A Variational Assimilation System for Nearshore Wave Modeling

MARK D. ORZECH, JAYARAM VEERAMONY, AND HANS NGODOCK

Naval Research Laboratory, Stennis Space Center, Mississippi

ABSTRACT

A variational data assimilation system is developed for the stationary, homogeneous portion of the wave model Simulating Waves Nearshore (SWAN). The system is based on a numerical adjoint constructed for the discrete forward SWAN code; its performance is compared to that of an earlier system based on a discretized analytical adjoint (Walker; Veeramony et al.). This paper describes the development and validation of individual numerical adjoint subroutines, followed by the testing and evaluation of the assimilation system as a whole with an idealized twin experiment and with data from Duck, North Carolina. In the twin experiment, the present system performs on par with that of Walker. Estimates of wave spectra and spectral statistics also compare well to measured spectral data at Duck, North Carolina. The error in these estimates is partly due to the exclusion of nonlinear source and sink terms from the adjoint and partly due to different spectral processing techniques used for different types of instruments.

1. Introduction

The wave model Simulating Waves Nearshore (SWAN) (Booij et al. 1999) solves the spectral action balance equation [Eq. (1)] to produce nearshore wave forecasts and climatologies. It is widely used by the coastal modeling community and is part of a variety of coupled ocean–wave–atmosphere model systems. In forecasting wave conditions for specific locations or events, boundary conditions for the local SWAN domains are generally obtained from regional or global simulations with WAVEWATCH III (Tolman 2009) and the Wave Model (WAM; WAMDI 1988). The accuracy of nearshore wave estimates is highly dependent on the quality of these boundary inputs. Even small errors in wave energy, period, and direction at the boundary can grow significantly as SWAN propagates a wave field into shallow water. With a system that utilizes measured data to correct the boundary condition for the SWAN model, such errors can be considerably reduced.

The principal spectral action balance equation in SWAN may be expressed as

$$\frac{\partial N}{\partial t} + \vec{\nabla} \cdot (\vec{\mathbf{C}}N) = \frac{S_{\text{tot}}}{\sigma}, \tag{1}$$

where N is the spectral action density, obtained by dividing the spectral energy density by an intrinsic representative wave frequency σ. The first term of Eq. (1) expresses the rate at which the action density is changing with time. In the second term, the propagation of action density in both physical (x, y) and spectral (σ, θ) space is represented by the dot product of vector gradient $\vec{\mathbf{V}} = (\partial/\partial x, \partial/\partial y, \partial/\partial\sigma, \partial/\partial\theta)$ with vector velocity $\vec{\underline{\mathbf{C}}} = (C_x, C_y, C_\sigma, C_\theta)$ and action density N. The vector $\mathbf{C}$ represents the propagation rate of wave energy and includes contributions from both wave group velocity and ambient current. On the right-hand side of Eq. (1), S_{tot} is a consolidated spectral source term that may include contributions from nonlinear wave–wave interactions, wind wave forcing, dissipation due to breaking, bottom friction, and whitecapping [see Booij et al. (1999) for additional details]. Once initialized with bathymetry, boundary wave data, initial conditions, and source term configuration, SWAN solves Eq. (1) to predict the evolution of the wave spectrum throughout the modeled coastal region.

While data assimilation has been a component of atmospheric modeling since the 1960s, it has only recently been introduced in ocean wave models. A number of nearshore studies have used assimilated wave data from in situ instruments or video to estimate a variety of nonwave nearshore parameters, including bottom bathymetry (Lippmann and Holman 1990; van Dongeren

Corresponding author address: Mark D. Orzech, Naval Research Laboratory, Code 7320, Stennis Space Center, MS 39529.
E-mail: mark.orzech@nrlssc.navy.mil

et al. 2008; Narayanan et al. 2004; Aarninkhof et al. 2005), alongshore currents (Chickadel et al. 2003), and bottom friction (Keen et al. 2007). Early attempts to improve wave predictions with data assimilation used relatively simple optimal interpolation methods with observed wave heights and periods to nudge model results toward more accurate values (Bauer et al. 1992). More recent efforts have improved model performance further by incorporating wave spectral parameters (Hasselmann et al. 1997; Aouf et al. 2006; Portilla 2009). However, these features have primarily been used in larger-scale models such as WAVEWATCH III and WAM. The public domain version of SWAN does not yet include data assimilation in any form.

Walker (2006) developed a data assimilation system that used measured data from inside the computational domain to correct offshore boundary inputs for SWAN. In this system, the adjoint of reduced SWAN equations was derived analytically and then discretized. This ''analytical adjoint'' was used to propagate the model–observation errors to the boundary, where only stationary (i.e., steady state) conditions were considered, and all nonlinear source and sink contributions were set to zero, as in

$$\vec{\nabla} \cdot (\vec{\mathbf{C}}N) = 0. \tag{2}$$

This approach assumes that energy-generating or dissipating effects like wind forcing, wave breaking, and bottom friction are negligible. The steady state, linear form of SWAN represented by Eq. (2) is commonly used to investigate longer-term effects of an average or relatively constant wave climate. Recent validation tests indicate that Walker's assimilation system leads to improved predictions of integrated spectral wave properties like significant wave height, peak period, and mean direction in deeper water, as well as limited improvements in estimates of the full 2D spectrum (Veeramony et al. 2010). However, the technique does not capture changes to the spectra that occur due to the contributions of nonlinear source terms described above. Since the algorithm only corrects boundary conditions offshore, an ad hoc treatment of lateral boundaries is required. These shortcomings generally lead to significant errors in predicted wave heights and currents in shallow water.

The analytical adjoint approach cannot be adapted to incorporate nonlinear wave transformations in a consistent manner. It has been demonstrated (Bücker et al. 2011; Giresse and Walther 2004) that discretizing the analytical adjoint model (as done in Walker 2006) leads to an erroneous numerical gradient of the cost function. As a result, an analytical adjoint system generally will

not pass the asymptotic gradient test of its cost function (cf. Järvinen 1998). The proper and consistent way to obtain the numerical gradient of the cost function is to derive the adjoint of the discretized forward model (hereafter called the ''numerical adjoint''). The difference in the numerical gradients obtained from the two adjoint types is small when working with a stationary linear system, but it can be significantly larger for adjoints to nonstationary and/or nonlinear models (Favennec 2005).

For environments where wind forcing, bottom friction, currents, or triad–quadruplet interactions are significant, effective adjoint-based assimilation of wave data thus requires a numerical approach. A numerical adjoint to SWAN can include adjoint subroutines for all nonlinear sources and sinks of wave energy in the model, regardless of their initial discrete form. An assimilation system featuring a complete numerical adjoint will enable SWAN users to improve spectral estimates with assimilated wave data from nearly any domain, including storm-generated chop and swell, shallow surf zones, and other nonlinear wave environments. Construction of such an adjoint is labor intensive, however, requiring each adjoint subroutine to be created and validated separately.

The present paper describes the development and testing of a variational data assimilation system based upon a numerical adjoint that is limited to the stationary, linear SWAN-governing equations [Eq. (2)]. The assimilation system described herein is suboptimal, implementing strong-constraint variational assimilation in which just the boundary condition is controlled. Although fully consistent four-dimensional variational data assimilation (4DVAR) systems have been built and validated for large-scale circulation models, such as the Navy Coastal Ocean Model (NCOM; Ngodock and Carrier 2013), the Océan Parallélisé (OPA) model (Weaver et al. 2003), and the Regional Ocean Model System (ROMS; Moore et al. 2004), this is the first time one has been created for a wave model. The purpose of this study is to introduce this new approach to the nearshore community, demonstrate its feasibility, and evaluate the initial suboptimal system through a direct comparison with the equivalently restricted analytical adjoint created by Walker (2006) and tested by Veeramony et al. (2010). Consequently, the same physical assumptions are adopted here as in Walker's analytical approach (i.e., operating only under stationary conditions and excluding nonlinear sources and sinks in the adjoint). The following two sections summarize the general structure and theory of the assimilation system, the development of the adjoint to the simplified model, and the initial validation tests. In section 4, the new system's performance is compared directly to that of the Walker (2006) system in a set of twin
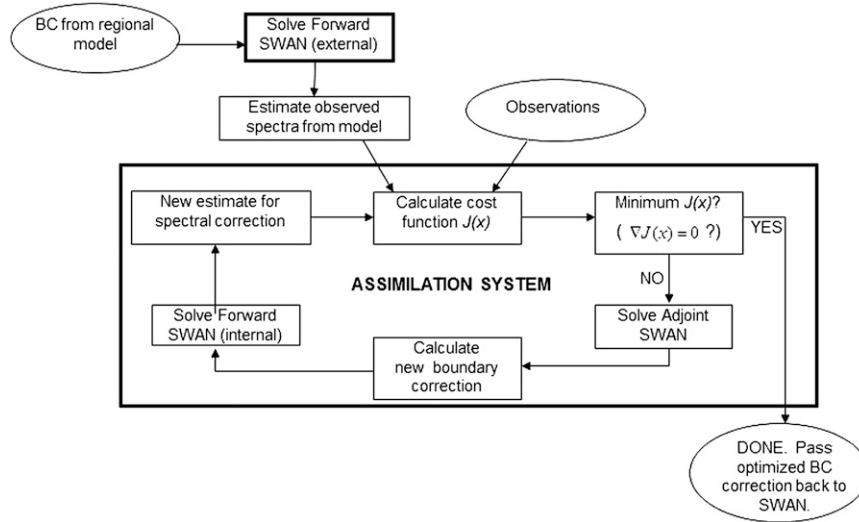
FIG. 1. Assimilation system flowchart.

experiment simulations, and then further examined using measured nearshore spectral data and bathymetry from Duck, North Carolina. Additional discussion is offered in section 5, followed by a brief summary of overall conclusions.

2. Assimilation system

As mentioned above, in the assimilation system presented here, the numerical adjoint encompasses only the stationary, homogeneous part of the SWAN model, including linear subroutines for the transport and refraction of wave action but no sources or sinks. The data assimilation procedure is built around an objective cost function $J(N)$, which is used to measure the aggregate model and measurement error at data locations within the model domain and has the following form:

$$J(N) = (\mathbf{BC} - \mathbf{BC}^b)^{\mathrm{T}}\mathbf{Q}_{\mathbf{BC}}^{-1}(\mathbf{BC} - \mathbf{BC}^b)$$

$$+ \sum_{k=1}^{K} (y_k - H_k N)^{\mathrm{T}}\mathbf{R}^{-1}(y_k - H_k N). \quad (3)$$

In Eq. (3), the first term on the right-hand side estimates the error in the boundary conditions as an actual value (**BC**) minus a modeled first guess (**BC**^{b}). The second term computes the weighted residual between the model and the observations, a sum over individual innovations $(y_k - H_k N)$ at each location k (e.g., Kalnay 2003; Park and Xu 2009; Lahoz et al. 2010; Swinbank et al. 2003). In this term, H_k is an observational operator, which projects the model solution onto the location of

observation y_k. Terms $\mathbf{Q}_{\mathrm{BC}}^{-1}$ and $\mathbf{R}^{-1}$ are weights based on the relative variance of each error source. The cost function is minimized by finding the set of controls (boundary conditions) for which the first variation (or gradient) of $J(N)$ is zero. An adjoint variable is introduced to facilitate the computation of the gradient of the cost function with respect to the controls. This leads to a set of coupled Euler–Lagrange equations that are solved iteratively (see appendix A for additional details). For the present analysis, only the boundary conditions are controlled, which for the steady-state system is equivalent to a strong-constraint approach to data assimilation, wherein the full model is assumed to generate no errors—that is, all errors are propagated into the domain from the boundary.

The assimilation system utilized here consists of the adjoint to linear homogeneous SWAN coupled to an internal copy of the full forward SWAN and supplemented by several additional modules that compute and track the cost function and optimize system convergence (Fig. 1). The system is wholly separate from the original forward SWAN, interfacing only at the beginning and end of the data assimilation process. To begin, the external forward SWAN is initialized with boundary data from regional models and used to compute estimated wave spectra at all measurement locations in the domain. An initial estimate for the cost function $J(N)$ is then obtained from the preliminary boundary conditions, SWAN spectra, and observed spectra, and its gradient is calculated at the offshore boundary using the Euler–Lagrange equations. A conjugate gradient technique (Polak and Ribière 1969; Walker 2006) is used to adjust the estimate toward its optimal value. If the cost

function is not yet minimized by the adjusted estimate, then the adjoint model is solved to obtain additional corrections to the boundary conditions, and the internal forward SWAN is run with the new boundary inputs to generate revised spectral estimates. The revised spectra are used to recompute the cost function, and the iterations continue in this manner until $J(N)$ is minimized and the optimal corrected boundary conditions are obtained. These are provided to the external forward SWAN, which computes the corrected solution for the entire domain.

Adjoint construction

Individual adjoint SWAN subroutines are built with the Parametric FORTRAN compiler (PFC) utility (Erwig et al. 2007), which uses automatic differentiation, applying standard rules in combination with user-specified active and dependent variables to linearize or transpose each forward subroutine. Active and dependent variables are identified in the forward subroutine and listed in a parameter file that is read by PFC. For example, if vector $\mathbf{c}$ is active in the iterative computation,

$$\text{do } i = 2 : \text{nlocs} - 1$$
$$\mathbf{a}(i + 1) = \mathbf{c}(i - 1) + \mathbf{B} \circ \mathbf{c}(i) + E$$
$$\text{enddo}, \tag{4}$$

where nlocs is the number of locations on a geographic grid, $\mathbf{a}$ is a dependent variable, $\mathbf{B}$ is a matrix of coefficients, and E is an independent fixed constant, then the (perturbation type) tangent linear version of Eq. (4) has essentially the same form, except that the unperturbed constant E is missing as shown:

$$\text{do } i = 2 : \text{nlocs} - 1$$
$$\delta\mathbf{a}(i + 1) = \delta\mathbf{c}(i - 1) + \mathbf{B} \circ \delta\mathbf{c}(i)$$
$$\text{enddo}. \tag{5}$$

The adjoint of Eq. (4) [and (5)] provides the sensitivity of the observation location with respect to the boundary control. It propagates a first derivative at a given observation location back to the boundary control, taking into account the model state and physics. It can be used to determine how the error in an estimate of $\mathbf{c}$ at a given grid location is related to a corresponding error in $\mathbf{c}$ at the boundary, in the presence of the field specified by Eq. (4). In the adjoint, time and spatial loops are reversed along with the relationships of the active variables, the matrix of coefficients is transposed, and the resulting expression for each adjoint variable becomes cumulative over the loop as shown:

$$\text{do } i = (\text{nlocs} - 1) : 2$$
$$\mathbf{ad\_c}(i - 1) = \mathbf{ad\_c}(i - 1) + \mathbf{ad\_a}(i + 1)$$
$$\mathbf{ad\_c}(i) = \mathbf{ad\_c}(i) + \mathbf{B}^T \circ \mathbf{ad\_a}(i + 1)$$
$$\text{enddo}. \tag{6}$$

In the adjoint model as a whole, subroutines themselves are also called in reverse order of those in the forward model. Variables that are not active or dependent are computed in the adjoint as they are in the forward model, and they must have the same values as their forward counterparts at the same times and locations. Equation (2) and its adjoint are presented in discrete form in appendix B.

Subroutines utilized by stationary homogeneous SWAN are generally already linear, so the tangent linear (TL) version of these subroutines is essentially unchanged from the original forward version [similar to Eqs. (4) and (5)]. Active and dependent variables in SWAN normally include the wave action and anything that depends on it. Examples of independent variables include the energy propagation velocity $\overline{C}$ and the water depth. As each numerical adjoint subroutine in this system is derived directly from its forward counterpart, the principal physics of the original subroutine are preserved and the adjoint model retains the associated properties of the discretized linear forward model. To improve overall accuracy, all real arrays in the assimilation system's adjoint and forward SWAN programs are compiled with double precision. Using these methods, a full set of adjoint subroutines is constructed for stationary homogeneous SWAN.

3. Validation

In matrix notation, the action of a given linearized SWAN subroutine upon an input array u, generating an output array v, may be expressed as $\mathbf{A}\mathbf{u} = \mathbf{v}$, while that of the corresponding adjoint subroutine can be written as $\mathbf{A}^T\mathbf{v} = \mathbf{u}$. For a real matrix $\mathbf{A}$ with transpose $\mathbf{A}^T$ and arbitrary real vectors $\mathbf{u}$ and $\mathbf{v}$, the following inner product is an identity:

$$\langle \mathbf{A}\mathbf{u}, \mathbf{v} \rangle \equiv \langle \mathbf{u}, \mathbf{A}^T\mathbf{v} \rangle. \tag{7}$$

For a given subroutine and its adjoint, this implies that if we initialize the forward subroutine with an arbitrary vector $\mathbf{u}$ (obtaining an output $\mathbf{A}\mathbf{u}$) and then initialize its adjoint with a second arbitrary vector $\mathbf{v}$ (obtaining an output $\mathbf{A}^T\mathbf{v}$), then an inner product of the forward routine output with the adjoint input should have the same value as an inner product of the adjoint output with the forward input.

This identity is also exhibited in the symmetry of the representer matrix (Bennett 2002), which provides a convenient method for testing the consistency of the complete adjoint program and its forward counterpart. For a given number m of observation locations (in the spatial–spectral domain), the representer function for each location is obtained by initializing the adjoint with a unit impulse centered at that location. The adjoint solution (at all grid points) is then used to force the forward model (running in stationary mode with no sources or sinks). The forward SWAN solutions at the m observation locations then fill one column of the m-by-m representer matrix ($\mathbf{HAA}^T\mathbf{H}^T$, where $\mathbf{H}$ is composed of m observational operators, H_k, for $k = 1$ to m). If the adjoint is consistent with the stationary linear forward model, then the representer matrix should be symmetric to machine precision (Ngodock and Carrier 2013).

Subroutines in the new adjoint are validated individually by means of the above-mentioned inner product test [Eq. (7)]. In each test, input vectors $\mathbf{u}$ and $\mathbf{v}$ are assigned different sets of random values, and inner product values are checked at multiple locations on an idealized bathymetry. For each individual subroutine or group of related subroutines tested, Eq. (7) is consistently satisfied to within machine precision. Specifically, for all such tests,

$$\frac{\langle \mathbf{Au}, \mathbf{v}\rangle - \langle \mathbf{u}, \mathbf{A}^T\mathbf{v}\rangle}{\langle \mathbf{Au}, \mathbf{v}\rangle} \leq 5 \times 10^{-15}.$$

Consistency of the full adjoint code is checked with an impulse test, in which the representer matrix is computed for a limited number of points, as outlined above, and its symmetry is checked. The tests are conducted using bathymetry from Santa Rosa Island, Florida (Fig. 2), and selecting four grid locations (TA1, TA2, SAB, SIB) matching those of actual instruments in a 2009 experiment (Edwards et al. 2009; Veeramony et al. 2010). Locations TA1 and SAB are closer to the shoreline, where bathymetry is alongshore uniform and some wave shoaling and breaking may occur, while TA2 and SIB are outside the surf zone, near the southwestern and southeastern grid boundaries, respectively. In four separate subtests, the adjoint is initialized with a unit impulse at a single frequency and direction for each of the four locations described above. For example, at location TA1 the initialization is

$$\nu(x_{TA1}, y_{TA1}, f = 0.066\,\text{Hz}, \theta = 50^\circ) = 1, \quad \nu = 0 \text{ elsewhere}.$$

In each subtest, the adjoint is allowed to run for a single iteration. Its output is used to initialize forward SWAN
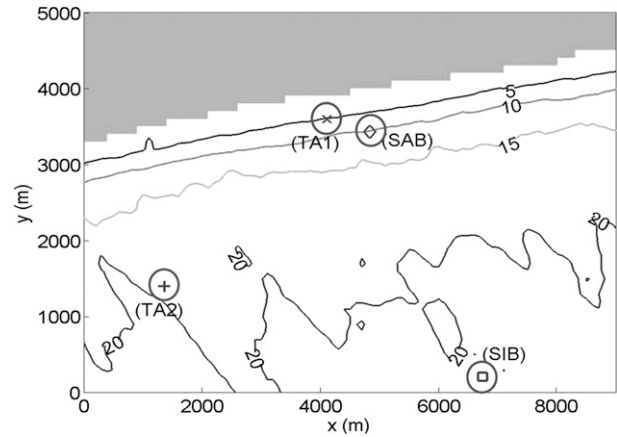

FIG. 2. Model bathymetry for inner product, impulse and twin experiment tests based on measurements at Santa Rosa Island, FL, recorded in 2009. Direction on the y axis is north. Shoreline is shaded. The four grid locations used for impulse and twin tests are circled and labeled.

(in stationary linear mode), whose output is in turn recorded at all four of the selected locations. The results are compiled into a (4×4) subsection of the representer matrix, which is found to be symmetric at all locations within model accuracy of order 10^{-12}, thus confirming that the full adjoint is consistent with the stationary linear forward model.

4. Simulations

Two sets of simulations are conducted to further examine the performance of the numerical adjoint. First, the numerical adjoint approach as a whole is evaluated by direct comparison with the Walker (2006) approach, using a twin experiment where artificial observations are generated by the full forward model with idealized boundary conditions. Following this, the assimilation system is evaluated with measured spectra from multiple instruments at Duck, North Carolina. For both tests, errors in boundary control are assumed to be completely spatially correlated and uniform. Final spectral estimates are compared to observations at boundary and interior grid locations.

a. Twin experiment

The idealized twin experiment validation tests are based directly on those conducted by Veeramony et al. (2010) as part of their original evaluation of the Walker (2006) adjoint, utilizing data and bathymetry from the aforementioned 2009 experiment at Santa Rosa Island, Florida. The fully nonlinear forward SWAN, initialized at seaward and lateral boundaries with a parametric Joint North Sea Wave Project (JONSWAP) wave spectrum, is
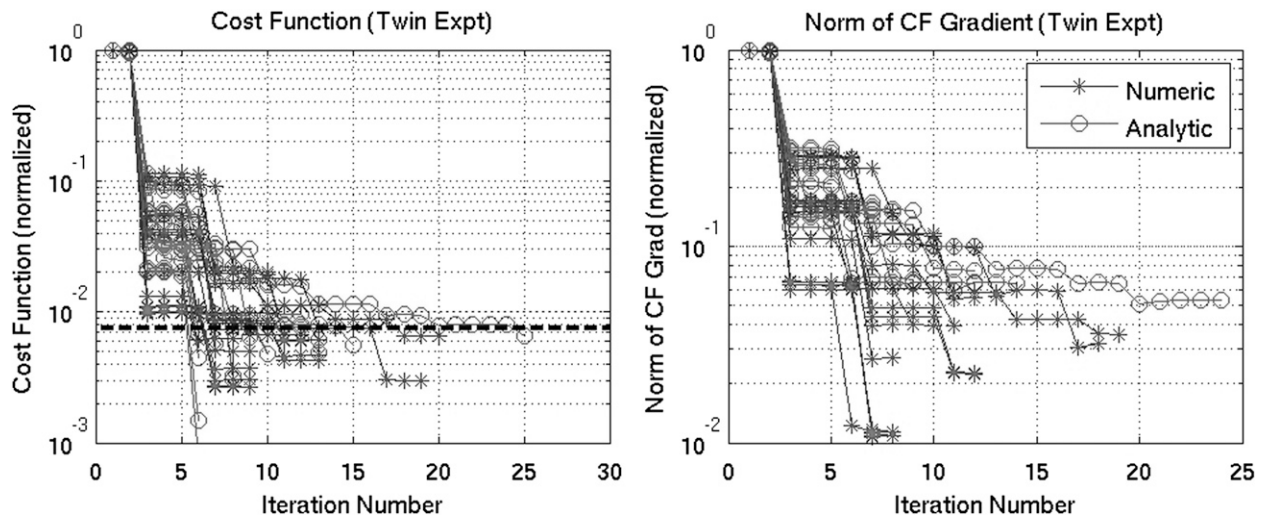
FIG. 3. Evolution of (left) cost function and (right) norm of gradient of cost function plotted vs assimilation iteration for all twin experiment tests. Asterisks and circles represent numeric and analytic adjoint system results, respectively. All values are normalized by the value at the first iteration. Heavy dashed line (left) marks the 0.75% criterion for the cost function below which each assimilation is stopped. Note log format on the y axis.

used to compute pseudo-observed spectra at the same four locations used in the 2010 analysis (Fig. 2). JONSWAP parameters are based on prevailing waves during the experiment, which were relatively mild—significant wave height (Hs) < 1.5 m—and approached generally from the southeast and normal to the offshore grid boundary. Four of the six cases from Veeramony et al. (2010) are rerun for both assimilation systems (cases 4 and 6, featuring local wind forcing nonlinearities, are excluded). For each case, four separate assimilation tests are conducted with each adjoint system. In the first of these tests, the adjoint is initialized with the innovation at grid location TA1; in the second test, at TA2; in the third test, at SAB; and in the fourth test, at SIB. The system assumes a zero-energy state as a first guess (i.e., spectra equal to zero at boundary and, consequently, throughout the grid), so the initializing innovation for a given test (i.e., the observed minus the estimated spectrum) is simply equal to the observed spectrum at that location.

In each test, adjoint-computed values of spectral density at the boundary nodes are averaged across all locations on the seaward boundary and provided as uniform offshore boundary conditions to the assimilation system's internal forward SWAN [with wave breaking and other source/sink subroutines activated, duplicating the setup of Veeramony et al. (2010) for each case]. Along lateral boundaries, spectra are set to a tapering fraction of the full offshore spectrum, ranging from 100% at the seaward end to zero at the shoreward end (SWAN Team 2011). A revised estimate for the

wave spectrum at the initializing location is obtained and compared with the original data spectrum. If the cost function is not yet minimized, then the adjoint is reinitialized with a recomputed innovation. Each assimilation system is allowed to run until its cost function declines to 0.75% of its original value. By this point the associated cost function gradients have been reduced to less than 5% of their initial values and are generally flattening out, as are the cost function values themselves (Fig. 3).

In each test, the two assimilation systems' relative performance is qualitatively evaluated by comparing how well they reproduce the observed spectra at the offshore boundary and all four instrument locations, including nonassimilated spectra as well as the selected innovation spectrum. For a more quantitative comparison, overall model accuracy is also evaluated using an RMS skill score computed from spectral densities as shown:

$$\text{skill} = 1 - \frac{\sqrt{\sum_{i,j}[S_{\text{mod}}(f_i, \theta_j) - S_{\text{obs}}(f_i, \theta_j)]^2}}{\sqrt{\sum_{i,j}[S_{\text{obs}}(f_i, \theta_j)]^2}}. \qquad (8)$$

In Eq. (8), S_{mod} is the model spectrum and S_{obs} is the observed spectrum (from nonlinear forward SWAN). Spectral energy densities are first squared and then summed over all frequencies (f_i) and directions (θ_j).

In general, spectral estimates from the present assimilation system are very similar to those from the Walker (2006) system. Both systems have errors of similar
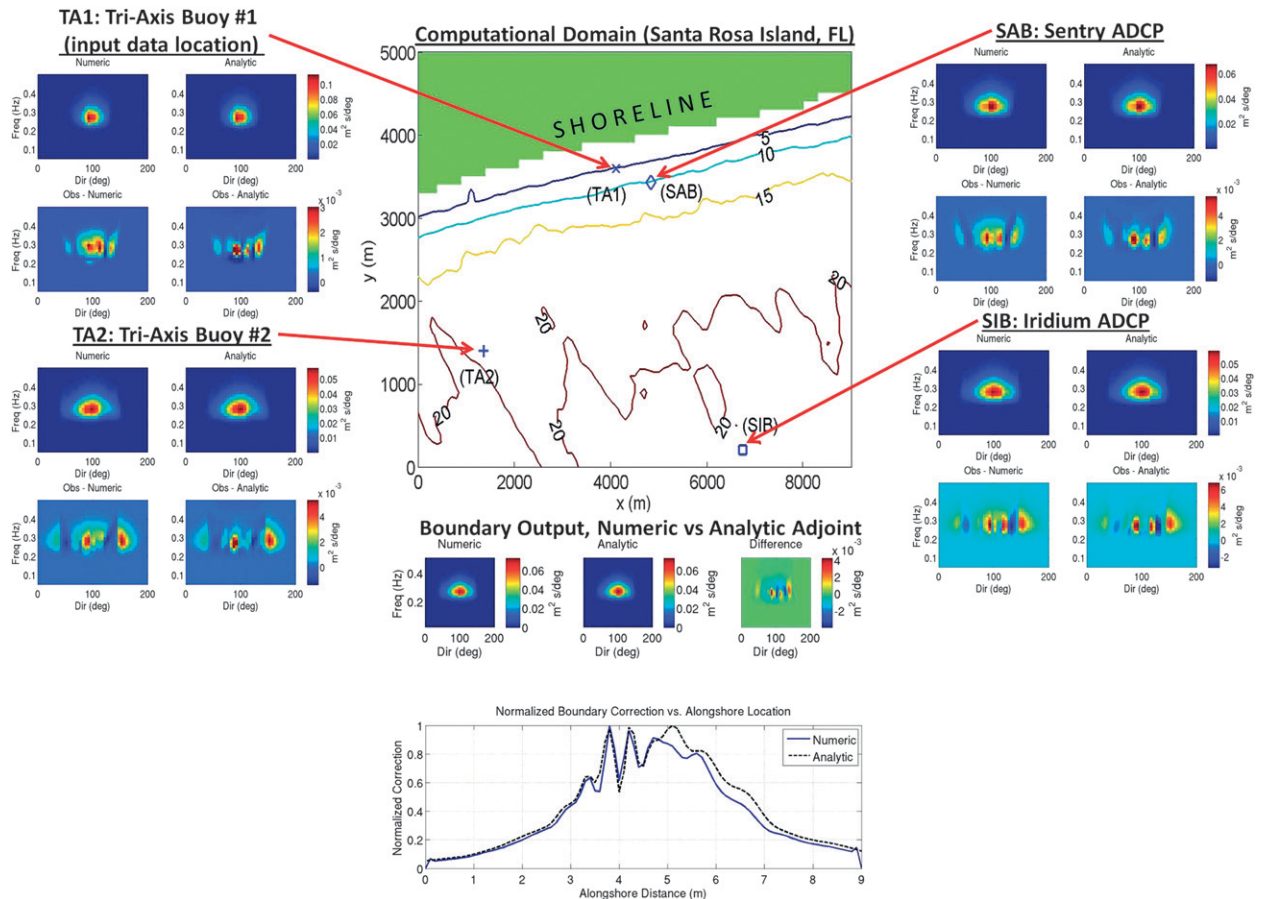
FIG. 4. (top panels) Twin experiment results, comparing stage 1 numerical adjoint to Walker (2006) analytical adjoint. Both are initialized with data from nonlinear forward SWAN at the TA1 location. Four-panel plots for each of the four interior locations (TA1, TA2, SAB, and SIB) include spectral results from numeric/analytic-adjoint-based assimilation systems (top two panels of each set) and the difference between those results and the original SWAN spectra (bottom two panels). The three panels below the grid map show mean predicted boundary spectrum: numeric, analytic, and numeric − analytic. (bottom) Comparison of normalized distribution of adjoint boundary energy along offshore boundary for the two adjoint types.

magnitude when compared to observed spectra (e.g., Fig. 4). When skill values are computed and averaged over all twin experiment tests (64 computations per adjoint), the numerical assimilation system achieves a mean skill of 0.914 (range = 0.82–0.98), roughly equivalent to the analytical adjoint system's mean of 0.911 (range = 0.78–0.97).

b. Observed data—Duck, North Carolina

The final evaluation of the simplified numerical adjoint tests how well the assimilation system can recapture actual measured spectra, starting from an essentially arbitrary first guess. Spectral wave datasets are obtained from the Army Corps of Engineers' Field Research Facility (FRF) in Duck, North Carolina, for several periods that also include nearshore bathymetry measurements. To minimize nonlinear effects, study dates are selected from relatively mild wave climates in

which significant wave heights are less than 1.5 m (Fig. 5). Bathymetry data recorded for the Duck "minigrid" on 16 April, 1 June, and 29 July 2010 are linearly extrapolated alongshore to 5 times the minigrid width and offshore to 17-m depth. Spectral data are available from up to six nearshore instruments (Fig. 6), including four Nortek acoustic wave and current (AWAC) profilers at depths of 5, 6, 8, and 11 m, and the FRF's 8-m array of 15 near-bottom-mounted pressure sensors. Boundary spectra for the model grid are provided by a Datawell Waverider 630 buoy, moored approximately 3 km offshore in roughly 17-m water depth. For the 16 April and 1 June simulations, four instruments are used (the 5 and 8 m AWAC profilers were not available). For 29 July, all six instruments are used.

Owing to the manner in which they are derived, spectra from the four AWACs have lower resolution than those from the 8-m array and the Waverider, and
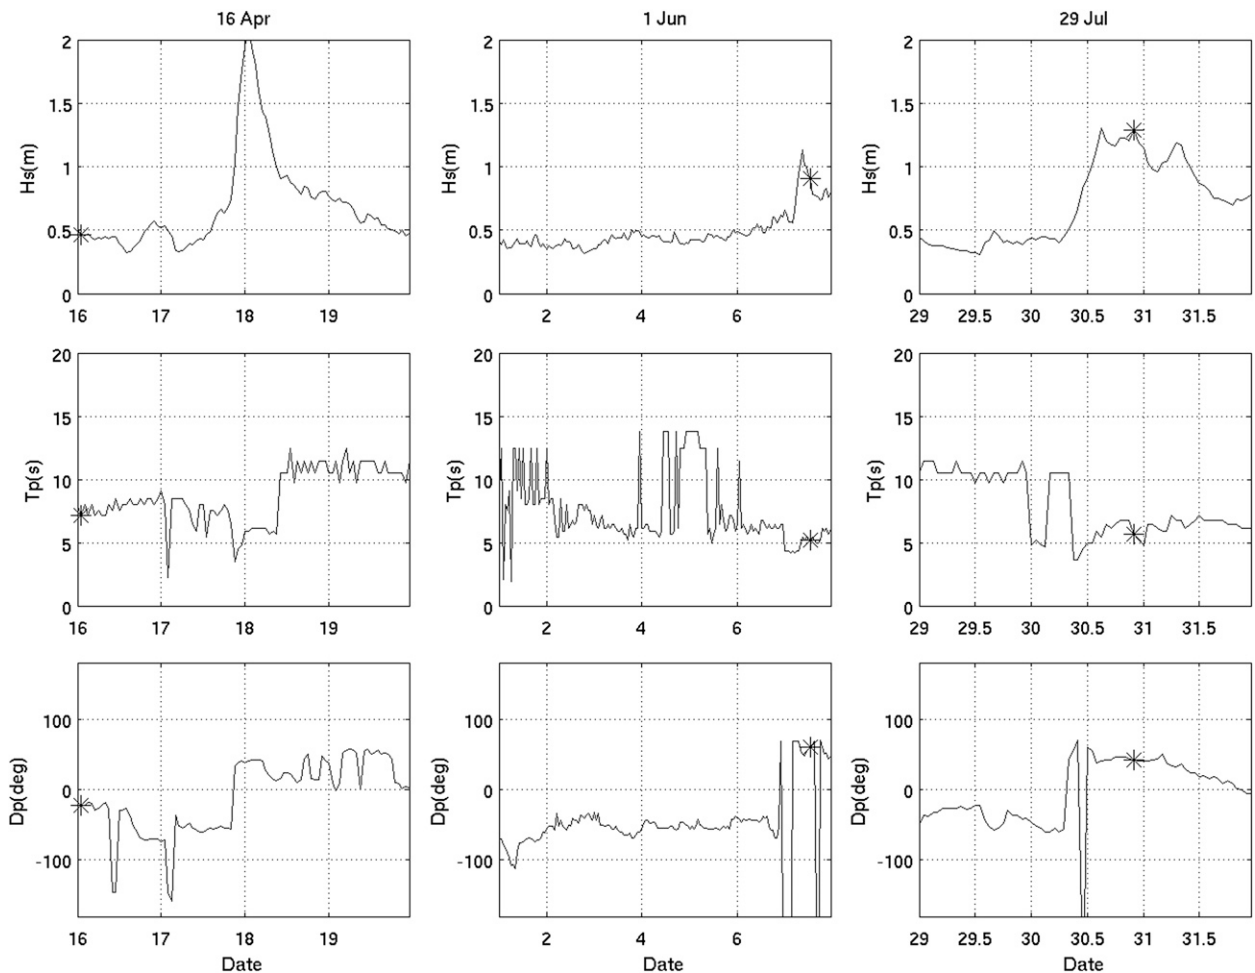
FIG. 5. (left to right) Time series of (top to bottom) significant wave height, peak period, and peak direction as measured at the 3-km Waverider buoy for periods roughly concurrent with the three bathymetry measurements included in these tests. Conditions at times of selected initializing spectra are marked with asterisks in each panel.

they also tend to have a significantly broader directional spread, with a small-to-moderate amount of wave energy directed offshore at some sensors. It is not known how much of this additional spreading is a result of instrument error and how much is due to unrepresented nonlinear effects such as wave breaking and reflection. The current assimilation system does not presently assimilate offshore-directed wave energy. To more effectively evaluate actual system capabilities, assimilated spectra are therefore limited to onshore directions in these tests (i.e., all offshore-directed spectral energy—generally much less than one percent of total energy—is set to zero in observed spectra). Even with this limitation, the varied directional spread among initializing instruments has a noticeable effect on model output that will be discussed further below.

The present tests are designed to evaluate the system's ability to assimilate real spectral data and investigate its

sensitivity to the observed data locations. For each date, the adjoint is initialized with innovations at one or more locations in nine different cases. In case 1, only spectra from the FRF 8-m array are used. In case 2, the adjoint is initialized at all available AWAC profiler locations, with each innovation equally weighted. Case 3 utilizes equally weighted spectra from all (three or five) interior grid locations. In case 4, all interior spectra are again used, but the 8-m array is given additional weight. Case 5 uses all interior spectra and instead gives additional weight to each of the AWAC profilers (distributed evenly among them). Cases 6, 7, 8, and 9 are initialized only at the 5, 6, 8, and 11 m AWAC profilers, respectively (where available).

As with the twin experiment, a zero-energy spectrum is again used for the first guess, so that at each location the innovation is just equal to the observed spectrum for the first iteration. Wave breaking is activated for the
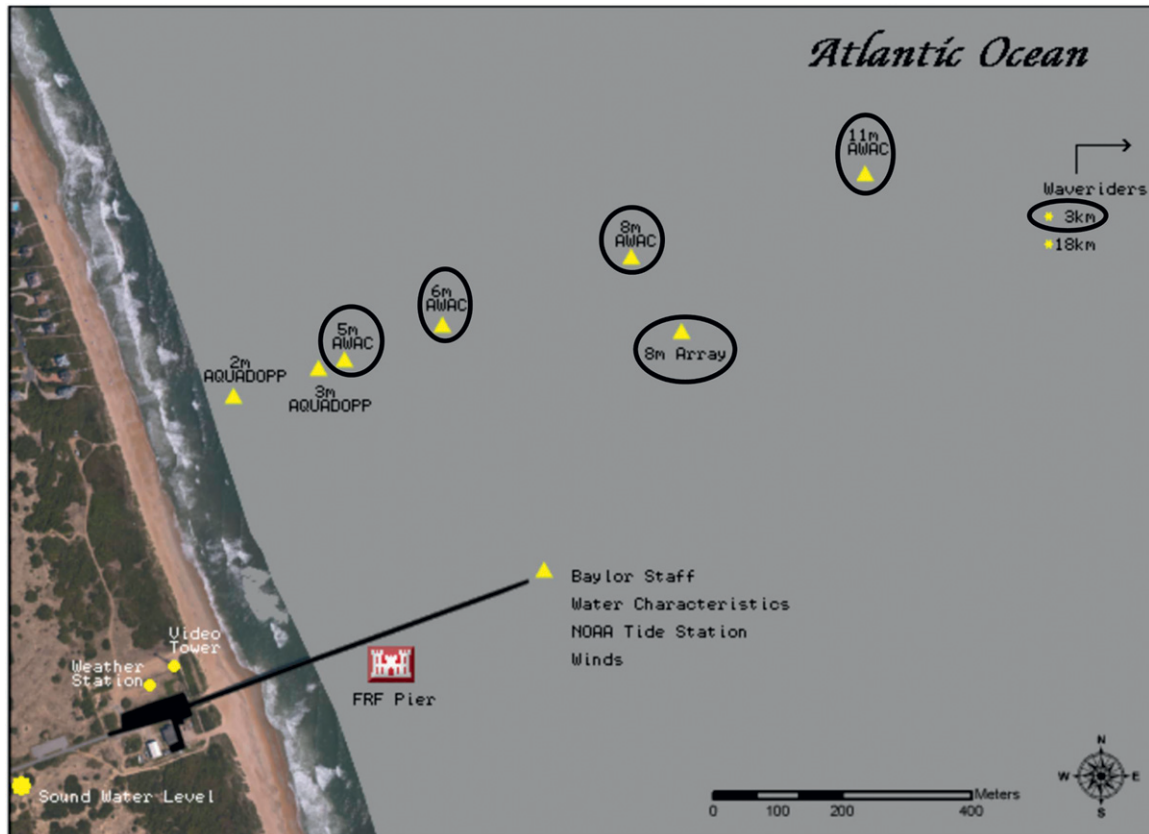
FIG. 6. Map of nearshore instrument locations at FRF (offshore Waverider buoys not shown). Instruments used in this study are circled. (Image courtesy of FRF, http://www.frf.usace.army.mil/frfzoom.shtml.)

system's internal forward SWAN model to maintain system stability, but other nonlinear sources and sinks are not activated. Each simulation is allowed to run until formal convergence has been achieved (i.e., the cost function value has become essentially constant). Final estimated spectra are recorded at the seaward boundary and at all observation locations, and wave statistics

including significant wave height, mean direction, mean period, and directional spread are computed from each modeled and observed spectrum.

Model skill [evaluated using Eq. (8)] ranges from a low near zero to a high of 0.92, with a mean skill of 0.52. Model results for the test series are summarized in Table 1. The highest skill values (averaging 0.73) are obtained for

TABLE 1. Adjoint skill scores for tests at Duck, NC. Tests were conducted for three separate dates in 2010 (16 Apr, 1 Jun, and 29 Jul 7). Where data were available from multiple dates, a range of skill scores is provided in the table. First column shows case number and second column shows initializing location(s) for the adjoint system. AWACs implies all available AWAC sensors were used; All implies all available sensors were used. Overwt is used for cases where one or more sensors were overweighted relative to other sensors used.

| Case	Initializing location	Output location				
		FRF	AWAC (5 m)	AWAC (6 m)	AWAC (8 m)	AWAC (11 m)
1	FRF	0.67–0.92	0.42	0.19–0.50	0.26	0.00–0.51
2	AWACs	0.32–0.61	0.53	0.61–0.67	0.52	0.63–0.67
3	All	0.50–0.73	0.55	0.60–0.67	0.48	0.48–0.64
4	All Overwt FRF	0.76–0.91	0.44	0.32–0.53	0.15	0.26
5	All Overwt AWACs	0.37–0.65	0.55	0.63–0.67	0.51	0.08–0.62
6	AWAC (5 m)	0.47	0.71	0.52	0.29	0.27
7	AWAC (6 m)	0.35–0.58	0.51	0.68–0.72	0.44	0.45–0.56
8	AWAC (8 m)	0.39	0.45	0.56	0.58	0.60
9	AWAC (11 m)	0.26–0.53	0.39	0.44–0.55	0.47	0.70–0.79
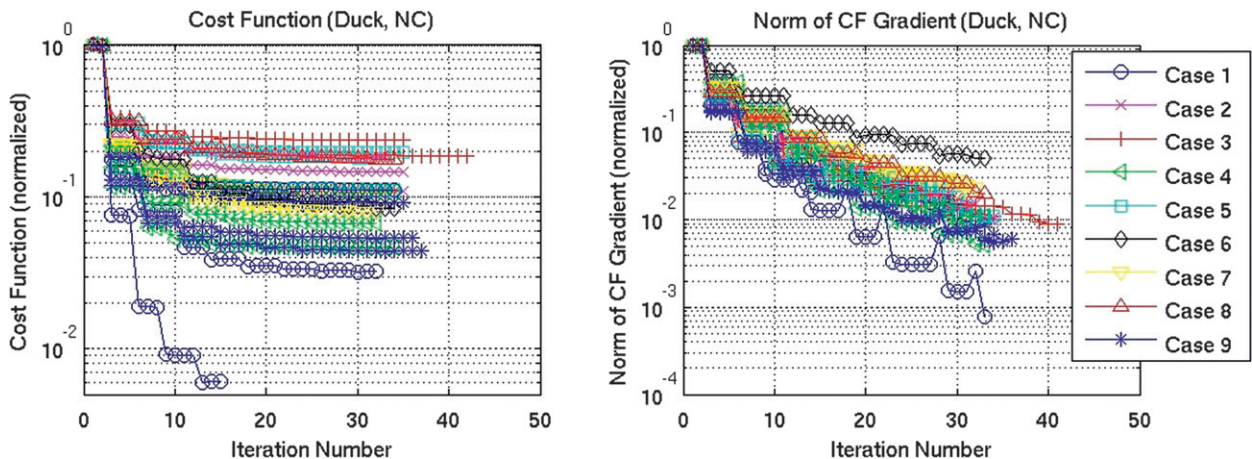
FIG. 7. Evolution of (left) cost function and (right) norm of gradient of cost function, plotted vs assimilation iteration for all adjoint simulations, with colors and symbols representing the nine different cases shown in Table 1. All values are normalized by the value at the first iteration. Note log format on *y* axis.

tests in which a single location (e.g., the 8-m array) is used for both input and output. In general, spectral estimates are poor for cases where the initializing instrument type differs from the instrument type at the output location. The worst results (average skill = 0.38) are obtained for spectral estimates at AWAC profiler locations in cases 1 and 4, when the model is initialized either solely at the 8-m array or at all locations with the 8-m array overweighted. The relative change in the cost function and its gradient is shown for all cases in Fig. 7. Cost functions and gradients decrease most rapidly for tests initialized at the farthest offshore sensors (cases 1 and 9), particularly when waves are small (16 April). The slowest convergence is obtained when using data from just the shallowest AWAC profiler (case 6) on the largest wave day (29 July), while the smallest reduction of the cost function is obtained for equally weighted tests initialized at all available sensors (case 3).

Modeled and observed spectral statistics are compared at all grid locations for all cases in Fig. 8. Regardless of initializing data and weights, the model consistently underpredicts significant wave height. The assimilation system appears to perform slightly better for cases with multiple initializing sensors. This is not surprising; using a larger number of initializing data locations of comparable quality will generally improve modeled wave heights. Although they have lower spectral resolution, the AWAC profilers capture spectral moments roughly as well as the 8-m array (K. Hathaway 2011, personal communication). Wave height results appear to be somewhat better for smaller waves, suggesting that neglected effects such as wind forcing, wave breaking, and triad interactions (per their formulations in the SWAN model) likely play a role in generating these errors.

Mean wave directions are recovered fairly well, although model estimates are slightly closer to shore normal (90°) than observations. Mean periods are captured quite well in nearly all cases, with a very high correlation and consistent 1:1 ratio of estimates to observations. Some error may result from neglected wind forcing and triad/quadruplet interactions, which shift observed wave energy to different frequencies (with different refractive properties). Directional spreading results are poor, largely owing to differences in the spectra produced by each type of instrument. Predicted directional spreads at AWAC profiler locations in cases 1 and 4 are consistently lower than observed values at the AWAC profilers. This is expected because of the stronger effects of the directionally narrower 8-m array spectra on model output in these cases. Inversely, in cases 2 and 5–9, modeled directional spreads are overestimated at the 8-m array location because the adjoint is more strongly influenced by the directionally broad AWAC profilers.

5. Discussion

Twin experiment spectral estimates from the data assimilation system that uses the numerical adjoint are consistently as good as those from the Walker (2006) system built around an analytical adjoint, and the two models show equal levels of skill. Results from the numerical adjoint tests with observed data at Duck, North Carolina, are good for small wave cases and become fair or worse as spectral energies increase. These comparisons illustrate some of the shortcomings of using the reduced forward model to derive the adjoint, particularly in estimates of significant wave height and directional spread.
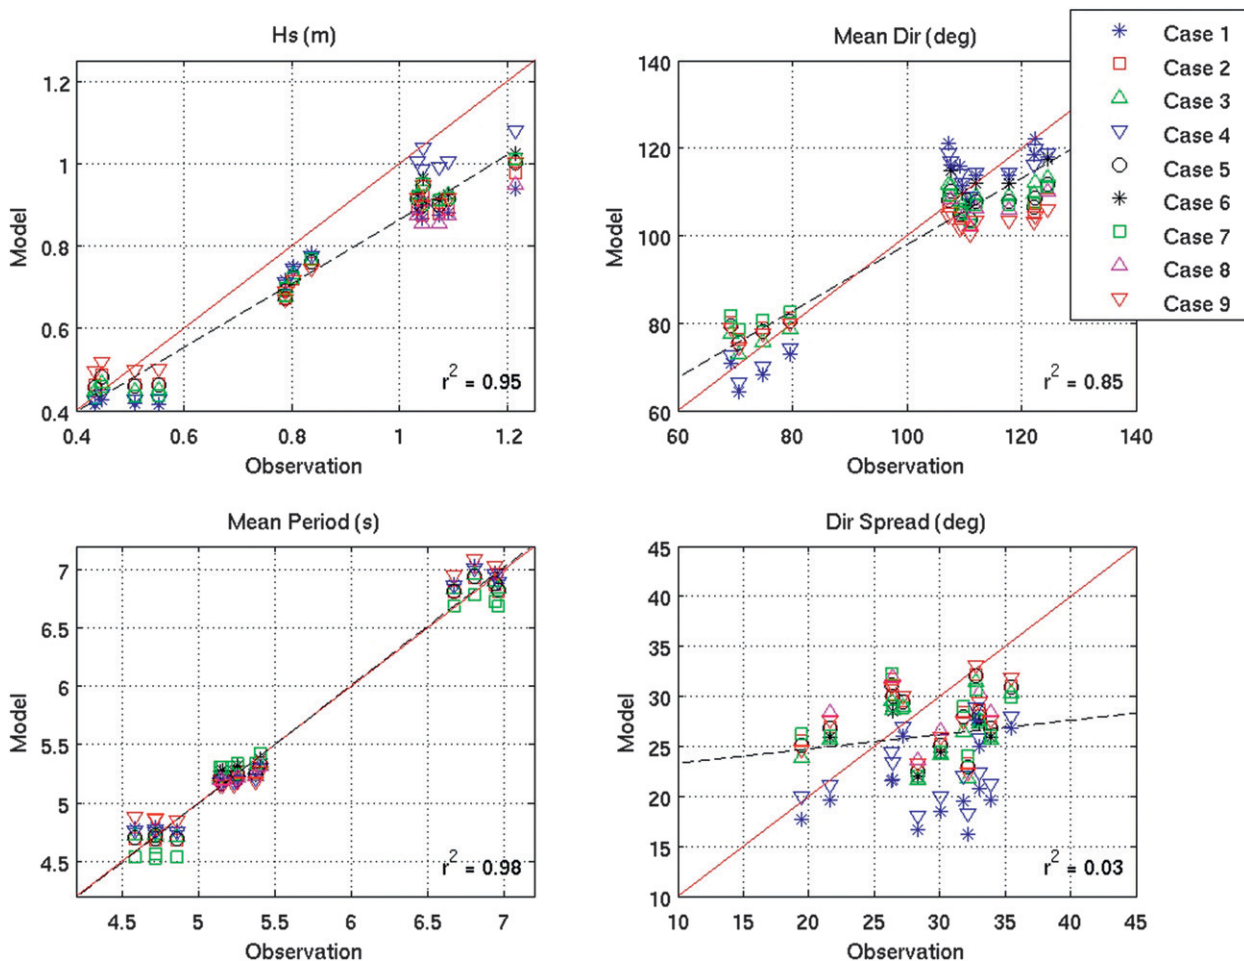
FIG. 8. Summary plots of modeled vs observed spectral statistics for all cases, showing results at interior and boundary data locations: (top) (left) Hs and (right) mean direction; and (bottom (left) mean period and (right) direction spread. Nine symbols represent each of the initializing scenarios described in the text. The solid line in each panel represents perfect agreement between model and data. The dashed line shows the best linear fit to the data; its r^2 value is provided in each panel. For the mean direction panel, shore normal is 90°.

Nonlinear sources of wave energy from wind forcing and wave–wave interactions are not represented in the adjoint code, although using the numerical approach to derive the adjoint will ultimately make this representation possible. Wave breaking was activated in the system's internal forward SWAN model, but its counterpart was not available in the adjoint. This mismatch between the forward and adjoint SWAN components of this suboptimal assimilation system may contribute to the underestimation of larger wave heights by the model. Directional spreading also appears to increase at the shallower AWAC profiler sensors, at least partially owing to triad and quadruplet interactions in or near the surf zone, both of which are also neglected in the numerical adjoint.

To a considerable extent, however, system performance is hampered by the use of spectra from different types of instruments. Spectral processing differences between the AWAC profilers and the 8-m array significantly limit the adjoint's performance in matching observed spectra (e.g., Fig. 9, left and center panels). Adjoint estimates in these tests might be improved if the directional spreads of all input spectra could be normalized (e.g., using the relative directional spreads measured concurrently at the 8-m array and the 8-m AWAC profiler). However, this modification would have to be based on an explicit mathematical representation of the complex relationship between the different techniques used to derive spectra at each instrument, which may or may not be derivable. Spectra from the FRF 8-m array of 15 pressure sensors are obtained with an iterative maximum likelihood estimator (IMLE) applied to 8192-s time series (USACE Field Research Facility 2011). In contrast, AWAC profiler spectra are generated using the combined
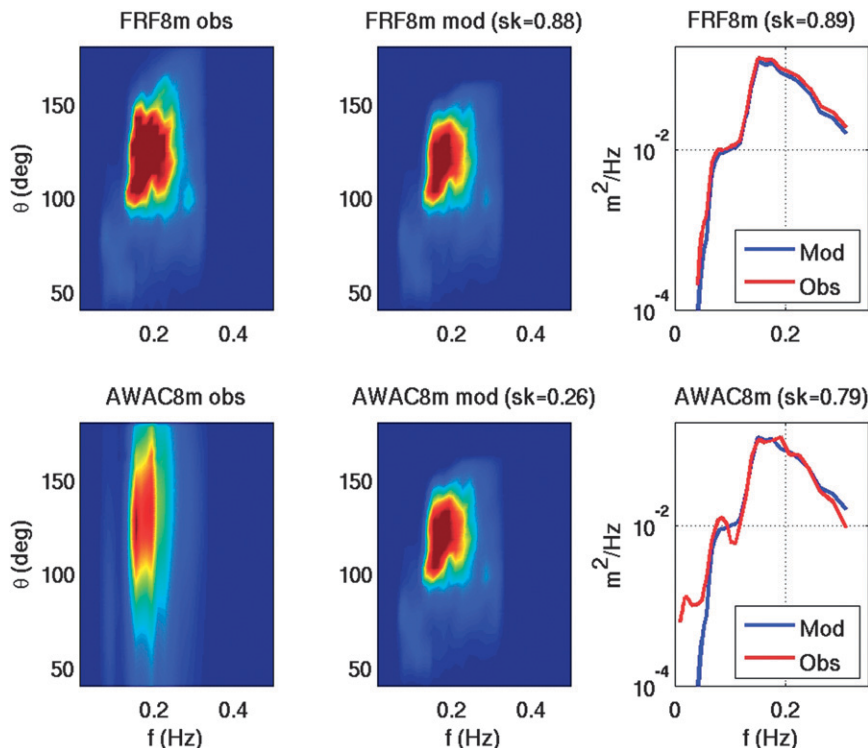
FIG. 9. (left) Observed and (middle) modeled frequency–directional spectra at the (top) FRF array and (bottom) 8-m AWAC profiler for 29 Jul 2010. (right) Corresponding modeled and observed frequency spectra for each instrument (log scale on *y* axis). Skill values for 2D and 1D estimates are provided in parentheses. Model results are for case 1, in which the adjoint was initialized only with a spectrum from the FRF 8-m array.

three-beam UV velocity data together with pressure measurements in another IMLE analysis, with slightly different methods used for low- and high-frequency ranges. The offshore Waverider buoy employs Fourier transform techniques with time series of vertical and horizontal buoy displacements, and an MLE analysis is applied to the resulting Fourier coefficients to obtain directional spectra (K. Hathaway 2011, personal communication).

To remove directional spreading from consideration in evaluating model performance, a second set of skill calculations is made using modeled and observed frequency spectra instead of frequency-directional spectra. The revised skill score expression is

$$\text{skill}_f = 1 - \frac{\sqrt{\sum_i [S_m(f_i) - S_o(f_i)]^2}}{\sqrt{\sum_i [S_o(f_i)]^2}}, \qquad (9)$$

where S_{mod} and S_{obs} from Eq. (8) have been summed over directions to obtain S_m and S_o above, respectively. The resulting skill scores (Table 2) are significantly better than those obtained with 2D spectra. For calculations based on Eq. (9), model skill at Duck, North Carolina, now ranges from 0.51 to 0.95 with a mean overall skill of 0.71. This result suggests that, although directional spreading differences are dominant, there are also some frequency domain variations between spectra from the 8-m array and the AWAC profilers. For the sample case in Fig. 9, a comparison of 1D frequency spectra (right two panels) illustrates the considerable improvement in model performance when directional spread is not considered. In this case, while model skill at the 8-m array improves marginally (from 0.88 to 0.89), there is a much more substantial increase at the 8-m AWAC profiler (from 0.26 to 0.79). Remaining spectral differences at the 8-m AWAC profiler are likely partly due to instrumental variations and in part to nonlinear effects that are neglected in the adjoint.

To further examine the importance of instrument type, the tests in section 4b are rerun, replacing observed data at each location with artificial spectra generated using fully nonlinear forward SWAN. The forward model is initialized with data from the offshore Waverider buoy, and wave breaking, triad interactions, wind

TABLE 2. As in Table 1, but for adjoint skill scores for tests at Duck, NC, computed in frequency space only using Eq. (9).

Case	Initializing location	Output location				
		FRF	AWAC (5 m)	AWAC (6 m)	AWAC (8 m)	AWAC (11 m)
1	FRF	0.83–0.92	0.56	0.59–0.67	0.79	0.51–0.77
2	AWACs	0.68–0.76	0.63	0.67–0.84	0.74	0.67–0.74
3	All	0.73–0.85	0.66	0.69–0.80	0.78	0.62–0.78
4	All Overwt FRF	0.82–0.95	0.70	0.62–0.76	0.85	0.54–0.83
5	All Overwt AWACs	0.71–0.78	0.66	0.70–0.83	0.77	0.66–0.77
6	AWAC (5 m)	0.67	0.74	0.69	0.74	0.72
7	AWAC (6 m)	0.57–0.83	0.64	0.74–0.76	0.73	0.55–0.71
8	AWAC (8 m)	0.63	0.55	0.60	0.71	0.66
9	AWAC (11 m)	0.51–0.78	0.63	0.66–0.76	0.72	0.78–0.83

forcing (10 m s^{-1}, shore-normal), and quadruplet interactions are all activated. Revised assimilation system estimates of wave statistics are plotted versus those from the new set of SWAN-based "uniform instrument" observations in Fig. 10. The results are significantly better than those presented in Fig. 8, especially with respect to directional spreading estimates (bottom-right panel). Using the idealized input spectra, model skill now ranges from 0.54 to 0.93 with a mean overall skill of 0.82 (Table 3), a considerable improvement upon the levels seen with the FRF dataset and also somewhat better, on

average, than the results obtained with 1D frequency spectra. When 1D skill values are computed for these results using Eq. (9), model skill ranges from 0.68 to 0.93, with a mean skill of 0.87.

6. Conclusions

A discrete numerical adjoint to the wave model SWAN has been constructed by creating individual numerical adjoints to all subroutines in the stationary, homogeneous part of the forward model. The adjoint
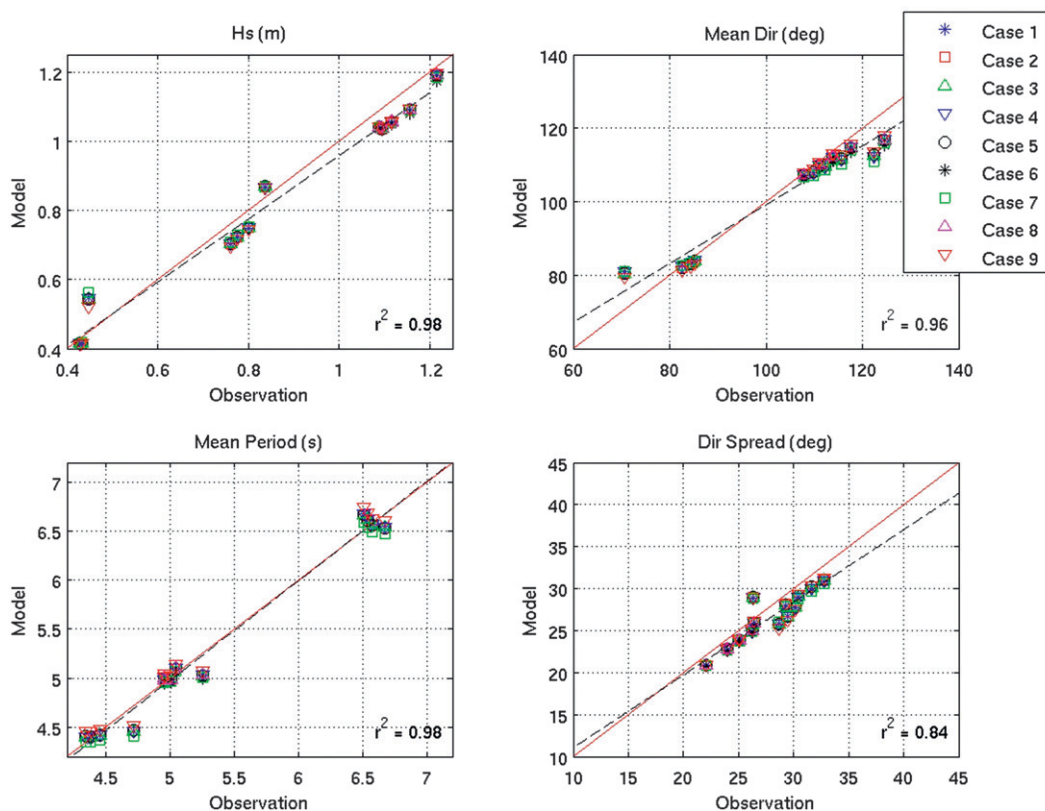

FIG. 10. As in Fig. 8, but for summary plots of modeled vs nonlinear SWAN-generated "observed" spectral statistics for all cases, showing results at interior and boundary data locations.

TABLE 3. As in Table 1, but for adjoint skill scores for tests at Duck, NC, using "same instrument" pseudodata observations generated by nonlinear forward SWAN in place of measured data at each location.

		Output location				
Case	Initializing location	FRF	AWAC (5 m)	AWAC (6 m)	AWAC (8 m)	AWAC (1 m)
1	FRF	0.82–0.91	0.93	0.88–0.91	0.89	0.76–0.86
2	AWACs	0.84–0.89	0.92	0.85–0.90	0.88	0.80–0.87
3	All	0.84–0.90	0.92	0.85–0.90	0.89	0.79–0.85
4	All Overwt FRF	0.84–0.89	0.93	0.84–0.91	0.89	0.77–0.86
5	All Overwt AWACs	0.84–0.89	0.92	0.85–0.90	0.88	0.80–0.87
6	AWAC (5 m)	0.84	0.90	0.87	0.84	0.80
7	AWAC (6 m)	0.78–0.88	0.92	0.82–0.90	0.87	0.71–0.83
8	AWAC (8 m)	0.89	0.93	0.91	0.89	0.86
9	AWAC (11)	0.83–0.91	0.93	0.79–0.92	0.91	0.83–0.91

and its subroutines are validated with standard dot-product and impulse tests, after which the adjoint is incorporated into an assimilation system together with a copy of the original forward SWAN code and modules for minimizing the cost function. The assimilation system is then tested in twin experiments with artificial observed spectra (generated by nonlinear SWAN) at Santa Rosa Island, Florida, and assimilations with actual measured spectra at Duck, North Carolina. In the twin experiments, the numerical adjoint performs comparably to the analytical adjoint. The mean skill level [Eq. (8)] for tests with artificial data is about 0.91, while the average skill in tests with measured data is 0.52. Differences in spectral resolution and quality at the different types of nearshore instruments in Duck play an important role in reducing skill levels for the latter tests. When skill scores at Duck are recomputed for 1D frequency spectra only, the average skill value increases to 0.71. Replacing measured spectra in the Duck tests with artificial "observations" from nonlinear SWAN improves the system's average 2D skill score to 0.82 and its average 1D skill score to 0.87. The suboptimal numerical adjoint system's performance is sufficient to justify the development of a complete adjoint to fully nonlinear, nonstationary SWAN, and such an effort is now underway.

APPENDIX A

Variational Data Assimilation

The following material is included for readers who are unfamiliar with representer-based data assimilation and follows the approach of Bennett (2002). Including errors in the forcing, the initial conditions, and the dynamics, the nonstationary, inhomogeneous SWAN model equation [Eq. (1)] may be written in one representative spatial dimension as

$$\frac{\partial N}{\partial t} + c\frac{\partial N}{\partial x} = F + f$$
$$u(x,0) = I(x) + i(x)$$
$$u(0,t) = B(t) + b(t), \qquad (A1)$$

where $N = N(x,t)$ and f; i and b represent the errors in the model dynamics, initial and boundary conditions respectively, with covariances $C_f(x,t,x',t')$, $C_i(x,x')$, and $C_b(t,t')$. Here, the idealized single spatial dimension x represents a combination of two spatial dimensions (X, Y) and two additional spectral dimensions (f, θ) from the "real world." It is assumed that, during the time interval $[0, T]$, observations of part of the state are collected, having the form

$$y_m = N(x_m, t_m) + \varepsilon_m, \quad 1 \leq m \leq M, \qquad (A2)$$

in which ε_m represents measurement error. These observations are collected at discrete points distributed in space and time. An observation operator [e.g., H_k in Eq. (3)] will not be utilized in this derivation in accordance with the preceding study, for which measurements are collocated with domain grid points and the observed quantity is the wave action N rather than another derived function of N. Because of various sources of error, the model trajectory generally deviates from the observations. The role of data assimilation is to minimize

the aggregate discrepancy over the given time interval. It is assumed that if the said discrepancy is minimized, then the optimal trajectory will yield a final state $N(x, T)$ for which a reliable forecast will be made. The aggregate model–data discrepancy is represented in a generalized cost function of the form

$$J = \int_0^T \int_0^L \int_0^T \int_0^L f(x,t)W_f(x,t,x',t')f(x',t')\,dx'\,dt'\,dx\,dt + \int_0^L \int_0^L i(x)W_i(x,x')i(x')\,dx'\,dx$$

$$+ \int_0^T \int_0^T b(t)W_b(t,t')b(t')\,dt'\,dt + \sum_{m=1}^M \sum_{n=1}^M [y_m - N(x_m,t_m)]R_{mn}^{-1}[y_n - N(x_n,t_n)], \tag{A3}$$

where W_f, W_i, and W_b are the respective inverses of C_f, C_i, and C_b defined in the convolution sense:

$$\int_0^T \int_0^L W_f(x,t,x',t')C_f(x',t',x'',t'')\,dx'\,dt' = \delta(x-x'')\delta(t-t'')$$

$$\int_0^L W_i(x,x')C_i(x',x'')\,dx' = \delta(x-x'')$$

$$\int_0^T W_b(t,t')C_b(t',t'')\,dt' = \delta(t-t'')$$

$$\tag{A4}$$

with the unit impulse $\delta(z) = 1$ for $z = 0$ but otherwise zero, and R_{mn} are the components of $\mathbf{R}$, the observation error covariance; that is, $\mathbf{R} = \overline{\varepsilon\varepsilon^T}$, with the overbar denoting the average in the ensemble sense. It is important to note that the cost function (A3) is quadratic in the errors; thus, J has one minimum. For the sake of clarity and without loss of generality, it is assumed that the errors are uncorrelated with constant variance. The cost function becomes

$$J = W_f \int_0^T \int_0^L f(x,t)^2\,dx\,dt + W_i \int_0^L i(x)^2\,dx + W_b \int_0^T b(t)^2\,dt$$

$$+ \omega \sum_{m=1}^M [y_m - N(x_m,t_m)]^2, \tag{A5}$$

where ω is the inverse of the data error variance [for the steady-state, homogeneous form of (A1), Eq. (A5) reduces to the form of section 2, Eq. (3)]. After some standard manipulation, in which we introduce the adjoint variable

$$\lambda = W_f \left(\frac{\partial N}{\partial t} + c\frac{\partial N}{\partial x} - F \right), \tag{A6}$$

we obtain

$$-\frac{\partial \lambda}{\partial t} - c\frac{\partial \lambda}{\partial x} + \omega \sum_{m=1}^M [N(x_m,t_m) - y_m]$$

$$\times \delta(x-x_m)\delta(t-t_m) = 0,$$

$$\lambda(L,t) = 0,$$

$$\lambda(x,T) = 0,$$

$$-c\lambda(0,t) + W_b[N(0,t) - B(t)] = 0, \quad \text{and}$$

$$-\lambda(x,0) + W_i[N(x,0) - I(x)] = 0. \tag{A7}$$

The equations in (A6) and (A7) constitute the Euler–Lagrange conditions for local extrema of the cost function.

While there are several techniques for obtaining a solution to (A6) and (A7), the present study makes use of the conjugate gradient method detailed in Walker (2006) and Polak and Ribière (1969), in combination with the method of representers. The optimal solution or "best estimate" N is obtained by iteratively solving the above equations, reorganized into the coupled system

$$(B) \begin{cases} -\dfrac{\partial \lambda}{\partial t} - c\dfrac{\partial \lambda}{\partial x} = -\omega \sum_{m=1}^M [N(x_m,t_m) - y_m]\delta(x-x_m)\delta(t-t_m) \\[2mm] \lambda(L,t) = 0 \\[1mm] \lambda(x,T) = 0 \end{cases}$$

$$\tag{A8}$$

$$(F) \begin{cases} \dfrac{\partial N}{\partial t} + c\dfrac{\partial N}{\partial x} = F + W_f^{-1}\lambda \\[2mm] N(x,0) = I + W_i^{-1}\lambda(x,0) \\[2mm] N(0,t) = B(t) + cW_b^{-1}\lambda(0,t) \end{cases} \tag{A9}$$

Note that we now also have the best estimate of the errors f, I, and b as the second terms in all three equations of (A9), respectively. The representer method uncouples the above system by introducing representer functions $r_m(x,t)$, $1 \le m \le M$. Each representer function has an adjoint $\alpha_m(x,t)$ that satisfies

$$(B_m) \begin{cases} -\dfrac{\partial \alpha_m}{\partial t} - c\dfrac{\partial \alpha_m}{\partial x} = \delta(x-x_m)\delta(t-t_m) \\[2mm] \alpha_m(L,t) = 0 \\[1mm] \alpha_m(x,T) = 0 \end{cases} \tag{A10}$$

$$(F_m) \begin{cases} \dfrac{\partial r_m}{\partial t} + c\dfrac{\partial r_m}{\partial x} = W_f^{-1}\alpha_m \\[2mm] r_m(x,0) = W_i^{-1}\alpha_m(x,0) \\[2mm] r_m(0,t) = cW_b^{-1}\alpha_m(0,t) \end{cases} \tag{A11}$$

Next, we seek the solution N of the form

$$N(x,t) = N_F(x,t) + \sum_{m=1}^{M} \beta_m r_m(x,t), \qquad \text{(A12)}$$

where the coefficients β_m are unknown constants and $N_F(x,t)$ is the prior estimate obtained as the solution of (A1) without the errors:

$$\frac{\partial N}{\partial t} + c \frac{\partial N}{\partial x} = F$$
$$N(x,0) = I(x)$$
$$N(0,t) = B(t). \qquad \text{(A13)}$$

We combine (A11), (A12), and (A13) to get

$$D(N) = D(N_F) + \sum_{m=1}^{M} \beta_m D[r_m(x,t)]$$
$$= F + W_f^{-1} \sum_{m=1}^{M} \beta_m \alpha_m(x,t), \qquad \text{(A14)}$$

in which $D = (\partial/\partial t) + c(\partial/\partial x)$. Now by definition of λ and by virtue of (A14),

$$\lambda = W_f[D(N) - F] = \sum_{m=1}^{M} \beta_m \alpha_m(x,t). \qquad \text{(A15)}$$

Applying the differential operator D on (A15) and using the first equation of (A8),

$$-D(\lambda) = -\sum_{m=1}^{M} \beta_m D(\alpha_m) = \sum_{m=1}^{M} \beta_m \delta(x - x_m)\delta(t - t_m)$$
$$= -\omega \sum_{m=1}^{M} [N(x_m, t_m) - y_m]\delta(x - x_m)\delta(t - t_m). \qquad \text{(A16)}$$

Equating the coefficients of β_m in (A16) yields the optimal choice of the representer coefficients:

$$\beta_m = -\omega[N(x_m, t_m) - y_m], \quad 1 \le m \le M. \qquad \text{(A17)}$$

The optimal representer coefficients still depend on the optimal solution N, so further manipulation is required. Substituting N in (A12) into (A17) gives

$$\beta_m = -\omega \left[N_F(x_m, t_m) + \sum_{l=1}^{M} \beta_l r_l(x_m, t_m) - y_m \right]$$
$$= -\omega \left(N_{Fm} + \sum_{l=1}^{M} \beta_l r_{l,m} - y_m \right). \qquad \text{(A18)}$$

Thus,

$$\sum_{l=1}^{M} (r_{l,m} + \omega^{-1}\delta_{l,m})\beta_l = h_m = y_m - N_{Fm}, \qquad \text{(A19)}$$

where $\delta_{l,m}$ is the Kronecker delta. In matrix notation, the M equations in (A19) for the representer coefficients β_m are

$$(\mathbf{R} + \omega^{-1}\mathbf{I})\boldsymbol{\beta} = \mathbf{h} = \mathbf{y} - \mathbf{N}_F. \qquad \text{(A20)}$$

The optimal solution is therefore obtained as

$$N(x,t) = N_F(x,t) + (\mathbf{y} - \mathbf{N}_F)^{\mathrm{T}}(\mathbf{R} + \omega^{-1}\mathbf{I})^{-1} r(x,t). \qquad \text{(A21)}$$

APPENDIX B

Discretized Equations

In expanded form, the stationary homogeneous wave action equation [Eq. (2)] may be written as

$$\frac{\partial C_x N}{\partial x} + \frac{\partial C_y N}{\partial y} + \frac{\partial C_\sigma N}{\partial \sigma} + \frac{\partial C_\theta N}{\partial \theta} = 0, \qquad \text{(B1)}$$

where all terms have been defined in section 2. Discretized using a first-order backward difference scheme for spatial dimensions and a central difference for spectral dimensions, this equation becomes

$$\frac{1}{\Delta x}(C_{x_i} N_{i,j,k,l} - C_{x_{i-1}} N_{i-1,j,k,l}) + \frac{1}{\Delta y}(C_{y_j} N_{i,j,k,l} - C_{y_{j-1}} N_{i,j-1,k,l})$$
$$+ \frac{1}{2\Delta\sigma}[(1-\nu)C_{\sigma_{k+1}} N_{i,j,k+1,l} + 2\nu C_{\sigma_k} N_{i,j,k,l} - (1+\nu)C_{\sigma_{k-1}} N_{i,j,k-1,l}]$$
$$+ \frac{1}{2\Delta\theta}[(1-\eta)C_{\theta_{l+1}} N_{i,j,k,l+1} + 2\eta C_{\theta_l} N_{i,j,k,l} - (1+\eta)C_{\theta_{l-1}} N_{i,j,k,l-1}] = 0, \qquad \text{(B2)}$$

where indices i, j, k, and l track the spatiospectral grid location in terms of coordinates x, y, σ, and θ, respectively. Here, the frequency and directional derivatives are expressed as weighted central differences with weights ν and η between zero and one. Taking the adjoint of (B2) gives

$$\frac{1}{\Delta x} C_{x_i}(\lambda_{i,j,k,l} - \lambda_{i+1,j,k,l}) + \frac{1}{\Delta y} C_{y_j}(\lambda_{i,j,k,l} - \lambda_{i,j+1,k,l}) + \frac{1}{2\Delta\sigma} C_{\sigma_k}[(1-\nu)\lambda_{i,j,k-1,l} + 2\nu\lambda_{i,j,k,l} - (1+\nu)\lambda_{i,j,k+1,l}]$$

$$+ \frac{1}{2\Delta\theta} C_{\theta_l}[(1-\eta)\lambda_{i,j,k,l-1} + 2\eta\lambda_{i,j,k,l} - (1+\eta)\lambda_{i,j,k,l+1}] = (y_m - N_m)\delta_{i_m}^i \delta_{j_m}^j \delta_{k_m}^k \delta_{l_m}^l, \tag{B3}$$

where $\lambda_{i,j,k,l}$ is defined as the residual variable or innovation (the difference between model estimate and observation) at the spatiospectral location given by indices i, j, k, and l, and the observation locations are denoted by m, where $1 \le m \le M$. Each δ is a Kroenecker delta for a given spatiospectral index, which is equal to one only at the index of an observation location and is zero otherwise. Note that in (B3), the energy propagation speeds $C_{x,y,\sigma,\theta}$ are assumed to be locally constant with respect to coordinates x, y, σ, and θ, which leads to a different form for the coefficients in (B3) compared to those in (B2). The adjoint to linear homogeneous SWAN essentially solves (B3), forced at selected interior grid observation locations $[x(i_m)$, $y(j_m)]$ with innovations $(y_m - N_m)$, each of which includes observed-minus-estimated spectral density for frequencies $\sigma(k_m)$ and directions $\theta(l_m)$.

REFERENCES

Aarninkhof, S. G. J., B. G. Ruessink, and J. A. Roelvink, 2005: Nearshore subtidal bathymetry from time-exposure video images. *J. Geophys. Res.,* **110,** C06011, doi:10.1029/2004JC002791.

Aouf, L., J.-M. Lefevre, and D. Hauser, 2006: Assimilation of directional wave spectra in the wave model WAM: An impact study from synthetic observations in preparation for the SWIMSAT satellite mission. *J. Atmos. Oceanic Technol.,* **23,** 448–463.

Bauer, E., S. Hasselmann, K. Hasselmann, and H. Graber, 1992: Validation and assimilation of Seasat altimeter wave heights using the WAM wave model. *J. Geophys. Res.,* **97** (C8), 12 671–12 682.

Bennett, A. F., 2002: *Inverse Modeling of the Ocean and Atmosphere.* Cambridge University Press, 234 pp.

Booij, N., R. C. Ris, and L. H. Holthuijsen, 1999: A third-generation wave model for coastal regions: 1. Model description and validation. *J. Geophys. Res.,* **104** (C4), 7649–7666.

Bücker, H. M., J. Willkomm, S. Gross, and O. Fortmeier, 2011: Discrete and continuous adjoint approaches to estimate boundary heat fluxes in falling films. *Optim. Methods Software,* **26,** 105–125.

Chickadel, C. C., R. A. Holman, and M. H. Freilich, 2003: An optical technique for the measurement of longshore currents. *J. Geophys. Res.,* **108,** 3364, doi:10.1029/2003JC001774.

Edwards, K. L., J. Veeramony, D. Wang, K. T. Holland, and Y. L. Hsu, 2009: Sensitivity of Delft3D to input conditions. *Proc. OCEANS '09 Conf.,* Biloxi, MS, MTS/IEEE, 8 pp. [Available online at http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5422127&isnumber=5422059.]

Erwig, M., Z. Fu, and B. Pflaum, 2007: Parametric FORTRAN: Program generation in scientific computing. *J. Software Maint. Evol. Res. Pract.,* **19,** 155–182.

Favennec, Y., 2005: Comparison between the discretized continuous gradient and the discrete gradient when dealing with nonlinear parabolic problems. *Numer. Heat Transfer,* **48B,** 345–362.

Giresse, R., and A. Walther, 2004: Evaluating gradients in optimal control: Continuous adjoints versus automatic differentiation. *J. Optim. Theory Appl.,* **122,** 63–86.

Hasselmann, S., P. Lionello, and K. Hasselmann, 1997: An optimal interpolation scheme for the assimilation of spectral wave data. *J. Geophys. Res.,* **102** (C7), 15 823–15 836.

Järvinen, H., 1998: Observations and diagnostic tools for data assimilation. Meteorological Training Course Lecture Series, ECMWF, 19 pp. [Available online at http://www.ecmwf.int/newsevents/training/rcourse_notes/pdf_files/Obs_and_diag_tools.pdf.]

Kalnay, E., 2003: *Atmospheric Modeling, Data Assimilation, and Predictability.* Cambridge University Press, 341 pp.

Keen, T. R., W. E. Rogers, J. Dykes, J. M. Kaihatu, and K. T. Holland, 2007: Determining heterogeneous bottom friction distributions using a numerical wave model. *J. Geophys. Res.,* **112,** C08008, doi:10.1029/2005JC003309.

Lahoz, W., B. Khattatov, and R. Ménard, Eds., 2010: *Data Assimilation: Making Sense of Observations.* Springer, 718 pp.

Lippmann, T. C., and R. A. Holman, 1990: The spatial and temporal variability of sand bar morphology. *J. Geophys. Res.,* **95** (C7), 11 575–11 590.

Moore, A. M., H. G. Arango, E. Di Lorenzo, B. D. Cornuelle, A. J. Miller, and D. J. Neilson, 2004: A comprehensive ocean prediction and analysis system based on the tangent linear and adjoint of a regional ocean model. *Ocean Modell.,* **7,** 227–258.

Narayanan, C., V. N. Rama Rao, and J. M. Kaihatu, 2004: Model parameterization and experimental design issues in nearshore bathymetry inversion. *J. Geophys. Res.,* **109,** C08006, doi:10.1029/2002JC001756.

Ngodock, H., and M. Carrier, 2013: A weak constraint 4D-Var assimilation system for the navy coastal ocean model using the representer method. *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications,* S. K. Park and L Xu, Eds., Vol. 2, Springer, 367–390.

Park, S. K., and L. Xu, Eds., 2009: *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications.* Springer, 476 pp.

Polak, E., and R. F. Ribière, 1969: Note sur la convergence de méthodes de directions conjugées. *Rev. Fr. Inf. Rech. Oper.,* **16,** 35–43.

Portilla, J., 2009: Buoy data assimilation in nearshore wave modeling. Ph.D. thesis, Katholieke Universiteit Leuven, 209 pp.

SWAN Team, 2011: SWAN cycle III version 40.85. Scientific and Technical Documentation, Delft University of Technology, Interpolation of Spectra (Section 3.14), 65–66.

Swinbank, R., V. Shutyaev, and W. Lahoz, Eds., 2003: *Data Assimilation for the Earth System.* NATO Science Series, Vol. IV, Springer, 390 pp.

Tolman, H. L., 2009: User manual and system documentation of WAVEWATCH III version 3.14. NOAA/NWS/NCEP/MMAB Tech. Note 276, 194 pp. + appendixes.

USACE Field Research Facility, cited 2011: Wind wave directional spectra. [Available online at http://140.194.28.33/8mArray/fdspec.html.]

Van Dongeren, A., N. Plant, A. Cohen, D. Roelvink, M. C. Haller, and P. Catalan, 2008: Beach Wizard: Nearshore bathymetry estimation through assimilation of model computations and remote observations. *Coastal Eng.,* **55,** 1016–1027.

Veeramony, J., D. Walker, and L. Hsu, 2010: A variational data assimilation system for nearshore applications of SWAN. *Ocean Modell.,* **35,** 206–214.

Walker, D. T., 2006: Assimilation of SAR imagery in a nearshore spectral wave model. GDAIS Tech. Rep. 200236-F, 33 pp.

WAMDI Group, 1988: The WAM model—A third generation ocean wave prediction model. *J. Phys. Oceanogr.,* **18,** 1775–1810.

Weaver, A. T., J. Vialard, and D. L. T. Anderson, 2003: Three- and four-dimensional variational assimilation with a general circulation model of the tropical Pacific Ocean. Part I: Formulation, internal diagnostics, and consistency checks. *Mon. Wea. Rev.,* **131,** 1360–1378.