Data Assimilation within the Advanced Circulation (ADCIRC) Modeling Framework for Hurricane Storm Surge Forecasting

T. BUTLER

Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, Texas

M. U. ALTAF

King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

C. DAWSON

Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, Texas

I. HOTEIT AND X. LUO

King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

T. MAYO

Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, Texas

(Manuscript received 12 May 2011, in final form 3 December 2011)

ABSTRACT

Accurate, real-time forecasting of coastal inundation due to hurricanes and tropical storms is a challenging computational problem requiring high-fidelity forward models of currents and water levels driven by hurricane-force winds. Despite best efforts in computational modeling there will always be uncertainty in storm surge forecasts. In recent years, there has been significant instrumentation located along the coastal United States for the purpose of collecting data—specifically wind, water levels, and wave heights—during these extreme events. This type of data, if available in real time, could be used in a data assimilation framework to improve hurricane storm surge forecasts. In this paper a data assimilation methodology for storm surge forecasting based on the use of ensemble Kalman filters and the advanced circulation (ADCIRC) storm surge model is described. The singular evolutive interpolated Kalman (SEIK) filter has been shown to be effective at producing accurate results for ocean models using small ensemble sizes initialized by an empirical orthogonal function analysis. The SEIK filter is applied to the ADCIRC model to improve storm surge forecasting, particularly in capturing maximum water levels (high water marks) and the timing of the surge. Two test cases of data obtained from hindcast studies of Hurricanes Ike and Katrina are presented. It is shown that a modified SEIK filter with an inflation factor improves the accuracy of coarse-resolution forecasts of storm surge resulting from hurricanes. Furthermore, the SEIK filter requires only modest computational resources to obtain more accurate forecasts of storm surge in a constrained time window where forecasters must interact with emergency responders.

1. Introduction

Prediction of coastal inundation due to hurricanes and tropical storms is a problem of great importance. Over

the past few decades, storm surge has been a major cause for deaths worldwide. In November 1970, the infamous Bhola cyclone made landfall in Bangladesh; the associated storm surge is estimated to have killed between 300 000 and 500 000 people (Murty et al. 1986). More recently, Hurricane Katrina made landfall in Louisiana and Mississippi in August 2005, directly killing at least 1200 people (Blake et al. 2011). The majority of deaths during hurricanes are caused by storm surge, and could

---

Corresponding author address: T. Butler, Institute for Computational Engineering and Sciences, The University of Texas at Austin, 1 University Station C0200, Austin, TX 78712.
E-mail: tbutler@ices.utexas.edu

be prevented with improved planning, warning systems, and emergency response. Accurate numerical forecasts of coastal flooding, delivered in real time, could result in more timely evacuations and help significantly with the deployment of first responders and emergency personnel. Storm surges have been the subject of much study, dating back several decades (Heap 1983). For example, in the southern North Sea a surge was responsible for the worst natural disaster to affect the United Kingdom, the Netherlands, and Denmark in recent times (Gerritsen et al. 1995; McRobie et al. 2005). Recent studies of severe U.S. hurricanes, including Hurricanes Katrina, Rita, Gustav, and Ike, can be found in Dietrich et al. (2010, 2011) and Kennedy et al. (2011).

Numerical models of storm surge will always be subject to significant uncertainty (Brown et al. 2007). To reduce these uncertainties, it is critical that storm surge forecast models use all of the data available to them, in particular real-time measurements of water levels and currents that are becoming readily available. One approach that incorporates measurement data into a model simulation to minimize these uncertainties is data assimilation (Ghil 1989; Malanotte-Rizzoli et al. 1989). Kalman filtering (Kalman 1960) is a commonly applied data assimilation procedure in which the model is not only used to forecast the system state, but also to determine the uncertainty of the estimate (Maybeck 1979).

Originally the Kalman filter was designed for linear systems. One difficulty that hampers the implementation of the classical Kalman filter for large-scale numerical models is the computational cost due to the propagation of the error covariance. For linear, time-invariant, and stable systems, a steady-state Kalman filter was developed by Heemink (1986) and has been proven efficient in many storm surge applications (see, e.g., Heemink and Kloosterhuis 1990; Sorensen et al. 2006; El Serafy and Mynett 2008). The limitation of this approach is that it requires a fixed observation network and does not address the nonlinearity issue. Recently, new methods have been developed to solve the nonlinearities and computational cost issues in implementing the Kalman filter. Important examples include the so-called reduced-rank approaches (Verlaan and Heemink 1997; Hoteit et al. 2002) and the ensemble Kalman filters (EnKFs) (Evensen 1994; Mitchell and Houtekamer 2000; Bishop et al. 2001; Hoteit et al. 2002). Numerically, these generic filter algorithms tend to be more robust for nonlinearities in the model than the conventional model-dependent approaches such as the extended Kalman filter (EKF) (Maybeck 1979). They also allow for efficient computation of the analysis step at a reasonable computational cost. Recent examples of applications in storm surge

forecasting are given by Verlaan and Heemink (1997) and Sorensen et al. (2004).

The goal of this paper is to describe recent advances in the science of storm surge forecasting and combine state-of-the-art, real-time data assimilation methods with accurate, forward storm surge modeling. The singular evolutive interpolated Kalman (SEIK) filter is the data assimilation method used in this study. The SEIK filter is an ensemble-based square root filter and uses the minimal number of ensemble members required to obtain second-order exact sampling of the mean and covariance of a feature space determined by an empirical orthogonal function analysis. The numerical model used in this study is the advanced circulation (ADCIRC) model. The model discretizes the shallow-water equations using finite element methods defined on unstructured meshes (Luettich and Westerink 2004).

We test the application of the SEIK filter, as described in Hoteit et al. (2002), to the ADCIRC model using wind fields and data collected from Hurricanes Ike and Katrina. Specifically, high-fidelity wind fields that are computed from wind data collected during these hurricanes are used with fine-resolution grids to compute hindcast studies. These hindcast studies have been validated against data of free-surface elevation that were gathered during these hurricanes and are used to simulate the data of free-surface elevation in our experiments. Thus, we are computing what are known as observing system simulation experiments (OSSEs), which are akin to twin experiments. However, the forecasts we generate are on a coarser grid than those used in the hindcast studies and use wind fields only as they would be computed during an actual forecast based on using the best track data, which are available from the National Oceanic and Atmospheric Administration (NOAA) archive (<ftp://ftp.tpc.ncep.noaa.gov/atcf/archive/>).

This paper is organized as follows: in section 2, we describe the storm surge model. The governing equations are given in section 2a. An outline of recent interest in the problem of collecting and utilizing data for extreme events is given in section 2b. In section 4 and its subsections, we describe the singular evolutive interpolated Kalman filter with an inflation factor. In section 5, we describe the setup and framework of the OSSEs in detail. Numerical results are presented in section 6, and concluding remarks are in section 7.

2. ADCIRC coastal circulation storm surge model

Mathematical models of storm surge are described by the shallow-water equations. These equations are derived from the incompressible Navier–Stokes equations under the assumption of hydrostatic pressure (Vreugdenhil

1994) and solve for water elevation (or water depth) and water currents (velocity). These models are defined on complex physical domains that include many scales, from the deep ocean through basins and up continental shelves to coastal inlands. Complex coastal features include barrier islands, rivers, and channels. In storm surge applications, structural features such as levees, raised roads, and railways should be included in the description of the domain where such data are available. The storm surge models are forced primarily by tides, winds, and wind waves; during a hurricane, wind stresses are the dominant forcing. Accurate knowledge of the bottom topography, coastal elevation, and friction characteristics, both underwater and on potentially inundated land, is needed but not always available.

Interest in numerical models of storm surge has increased dramatically since the devastating 2005 hurricane season. Over the past several years a multi-institutional research team has undertaken the development and application of the state-of-the-art advanced circulation model (Luettich and Westerink 2004). The ADCIRC model discretizes the shallow-water equations in space using finite element methods defined on unstructured meshes. Coupling of the ADCIRC model with a wind wave model for capturing wave-induced setup has recently been completed (Dietrich et al. 2011). The code has been parallelized for distributed memory and multicore computers and has been demonstrated to achieve excellent scalability on these platforms (Tanaka et al. 2011).

a. Governing equations and discretization

Starting with the Reynolds averaged Navier–Stokes equations, applying the hydrostatic and Boussinesq approximations, assuming the earth’s radius is large relative to the depth of the ocean, and averaging over the water depth H , we arrive at the two-dimensional shallow-water equations in spherical coordinates (λ, ϕ) :

- continuity,

$$\frac{\partial \zeta}{\partial t} + \frac{1}{R \cos \phi} \left[\frac{\partial(UH)}{\partial \lambda} + \frac{\partial(VH \cos \phi)}{\partial \phi} \right] = 0, \quad (2.1)$$

- and horizontal momentum,

$$\begin{aligned} \frac{dU}{dt} = fV - \frac{1}{R \cos \phi} \frac{\partial[g(\zeta - \alpha\eta) + p_s/\rho_0]}{\partial \lambda} \\ + \frac{\tau_{s\lambda}}{\rho_0 H} - \tau_{bf} U + m_\lambda \end{aligned} \quad (2.2)$$

$$\begin{aligned} \frac{dV}{dt} = -fU - \frac{1}{R} \frac{\partial[g(\zeta - \alpha\eta) + p_s/\rho_0]}{\partial \phi} \\ + \frac{\tau_{s\phi}}{\rho_0 H} - \tau_{bf} V + m_\phi. \end{aligned} \quad (2.3)$$

TABLE 1. Variable definitions for the continuity and momentum equations shown in (2.1)–(2.3).

Variable	Definition
t	Time
λ, ϕ	Degrees longitude, latitude
ζ	Free-surface elevation relative to the geoid
U, V	Depth-averaged horizontal velocity components
R	Mean radius of the earth ($6.378\,206\,4 \times 10^6$ m)
H	$\zeta + h$ = water depth
h	Bathymetric depth relative to the geoid
g	Gravitational acceleration
f	Coriolis coefficient
p_s	Atmospheric pressure at the free surface
η	Newtonian equilibrium tide potential
α	Effective earth elasticity factor
ρ_0	Reference density of water
$\tau_{s\lambda}, \tau_{s\phi}$	Applied free-surface stress
τ_{bf}	$C_f[(U^2 + V^2)^{1/2}/H]$ = bottom friction
C_f	Nonlinear bottom friction coefficient
$\frac{d}{dt}$	$\frac{\partial}{\partial t} + \frac{U}{R \cos \phi} + \frac{\partial}{\partial \lambda} + \frac{V}{R} + \frac{\partial}{\partial \phi}$
m_λ, m_ϕ	$\frac{\nu_T}{H} \left[\frac{\partial}{\partial \lambda} \left(\frac{\partial UH}{\partial \lambda} \right) + \frac{\partial}{\partial \phi} \left(\frac{\partial UH}{\partial \phi} \right) \right],$ $\frac{\nu_T}{H} \left[\frac{\partial}{\partial \lambda} \left(\frac{\partial VH}{\partial \lambda} \right) + \frac{\partial}{\partial \phi} \left(\frac{\partial VH}{\partial \phi} \right) \right]$
ν_T	Depth-averaged horizontal eddy viscosity

These equations are solved on a 2D spatial domain Ω and time interval $(t_0, T]$, with appropriate boundary and initial conditions. Refer to Table 1 for continuity and momentum equation variable definitions.

The ADCIRC model is based on a reformulation of the continuity equation into a generalized wave continuity equation (GWCE) (Lynch and Gray 1979; Kinnmark 1986), which replaces (2.1) by a second-order wave equation for the water elevation ζ . The domain Ω is discretized by triangular finite elements. A typical domain Ω is the Gulf of Mexico, possibly including the western North Atlantic, as seen in Fig. 1. The GWCE and momentum equations (2.2)–(2.3) are discretized in space using a standard Galerkin finite element method, with continuous, piecewise linear approximations of ζ , U , and V defined on the triangular mesh. The GWCE is discretized in time using a central difference scheme with three time levels, while the momentum equation uses a forward Euler discretization. The resulting scheme is essentially explicit in time, except for the solution of a mass matrix in the GWCE discretization. The complete numerical formulation is described in Luettich and Westerink (2004).

b. Data collection and hindcast studies

Since Hurricane Katrina in 2005, there have been substantial efforts by federal, state, and local agencies in

FIG. 1. Western North Atlantic domain and bathymetry in meters. The Gulf of Mexico is circled in black.

the United States to place instruments capable of measuring water levels, wind speeds, and wave heights throughout the most highly impacted areas of the coast, particularly along the Texas, Louisiana, and Mississippi coasts. These data collection efforts have been particularly useful for hindcasting recent hurricane events, for the purpose of better understanding the coastal impact, and designing improved coastal protection systems.

The ADCIRC model has undergone extensive verification and validation, in particular by comparison with data from previous storms, including Hurricanes Betsy (1965), Ivan (2004), Dennis (2004), Katrina (2005), and Rita (2005) (Westerink et al. 2008; Bunya et al. 2010; Dietrich et al. 2010), and more recently, Gustav (2008) (Dietrich et al. 2011) and Ike (2008) (Kennedy et al. 2011). These hindcasting studies after the event are extremely valuable for gaining physical understanding and guiding decisionmakers in future coastal development and coastal hazard mitigation. Hindcasting relies on highly accurate model inputs and accurate data. These studies have greatly benefited from vast amounts of data that are now collected during a storm event. The past five years has seen a large effort in deploying sensor networks throughout critical coastal regions to collect data on all aspects of the storm, including water levels, winds, and significant wave heights. These data are often available in real time. For example, the Texas Coastal Ocean Observation Network (TCOON) provides real-time water level and meteorological information at over 40 stations along the Texas coast and is available online (lighthouse.tamucc.edu/TCOON).

The extensive data collected and fundamental knowledge gained from these hindcasting studies have recently led to the development of a real-time forecasting system

based on the ADCIRC model, called the ADCIRC surge guidance system (ASGS) (Fleming et al. 2008). In this model, data on the hurricane track, forward speed, and wind characteristics (wind speed, central pressure, and radius of maximum winds) are obtained every 6 h from the National Weather Service. This input is used to generate a parametric wind field, which provides forcing to the ADCIRC model. When the storm is far from landfall the forecast track and intensity at landfall are highly uncertain; however, 48–72 h from landfall the forecast uncertainty begins to decrease. It is during this critical time that *accurate* forecasts of storm surge are required. To be useful to emergency managers, the storm surge model must compute predicted water levels along the coast within a 1–2-h time window after the most recent forecast is released from the National Hurricane Center (NHC). Typically the outputs desired are the maximum water elevations, measured over a given coastal region during the entire storm event, and time histories of water levels and/or currents at critical locations along the coast.

3. Data assimilation and extreme events

As noted above, the ADCIRC model has been used successfully in hindcast mode to study Hurricanes Betsy, Katrina, Rita, Gustav, and Ike. Because of these hindcast studies, we have developed very accurate descriptions of the Texas, Louisiana, and Mississippi coasts, which include accurate bathymetries, coastal features such as levees, raised roads, and railways, channels, wetlands and marshes, and land use characteristics. These hindcast models require highly resolved finite element meshes, with millions of degrees of freedom, and typically require small time steps to resolve wetting and drying fronts. While the hindcast models could be run in forecast mode, they require significant computational resources just to run one forecast. Running a large ensemble of forecasts, as is typically needed to capture uncertainty, is beyond the capabilities of current computational resources.

The goal is clear—to improve storm surge forecasts by a data assimilation strategy using only a small ensemble of states. We describe the so-called singular evolutive interpolated Kalman filter below as one such data assimilation strategy.

4. The singular evolutive Kalman filters

Various dimension reduction techniques have been proposed over the years to decrease the cost of the Kalman filter by projecting the state (of dimension n) onto a low-dimensional subspace (of dimension r) (Cane et al. 1996; Dee 1991; Fukumori and Malanotte-Rizzoli 1995; Hoang et al. 1998). We call this low-dimensional

subspace the *feature space* of the system. The singular evolutive extended Kalman (SEIK) filter was proposed by Pham et al. (1997) as a low-rank square root extended Kalman filter in which the error covariance matrix is approximated by a singular matrix of low rank $r \ll n$. This resulted in Kalman corrections only in the feature space spanned by the directions whose error is not sufficiently attenuated by the system. Noticing that ensemble representations can be efficiently used to represent square root covariance matrices (Tippett et al. 2003), Pham (2001) introduced an ensemble-based variant of the SEIK filter, called the singular evolutive interpolated Kalman filter. In the SEIK filter, an ensemble spanning the r -dimensional feature space replaces the linearization of the system required in the SEIK filter (Hoteit et al. 2005). The ensemble contains $r + 1$ *interpolating states* and is drawn pseudorandomly after the analysis step such that the analysis state vector and its error covariance are equal to the sample mean and sample variance, respectively, of the ensemble. This type of second-order exact sampling has parallels to the unscented Kalman filter, and we direct the interested reader to Julier and Uhlmann (1997) and Wan and van der Merwe (2000). Correcting the mean and resampling the ensemble are central ideas in widely applied deterministic square root ensemble Kalman filters [see, e.g., Tippett et al. (2003) and Sakov and Oke (2008) for more details].

In the following section, we adopt the abstract system notation

$$\mathbf{X}^t(t_k) = M(t_k, t_{k-1})\mathbf{X}^t(t_{k-1}) + \boldsymbol{\eta}(t_k), \quad (4.1)$$

where $\mathbf{X}^t(t)$ denotes the vector representing the true state of the system at time t , $M(t, s)$ is the state transition operator that takes as inputs the state at time s and outputs the state at time t , and $\boldsymbol{\eta}(t)$ is the system noise vector with covariance matrix $\mathbf{Q}_t$. At time t_k , the observed data vector of dimension p is given by

$$\mathbf{Y}_k = H_k \mathbf{X}^t(t_k) + \boldsymbol{\epsilon}_k. \quad (4.2)$$

Here, H_k is the observation operator and $\boldsymbol{\epsilon}_k$ is the observational noise with covariance matrix $\mathbf{R}_k$. The SEIK filter operates recursively beginning with an analysis, then a resampling step to generate the ensemble members, a forecast step with the model, and a correction (or analysis) step to update the forecasted ensembles with new available observations. The algorithm starts with an initialization step to determine an initial ensemble of states.

a. Initialization

We make use of an empirical orthogonal function (EOF) analysis as presented in Pham et al. (1997). The

EOF analysis requires storing a possibly large number of snapshots of the model state, which can be expensive. However, this is done only once, and the results can be used to initialize multiple runs of the SEIK filter. We are interested in the modeling of extreme events such as hurricanes, for which the time horizon of interest may only be a few days and the forcing wind data available for only a week.

Driving the ADCIRC model only by tidal forcing terms eliminates all transient behavior in the state within a few days. We run the ADCIRC model driven only by tidal forcing terms for 60 days and store the state every 5 h. The perturbation of these states from their mean value is used to define a sample covariance matrix from which the ensemble members can be initialized. In this way, we initialize ensemble members using a covariance from a physically meaningful feature space. Let $\mathbf{P}$ denote the sample covariance matrix of the sequence of state vectors from the long model run. Let $\mathbf{v}_1, \mathbf{v}_2, \dots$, denote the eigenvectors of $\mathbf{P}$ with associated eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots$. The vectors $\mathbf{v}_i$ are the EOFs, and we approximate $\mathbf{P}$ by the r -rank $\mathbf{P}^a = \mathbf{L}\mathbf{U}\mathbf{L}^T$, where

$$\mathbf{L} = (\mathbf{v}_1 \dots \mathbf{v}_r), \quad \mathbf{U} = \text{diag}(\lambda_1, \dots, \lambda_r). \quad (4.3)$$

The question remains as to how we choose appropriate r . As noted in Pham et al. (1997), the ratio $\sum_{j>r} \lambda_j / \text{Tr}(\mathbf{P})$ represents the relative error in the square $\mathbf{L}^2$ norm of using approximations to the state in the r -dimensional feature space and is useful in determining r given a prescribed error tolerance, which we also refer to as the percentage of inertia or system variability retained by the EOFs. In the experiments below, we chose to retain 90% of the spatial variability in this sequence of states using the EOF analysis, resulting in an ensemble size of $r + 1 = 10$.

b. Resampling step

At time t_{k-1} , the SEIK filter has produced the analysis state $\mathbf{X}^a(t_{k-1})$ and its corresponding error covariance matrix $\mathbf{P}^a(t_{k-1})$. We use the state and error covariance in producing an ensemble of $1 \leq i \leq r + 1$ interpolated states $\mathbf{X}_i^a(t_{k-1})$. We factor $\mathbf{P}^a(t_{k-1})$ into $\mathbf{L}_{k-1} \mathbf{U}_{k-1} \mathbf{L}_{k-1}^T$, where $\mathbf{L}_{k-1}$ is $n \times r$ and $\mathbf{U}_{k-1}$ is $r \times r$. We compute the square root of the symmetric positive definite matrix $\mathbf{U}_{k-1}^{-1}$ and use a random orthonormal matrix $\boldsymbol{\Omega}_{k-1}$ with zero column sums to sample the directions in the feature space spanned by $\mathbf{U}_{k-1}$. Specifically, we use the Cholesky decomposition $\mathbf{U}_{k-1}^{-1} = \mathbf{C}_{k-1} \mathbf{C}_{k-1}^T$ so that

$$\mathbf{P}^a(t_{k-1}) = \mathbf{L}_{k-1} (\mathbf{C}_{k-1}^{-1})^T \boldsymbol{\Omega}_{k-1}^T \boldsymbol{\Omega}_{k-1} \mathbf{C}_{k-1}^{-1} \mathbf{L}_{k-1}^T. \quad (4.4)$$

A procedure for computing $\mathbf{\Omega}_{k-1}$ using Householder matrices is detailed in Hoteit et al. (2002). The interpolating states are given by

$$\mathbf{X}_i^a(t_{k-1}) = \mathbf{X}^a(t_{k-1}) + \sqrt{r+1} \mathbf{L}_{k-1} (\mathbf{\Omega}_{k-1,i} \mathbf{C}_{k-1}^{-1})^T, \quad 1 \leq i \leq r+1, \quad (4.5)$$

where $\mathbf{\Omega}_{k-1,i}$ denotes the i th row of $\mathbf{\Omega}_{k-1}$. Drawing the interpolated states in this way ensures that the sample mean and sample variance of the ensemble are exactly $\mathbf{X}^a(t_{k-1})$ and $\mathbf{P}^a(t_{k-1})$, respectively.

c. Forecast stage

The computational model is initialized and executed $r+1$ times to integrate from t_{k-1} to t_k , the $r+1$ analyzed interpolated states $\mathbf{X}_i^a(t_{k-1})$ to the $r+1$ forecast states $\mathbf{X}_i^f(t_k)$; that is, $\mathbf{X}_i^f(t_k)$ is the forecasted state of $\mathbf{X}_i^a(t_{k-1})$. The forecast is taken as the average, that is,

$$\bar{\mathbf{X}}^f(t_k) = \frac{1}{r+1} \sum_{i=1}^{r+1} \mathbf{X}_i^f(t_k), \quad (4.6)$$

and the predicted error covariance is approximated by the sample covariance matrix, that is,

$$\mathbf{P}^f(t_k) = \frac{1}{r+1} \sum_{i=1}^{r+1} [\mathbf{X}_i^f(t_k) - \bar{\mathbf{X}}^f(t_k)] \times [\mathbf{X}_i^f(t_k) - \bar{\mathbf{X}}^f(t_k)]^T + \mathbf{Q}_k. \quad (4.7)$$

In general, $\mathbf{P}^f(t_k)$ is of rank r and can be rewritten from (4.7) as

$$\mathbf{P}^f(t_k) = \mathbf{L}_k [(r+1) \mathbf{T}^T \mathbf{T}]^{-1} \mathbf{L}_k^T + \mathbf{Q}_k, \quad (4.8)$$

where $\mathbf{L}_k$ is $n \times r$ with the i th column given by $\mathbf{X}_i^f(t_k) - \bar{\mathbf{X}}^f(t_k)$ and $\mathbf{T}$ is an $(r+1) \times r$ full-rank matrix with zero column sums. A convenient choice of $\mathbf{T}$ is given in Hoteit et al. (2002) as

$$\mathbf{T} = \begin{pmatrix} \mathbf{I}_{r \times r} \\ \mathbf{0}_{1 \times r} \end{pmatrix} - \frac{1}{r+1} \mathbf{1}_{(r+1) \times r}, \quad (4.9)$$

where $\mathbf{I}_{r \times r}$ is the $r \times r$ identity matrix, $\mathbf{0}_{1 \times r}$ is a $1 \times r$ vector of zeros, and $\mathbf{1}_{(r+1) \times r}$ is an $(r+1) \times r$ matrix of ones.

d. Correction stage

The new data are used to correct the forecast state $\bar{\mathbf{X}}^f(t_k)$ according to

$$\mathbf{X}^a(t_k) = \bar{\mathbf{X}}^f(t_k) + \mathbf{G}_k [\mathbf{Y}_k - H_k \mathbf{X}^f(t_k)], \quad (4.10)$$

where the Kalman gain matrix $\mathbf{G}_k$ is given by

$$\mathbf{G}_k = \mathbf{L}_k \mathbf{U}_k (\mathbf{H} \mathbf{L})_k^T \mathbf{R}_k^{-1}. \quad (4.11)$$

Here $(\mathbf{H} \mathbf{L})_k$ is $p \times r$ with the i th column given by $H_k \mathbf{X}_i^f(t_k) - H_k \mathbf{X}^f(t_k)$ and $\mathbf{U}_k$ given by (after factoring the various decompositions of the matrices above in the standard Kalman gain formula)

$$\mathbf{U}_k^{-1} = \{[(r+1) \mathbf{T}^T \mathbf{T}]^{-1} + (\mathbf{L}_k^T \mathbf{L}_k)^{-1} \mathbf{L}_k^T \mathbf{Q}_k \mathbf{L}_k (\mathbf{L}_k^T \mathbf{L}_k)^{-1}\}^{-1} + [(\mathbf{H} \mathbf{L})_k]^T \mathbf{R}_k^{-1} (\mathbf{H} \mathbf{L})_k. \quad (4.12)$$

The analyzed filter error covariance matrix is given by

$$\mathbf{P}^a(t_k) = \mathbf{L}_k \mathbf{U}_k \mathbf{L}_k^T. \quad (4.13)$$

e. The inflation factor

An inflation factor is a term used to increase (i.e., inflate) the background error covariance. The background error covariance estimated from an ensemble of states often underestimates the true forecast error partly because of the limited number of ensemble members (e.g., see Hamill et al. 2001) where the inflation factor helps to stabilize a filter when small ensemble sizes are used. The use of an inflation factor (taken as the reciprocal of the forgetting factor in some literature) in the SEIK filter is discussed thoroughly in Pham (2001) and was also used as a means of helping the filter tracking model instabilities in Hoteit et al. (2002). Covariance inflation is becoming a necessary part of any ensemble Kalman filter implementation (Hamill et al. 2009) and was recently related to the robust filtering theory by Luo and Hoteit (2011). The forgetting factor is denoted ρ with values between 0 and 1 and the corresponding inflation factor is given by ρ^{-1} . If $\rho < 1$, the inflation factor weighs recent data exponentially more than old data, essentially limiting the memory length of the filter. This is a particularly attractive attribute for our application of the filter to the modeling of extreme events and is one of the primary reasons we use it here. The use of the inflation factor does not change the computational cost of the filter and only updates (4.12) to

$$\mathbf{U}_k^{-1} = \{\rho^{-1} [(r+1) \mathbf{T}^T \mathbf{T}]^{-1} + (\mathbf{L}_k^T \mathbf{L}_k)^{-1} \mathbf{L}_k^T \mathbf{Q}_k \mathbf{L}_k (\mathbf{L}_k^T \mathbf{L}_k)^{-1}\}^{-1} + [(\mathbf{H} \mathbf{L})_k]^T \mathbf{R}_k^{-1} (\mathbf{H} \mathbf{L})_k. \quad (4.14)$$

In the filter results shown below, the inflation factor is set to $\rho = 1.42$. In practice, this term can be updated with the data assimilation, but we set it here to keep the focus

FIG. 2. (left) Track of Hurricane Ike through the Gulf of Mexico. The circles with annotations are the locations of landfall at 1110 UTC 29 Aug 2005 and the locations of the hurricane approximately 48 and 72 h before landfall. (right) Track of Hurricane Katrina through the Gulf of Mexico. The circles with annotations are the locations of landfall at 0710 UTC 13 Sep 2008 and the locations of the hurricane approximately 48 and 72 h before landfall.

on the ability to update water elevations using a fixed data assimilation framework.

5. Observation system simulation experiments

We performed high-resolution hindcast studies to simulate the data used in the assimilation experiments. These studies have been shown to match actual measured data from instruments located throughout the Gulf for a number of recent hurricanes, as described above. For this reason, we are computing what are known as observing system simulation experiments. We take as two particular case studies Hurricanes Ike and Katrina.

Hurricane Ike traveled through the Atlantic, Caribbean, and Gulf of Mexico in September of 2008, finally making landfall along the upper Texas coast in the early morning hours of 13 September (see left-hand plot in Fig. 2). Ike was a category-4 hurricane on 4 September 2008, and was classified as a category-2 hurricane when it made landfall near Galveston, Texas, at 0710 UTC 13 September 2008, resulting in what has been estimated by some as the third costliest hurricane (measured in inflation-adjusted U.S. dollars for 2008) ever to make landfall in the United States (Berg 2009). In the days leading up to and during landfall, the recorded wind intensity for Ike was on the high end of a category-2 hurricane (175 km h^{-1}) with a larger wind field than typical for a category-2 hurricane (Berg 2009). Not only did Ike cause incredible monetary damage, estimated at nearly \$29.6 billion in the United States alone, but it is also responsible for nearly 200 deaths (Berg 2009).

Hurricane Katrina formed in August of 2005 over the Bahamas, made its first landfall in southern Florida, traveled through the Gulf of Mexico, and made its

second landfall in southeast Louisiana on the morning of 29 August (see right-hand plot in Fig. 2). Katrina was a category-3 hurricane upon making its second landfall. Katrina became the costliest natural disaster in the history of the United States, causing an estimated \$81 billion in property damage (Knabb et al. 2005), and is also responsible for over 1800 confirmed deaths. Most of the destruction was due to the storm surge.

We give the specific details for each OSSE below, but the general outline is as follows:

- A hindcast of the hurricane is run and observation data are extracted/stored every 2 h
- The oceans are ramped up on the coarse grid of the Gulf of Mexico (see Fig. 3)
- The “ramped up” state is taken as the mean state of the system
- The first data assimilation cycle is computed on a 2-h forecast of the ensemble members, then
  - (i) the analyzed ensemble members are forecasted forward 2 h on the coarse grid
  - (ii) an analyzed ensemble of states is either computed (i.e., data are assimilated), or the forecasted ensemble members are used in place of the analyzed states to compute the next 2-h forecast (i.e., no data are assimilated)
  - (iii) repeat (i) and (ii) until the end of the storm is reached

a. Hindcast details

The ADCIRC model hindcast runs use a time step of 1 s on high-resolution grids of the same domain covering the Gulf of Mexico and the western North Atlantic seaboard (see Fig. 1). The hindcasts use data assimilated winds and atmospheric pressure fields provided by Ocean

FIG. 3. Discretization of the Gulf of Mexico domain containing 8006 nodes and 14 269 elements. Open boundaries are denoted by bold boundary lines. All other boundaries are land. The open boundaries are forced by the three most dominant diurnal and two most dominant semidiurnal tidal constituents in terms of amplitude. The land boundaries are reflective.

Weather, Inc. (OWI). These are the same wind fields used in previous hindcast studies of Katrina (Dietrich et al. 2010) and Ike (Kennedy et al. 2011). The hindcasts were run on 2038 CPUs of the Ranger parallel computer at the Texas Advanced Computing Center and required approximately 6 h to complete.

The Hurricane Ike hindcast was run on a grid of 3 322 439 nodes corresponding to 6 615 381 elements. The Hurricane Katrina hindcast was run on a grid of 5 035 113 nodes corresponding to 9 945 623 elements.

The measurement data obtained from these hindcast studies are, in practice, considered truth since the data can be collected continuously and the noise can be easily filtered. However, we err on the side of caution and set the measurement noise in the experiments assuming a 95% confidence interval of ± 0.01 m.

b. Simulation experiment details

For each hurricane, we compute several simulations. First, we compute the “no data assimilated” forecasts of the ensemble members for each hurricane where the initial ensemble members are never updated by the data (i.e., the forecasted states are stored every 2 h and are used as the “analysis” states by the model to generate another forecast). We then compute simulations where data are assimilated into the system every 2 h until approximately 48 h before the hurricane makes landfall, at which point the forecasted states are no longer updated by the data. Similarly, we compute simulations where data are assimilated into the system every 2 h until approximately 24 h before each hurricane makes landfall. Finally, we compute simulations where data are assimilated into the system every 2 h throughout the entire

timeline of each hurricane, which we refer to as the “full data assimilated” forecasts.

Each simulation is abbreviated by a tag for ease of reference. We use ISim-ND, ISim-48, ISim-24, and ISim-FD to denote Hurricane Ike simulations. Similarly, we use KSim-ND, KSim-48, KSim-24, and KSim-FD for Hurricane Katrina simulations. We introduce an additional simulation for Katrina, denoted KSim-36 (for data assimilated up to approximately 36 h before landfall), as no significant gains in data assimilation are recorded for the KSim-48 simulation. The suffix ND stands for the no data assimilated simulation, the 48 and 24 indicate the simulations where data are assimilated every 2 h until approximately 48 and 24 h before the hurricane makes landfall, and the FD stands for the full data assimilated experiment.

The forecasts for these simulations use a dynamic Holland wind model (Holland 1980), which generates a parametric wind field using hurricane data (central pressure, maximum wind speed, radius of maximum winds), using the best track data obtained from the NOAA archive (<ftp://ftp.tpc.ncep.noaa.gov/atcf/archive/>) where the radii of maximum winds are defined by the significant radii of the 1-m sustained 34-kt isotach ($1 \text{ kt} = 0.5144 \text{ m s}^{-1}$), given in the advisory. Thus the forecasts use a different wind field than the hindcast, which would most likely be the case in any actual hurricane event and is a source of error in the model.

The data assimilation experiments are run using a time step of 10 s on a grid of 8006 nodes and 14 269 elements covering the Gulf of Mexico (see Fig. 3). Since Hurricanes Ike and Katrina have distinct tracks (see plots in Fig. 2), we use separate arrays of observation stations for tracking Ike and Katrina (see section 6; Figs. 4, 5). These stations were chosen using a two-step process. First, we produced dense arrays of hypothetical stations whose data were extracted from the hindcast studies. Then, the stations used in the assimilation experiments were chosen from these arrays if they contained a clear signal of the surge (e.g., the station data shown in Fig. 6). Future studies may include the use of all hypothetical station data where data indicating a lack of storm surge are assimilated into the model to correct forecasts that overpredict water elevations in certain areas, but this is beyond the scope of this current study as we are primarily interested in improving maximum storm surge predictions.

By using separate domains, discretizations, wind forcing, and time steps for the simulation of true data versus the data assimilation experiments, we avoid many possible inverse crimes. We summarize the main differences between the hindcast simulations taken as truth and the simulations used in the data assimilation run in Table 2.

FIG. 4. The 371 hypothetical observation stations used for Hurricane Ike simulations. The red-filled circles denote an array of stations where data are simulated from the hindcast study. The black-filled circles correspond to 43 existing measurement sites placed by several local, state, and federal agencies, including the U.S. Geological Survey (USGS) and NOAA, and university research groups (Kennedy et al. 2011). These 43 observation stations collected data during the storm and were used to validate the hindcast study.

6. Numerical results

In the results presented below, we quantify the effect the data assimilation scheme (as discussed above) has on the forecasts of storm surge in several ways. One obvious metric is the rms error of the ensemble member forecasts for the various simulations. However, this one value does not entirely sum up the improvements to the forecasted states. As mentioned above, a particular quantity of interest is the prediction of maximum storm surge along coastal regions over the entire duration of a storm. We demonstrate this improvement graphically for each storm and simulation. There is also interest in the prediction of the storm surge at particular times, specifically in the hours before a hurricane makes

landfall. We also graphically demonstrate the improvement in the forecasts for the various simulations for each storm. To properly distinguish the gains of the data assimilation scheme in terms of improving forecasts at particular times or of maximum water elevations, we establish for each storm the “baseline error” defined as the difference between the truth (taken from the hindcast study) and the no data assimilated simulations (i.e., what the error would be if we simply make forecasts with no updating of the state). We then show plots of the difference between the various data assimilation simulations and the no data assimilated simulation for each storm, so that the improvement over this baseline prediction is made clear. In plots of state estimates, we chose a grayscale with discontinuities between small

FIG. 5. The 559 hypothetical observation stations used for Hurricane Katrina simulations. The red-filled circles denote an array of stations where data are simulated from the hindcast study.

FIG. 6. Plots of simulated station data in meters vs assimilation cycle where data are used to update the forecast. The stars mark data that are simulated from the hindcast study of Hurricane Ike (and then assimilated every 2 h beginning at 0200 UTC 10 Sep and ending at 0600 UTC 14 Sep 2008), the analyzed state (circles; from ISim-FD), and forecast state (plus signs; from ISim-FD) with 95% confidence intervals (vertical dashed lines centered at plus signs). Station locations given in (lat, lon): (top left) (29.63°N, 89.56°W), (top right) (29.53°N, 92.31°W), (bottom left) (27.9°N, 93.4°W), (bottom right) (29.1°N, 92.2°W).

ranges of hues to more clearly demonstrate the differences between the predictions in the various simulations.

a. Hurricane Ike

For all Hurricane Ike simulations, the ramp-up is set to 1 day starting at 0000 UTC 9 September 2008 and ending at 0000 UTC 10 September 2008, at which point all the ensemble members for all simulations are identically initialized. All Ike simulations end at 0600 UTC 14 September 2008, nearly 24 h after landfall occurred at 0710 UTC 13 September 2008. For ISim-48, data are assimilated every 2 h until 0800 UTC 11 September 2008, resulting in a total of 16 assimilation cycles being computed. For ISim-24, data are assimilated every 2 h until 0800 UTC 12 September 2008, resulting in a total of 28 assimilation cycles being computed.

In Fig. 6, we show the hydrographs of simulated data from the Ike hindcast study at four stations that represent a generic sampling of the type of data available to us

during the simulations. Two of the stations shown in this figure correspond to just two of the many existing stations that collected true data during Hurricane Ike and were used to verify and validate the hindcast study. In these hydrographs, the stars denote the true measurements at the assimilation times, the plus signs denote the forecasted measurement data, and the circles are the

TABLE 2. Summary of differences between simulations for hindcast (truth) simulation used to generate data and simulations used in data assimilation forecasting experiments.

	Data assimilation	Truth
Domain	Gulf of Mexico	Western North Atlantic
Avg mesh element size (km ²)	98	1.34
Time step (s)	10	1
Wind field	Dynamic Holland	OWI
Bottom friction formulation	Chezy	Hybrid

FIG. 7. Plots of error in max free-surface elevation predictions (in meters) from true state for (top left) ISim-ND, (top right) ISim-48, (bottom left) ISim-24, and (bottom right) ISim-FD. The top left plot is the baseline error case, and it represents the error of the existing system if the coarse grid was used to generate a sequence of 2-h forecasts with no data assimilation.

analyzed state's data. We also plot the 95% confidence intervals of the forecasted data, which are computed using (4.8). We see that in general there is a consistent underestimation of the forecast relative to the true data and the data assimilation corrects toward this truth. This is expected for a variety of reasons since the model is dissipative, and while the data generated in the hindcast study are from a fine grid with 1-s time steps, we are forecasting using a coarse grid and 10-s time steps, which explains a consistent underestimation. Even if we used the same spatial grids and time steps, we would not expect the forecast of data to match the true data since we are using a wind model consistent with what is available during a storm in the forecast computations, but the hindcast uses what may be considered the exact wind field.

One way to quantify the improvement is the use of rms error. We use the hindcast study to obtain the global true state of the system and compare the error of the

forecasted ensemble members. We are particularly interested in the ability to forecast the maximum coastal surge. We consider two rms error metrics. First, we consider the rms error of the maximum water elevation forecasts in the geographic area shown in Figs. 7–10, which contains the parts of the Texas coast around the landfall event that witnessed the largest storm surge. We see, in Table 3, that by assimilating data up to 24–48 h before landfall we obtain improvements in the coastal rms error of approximately 3%–5%, and we obtain approximately a 14% improvement in coastal rms error that the forecasts from ISim-FD give during this same time window of maximum surge. Second, we consider the rms error of the maximum water elevation forecasts at all nodes that recorded up to 60% of the maximum surge from the true forecast, that is, we consider the rms error at roughly 36.5% of the nodes where the maximum water elevation forecast of the hindcast study exceeded 3 m. Table 3 shows that improvement in the rms error of

FIG. 8. Plots of difference (in meters) between max free-surface elevation of (top) ISim-48 and ISim-ND, (middle) ISim-24 and ISim-ND, and (bottom) ISim-FD and ISim-ND. Observe that the bottom plot has a different color scale to better demonstrate the differences in the ISim-FD simulation.

FIG. 9. Plot of free-surface elevation error (in meters) at 0700 UTC 13 Sep 2008 from truth for (top) ISim-ND, (middle) ISim-24, and (bottom) ISim-FD. The top plot is the baseline error case, and it represents the error of the existing system if the coarse grid was used to generate a sequence of 2-h forecasts with no data assimilation.

FIG. 10. Plots of difference (in meters) of free-surface elevation forecasts at 0700 UTC 13 Sep 2008 between (left) ISim-24 and ISim-ND and (right) ISim-FD and ISim-ND. Observe that different scales are used to better demonstrate the differences between the plots.

the maximum water elevation at these nodes is approximately 5.8% for the ISim-48 and ISim-24 simulations and 15% for the ISim-FD simulation.

The rms metric only tells part of the story. We are interested in forecasts of maximum water elevations and forecasts of water elevations at particular times, specifically the times leading up to landfall. First, we look at forecasts of maximum water elevations. In Fig. 7, we see the errors from the true forecast of maximum water elevation in the ISim-ND, ISim-48, ISim-24, and ISim-FD forecasts of maximum water elevations. We show in the plots in Fig. 8 the improvements to the baseline (i.e., ISim-ND) prediction of maximum water elevation computed from ISim-48, ISim-24, and ISim-FD. We see in the top plot in Fig. 8 (and comparing the top right and top left plots in Fig. 7) that ISim-48 already offers some improvement to future forecasts in regions along the coast of up to 0.2 m. In the middle plot in Fig. 8 (i.e., the difference between the bottom left and top left plots in Fig. 7), we see that for the ISim-24 simulation there is a slight improvement from the ISim-48 simulation forecast where improvements to maximum elevation forecasts are as much as 0.25 m over some areas. Comparing the middle and bottom plots in Fig. 8, we see that ISim-FD more accurately captures the maximum surge within the bay area, with improvements of at least 0.5 m at many locations, as would be expected since data are used through the entire simulation to improve future forecasts (also, compare the differences between the bottom right and top left and the bottom left and top left plots in Fig. 7 to observe these differences). There exist a few small isolated regions where the ISim-48, ISim-24, and ISim-FD forecasts of maximum water elevation are approximately 0.01–0.05 m lower than the forecast of the ISim-ND simulation. These regions occur where the

error in forecasts is generally smaller compared to surrounding areas that show significant improvement.

We see in Fig. 9 the errors from the true forecast of water elevations at 0600 UTC 13 September 2008, which is an hour before Ike made landfall at 0710 UTC of the same day, in the ISim-ND, ISim-24, and ISim-FD forecasts. The differences between the ISim-48 and ISim-24 forecasts proved to be within ± 0.05 m over most of the coastal region at this particular time, so the results of the ISim-48 forecasts are omitted in the interest of space. In the plots in Fig. 10, we see the improvements to forecasts of water elevation using data assimilation simulations compared to the baseline prediction at this time. The ISim-24 simulation (see top plot in Fig. 10 and also compare to middle plot in Fig. 9) predicts up to approximately 0.2-m-higher water elevation in some areas across the coastline. The ISim-FD simulation (see bottom plot in Fig. 10 and also compare to top and bottom plots in Fig. 9) does the best job of improving water elevation predictions in inlet and bay areas by improving the error by at least 0.5 m over most of the bay and up to

TABLE 3. The average rms errors of the max water elevation forecasts are shown for the four distinct Hurricane Ike simulations. The second column shows the average rms errors of the max water elevations forecasted in the geographic area shown in Fig. 8. The third column shows the average rms errors of the forecasted max water elevations at the 2927th of 8006 nodes where the hindcast study recorded max water elevations greater than 3 m.

Simulation	Coastal rms error	Rms error for surge >3 m
ISim-ND	1.92	1.91
ISim-48	1.86	1.81
ISim-24	1.82	1.80
ISim-FD	1.65	1.62

TABLE 4. The average rms errors of the forecasted max water elevations. The second column shows the average rms errors at the 405th of 8006 nodes where the hindcast study recorded surge greater than 4 m. The third column shows the average rms errors at the 163rd of 8006 nodes where the hindcast study recorded surge greater than 5 m.

Simulation	Rms error for surge >4 m	Rms error for surge >5 m
KSim-ND	2.43	2.94
KSim-36	2.03	2.48
KSim-24	1.99	2.48
KSim-FD	1.92	1.89

0.25 m on the coastline area that is directly west of the bay area.

b. Hurricane Katrina

For all Hurricane Katrina simulations, the ramp-up is set to 6 h starting at 0000 UTC and ending at 0600 UTC 25 August 2005, at which point all the ensemble members for all simulations are identically initialized. All Katrina simulations end at 1200 UTC 30 August 2005, approximately 24 h after landfall occurred at 1110 UTC 29 August 2005. For KSim-48, data are assimilated every 2 h until 1200 UTC 27 August 2005, resulting in a total of 26 assimilation cycles being computed. However, we saw minimal to no improvement from these simulations compared to the KSim-ND simulation. Therefore, we assimilated data another 12 h (at 2-h intervals) to generate a simulation of forecasts where data are assimilated approximately 36 h before landfall, and we denote this simulation as KSim-36. For KSim-24, data are assimilated every 2 h until 1200 UTC 28 August 2005, resulting in a total of 26 assimilation cycles being computed.

As before, we first quantify the improvement using the rms error, and we are again particularly interested in

the ability to forecast the maximum coastal surge. Unlike Hurricane Ike, which had a 3–5-m storm surge over a large geographic area, Katrina saw a significantly higher storm surge, from 4 to approximately 6.6 m over a smaller area as determined by the hindcast. Specifically, for Hurricane Ike, 36.5% of the nodes recorded water elevations within 60% of the overall maximum value, but for Katrina, only 5.1% of the nodes recorded water elevations within 60% of the overall maximum value. Accurate estimates of extreme storm surge over small areas are of particular importance. We see, in Table 4, that on average the forecasts from KSim-36, KSim-24, and KSim-FD decrease the rms errors of the forecasted maximum storm surge at nodes with at least 4 m of surge by approximately 16%–21%. Also from this table we see that at nodes where the maximum storm surge exceeded 5 m the error was reduced by 16%–36%. The rms errors for these simulations are more informative as to the significant improvement in capturing maximum storm surge than the plots below indicate because of the small number of nodes recording the largest surge.

In Fig. 11, we show the errors from the true forecast of maximum water elevation in the KSim-ND and KSim-36 forecasts of maximum water elevations. The extreme localization of the maximum surge makes it difficult to visually identify significant differences between the KSim-24/KSim-FD and the KSim-36 forecasts as can be seen in the rms metrics discussed above, so these are omitted in the interest of space. We see in the plot in Fig. 12 the improvements to the baseline prediction of maximum water elevation computed from the KSim-36 simulation. In the nearshore areas where the storm surge error in the baseline case was greater than 2 m, the KSim-36 simulation often improves the error by approximately 0.5 m. In areas where the error in the

FIG. 11. Plots of error in max free-surface elevation predictions (in meters) from true state for (left) KSim-ND and (right) KSim-36. The left plot is the baseline error case, and it represents the error of the existing system if the coarse grid was used to generate a sequence of 2-h forecasts with no data assimilation.

FIG. 12. Plot of difference (in meters) between max free-surface elevation of KSim-36 and KSim-ND.

storm surge was less than 1.75 m, which generally correspond to areas of smaller surge, the KSim-36 produced forecasts of maximum water elevation very similar to the KSim-ND results.

We now consider forecasts of water elevations at the specific time of 0800 UTC 29 August. In Fig. 13, we see the errors from the true forecast in the KSim-ND and KSim-36 forecasts at this particular time. In the plots in Fig. 14, we see the improvement to the forecast error in the KSim-ND forecast using the KSim-36 simulation. We see that the KSim-36 simulation most improves the error in the area of greatest storm surge and forecast error. We again observe that the largest errors and storm surge are recorded within a small localized area within the domain.

c. Computational cost

The simulations and data assimilation were executed on the local cluster Bevo2 at The University of Texas at Austin. Bevo2 is a 23-node compute cluster made up of Dell PowerEdge servers that house $2 \times$ quad core 2.66-GHz Intel Xeon processors for a total of 184 cores. Each node has 16 gigabytes of RAM, dual-gigabit ethernet ports, and a single-port Mellanox III Lx Infiniband adapter attached to a QLogic SilverStorm Infiniband 24-port switch capable of up to 20 Gb s^{-1} . Each ensemble member was executed in parallel using only eight processors. The ensemble member jobs were submitted to the machine sequentially. The data assimilation was performed in Matlab. The total wall time for the experiment including submitting jobs to the queue was under 90 min for all Katrina and Ike experiments.

7. Conclusions

We have successfully implemented the square root ensemble Kalman filter known as the SEIK filter to the ADCIRC model formulation for two hurricanes: Ike and Katrina. Validated hindcast studies over the western North Atlantic and Gulf of Mexico were used to generate the true states and observations, and a coarser resolution for the Gulf of Mexico was used for the data assimilation simulations. We have shown that there is promise in improving forecasts up to 48 h before landfall in some cases. Specifically, decreased rms error, improved forecasts of maximum water elevation, and improved forecasts of surge in the hours before landfall were all demonstrated for simulations where data were assimilated only up to 24–48 h before landfall. Furthermore, the computational cost was quite modest in

FIG. 13. Plot of free-surface elevation error (in meters) at 0800 UTC 29 Aug 2005 from truth for (left) KSim-ND and (right) KSim-36. The left is the baseline error case, and it represents the error of the existing system if the coarse grid was used to generate a sequence of 2-h forecasts.

FIG. 14. Plot of difference (in meters) of free-surface elevation forecasts at 0800 UTC 29 Aug 2005 between KSim-36 and KSim-ND.

comparison with running the hindcast studies in forecast mode.

The approach that we describe in this paper is general and can be applied to storms of various sizes and intensities. To perform a data assimilated forecast, one needs the ADCIRC model input files, primarily the finite element mesh, bathymetry, and other parameters (bottom friction coefficients, eddy viscosity, etc.); these data exist for many areas of the coastal United States. One also needs wind data, generally available from NHC, and of course the measurement data to assimilate. One concern with storms of varying size would be whether the resolution of the finite element mesh is sufficient to capture the storm. For the mesh used in the Katrina and Ike simulations presented here, the average mesh diameter was approximately 10 km. Note that the same mesh was used for the assimilation in both storms, unlike the hindcast runs where different meshes were used that were highly resolved in either Texas or Louisiana. However, a narrower or weaker storm might require a higher-resolution mesh to accurately capture the hurricane-force winds and the more local nature of the storm surge. In addition, different meshes require a different EOF analysis. The computational cost of the EOF analysis may be significant, but for each mesh, this analysis can be done once, stored, and used for future simulations. Furthermore, how well the assimilation would improve results for storms without a significant surge signal in the data is of some interest. We will further explore the capabilities of the assimilation framework we have developed for storms of varying size, location, and intensity in future work.

The results show much promise moving forward with the goal of implementing data assimilation methodologies into real-time forecasts of storm surge. Future work

includes producing a more complete and robust hurricane forecasting system, whereby we assimilate not only water levels, but wind and significant wave heights.

Acknowledgments. This work was supported by an Academic Excellence Alliance Grant with KAUST, the King Abdullah University of Science and Technology. This support is gratefully acknowledged. The authors thank Dr. Joel C. Dietrich and Dr. Joannes J. Westerink for providing data, the ADCIRC model code, and guidance for the hindcast simulations of Hurricanes Ike and Katrina used in this paper.

REFERENCES

- Berg, R., 2009: Tropical cyclone report: Hurricane Ike (AL092008) 1–14 September 2008. National Hurricane Center Rep., 55 pp.
- Bishop, C. H., B. J. Etherton, and S. J. Majumdar, 2001: Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Wea. Rev.*, **129**, 420–436.
- Blake, E. S., C. W. Landsea, and E. J. Gibney, 2011: The deadliest, costliest, and most intense United States tropical cyclones from 1851 to 2010 (and other frequently requested hurricane facts). NOAA Tech. Memo. NWS NHC-6, 49 pp.
- Brown, J. D., T. Spencer, and I. Moeller, 2007: Modeling storm surge flooding of an urban area with particular reference to modeling uncertainties: A case study of Canvey Island, United Kingdom. *Water Resour. Res.*, **43**, W06402, doi:10.1029/2005WR004597.
- Bunya, S., and Coauthors, 2010: A high-resolution coupled riverine flow, tide, wind, wind wave, and storm surge model for southern Louisiana and Mississippi. Part I: Model development and validation. *Mon. Wea. Rev.*, **138**, 345–377.
- Cane, M., A. Kaplan, R. Miller, B. Tang, E. C. Hackert, and A. Busalacchi, 1996: Mapping tropical Pacific sea level: Data assimilation via a reduced state space Kalman filter. *J. Geophys. Res.*, **101** (C10), 599–617.
- Dee, D. P., 1991: Simplification of the Kalman filter for meteorological data assimilation. *Quart. J. Roy. Meteor. Soc.*, **117**, 365–384.
- Dietrich, J. C., and Coauthors, 2010: A high-resolution coupled riverine flow, tide, wind, wind wave, and storm surge model for southern Louisiana and Mississippi. Part II: Synoptic description and analysis of Hurricanes Katrina and Rita. *Mon. Wea. Rev.*, **138**, 378–404.
- , and Coauthors, 2011: Hurricane Gustav (2008) waves and storm surge: Hindcast, synoptic analysis, and validation in southern Louisiana. *Mon. Wea. Rev.*, **139**, 2488–2522.
- El Serafy, G. Y. H., and A. E. Mynett, 2008: Improving the operational forecasting system of the stratified flow in Osaka Bay using an ensemble Kalman filter-based steady state Kalman filter. *Water Resour. Res.*, **44**, W06416, doi:10.1029/2006WR005412.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99** (C5), 10 143–10 162.
- Fleming, J., C. Fulcher, R. Luettich, B. Estrade, G. Allen, and H. Winer, 2008: A real time storm surge forecasting system using ADCIRC. *Estuarine and Coastal Modeling: Proceedings of the Tenth International Conference*, M. L. Spaulding, Ed., ASCE, 893–912.
- Fukumori, I., and P. Malanotte-Rizzoli, 1995: An approximate Kalman filter for ocean data assimilation: An example with an

- idealized Gulf Stream model. *J. Geophys. Res.*, **100** (C4), 6777–6793.
- Gerritsen, H., H. de Vries, and M. Philippart, 1995: The Dutch Continental Shelf Model. *Quantitative Skill Assessment for Coastal Ocean Models, Coastal and Estuarine Studies*, Vol. 47, D. R. Lynch and A. M. Davies, Eds., Amer. Geophys. Union, 425–467.
- Ghil, M., 1989: Meteorological data assimilation for oceanographers. Part I: Description and theoretical framework. *Dyn. Atmos. Oceans*, **13**, 331–376.
- Hamill, T., J. Whitaker, and C. Snyder, 2001: Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.*, **129**, 2776–2790.
- , —, J. L. Anderson, and C. Snyder, 2009: Comments on “Sigma-point Kalman filter data assimilation methods for strongly nonlinear systems.” *J. Atmos. Sci.*, **66**, 3498–3500.
- Heap, N. S., 1983: Storm surges 1967–1982. *Geophys. J. Roy. Astron. Soc.*, **74**, 331–376.
- Heemink, A. W., 1986: Storm surge prediction using Kalman filtering. Ph.D. thesis, Twente University of Technology, 194 pp.
- , and H. Kloosterhuis, 1990: Data assimilation for non-linear tidal models. *Int. J. Numer. Methods Fluids*, **11**, 1097–1112.
- Hoang, S., R. Baraille, O. Talagrand, X. Carton, and P. D. Mey, 1998: Adaptive filtering: Application to satellite data assimilation in oceanography. *Dyn. Atmos. Oceans*, **27** (1–4), 257–281, doi:10.1016/S0377-0265(97)00014-6.
- Holland, G., 1980: An analytic model of the wind and pressure profiles in hurricanes. *Mon. Wea. Rev.*, **108**, 1212–1218.
- Hoteit, I., D. Pham, and J. Blum, 2002: A simplified reduced order Kalman filtering and application to altimetric data assimilation in tropical Pacific. *J. Mar. Syst.*, **36**, 101–127.
- , G. Korres, and G. Triantafyllou, 2005: Comparison of extended and ensemble based Kalman filters with low and high-resolution primitive equations ocean models. *Nonlinear Processes Geophys.*, **12**, 755–765.
- Julier, S., and J. Uhlmann, 1997: A new extension of the Kalman filter to nonlinear systems. *Proc. AeroSense: 11th Int. Symp. on Aerospace/Defence Sensing, Simulation and Controls*, Orlando, FL, SPIE, 182–193.
- Kalman, R. E., 1960: A new approach to linear filtering and prediction problems. *J. Basic Eng.*, **82**, 35–45.
- Kennedy, A. B., and Coauthors, 2011: Origin of the Hurricane Ike forerunner surge. *Geophys. Res. Lett.*, **38**, L08608, doi:10.1029/2011GL047090.
- Kinnmark, I., 1986: *The Shallow Water Wave Equations: Formulation, Analysis, and Applications*. Vol. 15, *Lecture Notes in Engineering*, Springer-Verlag, 187 pp.
- Knabb, R. D., J. R. Rhome, and D. P. Brown, 2005: Tropical cyclone report: Hurricane Katrina, 23–30 August 2005. National Hurricane Center Rep., 43 pp.
- Luettich, R., and J. Westerink, 2004: Formulation and numerical implementation of the 2D/3D ADCIRC finite element model version 44.XX. ADCIRC Tech. Rep., 74 pp.
- Luo, X., and I. Hoteit, 2011: Robust ensemble filtering and its relation to covariance inflation in the ensemble Kalman filter. *Mon. Wea. Rev.*, **139**, 3938–3953.
- Lynch, D. R., and W. G. Gray, 1979: A wave equation model for finite element tidal computations. *Comput. Fluids*, **7**, 207–228.
- Malanotte-Rizzoli, P., R. E. Young, and D. B. Haidvogel, 1989: Initialization and data assimilation experiments with a primitive equation model. *Dyn. Atmos. Oceans*, **13**, 349–378.
- Maybeck, P. S., Ed., 1979: *Stochastic Models, Estimation, and Control*. Mathematics in Science and Engineering Series, Vol. 141, Academic Press, 411 pp.
- McRobie, A., T. Spencer, and H. Gerritsen, 2005: The Big Flood: North Sea storm surge. *Philos. Trans. Roy. Soc. London*, **A363**, 1263–1270.
- Mitchell, H. L., and P. L. Houtekamer, 2000: An adaptive ensemble Kalman filter. *Mon. Wea. Rev.*, **128**, 416–433.
- Murty, T. S., R. A. Flather, and R. F. Henry, 1986: The storm surge problem in the Bay of Bengal. *Prog. Oceanogr.*, **16**, 195–233.
- Pham, D., 2001: Stochastic methods for sequential data assimilation in strongly nonlinear systems. *Mon. Wea. Rev.*, **129**, 1194–1207.
- , J. Verron, and M. Roubaud, 1997: Singular evolutive Kalman filter with EOF initialization for data assimilation in oceanography. *J. Mar. Syst.*, **16**, 323–340.
- Sakov, P., and P. R. Oke, 2008: A deterministic formulation of the ensemble Kalman filter: An alternative to ensemble square root filters. *Tellus*, **60A** (2), 361–371, doi:10.1111/j.1600-0870.2007.00299.x.
- Sorensen, J. V. T., H. Madsen, and H. Madsen, 2004: Efficient Kalman filter techniques for the assimilation of tide gauge data in three-dimensional modeling of the North Sea and Baltic Sea system. *J. Geophys. Res.*, **109**, C03017, doi:10.1029/2003JC002144.
- , —, and —, 2006: Parameter sensitivity of three Kalman filter schemes for assimilation of water levels in shelf sea models. *Ocean Modell.*, **11**, 441–463.
- Tanaka, S., S. Bunya, J. J. Westerink, C. Dawson, and R. A. Luettich Jr., 2011: Scalability of an unstructured grid continuous Galerkin based hurricane storm surge model. *J. Sci. Comput.*, **46**, 329–358.
- Tippett, M., J. Anderson, C. Bishop, T. Hamill, and J. Whitaker, 2003: Ensemble square root filters. *Mon. Wea. Rev.*, **131**, 1485–1490.
- Verlaan, M., and A. W. Heemink, 1997: Tidal flow forecasting using reduced rank square root filters. *Stochastic Hydrol. Hydraul.*, **11**, 349–368.
- Vreugdenhil, C. B., 1994: *Numerical Methods for Shallow-Water Flow*. Kluwer Academic, 261 pp.
- Wan, E. A., and R. van der Merwe, 2000: The unscented Kalman filter for nonlinear estimation. *Proc. Adaptive Systems for Signal Processing, Communications, and Control Symp.*, Lake Louise, AB, Canada, IEEE, 153–158.
- Westerink, J. J., and Coauthors, 2008: A basin- to channel-scale unstructured grid hurricane storm surge model applied to southern Louisiana. *Mon. Wea. Rev.*, **136**, 833–864.