at the target times and validation time (section 7.2). The same analysis was conducted for measurements taken on the floodplain (section 7.3) and measurements with different variances (section 7.4). These results are followed by a discussion and some conclusions.

Targeting measurements to previously unmonitored locations is possible because the RC model simulates stage and flow at distributed locations, where state measurements need not be available. A number of assumptions were made when implementing the tests in sections 7.2 to 7.4. These include:

1. The RC model simulations of error (co)variance were equally accurate at all locations and between all locations in the model domain (e.g. the RC model and EnKF parameters were not tuned to measurement locations by the calibration process in section 5.2.2).
2. The same measurement error statistics would apply at all locations and to all sensors.
3. The errors in theoretical measurements were uncorrelated with each other and the system state.
4. A linear transform of the forecast state perturbations was a reasonable analogue for the error propagation of an ensemble of RC model simulations.


In chapter 4 assumption 1 was shown to be untrue because error variance forecasts on the floodplain were too small and the covariance between channel and floodplain errors was too large. Nevertheless, the approach did forecast increases in analysed state uncertainty with distance from measurements that were consistent but not equal to what was observed. Despite the approach requiring further sources of uncertainty it was more representative than assuming a flat water surface with equal uncertainty everywhere. Furthermore, in later chapters, roughness uncertainty was included in RC model simulations. In the context of this study, assumptions 2 and 3 are likely to be met because each location would be monitored by a different sensor, although it may be possible to argue that sensor error statistics are related to the depth, temperature and salinity of the water being measured and thus not equal at all locations. Assumption 4 was tested in chapter 5. The linear transform was not an exact predictor of the EnKF approach due to RC model instability and nonlinear error propagation. However, it did exhibit the same pattern behaviour to the EnKF approach in terms of error and signal variance, when assimilating measurements. These assumptions are revisited in chapter 8.

7.2 Signal variance for in-channel sensors


In this experiment, the ETKF was used to calculate expected target time signal variance $s_j(t_m | \mathbf{H}_m)$ at the validation location j and target time t_m (see Fig. 6.4) for a measurement at the location of any of the RC model channel computational nodes using Eq. 3.65. The analysis was conducted at nine target times equally spaced between 4 hours and 1 hour and 20 minutes prior to the validation time. The validation time signal variance at the validation location $s_j(t_h | \mathbf{H}_m)$ was also calculated. A fixed measurement variance of 0.005 m was adopted for each possible measurement. Only one measurement was used to calculate signal variance in this case.

Figs. 7.1.a-i show mean ensemble state (stage and flow) and state uncertainty (variance) during tide 26 for each channel node in the RC model at nine 20 min time steps starting 4 hours prior to the validation time, which was at high tide. All distances on the x -axis are in km from the upstream model boundary at Wickford. A map of RC model channel nodes and a long section of the river thalweg can be seen in Fig. 7.2. The tide was coming in from east to west (right to left) at all nine target times in this example, as indicated by the negative flow values at the downstream boundary (Fig 7.1). In Fig 7.1.a there was positive flow between 0 and 2.5 km downstream (BNG easting 577000 and 579000 respectively) because the incoming tide has yet to inundate this section of the estuary. As time progresses towards high tide this area of positive flow became progressively smaller until the whole domain was affected by tidal flow in Fig 7.1.f, 2 hours 20 minutes before high tide. The estimated confluence between incoming tidal and fluvial flows during this time was demarked by the lowest points on the mean stage plot. Values of ensemble mean stage ranged between -0.79 m in Fig. 7.1.a and 2.24 m in Fig. 7.1.i, whilst mean flow ranged from almost -2000 cumecs to 0.35 cumecs.


(a)




(b)




(c)




(d)




(e)




(f)



(g)



(h)



(i)




Fig 7.1.a-i: Plots of model states (stage and flow) and their uncertainties (variance) for channel cross-sections, with log signal variance at validation location for validation and target times given a measurement with variance 0.005. Note: all axes are constant between plots except log signal variance. a) Plots of states and signal 4 hours prior to validation time. b) Plots of states and signal 3 hours 40 minutes prior to validation time. c) Plots of states and signal 3 hours 20 minutes prior to validation time. d) Plots of states and signal 3 hours prior to validation time. e) Plots of states and signal 2 hours 40 minutes prior to validation time. f) Plots of states and signal 2 hours 20 minutes prior to validation time. g) Plots of states and signal 2 hours prior to validation time. h) Plots of states and signal 1 hour 40 minutes prior to validation time. i) Plots of states and signal 1 hour 20 minutes prior to validation time.


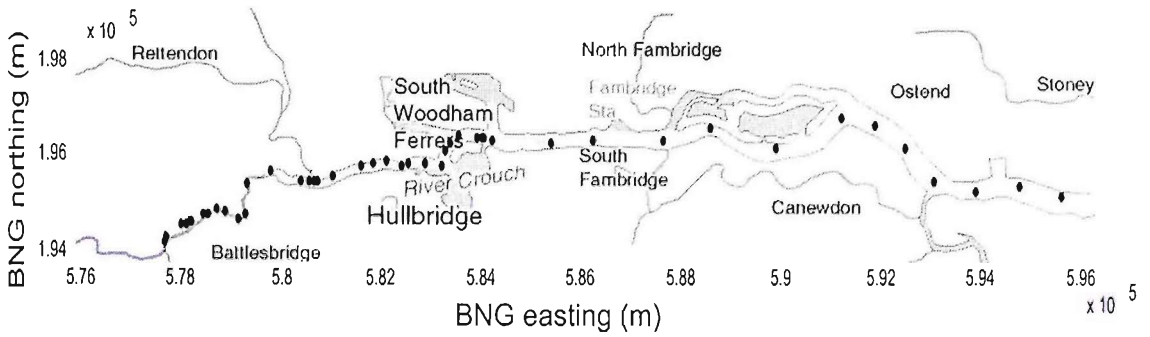


Fig. 7.2: Top: Map of RC model channel nodes (black dots). The study site is at Hullbridge. Bottom: Long section of river thalweg and bank heights.

State variance increased between the tidal boundary at 20.3 km downstream (BNG easting 595000) and the validation location at 6.3 km downstream (BNG easting 583500). The increase in stage variance was more pronounced at lower stage values, with almost no increase after 2 hours prior to high tide. This increase in variance is believed to be predominantly due to the nonlinear manner in which the model simulates flow along the river for different stages and the affect of variations in the value of Manning's roughness coefficient over the ensemble. Variance at the downstream boundary of the RC model was roughly constant at around 0.1 m. State variance rose to 0.137 m (a change of 0.037 m from the downstream boundary) at the

validation location at the first target time (Fig. 7.1.a). At the ninth target time (Fig. 7.1.i) the change in variance over this distance was less than 0.001 m with variance falling slightly at some intermediate RC model nodes. The greatest values of stage variance were upstream of the study site at the confluence of fluvial and tidal flows. As such the region of greatest stage uncertainty in the model domain migrated upstream with the incoming tide. The greatest stage variance estimate at any of the nine discrete time steps was 0.274 m and occurred at 1.4 km downstream (BNG easting 578800) in Fig. 7.1.f. Upstream of the confluence stage variance was approximately two orders of magnitude smaller, which was expected as no errors were added to the upstream boundary condition. Therefore, any stage variance in this region of the model domain was due to the variation of channel roughness over the ensemble. The variance in flow increased with mean flow, thus, the largest variances were towards the tidal boundary, where flow magnitude was greatest. Unlike stage variance, flow variance was not greatest at the confluence of fluvial and tidal flows. However, the standard deviation in flow as a proportion of flow was greater where stage variance was greatest. Flow variance ranges between 0 cumecs upstream of any tidal influence and 19740 cumecs at the downstream boundary in Fig 7.1.i.

Figs. 7.1.a-i plot estimated signal variance at the validation location for a theoretical measurement at each channel node in the RC model. Signal variance at the validation location is shown at the target time and the validation time. Note that the distribution of signal variance on the x -axis is determined by the location of the measurement not the validation location, which remained constant.

Target time signal variance at the validation location varied with the spatial location of measurements. Target time signal variance ranged from 1.4×10^{-4} m upstream of the confluence to 0.145 m close to the maximum stage variance in Fig 7.1.a. Both of these extremes occurred upstream of the validation location. The smallest values of signal variance occurred upstream of the flow confluence, indicating that fluvial flows had very little influence over stage variance at the validation location, according to the RC model setup used to generate the background ensemble. In Fig 7.1.a the largest values of target time signal variance did not occur at the validation location but sat between the validation location and the location of peak stage

variance. Closer (temporally) to the validation time, such as in Fig 7.1.f, this relationship was no longer as prominent, with the peak target time signal variance closer to the validation location relative to the peak stage variance, despite peak target time stage variance being greater. In fact target time signal variance became spatially more correlated over the length of the river as the tide came in, whilst decreasing with falling stage variance. Downstream of the validation location the signal is seen to decrease with distance from the validation location but not as dramatically as upstream of the validation location. This was reflected by a range of target time signal variance of 0.088 m to 0.142 (Fig. 7.1.i and Fig. 7.1.a respectively). The range in stage variance was also less than upstream of the validation location.

The pattern in target time signal variance occurred due to a combination of the sensitivity of the state ensemble at the target location to measurement data and the covariance structure between the target location and validation location. This relationship is illustrated by the next series of plots. These plots contain points for all RC model nodes at 8 target times between 3:40 and 1:40 hours prior to the validation time calculated over 39 tides.

Fig. 7.3 plots the target time signal variance against target time stage variance at the target location. The plot indicates a positive relationship between signal variance and the variance at the target location, which was relatively straightforward (linear) until stage variance reaches a value of just less than 0.1 m, which coincides with the minimum state variance seen downstream of the validation location. Above a stage variance of 0.1 m the signal variance trace splits into two clusters, the first shows little increase in signal variance with variance. The second trace shows signal variance varying between 0.1 m and 0.18 m, of which a proportion could be attributed to stage variance. Fig. 7.4 compares the distance between the target location and the validation location with stage variance. This plot supports the pattern in Figs. 7.1.a-i which showed that stage variance was greatest towards the front edge of the incoming tide and was greater in magnitude the further the flow wave had travelled upstream. The state variance was very small upstream of the incoming tide principally due to the manner in which the ensemble was generated. A

comparison of the target time signal variance at validation location with the error covariance between target location and validation location is shown in Fig. 7.5. There was a stronger positive relationship between these two variables than between signal variance and stage variance in Fig 7.3. Nevertheless two distinct point clusters can be seen above a signal variance of 0.08 m. Covariance decreased with distance from the validation location as illustrated by Fig. 7.8. The smooth and relatively uniform relationship in the downstream section of the river contrasted markedly with that of the upstream section. A 3D plot of signal variance, stage variance and covariance (Fig. 7.6) again demonstrates the two traces present in Figs 7.3 and 7.5. In this plot the nodes downstream of the validation location have been highlighted in red. These red dots occupy a relatively small range of state variances compared to the upstream nodes, with signal variance mainly determined by the covariance and the stage variance at the validation location. Viewed from a different angle in Fig. 7.7 the scatter of points forms a near flat surface in the covariance/signal variance space. Referring back to plot 7.1.a-i the cause of the peak in signal variance just above the validation location was believed to be the combination of the covariance and state variance upstream of the validation location. The point cloud on Fig. 7.6 supports this conclusion as, although, the covariance between validation location and target location was the dominant factor in determining signal variance stage variance did have a secondary effect. Values of stage variance above around 0.15 m were split into two distance signal variance clusters depending on covariance, indicating that covariance decreased rapidly at the edge of the incoming tide relative to stage variance.




Fig. 7.3: Plot of variance in stage at the target location and time against the signal variance at the validation location and target time. Results are shown for 39 consecutive tidal events at each channel node in the RC model.




Fig. 7.4: Plot of stage variance at the target location against distance from the validation location. Results are shown for 39 consecutive tidal events at each channel node in the RC model.




Fig. 7.5: Plot of target time covariance between target location and validation location against target time signal variance. Results are shown for 39 consecutive tidal events at each channel node in the RC model.




Fig. 7.6: Plot of target time signal variance against target time stage variance at the target location and target time error covariance between the target location and validation location. Results are shown for 39 consecutive tidal events at each channel node in the RC model.




Fig. 7.7: Plot of target time signal variance against target time stage variance at the target location and target time error covariance between the target location and validation location, from an alternative angle. Results are shown for 39 consecutive tidal events at each channel node in the RC model.




Fig. 7.8: Plot of error covariance between target location and validation location against distance from the validation location. Results are shown for 39 consecutive tidal events at each channel node in the RC model.

In Figs. 7.1.a-i the validation time signal variance was less than that at the target time at all eight target times represented. In general, the difference between the signal variance at the validation and target times decreased with the lead time between the two. This decrease resulted from two factors. Firstly, the signal variance at the target time decreased as the tide came in because of reducing stage variance at locations that strongly co-varied with the validation location. Secondly, signal variance at the validation time tended to increase as lead time reduced (Fig. 7.1.a-i).

Another difference between validation and target time signal variances was that the peak in validation time signal variance occurred downstream of the validation location, rather than at the location with the most favourable combination of target time variance and covariance. Therefore, peak validation time signal variance calculated here acknowledges the principal source of errors in the RC model at the downstream boundary condition, recognising that, in this model, error propagation is dependent on the flow of errors upstream from the downstream boundary.

If error propagation was sufficiently dependent on errors propagating upstream (a flow of errors) from the model boundary then it should be possible to create a situation where the signal variance at the validation location from a downstream (or up-flow) sensor is greater at the validation time than the target time. This is not the case for any of the examples in Figs 7.1.a-i. Therefore, an experiment was conducted in the next section with the validation location moved upstream.

7.2.1 Flow dependent error propagation and signal variance


Floods propagate through space over time, as a wave. As a consequence, an adaptive sampling approach should ideally be able to target a 'flood wave' before any affect from that wave is noticed at the validation location. The ability of the ETKF to predict signal variance before the incoming tide has reached the validation location was tested by moving the validation location to the upstream sensor location (see Table 7.1 for definition of upstream location). Signal variance was then calculated for theoretical sensors at each of the RC model channel nodes at the nine target times used in the previous section. The results are plotted in Fig. 7.9.a-i.

Table 7.1: Table describing example sensor locations. Distances upstream of the validation location are denoted by negative numbers.


Name of theoretical sensors	Distance downstream (m)	Distance from validation location (m)	BNG Easting	BNG Northing
Upstream	1,940	-4,367	579216	195470
Validation location	6,304	0	582758	195902
downstream	18,576	12,272	593814	593814

For Figs. 7.9.c-i the pattern of validation and target time signal variances was similar to those in Figs 7.1.a to 7.1.i, in that, the target time signal variance was greater than the validation time signal variance at all locations along the river. Interestingly, in Figs. 7.9.c-i the validation time signal variance was similar to the equivalent times in Figs. 7.1.c-i, whilst the target time signal variance was greater. This occurred because the stage variance at the upstream validation location was greater at these target times than that at the study site, but, similar at the validation time. At the time of Fig. 7.9.a the incoming tide had yet to reach the upstream validation location. Therefore, collecting data at almost any location along the river resulted in very little target time signal variance. Nevertheless, the validation time signal variance was similar to the signal variance in Fig 7.1.a.


(a)




(b)




(c)




(d)




(e)




(f)



(g)



(h)



(i)




Fig 7.9.a-i: Plots of model states (stage and flow) and their uncertainties (variance) for channel cross sections, with measurement with variance 0.005. Validation location was 1.9 km downstream. a) Plots of states and signal 4 hours prior to validation time.

b) Plots of states and signal 3 hours 40 minutes prior to validation time. c) Plots of states and signal 3 hours 20 minutes prior to validation time. d) Plots of states and signal 3 hours prior to validation time. e) Plots of states and signal 2 hours 40 minutes prior to validation time. f) Plots of states and signal 2 hours 20 minutes prior to validation time. g) Plots of states and signal 2 hours prior to validation time. h) Plots of states and signal 1 hour 40 minutes prior to validation time. i) Plots of states and signal 1 hour 20 minutes prior to validation time.


In Fig. 7.9.b some of the ensemble members predicted that the tide had reached the validation location whilst others did not. Under this circumstance, the peak target time signal variance was at the validation location. Target time signal decreased between the validation location and 5.5 km downstream but then increased slightly towards the downstream boundary. Validation time signal variance was similar to

that in Fig. 7.1.b due to the large covariance between all the locations in the RC model at high tide.


7.2.2 Signal variance after one measurement

A scenario was set up where the one measurement r was assimilated 4 hours prior to the validation time (referred to as the measurement time). This was achieved by using Eq. 3.60 and 3.61 to calculate the transform matrix \mathbf{T} . This transform was then used to update background ensemble perturbations $\mathbf{Z}(t)$ at eight target times and the validation time. This was repeated for 39 consecutive tides. The updated ensemble perturbations $\mathbf{Z}(t|\mathbf{H}^r)$ were then used as the background forecast in equation 3.65 to calculate signal variance for a theoretical measurement at each channel node in the RC model (see Fig. 7.2). There were 161 possible locations in the RC model where the initial measurement used to calculate \mathbf{T} could have been taken. This series of calculations took less than 0.1 seconds on a 2.4 GHz PC for a possible sensor location and time. However, the computation of signal variance at 161 nodes and eight target times for 39 tides at both the target and validation time took around one hour. The number of data was also difficult to summarise concisely. Therefore, each example shown in Fig. 7.10.a-h used a ‘measurement’ taken from one of the three locations called the upstream, validation and downstream location (Table 7.1), during tide 26 only. Since no measurements were available at the upstream or downstream sensor location the variance of the measurements from these locations was assumed to be the same as that used for the validation location sensor.


(a)




(b)



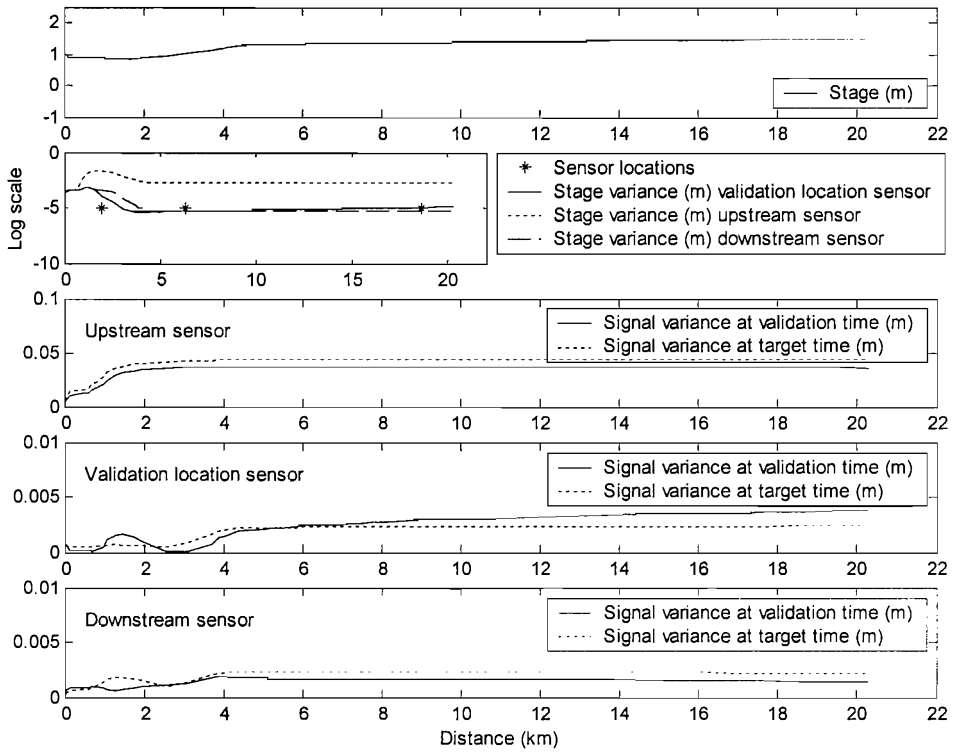
(c)




(d)




(e)



(f)



(g)



(h)




Fig 7.10.a-h: Plots of model states (stage and flow) and their uncertainties (variance) for channel cross sections, with measurement with variance 0.005. After assimilating one measurement at one of three locations 4 hours prior to the validation time with the ETKF. a) Plots of states and signal 3 hours 40 minutes prior to validation time. b) Plots of states and signal 3 hours 20 minutes prior to validation time. c) Plots of states and signal 3 hours prior to validation time. d) Plots of states and signal 2 hours 40 minutes prior to validation time. e) Plots of states and signal 2 hours 20 minutes prior to validation time. f) Plots of states and signal 2 hours minutes prior to validation time. g) Plots of states and signal 1 hour 40 minutes prior to validation time. h) Plots of states and signal 1 hour 20 minutes prior to validation time.

Mean stage estimates from the background ensemble were plotted for each channel node in the RC model domain in the uppermost plot of Figs. 7.10.a-h. Estimates of target time stage variance using the above method are displayed for each of the three example sensors on the second plot in each figure. Target time signal variances for each target time and location were plotted for the example measurement locations in Table 7.1. The RC model predicted that stage variance would increase as the target time moved away from the measurement time. This increase was more pronounced for the downstream and validation location measurements, reflecting the greater signal variance available from measurements downstream and at the validation location. Disregarding stage variance in the upper 2.5 km of the river, the lowest stage variance associated with the downstream and validation location measurements always occurred close to the measurement location, although, this became less apparent as the time between the measurement time and target time increased.

Signal variances after assimilating a measurement from the upstream sensor were similar to those in Fig 7.10.a-h. This was expected given the relatively small signal variance predicted at this node (see Fig. 7.1.a) and supported by the relatively large stage variances (relative to the other two locations) after assimilating the synthetic measurement from this location. The target and validation time signal variances after assimilating a measurement at the validation location sensor indicated a decrease in signal variance of approximately one to two orders of magnitude, compared with


Figs. 7.1.b-i. Decreases in signal variance were not constant over either space or time. For the first target time after the measurement time (Fig. 7.10.a) target time signal variance was greater than validation time signal variance at most locations between 0 km and 12 km downstream. This trend was reversed between 12 km and 21 km downstream as validation time signal variance increased towards the tidal boundary of the RC model. As time progressed away from the measurement time, the crossover described above moved upstream, until reaching 5 km downstream 1 hour 40 minutes after the measurement time (Fig. 7.10.e). From this time onwards the crossover remained constant around 8 to 9 km downstream. The downstream sensor yielded the greatest validation time signal variance of the three locations tested. In Figs. 7.10.a-c target time signal variances peaked between 3 and 5 km downstream, whilst validation time signal variance was greater towards the tidal boundary. For Figs. 7.10.d-h target time signal variances were slightly larger than at the validation time for all locations greater than 3 km downstream. However, there was little to choose between possible measurement locations greater than 3 km downstream at the target or validation time in these plots.

7.2.3 Signal variance and time series data


Once a sensor has been deployed it is unlikely to collect a single measurement. In the case of a pressure transducer measuring water depth the sensor is able to collect a time series of data at almost any sampling interval, provided sufficient power and data storage is available on the sensor node. Therefore, it is imperative that any method used to determine the location of a sensor can accommodate time series data in addition to data points. The signal variance of a measurement in addition to a time series of measurements was estimated by using the ETKF to update RC model ensemble perturbations with a theoretical time series of measurements from one of three locations at eight target times, using Eq. 2.59-64. Since chapter 4 did not compare ETKF and EnKF estimates of error variance given a time series of measurements the EnKF forecasting model was run in parallel with the ETKF experiment, using a time series of measurements from the validation location.

Fig. 7.11.a-h shows stage variance and signal variance assuming a time series of measurements were collected and assimilated from BNG easting 581000 (upstream, short dashes) or BNG easting 595500 (downstream, long dashes). Predictions of mean stage, mean flow and flow variance were not updated. For Figs. 7.11.a-h assimilating the upstream measurement time series resulted in a greater average stage variance over the whole river, compared to the time series from the downstream and validation location sensors. For all measurement locations stage variance increased with distance from the sensor location.


(a)




(b)




(c)




(d)




(e)



(f)



(g)



(h)




Fig. 7.11.a-h: Results of assimilating a time series of synthetic measurements at 3 locations within the RC model domain. The top plot in each figure displays mean stage estimates from the background ensemble. Estimated stage variance after assimilating a measurement 4 hours prior to the validation time is displayed for each of the three example sensors. Signal variance for the target time and locations was plotted for each example sensor. Flow, flow variance and stage were not updated during the ETKF analysis. Note: all axes are constant between plots except signal variance. The lower three plots on each figure show the target and validation time expected signal variance given either the upstream, validation or downstream sensor. a) Plot of signal variance 3 hours 40 minutes prior to the validation time. b) Plot of signal variance 3 hours 20 minutes prior to the validation time. c) Plot of signal variance 3 hours prior to the validation time. d) Plot of signal variance 2 hours 40 minutes prior to the validation time. e) Plot of signal variance 2 hours 20 minutes prior to the validation time. f) Plot of signal variance 2 hours prior to the validation time. g) Plot of signal variance 1 hour 40 minutes prior to the validation time. h) Plot of signal variance 1 hour 20 minutes prior to the validation time.

For the upstream measurement time series, target and validation time signal variance increased with distance from the sensor location. Since this sensor location resulted in the smallest fall in state variance the signals available from other locations are generally greater than those after assimilating the validation and downstream sensor time series.

After assimilating the validation location measurement time series, target time signal variance was greatest just upstream of the validation location in Figs 7.11.a-c. For Figs 7.11.d-h the greatest target time signal variances occurred towards the downstream boundary. Validation time signal variance tended to increase towards the downstream boundary, but, also spiked in conjunction with the greatest state variance upstream of the validation location. This is a potentially undesirable trait of this approach as there seems little reason to believe that flows upstream of the validation location could increase forecast accuracy at the validation location by such an amount in this particular model, even though the greater state variance upstream

of the validation location will make state at these location more sensitive to measurements.

After assimilating the downstream measurement time series, target time signal variance was predicted to increase with distance from the downstream sensor until the upstream edge of the incoming tide. The magnitude of this increase decreased for target times closer to the validation time as model error covariance increased between the locations of the downstream and validation location sensors. Validation time signal variance was less than 0.001 m at all locations and for all eight target times. These signal variances are lower than those predicted after assimilating the upstream and validation location time series. Therefore, according to the ETKF algorithm, the optimal location (within the RC model domain) for a sensor collecting a time series of measurements given the validation location and target times used here was the downstream boundary.

An experiment was conducted on tide 26 to compare the ETKF predictions of stage error variance with EnKF predictions of stage error variance. The EnKF was used to assimilate a time series of measurements from sensor 2, which was also the validation location. After each update time (target time) the RC model simulated a set of updated state estimates at subsequent target times and the validation time (high tide). The forecasts issued after each target time are shown in Fig. 7.12 as dark lines overlaying RC model background predictions (lighter lines) made before any data assimilation. Fig. 7.13 contains a histogram of background state at the validation time (top left), followed by histograms of validation time forecast state after assimilating one through to six measurements from sensor 2. The bottom right plot of Fig. 7.13 plots the changes in the ensemble mean at the validation time and location as more measurements were assimilated (left hand y-axis, darker line). The lighter line and right hand axis plots the changes in error variance as more measurements were assimilated by the EnKF. The dots are corresponding error variance predictions made by the ETKF. The ensemble perturbations used by the ETKF were calculated from an ensemble of RC model simulations made after assimilating the measurement at target time 1 with the EnKF.




Fig. 7.12: Forecast state at sensor 2 at six target times (dark lines). Light lines show background simulations made before data assimilation, whilst black stars are measurements at five minute intervals from sensor 2.

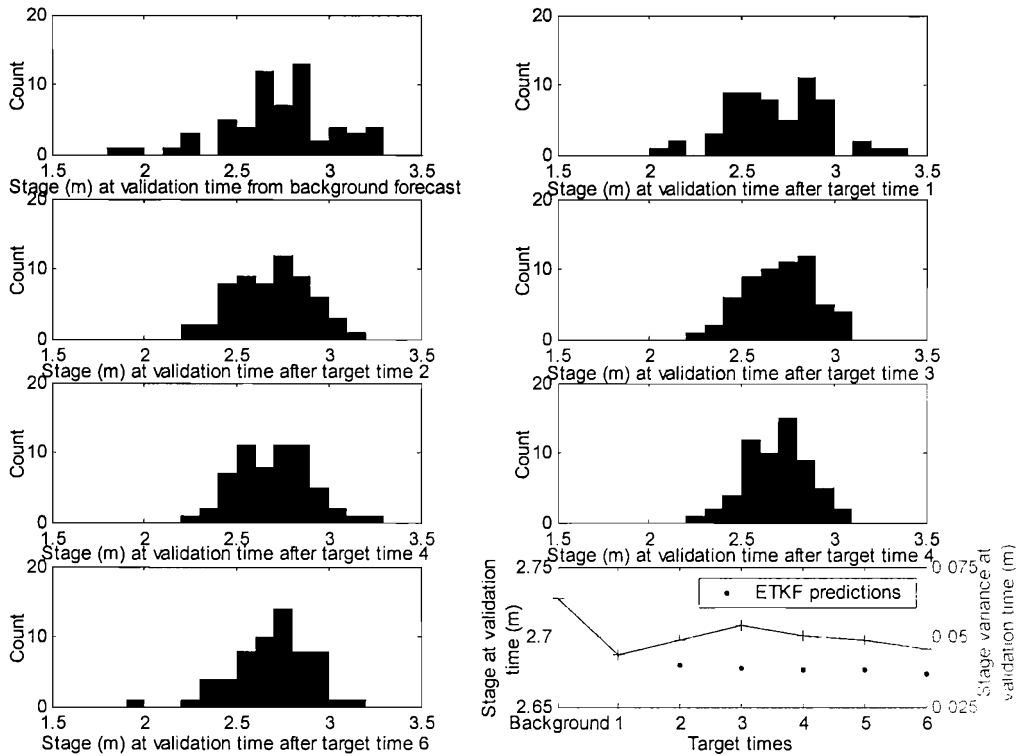


Fig. 7.13: Histogram of background validation time simulated state (top left), followed by histograms of validation time forecast state after assimilating 1 through to 6 measurements from sensor 2 with the EnKF approach. The bottom right plot shows the mean EnKF simulations of state (dark line) and state error variance (lighter line). Also shown are the corresponding ETKF predictions of state error variance.

The ETKF estimates of error variance were less than those made by the EnKF after measurements 2, 3, 4 and 6. As seen previously in chapters 5 and 6 when assimilating measurements at multiple times, stage error variance estimates made using the EnKF approach both increase and decrease after each update, whereas the ETKF always causes error variance to decrease as more measurements are assimilated.

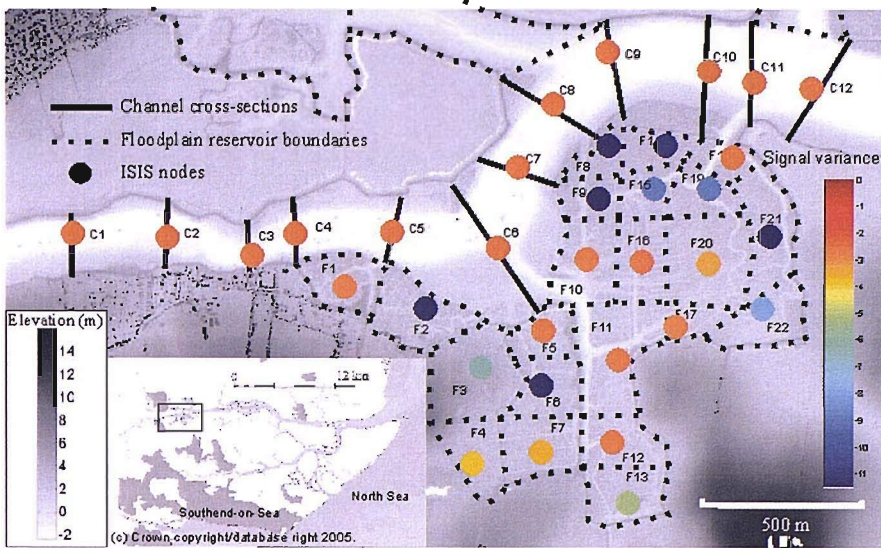
Interestingly, the forecast made after target time 5 was closer to a normal distribution than the other forecasts, where either a positive or negative skew was more evident. Furthermore, the difference between EnKF and ETKF estimates of target time error variance was less after this target time than the other four. This supports the

conclusions in chapters 5 and 6 that the majority of the difference between EnKF and ETKF estimates of error variance was due to RC model instability.

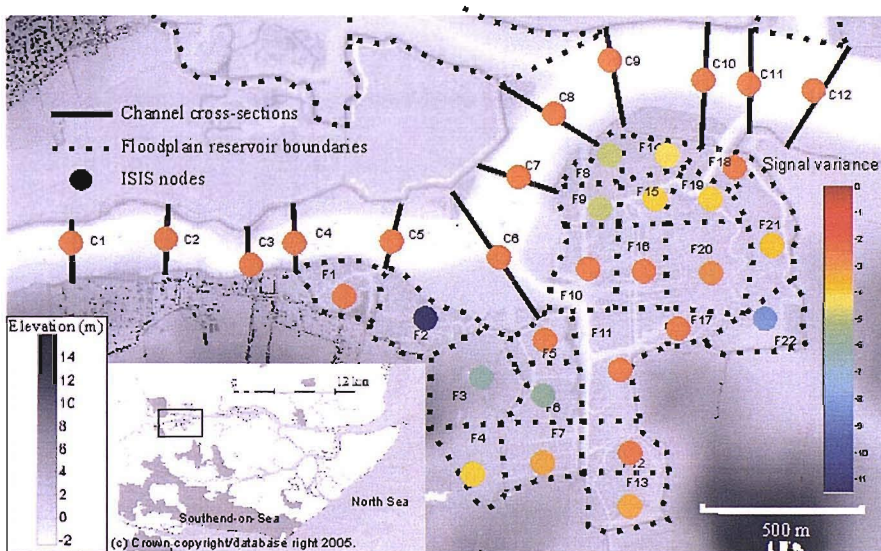
7.3 Signal variance at the study site

In this section, areas of floodplain were considered as possible locations for sensors. Figs 7.14.a-f track changes in signal variance over time, for hypothetical measurements at RC model nodes in and around the Brandy Hole study site. RC model nodes were represented as coloured dots, overlaying a topographic image, located according to their BNG references and labelled with either a C or F indicating a channel section or floodplain reservoir, respectively. Channel nodes are labelled 1-12, upstream to downstream, with floodplain nodes labelled 1 to 22. Dot colour represents validation time signal variance, with greater signal variances having colours towards the red end of the spectrum. In some cases, RC model nodes were dry and gave zero signal variance because the ensemble predicted zero state variance at these nodes. Each image in Fig.7.14 is 20 minutes apart starting at 3 hours 40 minutes prior to the validation of tide 26. For this experiment, the validation location was set at sensor 8 on the floodplain (see Fig. 6.3). This location corresponds with node F17. The background state simulations used to calculate ensemble state perturbations were the same as those in the previous section. Thus, the background forecasts in Fig. 7.12 correspond with the location of node C6. Note that that the coloured dots on Fig 7.14.a and 7.14.b are displayed on a log scale whilst the latter plots are linear.

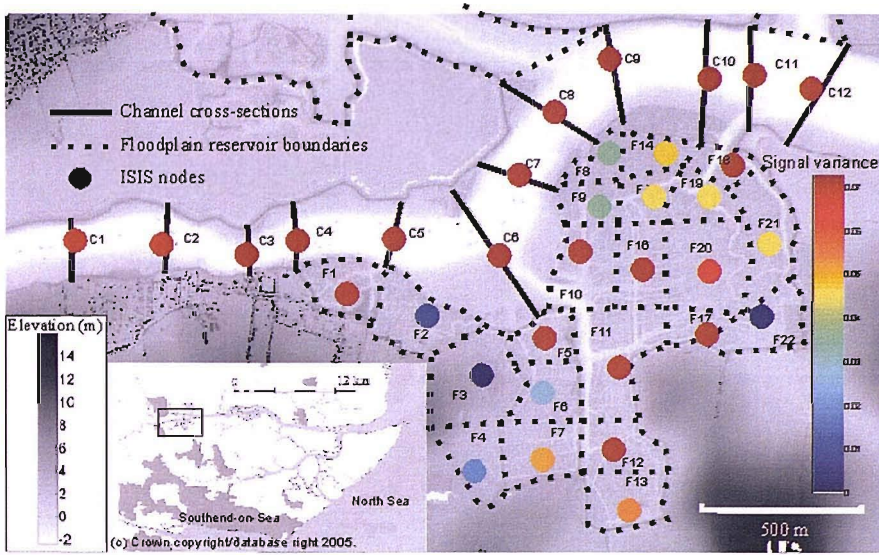
(a)



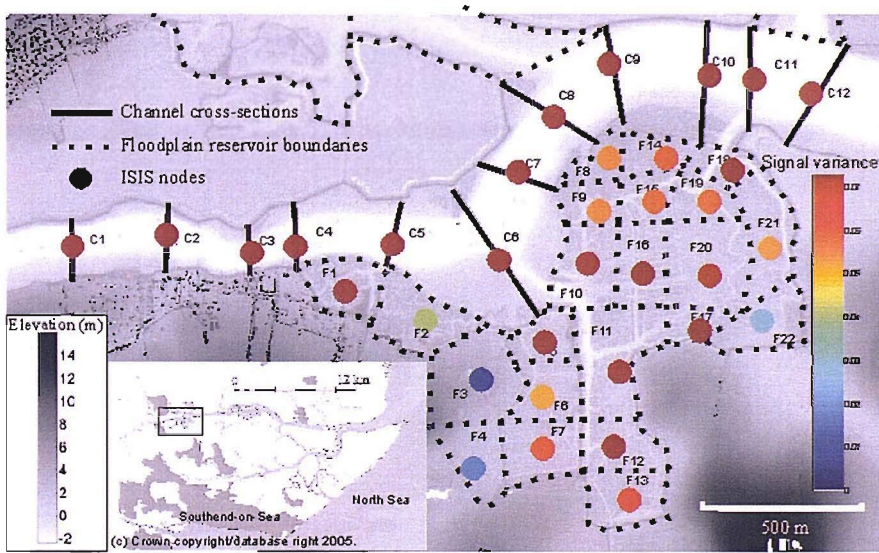
(b)



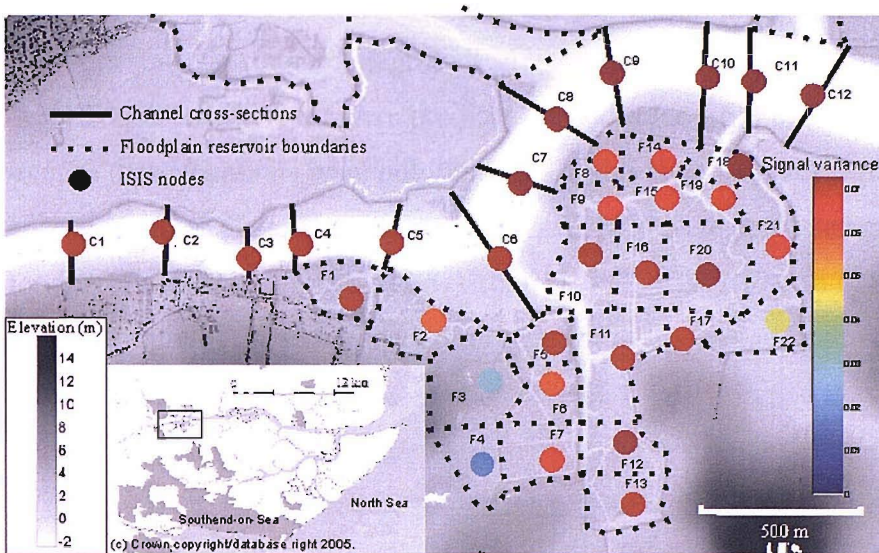
(c)



(d)



(e)



(f)

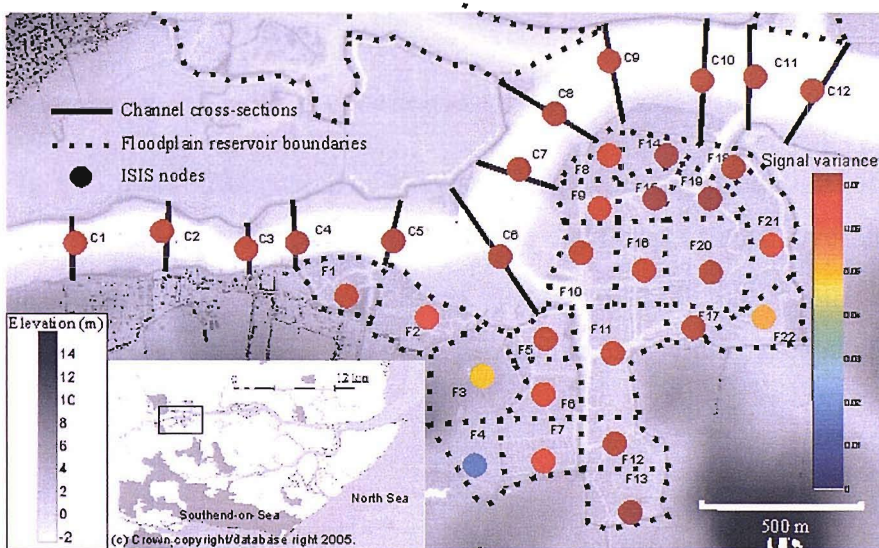


Fig. 7.14.a-f: Signal variance at the validation time and sensor 8 (RC model node F17) during tide 26. Signal variance is shown for theoretical sensors at 22 locations on the floodplain and 12 locations in the channel for six target times. The background image is a DEM derived from LiDAR data. a) Log signal variance 3 hours 40 minutes prior to the validation time. b) Log signal variance 3 hours 20 minutes prior to the validation time. c) Signal variance 3 hours prior to the validation time. d) Signal variance 2 hours 40 minutes prior to the validation time. e) Signal variance 2 hours 20 minutes prior to the validation time. f) Signal variance 2 hours prior to the validation time.

The channel nodes at the downstream end of the study site had the greatest signal variances. Relative to the channel nodes the floodplain nodes had a range of signal variances several orders of magnitude greater. In Fig. 7.14.1 signal variances of less than 10^{-11} essentially zero were predicted at nodes F2, F6, F8, F9, F14 and F21 because these locations were not inundated at this time. Weaker signal variances than those seen in the channel of between 10^{-2} and 10^{-10} were predicted at the other floodplain nodes, except nodes F5, F10, F11, F12, F16, F17 and F18 where signal variances were of the same order of magnitude as the channel nodes. These nodes are distinct from other floodplain nodes with weaker signals in that they contain deep channels and have topographic connectivity with the channel. As time progresses and the tide came in, as captured by Figs. 7.14.a-f, more floodplain nodes yielded similar signal variances to the channel nodes. By Fig.7.14.f only nodes F3, F4 and F22 on the edge of the floodplain yielded signal variances that differed from the channel nodes by more than 0.02 m. The validation location never produced the largest signal.

Fig 7.15 plots validation time signal variance with the validation location at the upstream sensor location (Table 7.1) over eight target times. Overall, signal variances at all locations were lower than those on the floodplain. However, the pattern of signals was similar with channel nodes yielding greater signals than floodplain nodes.

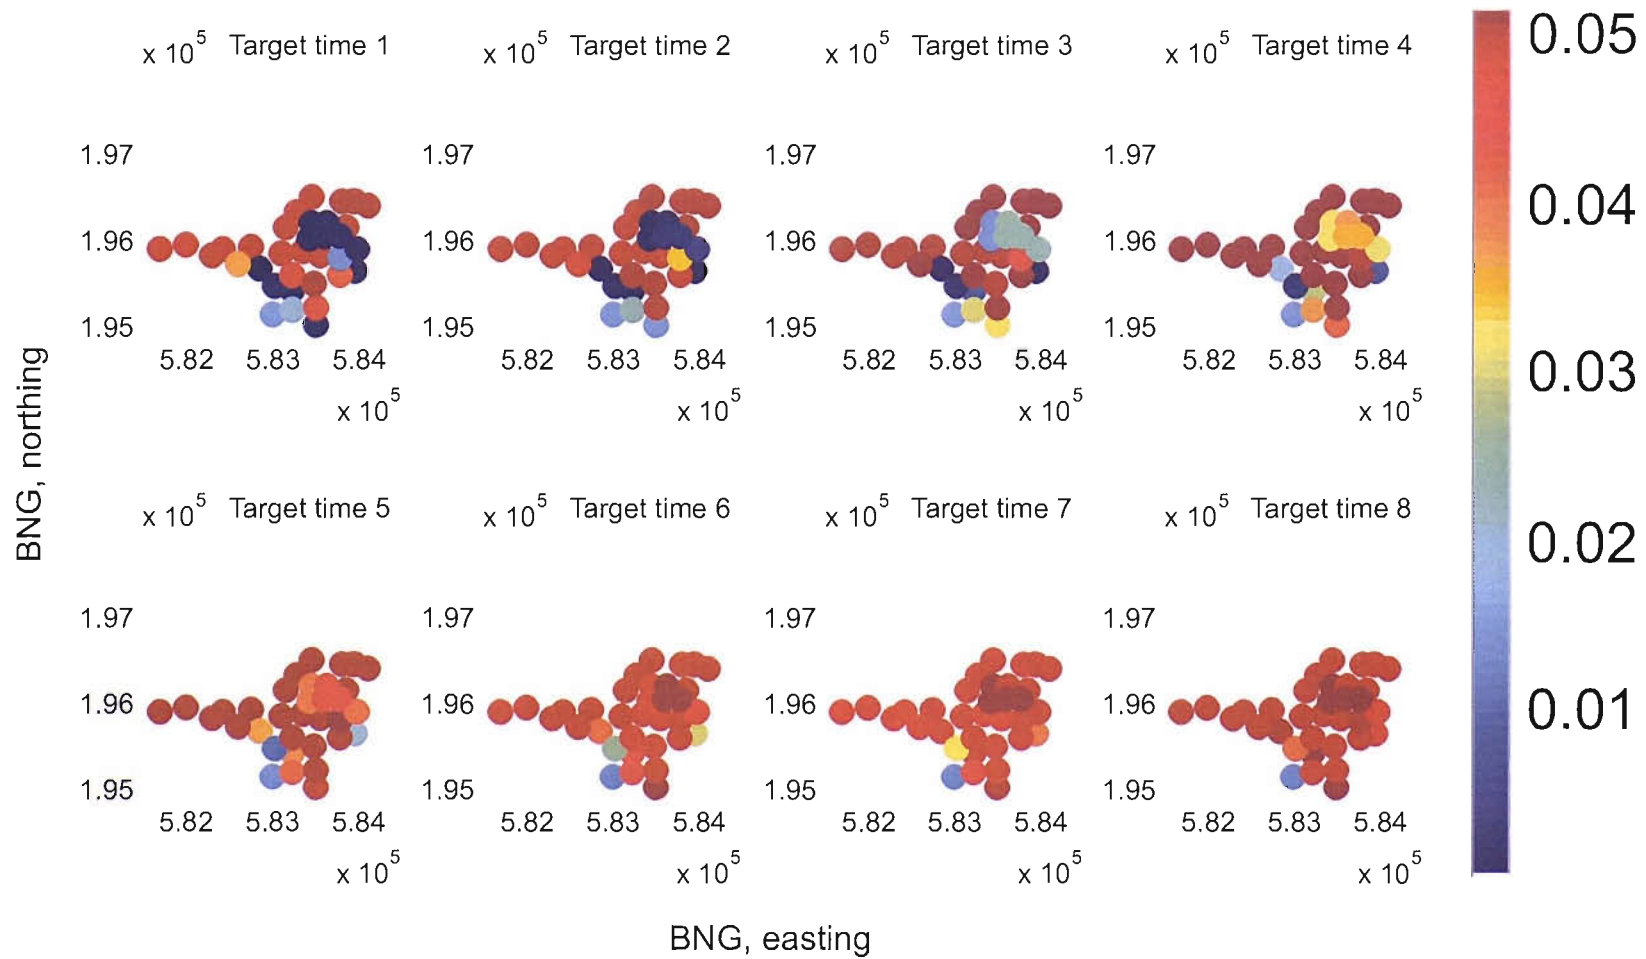


Fig. 7.15: Signal variance (coloured dots) at the validation time and upstream validation location during tide 26. Signal variance is shown for theoretical sensors at 22 locations on the floodplain and 12 locations in the channel for six target times.

To summarise, floodplain nodes that were dry provided no signal, as would be expected. As a floodplain node was inundated signal variance progressively rose towards a value that was similar to the nearest channel node. Thus, as the floodplain filled up the covariance between floodplain nodes and channel nodes increased, causing sensors at these locations to yield similar signals.

7.4 The influence of measurement variance on signal variance

The previous sections explored the links between possible sensor location and the signal variance at a validation location using an ETKF. In this section, the affect of measurement variance on this relationship will be examined with the same algorithm. The setup of the RC model and the ensembles of state perturbations are identical to those used in section 7.2, with the exception that measurement variance was varied between 10 m and 0.0001 m. This range is intended to cover the feasible range of measurement variance which might be expected from a variety of sensors. The measurement variance values tested were: 10 m, 5 m, 1 m, 0.8 m, 0.5 m, 0.1 m, 0.08 m, 0.05 m, 0.01 m, 0.008 m, 0.005 m, 0.001 m, 0.0008 m, 0.0005 m and 0.0001 m.

Fig. 7.16 compares the scatter of signal variance values when prescribing different variances to measurements. Values are shown for all RC model channel nodes and 39 consecutive tides at each of the nine times when target and validation time signal variance estimates were made in Figs 7.1.a-i. The x -axis of these plots displays the signal variances obtained from measurements with a variance of 1 m. This is then plotted against signal variances obtained using measurement variances of 0.5 m, 0.1 m, 0.05 m, 0.01 m, 0.005 m, 0.001 m, 0.0005 m and 0.0001 m. Peak values of signal variance vary between 0.024 m for a measurement variance of 1 m and 0.12 m for a measurement variance of 0.0001 m. The plots show an inverse relationship between signal variance and measurement variance, with the spread of the point cloud becoming greater with the disparity between measurement variances. Therefore, the affect of measurement variance differs from location to location with signal variance

at some locations being more sensitive in terms of signal variance than other locations.





Fig. 7.16: Plots showing signal variance estimates when assimilating one measurement during any of the 39 consecutive tidal events or at any channel node in the RC model assuming a measurement variance of 1 m on the x-axis against the signal variance assuming the measurement error to be either 0.5, 0.1, 0.05, 0.01, 0.005, 0.001, 0.0005, 0.0001.

In Fig. 7.17a, log measurement variance is plotted against the mean target and validation time signal variance for the three sensor locations described in Table 7.1. In this plot, the signal variance is always greater at the target time than validation time when compared to a sensor at the same location and with the same measurement variance. The maximum mean signal variance of 0.108 m was achieved by the validation location sensor at the target time with a measurement variance of 0.0001 m. The smallest signal variance of 0.0009 m was achieved by the downstream sensor at the validation time with a measurement variance of 10 m.

(a)



(b)




Fig. 7.17.a-b: a) Plot of measurement variance against mean signal variance at the target time and validation time at the validation location. The mean is taken from 39 consecutive tides and all channel nodes in the RC model. b) Plot of measurement variance against mean signal variance at the target time and validation time at the validation location (log scales).

The plot indicates that much of the variation in signal variance as a result of measurement variance occurred between a measurement variance of 1 m and 0.001 m. When measurement variance is above 1 m the signal variance decreased into the lower 5% of the distribution of these results. For variances above 1 m the signal decreases log linearly with increasing variance (Fig. 7.17b). When measurement variance was below 0.001 m the gain in signal variance with increasing measurement accuracy decreased exponentially. Thus, the signal has begun to saturate at this level of measurement accuracy.

7.5 Discussion

Figs 7.9.a-i has shown that the region of the RC model with the greatest stage variance was close to the confluence of fluvial and tidal flows and, thus, migrated upstream with the incoming tide. The increase in stage variance corresponded with increasing distance from the downstream boundary condition, along with decreasing channel width and depth (Fig.7.2). The variance increase lead to RC model nodes upstream of the validation location being more sensitive to measurement data than those downstream of the validation location, to the extent that target time signal variance was greatest upstream of the validation location rather than at the validation location. Note that because this is a tidal estuary upstream of the validation location was down flow from the validation location. Since roughness was the only internal RC model parameter to which errors were added (see chapter 5), it seems appropriate to assume that at least some of this variance increase was due to the increasing influence of this parameter on stage as distance from the boundary condition increased and the channel cross-sectional area decreased. Consequently, the propagation of errors was partially dependent on the RC model parameterisation

decisions made in chapter 4, which were based on measurement data from the study site rather than measurements distributed over the RC model domain. It was not possible to test if the stage error variance dynamics were appropriate at node locations, other than those near the study site (e.g. the analysis conducted in section 4.4.3), due to a lack of measurements distributed along the river with which to validate the way in which spatio-temporal errors evolved within the model domain.

Regardless of the validity of the RC model state ensembles, the ETKF was used to identify optimal sampling location away from the study site. Strictly speaking, adaptive sampling or measurement targeting can only be applied with confidence when the validity of the state ensembles is known. Further research requires more distributed measurement data to validate the RC model ensemble on which the ETKF depends. If the ETKF is used to target measurements to previously unmonitored locations, where prior model calibration is not possible, a detailed understanding of error sources and propagation from auxiliary data will be required.

In section 7.2.2, the validation location was placed at the upstream sensor location, as defined in Table 7.1. At target time one the incoming tide had not reached the validation location, meaning almost no target time signal variance was available from locations downstream of the validation location. Nevertheless, measurements from locations downstream of the validation location did yield validation time signal variance. This demonstrated an important capability of the ETKF in that it made an estimate of validation time signal variance prior to the validation location being inundated, whilst recognising that the target time signal variance will be low. Thus, the ETKF can be applied when errors propagate over space and when the state anomaly of interest was spatially separated from the validation location. This would be crucial in an operational river flood forecasting context where it is desirable to have as long a lead time as possible between observations of the flood wave and the arrival of the flood wave at the validation location. One aspect of a likely operational scenario, such as the river Rhine example in Sprokkereef (2001), which could not be simulated by the RC model, was the distance between the region of flow generation and the validation location. Therefore, it is worth testing the ETKF on a set of ensembles generated on a much larger river.

When the validation location was on the edge of the incoming tide in Fig. 7.9.b, only measurements at the edge of the incoming tide gave large signals. If the objective of the data assimilation process is to map the flood edge; measurements at the flood edge will be of most use. However, they will be less useful in a forecasting context where measurements up-flow of the validation location resulted in larger signals.

In this thesis, stage measurements were taken by pressure transducers linked to a digital sensor node. The sensor nodes collected time series of measurements at a fixed location. Therefore, for this type of sensor it is imperative that an adaptive sampling algorithm can account for a time series of data as well as the simpler single measurement scenario. Section 7.2.3 used the ETKF to assimilate a time series of measurements from a single location and then target an additional measurement.

The results indicated that target time and validation time signal variance was an order of magnitude or two less than the single measurement example. Thus, the value of an additional measurement was predicted to be small relative to the previous sections. Overall the ETKF was able to deal with the time series of data with no evidence of numerical instability in estimates of signal or stage variance. However, after assimilating the time series of measurements from the validation location sensor validation time signal variance was expected to be greatest upstream of the validation location in Figs 7.11.d-h. Although, the signal variances are at least an order of magnitude less than those after assimilating one measurement and backwatering from upstream of the validation location may have some effect on validation time stage. It is difficult to justify signal variance from measurements upstream of the validation being greater than those at the downstream boundary based on the findings in section 7.3.2. Therefore, this anomaly may represent a limitation of using a covariance model to target measurements when the principal dimensions of the covariance matrix have been reduced through data assimilation. It may be necessary to incorporate additional rules into the targeting algorithm such that spurious weak correlations that have no physical justification do not influence measurement targeting. Alternatively, the first occurrence of validation time signal variance being greatest over a region of the model domain known to have no physically justifiable

association with the validation location may indicate a reasonable time to stop targeting measurements with this method.

EnKF and ETKF validation time stage error variance estimates were similar after each measurement in the time series was assimilated. Although, the ETKF tended to underestimate the EnKF, the ETKF would seem to have reasonable predictive ability as supported by chapters 5 and 6. As in chapters 5 and 6 RC model instability was a significant factor in the differences between the two algorithms, making it difficult to assess the linear error propagation assumptions made by the ETKF when calculating the transform matrix.

In section 7.3, the validation time signal variances from measurements on the floodplain at the Brandy Hole study site were calculated. Signal variance increased as the tide came in and as the connectivity between the channel and floodplain increased. Signal was always greater at channel nodes even though the validation location was on the floodplain. Therefore, at this study site, when the validation time was at least 2 hours prior to the validation time stage measurements in the channel and towards the RC model downstream boundary yielded greater validation time signal variance than measurements on the floodplain.

If the validation location is moved to the upstream sensor location channel nodes continue to yield the greatest signal variance (Fig. 7.15), indicating that uncertainty in floodplain level has little influence on down flow stage forecasts at this site. However, this may not be representative of other scenarios, especially as the area of floodplain at the study site (approximately 2 km²) is not especially large relative to the length of the river (21 km). For situations where floodplain storage is much greater and has a significant influence on down flow stage, floodplain measurements may become more useful, particularly as state variance is generally greater on the floodplain and at the flood edge than in the channel.

The ETKF could deal with a realistic range of measurement variances and, thus, has the potential to deal with alternative measurement types, as shown by Fig. 7.17a-b. However, errors in measurement data were assumed to be approximated well by a

normal distribution, which may not be the case for measurements errors from certain sensor types. No attempt was made to propose a framework or conduct an analysis of sensor cost verses signal variance. However, the ETKF may provide a means of quickly deriving the signal from a number of different sensor deployment options. For example, are three sensors with variances of 0.005 m likely to yield a greater validation time signal variance than two sensors with variance of 0.001 m, and what are the optimal sensor locations in each case?

7.6 Conclusions

The ETKF could be used estimate target time and validation time signal variance at previously unmonitored locations at any RC model node. However, confidence in the targeting was partly dependent on the accuracy of the RC model simulation of error variance. As would be expected, when the validation time was not at the target time, signal variance was greatest from measurements up-flow of the validation location. Therefore, measurements from these locations were of more use for the purpose of flood forecasting than measurements at the validation location.

Target time signal variance was greatest at or close to the validation location depending on the structure of the covariance matrix. Covariance between RC model nodes was greatest towards the downstream model boundary where the river was deeper and wider, and where the wave is strongest. Water surface slope and the influence of the uncertain roughness parameter was also less in this region of the model domain during the test period, although, more distributed validation data is required to calibrate and validate variance estimates.

Target locations could be selected at target times when the validation location was unaffected by the anomaly of interest. Thus, the ETKF could deal with flow dependent errors, suggesting that the ETKF can be applied to rivers where the region of flow generation is some distance from the validation location. However, as the test conducted using the RC model covered a river only 21 km long, further work is required, on a larger river, to assess the suitability of error covariance as a means of

targeting measurements and the impact of the linear error propagation assumed by the ETKF, in likely operational scenarios.

As in chapters 4, 5 and 6 RC model instability prevented a more comprehensive comparison of the EnKF and ETKF error variance estimates.

At the Brandy Hole field site measurements of floodplain stage were less useful than measurements of channel stage for forecasting error variance at future validation times, even when validation location was on the floodplain. However, this may not be the case in locations where errors in floodplain stage have a significant influence on down flow stage forecasts.

The ETKF could consider a range of measurement error variance values, meaning that the signal variance of measurements from different sensor hardware could be considered as part of the sampling strategy, provided that the error statistics of the sensor measurements can be approximated well by a normal distribution.

7.7 Summary

The location, timing and accuracy of measurements are critical to the performance of data assimilation algorithms such as the Kalman filter. Near-optimal sequential sampling frameworks can be designed for non-linear models using an ensemble transform Kalman filter, provided a linear transform of error perturbations is reasonable. Here, an ensemble transform Kalman filter was applied to an ensemble of simulations from a one-dimensional hydrodynamic model of a 21 km reach of the River Crouch, Essex, UK. This model-based targeting method was able to estimate the ‘signal variance’ of theoretical measurements at any computational node in the hydrodynamic model. Time series data, different sensor types and measurements of floodplain stage could all be taken into account either as part of the targeting process or prior to measurement targeting. The River Crouch provided a useful test scenario. However, it is significantly shorter in physical length than a likely operation scenario, meaning that further tests are required at a different study site. Confidence in measurement targeting was partly dependent on the accuracy of the hydrodynamic model ensemble’s simulation of error variance, which could not be validated due to a lack of measurement data.

8 Discussion

8.1 Updating hydrodynamic model state

Increasingly, stochastic predictions of state variables and the delivery of probabilistic rather than deterministic predictions have been sought by environmental managers interested in hydrological forecasting (Krzysztofowicz, 2001). It follows that in situations where the latest measurement data is used to update simulation models to increase simulation accuracy in the near future, a similar product would be desirable to decision makers. The KF provides an optimal sequential data assimilation framework by minimising the mean squared error between measurement data and the simulated state of a system. In chapter 4 a Monte Carlo version of the KF known as the EnKF was used to update the state (stage, flow) and boundary condition error of a hydrodynamic model of the River Crouch (RC model). The EnKF analysis was computationally quick (<1 second) to implement on the 2.4 GHz Pentium 4 processor and posed no problems in terms of memory requirements. However, there were only 323 states in the RC model state vector, 60 ensemble members and up to 4 measurements at any single update time, making this a small problem relative to those commonly found in meteorology and oceanography. The propagation of the ensemble to the next measurement time step using the RC model was by far the most demanding computational task required by the RC forecasting model. However, running multiple model simulations, of identical length, can be easily posed in a grid computing framework.

In the RC model example the EnKF was an efficient means of reducing state error variance with measurement data. However, the EnKF update caused state instabilities in some ensemble members, which generated unwanted perturbations in simulated state. These instabilities did not occur when the RC model was run without the update. More efficient 'square root' versions of the EnKF which do not introduce error through the use of perturbed measurements (Anderson 2001; Whitaker and Hamill, 2002; Tippett, 2003; and Evensen 2004) have been developed and should

supersede the Evensen (2003) version of the filter used here. These square root filters may help to reduce post update model instability by not introducing errors into the water surface from perturbed measurements.

Mass was not conserved by the EnKF update of RC model states. This implies that the filter can adjust the volume of water within the model domain in addition to its distribution. This seems a reasonable assumption in the case of the RC model given that the volume of water entering the model domain is one of the principal sources of state uncertainty. However, when the coverage of measurements over the model domain is far from comprehensive (such as the case with the FloodNet sensor nodes) but, the flows into and out of the model domain are well known (such as the situation where the upstream and downstream boundaries of the model domain are gauged) a redistribution of water may be more appropriate than a change in the overall volume. In other words, if the principal source of state uncertainty is the process representation within the model rather than the flows into and out of the model then changing the volume of water stored becomes less relevant. An update that conserves the volume of water within the model domain within certain bounds, in affect, adding a constraint on the update would be of potential value. One possible approach might be to move water into and out of floodplain storage, although this would pose implementation issues within the statistical constraint of the data assimilation tools available. However, the accuracy of such an update could be easily assessed by comparison with the EnKF based forecasting model.

Flood inundation applications that use two-dimensional simulation models or cover much larger spatial domains may raise some additional issues not covered by this thesis. For instance, in atmospheric applications of ensemble data assimilation methods it is often necessary to limit the impact of observations to a subset of model state variables (Anderson, 2004). This often entails only allowing state variables to be influenced by measurements that are physically close, such that errors are not introduced by spurious linear relations between distant states. A similar approach may be necessary for river simulation models that cover larger domains or simulate state variables at finer resolution than the RC model. It may also be necessary to

introduce rules that prevent measurements on one tributary influencing state variable on another.

There are many types of simulation model that could be used to simulate river stage in place of the distributed hydrodynamic model used here, these including statistical or data based models such as transfer functions (Lees *et al.*, 1994), database mechanistic model (Young, 2002; 2003) and neural networks. These data driven methods are computationally less demanding than the hydrodynamic model used here, have fewer state variables, have been shown to perform well where sufficient data are available and are not prone to the same state instabilities. Furthermore, they do not require the costly topographic and land cover data required by distributed hydrodynamic models (see section 2.1.1). Thus, in situations where forecasts are required at locations where sufficient data are available there is a compelling case of using these approaches. However, these methods only simulate states at measurement points, thereby not directly estimating flood inundation; meaning flood inundation must be inferred by comparing forecasted state with offline simulations of flood inundation. Other advantages to using the distributed process based approach used in this study include:

1. The simulation of state dynamics is governed by a physical process representation, which in theory can provide a means of extrapolating to previously unseen situations with greater reliability than a data driven method. However, the parameter and structural uncertainties outlined in section 2.1.2 that are inherent in the process of simplifying reality will introduce errors.
2. Historical records of the system's state dynamics are not required to build the model and as such a monitoring program need not be in place prior to model building. However, prior state data would be beneficial for calibrating model parameters and necessary if model simulations are to be validated.
3. State is simulated at unmonitored locations and can be updated as part of the forecasting process.

4. Changes to the geomorphology of the model domain (e.g. dredging, meander cut-off etc.) can be accounted for explicitly by adjusting the process model. The propagation errors, error sources and parameter values may all be affected by change in the geomorphology of the model domain.

5. The combination of both a distributed process model and KF allows measurements to be incorporated on an *ad-hoc* basis from previously unmonitored locations where model state estimates are made (see chapter 7 for a more detailed discussion). Thus, data from opportunistic or infrequent sources can be incorporated into forecasts as and when they are available. A consideration of measurement error also allows multiple sensor types and noise corrupted data to be assimilated.

The final point listed led to the research conducted on the collection of *ad-hoc* data via an adaptive sampling methodology presented in chapters 5 and 6, and 7 and will be returned to later in this discussion.

8.1.1 Error sources and parameter uncertainty

In chapter 4, the autoregressive model at the downstream boundary was the only source of model error modelled, thus, contributions from three error sources, namely the model structure, parameters and upstream flow boundary were neglected. The coefficients of this model were manually calibrated such that the error forecasts made by the RC model ensemble were consistent with the errors observed between forecasts and measurements from the key data gathering source (sensor 1). However, the ensemble error statistics were not suitable at a location on the floodplain. Thus, the model had been tuned to a single measurement location and as a consequence did not have the same predictive ability at other locations or a framework for introducing errors between sensor locations. In chapter 5, roughness uncertainty was incorporated into the ensemble simulations using GLUE. However, the approach used here (as opposed to the methodology in general) was inconsistent with the data assimilation process not least because measurement errors were not considered and the parameter uncertainty relevant to a short forecast may have been different to those that dominated over the calibration period. It may be beneficial for the

calibration of model parameters to be directly integrated with the KF, such as seen in Vrugt *et al.* (2005) where an ensemble Kalman filter was integrated with a global parameter estimation scheme, although this was not carried out within the scope of this study. As a result further research was required into the identification of the sources of model error and their quantification.

8.2 Adaptive sampling

Chapter 6 compared two adaptive sampling methods to inform the real-time measurement reporting of a network of wireless sensor nodes, with limited power. These were called Method 1 and Method 2. Method 1 used a systematic sampling framework with a stopping criteria based on the probability of validation location stage exceeding a threshold over a validation period. The KF's stochastic framework is well suited to this type of adaptive sampling strategy as it allows forecasts to be posed in a probabilistic context with little additional effort. To convert state forecasts to probabilistic forecasts managers and stakeholders, who might use the forecasts produced by a KF based forecasting model, need only decide which region(s) of model state space are of interest, the value of the threshold and the period of time over which they wish to forecast. The sequential nature of the KF also fits well with a system that regularly collects measurements until some stopping criteria is met.

For the River Crouch case study no criteria were devised to decide when to restart real-time reporting after the decision to stop. Instead, each tide was treated in isolation from previous tides, such that, reporting always began four hours prior to the validation time (high tide). In a river system where the tide does not dominate the flow pattern the regularity of stage fluctuations required by this approach would not be suitable. One alternative might be to periodically issue longer forecasts after the decision to stop reporting and use them to decide when to restart reporting measurements in real-time. This would require a method to decide when measurements at a particular target time would first become relevant to the time at which the probability of exceeding the threshold becomes significant, such as, Method 2.

Adaptive sampling Method 2 targeted measurements based on their expected signal variance, quickly calculated by transforming ensemble state perturbations, using an ETKF. For the RC model case study a minimum forecast lead time was also required because the signal variance expected from sensors at the study site increased as time to the validation time decreased. Method 2 was more efficient in terms of signal variance yield per measurement collected. However, unlike Method 1, the stopping criteria could not be probabilistic, since the sampling decisions were made before data collection. Therefore, Method 2 was best suited to targeting situations where a flexible but finite sampling resource exists or when sampling decisions need to be made before the ‘event’ of interest begins. For example, situations where sensors are mobile, directional or it is expensive to measure. For the River Crouch case study the ETKF indicated that a few samples (<5 see section 6.3.2) close to the time at which forecasts were issued could be almost as effective in terms of signal variance as collecting data systematically over eight target times because signal had saturated.

One problem with both the above adaptive sampling methods is that state dynamics not anticipated by the simulation model may be missed. This could also introduce bias into the measurement selection and consequently bias in simulated state. In the context of wireless sensor nodes this problem could be solved by engineering the sensor nodes such that real-time reports are sent at regular intervals, hence avoiding the need to adapt the sensor node reporting in real-time. However, this may be undesirable when, for example, sensors are battery powered or the network has other tasks to perform. Since each FloodNet sensor node was capable of local processing it may be worth seeking further innovations that allow sensors to perform simple operations, which could influence real-time reporting, such as spotting state anomalies. Alternatively, computationally ‘lightweight’ forecasting models such as those discussed in section 8.1 could be deployed on the sensor network, as discussed by Hughes *et al.*, (2006).

8.2.1 Adaptive sampling with the ETKF

In chapter 5, the ETKF was used to rapidly obtain estimates of system model error covariance at measurement times and validation times using estimates of measurement error covariance. Due to its speed of execution it was possible to test many possibly measurement sets and build up near-optimal sampling strategies using serial assimilation. Given that this method was originally applied to a meteorological application (Majumdar, 2002) it seems likely that its computational burden will not pose a significant barrier to its use in flood forecasting, where the size of the model state vector and number of measurements tend to be less. The ETKF was easy to implement provided suitable state ensembles and estimates of error covariance were available, as all the information required by the filter is contained in the state ensemble.

The EnKF was found to be strongly dependent on the error statistics generated from the state ensemble, supporting the conclusion made by Refsgaard (1997) that KF-based data assimilation is strongly dependent on simulation model accuracy. This conclusion was equally applicable to the ETKF, meaning that the accuracy of the RC model was integral to the performance of both the forecasting model and the adaptive sampling. Instability in the RC model after data assimilation did not effect the ETKF forecasts of error or signal variance, unless the background forecasts used contained instabilities relating to measurements previously assimilated by the EnKF. However, the comparisons of EnKF and ETKF forecast error variances made in chapter 5 were affected by these instabilities to the extent that ETKF error variance predictions were generally lower than those based on the EnKF. This made it difficult to assess the disadvantages of using the linear transform of ensemble state perturbations at validation times rather than the RC simulation model. Therefore, simulation stability after data assimilation would be of utmost importance to further research with the EnKF and ETKF with the RC model. Nevertheless, the ETKF made reasonable estimates of forecast error variance at this site despite the presence of instabilities in some of the RC model simulations.

The performance of the ETKF as a targeting tool and specifically the suitability of using a linear transform of error perturbations on longer or non-tidal rivers are not well understood. Further research is required to assess the applicability of the ETKF in situations when the distance between validation location and target location is greater than 21 km, as it is possible that error propagation will be more complex as distance increases. The influence of tributaries entering the river between the target location and validation location on error perturbations is also unknown, as is the effect of areas of floodplain adjacent to the river channel that are larger than those on the River Crouch. Furthermore, the RC model does not incorporate inputs from baseflow or outputs due to infiltration or evaporation, which were believed to be insignificant on the River Crouch relative to tidal forcing.

It should be possible to generate state ensembles prior to the deployment of any monitoring equipment in most cases. For example, in a situation where boundary conditions are generated from uncertain rainfall fields and translated to river flows by a hydrological model. Therefore, it may be possible to use the ETKF as a means of designing a monitoring network prior to its implementation, especially as different sensor types can be considered (see section 7.4 and subsequent discussion). Alternatively, an overly dense monitoring network could be deployed during an initial calibration phase, before being reduced in size after careful consideration of the signal variance and data redundancy.

8.3 Measurement data issues

In chapter 4, measurement bias was found to be a significant problem. The pressure transducers used by the FloodNet sensor nodes are designed to self calibrate to atmospheric pressure in order to prevent sensor drift. However, errors in measurements of sensor height (differential global positioning system readings in this case) will be constant over time. Measurement bias was identified in chapter 4 due to the close proximity of other sensors, which it was believed should have yielded similar water levels. However, it may be more difficult to distinguish measurement bias from system model bias in situations where the spatial density of

measurements is less. Nevertheless, methods for using the KF to identify measurement and system model bias have been developed.

As introduced in section 2.1.1 the development of LiDAR technology has allowed floodplain topography to be mapped at finer spatial resolutions than the grid, mesh or node spacing of hydrodynamic models (Marks and Bates, 2000; Cobby *et al.*, 2001; French, 2003), whilst techniques for flood edge delineation have been developed using synthetic aperture radar and optical remote sensing imagery (Hunter *et al.*, 2005; Bates *et al.*, 2006). It follows that it should be possible to assimilate estimates of stage at the flood edge derived from remotely sensed data and a fine spatial resolution DEM, where the intersection of flood edge and DEM leads indirectly to a stage measurement. This is especially true given that in section 8.1 one of the benefits of using a hydrodynamic model as a component of an EnKF based forecasting model was argued to be an ability to assimilate measurements collected on an *ad-hoc* basis from previously unmonitored locations. The benefits of such measurements are not well understood and would probably depend on the sensitivity of validation location error variance to error variance at the flood edge and how accurately the flood edge could be delineated. The accuracy of flood edge delineation is likely to depend on the accuracy of the DEM, the spatial resolution of the image, floodplain topography (specifically the slope of the floodplain relative to the range of possible stage values or range of flood inundation estimates) and above surface features on the floodplain such as vegetation and buildings. A further problem that would need to be addressed might include the risk of introducing bias into state estimates and correlation between errors in multiple estimates of stage from a single image.

Before the above is attempted, further research is required into the benefit of assimilating flood inundation data into simulation models of a number of study sites and at a range of different scales. Especially, given that at the River Crouch site:

1. Stage measurements at the flood edge were found to yield less signal variance than those in the channel.

2. Floodplain measurements were found to yield very small signals when considered in addition to channel stage measurements, despite stage error variance being greater towards the floodplain edge (see discussion in chapter 7).
3. Error in floodplain stage had little influence on down flow error propagation.

Nevertheless, these properties may not be repeated on alternative rivers with much larger areas of floodplain storage relative to the channel, where present storage has an impact on downstream stage errors.

With the growing availability of multi-temporal satellite and airborne data from which flood inundation can be estimated (Smith, 1997) it may be possible to use adaptive sampling techniques to target or design image acquisition during flood events to optimise signal variance, especially given the ongoing research concerning the assimilation of swath-altimetry into hydrodynamic models (Andreadis *et al.*, 2007). In theory, the ETKF could be used to estimate the signal variance expected from assimilating stage estimates made from a remotely sensed image prior to its acquisition, provided that a linear transform of error perturbations is appropriate and the errors in measurements are close to normally distributed. Given these assumptions it would also be possible to use the ETKF to assess the likely benefit in terms of signal variance of building a forecasting model to assimilate remotely sensed flood inundation or any other data source prior to building the forecasting model. Obviously, an ensemble of model simulations would be required, but, the EnKF would not need to be implemented, avoiding the associated risk or model instability. In fact the ETKF could be used to optimise a sensor network deployment, the measurements from which could then be used by a forecasting model not based on a hydrodynamic model. However, it may be difficult to justify the simulated error perturbations used by such an approach without calibration or validation data.

9 Conclusions

In section 1.1 the aims of this thesis were set out along with some more specific aims for each of the four research chapters. By considering the results from each of the four research chapters some concluding remarks will now be made in reference to each of these thesis aims:

1. Build a forecasting model based on:

- a. **A simulation model capable of simulating river states at unmonitored locations.**
- b. **A data assimilation algorithm that can utilise real-time measurement data that are spatiotemporally discontinuous and corrupted by errors.**

The EnKF provided a framework for assimilating stage measurements into stage and flow simulations made by a one-dimensional non-linear hydrodynamic model with a quasi two-dimensional representation of floodplain flow dynamics. In contrast to previously published forecasting models of this type, floodplain storage cells were updated in addition to the channel nodes. Explicitly accounting for areas of floodplain storage using the ISIS storage cell technique allowed measurements of floodplain stage to be assimilated with model states, without having to assume a flat water surface between channel and floodplain. The ability to assimilate measurements of floodplain stage may, in the future, generate opportunities to assimilate measurements derived from remotely sensed data sources, including images of flood inundation. Further simulation model developments might include changing the representation of floodplain processes through the use of more complex floodplain flow equations or including smaller scale floodplain features that could not be accounted for by the resolution of the storage cells used here.

The principal difficulties with the approach used here were integrating the analysed state estimates with the hydrodynamic model and simulating stage and flow uncertainty rather than the estimation problem itself. Thus if the lessons of this case study are applicable in a general sense, the data assimilation tools currently available

would appear more than adequate given the difficulties associated with generating realistic state ensembles based on the present degree of understanding of error propagation in river flow models. However, as discussed in chapter 8 this may not be the case in situations where the simulation model has many more state variables or the ensemble size is smaller because of errors introduced by spurious long distance relationships between state variables and the finite ensemble size.

2. Assess the ability of the forecasting model to simulate state error propagation.

In chapter 4 the parameters that introduced errors into the RC model at the downstream tidal boundary were calibrated by trial and error such that the forecasting model made reasonable estimates of forecast error variance over a 4 hour forecast period. However, the RC forecasting model was unable to make variance forecasts that were as accurate at a location on the floodplain. It was suggested that errors introduced by uncertainty in internal model structure, roughness parameters and measurement bias may have been responsible for this. Uncertainty in the global roughness parameter was subsequently incorporated into the RC forecasting model. However, the close proximity of the sensor nodes made it difficult to identify weaknesses in the model structure or parameterisation since each sensor collected very similar information. A conclusion subsequently supported by the rapid decrease signal variance as measurements from the FloodNet sensors were assimilated in chapter 6.

Forecasting over periods of time longer than 4 hours was limited because this essentially depended on the errors introduced first order autoregressive model at the downstream tidal boundary. An improvement on the experiments conducted here might include running the simulation model on larger river and spreading the sensor nodes out over a larger area.

3. Assess the ability of an ETKF to quickly estimate forecast error variance simulated by the forecasting model.

Estimates of forecast error variance made by an ETKF were compared with those made by the EnKF based RC forecasting model in situations where the measurement data available to the RC forecasting model changed and did not change the ensemble mean state. For both situations the ETKF made estimates of error variance that were correlated significantly with those from the forecasting model. Correlations were higher when the ensemble mean did not change, particularly towards the edge of the incoming tide where forecast uncertainty was highest and most sensitive to changes in ensemble mean. Although, the decrease in correlations when the ensemble mean changed was also partly due to the greater prevalence of model state instability when the ensemble mean changed. The ETKF was therefore a useful tool for quickly estimating forecast error variance after assimilating a given set of measurements. This supports conclusions made previously in an atmospheric sciences context (e.g Bishop *et al.*, 2001; Majumdar *et al.*, 2002; Bishop *et al.*, 2006).

4. Develop and evaluate an adaptive sampling framework that could be applied to a network of wireless sensor nodes with limited power.

Two approaches to adaptive sampling with wireless sensor nodes were implemented in this thesis. The first (Method 1) used the probability of flooding derived from an ensemble forecast to determine if further measurements were required, thereby updating the sampling framework systematically after each forecast was issued. The second (Method 2) used an ETKF to calculate the expected signal variances that would result from Q sampling options. The approach based on an ETKF facilitated the development of a more efficient sampling framework, in terms of signal variance per measurement, than the systematic approach. It also allowed the sampling framework to be planned in advance of any data collection.

Targeting measurements with a simulation model has a number of drawbacks. Numerical process models that inherently ignore or simplify various processes that influence system dynamics may lead to bias into model state simulations. A filter using measurements collected systematically or randomly has a fair chance of correcting such errors, depending on the sampling density and sensitivity of state to the unrepresented process. Targeted measurements preferentially update state errors

based on the setup of the simulation model and therefore may actively under correct for un-modelled processes in accordance with the limitations of the simulation model. This is important since one of the principal reasons for updating is that errors in simulated state often build up over time because of processes not accounted for by the simulation model. Similarly, targeting measurements based on the assumptions made about error sources will lead to sampling bias based on the models prior or assumed representation of error sources and propagation.

5. Develop and test a methodology to design optimal sampling frameworks for flood forecasting.

The ETKF used by adaptive sampling Method 2 can be used to design optimal spatio-temporal sampling frameworks in much the same way that geostatistical methods can be used to design spatial sampling frameworks. Namely sensors can be placed in locations where their measurements yield the greatest signal variance when assimilated with state estimates from a simulation model. However, the scope of the tests carried out at the River Crouch field site was limited by the size of the river network, the size of the floodplain storage, the spatial extent of validation data and a lack of significant inflow from tributaries. Thus, whilst it was possible to demonstrate the methodology, assess its numerical stability, conclude that the quality of the simulation model ensemble is likely to be a key factor and conduct a limited comparison with the EnKF. Anything approaching a comprehensive assessment of the accuracy of error variance estimates made by the ETKF was not possible in this case.

10 References

- Andreadis, K. M., Clark, E. A., Lettenmaier, D. P., Alsdorf, D. E., 2007. Prospects for river discharge and depth estimation through assimilation of swath-altimetry into a raster-based hydrodynamics model. *Geophysical Research Letters*, 34, L10403, doi:10.1029/2007GL029721.
- Anderson, J.L., Anderson, S.L., 1999. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review* 127, 2741-2758.
- Anderson, J.L., 2001. An ensemble adjustment filter for data assimilation. *Monthly Weather Review* 129, 2884-2904.
- Anderson, M.G., Walling, D.E. & Bates, P.D. (Eds.). 1996. *Floodplain Processes*, John Wiley. ISBN: 0471966797
- Applied data systems. BitsyX users manual. URL:
<http://www.applieddata.net/support/> (Accessed: May 2007)
- Aronica G, Hankin B.G., and Beven K.J., 1998. Uncertainty and equifinality in calibrating distributed roughness coefficients in a flood propagation model with limited data. *Advances in Water Resources*, 22 349-365.
- Aronica, G, Bates, P.D., and Horritt, M.S., 2002. Assessing the uncertainty in distributed model predictions using observed binary pattern information within GLUE. *Hydrological Processes*, 16 2001-2016.
- Atkinson, P.M., and Tate, N.J., 2000. Spatial scale problems and geostatistical solutions: A review. *Professional Geographer*, 52 607-632.

- Aubert, D., Loumagne, C., Oudin, L., 2003. Sequential assimilation of soil moisture and streamflow data in a conceptual rainfall-runoff model. *Journal of Hydrology* 280, 145-161.
- Babovic, V., Canizares, R., Jensen, H.R., Klinting, A., 2001. Neural Networks as Routine for Error Updating of Numerical Models. *Journal of Hydraulic Engineering* 127 181-193.
- Bates, P.D., Anderson, M.G., Baird, L., Walling, D.E., and Simm, D., 1992. Modelling floodplain flows using a two-dimensional finite element model. *Earth Surface Processes and Landforms*, 17 575-588.
- Bates, P.D., Anderson, M.G., Hervouet, J-M., and Hawkes, J.C., 1997. Investigating the behaviour of two-dimensional finite element models of compound channel flow. *Earth Surface Processes and Landforms*, 22 3-17.
- Bates, P.D., Horritt, M.S., and Hervouet, J-M., 1998. Investigating two-dimensional, finite element predictions of floodplain inundation using fractal generated topography. *Hydrological Processes*, 12 1257-1277.
- Bates, P.D., and De Roo, A.P.J., 2000. A simple raster-based model for flood inundation simulation. *Journal of Hydrology*, 236 54-77.
- Bates, P.D., Marks, K.J. & Horritt, M.S., 2003. Optimal use of high-resolution topographic data in flood inundation models. *Hydrological Processes*, 17 5237-5257.
- Bates, P.D., Horritt, M.S., Hunter, N.M., Mason, D.C., and Cobby, D.M., 2005. Numerical modelling of floodplain flow. In: Bates, P.D., Lane, S.N., and Ferguson, R.I. (Editors), *Computational fluid dynamics: Applications in environmental hydraulics*. John Wiley & Sons, Chichester pp. 271-304.

- Bates, P.D., Wilson, M.D., Horritt, M.S., Mason, D., Holden, N. & Currie, A., 2006. 'Reach scale floodplain inundation dynamics observed using airborne Synthetic Aperture Radar imagery: data analysis and modelling. *Journal of Hydrology*, 308 306-318.
- Beven, K.J., 1989. Changing ideas in hydrology: the case of physically based models. *Journal of Hydrology*, 105 157-172.
- Beven, K.J., 2000. *Rainfall-Runoff modelling: A primer*. Wiley, Chichester.
- Beven, K.J., and Binley, A., 1992. The Future of distributed models: Model calibration and uncertainty prediction. *Hydrological Processes* 6, 279-298.
- Beven, K.J., and Freer, J., 2001. Equifinality, data assimilation, and uncertainty estimation in mechanistic modelling of complex environmental systems using the GLUE methodology. *Journal of Hydrology*, 249 11-29.
- Bishop, C.H., Etherton, B.J., and Majumdar, S.J., 2001. Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Monthly Weather Review*, 129 420-436.
- Bishop, C.H., Etherton, B.J., and Majumdar, S.J., 2006. Verification region selection and data assimilation for adaptive sampling. *Quarterly Journal of the Royal Meteorological Society*, 132 915-933.
- Box, G.E.P., and Tiao, G.C., 1973. *Bayesian inference in statistical analysis*. Reading, Mass: Addison-Wesley.
- Bras, R.L. and Rodriguez-Iturbe, I., 1985. *Random functions and Hydrology*. Dove, New York.
- Burgers, G., Leeuwen, P.J.V., Evensen, G., 1998. Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review* 126, 1719-1724.

- Burgess, T.M., and Webster, R., 1980. Optimal interpolation and isarithmic mapping of soil properties I: The semi-variogram and punctual kriging. *Journal of Soil Science*, 31 315-331.
- Butts, M.B., Falk, A.K., Hartnack, J., Madsen, H., Klinting, A., Kalken, T. van, Cadman, D., Price, D. 2005. Ensemble-based methods for data assimilation and uncertainty estimation in the FLOOD RELIEF project. Proceedings of the International Conference “innovation, advances and implementation of flood forecasting technology, Tromsø, Norway, 17-19 October, 2005.
- Butts, M.B., Payne, J.T., Kristensen, M., and Madsen, H., 2004. An evaluation of the impact of model structure on hydrological modelling uncertainty for streamflow simulation. *Journal of Hydrology*, 298 242-266.
- Cahill, A.T., Ungaro, F., Parlange, M.P., Mata, M., and Nielsen, D.R., 1999. Combined spatial and Kalman filter estimation of optimal soil hydraulic properties. *Water Resources Research*, 35 1079-1088.
- Cameron, D., Kneale, P., and See, L., 2002. An evaluation of a traditional and neural net modelling approach to flood forecasting for an upland catchment. *Hydrological Processes*, 16 1033-1046.
- Canizares, R., Heemink, A.W., Vested, H.J., 1998. Applications of advanced data assimilation methods for the initialisation of storm surge models. *Journal of Hydraulic Research* 36 655-674.
- 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.
- Carsell, K.M., Pingel, N.D., and Ford, D.T., 2004. Quantifying the Benefit of a Flood Warning System. *Naval Research Review*, 5 131-140.

Chow, V.T., 1988. Open channel Hydraulics. McGraw-Hill, New York.

Cobby, D.M., Mason, D.C., Davenport, I.J., 2001. Image processing of airborne scanning laser altimetry data for improved river food modelling. *ISPRS Journal of Photogrammetry & Remote Sensing*, 56 121-138.

Cobby, D.M., Mason, D.C., Horritt, M.S., and Bates, P.D., 2003. Two-dimensional hydraulic flood modeling using a finite-element mesh decomposed according to vegetation and topographic features derived from airborne scanning laser altimetry. *Hydrological Processes*, 17 1979-2000.

Connell, R., Painter, D.J., and Beffa, C., 2001. Two-dimensional flood plain flow. II: model Validation. *Journal of Hydrologic Engineering*, 6 406-415.

Cunge, J. A., 1980. Practical aspects of computational river hydraulics. Pitman publishing program, London.

Cunge, J. A., 2003. Of data and models. *Journal of Hydroinformatics* 5 75-98.

Curran, P.J. and P.M. Atkinson, 1998. Remote sensing and geostatistics. *Progress in Physical Geography*, 22 61-78.

Curtin, T., Bellingham, J., Catipovic, J., & Webb, D., (1993). Autonomous oceanographic sampling networks. *Oceanography*, 6(3).

Danish Hydraulic Institute. MIKE FLOOD WATCH. URL:

<http://www.dhigroup.com/Software/WaterResources/MIKEFLOODWATCH.aspx> (Accessed: May 2007)

Dee, D.P., 1994. On-line estimation of error covariance parameters for atmospheric data assimilation. *Monthly Weather Review*, 123 1128-1145.

- De Roo, A.P.J., Gouweleeuw, B., Thieelen, J., Bartholmes, J., Bongiannini, P., Todini, E., Bates, P.D., Horritt, M.S., Hunter, N., Beven, K.J., Pappenberger, F., Heise, E., Rivin, G., Hils, M., Hollingsworth, A., Holst, B., Kwadijk, J., Reggiani, P., Van Dijk, M., Sattler, K., Sprokkereef, E., 2003. Development of a European flood forecasting system. *International Journal of River Basin Management* 1 49-59.
- De Roure, D., 2005. FloodNet: A new flood warning system. *Royal Academy of Engineering*, 23 49-51.
- DEFRA, 2004. Making space for water: Developing new government strategy for flood and coastal erosion risk management in England: Flood warning and forecasting. DEFRA, London.
- Environment Agency, 2003. Flood forecasting – real time modelling. R&D technical report W5C-013/5/TR.
- Environment Agency. Station Summary Sheets: Crouch at Wickford. URL: http://www.nwl.ac.uk/ih/nrfa/station_summaries/037/031.html (Accessed: May 2007)
- Dobson, C., and Davies, G.P., 1990. Integrated real-time data retrieval and flood forecasting using conceptual models. Institute of Hydrology, Wallingford.
- Drecourt, J-P., and Madsen, H., 2002. Uncertainty estimation in groundwater modelling using Kalman filtering. ModelCARE 2002, Proceedings of the 4th International Conference on Calibration and Reliability in Groundwater Modeling, Prague, Czech Republic (Eds. K. Kovar and Z. Hrkal), Acta Universitatis Carolinae - Geologica 2002, 46(2/3), 306-309.
- Evensen, G., 1992. Using the extended Kalman filter with a multilayer quasi-geostrophic ocean model. *Journal of Geophysical Research-Oceans* 97 (C11), 17905-17924.

- Evensen, G., 1994. Sequential data assimilation with a nonlinear quasi-geostrophic 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.
- Evensen, G., 2004. Sampling strategies and square root analysis schemes for the EnKF. *Ocean Dynamics* 54 539-560.
- Franks, S.W., Gineste, P., Beven, K.J., and Merot, P., 1998. On constraining the prediction of a distributed model: The incorporation of fuzzy estimates of saturated areas into the calibration process. *Water Resources Research*, 34 787-797.
- Freer, J., and Beven, K.J., 1996. Bayesian estimation of uncertainty in runoff prediction and the value of data: An application of the GLUE approach. *Water Resources Research*, 32 2161-2173.
- French, J.R., 2003. Airborne LiDAR in support of geomorphological and hydrological modelling. *Earth Surface Processes and Landforms*, 28 321-335.
- GE Druck. 1830/1840 series high performance level pressure sensors. URL: <http://www.gesensing.com/products/resources/datasheets/1830.pdf>
(Accessed: May 2007)
- Georgakakos, K.P., Seo, D., Gupta, H., Schaake, J., and Butts, M.B., 2004. Towards the characterization of streamflow simulation uncertainty through multimodel ensembles. *Journal of Hydrology*, 298 222-241.
- Glenn, S.M., Grassle, J.F., and Alt, C.J., 2000. The LEO Approach. *Oceanus*, 42 28-30.

- Hankin, B.G., Hardy, R., Kettle, H., and Beven, K.J., 2001. Using CFD in a GLUE framework to model the flow and dispersion characteristics of a natural fluvial dead zone. *Earth Surface Processes and Landforms*, 26 667-687.
- Hartnack, J., Madsen, H., 2001. Data assimilation in river flow modelling. Proceedings of the fourth DHI software conference (June 2001, Helsingør, Denmark).
http://www.dhisoftware.com/uc2001/abstracts_proceedings/Proceedings/water_resources_track.htm.
- Hervouet, J-M., and Van Haren, L., 1996. Recent advances in numerical methods for fluid flows. In: Anderson M.G., Walling D.E., and Bates P.D. (Editors), *Floodplain processes*. John Wiley and Sons, Chichester pp. 183-214.
- Heuvelink, G.B.M., and Webster, R., 2001. Modeling soil variation: past, present, and future. *Geoderma*, 100 269-301.
- HR Wallingford. ISIS Flow / Hydrology. URL:
<http://www.wallingfordsoftware.com/products/isis/> (Accessed: May 2007)
- Horritt, M.S., and Bates, P.D., 2001. Predicting floodplain inundation: raster based modeling versus the finite element approach. *Hydrological Processes*, 15 825-842.
- Horritt, M.S., and Bates, P.D., 2002. Evaluation of 1D and 2D numerical models for predicting river flood inundation. *Journal of Hydrology*, 268 87-99.
- Horritt, M.S., Mason, D.C., and Luckman, A.J., 2001. Flood boundary delineation from synthetic aperture radar imagery using a statistical active contour model. *International Journal of Remote Sensing*, 22 2489-2507.

- Horritt, M.S., 2004. Development and testing of a simple 2-D finite volume model for sub-critical shallow water flow. *International Journal for Numerical Methods in Fluids*. 44 1231-1255.
- Hughes, D., 2006. An intelligent and adaptable grid-based flood monitoring and warning system. *Proceedings of the UK e-Science All Hands Meeting 2006* 18 - 21st September, Nottingham UK
- Hunter, N., Bates, P.D., Horritt, M.S., De Roo, A.P.J., Werner M., 2005. Utility of different data types for calibrating flood inundation models within a GLUE framework. *Hydrology and Earth System Sciences* 9 412-430.
- Ide, K., Courtier, P., Ghil, M., and Lorenc, A.C., 1997. Unified notation for data assimilation: Operational, sequential and variational. *Journal of the Meteorological Society of Japan*, 75 181-189.
- Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME-Journal of Basic Engineering* 82D 35-45.
- Kaun, E.L., 2005. A Study into the Practical Issues related to a Deployed Ad Hoc Wireless Sensor Network. In: New Orleans, LA USA.
- Khatibi, R., Stokes, R., Ogunyoye, F., Solheim, I., and Jackson, D., 2003. Research issues on warning lead-time and synergy in flood mitigation measures. *International Journal of River Basin management*, 1 331-346.
- Knebl, M.R., Yang, Z.L., Hutchison, K., Maidment, D.R., 2005. Regional scale flood modelling using NEXRAD rainfall, GIS, and HEC-HMS/RAS: a case study for the San Antonio River Basin Summer 2002 storm event. *Journal of Environmental Management* 75, 325-336.

- Knight, D.W., and Shiono, K., 1996. River channel and floodplain hydraulics. In: Anderson M.G., Walling D.E., and Bates P.D. (Editors), Floodplain Processes. John Wiley & Sons Ltd, Chichester pp. 139-181.
- Knott, A., 1999. Telemetry and Forecasting Systems for Managing Rivers and the Environment. URL:
http://www.pbworld.com/news_events/publications/network/tools/print_article.asp?referrer=/news_events/publications/network/issue_45/45_25_KnottA_TelemetryForecastingSys.asp (Accessed: May 2007)
- Krzysztofowicz, R., 2001. The case for probabilistic forecasting in hydrology. Journal of Hydrology, 249 2-9.
- Lamb, R., Beven, K.J., and Myrabo, S., 1998. Use of spatially distributed water table observations to constrain uncertainty in a rainfall-runoff model. Advances in Water Resources, 22 305-317.
- Lane, S.N., 2005. Roughness - time for a re-evaluation? Earth Surface Processes and Landforms, 30 251-253.
- Lee, Y.H., Singh, V.P., 1998. Application of the Kalman Filter to the Nash model. Hydrological Processes 12, 755-767.
- Lees, M.J., Young, P.C., Ferguson, S., Beven, K.J., Burns, J., 1994 An adaptive flood warning scheme for the River Nith at Dumfries, in W.R. White and J. Watts (eds.), 2nd International Conference on River Flood Hydraulics, p65-77, J. Wiley & Sons, Chichester.
- Lermusiaux, P.F.J., Evangelinos, C., Tian, R., Haley, P.J., McCarthy, J.J., Patrikalakis, N.M., Robinson, A.R., and Schmidt, H., 2004. Adaptive coupled physical and biogeochemical ocean predictions: A conceptual basis. Computational Science - Iccs 2004, Pt 3, Proceedings, 3038 685-692.

- Lermusiaux, P.F.J., 2006. Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *Journal of Computational Physics*, 217 176-199.
- Livings, D., 2005. Aspects of the ensemble Kalman filter. Unpublished master's dissertation. Department of Mathematics, University of Reading.
- Lloyd, C.D., and Atkinson, P.M., 2002. Deriving DSM's from LiDAR data with kriging. *International Journal of Remote Sensing*, 23 2519-2524.
- Lloyd, C.D., and Atkinson, P.M., 2006. Deriving ground surface digital elevation models from LiDAR data with geostatistics. *International Journal of Geographical Information Science*, 20 535-563.
- Madsen, H., Cañizares, R., 1999. Comparison of extended and ensemble Kalman filter for data assimilation in coastal area modelling. *International Journal for Numerical Methods in Fluids* 31, 961-981.
- Madsen, H., Skotner, C., 2005. Adaptive state updating in real-time river flow forecasting - a combined filtering and error forecasting procedure. *Journal of Hydrology* 308, 302-312.
- Majumdar, S.J., Bishop, C.H., Etherton, B.J., and Toth, Z., 2002. Adaptive sampling with the ensemble transform Kalman filter. Part II: Field program implementation. *Monthly Weather Review*, 130 1356-1369.
- Mantovan, P., Todini, E., 2006. Hydrological forecasting uncertainty assessment: incoherence of the GLUE methodology. *Journal of Hydrology*. 330 368-381.
- Marks K., Bates P. D., 2000. Integration of high-resolution topographic data with floodplain flow models. *Hydrological Processes* 14, 2109-2122.
- Mason, D.C., Cobby, D.M., Horritt, M.S., and Bates, P.D., 2003. Floodplain friction parameterization in two-dimensional river flood models using vegetation

heights derived from airborne scanning laser altimetry. *Hydrological Processes*, 17 1711-1732.

Mason, D.C., Horritt, M.S., Hunter, N.M. & Bates, P.D., 2007. Use of fused airborne scanning laser altimetry and digital map data for urban flood modelling, *Hydrological Processes*, 21 1436-1447.

Maybeck, P.S., 1979. *Stochastic models, estimation and control*. 1. Academic Press, New York.

McLaughlin, D.B., 2002. An integrated approach to hydrologic data assimilation: interpolation, smoothing, and filtering. *Advances in Water Resources* 25, 1275-1286.

Moore, R.J., Jones, D.A., Bird, P.B., Cottingham, M.C., 1990. A basin-wide flow forecasting system for real-time flood warning, river control, and water management. In *River flood hydraulics* White W. R. and Watts J. (Eds), Wiley, Chichester.

Moore, R.J., Jones, A.E., Jones, D.A., Black, K.B., Bell, V.A., 2004. Weather radar for flood forecasting. *Sixth international symposium on hydrological applications of weather radar*, Melbourne, Australia.

Moradkhani, H., Sorooshian, S, Gupta, H., and Houser, P.R., 2005. Dual state-parameter estimation of hydrological models using ensemble Kalman filter. *Advances in Water Resources*, 28 135-147.

Moussa, R., and Bocquillon, C., 1996. Criteria for the choice of flood-routing method in natural channels. *Journal of Hydrology*, 186 1-30.

Nadis, S., 1997. 'Real-time' oceanography adapts to sea changes. *Science*, 275 1881-1882.

- Nash, J. E. and J. V. Sutcliffe (1970), River flow forecasting through conceptual models part I - A discussion of principles, *Journal of Hydrology*, 10 282–290.
- Nicholas, A.P., and Mitchell, C.A., 2003. Numerical simulation of overbank processes in topographically complex floodplain environments. *Hydrological Processes*, 17 727-746.
- Page, T., Whyatt, J.D., Beven, K.J., and Metcalfe, S.E., 2004a. Uncertainty in modelled estimates of acid deposition across Wales: a GLUE approach. *Atmospheric Environment*, 38 2079-2090.
- Page, T., Beven, K.J., and Whyatt, D., 2004b. Predictive capability in estimating changes in water quality: long term responses to atmospheric deposition. *Water, Air, and Soil Pollution*, 151 215-244.
- Pappenberger, F., Beven, K.J., Horritt, M.S., Blazkova, S., 2005a. Uncertainty in the calibration of effective roughness parameters in HEC-RAS using inundation and downstream level observations. *Journal of Hydrology* 302, 46-69.
- Pappenberger, F., Beven, K.J., Hunter, N., Gouweleeuw, B., Thielen, J., De Roo, A.P.J., 2005b. Cascading model uncertainty from medium range weather forecasts (10 days) through a rainfall-runoff model to flood inundation predictions within the European Flood Forecasting System (EFFS). *Hydrology and Earth System sciences* 9 381-393.
- Rahimi, M., Pon, R., Kaiser, W.J., Sukhatme, G.S., Estrin, D., and Srivastava, M., 2003. Adaptive sampling from environmental robotics. UCLA Center for Embedded Networked Sensing, Technical Report 29, November 2003.
- Refsgaard, J.C., 1997. Validation and intercomparison of different updating procedures for real-time forecasting. *Nordic Hydrology* 28, 65-84.

- Reichle, R.H., McLaughlin, D.B., Entekhabi, D., 2002. Hydrologic data assimilation with the ensemble Kalman filter. *Monthly Weather Review* 130, 103-114.
- Robinson, A.R. and S.M. Glenn, 1999. Adaptive sampling for ocean forecasting, *Naval Research Reviews*, to appear. Available via the WWW at http://www.deas.harvard.edu/~robinson/PAPERS/Naval_Review.html
- Romanowicz, R., Beven, K.J., 1998. Dynamic real-time prediction of flood inundation probabilities. *Hydrological Sciences-Journal-des-Sciences Hydrologiques* 43 181-196.
- Romanowicz, R., and Beven, K.J., 2003. Estimation of flood inundation probabilities as conditioned on event inundation maps. *Water Resources Research*, 39 1073.
- Romanowicz, R., Beven, K.J., and Tawn, J., 1996. Bayesian calibration of flood inundation models. In: Anderson M.G., Walling D.E., and Bates P.D. (Editors), *Floodplain Processes*. John Wiley & Sons Ltd., New York pp. 333-360.
- Schreider, S.Y., Young, P.C., Jakeman, A.J., 2001. An application of the Kalman filtering technique for streamflow forecasting in the upper Murray basin. *Mathematical and Computer Modelling* 33, 733-743.
- Shiiba, M., Laurenson X., Tachikawa Y., 2000. Real-time stage discharge estimation by a stochastic-dynamic flood routing model. *Hydrological Processes* 14, 481-495.
- Singh, V.P., 2001. Kinematic wave modelling in water resources: a historical perspective. *Hydrological Processes*, 15 671-706.

- Snepvangers, J.J.J.C., Heuvelink, G.B.M., and Huisman, J.A., 2003. Soil water content interpolation using spatio-temporal kriging with external drift. *Geoderma*, 112 253-271.
- Sprokkereef, E., 2001. Extension of the Flood Forecasting Model FloRIJN. NCR Publication 12-2001, ISSN no. 1568234X.
- Tippett, M.K., Anderson, J.L., Bishop, G.H. Hamill, T.M., and Whitaker J.S., 2003. Ensemble Square Root Filters. *Monthly Weather Review*, 131, 1485-1490.
- Todini, E., 1994. The ARNO rainfall-runoff model. *Journal of Hydrology* 175, 339-382.
- Troch, P.A., Paniconi, C., and McLaughlin, D., 2003. Catchment-scale hydrological modelling and data assimilation. *Advances in Water Resources*, 26 131-135.
- Verlaan, M., Heemink, A.W., 1997. Tidal flow forecasting using reduced rank square root filter. *Stochastic Hydrology and Hydraulics*, 11, 349-368.
- Walker, J.P., Houser, P.R., Reichle, R.H., 2003. New Technologies require advances in Hydrologic data assimilation. *EOS*, 84 545-551.
- Wang, X.G. and Bishop, C.H., 2003. A comparison of breeding and ensemble transform Kalman filter ensemble forecast schemes. *Journal of the Atmospheric Sciences*, 60 1140-1158.
- Wang,X.G., Bishop,C.H., and Julier,S.J., 2004. Which is better, an ensemble of positive-negative pairs or a centered spherical simplex ensemble? *Monthly Weather Review*, 132 1590-1605.
- Welch, G., Bishop, G., 2002. An introduction to the Kalman filter. Technical Report 95-041, University of North Carolina at Chapel Hill, Department of Computer Science.

- Werner, M., Reggiani, P., De Roo A.D., Bates P.D., and Sprokkereef, E., 2005. Flood forecasting and warning at the river basin and European scale. *Natural Hazards*, 36 25-42.
- Whitaker, J.S., Hamill, T.M., 2002. Ensemble data assimilation without perturbed observations. *Monthly Weather Review* 130, 1913-1924.
- Wilson, M.D., 2004. Evaluating the effect of data and data uncertainty on predictions of flood inundation. Unpublished PhD Thesis. University of Southampton.
- Wilson M.D., and Atkinson, P.M., 2005. The use of elevation data in flood inundation modelling: a comparison of ERS interferometric SAR and combined contour and differential GPS data. *International Journal of River Basin Management*, 3 3-20.
- Wilson M.D. and Atkinson P.M., (2007). The use of remotely sensed land cover to derive floodplain friction coefficients for flood inundation modelling. *Hydrological Processes*.
- WMO, 1992. Simulated real-time intercomparison of hydrological models, WMO operational hydrology report no 38, WMO no 779, World Meteorological Organisation, Geneva.
- Young, P.C., 1984. *Recursive estimation and time series analysis*. Springer-Verlag, Berlin.
- Young, P.C., 2002. Advances in real-time flood forecasting. *Philosophical Transactions of the Royal Society of London Series A-Mathematical Physical and Engineering Sciences*, 360 (1796): 1433-1450.
- Young, P.C., 2003. Top-down and data-based mechanistic modelling of rainfall-flow dynamics at the catchment scale. *Hydrological Processes*, 17 2195-2217.

- Yu, D., and Lane, S.N., 2006. Urban fluvial flood modeling using a two-dimensional diffusion-wave treatment, part 1: mesh resolution effects. *Hydrological Processes*, 20 1541-1565.
- Zak, S.K., and Beven, K.J., 1999. Equifinality, sensitivity and predictive uncertainty in the estimation of critical loads. *The Science of the Total Environment*, 236 191-214.
- Zhou, J., and De Roure, D., 2007. FloodNet: Coupling adaptive sampling with energy aware routing in a flood warning system. *Journal of Computer Science and Technology*, 22 121-130

Appendices

Appendix A: Simple KF example: bucket water levels

To help clarify KF analysis scheme in conjunction with section 3.3.4 a hypothetical example for a single point based on a bucket filling with rainwater is presented here.

The TV weatherman states that 4 cm of rainfall is expected to fall overnight ($x = 4$ cm). However, he also states that the forecast is not certain but decides to show off by stating that it has a variance of 1.5 cm ($\sigma_x^2 = 1.5$ cm). By observing a bucket placed outside overnight by eye it appears that around 5 cm of rain has fallen ($y = 5$ cm) However, without a more accurate instrument to take the measurement you assume that your measurement has a variance of 1 cm ($\sigma_y^2 = 1$ cm). The Kalman filter analysis from Eq. 3.24 can give you a more accurate estimate of how much rain fell into the bucket by combining these two sources of information. Recalling Eq. 3.23-24:

$$\begin{aligned} \mathbf{K} &= \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} \\ \mathbf{x}^a &= \mathbf{x}^f + \mathbf{K}(\mathbf{y} - \mathbf{H}\mathbf{x}^f) \end{aligned} \tag{A1.1-2}$$

Under the single point circumstance the observation operator \mathbf{H} becomes 1, as the measurement maps directly onto the model state. This effectively means the innovation matrix ($\mathbf{y} - \mathbf{H}\mathbf{x}^f$) becomes the difference between the measured and simulated state ($d = y - x = 1$). For a point model the Kalman gain \mathbf{K} simplifies to $k = \sigma_x / \sigma_x + \sigma_y = 0.6$. Thus by using equation 3.24/A1.2 the analysed estimate of x is:

$$\begin{aligned} x^a &= x^f + kd \\ x^a &= 4 + 0.6 * 1 \\ x^a &= 4.6 \end{aligned} \tag{A1.3-5}$$

The variance of the analysed state estimate can also be calculated by recalling Eq. 3.25:

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{KH})\mathbf{P}^f (\mathbf{I} - \mathbf{KH})^T + \mathbf{K}\mathbf{R}\mathbf{K}^T \quad (\text{A1.6})$$

Which when applied to the single point bucket example simplifies to become

$$\begin{aligned} \sigma^a &= \sigma_x^f - k\sigma_x^f \\ \sigma^a &= 1.5 - 0.6 * 1.5 \\ \sigma^a &= 0.6 \end{aligned} \quad (\text{A1.7-9})$$

Thus the analysed estimate of the bucket water level is 4.6 cm variance 0.6 cm, which is less than both the forecasted and measured values.

Appendix B: Ensemble transform Kalman filter code.

This appendix contains the Matlab 6 (R12) code used to implement the ETKF (Note: there is no reason why it shouldn't also work in Matlab 7). Assuming there are N model states, K ensemble members and M measurements the following matrices must be specified:

Z	N -by- K	Ensemble perturbations at target time scaled by the square root of $K-1$
Zval	N -by- K	Ensemble perturbations at validation time scaled by the square root of $K-1$
H	M -by- K	Measurement operator
R	M -by- M	Measurement error covariance
I	K -by- K	Identity matrix

The code below calculates the following:

T	K -by- K	Ensemble transform matrix
Za	N -by- K	Analysed ensemble perturbations
Zaval	N -by- K	Analysed ensemble perturbations at validation time
Pa	N -by- N	Analysed error covariance
Paval	N -by- N	Analysed error covariance at validation time
S	N -by- N	Signal covariance
Sval	N -by- N	Signal covariance at validation time
A	K -by- K	Prediction error covariance matrix
C	K -by- K	Eigenvalues of A
Y	K -by- K	Eigenvector of A

Using the function:

```
function [T, Za, Zaval, Pa, Paval, S, Sval] = ETKF (Z, Zval, H, R, I)
% Ensemble transform Kalman filter for Matlab 6 (R12)
```

```
% Calculate the prediction error covariance matrix associated with the measurements
H2 = (R^-0.5)*H*Z;
```

$A = H_2' * H_2;$

% Calculate the eigenvalues and eigenvectors of the matrix A
 $[C, Y] = \text{eig}(A);$

% Calculate the transform matrix T
 $T = C * ((Y+I)^{-0.5}) * C';$

% Calculate updated ensemble perturbations at the target time and validation time
 $Z_a = Z * T;$
 $Z_{val} = Z_{val} * T;$

% Calculate analysed error covariance at the target and validation time (optional)
 $P_a = Z * Z';$
 $P_{val} = Z_{val} * Z_{val}';$

% Calculate the signal covariance at the target time and validation time (optional)
 $S = (Z * C) * Y * ((Y+I)^{-1}) * (Z * C)';$
 $S_{val} = (Z_{val} * C) * Y * ((Y+I)^{-1}) * (Z_{val} * C)';$