



Article (refereed) – Published version

Wilson, Chris; Horsburgh, Kevin J.; Williams, Jane; Flowerdew, Jonathan; Zanna, Laure. 2013 Tide-surge adjoint modeling: A new technique to understand forecast uncertainty. *Journal of Geophysical Research: Oceans*, 118 (10). 5092-5108. [10.1002/jgrc.20364](https://doi.org/10.1002/jgrc.20364)

This version available at <http://nora.nerc.ac.uk/504367/>

NERC has developed NORA to enable users to access research outputs wholly or partially funded by NERC. Copyright and other rights for material on this site are retained by the rights owners. Users should read the terms and conditions of use of this material at <http://nora.nerc.ac.uk/policies.html#access>

AGU Publisher statement: An edited version of this paper was published by AGU. Copyright (2013) American Geophysical Union. Further reproduction or electronic distribution is not permitted.

Wilson, Chris; Horsburgh, Kevin J.; Williams, Jane; Flowerdew, Jonathan; Zanna, Laure. 2013 Tide-surge adjoint modeling: A new technique to understand forecast uncertainty. *Journal of Geophysical Research: Oceans*, 118 (10). 5092-5108. [10.1002/jgrc.20364](https://doi.org/10.1002/jgrc.20364)

To view the published open abstract, go to <http://dx.doi.org/10.1002/jgrc.20364>

Contact NOC NORA team at
publications@noc.soton.ac.uk

Tide-surge adjoint modeling: A new technique to understand forecast uncertainty

Chris Wilson,¹ Kevin J. Horsburgh,¹ Jane Williams,¹ Jonathan Flowerdew,² and Laure Zanna³

Received 5 February 2013; revised 13 August 2013; accepted 15 August 2013; published 8 October 2013.

[1] For a simple dynamical system, such as a pendulum, it is easy to deduce where and when applied forcing might produce a particular response. However, for a complex nonlinear dynamical system such as the ocean or atmosphere, this is not as obvious. Knowing when or where the system is most sensitive, to observational uncertainty or otherwise, is key to understanding the physical processes, improving and providing reliable forecasts. We describe the application of adjoint modeling to determine the sensitivity of sea level at a UK coastal location, Sheerness, to perturbations in wind stress preceding an extreme North Sea storm surge event on 9 November 2007. Sea level at Sheerness is one of the most important factors used to decide whether to close the Thames Flood Barrier, which protects London. Adjoint modeling has been used by meteorologists since the 1990s, but is a relatively new technique for ocean modeling. It may be used to determine system sensitivity beyond the scope of ensemble modeling and in a computationally efficient way. Using estimates of wind stress error from Met Office forecasts, we find that for this event total sea level at Sheerness is most sensitive in the 3 h preceding the time of its unperturbed maximum level and over a radius of approximately 300 km. We also find that the pattern of sensitivity follows a simple sequence when considered in the reverse-time direction.

Citation: Wilson, C., K. J. Horsburgh, J. Williams, J. Flowerdew, and L. Zanna (2013), Tide-surge adjoint modeling: A new technique to understand forecast uncertainty, *J. Geophys. Res. Oceans*, 118, 5092–5108, doi:10.1002/jgrc.20364

1. Introduction

[2] For extreme storm surge events, such as the 1 February 1953 North Sea flooding disaster, the 9 November 2007 UK east coast floods and the flooding of New York by tropical storm Sandy on 29 October 2012, the capability of the forecast system is critically tested. Coastal sea level forecasting plays a central role in coastal defence, both in active defence in guidance of sea wall design and in reactive response guidance when flooding is inevitable and evacuation may be required. Although operational surge forecasting is a mature discipline, and many countries have effective operational systems [Horsburgh *et al.*, 2011], any improvement in skill in the observational and forecast system can translate into many more lives protected and infrastructure saved. The main challenges are to determine where, when, and what to measure and to understand the fundamental dynamical processes involved in storm surge generation and modeling. The pres-

ence of nonlinear dynamical processes, such as advection, and the large degrees of freedom of the atmosphere-ocean system can pose a problem for designing an experimental strategy using traditional methods, such as perturbed-member ensemble modeling, so we must seek new approaches.

[3] Ocean modeling techniques often evolve to reflect those first used in atmospheric modeling where the system has better observational constraints and where computer models are more capable of resolving the energetic and nonlinear synoptic scales (equivalent to ocean mesoscales). The typical method for tide-surge forecasting [e.g., Flather, 1976; Verlaan *et al.*, 2005] is to force a nonlinear shallow water ocean model at its open lateral boundaries with tides and at its surface with meteorological forcing. This deterministic model is integrated forward in time to produce a single sea level forecast of tide and surge. However, the chaotic nature of the atmosphere and ocean means that small uncertainties about the system (because observations are not perfect and they are interpolated onto the model grid) at the initial time and during the forecast may grow to produce large errors near the end of the forecast, meaning that a single forecast may be atypical and have low skill on occasion [e.g., Lorenz, 1963; Penduff *et al.*, 2011]. Initial condition sensitivity, together with sensitivity to boundary conditions and physical parameters, are characteristics of chaotic systems. To address the atmospheric component of this dynamical systems problem, weather forecasting has developed the application of singular vector analysis and ensemble forecasting techniques [e.g., Lorenz, 1963;

¹National Oceanography Centre, Liverpool, UK.

²Met Office, Fitzroy Road Exeter, UK.

³Dept. of Atmospheric, Oceanic and Planetary Physics, University of Oxford, Oxford, UK.

Corresponding author: C. Wilson, National Oceanography Centre, Joseph Proudman Building, 6 Brownlow St., Liverpool L3 5DA, UK. (cwi@noc.ac.uk)

©2013. American Geophysical Union and Crown copyright. This article is published with the permission of the Controller of HMSO and the Queen's Printer for Scotland.
2169-9275/13/10.1002/jgrc.20364

Epstein, 1969; Blumenthal, 1991; Palmer et al., 1994; Buizza and Palmer, 1998]. Carefully chosen, small perturbations are added to the system during the forecast to produce a set of forecasts which gives a probabilistic estimate of the future state and which is more representative of the chaotic variability. *Flowerdew et al. [2010]* extend this to ocean forecasts of coastal sea level by using an atmospheric ensemble system, MOGREPS, to drive a tide-surge model, producing an ensemble of coastal sea level forecasts.

[4] In this paper, we will focus on understanding the *sensitivity* of the coastal sea level system to uncertainties in its forcing, specifically to small perturbations in the wind stress forcing. This is different to the ensemble forecasting approach, since we want to determine the prior pattern of wind stress sensitivity which gives a certain perturbation to sea level at a particular time, rather than considering what is the envelope of possible sea level perturbations that we may obtain by perturbing the model earlier during its forecast. The sensitivity question is relevant to understanding the physical processes, time, and space scales which lead to the event. In this sense, it also guides our observational strategy; improved observational sampling at particular places and times of high sensitivity should lead to better prediction of sea level (through model validation or data assimilation) or better understanding of the sea level processes.

[5] Another technique, adjoint modeling, which was first developed in the context of atmospheric dynamics [e.g., *Lewis and Derber, 1985*], has begun to be applied to ocean sensitivity questions. This has been done predominantly for open ocean, large-scale problems such as the sensitivity of the North Atlantic poleward heat transport to surface heat flux forcing [*Marotzke et al., 1999*], but there are also some examples of its application to the sensitivity of coastal sea surface temperature, eddy kinetic energy, and baroclinic instability to surface forcing [e.g., *Moore et al., 2009; Zhang et al., 2009*].

[6] In a different context, within a data assimilative model, *Kurapov et al. [2009]* use the adjoint model in a 2-D cross-shore slice to determine “representers.” These may be interpreted as the prior model error covariance between the observed quantity and all the elements of the multivariate ocean state vector. Crucially, the “representer” is defined as being related to both the sensitivity and to an assumed structure for the error covariance.

[7] There are other applications of adjoint modeling in the context of data assimilation, such as *Peng et al. [2013]*, but these are also quite distinct as they do not assess the sensitivity.

[8] As far as we are aware, our study is the first application of adjoint modeling to the estimation of the sensitivity of shelf-wide coastal sea level dynamics. Unlike many similar studies, we directly examine the adjoint sensitivity and our assumptions are limited to the diagnosis and interpretation of wind stress error, rather than to any prior assumptions about error covariance.

[9] In section 2, we outline the concept of adjoint modeling and how it may be used to determine the sensitivity of aspects of the tide-surge model. The construction of the tide-surge forward model is detailed in section 3, and the tide-surge adjoint model is covered in section 5. Section 4 discusses a range of estimates of wind stress error, used alongside the adjoint sensitivity to estimate the impact on

sea level in section 6. Finally, we discuss the results in section 7.

2. Adjoint Modeling

2.1. Aim

[10] The concept of adjoint modeling is quite straightforward although the mathematics can often appear complex. The central aim is to determine the sensitivity of the system, or what prior perturbations around the background state lead to a given perturbation in a quantity of interest at a later time. Rather than recite a formal, mathematical description, we refer the reader to *Marotzke et al. [1999]*, which is also especially relevant for the MITgcm application. Instead, we describe heuristically the main elements in more familiar terms. These elements involve a dynamical systems view, where the model background state and perturbations to it are uniquely described and evolve in time under nonlinear, deterministic governing equations. The evolution may be approximated by linearization of the governing equations, analogous to a Taylor series expansion.

2.2. State Vector and Tangent Linear Model

[11] The dynamical systems description depends on the model physics being represented uniquely by a *state vector*, whose components represent measurements such as velocity or sea level at each location at the given time. The (generally nonlinear) model governing equations determine the temporal evolution of the state vector to form a trajectory in state space (equivalently phase space). The dynamical system may be simplified by linearization, i.e., using the Jacobian matrix of the governing equations at each time step to update the state vector, producing the tangent linear model. The nonlinear trajectory is curved and the tangent linear model trajectory is equivalent to taking the tangent line to this curved trajectory at each time step. Therefore, the ability of the tangent linear model to evolve the system from an earlier state to a later state depends on how curved the nonlinear trajectory is and the size of the time step. The evolution of the unperturbed, background state is a special case (with zero perturbation) of the more general situation where we consider the evolution of perturbations to the background state. In this case, the direction of the perturbations relative to the background evolution is also relevant. The standard procedure for ensuring a good approximation is to switch off highly nonlinear processes (such as the K-Profile Parameterization of vertical mixing, where relevant) and to choose a small tangent linear or adjoint model time step, usually no greater than the forward model time step.

2.3. Reducing the Complexity and Definition of the Cost Function

[12] The state vector typically contains $D = 10^5 - 10^6$ degrees of freedom (e.g., velocity and sea level at every grid point). If we were to examine the sensitivity of the state vector, $O(D)$, to perturbations of the state at earlier times, $O(D)$, we would have a massively complex problem, $O(D^2)$. Instead, we simplify by posing a question around a scalar measure of the system state, e.g., what is the sensitivity of the area-mean velocity variance, potential energy or maximum speed to the state at earlier times? By applying a




Figure 1. Tidal amplitude (shaded, 10^{-2} m) and phase (contours, degrees) from a 29 day, 26-harmonic tidal analysis beginning 0000 UTC, 14 October 2007. M2 from (a) MITgcm and (b) the CS3X operational tide-surge model. O1 from (c) MITgcm and (d) the CS3X operational tide-surge model. Sheerness is labeled “S” and its location is shown with a red dot. White areas are treated as land due to minimum depth constraints and the Mediterranean Sea effectively closed because it has no influence.

function to measure the state vector to produce a real scalar output, we reduce the size of the sensitivity problem from $O(D^2)$ to $O(D)$. This scalar measure is often referred to as a cost function, J . In our particular application, the cost function is merely a subsampling of the sea level at the model point representing Sheerness at a given time, since we are interested in what drives sea level perturbations about the background evolution at this important location. Sheerness is chosen because it is the location (Figure 1, red dot) in the UK tide gauge network which is closest to the Thames Flood Barrier that protects London from storm tides; forecast sea level at Sheerness is the critical factor that determines the closure of the barrier.

2.4. Relationship Between the Cost Function, Tangent Linear, and Adjoint Models

[13] The next steps are justified in mathematical detail in Marotzke *et al.* [1999], but involve (i) the definition of a scalar cost function, J , in terms of the final state vector, \mathbf{X}_N : $J = f(\mathbf{X}_N)$; (ii) the application of the nonlinear model as a compound sequence of operations, Ψ_n : $\mathbf{X}_N = [\Psi_N[\Psi_{N-1} \dots [\Psi_1(\mathbf{X}_0) \dots]]]$ linking the initial state, \mathbf{X}_0 , to the final state; (iii) the construction of the derivative of the cost function with respect to the initial state by using the chain rule of calculus applied to this sequence of operations, including the tangent linear model approximation,

$(\Psi'_n)_{ij} = \frac{\partial(\Psi_n)_i}{\partial(\mathbf{x}_n)_j}$, for the i, j th element of the Jacobian matrix at time step $n \in [1, N]$; and (iv) the use of a matrix transpose or adjoint operation on the expression for this derivative (the *sensitivity* of the cost function, $\frac{\partial J}{\partial \mathbf{x}_0}$) to produce a counterpart, $\left(\frac{\partial J}{\partial \mathbf{x}_0}\right)^T$.

[14] The algorithmic evaluation of this counterpart is essentially the adjoint model, as described by *Marotzke et al.* [1999]. (In the context of an optimization problem, the cost function sensitivity is related to the product of the adjoint and Lagrange multipliers (see *Wunsch and Heimbach* [2007] for more details)—used to find the minimum cost.) By taking the transpose of the sequence of operations defined in $\frac{\partial J}{\partial \mathbf{x}_0}$, their order is reversed. This means that the first operation to be evaluated, which contains the evaluation of the cost function on the final state, now immediately reduces the problem from $O(D^2)$ to $O(D)$, to be followed by further computational steps of similar expense. In contrast, the equivalent expression without the transpose is evaluated as a sequence of $O(D^2)$ matrix multiplications followed by the evaluation of the cost function on the final state, so is far more computationally expensive.

[15] It is noteworthy that the name “adjoint” comes from earlier, more abstract, descriptions of this process. However, in practice we are dealing with a special case of a real matrix in a Euclidean space, for which the adjoint is simply the matrix transpose, as highlighted by *Marotzke et al.* [1999].

[16] The description of the model state vector may be augmented to include information about initial and boundary conditions and physical parameterizations and all of the above may be repeated for this augmented “control vector” at time n , \mathbf{X}_n . The adjoint model then produces the sensitivity of the cost function, J , taken here to be the sea level at Sheerness at a particular time of interest, to the control vector at earlier times, i.e., $\frac{\partial J}{\partial \mathbf{x}_n}$. This derivative is a vector and includes the spatial pattern of sensitivity at each prior time. If we are interested in the sensitivity of J to wind stress, we augment the model state with the wind stress boundary condition in the model formulation and then examine the relevant part of $\frac{\partial J}{\partial \mathbf{x}_n}$ after adjoint model integration.

[17] We describe the application of the adjoint framework to the algorithm of the MITgcm model in section 5. The algorithmic framework of MITgcm was designed so that its structure allows algorithmic differentiation, hence its choice for this study. The following section describes the setup and testing of the MITgcm forward model simulation.

3. Construction of the MITgcm Forward Model

[18] The Massachusetts Institute of Technology General Circulation Model (MITgcm; see <http://mitgcm.org>) [*Marshall et al.*, 1997; *Marotzke et al.*, 1999] is designed for the simulation of a wide variety of fluid flows in the ocean and atmosphere. Its numerical code has been specially constructed to allow the application of automatic differentiation tools to calculate the tangent linear and adjoint model counterparts to the forward model. The model code is open source and freely available, has a large and broad user base. For this application, the model domain covers

the shelf seas surrounding the UK and matches the grid and resolution of the UK operational tide-surge Model, CS3X [*Flather*, 1976; *Horsburgh et al.*, 2008], at $\frac{1}{6}^\circ$ longitude and $\frac{1}{9}^\circ$ latitude, 20°W to 13°E , 40°N to 63°N . MITgcm is forced at the lateral boundaries with the same tidal boundary conditions used by CS3X, consisting of 26 tidal harmonics from an analysis of a broader-area model of the Northeast Atlantic [*Flather*, 1981] and with operational meteorological forecast wind from the Met Office. For simplicity, the inverse barometer effect is excluded and atmospheric surface pressure is assumed to be an arbitrary constant everywhere at all times. This ensures that errors in forecast pressure are not considered in this study.

[19] As for the operational tide-surge model, MITgcm is set up with nonlinear shallow water governing equations for momentum and continuity:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla_S \mathbf{u} + \left(f + \frac{u}{R} \tan \phi\right) \mathbf{k} \times \mathbf{u} \\ = -g \nabla_S \eta + \frac{\boldsymbol{\tau}}{\rho_0 H} + A_h \nabla_S^2 \mathbf{u} - c \frac{\mathbf{u}|\mathbf{u}|}{H} \end{aligned} \quad (1)$$

$$\frac{\partial \eta}{\partial t} + \nabla_S \cdot (\mathbf{u}H) = 0 \quad (2)$$

where $\mathbf{u} = (u, v)$, η are the horizontal velocity and sea level; $f = f_0 + \beta R \sin(\phi - \phi_0)$ is the Coriolis parameter where $f_0 = 10^{-4} \text{s}^{-1}$ and $\beta = 10^{-11} \text{m}^{-1} \text{s}^{-1}$, R is the radius of the Earth, (χ, ϕ) are longitude and latitude in spherical coordinates, and ϕ_0 is the latitude of the southern edge of the domain; \mathbf{k} is the unit vertical vector, $g = 9.81 \text{ms}^{-2}$ is gravitational acceleration, $\boldsymbol{\tau}$ is the wind stress, $\rho_0 = 10^3 \text{kg m}^{-3}$ is the background density, $H = h_0 + \eta$ is the total water depth, h_0 is the undisturbed water depth, $A_h = 4 \times 10^{-2} \text{m}^2 \text{s}^{-1}$ is the viscosity, and $c = 3 \times 10^{-3}$ is the bottom drag coefficient. For a vector, $\mathbf{V} = (V_\chi, V_\phi)$, its horizontal divergence in spherical coordinates is $\nabla_S \cdot \mathbf{V} = \frac{1}{R \cos \phi} \left(\frac{\partial V_\chi}{\partial \chi} + \frac{\partial}{\partial \phi} (V_\phi \cos \phi) \right)$. For a scalar, C , its gradient is $\nabla_S C = \left(\frac{1}{R \cos \phi} \frac{\partial C}{\partial \chi}, \frac{1}{R} \frac{\partial C}{\partial \phi} \right)$.

[20] The MITgcm bathymetry is identical to that of CS3X and is bilinearly interpolated from ETOPO5 (Data Announcement 88-MGG-02, Digital relief of the Surface of the Earth. NOAA, National Geophysical Data Center, Boulder, Colorado, 1988) to the model grid.

[21] A significant storm surge event in the North Sea on 9 November 2007 was chosen as a test case. The forward model was spun-up from rest and geostrophic balance, on 0000 coordinated universal time (UTC), 7 October 2007, as a tide-only run driven by tidal boundary forcing as described above. To validate the tide in MITgcm, a 29 day, 26-harmonic tidal analysis beginning 0000 UTC, 14 October 2007 was performed and compared with that from CS3X. Figure 1 shows the comparison between MITgcm and CS3X for key semidiurnal (M2) and diurnal (O1) tidal components. Generally, the amplitude and phase of the tide is comparable over the NW European continental shelf, despite the operational model CS3X having been extensively tuned to optimize its tidal simulation, unlike MITgcm. In CS3X, the quadratic bottom drag coefficient, tuned for the particular setup of the model, is 2.5×10^{-3} , so is 17%




Figure 2. Sheerness sea level time series from 1200 UTC, 8 November 2007 to 1300 UTC, 9 November 2007 showing tide (from MITgcm in solid red; from tide gauge observations with offset applied in dashed red), total sea level (from MITgcm in solid black; from tide gauge observations with offset applied in dashed black), and nontidal residual (from MITgcm in solid blue; from tide gauge observations in dashed blue). The residual peaks at 2.40 m at 0630 UTC for MITgcm and at 2.40 m at 0715 UTC for the observations. The observed tide has had an offset of 3.09 m subtracted, relative to its Admiralty Chart Datum, to ensure that it has the same mean over 1–11 November as the modeled tide to aid comparison. Note that the modeled total sea level peaks at 1200 UTC on the 9 November.

weaker than in MITgcm. For MITgcm, M2 (Figure 1a) is slightly weaker in amplitude than for CS3X (typically by order 30 cm on the North Sea coast of the UK; Figure 1b) and the M2 tidal phase is advanced by approximately 30° , or one twelfth of the period. This phase difference of order 1 h could be relevant to our later results only if tide-surge interaction was shown to be relevant. The position of M2 tidal amphidromes are similar in both models. For MITgcm, O1 (Figure 1c) is a few centimeters stronger in amplitude than for CS3X (Figure 1d) along the North Sea coast. The O1 tidal phases are similar, with the phase in MITgcm being advanced by up to 30° . Examining the quantitative differences at the model point corresponding to Sheerness, for M2, MITgcm has amplitude of 1.37 m and phase of 30° , and CS3X has amplitude of 1.45 m and phase of 4° ; for O1, MITgcm has amplitude of 0.12 m and phase of 223° , and CS3X has amplitude of 0.12 m and phase of 197° .

[22] The models differ in their treatment of the near coast: CS3X has an active wetting and drying scheme allowing a more realistic change to the coastline dependent on local water depth [Flather and Heaps, 1975], but MITgcm does not have such a scheme and instead treats points with undisturbed depth less than 5 m as dry land. The implementation of a wetting and drying scheme in MITgcm and an associated tuning of bottom drag coefficient is beyond the scope of this preliminary study on the adjoint technique, but may be relevant to subsequent investigations into the details of adjoint modeling for coastal sea level.

[23] In addition to a tide-only integration of MITgcm, an integration with both tidal forcing and surface meteorological forcing (wind taken from the Met Office deterministic

forecast used to drive the operational model CS3X) was calculated. Both integrations formed continuations to the spin-up, and covered the period 0000 UTC, 1 November to 2300 UTC, 11 November 2007 which included the storm surge event that peaked at Sheerness on 9 November. The sea level time series at Sheerness is shown in Figure 2. The total sea level is composed of a tidal component and a nontidal residual, usually associated with a storm surge (but which also contains observational or modeling errors, etc.). The tidal component from the tide-only run may be subtracted from the total sea level in the tide-plus-meteorology run to obtain the nontidal residual, hereafter called the surge. Figure 2 compares the tide and surge in MITgcm with observations from the Sheerness tide gauge. The tidal component of sea level from the tide gauge is obtained by harmonic analysis. To compare the model tide with tide gauge observations at Sheerness, an offset of 3.09 m has been subtracted from the observations so that they have the same mean sea level as the model over the period 1–11 November 2007. This is the usual procedure for comparing tide measurements relative to two different datums. Here, MITgcm underestimates the tidal amplitude by 50–60%, although its phase is accurately reproduced.

[24] For the surge, MITgcm compares very closely to the observations (since this is a residual, no offset need be applied). The peak surge is 2.40 m at 0630 UTC for MITgcm and 2.40 m at 0715 UTC for the observations. The observed surge decays more slowly in MITgcm than observed. That the surge is modeled so skillfully, despite the local weakness of the modeled tide, suggests that tide-surge interaction is insignificant at this location for this event. We would expect that the local modeled tide and the decay rate of the surge at Sheerness would be improved

by parameter tuning and wetting-and-drying, as is often the case in operational modeling, but here we assume that our examination of the dynamics of the system is not fundamentally affected. As we have stated, our focus is on the sensitivity of total sea level at Sheerness at the time of its maximum to the prior wind stress. Figure 2 shows that this maximum of total sea level occurs at 1200 UTC on 9 November 2007. Before 1000 UTC on 9 November, total sea level in MITgcm is generally of smaller amplitude than the tide gauge observations, reflecting the weaker modeled tide and the close agreement between observed and modeled surge. From 1000 UTC to 1200 UTC on 9 November, the modeled total sea level is within a few centimeters of observations, reflecting instead the error due to a slower rate of decay of the surge (order 1 m at 1200 UTC) in MITgcm canceling the error in modeled tide (order -1 m at 1200 UTC). Our adjoint cost function, J , is simply the subsampling of the model state vector to extract the total sea level at the grid cell corresponding to Sheerness at this time of maximum. Perturbations to J , δJ , are then caused by perturbations to the model state at time step $n \in [1, N]$, $\delta \mathbf{X}_n$, combining with the sensitivity, $\frac{\partial J}{\partial \mathbf{X}_n}$ and accumulated over the time steps of the adjoint model:

$$\delta J_N = \sum_{n=1}^N \delta \mathbf{X}_n \cdot \frac{\partial J}{\partial \mathbf{X}_n} \quad (3)$$

[25] Since our focus is on the sensitivity of J to wind stress, it is only this component of the control vector perturbation, $\delta \mathbf{X}_n$, and the sensitivity $\frac{\partial J}{\partial \mathbf{X}_n}$ which we are interested in. Therefore, we need to establish good estimates of typical wind stress perturbations associated with forecast error. If the wind stress error vector field locally aligns with the adjoint sensitivity vector field, then it may cause a perturbation to the dynamics of the system at that particular place and time *which subsequently contributes to a perturbation to sea level at Sheerness at 1200 UTC on 9 November*. In the following section, we discuss measures of wind stress error.

4. Estimating Wind Stress Error

[26] The simplest estimate of wind stress error associated with a forecast is derived from the Met Office deterministic forecast minus the hindcast. The hindcast (assumed to be the truth) is from a 4DVar assimilative run which is optimized to be closest to atmospheric observations over the hindcast period, here 0600 UTC to 1200 UTC, 9 November 2007. The hindcast integration is performed at 0000 UTC, 0600 UTC, 1200 UTC, and 1800 UTC. The forecast used was that given at 0600 UTC for the following 6 h. We should note that deterministic forecast minus hindcast probably underestimates the magnitude of the error at short lead times, because of the autocorrelation between the error of a 6 h forecast and the analysis constructed from it.

[27] As discussed in the section 1, ensemble forecasting is also commonly used to give insight into the chaotic dynamics, so we additionally consider the Met Office Global and Regional Ensemble Prediction System (MOGREPS) [Bowler *et al.*, 2008]. MOGREPS is used by *Flowerdew et al.* [2010] in a storm surge predictability context,

but here we use it to gain insight into the issue of adjoint sensitivity. There are $K=24$ members in the ensemble, including the control. The wind stress error for each member minus ensemble mean may be used in an ensemble of equation (3) to estimate a distribution of expected sea level perturbations. We examine this later in the paper. We also consider the vector standard deviation of the MOGREPS ensemble forecast wind stress,

$$\sigma_{\tau} = \left(\frac{1}{K-1} \sum_{k=1}^K \left((\tau_{x_k} - \bar{\tau}_x)^2 + (\tau_{y_k} - \bar{\tau}_y)^2 \right) \right)^{\frac{1}{2}} \quad (4)$$

where $\tau = (\tau_x, \tau_y)$ is the wind stress, the subscript k refers to ensemble member and the overbar refers to the ensemble mean, $\frac{1}{K} \sum_{k=1}^K (\cdot)_k$.

[28] The ensemble vector standard deviation predicts the typical magnitude of departures of truth from the ensemble mean wind stress: in other words, it predicts the root-mean-square magnitude of the wind stress error that would be seen if you averaged over many similar cases and the ensemble were perfect. The deterministic forecast-minus-hindcast is measuring the actual error of the deterministic forecast in this particular case (assuming the hindcast is perfect). So they are measuring the same kind of thing: in a perfect system, the forecast-minus-hindcast is one realization of the distribution whose standard deviation the ensemble predicts. Hereafter, we will refer to the forecast-minus-hindcast wind stress error simply as the wind stress error, and to the ensemble vector standard deviation of wind stress as the root-mean-square wind stress error.

[29] We will discuss these specific estimates in more detail in section 6.

5. Construction and Validation of the Adjoint Model

5.1. Construction

[30] To estimate the adjoint sensitivity, we need to consider the derivative of the cost function with respect to the model state vector. For the forward (tangent linear) model, or the reverse (adjoint) model, the evaluation of this derivative must be done accurately and efficiently. We rely on a technique called automatic (or algorithmic) differentiation [e.g., *Griewank*, 1992; *Heimbach et al.*, 2002], in particular the Transformation of Algorithms in Fortran (TAF) software [*Giering and Kaminski*, 1998]. *Heimbach et al.* [2002] describe the application of the predecessor of TAF, tangent linear and adjoint model compiler (TAMC), to the MITgcm model. TAF is widely used with MITgcm in the open ocean adjoint modeling community. Automatic differentiation analyses the operation of every line of the model code making up the algorithm and decomposes the overall governing equations or model operator into elementary functions. The derivatives of these elementary functions, the local Jacobians, can be derived following simple rules. As described in section 2, the adjoint sensitivity may then be calculated in an efficient way by evaluating the computational operations in a particular order.

[31] The practical application of TAF to MITgcm involves specification of the cost function and control vector in the MITgcm code and the choice of specific

compilation flags which send the precompiled forward model for separate automatic differentiation of the code by TAF, which returns the adjoint code for compilation to be completed locally. TAF has been extensively tested with a variety of numerical model codes, including MITgcm (see <http://mitgcm.org> for a list of publications on adjoint results).

5.2. Validation

[32] It should be emphasized that it is not generally possible to verify the adjoint sensitivity results a priori with a forward model ensemble simulation due to the large degrees of freedom inherent in this system. Again, this contrasts with the ensemble forecasting problem where the aim is to produce a range of estimates which reflect the sensitive dependence on initial (or boundary, etc.) conditions. Such initial condition perturbations are generally composed of a small number (fewer than 100) of vectors which are a sum of the basis vectors of the state space (numbering 10^5 or greater), so containing nonzero perturbations in each basis direction which may grow if unstable. For the adjoint sensitivity problem, the analogous operation would be to perturb an ensemble in each of the basis directions individually, a much larger and intractable problem.

[33] However, we may validate the adjoint results a posteriori by using an ensemble of nonlinear forward model experiments driven by perturbations based on the pattern of the adjoint sensitivity. We examined the 6 h from 0600 UTC to 1200 UTC on the 9 November and perturbed the initial conditions of background wind stress forcing by using the pattern of the adjoint sensitivity vector field from Figure 8b, applied for one model time step at 0600 UTC. This pattern was normalized to produce the same area-mean magnitude as that of the wind stress error in Figure 8c over the region bounded by $|\frac{\partial \zeta}{\partial \tau}| > 10^{-5} \text{m Pa}^{-1}$. This region was chosen to exclude the effect of large wind stress error, in places where the sensitivity is small, on the normalization. The normalization factor was found to be $9.30 \times 10^3 \text{Pa}^2 \text{m}^{-1}$. Each of these normalized patterns was multiplied by a factor from $\{-1, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1\}$ to produce the wind stress perturbation in the eight-member ensemble. This ensemble was then integrated forward in time. From these nonlinear model experiments, the response of sea level at Sheerness at 1200 UTC was found to have a very strong linear relationship to the strength of this *particular* perturbation pattern, with an inferred sensitivity of sea level at Sheerness of $6.5 \times 10^{-6} \text{m Pa}^{-1}$ (not shown). Note that although the normalization factor calculation considers a region bounded by a threshold of sensitivity magnitude, the wind stress perturbation is constructed using the sensitivity pattern over the whole domain (therefore including regions of weak sensitivity), consisting of approximately 2.4×10^4 ocean grid cells. The area mean of the magnitude of the sensitivity from Figure 8b is $1.9 \times 10^{-6} \text{m Pa}^{-1}$, and this gives a check that the two results are of the same order of magnitude. One might imagine that applying the *particular* perturbation described above to an infinite set of sensitivity fields described only as having area-mean magnitude of $1.9 \times 10^{-6} \text{m Pa}^{-1}$ would generally lead to a lower correlation in the scalar product, hence a lower response in sea level. This is why we might expect our inferred sensitivity, as

described above, to be larger than the rough estimate of $1.9 \times 10^{-6} \text{m Pa}^{-1}$. It is therefore always worth remembering that the adjoint sensitivity of sea level to wind stress is best considered as a full vector field within a scalar product and that care has to be taken when using magnitude information alone for estimation.

[34] The linear relationship found between the response of these nonlinear forward model experiments and the *particular* perturbations demonstrates that the linear assumptions that have both been (a) embedded in the construction and solution of the adjoint sensitivity pattern and (b) tested by perturbing the nonlinear forward model with this derived pattern, are valid over the 6 h in question.

6. Evolution of the System Preceding the Extreme Sea Level Event

6.1. Spatial Comparison of the Forward Model Total Sea Level With Adjoint Sensitivity and Wind Stress Error

[35] We will now describe the main results of this study, a sequence of total sea level and adjoint sensitivity from MITgcm, with wind stress error information from Met Office simulations over the 6 h preceding the peak total sea level at Sheerness (shown in Figures 3–8 in reverse-time order.) By examining the spatial structure of these fields at each time, we can examine how the fields and their scalar product evolve both locally and remotely with respect to Sheerness.

6.1.1. At 1100 UTC, 1 h Before Maximum Total Sea Level at Sheerness—Maximum Sensitivity is Localized at Sheerness; Wind Stress Error Contains Larger-Scale Structure

[36] Figure 3 corresponds to 1100 UTC on 9 November 2007. From the time series in Figure 2 (black line), this is 1 h before maximum total sea level at Sheerness. The spatial distribution of total sea level is shown in Figure 3a. A large surge is evident in the southern North Sea. Total sea level, with respect to undisturbed model sea level, exceeds 3.1 m. The adjoint sensitivity of sea level at Sheerness at 1200 UTC with respect to wind stress at 1100 UTC is very localized and circular, with radius of order 100 km (Figure 3b), with its magnitude shown by shading and direction by vectors of uniform length, plotted only above a small magnitude threshold. The direction of the sensitivity vector field is convergent on Sheerness.

[37] The wind stress error from the Met Office deterministic forecast-minus-hindcast is shown in Figure 3c. The wind stress error vector field is generally directed southeasterly and parallel to the coast in the North Sea, meaning that the forecast has overestimated the wind stress in this sense relative to the hindcast. The spatial distribution of the wind stress error is different from that of adjoint sensitivity, being more elongated, reaching a scales of order 1000 km by 100 km. Its maximum magnitude does not coincide with the region of maximum sensitivity. It must be remembered that this is the error relative to the “true” wind stress, and over much of the domain the southeasterly error corresponds to the forecast having weaker northwesterly wind stress than the hindcast (which is typically around 1 Pa and northwesterly). The error reaches 0.7 Pa just offshore of the UK North Sea coast and up to 0.9 Pa near the Danish coast.




Figure 3. Diagnostics 1 h before maximum total sea level at Sheerness (circled), 1100 UTC, 9 November 2007: (a) Total sea level (m) from MITgcm. (b) The MITgcm adjoint sensitivity ($10^{-5} \text{ m Pa}^{-1}$) of sea level at Sheerness at 1200 UTC to wind stress at the stated time. Magnitude is shaded on a partial log scale and unit vectors show direction. (c) Wind stress error (Pa), estimated from Met Office deterministic forecast minus hindcast. (d) A measure of the root-mean-square wind stress error, the ensemble vector standard deviation (Pa), from the MOGREPS ensemble.

The vector standard deviation of the wind stress from the MOGREPS ensemble forecast with base time 1800 UTC on 8 November, shown in Figure 3d and defined in equation (4), indicates relatively large root-mean-square wind stress error at this time, reaching 0.3 Pa along the North Sea coast of the UK and up to 0.4 Pa near the Danish coast. Therefore we should be cautious in our interpretation of the wind stress error estimate given in Figure 3c since it may not be typical. Along the North Sea coast of the UK, the root-mean-square wind stress error is about 30% of the wind stress error. The root-mean-square wind stress error pattern has broader spatial scale than that of wind stress error at this time. However, given that it is difficult to display the wind stress error for all the members of the MOGREPS ensemble, we rely on Figures 3c and 3d to show

conveniently a single, yet possibly atypical, estimate and the spread of alternative realizations, respectively, in order to make broad comparison between wind stress error and adjoint sensitivity at each time. Later, we explore the use of the wind stress error from the full MOGREPS ensemble in a condensed diagnostic of inferred sea level error.

6.1.2. At 1000 UTC, 2 h Before Maximum Total Sea Level at Sheerness—Sensitivity Pattern Has Spread Outward From Sheerness; Wind Stress Error Remains Similar

[38] At 1000 UTC, there is still a large surge in the southern North Sea and the total sea level pattern (Figure 4a) is very similar to that in Figure 3a. At Sheerness, total water level is changing relatively slowly (Figure 2 (black line)) since it is near high tide. The adjoint sensitivity




Figure 4. As for Figure 3 but for 2 h before maximum total sea level at Sheerness (circled), 1000 UTC, 9 November 2007.

pattern (Figure 4b) is larger in extent than at 1100 UTC, being approximately circular and of order 200 km in radius, with direction still convergent on Sheerness. Wind stress error (Figure 4c) is very similar to that at 1100 UTC, with only subtle changes to magnitude and direction, and the pattern on the North Sea coast of the UK does not extend quite as far south, remaining just outside the expanded region of high sensitivity. There is a broad reduction in root-mean-square wind stress error (Figure 4d), especially near the Danish coast where the standard deviation is of order 0.3 Pa. The distribution of root-mean-square wind stress error does not match that of wind stress error.

6.1.3. At 0900 UTC, 3 h Before Maximum Total Sea Level at Sheerness—Sensitivity has Spread Outward and has Translated Northward; Wind Stress Error has Locally Weakened and Remotely Strengthened

[39] The total sea level is now changing rapidly at Sheerness (Figure 2 (black line)) and the pattern (Figure 5a) is significantly different from that at 1000 UTC. There is a

maximum in total sea level, apparent as a coastally trapped Kelvin wave, approximately 200 km north of Sheerness along the UK coast. There are also extrema in total sea level at other coastal locations which have generally evolved in the sense of a coastally trapped Kelvin wave propagating backward in time with respect to 1000 UTC. The adjoint sensitivity (Figure 5b) has expanded further in extent compared to 1000 UTC, but is now asymmetric about Sheerness and has extended further northward up the UK coast. Its pattern is like that at 1000 UTC, except for this translation, which is similar to that of the local maximum of total sea level. Although the sensitivity pattern is quite circular, there are additional features embedded within this simple structure. Compared to the wind stress error at 1000 UTC, there has been a slight weakening by around 0.1 Pa along the North Sea coast of the UK and a strengthening by a similar amount off the Danish coast, without any notable changes in direction (Figure 5c). There is still a mismatch in the regions of largest wind stress error




Figure 5. As for Figure 3 but for 3 h before maximum total sea level at Sheerness (circled), 0900 UTC, 9 November 2007.

and the regions of largest sensitivity. The root-mean-square wind stress error (Figure 5d) has increased by order 0.1 Pa off the Danish coast and has reduced off the Norwegian coast, compared to 1000 UTC.

6.1.4. At 0800 UTC, 4 h Before Maximum Total Sea Level at Sheerness—Sensitivity has Evolved into a Circular Ring, Translated Further Northward and Weakened; Wind Stress Error has Weakened Locally

[40] The total sea level (Figure 6a) shows that the Kelvin wave is yet further north up the UK coast relative to 0900 UTC, the sensitivity (Figure 6b) has evolved as a circular ring with larger radius than at 0900 UTC and which also has a center close to the maximum total sea level, about 300 km north of Sheerness. Compared to earlier times, the sensitivity is weaker and the majority of the distribution has magnitude smaller than $10^{-3} \text{ m Pa}^{-1}$. The direction of the sensitivity converges toward this center. Again, the regions of large wind stress error (Figure 6c) do not match those of large sensitivity. Relative to 0900 UTC, the wind

stress error along the North Sea coast of the UK has further weakened, by around 0.1 Pa, and the maximum off the Danish coast has moved northward. In the region of high sensitivity, there have been some small changes to the direction of wind stress error, but there has also been an increase in the root-mean-square wind stress error in the southern North Sea by approximately 0.2 Pa—this is comparable to the magnitude of the wind stress error in that region.

6.1.5. At 0700 UTC, 5 h Before Maximum Total Sea Level at Sheerness—Sensitivity Continues to Spread, Translate Northward and Weakens, Becoming More Asymmetric; Wind Stress Error Remains Similar, With Some Remote Overlap With Sensitivity

[41] There is more evidence that the sensitivity (Figure 7b) continues to expand as a circular ring, to weaken and to translate with the maximum in total sea level (Figure 7a), which is now approximately 400 km north of Sheerness. The direction of the sensitivity still broadly converges




Figure 6. As for Figure 3 but for 4 h before maximum total sea level at Sheerness (circled), 0800 UTC, 9 November 2007.

toward the center of this ring but is asymmetric. The wind stress error (Figure 7c) is very similar to that at 0800 UTC but the sensitivity pattern has now expanded sufficiently to begin to overlap with the large wind stress error off the Danish coast. The root-mean-square wind stress error at this time (Figure 7d) is quite similar to that at 0800 UTC, but is larger near eastern Scotland, approaching 0.3 Pa, or 50% of the wind stress error. It is therefore possible that the pattern of wind stress error is quite atypical.

6.1.6. At 0600 UTC, 6 h Before Maximum Total Sea Level at Sheerness—Further Spreading and Weakening of Sensitivity Pattern, With More Complex Structure; Local Wind Stress Error is Weaker

[42] The story is similar at 0600 UTC, with a continued reverse-time propagation of the coastally trapped Kelvin wave northward up the coast (Figure 8a), a pattern of sensitivity (Figure 8b) which is order 10 times weaker than at 0700 UTC, with a greater radius and that has translated northward matching the Kelvin wave. The direction of sen-

sitivity within the ring-like region of maximum magnitude is yet more complex than at later times. The wind stress error (Figure 8c) is again weaker in the region of high sensitivity compared to later times and its maxima do not align with maxima of sensitivity. Due to the weaker wind stress error and the similarity of the root-mean-square wind stress error (Figure 8d), there are regions in the southern North Sea, coincident with the highest sensitivity, where the root-mean-square wind stress error is of comparable size to the wind stress error.

6.2. The Pattern of Adjoint Sensitivity

[43] The time at which a shallow water wave emanating from Sheerness reaches a certain distance, x , may be estimated by $\int \frac{dx}{\sqrt{gH(x)}}$, where the indefinite integral begins at Sheerness and \sqrt{gH} is the wave speed. Based on the model topography, $H(x)$, along the line of latitude corresponding to Sheerness, this diagnostic (not shown) suggests that the




Figure 7. As for Figure 3 but for 5 h before maximum total sea level at Sheerness (circled), 0700 UTC, 9 November 2007.

sensitivity pattern is similar to such a disturbance propagating backward in time with speed of order \sqrt{gH} . In the 2 h preceding 1200 UTC, the pattern of adjoint sensitivity approximates this backward-propagating shallow water wave distribution—Figures 3a and 4a. At times between 6 and 2 h before 1200 UTC, the adjoint sensitivity contains a similar reverse-time wavefront signature, but this signature also translates along the coast with the maximum of total sea level, a coastally trapped Kelvin wave, on the east coast of the UK which is eventually associated with the peak in sea level at 1200 UTC.

[44] The propagating coastally trapped wave seen in sea level is a balance between the horizontal pressure gradient provided by the coastal boundary and the Coriolis force, relevant on time scales of order the inertial period, $f^{-1} = 10^4$ s or 3 h and longer. In a nonrotating system we might hypothesize that the adjoint sensitivity of the reverse-time wavefront signature may continue spreading backward in time, but becoming more complex as topographic reflections are introduced. In our rotating system, we add the

reverse-time propagation of the coastally trapped Kelvin wave northward up the east coast to our reverse-time ripple of the shallow water wave caused by a point perturbation at Sheerness. The combination of (a) spreading point perturbation “ripple” at speed \sqrt{gH} , (b) topographic reflection of this ripple, and (c) advection of the disturbance by the coastally trapped wave, all in the reverse-time direction, seem to fit our diagnosed distribution of the magnitude of adjoint sensitivity. The direction of the sensitivity vector field is simply a convergence on Sheerness at times of up to 2 h prior to peak total sea level. At earlier times, the direction becomes more complex.

[45] Translating this into more familiar language, the total sea level at Sheerness at 1200 UTC on 9 November is most sensitive to earlier wind stress perturbations which follow a certain pattern that evolves in time and space in the preceding hours. This pattern is composed of two components: a circular, wave-like structure that converges with time (similar to that seen when throwing a pebble into the shallow water of a pond, but in reversed time), and a




Figure 8. As for Figure 3 but for 6 h before maximum total sea level at Sheerness (circled), 0600 UTC, 9 November 2007.

translation of this structure in a sense that follows the maximum elevation of the coastally trapped tidal Kelvin wave initially north of Sheerness and responsible for high tide at 1200 UTC. The pattern generally becomes more localized and symmetric approaching the time and location at which the cost function is evaluated, and the sensitivity grows rapidly. It is not obvious which of these effects will dominate in the integral over the spatial domain. Nor is it obvious how the wind stress error might combine with the adjoint sensitivity to produce an overall measure of the expected error in forecast total sea level. As we have seen in Figures 3a–3c to 8a–8c, there is no obvious correlation between the large-scale patterns of wind stress error and adjoint sensitivity or total sea level at each time.

6.3. The Instantaneous, Local Combination of Wind Stress Error and Adjoint Sensitivity to the Total Sea Level Perturbation

[46] The scalar product of the wind stress error with the adjoint sensitivity at each grid cell shows how effectively

the wind stress error may contribute over the local space-time domain to perturbations in the cost function (i.e., whether the wind stress error may push in the right direction to achieve a particular sea level response at the later time). When accumulated over time, as described in equation (3), this gives the expected perturbation to the cost function, the total sea level at Sheerness at 1200 UTC. Obviously, it is difficult to estimate this scalar product visually from the vector fields in, e.g., Figures 3b and 3c, but we may diagnose the pattern of the components of this scalar product made by the “local” scalar products of wind stress error and sensitivity vectors at each grid point (these “local” scalar products are then summed to give the scalar product defined in equation (3)).

[47] The pattern of these “local” scalar products (Figure 9) is very localized to Sheerness at 1100 UTC (Figure 9a), consisting of a dipole. The wind stress error is directed northeastward near Sheerness, but the sensitivity converges on Sheerness, hence forming the scalar product ingredients for the dipole. The rapid decay of the adjoint sensitivity




Figure 9. Time sequence showing the local contributions to the scalar product $\delta\mathbf{X}_n \cdot \frac{\partial J}{\partial \mathbf{X}_n}$ (10^{-6} m) at times of (a) 1 h, (b) 2 h, (c) 4 h, and (d) 6 h before maximum sea level at Sheerness, 1200 UTC, 9 November 2007. See equation (3). The value at each location corresponds to the scalar product between two vectors at that point: the wind stress error (Figure 8c) and the adjoint sensitivity (Figure 8b), which are each, respectively, subsampled from the vectors $\delta\mathbf{X}_n$ and $\frac{\partial J}{\partial \mathbf{X}_n}$ which cover the whole domain. For example, $\delta\mathbf{X}_n = [\delta\tau_1; \delta\tau_2; \delta\tau_3; \dots; \delta\tau_k; \dots; \delta\tau_A]$, where the $\delta\tau_k$ refer to wind stress error vector at a spatial point, k . Note the partial logarithmic scale. The sum of these local scalar products over the whole spatial domain equals $\delta\mathbf{X}_n \cdot \frac{\partial J}{\partial \mathbf{X}_n}$. When accumulated in time, this makes up the red line in Figure 10. Sheerness is circled.

with distance from Sheerness is the limiting element of the extent of this “local” scalar product pattern, since the distribution of wind stress error magnitude does not vary as extremely and there are no features in the “local” scalar product which correspond to large wind stress error off the Danish coast. At 1000 UTC, the “local” scalar product has extended in scale to order 200 km radius and is more negative on average—the sum of this pattern over the domain equates to the significant decrease in the red line in Figure 10 at this time. At 0800 UTC, 4 h before peak total sea level

at Sheerness, the “local” scalar product is of similar scale to the sensitivity, having expanded to order 300 km radius and translated north of Sheerness. It contains both positive and negative values of up to 10^{-3} m. At 6 h before 1200 UTC, the “local” scalar product is more complex in structure and extends further spatially, reaching 400 km from the coast. However, when integrated to form the full scalar product of wind stress error and adjoint sensitivity, the 3 h before 1200 UTC are most important and there is cancellation in these “local” scalar product patterns at earlier times.

6.4. Temporal Contribution of Integrated Wind Stress Error and Adjoint Sensitivity to the Total Sea Level Perturbation

[48] We finally evaluate the temporal contribution to the perturbation of total sea level at Sheerness at 1200 UTC, $\delta\eta_N = \delta J_N$, inferred by our combined prior estimates of wind stress error and adjoint sensitivity, as given by equation (3) and shown in Figure 10. This is the accumulation of wind stress error on the model sensitivity over the whole spatial domain from 0600 UTC to 1200 UTC on 9 November 2007, the time of maximum total sea level at Sheerness. Equivalently, for the red line in Figure 10 it is the sum of the “local” scalar product over the domain, shown in Figure 9, accumulated in time. We first describe the estimate of $\delta\eta$ (Figure 10) and then compare it to other estimates from the MITgcm forward model and the CS3X operational model.

[49] The wind stress error from the deterministic forecast-minus-hindcast infers a negative perturbation $\delta\eta$ (Figure 10, red line). This implies that, relative to the hindcast model state (assumed to be the truth), the forecast wind stress error would have likely caused an underestimate of total sea level which accumulated monotonically in time up to approximately 1100 UTC, after which it weakened, eventually reaching around 1 m below assumed truth.

[50] We may also estimate the distribution of $\delta\eta$ inferred by using wind stress error estimates from MOGREPS ensemble members minus the ensemble mean. Here we use the two nearest forecasts of MOGREPS, one with base time of 1800 UTC on 8 November (as used in Figures 3–8) and another with base time of 0600 UTC on 9 November. The wind stress diagnostics used for this case study are not available until 3 h after the base time and are output every 3 h. The adjoint model sensitivity is combined with the wind stress error by linear interpolation onto a 20 min temporal grid. The MOGREPS-based estimates are shown in Figure 10 (black lines for the earlier forecast and blue lines for the later forecast).

[51] Note that the spread of the MOGREPS-based estimates of $\delta\eta$ for the earlier forecast (which is more relevant for operational coastal defence, due to its greater lead time) is much larger than that for the later forecast (Figure 10, thick black and thick blue lines, 1.7 and 0.8 m at 1200 UTC, respectively), reflecting that the atmospheric ensemble perturbations have had longer to grow in the former case. The deterministic forecast-minus-hindcast estimate (Figure 10, red) is contained within the spread of the earlier MOGREPS-based estimate.

[52] It is natural to question how representative these adjoint-derived sea level error estimates are. The answer depends on the validation of the adjoint and the linearity assumption over the time period of interest, the size of the wind stress error estimate and comparison against typical sea level errors.

[53] For MITgcm, at 1200 UTC, 9 November, the residual sea level is overestimated with respect to tide gauge observations by 0.98 m, as shown in Figure 2. This is quite large in relation to the error in the preceding few hours, where MITgcm performs exceptionally well, with residual error typical of order 0.1 m. For this event, the operational model, CS3X, performed better and for the skew surge near 1200 UTC overpredicted by approximately 0.25 m (0.21 m




Figure 10. Time series of the inferred sea level perturbation based on wind stress error, $\delta\eta_n \equiv \delta J_n = \sum_{n'=1}^n \delta \mathbf{X}_{n'} \cdot \frac{\partial J}{\partial \mathbf{X}_{n'}}$ (meter) over the 6 h before maximum sea level at Sheerness, 0600 UTC to 1200 UTC, 9 November 2007. n is the adjoint model diagnostic time step index within this window. The sea level perturbation estimate may be derived for various estimates of wind stress error, $\delta\tau \equiv \delta \mathbf{X}$. The integrations are: deterministic forecast-hindcast (red); MOGREPS ensemble member-ensemble mean for 8 November, 1800 UTC forecast (black); 9 November, 0600 UTC MOGREPS forecast (blue). The wind stress diagnostics used for this case study are not available until 3 h into the forecast. The ensemble standard deviation of $\delta\eta$ for each ensemble estimate is shown by thicker lines of the same color. An estimate of the actual error from the deterministic forecast with respect to the Sheerness tide gauge observation at 1200 UTC is 0.21 m, i.e., the observed total sea level was smaller than the deterministic CS3X forecast at that time.

for the usual residual) with respect to observations for the two nearest deterministic-forced estimates and by 0.08 m for the hindcast-forced estimate [Horsburgh *et al.*, 2008]. All of these sea level forecast/hindcast error estimates fit within the envelope of adjoint-derived estimates shown in Figure 10. As we have discussed in the previous section, the adjoint model validation is robust over the 6 h from 0600 UTC to 1200 UTC and the response of the nonlinear forward model to wind stress perturbations matching the sensitivity pattern is highly linear in nature. This is true even for wind stress perturbations typically reaching 0.5 Pa in magnitude and may possibly hold for larger perturbations.

[54] There is therefore justification that the MITgcm adjoint sensitivity solutions, when combined with typical wind stress errors derived from either the MOGREPS ensemble or the deterministic forecast-minus-hindcast wind stress fields, give us predictions in sea level error which are of the right order of magnitude. So we might conclude that both quantitatively and qualitatively within MITgcm, this calculation seems consistent. However, the fact that MITgcm has a larger residual sea level error with respect to observations than the operational model, means that it is likely that the sensitivity of sea level at Sheerness at 1200 UTC to wind stress would be weaker for the operational model than the results presented in this paper.

[55] Since the MOGREPS diagnostics available for this case study were limited to a three-hourly temporal resolution, any finer structure evident in Figure 10 is due to the 20 min resolution of the adjoint sensitivity development. From the earlier MOGREPS-based estimate (Figure 10, black lines) the spread grows linearly over the last 6 h, but embedded within the ensemble is additional structure showing more rapid changes for certain members over the final few hours. For the later MOGREPS-based estimate (Figure 10, blue lines) the spread again grows approximately linearly over the final 3 h, and at a similar rate. However, several members do appear to provide estimates of $\delta\eta$ which grow rapidly in the final 3 h. This implies that the final 3 h is the most important period, both in terms of the system being most sensitive and the typical wind stress errors being able to force this sensitivity. There is some suggestion of a reduction in growth of $\delta\eta$ magnitude within the final 20 min before 1200 UTC. We hypothesize that this may be related to the adjoint sensitivity occupying a region smaller (~ 20 – 60 km radius) than features in the wind stress error, but realize that further experiments, with greater spatial and temporal resolution, beyond the scope of this study may be required to examine this hypothesis in detail.

[56] We have outlined how the technique of adjoint modeling may be used with estimates of wind stress error to derive estimates of total sea level error. These sea level error estimates contain information about the rate of error growth, statistical distribution and clustering, as well as highlighting the potential importance of high temporal frequency meteorological information. They span a range that matches the MITgcm forward model estimate and encloses operational forecast model estimates.

[57] The chosen adjoint sensitivity diagnostic isolates the effects and sensitivity surrounding wind stress only. There may be other processes, for example sea level pres-

sure or tidal boundary conditions, whose sensitivities could act to cancel. An adjoint modeling framework is ideal for studying potentially opposing sensitivities, and we will consider this extension for a future study. Despite these caveats, qualitatively Figure 10 implies the importance of minimizing wind stress error over the few hours preceding 1200 UTC and also that a single, deterministic weather forecast may give an atypical response in coastal sea level. This strengthens the motivation for using the MOGREPS ensemble in operational tide-surge forecasting, as outlined in Flowerdew *et al.* [2010].

7. Discussion

[58] By adapting the framework of adjoint modeling to coastal sea level on scales of tide-surge events, we have begun to examine the relevant physical scales of sensitivity for a North Sea extreme storm surge event on 9 November 2007. Adjoint modeling allows computation of the sensitivity of a scalar measure of the system to the state at prior times. To demonstrate the procedure here we focused on one aspect, the sensitivity of sea level at Sheerness at the time of peak total sea level for this event to wind stress perturbations at prior times. We found that the pattern of sensitivity may be described by a spreading point perturbation “ripple” emanating from Sheerness at speed \sqrt{gH} in the reverse-time direction at short times and, at longer times prior to 3 h, modified by topographic reflection of this ripple and reverse-time advection of the disturbance by the coastally trapped tidal Kelvin wave which most significantly contributes to total sea level at Sheerness at peak time in the unperturbed system. We also highlight the possible analogy to “Proudman Resonance” [Proudman, 1929], where a moving sea level pressure anomaly may combine with a propagating shallow water wave to resonate, with maximum resonance as the atmospheric anomaly propagation speed approaches \sqrt{gH} . There are limitations to this theory, and also it refers to sea level pressure rather than wind stress, but the concept of resonance dependent on coupled propagation of disturbances is remarkable.

[59] The sensitivity of the system to wind stress perturbations is then explored with estimates of wind stress error and uncertainty from Met Office forecasts. The inferred perturbation to total sea level at Sheerness by integration of the scalar product of these wind stress error estimates with the adjoint sensitivity gives a broad range, but a strong suggestion that the details of wind stress error matter and particularly so in the 3 h preceding the peak of total sea level associated with the 9 November 2007 event. This motivates the study of further extreme storm surge events to determine whether this behavior is typical, whether there are characteristics of wind stress error and model sensitivity that cluster at particular scales or whether there is a continuous distribution. Also, one might expect that the pattern of sensitivity examined here might follow a similar structure of a reverse-time spreading ripple, advected along the coast, were the cost function to be evaluated at other coastal locations of interest. Examination of the instantaneous scalar product of wind stress error with adjoint sensitivity in the 3 h preceding the event defines a near-coastal region of typical extent less than 300 km from Sheerness, determined primarily by the extent of the adjoint

sensitivity. Our study suggests that it is over this space-time scale that one should focus efforts on improving atmospheric forecasts of wind stress (possibly involving model improvements, reduction of observational error, improvement of observational coverage or improvement of the data assimilation scheme) in order to improve the constraints on forecast total sea level. There are shortcomings in the exact reproduction of the operational model tidal harmonics, most likely due to lack of a wetting-and-drying advection scheme, beyond the scope of this paper. The background M2 phase difference equates to up to 1 h, which should be remembered when interpreting the results, as should the suggestion from Figure 2 that tide-surge interaction is weak.

[60] It seems plausible that the extent of agreement between the temporal and spatial scales of wind stress error (affected by large-scale atmospheric geostrophic balance but modified by near-surface effects) and adjoint sensitivity (likely affected by smaller-scale oceanic geostrophic balance in this example) determines the perturbation to coastal sea level. It is therefore very unlikely that wind stress error will align with the pattern of sensitivity except for small regions or short times. However, one might speculate on the construction of an extreme worst-case scenario by assuming that the worst-case wind stress error is parallel to the adjoint sensitivity but is constrained also by (for instance) the outer quantiles of the MOGREPS ensemble. Perhaps further studies on extreme storm surge events might be able to draw upon the information contained in the adjoint model in this way. For example, the adjoint sensitivity of historic extreme surge events such as that in the North Sea in 1953 might be useful for constructing a new type of worst-case scenario estimates.

[61] **Acknowledgments.** We thank Chris Hughes, Peter Stansby, and Judith Wolf for fruitful discussions and Andy Lane for advice on aspects of harmonic tidal analysis. Thanks to Patrick Heimbach for technical advice on MITgcm and the adjoint, to Javier Zavala-Garay, two anonymous reviewers, and the editor for constructive comments on the paper. We acknowledge the use of the TAF software provided by FastOpt. The observed tide gauge data were supplied by the British Oceanographic Data Centre, hosted by the National Oceanography Centre, Liverpool, and funded by the Environment Agency and the Natural Environment Research Council (NERC). C.W. and K.J.H. were partly funded by the Engineering and Physical Science Research Council as part of the Flood Risk Management Research Consortium (FRMRC) and also by the NERC.

References

- Blumenthal, M. B. (1991), Predictability of a coupled ocean-atmosphere model, *J. Clim.*, *4*, 766–784.
- Bowler, N. E., A. Arribas, K. R. Mylne, K. B. Robertson, and S. E. Beare (2008), The MOGREPS short-range ensemble prediction system, *Q. J. R. Meteorol. Soc.*, *134*(632), 703–722.
- Buizza, R., and T. Palmer (1998), Impact of ensemble size on ensemble prediction, *Mon. Weather Rev.*, *126*(9), 2503–2518.
- Epstein, E. (1969), Stochastic dynamic prediction, *Tellus*, *21*(6), 739–759.
- Flather, R. (1976), A tidal model of the northwest European continental shelf, *Mem. Soc. R. Sci. Liege*, *10*(6), 141–164.
- Flather, R. (1981), Results from a model of the north east Atlantic relating to the Norwegian Coastal Current, in *The Norwegian Coastal Current, Proceedings of the Norwegian Coastal Current Symposium*, edited by R. Saetre and M. Mork, pp. 427–458, University of Bergen, Bergen.
- Flather, R., and N. Heaps (1975), Tidal computations for Morecambe Bay, *Geophys. J. R. Astron. Soc.*, *42*(2), 489–517.
- Flowerdew, J., K. Horsburgh, C. Wilson, and K. Mylne (2010), Development and evaluation of an ensemble forecasting system for coastal storm surges, *Q. J. R. Meteorol. Soc.*, *136*(651), 1444–1456.
- Giering, R., and T. Kaminski (1998), Recipes for adjoint code construction, *ACM Trans. Math. Software*, *24*(4), 437–474.
- Griewank, A. (1992), Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation, *Optim. Methods Software*, *1*(1), 35–54.
- Heimbach, P., C. Hill, and R. Giering (2002), Automatic generation of efficient adjoint code for a parallel Navier-Stokes solver, *Comput. Sci.*, *2330*, 1019–1028.
- Horsburgh, K., J. Williams, J. Flowerdew, and K. Mylne (2008), Aspects of operational forecast model skill during an extreme storm surge event, *J. Flood Risk Manage.*, *1*(4), 213–221.
- Horsburgh, K. J., et al. (2011), *Guide to Storm Surge Forecasting*, World Meteorol. Organ., Geneva.
- Kurapov, A. L., G. D. Egbert, J. S. Allen, and R. N. Miller (2009), Representative-based analyses in the coastal upwelling system, *Dyn. Atmos. Oceans*, *48*(1–3), 198–218.
- Lewis, J. M., and J. C. Derber (1985), The use of adjoint equations to solve a variational adjustment problem with advective constraints, *Tellus, Ser. A*, *37*, 309–322.
- Lorenz, E. (1963), Deterministic nonperiodic flow, *J. Atmos. Sci.*, *20*(2), 130–141.
- Marotzke, J., R. Giering, K. Zhang, D. Stammer, C. Hill, and T. Lee (1999), Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity, *J. Geophys. Res.*, *104*(29), 529–548.
- Marshall, J., C. Hill, L. Perelman, and A. Adcroft (1997), Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling, *J. Geophys. Res.*, *102*(C3), 5733–5752.
- Moore, A. M., H. G. Arango, E. Di Lorenzo, A. J. Miller, and B. D. Cornuelle (2009), An adjoint sensitivity analysis of the Southern California current circulation and ecosystem, *J. Phys. Oceanogr.*, *39*(3), 702–720.
- Palmer, T., R. Buizza, F. Molteni, Y. Chen, and S. Corti (1994), Singular vectors and the predictability of weather and climate, *Philos. Trans. R. Soc. London A*, *348*(1688), 459–475.
- Penduff, T., M. Juza, B. Barnier, J. Zika, W. K. Dewar, A.-M. Treguier, J.-M. Molines, and N. Audiffren (2011), Sea level expression of intrinsic and forced ocean variabilities at interannual time scales, *J. Clim.*, *24*(21), 5652–5670.
- Peng, S., Y. Li, and L. Xie (2013), Adjusting the wind stress drag coefficient in storm surge forecasting using an adjoint technique, *J. Atmos. Oceanic Technol.*, *30*(3), 590–608.
- Proudman, J. (1929), The effects on the sea of changes in atmospheric pressure, *Geophys. J. Int.*, *2*, 197–209.
- Verlaan, M., A. Zijderfeld, H. de Vries, and J. Kroos (2005), Operational storm surge forecasting in the Netherlands: Developments in the last decade, *Philos. Trans. R. Soc. A*, *363*(1831), 1441–1453.
- Wunsch, C., and P. Heimbach (2007), Practical global oceanic state estimation, *Physica D*, *230*(1–2), 197–208.
- Zhang, W. G., J. L. Wilkin, J. C. Levin, and H. G. Arango (2009), An adjoint sensitivity study of buoyancy- and wind-driven circulation on the New Jersey inner shelf, *J. Phys. Oceanogr.*, *39*(7), 1652–1668.