Portland State University
PDXScholar

Civil and Environmental Engineering Faculty
Publications and Presentations

Civil and Environmental Engineering

9-1-2008

The Inverse Ocean Modeling System. Part II: Applications

J. C. Muccino
Portland State University

H. G. Arango
State University of New Jersey

Andrew F. Bennett
Oregon State University

B. S. Chua
Oregon State University

Julia Muccino Cornuelle
University of California - San Diego

See next page for additional authors

Let us know how access to this document benefits you.

Follow this and additional works at: http://pdxscholar.library.pdx.edu/cengin_fac

Part of the Civil and Environmental Engineering Commons

Citation Details

Muccino, J. C., Arango, H. G., Bennett, A. F., Chua, B. S., Cornuelle, B. D., Di Lorenzo, E. E., & ... Zaron, E. D. (2008). The Inverse Ocean Modeling System. Part II: Applications. Journal Of Atmospheric & Oceanic Technology, 25(9), 1623-1637.

Authors

J. C. Muccino, H. G. Arango, Andrew F. Bennett, B. S. Chua, Julia Muccino Cornuelle, E. Di Lorenzo, Gary D. Egbert, D. Haidvogel, J. C. Levin, H. Luo, A. J. Miller, A. M. Moore, and Edward D. Zaron

The Inverse Ocean Modeling System. Part II: Applications

J. C. Muccino,* H. G. Arango,[+] A. F. Bennett,[#] B. S. Chua,[#] B. D. Cornuelle,[@]
E. Di Lorenzo,[&] G. D. Egbert,[#] D. Haidvogel,[+] J. C. Levin,[+] H. Luo,* A. J. Miller,[@]
A. M. Moore,** AND E. D. Zaron[++]

*Department of Civil and Environmental Engineering, Arizona State University, Tempe, Arizona
[+]Institute of Marine and Coastal Sciences, Rutgers, The State University of New Jersey, New Brunswick, New Jersey
[#]College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, Oregon
[@]Scripps Institution of Oceanography, University of California, San Diego, La Jolla, California
[&]School of Earth and Atmospheric Science, Georgia Institute of Technology, Atlanta, Georgia
**Department of Ocean Sciences, University of California, Santa Cruz, Santa Cruz, California
[++]College of Civil and Environmental Engineering, Portland State University, Portland, Oregon

ABSTRACT

The Inverse Ocean Modeling (IOM) System is a modular system for constructing and running weak-constraint four-dimensional variational data assimilation (W4DVAR) for any linear or nonlinear functionally smooth dynamical model and observing array. The IOM has been applied to four ocean models with widely varying characteristics. The Primitive Equations Z-coordinate-Harmonic Analysis of Tides (PEZ-HAT) and the Regional Ocean Modeling System (ROMS) are three-dimensional, primitive equations models while the Advanced Circulation model in 2D (ADCIRC-2D) and Spectral Element Ocean Model in 2D (SEOM-2D) are shallow-water models belonging to the general finite-element family. These models, in conjunction with the IOM, have been used to investigate a wide variety of scientific phenomena including tidal, mesoscale, and wind-driven circulation. In all cases, the assimilation of data using the IOM provides a better estimate of the ocean state than the model alone.

1. Introduction

The Inverse Ocean Modeling (IOM) system is a modular system for constructing and running weak-constraint, four-dimensional variational data assimilation (W4DVAR) for any linear or nonlinear functionally smooth dynamical model and observing array. Details of the IOM are described in a companion paper (Bennett et al. 2008) and only briefly summarized here. The objective of this paper is to demonstrate the flexibility, power, and usefulness of the IOM.

Implementation of the IOM in conjunction with four ocean models is described: Primitive Equation Z-coordinate-Harmonic Analysis of Tides (PEZ-HAT), Regional Ocean Modeling System (ROMS), Advanced Circulation in 2D (ADCIRC-2D), and Spectral Element Ocean Model in 2D (SEOM-2D). These models are diverse in their character, formulation and scientific applications. Two of the models (PEZ-HAT and

SEOM-2D) are new models; they have been developed simultaneously with the IOM and thus their structure was readily matched to the IOM interface. The other two models (ROMS and ADCIRC-2D) have been in widespread use long before the IOM came into existence. The IOM is flexible enough to interface with these models with only minor alterations to their structure.

The models also differ in their numerics, as detailed in Table 1. Both PEZ-HAT and ROMS are three-dimensional, primitive equation, finite-difference models. The numerics in PEZ-HAT are based on the 1991 Killworth free-surface version (Killworth et al. 1991) of the Bryan and Cox (1968) model, as recoded by Bennett and Chua (A. Bennett 2003, personal communication), and include modifications for extensive run-time (rather than compile time) model configuration, more accurate baroclinic–barotropic time stepping (utilizing a subset of the developments described in Shchepetkin and McWilliams 2005), partial cell bottom topography (Pacanowski and Gnanadesikan 1998), harmonic analysis, and open boundary conditions. ROMS is a free-surface, split-explicit, terrain-following primitive equa-

Corresponding author address: J. Muccino, Department of Civil and Environmental Engineering, P.O. Box 875306, Arizona State University, Tempe, AZ 85287-5306.
E-mail: jmuccino@asu.edu

TABLE 1. Numerical models.

Model attributes	PEZ-HAT	ROMS	ADCIRC-2D	SEOM-2D
Dynamical equations	Primitive equations hydrostatic, free surface	Primitive equations hydrostatic, free surface	Shallow water (wave equation)	Shallow water
Spatial dimensions	3D	3D	2D[a]	2D[b]
Adjoint accuracy	Approximate	Exact	Approximate (discrete in space, continuous in time)	Exact
Adjoint metric	Riemann	Euclidean	Riemann	Euclidean
Mesh	Structured, spherical	Structured, curvilinear, nested/composed	Unstructured triangles	Unstructured quadrilaterals
Equation of state	Simplified	Full UNESCO	—	—
Horizontal discretization	Finite differences, 2nd order	Finite differences, 2nd order	Finite elements, linear	Spectral elements, any order
Vertical discretization	Finite volume/finite difference, geopotential levels	Finite volume–finite difference, terrain-following levels	—	—
Bathymetry treatment	Unsmoothed	Smoothed	Unsmoothed	Unsmoothed[c]
Time stepping	2nd order, split-explicit barotropic and baroclinic modes	3rd order, predictor–corrector, split–explicit barotropic and baroclinic modes	Crank Nicolson, w/3 time levels for surface gradient	3rd order Adams-Bashford
Horizontal mixing	Smagorinsky, deviatoric stress tensor, geopotential	Deviatoric stress tensor ($\mathbf{u}$) Geopotential–isopycnic (T, S)	Eddy viscosity	Eddy viscosity, spectral filter
Vertical mixing	KPP	KPP, MY-2.5, GLS	—	—
Surface conditions	Wind stress, heat fluxes	Wind tress, heat fluxes, air–sea BL model, COARE	Wind stress, atmospheric pressure	Wind stress
Bottom parameterization	Linear–quadratic Drag, partial cells	Linear–quadratic–log drag, BBL models	Linear–quadratic–hybrid drag	Linear drag
Coupled models	No	Biology, sediment, sea ice	NCEP, COAMPS, MM5, WindGen	No
User community	No	Yes, over 800 users worldwide	Yes, over 300 users worldwide	Yes
Web site	http://ocean.cee.pdx.edu/ ezaron/ PEZHAT.html	http: //www.myroms.org	http://www.adcirc.org	http://marine. rutgers.edu/seom
Architecture	Parallel–MPI	Coarse grain parallelization (OpenMP, MPI)	Parallel–MPI with Metis Domain decomposition	Parallel–MPI
Interface–configuration	CPP, Makefile, XML, OpenDAP	F90–F95, CPP, Makefile	Fortran preprocessing, OpenDAP, SMS-GUI	CPP, Makefile

[a] Forward ADCIRC available in 3D.
[b] Forward SEOM available in 3D.
[c] Smoothing required in 3D.

tions ocean model with several vertical mixing schemes, multiple levels of nesting, and composite grids and other advanced numerical features (Shchepetkin and McWilliams 2003, 2005).

Both ADCIRC-2D and SEOM-2D are shallow-water models belonging to the finite-element family. The use of unstructured grids permits a better geometric description of complex coastlines, accommodates spatially

variable resolution for localized small-scale processes, and therefore allows multiscale simulations within the framework of a single model. ADCIRC-2D uses linear interpolation on triangular elements and a wave equation formulation of the shallow-water equations. SEOM-2D uses higher-order interpolation on quadrilateral elements in conjunction with the standard formulation of the shallow-water equations. It offers dual approaches to achieving convergence: algebraic (elemental grid refinement) and exponential (increase in the order of intraelement interpolation).

a. The scientific applications of the models

The four models have a wide range of current and potential uses, including realistic simulation and prediction, as well as idealized process studies.

PEZ-HAT was developed to quantify and understand the impact of the tides on mixing in the abyssal ocean. It has been used to model the internal tides around the Hawaiian Ridge (Zaron and Egbert 2007) and is currently being used in conjunction with the IOM to globally map the energy flux out of the tidal internal wave field using observations from satellite altimetry.

ROMS is used for a wide range of applications from the basin to coastal and estuarine scales, and includes several coupled models for biogeochemical, bio-optical, sediment, and sea ice applications. W4DVAR studies using ROMS are currently being used for predictability studies on the East and West Coasts (United States), the Gulf of Mexico, the Caribbean Sea, and the east Australia Current. A detailed description of the W4DVAR version of ROMS appears in Di Lorenzo et al. (2007).

ADCIRC-2D is used to model tidal and wind-driven circulation, to forecast hurricane storm surge and flooding, to study larval transport, material disposal, and dredging feasibility. Recently, ADCIRC-2D has been used to successfully forecast the flooding of New Orleans during Hurricane Katrina (Travis 2005; Duffy 2005) and the Federal Emergency Management Agency is using ADCIRC-2D to regenerate its Flood Insurance Rate Maps (FIRMs).

SEOM-2D has been used for multiscale modeling of tidal and wind-driven flows (Wunsch et al. 1997; Levin et al. 1997); and for modeling of abyssal flows (Curchitser et al. 1998, 1999). SEOM-2D has been also used for highly accurate process studies (Haidvogel and Beckmann 1999).

b. The hypotheses

The IOM allows all dynamical and observational constraints to be weak, tacitly assuming that errors are normally distributed. Thus, the first and second moments of the error fields must be hypothesized. While these errors may be jointly normally distributed, the IOM interface permits the specification and parameterization of the model forcing error variance within a class of spatially and temporally correlated error models. By default, the errors are assumed to be separable in space and time, with spatial covariances having a "bell-shaped" form:

$$C(x, x') \propto \exp\left(\frac{-|x - x'|^2}{X^2}\right), \qquad (1)$$

where X is the decorrelation length scale. Weaver and Courtier (2001) offer more complex spatial covariances; this will be included in future releases of the IOM. Temporal covariances have a Markovian form:

$$C(t, t') \propto \exp\left(\frac{-|t - t'|}{s}\right), \qquad (2)$$

where s is the decorrelation time scale. In general, the IOM estimate of the dynamical error, λ, is related to the corresponding adjoint variable, $\hat{\lambda}$, by convolution with the appropriate error covariance, C:

$$\lambda = \int C \circ \hat{\lambda} \, dt = \iint C\lambda \, dx \, dt, \qquad (3)$$

where the "∘" indicates integration over the spatial domain. The convolutions in (3) can be extremely expensive to compute directly, but the IOM offers efficient alternatives (Chua and Bennett 2001). The IOM also allows the user to implement other error covariances through a user-supplied covariance kernel. For example, both the ROMS and PEZ-HAT applications described below have implemented non-bell-shaped, spatially inhomogeneous covariance functions along the lines of Weaver and Courtier (2001) and Purser et al. (2003), respectively.

Each model will be briefly described in the following sections; the reader is referred to other publications for more information. It is, of necessity, assumed that the reader is broadly familiar with variational assimilation and with the definition of an adjoint model (e.g., Bennett 2002). For each model, the governing equations, numerical formulation, penalty, and practical details of the IOM implementation will be described. In addition, some representative results obtained with the inversion of each model will be discussed.

2. PEZ-HAT

PEZ-HAT estimates the state of the ocean as a field over space at selected tidal frequencies. The forward model solves for these in two stages: a forward integration of the time-dependent governing equations from

some initial condition and then harmonic analysis of those fields to yield the complex tidal amplitudes. The IOM can accept both time and frequency domain models; the domain of the fields is no limitation.

a. PEZ-HAT: The governing equations

The governing equations are the primitive equations and conservation of potential temperature together with an equation of state. In these open-ocean tidal applications, the primitive equations are expressed in terms of complex tidal amplitudes and linearized around a motionless and laterally homogeneous background state. Kinematic boundary conditions are enforced at the material surface of the fluid and at the bottom. Turbulent mixing of momentum and the active tracer, potential temperature, are parameterized with down-gradient diffusion. Open boundary conditions are posed separately for the baroclinic and barotropic components of flow, as described by Zaron and Egbert (2006).

b. PEZ-HAT: Defining a penalty

The integral penalty function may be written as

$$J(\mathbf{u}, w, p, \rho, \theta) = \hat{\lambda}_M \circ \lambda_M + \hat{\lambda}_\theta \circ \lambda_\theta$$
$$+ [\mathcal{L}(\mathbf{u}) - \mathbf{d}]^T \mathbf{V}^{-1} [\mathcal{L}(\mathbf{u}) - \mathbf{d}], \quad (4)$$

where $\mathbf{u}$ is the horizontal velocity, w is the vertical velocity, p is the pressure, ρ is the density, θ is the potential temperature, λ_M and λ_θ are residuals in the momentum and the potential temperature equations, respectively, $\mathcal{L}$ is the vector of measurement operators that are defined by integration kernels, $\mathbf{d}$ is the vector of scalar observations, $\mathbf{V}$ is a diagonal data error covariance matrix, and superscript T indicates Hermitian transpose. The adjoint variable $\hat{\lambda}$ corresponds to λ; the subscripts on these variables indicate the dynamical component. The domain of integration extends over the three-dimensional volume of the domain, and, in this context, the operator $\circ$ takes the complex conjugate of its first argument.

The adjoint code for PEZ-HAT utilizes both exact adjoints (i.e., finite-difference operators that are adjoint with respect to the Riemann sum approximation to the appropriate inner product) and approximate adjoints (i.e., finite-difference operators that converge to the exact operators in the limit of increasing model resolution). Experiments have shown that at the high resolution used in PEZ-HAT applications, the error in the representer functions computed using the approximate adjoints is comparable to the truncation error in the time-stepping/harmonic analysis, and it may be neglected.

c. PEZ-HAT: Details of IOM implementation

Several decisions were made during the development of PEZ-HAT that were influenced by its intended role as an IOM client. The architecture of the IOM dictates that each component for the variational assimilation be implemented as a separate command-line executable command, so PEZ-HAT was built this way. The IOM utilizes Network Common Data Form (NetCDF) (Davis and Rew 1990) as its native file type, and this was utilized by PEZ-HAT. The default IOM implementation of (1) and (3) involves time stepping a pseudodiffusion equation, which does not scale well as the resolution increases, so a custom routine (a command-line executable) was written for this purpose.

The "-meas" and "-comb" generation features of the IOM (Bennett et al. 2007) are tied to a specific design for the representation of measurement functionals; PEZ-HAT already had a similar infrastructure for model–data comparisons, and it was simple to adapt this to the separate executable model. Other features of the IOM (e.g., the solvers to obtain the optimal solution, the conditioning, and the error analysis algorithms) were used directly through the separate executable architecture.

There are some trade offs in the IOM design. The separate executable architecture may prove a limitation for modeling systems that are severely Input/Output (I/O) bound. For PEZ-HAT this was not an issue, since the computation time is much larger than the I/O time on current hardware. Another issue is the complexity of managing the "IOM runtime." When a user executes a data-assimilation task via the IOM, the controlling process is a perl script that runs sequentially on a master computer. For PEZ-HAT, the separate executable modules consist of shell scripts that submit the appropriate programs for execution on a variety of hardware architectures, including clusters of parallel processing computers. Managing the individual IOM runtime configurations among different hardware and job submission architectures proved a challenging task; however, this issue would likely have been even more problematic with the custom code. The IOM has the great benefit of declaring a relatively simple interface, and letting the user adopt as much or as little as necessary to accomplish their task. The nature of the interface does not change among underlying hardware platforms, only the detailed implementations of the separate executables.

During times of active code development, the PEZ-HAT code base expands at a rate of around 1000 non-comment lines per month. Furthermore, periods of significant code generation are always followed by periods
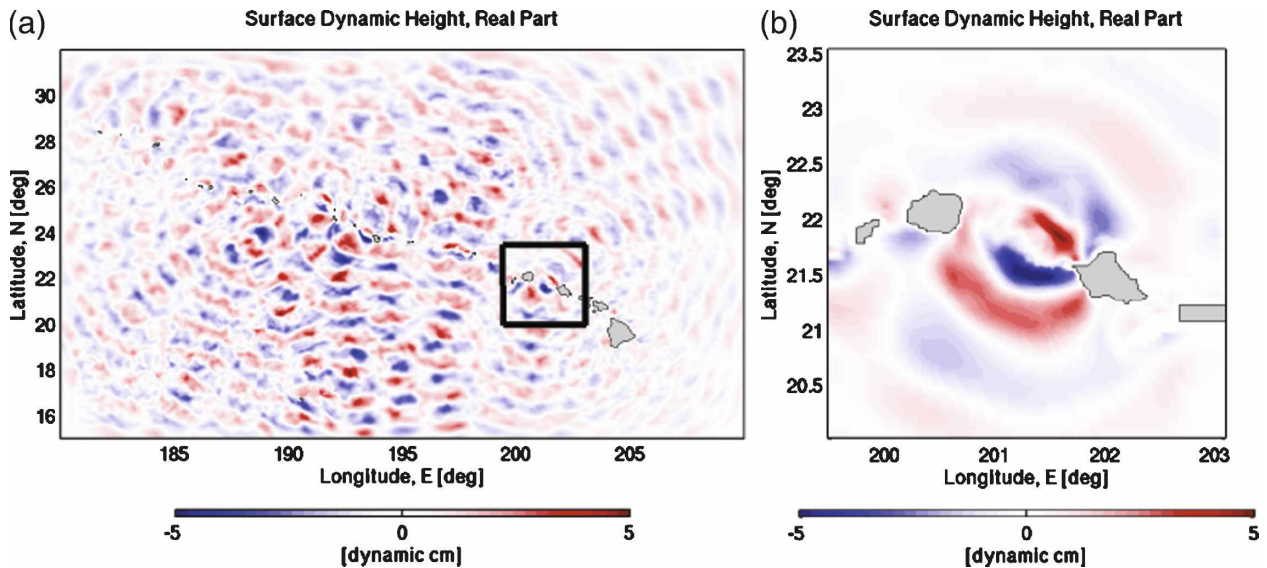
FIG. 1. Model domains: (a) geographic context for (b) the smaller region. Surface dynamic topography (essentially, vertically integrated specific volume anomaly) from the prior model shows the internal tide generated by the Hawaiian Ridge when the barotropic (surface) tide impinges on topographic features. The data assimilation discussed in the text uses the domain in (b), centered on the Kauai Channel, between the islands of Oahu and Kauai.

of code changes, keeping the number of lines relatively constant (i.e., debugging). The IOM functionality used by PEZ-HAT (which excludes most of the parametric FORTRAN program generation features) is encoded in roughly 23 000 lines of FORTRAN 90, which easily represents a savings of 2+ yr of PEZ-HAT developer time. In fact, interfacing and testing PEZ-HAT with the IOM was accomplished in approximately 2 months.

d. PEZ-HAT: Assimilation of HF radar data at Kaena Ridge, Hawaii

Surface velocity data obtained from a high-frequency phased-array radar system (as described in C. Chavanne et al. 2007, unpublished manuscript) have been assimilated to infer the unobserved subsurface fields and investigate mesoscale tidal interactions. Two radars were deployed on the west side of the island of Oahu, Hawaii, from September 2002 to May 2003 overlooking the Kauai Channel (C. Chavanne et al. 2006, unpublished manuscript). This site is unique among the many that have been observed with coastal radars because of the large range of bottom depth (from 0 to nearly 5000 m) and strong cross-isobath tidal flows that occur in the observation area. Together, these characteristics cause Kauai Channel to be a site of intense barotropic to baroclinic kinetic energy conversion, with conversion rates in excess of 1 W m^{-2}. Figure 1b shows the domain of PEZ-HAT, which is nested within a larger domain for context (Fig. 1a). The resolution of the numerical grid is 2 km in the central portion of the domain,

stretching to 10 km in the region within 1° of the domain boundary. There are 60 vertical levels, logarithmically stretched from 30 m at the ocean surface to 500 m at depth.

Each radar measures the ocean surface velocity component in the direction of the antenna array within an area of roughly 100 km in diameter, as shown in Fig. 2.
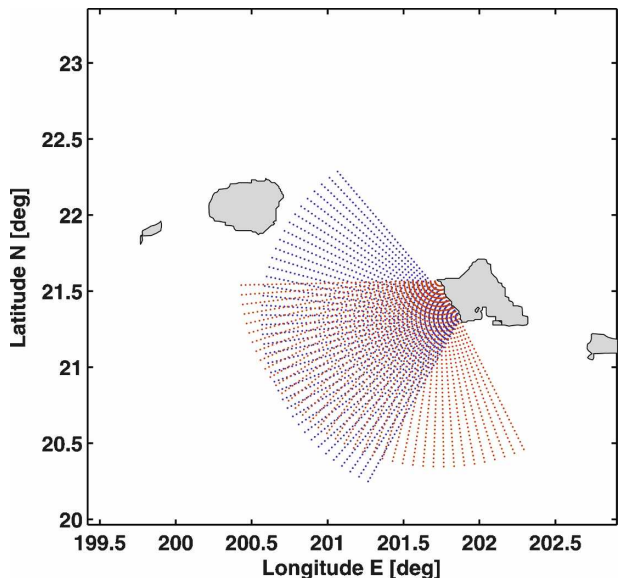

FIG. 2. HF-radar measurement sites. Radial velocities used in the assimilation are shown. Red: Kaena Point antenna array; blue: Ko'Olina antenna array.
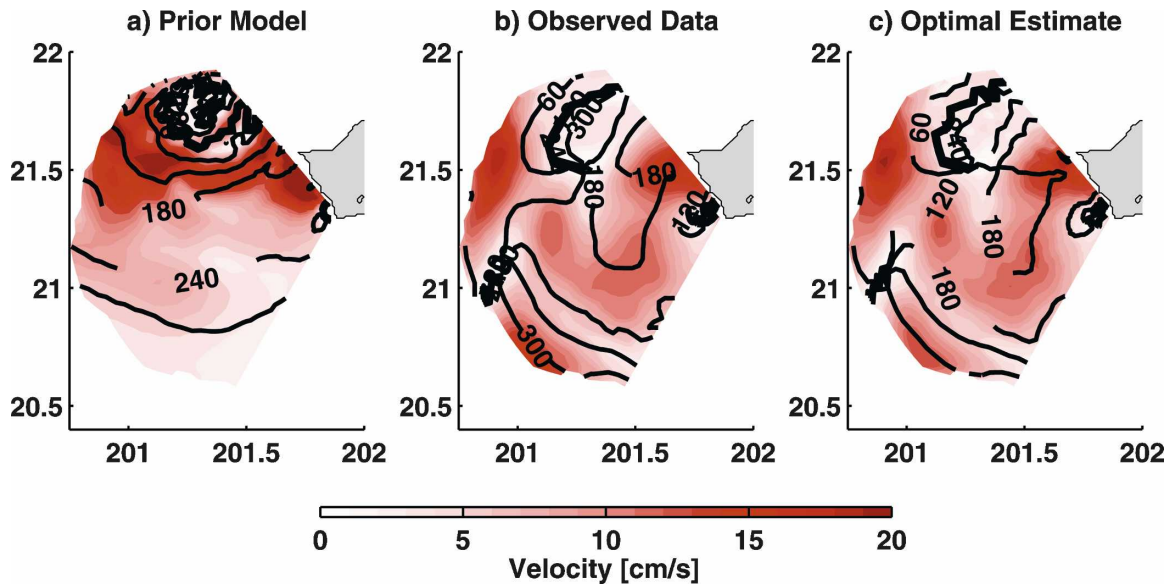
FIG. 3. Amplitude and phase of radial velocities, Ko'Olina. (a) Prior model predicts larger velocities and more uniform southward phase propagation than (b) the observed data. (c) Optimal estimate is in better agreement with the data. (Contour units: degree of phase.)

The spatial resolution of the data (approximately 2 km) is shown by the red dots (Kaena Point antenna array) and the blue dots (Ko'Olina antenna array) in the figure. Signals in the radar data represent a range of phenomena at different time and space scales, but in this application they have been harmonically analyzed and only the amplitude and phase of the dominant M_2 semidiurnal tidal component is assimilated. The nominal measurement error is between 1 and 5 cm s^{-1}, varying spatially, which is estimated from the residual in the harmonic analysis of radial velocity time series.

The measurement functional associated with each datum consists of a weighted average of the velocity components of the harmonically analyzed velocity field at the ocean surface. The coefficients in the weighted average are obtained by projecting the velocity at the model grid nodes into the local radial direction, and bilinear interpolation is used to compute the measurement at the actual measurement site (which normally falls between computational grid nodes).

To conduct the inversion, it is necessary to specify the second moments of the model forcing errors. It is assumed that the error is dominated by neglected advection and convection terms in the momentum and tracer conservation equations, respectively, and that the tidal fields are passively advected by an unknown mesoscale flow field that decays downward from the ocean surface like N, the background buoyancy frequency, that is correlated in scales of 50 km and 250 m in the horizontal and vertical, respectively. For simplicity, the errors in the momentum and tracer equations are assumed to be

uncorrelated. The spatial correlations are parameterized by correlation length scales and spectral roll-off rates and implemented using dimensional splitting with tridiagonal operators (Zaron 2006).

The indirect representer method, as implemented by the IOM, is used because of the large number of observations (5442) and the computational demands of PEZ-HAT. The optimal solution is computed with just 30 adjoint/forward model integrations.

Figures 3a–c compare the radial velocity from the prior model solution, the observed data, and the optimal solution, respectively, for the Ko'Olina antenna array. The prior solution has larger amplitude and sharper gradients than the observed data. The southward phase propagation and location of maximum amplitude are somewhat different as well, and the quantitative agreement between the prior solution and data is not good. In contrast, there is excellent agreement between the optimal estimate and the data, with the discrepancy being about 2 times the assumed measurement error. The results for the Kaena Point array are similar (not shown).

The generally weaker gradients in the data, as compared with the prior, point to deficiencies in PEZ-HAT, specifically the neglect of nonlinearity. A detailed study of the optimal model forcing (which is an estimate of these nonlinearities) suggests that it is compatible with the assumed level of mesoscale–tidal interactions. Mesoscale and tidal self-interactions do influence the time-averaged barotropic tidal conversion, and partitioning the energy into locally dissipated and

propagating components remains the subject of ongoing research.

3. ROMS

ROMS is a hydrostatic, Boussinesq, free-surface primitive equations model that incorporates state-of-the-art physics and numerical algorithms (Shchepetkin and McWilliams 2003, 2005; Warner et al. 2005). ROMS is modular and uses C preprocessing to activate the various physical, numerical, and open boundary condition options. The parallel framework is coarse grained with both shared-memory (OpenMP) and distributed-memory (MPI) paradigms coexisting in the same code. The parallelization of the adjoint model is only available for MPI.

a. ROMS: Defining a penalty

The ROMS tangent linear and adjoint models predate the interfacing of ROMS with the IOM (Moore et al. 2004). The exact adjoint is defined relative to the Euclidean ℓ_2 norm following Giering and Kaminski (1998). The implied penalty functional is

$$J(\mathbf{u}, \theta, S, \zeta) = \hat{\boldsymbol{\lambda}}_M^T \boldsymbol{\lambda}_M + \hat{\boldsymbol{\lambda}}_\theta^T \boldsymbol{\lambda}_\theta + \hat{\boldsymbol{\lambda}}_S^T \boldsymbol{\lambda}_S + \hat{\boldsymbol{\lambda}}_\zeta^T \boldsymbol{\lambda}_\zeta$$
$$+ [\mathcal{L}(\boldsymbol{\Phi}) - \mathbf{d}]^T \mathbf{V}^{-1} [\mathcal{L}(\boldsymbol{\Phi}) - \mathbf{d}], \qquad (5)$$

where $\boldsymbol{\Phi} = (\mathbf{u}, \theta, S, \zeta)^T$, $\mathbf{u}$ is horizontal velocity, θ is potential temperature, S is salinity, and ζ is the free-surface elevation.

b. ROMS: Details of IOM implementation

The single executable framework of the IOM initially caused some difficulties because of the sophisticated I/O data structure of ROMS. For this reason the early version linked ROMS to the IOM core programs using MATLAB scripts rather than the IOM graphical user interface (GUI). The MATLAB scripts enabled NetCDF data exchange between ROMS and IOM without restructuring the ROMS I/O. This step proved not only to be successful for obtaining the first inversion of the representer matrix (Di Lorenzo et al. 2007), but it also allowed the use of IOM symmetry checks to perform a more thorough debugging of the ROMS adjoint and tangent linear components.

The I/O of the ROMS submodels has now been restructured to be compatible with the I/O requirements of the IOM. An important advantage of this compatibility is the ease with which the user can build new assimilation configurations from existing ROMS applications. The user needs only to compile the ROMS single executables and provide the IOM with the observations to assimilate. Given that IOM allows the assimilation of a wide range of observation types and that ROMS is widely used in coastal application, this IOM–ROMS framework could have a strong impact on the modeling components of the regional coastal observing systems initiatives, which is a component of the Integrated Ocean Observing System (IOOS; http://www.ocean.us). There are still some pending issues associated with the convolution operations of the adjoint fields. Preliminary experiments that assimilate real California Cooperative Oceanic Fisheries Investigations (CalCOFI) observations in the Southern California Current show that such an assumption of bell-shaped covariance function limits the convergence of the inner loop and degrades the quality of the solution. An efficient, customized convolution routine will eventually be required if a different covariance is implemented.

c. ROMS: Results

Preliminary assimilation experiments are conducted in the Southern California Bight (SCB), a region characterized by strong mesoscale eddy variability. Over the last 55 yr, the SCB has been extensively sampled by the CalCOFI program, which provides an unparalleled time series of physical, chemical, and biological data, well suited for data assimilation. The core physical observations consist of upper-ocean (0–500 m) CTD casts, collected 4 times per year over a 7–10-day period. The horizontal sampling array has a resolution of approximately 70 km, which is inadequate to fully resolve the mesoscale eddy field, as confirmed by recent higher-resolution observations of sea surface temperature and height (SST and SSH) from satellites and acoustic Doppler current profiler measurements between CalCOFI stations. Therefore, the assimilation of CalCOFI observations in ROMS using IOM may provide a means to better resolve and diagnose the ocean mesoscale. To test this hypothesis, the IOM is used to assimilate synthetic CalCOFI-like observations (Fig. 4a) in a model twin experiment. The accuracy with which the IOM is able to reconstruct the ocean mesoscale from spatially and temporally aliased observations is examined. The goal is not to reconstruct the ocean state between the four ocean surveys, which are clearly too far apart in time, but rather to provide ocean state estimates during the surveys themselves since the spatial resolution of the surveys is not sufficient to resolve the ocean mesoscale. In this context, the assimilation is performed to initialize the ocean model over the 7–15-day data collection period. After the ocean model has been initialized using the individual surveys, it can be used to
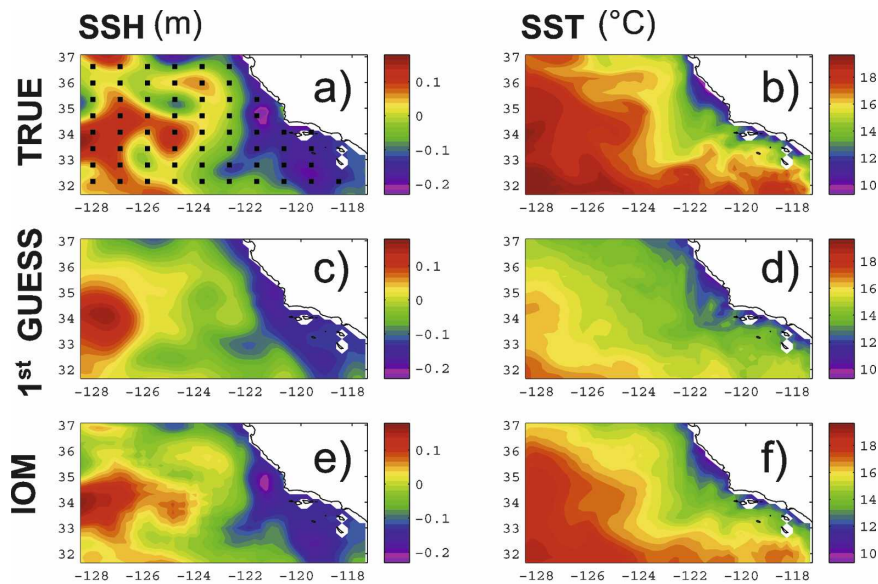
FIG. 4. (left) SSH and (right) SST at the end of the assimilation period. (a), (b) Truth, (c), (d) prior, and (e), (f) solution obtained using the IOM weak-constraint assimilation. The black dots on the SSH map in (a) indicate the location of the sampling array used to produce the synthetic observations.

diagnose the dynamical balances and perhaps exploit (in future studies) the predictability associated with the slow-evolving component of the coastal ocean (e.g., large eddies).

The synthetic observations are generated with a forward integration of nonlinear ROMS using a high-resolution grid of the SCB. The model domain is constructed using real coastlines and a smooth version of satellite bottom topography (Sandwell and Smith 1997). The model is forced with the National Centers for Environmental Prediction (NCEP) heat and freshwater fluxes and high-resolution wind stresses from a downscaled version of the NCEP winds obtained with a regional spectral atmospheric model. At the three open boundaries, a clamped open boundary condition is specified using the output of a coarse-resolution California Current ROMS configuration. It takes approximately 2 weeks for an oceanographic cruise to cover the CalCOFI sampling grid. To generate the synthetic observations, nonlinear ROMS is integrated for 10 yr and sampled 1–14 January of Year 11 using a CalCOFI-like sampling array. The observations are collected in the south 1–5 January, in the center of the domain 6–9 January, and in the north 10–14 January. Observations consist of temperature and salinity from 0 to 500 m. A snapshot of the true model fields on 14 January of Year 11, for SSH and SST, reveals the typical meandering pattern of the California Current and its mesoscale structure (Fig. 4, row 1).

Prior to the assimilation, the ocean fields were generated by integrating the nonlinear ROMS initialized with the January climatological conditions. A comparison of the first guess with the true fields on 14 January (Fig. 4, row 2) shows an overall agreement in the mean SSH and surface velocity. However, the first guess fails to capture some of the eddy scale structures in the north and close to the coast and leads to an overall weaker SST gradient. Then, two assimilation experiments are performed. In the first experiment, adjustments are allowed only in the initial conditions; this is the strong-constraint case. In the second experiment, adjustments are also allowed in the forcing on the right-hand side of the dynamical equations; this is the weak-constraint case. The evolution of the mesoscale flow over a period of 10 days in January is primarily sensitive to errors in initial conditions rather than surface forcing. For this reason, surface forcing is not used as a control variable.

Both assimilation experiments lead to a substantial improvement when compared to the first guess (Fig. 4). The strong-constraint experiment successfully reduces the initial model–observation misfit by 75% (Fig. 5) and improves the model fields at locations where observations are not assimilated. In the weak-constraint experiment the initial model–observation is reduced by 89%. To verify the quality of the assimilation solution the model was run beyond the assimilation window for an additional 7 days to measure its predictive skill
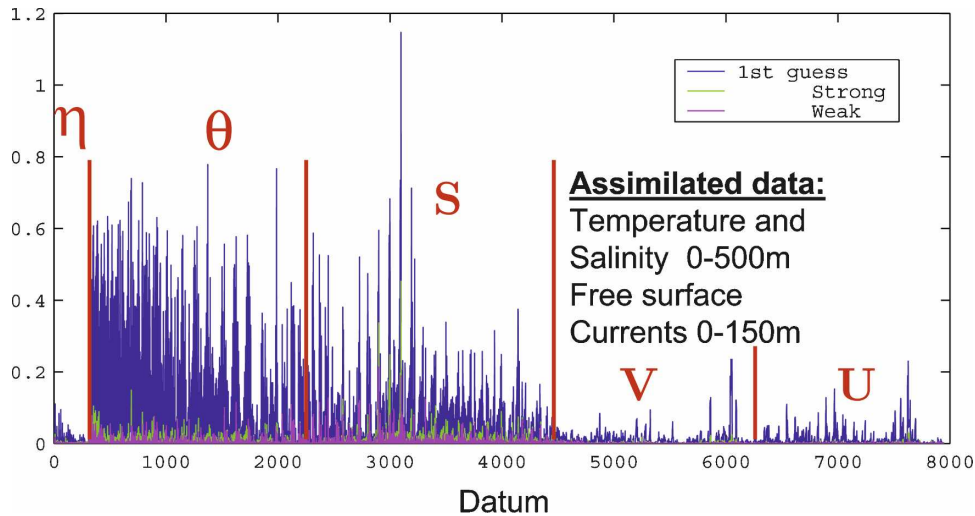
FIG. 5. Normalized misfits between model and the synthetic data. The std dev are used to normalize the misfits. The blue line represents the misfit for the prior (climatology in this case), the green line is for the model solution obtained using the strong-constraint assimilation in inverse ROMS, and the red line is the weak-constraint model solution. The misfit array includes free-surface, temperature, salinity, and horizontal velocity.

against independent observations. Both the strong and weak-constraint case show forecast skill greater than persistence and climatology for SST. However, SSH, velocities, and subsurface data do not show real forecast skill. Factors that limit forecast skill are linear instabilities that develop in the tangent linear model (Di Lorenzo et al. 2007) and the choice of Gaussian background covariance. Future assimilation experiments are planned to address both these factors.

A more in-depth study of the SCB assimilation case using real observations is currently being performed. These preliminary experiments suggest that the IOM is capable of extracting dynamical information from the spatially and temporally aliased CalCOFI observations and improving our diagnosis of the mesoscale ocean circulation in the Southern California coastal ocean.

4. ADCIRC-2D

ADCIRC-2D is a finite-element circulation model for shelves, coasts, and estuaries (Kolar et al. 1994; Luettich et al. 1991; Luettich and Westerink 2000) that is used extensively by Naval Oceanographic Office (NAVOCEANO), the Office of Naval Research (ONR), the U.S. Army Corps of Engineers, the National Oceanic and Atmospheric Administration (NOAA), the U.S. Geological Survey (USGS), and a wide range of academic researchers. The extensive use of ADCIRC-2D, operational and otherwise, together with the derived nature of the governing equations,

lends a particular challenge to working with the IOM as will be described in the following sections.

a. ADCIRC-2D: Governing equations

The governing equations are the two-dimensional continuity and conservation of momentum equations, including advection, constant Coriolis, eddy viscosity, and nonlinear bottom friction. The momentum equation is forced by surface gradients, atmospheric pressure gradients, tidal potentials, and surface stresses. Lateral boundary conditions are either specified flow, specified elevation, or some combination. The solution to these equations using finite elements on linear triangular elements leads to an anomalous, folded dispersion relationship (i.e., for at least one wave of low frequency, there is another of high frequency with the same phase speed; the latter is a numerical artifact that causes spurious oscillations in the solution). These equations can be reformulated by subtracting the divergence of the momentum equation from the time derivative of the continuity equation. The result is the well-known second-order wave continuity equation (WCE) with the attractive feature that it "unfolds" the dispersion relationship (Lynch and Gray 1979; Kinnmark 1986).

b. ADCIRC-2D: Defining a penalty

The obvious approach to defining the penalty would be to penalize the residuals in the wave equation, mo-

mentum equation, boundary conditions, initial conditions, and data. However, the null hypothesis must include first and second moments of the errors in each of these components. Given that the wave equation is a derived equation, characterization of its residuals is extremely difficult from a scientific perspective. It is more scientifically feasible to characterize errors in the continuity equation rather than the wave equation; in fact, conservation of mass could be enforced as a strong constraint. This is an extremely attractive feature in a model that otherwise does not conserve mass locally. Thus, the integral penalty is defined in terms of residuals in the continuity equation (rather than WCE), momentum equation, boundary conditions, initial conditions, and data:

$$J(\zeta, \mathbf{u}) = \sum_k \lambda_C^k \circ \hat{\lambda}_C^k \Delta t + \sum_k \boldsymbol{\lambda}_M^k \circ \hat{\boldsymbol{\lambda}}_M^k \Delta t$$

$$+ \cdots + i \circ \hat{\imath} + \sum_k \int_{\Gamma_o} b^k \hat{b}^k \Delta t$$

$$+ [\mathcal{L}(\boldsymbol{\Phi}) - \mathbf{d}]^{\mathrm{T}} \mathbf{V}^{-1} [\mathcal{L}(\boldsymbol{\Phi}) - \mathbf{d}], \qquad (6)$$

where λ_C^k is the residual in the continuity equation, i is the residual in the initial condition, b^k is the residual in the boundary conditions, Γ_o is the open ocean boundary, $\boldsymbol{\Phi} = (\mathbf{u}, \zeta)$, Δt is the time step, and the superscript k indicates the time steps. The ellipsis in (6) represents terms arising from the "startup" equations required for the multilevel time-stepping scheme. Closed boundaries are imposed as strong constraints. Note that the penalty is discrete in time, but continuous in space.

c. ADCIRC-2D: Details of IOM implementation

Because the ADCIRC I/O structure is relatively simple, the IOM requirement that the various components be separate executables was easy to accomplish. The tangent-linear model was developed first (Muccino and Luo 2005). Developing the adjoint model was the most difficult task of the IOM implementation. As will be described in the next section, the derivation of the adjoint was not straightforward. It is important to emphasize here that these difficulties are a result of the derived nature of the ADCIRC-2D governing equations and not of the representer method or the IOM.

The convolutions were computed by direct summation. The ocean observing system here is simply a collection of point measurements of surface elevation and velocity; this simplest of measurement kernels was supplied to the IOM.

The IOM facilitated the ADCIRC study not only by generating the infrastructure for effecting the "data space" search, but also by generating scripts to run the

assimilation, to perform adjoint symmetry checks and other algebraic consistency checks, and to generate the error fields required to generate the synthetic data described below. Each of these tasks would otherwise have required extensive, intricate coding.

d. ADCIRC-2D: Quantifying suboptimality

The disadvantage of the formulation in (6) is that the resulting adjoint system has the same folded dispersion relationship as the forward system; it must therefore be reformulated into an adjoint wave equation. This apparently benign step is the root of several difficulties, each of which causes the adjoint to be inconsistent with the forward equations, leading ultimately to suboptimality of the inverse solution.

Ideally, the penalty should be defined using discretized governing equations; if the continuous form of the governing equation is used, some of the adjoint boundary conditions will be incorrect. However, the penalty functional used here is continuous in space because it is algebraically impossible to develop the adjoint wave equation from an adjoint system that is already discretized in space. The spatial discretization is therefore performed after the adjoint wave equation is derived. If the resolution is sufficiently fine, the resulting suboptimality is small, but present nonetheless.

The WCE is a reformulation of the continuity and conservation of momentum equations. Thus, many of the terms in the WCE have their origins in conservation of momentum equation. ADCIRC-2D uses different time-stepping schemes for the terms in the WCE (three level) and terms in the momentum equation (Crank–Nicolson), resulting in a numerical inconsistency. If the time step is sufficiently small, the resulting suboptimality should be inconsequential in any practical sense. Although the time-stepping scheme could have been made consistent for this work, given the widespread operational use of ADCIRC-2D and the minor impact of the time-stepping inconsistency on the forward model solutions, the original formulation was maintained.

ADCIRC-2D solves the full WCE system of equations, but lumps the mass matrix of the discrete momentum equations.[1] Levin et al. (2006) have shown that to maintain adjoint consistency with lumping schemes, the mass matrix of the discrete adjoint continuity equation should be lumped and the mass matrix of the discrete adjoint momentum equation should not. How-

---

[1] Sum each row of the matrix, zero the row, and place the sum on the diagonal, thus eliminating the need to solve a $2N \times 2N$ system of equations.

ever, since the wave equation is solved here, the approach of Levin et al. (2006) cannot be used. Based upon heuristic reasoning and experimentation, the best option from both an accuracy and efficiency perspective is to lump the mass matrix of both the discrete adjoint WCE and momentum equations.

As stated before, these difficulties manifest themselves in an adjoint wave equation that is not consistent with the forward wave equation. Any data assimilation algorithm that relies on the adjoint, such as the popular gradient search algorithm (e.g., Stammer et al. 2003), will yield a suboptimal solution. There are many instances in the literature that use adjoint models that are inconsistent with the forward models. For example, to reduce strong sensitivities in the adjoint model, Marotzke et al. (1999) and Stammer et al. (2003) often turn off mixing schemes such as K-profile parameterization (KPP) (Large et al. 1994) or Gent–McWilliams (Gent and McWilliams 1990) in the adjoint model. Some of the European Centre for Medium-Range Weather Forecasts (ECMWF) 4DVAR implementations have run the adjoint at lower resolution than the forward run, also with the goal of suppressing strong sensitivities arising from nonlinearities (Leidner et al. 2003). In these cases the descent methods used in strong-constraint 4DVAR produced improved solutions for the unaltered forward model, and so the inconsistency between the adjoint model and the forward model was deemed inconsequential. In all of these cases, the adjoint inconsistencies are introduced by a conscious choice of the modeler; in the case of ADCIRC, however, the inconsistencies are born of the governing equations themselves and cannot be avoided.

While disturbing from a theoretical and programming perspective (perfect adjoint symmetry provides valuable debugging tools), the difficulties described above do not necessarily diminish the value of the assimilation algorithm in a practical sense. One approach for quantifying the impact is a statistical hypothesis test. When the null hypothesis is correct (i.e., the errors in the model and data are hypothesized correctly) and the adjoint and forward models are consistent, the minimum value of the penalty, $\hat{J}$, has a chi-square distribution, χ^2_M, where M is the number of scalar data (Bennett 2002; Muccino et al. 2004). The IOM can synthesize data with prescribed moments; when the same statistical hypothesis is used to synthesize and to assimilate the data, the hypothesis test will only fail if there is an inconsistency between the adjoint and forward models. Numerical experiments on several domains show that the internal inconsistencies described above do not impact the $\hat{J}$ distribution if the spatial and temporal domains are sufficiently resolved.

e. ADCIRC-2D: Results

The Bight of Abaco is characterized by highly non-linear dynamics (Muccino and Luo 2005). The shallow domain is bounded by islands of the Bahamas except along the southwest open boundary (Fig. 6), where the tide is driven by O_1, K_1, N_2, M_2, and S_2 tidal constituents.

The purpose of the assimilation is to test the internal consistency of the algorithm. The simulation is spun up from rest for one day to avoid excitation of resonant frequencies and then the simulations are run for another day. The parameters related to the hypothesis are estimated from the scales of the dynamics. However, the actual values are insignificant because the purpose of the test is to examine the consistency of the algorithm. The important point is that the parameters in the hypothesis used to synthesize the data are the same as the parameters in the hypothesis used by the assimilation. One hundred datasets, each with 54 synthesized data (16 elevation data and 32 scalar velocity data), are assimilated; each dataset yields one value of $\hat{J}$.

According to theory, the mean value of $\hat{J}$ should be $M = 54$ and the standard deviation should be $\sqrt{2M} = 10.4$. In this experiment, the mean value of the 100 $\hat{J}$ values was 55.2 (reduced from $J_{\text{prior}} = 4.8 \times 10^4$), just 0.1 standard deviation from the theoretical value. Although the adjoint is not consistent with the forward model, that inconsistency is not statistically significant here. However, if the forward model has insufficient resolution in space or time, the inverse solution is quite suboptimal (not shown); of course, in such a case even the forward model solution will be meaningless. The resolution requirements of the assimilation algorithm appear to be no more stringent than those for the forward model. Figures 7a–c show the prior elevation solution, the true solution, and the optimal solution, respectively, for one of the synthetic datasets. Throughout the domain, the optimal solution is much closer to the true solution than the prior solution is.

5. SEOM-2D

The governing equations are the two-dimensional continuity and conservation of momentum with Coriolis acceleration approximated on a beta-plane, reduced-gravity, and linear bottom friction representation. The discretization is based on a Galerkin formulation, where the basis functions and test functions are Legendre Cardinal functions defined locally within each element. The sole global requirement is that the interpolation be continuous across element boundaries. The continuity of the normal derivative is enforced weakly
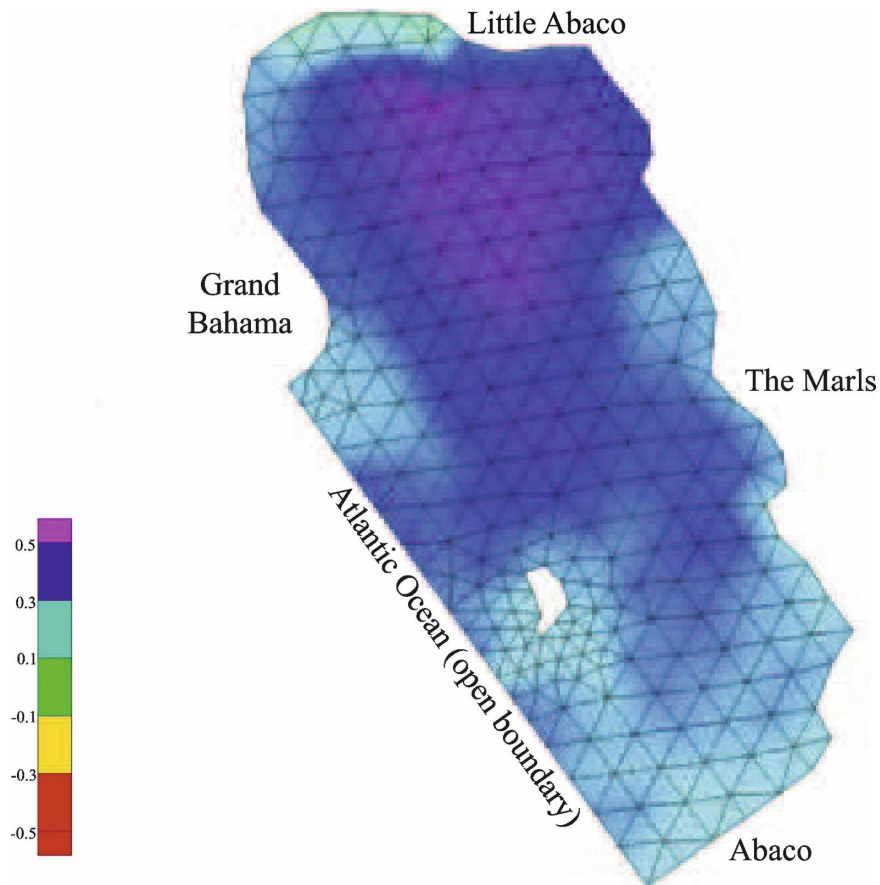
FIG. 6. The computational grid for the Bight of Abaco. The domain is bounded by the islands of the Bahamas except for the southwest open boundary, where the tide is driven by five tidal constituents. The bathymetry (shown in m) is shallow and the dynamics are very nonlinear. The grid has 259 and 440 elements with a relatively constant node spacing of approximately 2.7 km.
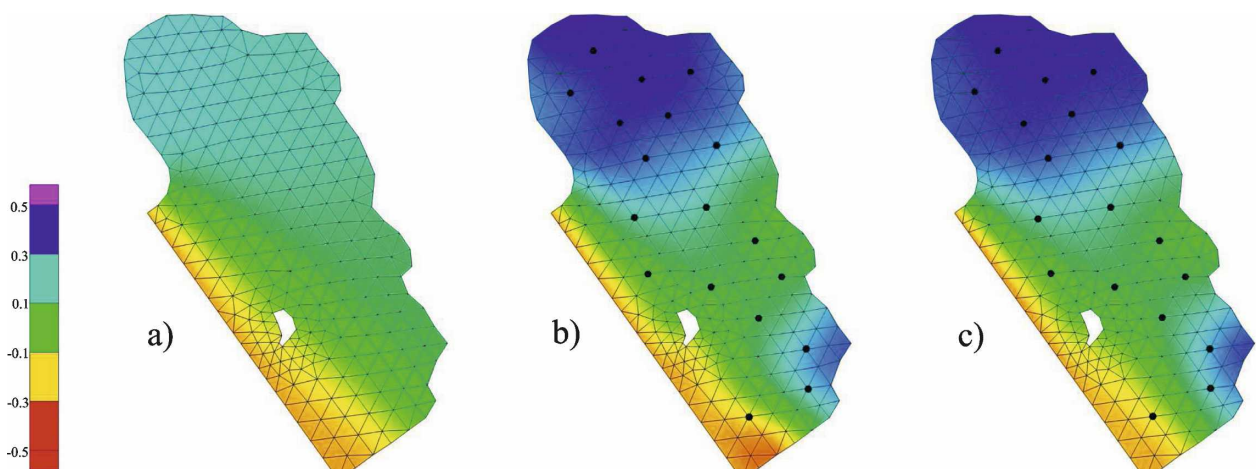

FIG. 7. Results of assimilation. (a) Prior solution, (b) true solution, and (c) optimal solution; black circles indicate locations where synthetic data (surface elevation and scalar components of velocity) are collected at three time levels. Assimilation of the data has led to an optimal solution that is much closer to the true solution than the prior.

by the variational formulation, that is, the jump between two different representations of the normal derivative, arising from two neighboring elements, is allowed on an elemental interface. The jump is accounted for through Galerkin integration. More detail on the spectral element discretization and the choice of basis functions is given in Iskandarani et al. (1995). SEOM-2D is used as a "barotropic" engine in SEOM-3D (Iskandarani et al. 2003).

The tangent-linear and adjoint versions of SEOM, together with the convolution and interpolation operators, are incorporated into the IOM as separate executables. Both of the latter had already been developed for other purposes, thus their implementation did not impose an additional coding burden. The adjoint code was obtained by differentiating the discrete cost function, following procedures described in Bennett (2002).

a. SEOM-2D: Defining a penalty

The discrete adjoint equations are obtained from the following Euclidean penalty function:

$$J(\zeta, \mathbf{u}) = \sum_m \sum_k \lambda_{C_m}^k \hat{\lambda}_{C_m}^k \Delta t + \sum_m \sum_k \boldsymbol{\lambda}_{M_m}^k \hat{\boldsymbol{\lambda}}_{M_m}^k \Delta t$$
$$+ \sum_m i_m \hat{\imath}_m + [\mathcal{L}(\boldsymbol{\Phi}) - \mathbf{d}]^{\mathrm{T}} \mathbf{V}^{-1} [\mathcal{L}(\boldsymbol{\Phi}) - \mathbf{d}],$$
(7)

where the subscripts and superscripts on the residual and adjoint variables represent the collocation point and time step, respectively, and $\boldsymbol{\Phi} = (\mathbf{u}, \zeta)$. The boundary conditions are no-flow and are imposed as strong constraints.

Derivation of the discrete adjoint system is given in Levin et al. (2006), where it is shown that it is possible to obtain an exact discrete adjoint system from a continuous adjoint system by interchanging the two different forms of spatial derivatives that are allowed in Galerkin formulation. In a Galerkin integral, a spatial derivative can be applied using one of two different approaches. The first involves application of the derivative to the field itself (i.e., to each of the basis functions used in its discrete representation); this is known as the "strong" form of Galerkin derivative. The second involves application of the derivative to the test function instead, and the Galerkin integral is obtained by integration by parts; this is known as the "weak" form.[2] If

a weak form of the horizontal velocity divergence operator is used in the Galerkin formulation of the forward system, then a strong form of the pressure gradient operator must be used in the Galerkin formulation of the adjoint system, and vice versa: the strong form of the divergence operator in the forward system leads to a weak form of the gradient operator in the adjoint system. Similarly, the strong form of the gradient operator in the forward system leads to a weak form of the divergence in the adjoint system. And finally, the strong form of the advection operator in the forward system leads to a weak form of the advection operator in the adjoint system.

b. SEOM-2D: Results

The inversion has been tested using the canonical problem of a wind-driven, double-gyre circulation in a midlatitude ocean as described in Haidvogel et al. (1992), Holland (1978), and McCalpin and Haidvogel (1996). The goal is to study the impact of the quantity and quality of data on the optimal solution. Synthetic observations of interface displacement are computed from a twin experiment. The dynamics and the initial condition are imposed as weak constraints, and the former are uncorrelated in time. Spatial convolutions are computed using the default IOM implementation in conjunction with the SEOM diffusion operator.

The linearization error is small, so only one outer iteration is performed; the background state for the first iteration of the tangent linearization is the field computed by the nonlinear forward model, and the indirect representer algorithm is used. The indirect representer algorithm, together with prior and posterior error analysis, is supplied by the IOM. The solution is obtained with nine inner iterations when assimilating 7081 data and with eleven inner iterations when assimilating 14 162 data. In all experiments, the interface displacement and velocity fields are significantly closer to the true solution after the assimilation of data. A better fit is obtained with either more or better quality data. More detail on the twin experiment results can be found in Levin et al. (2006).

6. Conclusions

Four ocean models, PEZ-HAT, ROMS, ADCIRC-2D, and SEOM-2D have been used successfully in conjunction with the new Inverse Ocean Modeling system (Bennett et al. 2006). These models are very different in maturity, physics, numerics, and structure and have been used to investigate a wide variety of scientific phenomena. While the implementation of the IOM in con-

---

[2] Use of "strong" and "weak" here should not to be confused with strong and weak forms in variational assimilation used elsewhere, including in this paper.

junction with these models is important, even more important is the demonstration of the flexibility, power, and usefulness of the IOM; the goal of the project has been to enable the IOM to interface with *any* numerical model and observation system. Indeed, the model need not even be an ocean model, or even a fluid dynamical model. Other types of models, such as the Geophysical Finite Element Simulation Tool (GeoFEST), a crustal geodynamics model (Lyzenga and Parker 2006), are currently also being interfaced with IOM, enabling the assimilation of new data sources [e.g., GPS and Interferometric Synthetic Aperture Radar (InSAR)]. Given the drive toward integrating a wide variety of observations into numerical models, the development of a model-independent, modular assimilation system such as the IOM is particularly timely for the broader geoscience community.

REFERENCES

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

——, B. S. Chua, B. L. Pflaum, M. Erwig, Z. Fu, R. D. Loft, and J. C. Muccino, 2006: Inverse Ocean Modeling System user manual. IOM Tech. Doc. 6, 221 pp. [Available online at ftp://ftp.coas.oregonstate.edu/dist/chua/iom/tech_doc/IOM_TD6.pdf.]

——, ——, ——, ——, ——, ——, and ——, 2008: The Inverse Ocean Modeling System. Part I: Implementation. *J. Atmos. Oceanic Technol.,* **25,** 1608–1622.

Bryan, K., and M. D. Cox, 1968: A nonlinear model of an ocean driven by wind and differential heating. Part I: Description of the three-dimensional velocity and density fields. *J. Atmos. Sci.,* **25,** 945–967.

Chua, B. S., and A. F. Bennett, 2001: An inverse ocean modeling system. *Ocean Modell.,* **3,** 137–165.

Curchitser, E. N., M. Iskandarani, and D. B. Haidvogel, 1998: A spectral element solution of the shallow water equations on multiprocessor computers. *J. Atmos. Oceanic Technol.,* **15,** 510–521.

——, D. Haidvogel, and M. Iskandarani, 1999: On the transient adjustment of a mid-latitude abyssal ocean basin with realistic geometry: The constant depth limit. *Dyn. Atmos. Oceans,* **29,** 147–188.

Davis, G., and R. Rew, 1990: NetCDF: An interface for scientific data access. *IEEE Comput. Graph. Appl.,* **10,** 76–82.

Di Lorenzo, E., A. M. Moore, H. G. Arango, B. D. Cornuelle, A. J. Miller, B. Powell, B. S. Chua, and A. F. Bennett, 2007: Weak and strong constraint data assimilation in the inverse Regional Ocean Modeling System (ROMS): Development and application for a baroclinic coastal upwelling system. *Ocean Modell.,* **16,** 160–187.

Duffy, B., 2005: Anatomy of a disaster: 5 days that changed a nation. Saturday, August 27. *US News World Rep.,* **139,** 24–43.

Gent, P. R., and J. C. McWilliams, 1990: Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.,* **20,** 150–155.

Giering, R., and T. Kaminski, 1998: Recipes for adjoint code construction. *ACM Trans. Math. Software,* **24,** 437–474.

Haidvogel, D. B., and A. Beckmann, 1999: *Numerical Ocean Circulation Modeling.* Imperial College Press, 300 pp.

——, J. C. McWilliams, and P. R. Gent, 1992: Boundary current separation in a quasigeostrophic, eddy-resolving ocean circulation model. *J. Phys. Oceanogr.,* **22,** 882–902.

Holland, W. R., 1978: The role of mesoscale eddies in the general circulation of the ocean—Numerical experiments using a wind-driven quasi-geostrophic model. *J. Phys. Oceanogr.,* **8,** 363–392.

Iskandarani, M., D. B. Haidvogel, and J. P. Boyd, 1995: A staggered spectral element model with application to the oceanic shallow water equations. *Int. J. Numer. Methods Fluids,* **20,** 393–414.

——, ——, and J. G. Levin, 2003: A three-dimentional spectral element model for the solution of hydrostatic primitive equations. *J. Comput. Phys.,* **186,** 397–425.

Killworth, P. D., D. Stainforth, D. J. Webb, and S. M. Paterson, 1991: The development of a free-surface Bryan–Cox–Semtner ocean model. *J. Phys. Oceanogr.,* **21,** 1333–1348.

Kinnmark, I., 1986: *The Shallow Water Wave Equations: Formulation, Analysis and Application.* Vol. 15, *Lecture Notes in Engineering,* Springer-Verlag, 187 pp.

Kolar, R. L., W. G. Gray, J. J. Westerink, and R. A. Luettich, 1994: Shallow water modeling in spherical coordinates: Equation formulation, numerical implementation, and application. *J. Hydraul. Res.,* **32,** 3–24.

Large, W. G., J. C. McWilliams, and S. Doney, 1994: Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameterization. *Rev. Geophys.,* **32,** 363–403.

Leidner, S. M., L. Isaksen, and R. N. Hoffman, 2003: Impact of NSCAT winds on tropical cyclones in the ECMWF 4DVAR assimilation system. *Mon. Wea. Rev.,* **131,** 3–26.

Levin, J. C., M. Iskandarani, and D. B. Haidvogel, 1997: A spectral filtering procedure for eddy-resolving simulations with a spectral element ocean model. *J. Comput. Phys.,* **137,** 130–154.

——, D. B. Haidvogel, B. Chua, A. F. Bennett, and M. Iskandarani, 2006: Euler–Lagrange equations for the spectral element shallow water system. *Ocean Modell.,* **12,** 348–377.

Luettich, R. A., and J. J. Westerink, 2000: ADCIRC: A (parallel) advanced circulation model for oceanic, coastal, and estuarine waters: Consolidated and updated theory and users manual. [Available online at http://adcirc.org/.]

——, ——, and N. W. Scheffner, 1991: ADCIRC: An advanced three-dimensional circulation model for shelves, coasts, and estuaries. Report 1: Theory and methodology of ADCIRC-2DDI and ADCIRC-3DL. Dredging Research Program Tech. Rep. DRP-92–6, 143 pp.

Lynch, D. R., and W. G. Gray, 1979: A wave equation model for finite element tidal computations. *Comput. Fluids,* **7,** 207–228.

Lyzenga, G., and J. Parker, 2006: GeoFEST: Geophysical finite

element simulation tool. Caltech/Jet Propulsion Laboratory. [Available online at http://www.openchannelfoundation.org/projects/GeoFEST.]

Marotzke, J., R. Giering, K. Q. 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,** 29 529–29 547.

McCalpin, J. D., and D. B. Haidvogel, 1996: Phenomenology of the low-frequency variability in a reduced-gravity, quasigeostrophic double-gyre model. *J. Phys. Oceanogr.,* **26,** 739–752.

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.

Muccino, J. C., and H. Luo, 2005: Picard iterations for a finite element shallow water equation model. *Ocean Modell.,* **10,** 316–341.

——, A. F. Bennett, and N. F. Hubele, 2004: Significance testing for variational assimilation. *Quart. J. Roy. Meteor. Soc.,* **130,** 1815–1838.

Pacanowski, R. C., and A. Gnanadesikan, 1998: Transient response in a *z*-level ocean model that resolves topography with partial cells. *Mon. Wea. Rev.,* **126,** 3248–3270.

Purser, R. J., W.-S. Wu, D. F. Parrish, and N. M. Roberts, 2003: Numerical aspects of the application of recursive filters to variational statistical analysis. Part II: Spatially inhomogeneous and anisotropic general covariances. *Mon. Wea. Rev.,* **131,** 1536–1548.

Sandwell, D. T., and W. H. F. Smith, 1997: Marine gravity anomaly from Geosat and ERS 1 satellite altimetry. *J. Geophys. Res.,* **102,** 10 039–10 054.

Shchepetkin, A. F., and J. C. McWilliams, 2003: A method for computing horizontal pressure-gradient force in an oceanic model with a nonaligned vertical coordinate. *J. Geophys. Res.,* **108,** 3090, doi:10.1029/2001JC001047.

——, and ——, 2005: The regional oceanic modeling system (ROMS): A split-explicit, free-surface, topography-following coordinate oceanic model. *Ocean Modell.,* **9,** 347–404.

Stammer, D., and Coauthors, 2003: Volume, heat, and freshwater transports of the global ocean circulation 1993–2000, estimated from a general circulation model constrained by World Ocean Circulation Experiment (WOCE) data. *J. Geophys. Res.,* **108,** 3007, doi:10.1029/2001JC001115.

Travis, J., 2005: Scientists' fears come true as hurricane floods New Orleans. *Science,* **309,** 1656–1658.

Warner, J. C., C. R. Sherwood, H. G. Arango, R. P. Signell, and B. Butman, 2005: Performance of four turbulence closure models implemented using a generic length scale method. *Ocean Modell.,* **8,** 81–113.

Weaver, A., and P. Courtier, 2001: Correlation modelling on the sphere using a generalized diffusion equation. *Quart. J. Roy. Meteor. Soc.,* **127,** 1815–1846.

Wunsch, C., D. B. Haidvogel, M. Iskandarani, and R. Hughes, 1997: Dynamics of the long-period tides. *Prog. Oceanogr.,* **40,** 81–108.

Zaron, E. D., 2006: A comparison of data assimilation methods using a planetary geostrophic model. *Mon. Wea. Rev.,* **134,** 1316–1328.

——, and G. D. Egbert, 2006: Verification studies for a *z*-coordinate primitive-equation model: Tidal conversion at a mid-ocean ridge. *Ocean Modell.,* **14,** 257–278.

——, and ——, 2007: The impact of the M_2 internal tide on data-assimilative model estimates of the surface tide. *Ocean Modell.,* **18,** 210–216, doi:10.1016/j.ocemod.2007.05.001.