Available online at www.sciencedirect.com
ScienceDirect
Coastal Engineering 54 (2007) 1–15
Coastal Engineering
An International Journal for Coastal, Harbour and Offshore Engineers
www.elsevier.com/locate/coastaleng

Forecasting and hindcasting waves with the SWAN model in the Southern California Bight

W. Erick Rogers [a,*], James M. Kaihatu [a], Larry Hsu [a], Robert E. Jensen [b], James D. Dykes [a], K. Todd Holland [c]

[a] *Oceanography Division, Code 7320, Naval Research Laboratory, Stennis Space Center, MS, USA*
[b] *U.S. Army Engineer Research and Development Center, Vicksburg, MS, USA*
[c] *Marine Geosciences Division, Code 7440, Naval Research Laboratory, Stennis Space Center, MS, USA*

Received 20 April 2005; received in revised form 27 June 2006; accepted 28 June 2006
Available online 28 August 2006

Abstract

The Naval Research Laboratory created a wave forecasting system in support of the Nearshore Canyon Experiment (NCEX) field program. The outer nest of this prediction system encompassed the Southern California Bight. This forecasting system is described in this paper, with analysis of results via comparison to the extensive buoy network in the region. There are a number of potential errors, two of which are poor resolution of islands in the Bight—which have a strong impact on nearshore wave climate—and the use of the stationary assumption for computations. These two problems have straightforward solutions, but the solutions are computationally expensive, so an operational user must carefully consider their cost. The authors study the impact of these two types of error (relative to other errors, such as error in boundary forcing) using several hindcasts performed after the completion of NCEX. It is found that, with buoy observations as ground truth, the stationary assumption leads to a modest increase in root-mean-square error; this is due to relatively poor prediction of the timing of swell arrivals and local sea growth/decay. The model results are found to be sensitive to the resolution of islands; however, coarse resolution does not incur an appreciable penalty in terms of error statistics computed via comparison to buoy observations, suggesting that other errors dominate. Inaccuracy in representation of the local atmospheric forcing likely has a significant impact on wave model error. Perhaps most importantly, the accuracy of directional distribution of wave energy at the open ocean boundaries appears to be a critical limitation on the accuracy of the model-data comparisons inside the Bight.
Published by Elsevier B.V.

Keywords: Wave modeling; Wave forecasting; Wave hindcasting; Southern California Bight; Swells; Swell forecasting

1. Introduction

Wave forecasting systems are run routinely by the operational Navy for a number of coastal areas around the world. The operational Navy (specifically the Naval Oceanographic Office, NAVO), typically uses WAM ("WAve Model", WAMDI Group, 1988; Günther et al., 1992; Komen et al., 1994) to model sub-regional scale domains (e.g. the size of the domain depicted in Fig. 1A) and the SWAN model ("Simulating WAves Nearshore"; Booij et al., 1999) for a nearshore region such as the one depicted in Fig. 1C. [These grids will be introduced in detail later in this paper.]

The Nearshore Canyon Experiment (NCEX) was a coastal research project with field operations Sept. 16 through Dec. 15, 2003, located at and near La Jolla, California. This paper deals with the application of the SWAN model at both sub-regional and nearshore scale, for the area of the Southern California Bight, during the duration of the field experiment. In this environment, the key challenge is to accurately represent the propagation/blocking of swell energy through/by the islands of the Bight as they approach the NCEX area.

There has been previous work related to wave modeling in the Southern California Bight. The reader is referred to O'Reilly and Guza (1998) and references therein. Also, as of June 2006, there are active relevant websites run by the Coastal Data Information Program (CDIP).

An objective of this article is to learn ways in which a modeler might minimize errors in a forecasting system. However, the

* Corresponding author. Fax: +1 228 688 4759.
E-mail address: rogers@nrlssc.navy.mil (W.E. Rogers).

0378-3839/$ - see front matter. Published by Elsevier B.V.
doi:10.1016/j.coastaleng.2006.06.011

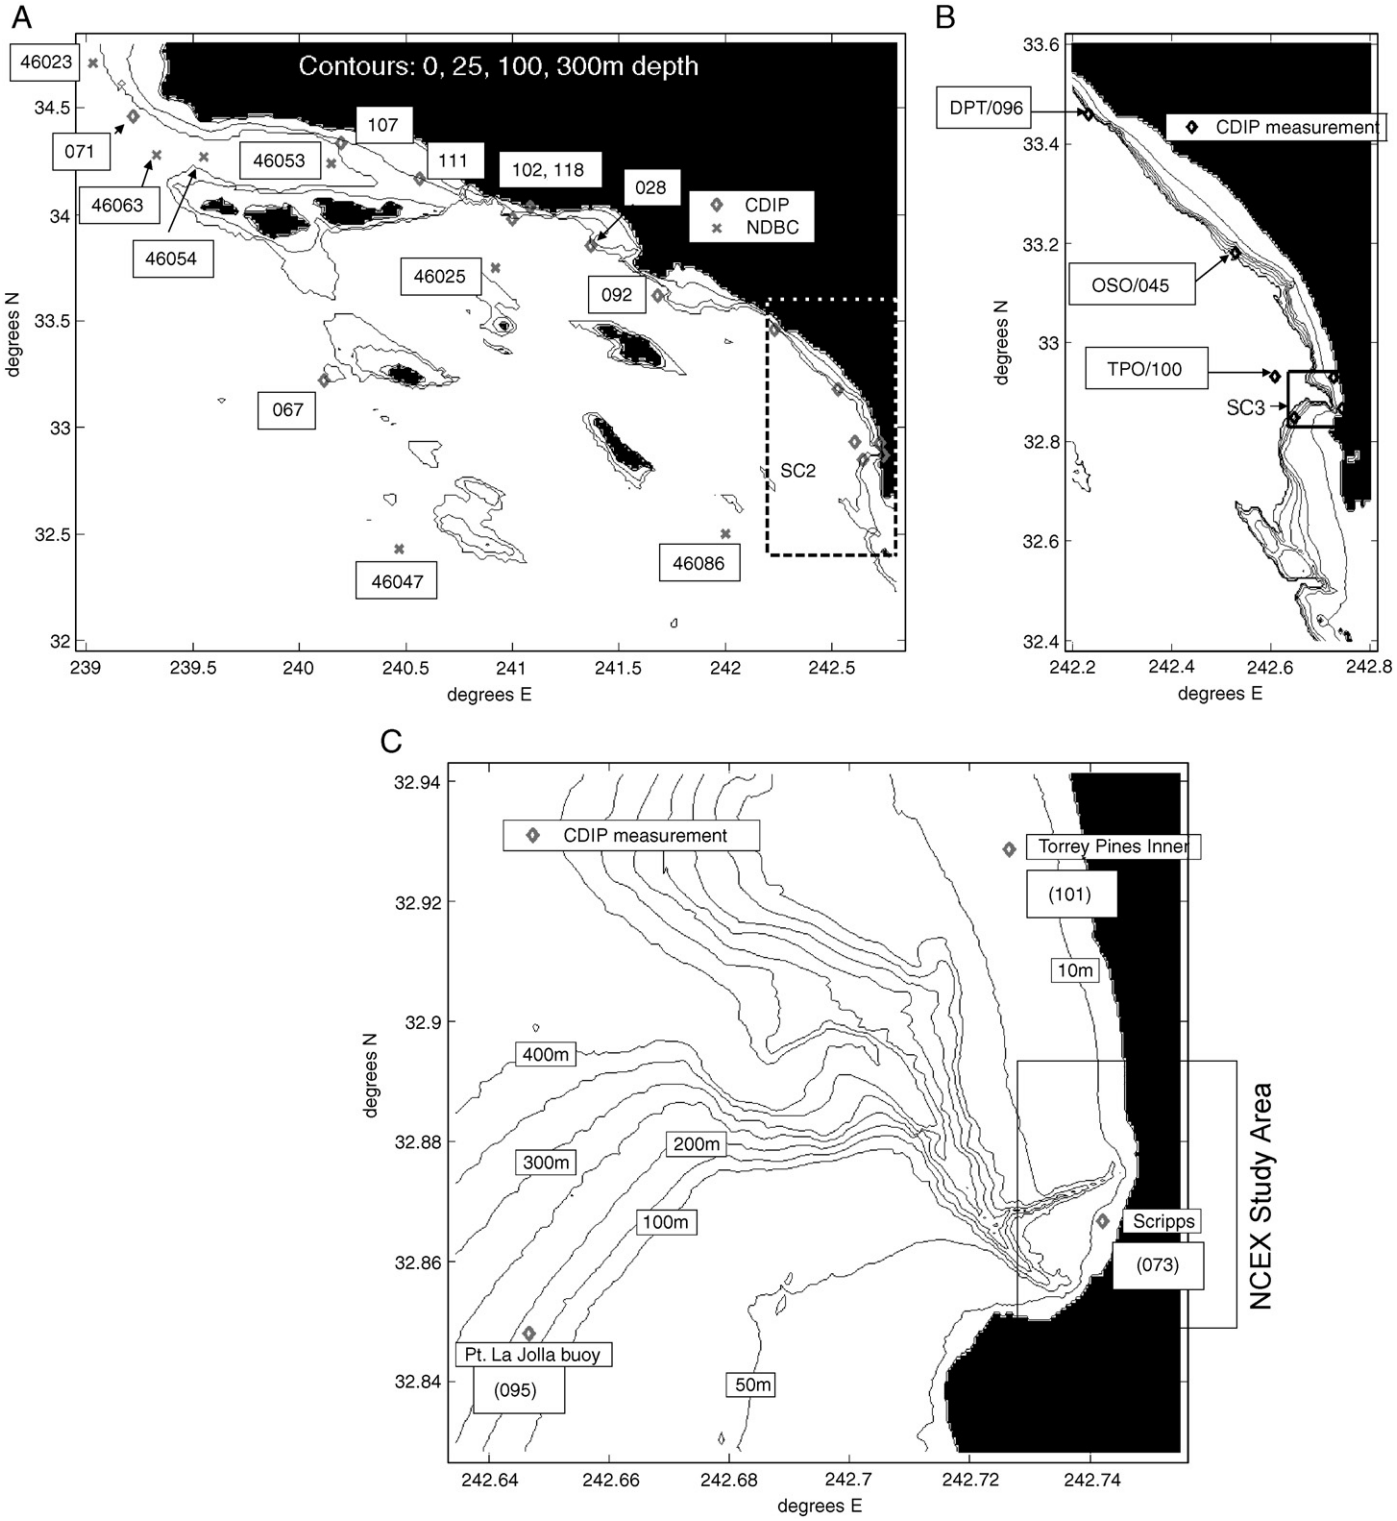

Fig. 1. A. The SC1 grid, with bathymetry. The 0, 25, 100, and 300 m depth contours are indicated. B. The SC2 grid, with bathymetry. The depth contours drawn at 50 m intervals, out to 400 m. C. The SC3 grid, with bathymetry.

article is not about tuning. Rather, it is about describing and discussing some key aspects (and problems) of wave modeling system design. There are a number of questions that a wave modeler is faced with for which there are no ready answers. Usually, the modeler makes decisions on these questions with a mix of experience and guesswork. Two of these questions are:

1) Use of stationary assumption for a large computational region can lead to poor timing of swell arrivals and temporal description of local growth and decay. Use of nonstationary computations for small regions is computationally wasteful. So, where should one make the hand-off from the larger-scale nonstationary model to the smaller-scale stationary model?

2) What geographic resolution is necessary for the outer nests? Is it better to spend CPU cycles on something other than high geographic resolution?

These two questions are essentially considerations of whether to apply two computational "shortcuts". The objectives of this study are to

- Evaluate the sensitivity of the model results to these two computational shortcuts.
- Determine the relative importance of other errors, such as boundary forcing and wind forcing.
- Evaluate the feasibility of operational application of the SWAN model for a region the size of the Southern California Bight.
- Identify and discuss special considerations for modeling waves in this (and similar) regions.

In evaluating the sensitivity of the model to these two computational shortcuts, we distinguish between two types of sensitivity: model-based and observation-based. Model-based sensitivity to a shortcut is simply taken from the difference between simulations with and without the shortcut. Observation-based sensitivity is based on comparing the errors statistics (bias, RMS error) for these two simulations, with the error statistics being based on comparison to observations. The error statistics include errors associated with other problems, for example, the accuracy of boundary forcing and wind forcing. Thus the observation-based sensitivity to a shortcut may turn out to be small if these other, less easily controlled, sources of error are large. Observation-based sensitivity instructs us on the expected immediate benefit of not using a computational shortcut, i.e. whether the added computational cost is justified. Model-based sensitivity instructs us on the expected benefit of not using a computational shortcut after other sources of error are reduced.

This paper is organized as follows: Section 2 briefly describes the SWAN model and these two computational shortcuts. Section 3 describes the Nearshore Canyon Experiment (NCEX) and a realtime wave modeling system designed to support that experiment. Section 4 describes idealized cases designed to study one of the two computational shortcuts (the stationary assumption). Section 5 presents hindcasts for the Southern California Bight, similar to the realtime wave modeling system, designed to study the two shortcuts. Discussion is given in Section 6, and Conclusions in Section 7.

2. Description of model and computational "shortcuts"

For this investigation, we used a beta version of SWAN ("Simulating WAves Nearshore"; Booij et al., 1999) which can be considered intermediate between the official versions 40.20 (released in June 2003) and 40.31 (released in February 2004). SWAN is a third generation wave action model designed to overcome traditional difficulties of applying wave action models such as WAM in coastal regions. It uses typical formulations for wave growth by wind, wave dissipation by whitecapping, and four wave nonlinear interactions ("quadruplets" or "quads"). It also includes physical processes associated with intermediate-depth and shallow water (e.g. bottom friction, depth-limited breaking).

The governing equation of SWAN and other third generation wave action models is the action balance equation. In Cartesian coordinates, this is:

$$\frac{\partial N}{\partial t}+\frac{\partial C_{g,x}N}{\partial x}+\frac{\partial C_{g,y}N}{\partial y}+\frac{\partial C_{g,\sigma}N}{\partial \sigma}+\frac{\partial C_{g,\theta}N}{\partial \theta}=\frac{S}{\sigma}. \quad (1)$$

where σ is the relative (intrinsic) frequency (the wave frequency measured from a frame of reference moving with a current, if a current exists), N is wave action density, equal to energy density divided by relative frequency ($N=E/\sigma$), θ is wave direction, C_g is the wave action propagation speed in (x, y, σ, θ) space, and S is the total of source/sink terms expressed as wave energy density. In deep water, the right hand side of Eq. (1) is dominated by three terms, $S\approx S_{\mathrm{in}}+S_{\mathrm{nl}}+S_{\mathrm{ds}}$ (input by wind, four wave nonlinear interactions, and dissipation, respectively). Source term formulations used in wave models are by no means universal, but the default formulations used in SWAN are a fair representation of the mainstream.

Boundary conditions for the outer nest SWAN models used herein are taken from nowcasts/forecasts from operational implementations of the WAVEWATCH-III model (Tolman, 1991, 2002a). [We use the word "operational" to indicate "realtime and not experimental".] Like SWAN, WAVEWATCH-III (henceforth denoted "WW3") is governed by the action balance equation. WW3 tends to be more efficient at global scales (due to resolution), whereas SWAN holds the advantage at smaller scales (e.g. grid spacing less than 2 km). Both SWAN and WW3 can be solved in either Cartesian or spherical coordinates and both are finite difference models. Because of their similarities, they complement each other nicely. [WAM, a predecessor of both SWAN and WW3, is another third generation model; it is not used in this study.]

2.1. The stationary assumption

In nonstationary applications of conditionally stable models, the time step must be small enough that a packet of wave energy does not travel a distance of more than one grid cell (or some fraction thereof) during any given time step. With the unconditionally stable nonstationary scheme of SWAN this requirement is removed, but accuracy of the scheme falls off considerably when the wave energy travels much more than 2–4 grid cells per time step (see Rogers et al., 2002). With high geographic resolution (say higher than 1/30° or 3 km), this might correspond to a time step of 5 min, or 144 time steps for each 12 h increment in a forecast, which can be computationally oppressive. Fortunately, SWAN can optionally compute using the assumption of stationarity. Computed in this manner, there are no time steps, though some iterating is required: 5–10 iterations per 12 h increment in the forecast would be typical, resulting in time saving of a factor 15–30 in this example.

However, the stationary assumption implies instantaneous wave propagation across the domain, as well as instantaneous wave response to changes in the wind field. These restrictions are

not unreasonable for a smaller domain. This is particularly true if the cross-domain wave propagation occurs at a faster rate than the change in offshore forcing at the domain's boundary. Furthermore, for these smaller areas, wave growth internal to the domain is fetch-limited, so the stationary model can represent wave growth faithfully.

However, these same restrictions are obviously inaccurate for global or basin-scale models. Even for intermediate-scale domains like the Southern California Bight, use of the stationary assumptions might lead to consistent predictions swells arriving too early and/or too rapid response to wind changes.

Since a simple phase-shift in a time series will affect RMS error but not bias, the impact of the stationary assumption should be more noticeable in the former statistic; however, because of the wind response error, the effect on bias should not necessarily be zero.

Since the stationary assumption implies an assumption of infinite duration, one might expect that local windsea in a stationary model will always be more energetic than that in a nonstationary model. However, this is not the case: the nonstationary model is affected by prior wind speeds, which may be higher than the present wind speed.

2.2. Coarse geographic resolution

The primary benefit of increased geographic resolution in the Southern California Bight is to better represent the blocking of wave energy by islands in the Bight. This blocking has a dominant impact on the wave climate at most of the coasts inside the Bight. The word "blocking" here implies that an island is completely blocking wave energy from some direction. Blocking is not the only problem associated with geographic resolution, of course: the submerged part of an island will scatter, focus, defocus, dissipate, and shoal energy. In the regional scale domains with narrow continental shelf, these effects are expected to be secondary to blocking.

3. Realtime Southern California Bight modeling system

The Nearshore Canyon Experiment (NCEX) was a coastal research project with field operations Sept. 16 through Dec. 15, 2003, located at and near La Jolla, California. A variety of instruments were deployed by scientists from several institutions to monitor the coast from water, land, and air. Quoting a University of California, San Diego press release, "NCEX is designed to determine the effects of submarine canyons and other complex seafloor formations on waves and currents. Understanding such processes is important to answer scientific questions and to address public safety issues such as rip currents." The Naval Research Laboratory (NRL) created a wave forecasting system for this region. There were several motivations:

1) Supporting the NCEX field program: to assist in planning of instrument deployment and anticipate the arrival of scientifically interesting wave conditions. This motivation is diminished somewhat by existing systems for forecasting waves in the Bight, but the NRL system is the first full application of third generation wave models to realtime forecasting of combined wind sea (generation, dissipation, propagation) and swell (dissipation, propagation) in the Bight.
2) To get "hands on" knowledge and experience with modeling waves in realtime in a challenging environment (the primary challenge being associated with the sheltering effect of islands in the Bight). This experience is valuable for future wave modeling exercises by the operational Navy.
3) The quantity of wave data in this region is probably the highest concentration anywhere in the U.S. This is of great benefit to validation and for determining sources of model errors.

3.1. System description

We created several competing wave nowcast/forecast systems for the NCEX experiment. The earliest system started producing forecasts on 26 September 2003. All systems stopped producing forecasts on or before 15 December. Since we compare different modeling methods in other sections using hindcasts, we will present only one of the competing wave nowcast/forecast systems here. Within this system, there are three SWAN grids. The second (denoted "SC2") is nested within the first (denoted "SC1") and the third (denoted "SC3") is nested within the second, SC2. The SC3 grid corresponds to the vicinity of the NCEX experiment. The three grids are shown in Fig. 1A–C. All were solved in a spherical coordinate system. Table 1 lists some details of the modeling system. [In this table, the 5-digit output locations are NDBC buoys locations; the three-digit output locations are locations of CDIP instruments (all buoys, except for 073). Some CDIP locations are referred to by three-letter identifiers, which are given in parentheses here.]

Table 1
Details of realtime system for the Southern California Bight

GRID	SC1	SC2	SC3
Δx (longitude)	2.0′ or 3087 m	0.4′ or 621 m	1.5″ or 38.9 m
Δy (latitude)	1.67′ or 3087 m	0.4′ or 741 m	2.25″ or 69.5 m
Origin (° E, ° N)	239.0, 32.0	242.2, 32.4	242.634, 32.828
# x-cells	121	121	290
# y-cells	109	181	181
Bathymetry	6″×6″	6″×6″	1.5″×2.25″
Boundary forcing	NCEP ENP	SC1	SC2
Wind forcing	NCEP ENP	NCEP ENP	None
Computation	Nonstationary	Stationary	Stationary
Execution method	Parallel (8 threads on 1.3 GHz IBM-P4)	Serial (2.4 GHz Lintel)	Serial (2.4 GHz Lintel)
Computation time	100 min	20–30 min	40 min
Output interval (h)	3	24 (effectively 12)	24 (effectively 12)
Output locations	46047, 46025, 46053, 46054, 46063, 46011, 46023, 46086, 067, 092, 028, 102, 118, 111, 107, 071	096 (DPT), 045 (OSO), 100 (TPO)	095 (PLJ), 101 (TPI), 073 (SCP)

Boundary forcing for the outer SWAN grid was taken from the NCEP (National Centers for Environmental Prediction) ENP (Eastern North Pacific) WW3 implementation (see http://polar.ncep.noaa.gov/waves/implementations.html). Realtime spectral output from that WW3 model was available from the ftp site of NCEP at two locations near the boundary of SC1, corresponding to the locations of NDBC buoys 46063 and 46047. WW3 spectra for the location of 46063 were applied to the north and west boundary of SC1; WW3 spectra for the location of 46047 were applied to the south boundary of SC1. These spectra were given in files which included recent hindcasts, the analysis period, and forecasts out to 7 days at 3 h intervals.

Wind forcing for the SWAN models were taken from fields provided by NCEP corresponding to the computational grid of the WW3 ENP model. These winds are from the NCEP Global Forecast System (GFS). As with the ENP spectra, the wind fields included forecasts out to 7 days at 3 h intervals. Global winds were used rather than those from a regional model (such as COAMPS, Hodur, 1997; Hodur et al., 2002) because of the longer forecast period.

The default bottom friction formulation of SWAN was used, though it is not expected to play a significant role in the Southern California Bight due to the relatively narrow continental shelf. For the three deepwater source terms, $S \approx S_{in}+S_{nl}+S_{ds}$, default formulations were used, except for the dissipation term, where the integer used for the weighting of relative wavenumber was increased by 1.0 (from Rogers et al., 2003 and Janssen, 1989) (this is to correct a tendency to underpredict the mean wave period of wind sea). In all three grids, 36 directional bins are used ($\Delta\theta=10°$), and 35 frequencies are used, with logarithmic spacing from 0.05 to 1.00 Hz. [To use a lowest frequency of 0.05 Hz may lead to problems when modeling the Pacific basin, so this was changed to 0.0418 for the hindcasts (Section 5).]

Due to the integration in four dimensions with computation of four-wave interactions and an implicit propagation scheme, SWAN can be computationally demanding. Each 7 day forecast computation for the SC1 grid would have taken an estimated 28 h in serial mode on the workstation used for the SC2 and SC3 computations, clearly infeasible for a realtime system. Therefore, the SC1 simulations were computed on a parallel computing platform, utilizing the OpenMP modifications of the code made by Campbell et al. (2002). Each seven-day SC1 forecast typically required 100 min of computation time on that platform. To our knowledge, this represents the first use of the OpenMP capabilities of SWAN for realtime forecasting, and is a strong demonstration of the expanded utility of the SWAN model for such purposes.

Fields of wave height and peak direction were output for graphical display on a web site. Wave spectra were saved at locations where NDBC and CDIP instruments were deployed. The system was launched every 12 h. Individual SC2 and SC3 simulations produced output at 24 h intervals—which is a fairly coarse interval—due to computation time constraints. Since the system ran every 12 h, there was SC2 and SC3 output at staggered 12 h intervals: still a coarse interval, but better than 24 h. In the case of the SC1 model, the output interval has only a very slight effect on computation time (usage of disk space is the greater constraint), so a 3 h output interval was used. The SC1 model used a 5 min time step for computations. All three SWAN models produced output out to 7 days.

For the SC1 and SC2 grids, a 6″ bathymetry provided by Dr. W.C. O'Reilly (Scripps Institution of Oceanography, "SIO") was used. For the SC3 grids, bathymetry provided on the SIO NCEX website was used. The latter bathymetry data set was developed specifically for the NCEX experiment.

3.2. Results

For realtime comparison, CDIP data at the three SC3 instrument locations were downloaded during every modeling cycle and plotted along with time series of wave height, peak period, and mean direction from the SC3 model at those locations. An example time series plot similar to the ones displayed on the web page is shown in Fig. 2. Plots of fields of wave height and direction for each of the three grids for various forecast times were also shown on the webpage, but are not reproduced here.

Calculations of error—with NDBC and CDIP data as ground truth—are given in Table 2. All dates are in 2003. The bias and root-mean-square error "RMSE" are calculated over the time interval shown, which varies due to inconsistent archiving of model output and data outages. The error metrics are calculated for the analyses of each realtime SWAN simulation (error metrics for the forecasts are not reported here, due to limitations on space). In the table, we organize the instruments locations into three groups, in order to better detect any correlation of error and location. The three groups are a) locations relatively unsheltered from swells from the open ocean, b) locations along the northern shoreline of the Bight, and c) locations that fall within the SC2 and SC3 grids. We give averages for each grouping and also an average of all locations. In the averaging, each location is weighted equally even though the duration of time intervals for comparisons are different in many cases. When comparing model output to data, we pass the data through a three-hour running-average type filter.

To put these numbers in context, the magnitude of bias of analyses of global wave models (at any given location) tend to be 0.15 m or lower and RMS errors tend to be 0.4–0.6 m (e.g. Tolman, 2002b). For energetic, but enclosed areas (e.g. the Great Lakes), 0.18 m RMS error and negligible bias is possible in blindfold hindcasts.

4. Idealized case: impact of the stationary assumption

In this section, we present idealized cases. The strategy is to create simplified model scenarios so that we can—without excessive runtimes—test the effect of the stationary assumption. Using these tests, this source of error is isolated from other sources of error. Also, by including a test case for another environment (the Gulf of Maine during a similar time period), insight is gained regarding how the error might vary with climate.

Actual buoy data are used in the design of these idealized cases. In canonical tests which follow, the results are sensitive to the time scale of variation of input. The boundary-forced case is sensitive to the group velocity of energy parcels. The wind-

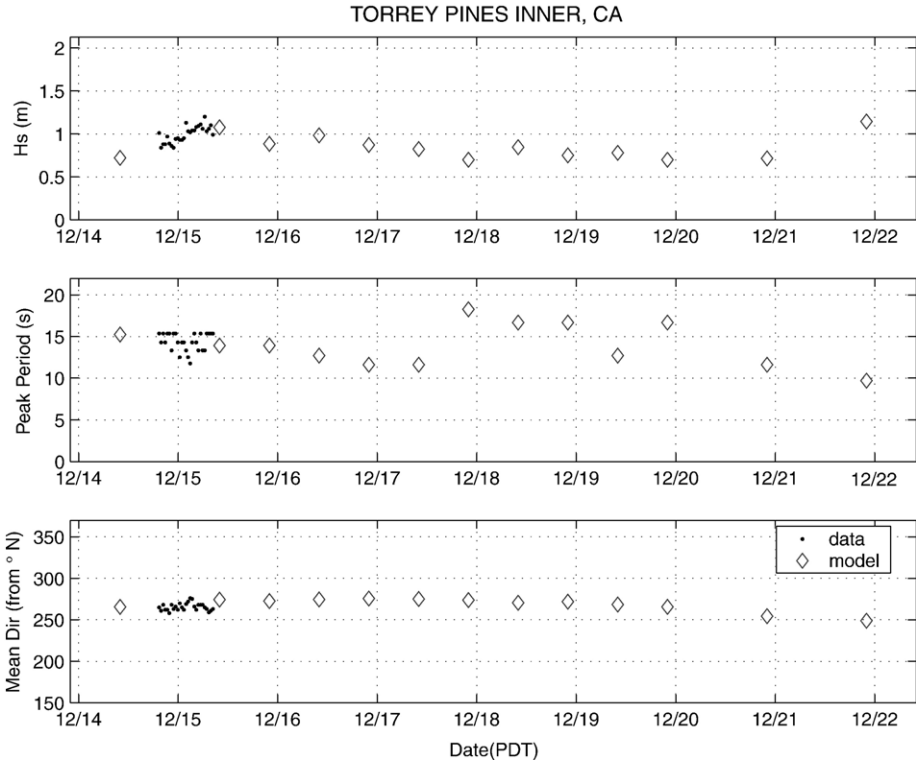

Fig. 2. Comparison of near-realtime CDIP data to model analyses and model forecasts at location TPI (CDIP 101). This plot is essentially the same as a plot that was made by the realtime system at 0938 PDT 15 December (the realtime plot included output from two models that are not described in this paper and are therefore not shown here).

forced case is sensitive to the magnitude of the wind. Thus, we want the input to be as realistic as possible. This is the primary motivation for using buoy data for forcing.

4.1. Idealized tests

Simplified long-duration simulations using stationary computations are conducted to study the impact of the stationary assumption. Wave height root mean square (RMS) error is calculated over the entire model domain, using simulations with nonstationary computations for "ground truth". For forcing, we use actual buoy data. Thus with these tests, we get an estimate of the typical levels of error under realistic forcing conditions. The time period of the NCEX experiment is used. To get an idea of the impact of local wave climate, we conduct tests using the climate for the Gulf of Maine as well as the Southern California Bight. Characteristics common to all idealized simulations regarding effect of stationary assumption are:

1) One dimensional simulations
2) Domain size = 300 km
3) $\Delta x = 0.6$ km
4) Directional resolution = 10°
5) 33 frequencies in logarithmic distribution
6) Deep water
7) Whitecapping identical to that used in other simulations in this study
8) Wave height output over the entire model domain, every 3 h, from 1800 UTC 14 October 2003 to through 2100 UTC 15 December 2003.

Characteristics common to all the idealized simulations with stationary computation:

1) 15 iterations per computation
2) One computation every 3 h simulated.

Characteristics common to all the idealized simulations with nonstationary computation:

1) Initialized at state of rest at 1800 UTC 13 October 2003.
2) $\Delta t = 2.5$ min

4.2. Idealized wind-forced test

Characteristics common to all the wind-forced idealized simulations:

1) These simulations used winds taken from buoy data, converted to 10 m elevation. The scalar buoy wind speed is used for the along-axis wind speed, U_x. The cross-axis wind speed U_y is set to zero. The sign of U_x is preserved, so wind direction here is binary (westerly or easterly)
2) Homogeneous winds are used.
3) No boundary forcing is used.

Table 2
Error calculations of realtime model results vs. measurements

	Bias (m)	RMSE (m)	Begin time	End time
Open locations				
B46063	0.06	0.34	0600 UTC 21-Oct	0600 UTC 16-Dec
B46054	0.12	0.30	0600 UTC 21-Oct	0600 UTC 16-Dec
B46023	0.01	0.38	1800 UTC 14-Nov	0600 UTC 16-Dec
B071	0.14	0.40	1800 UTC 14-Nov	0600 UTC 16-Dec
B46047	0.03	0.41	0600 UTC 21-Oct	0600 UTC 16-Dec
B46086	0.05	0.32	1800 UTC 14-Nov	0600 UTC 16-Dec
B067	0.16	0.42	1800 UTC 14-Nov	0600 UTC 16-Dec
average	0.08	0.37		
"North Shore" locations				
B107	−0.11	0.24	1800 UTC 14-Nov	0600 UTC 16-Dec
B46053	−0.03	0.22	0600 UTC 21-Oct	0600 UTC 16-Dec
B102	0.07	0.19	1800 UTC 14-Nov	0600 UTC 16-Dec
B111	0.10	0.21	1800 UTC 14-Nov	0600 UTC 16-Dec
B092	−0.05	0.28	1800 UTC 14-Nov	0600 UTC 16-Dec
B46025	−0.03	0.32	0600 UTC 21-Oct	0600 UTC 16-Dec
B028	0.01	0.25	1800 UTC 14-Nov	0600 UTC 16-Dec
average	−0.01	0.24		
Locations inside SC2 and SC3				
DPT	0.01	0.25	0600 UTC 21-Oct	0600 UTC 16-Dec
OSO	−0.04	0.21	0600 UTC 21-Oct	0600 UTC 16-Dec
TPO	0.05	0.21	0600 UTC 21-Oct	0600 UTC 16-Dec
PLJ	−0.04	0.22	0600 UTC 22-Oct	0600 UTC 16-Dec
TPI	−0.01	0.19	0600 UTC 22-Oct	0600 UTC 16-Dec
SCP	0.08	0.21	0600 UTC 22-Oct	0600 UTC 16-Dec
average	0.01	0.21		
All locations				
average	0.03	0.28		

In this canonical wind-forced case, we are treating the region in question (the Southern California Bight or the Gulf of Maine) as a rectangular lake of arbitrary north–south dimension, and east–west dimension comparable to that of the actual region. All energy is generated internally, as opposed to the actual situation where local winds add energy to that coming in through the boundaries. The wind-forced idealized simulations with nonstationary computation included linear wave growth physics, so that simulations could start from rest.

For the Southern California Bight idealized wind-forced test, 46047 buoy data is used. In this canonical case, west-to-east airflow is predominant. For the Gulf of Maine idealized wind-forced test, 44005 buoy data is used.

4.3. Idealized boundary-forced test

Characteristics common to all the boundary-forced idealized simulations:

1) No wind forcing is used.
2) Nonlinear interactions are disabled.
3) A JONSWAP spectrum with peak enhancement factor of 3.3 (see Holthuijsen et al., 2003) is used for boundary forcing along the open-ocean side of grid.

In the stationary results, the only cause of x-wise variation (not shown) is dissipation; whereas the nonstationary model results (properly) also vary due to time-history of the boundary forcing.

Note that if all energy is traveling the same speed c, and if $\frac{\partial H_{x=0}}{\partial t}$ is constant, then the wave height error E_H at any point x is $E_H = \frac{\partial H_{x=0}}{\partial t}\frac{x}{c}$.

4.3.1. Southern California Bight case

Boundary forcing is based on measurements with CDIP buoy 71 (the "Harvest" buoy) during the time period 0000 UTC 1 October 2003–0000 UTC 1 January 2004. During the infrequent gaps in data from this buoy, spectra are taken from CDIP buoy 67 (the "San Nicholas Island" buoy). The spectra used here are described at a 3 h interval, determined by 3 h moving average of data provided by CDIP, which are described at a 0.5 h interval. These 3-hour interval combined wave spectra are denoted in this paper as "CDIP/071/067". [We use CDIP data rather than NDBC data here due to directional information in the CDIP data.]. Time series of wave height, peak period, and directional spreading are calculated from these spectra. However, the mean direction is always "from west". Wave conditions are passed to the SWAN model in this parameterized form.

4.3.2. Gulf of Maine case

Time series of wave height, peak period are taken from data for NDBC buoy 44005. However, the mean direction is always "from east" and directional spreading (see Holthuijsen et al., 2003) is always 47.7° (taken from the mean of the Southern California Bight case).

4.4. Results

Results are shown in Fig. 3A–D. In these four figures, there are four curves and each is presented twice: Fig. 3A,B contrast the different forcing set (boundary forcing vs. wind-forcing); Fig. 3C,D contrast the different climates (Southern California Bight vs. Gulf of Maine).

In the wind-forced canonical cases, if one inspects individual cases where the wind shifts directions, the stationary model is especially inaccurate because it responds too quickly to the shift, creating new energy and destroying old. Additionally, the wind speeds reach greater extremes over shorter time intervals in the Gulf of Maine case than in the Southern California Bight case; this variability would not be represented particularly well in the stationary runs.

Of the boundary-forced canonical cases, error is worse in the Gulf of Maine case. This may be simply due to larger wave heights (all else being equal, RMS error will tend to be greater in more energetic wave climates).This may also be affected by the travel speed of wave energy (the Gulf of Maine tends to experience shorter waves, Fig. 4, which will tend to be less well represented by the assumption of instantaneous propagation).

5. Hindcasts

Here we build on what was learned in the idealized simulations using long term hindcasts comparable to the realtime

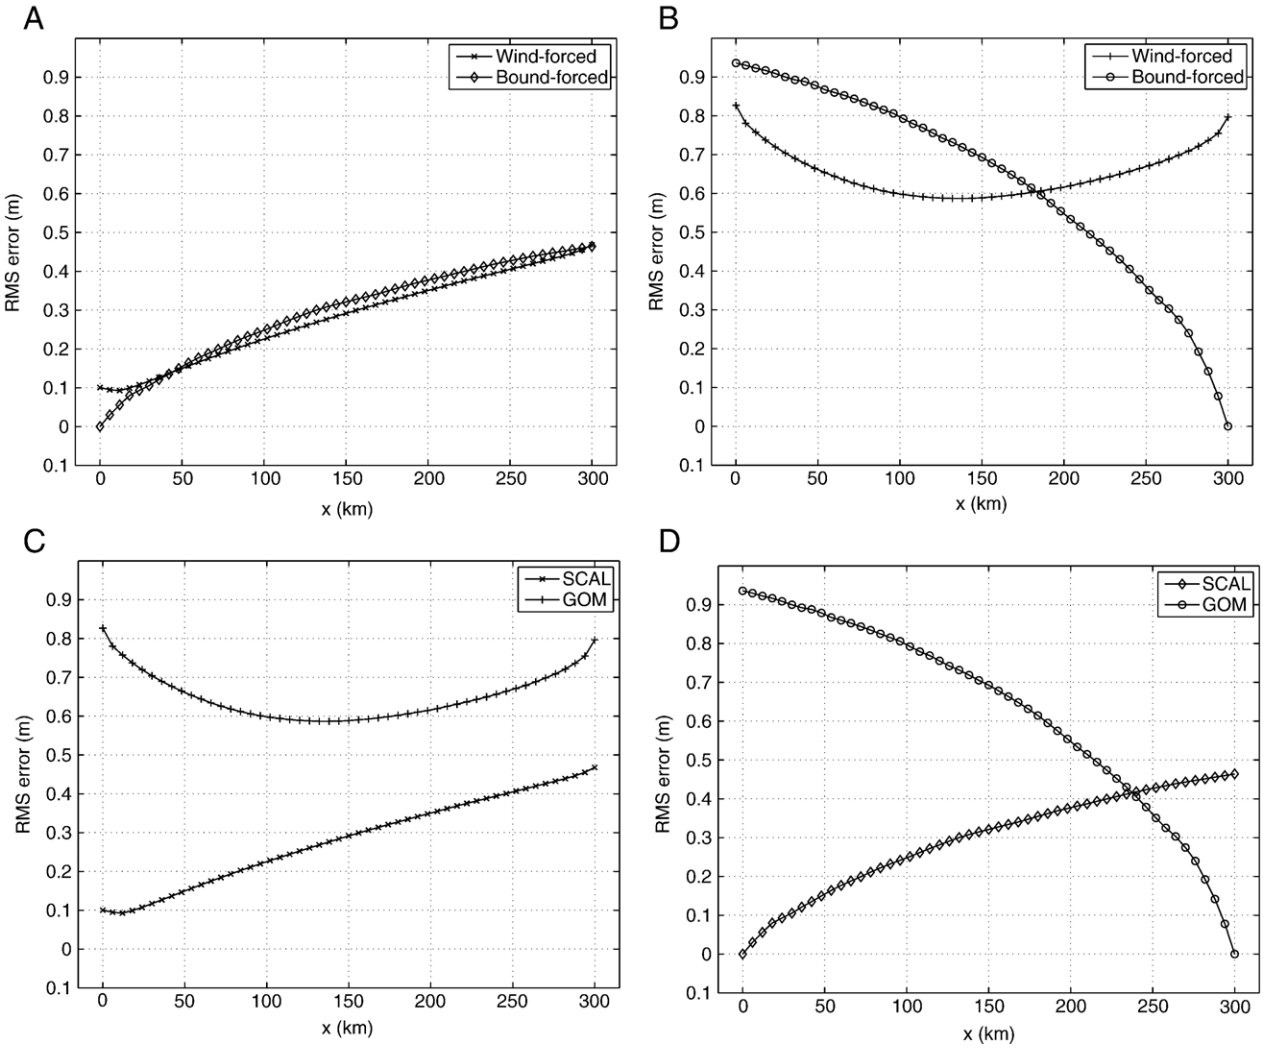

Fig. 3. A.Wave height RMS error computations for the wind-forced idealized simulation with stationary computations (simulation with nonstationary computations are taken as ground truth). Cases with forcing corresponding to the Southern California Bight are shown. B. Cases with forcing corresponding to the Gulf of Maine. C. Cases with wind forcing. D. Cases with boundary forcing.

system. With the hindcast mode, we have a few advantages over the realtime mode:

1) In hindcast mode, computation time is not a major constraint on model design.
2) The realtime system was subject to problems with forcing arriving late and having to use forecast winds as analyses.
3) In hindcast mode, we can test/estimate the accuracy of various forcing methods (including forcing with observations, which wouldn't be possible in a forecast system) and choose one.
4) In hindcast mode, we can pay more careful attention to numerical issues, etc.
5) Some settings (such as the garden sprinkler effect correction [Booij and Holthuijsen, 1987]) were changed midway during the lifetime of the realtime system. With the hindcast system, settings are uniform for the duration, leading to more meaningful comparisons.

5.1. Hindcast descriptions

The hindcasts are designed to investigate the practical effect of two computational "shortcuts":

1) Use of stationary computations for the SC1 region
2) Use of coarse geographic resolution for the SC1 region

These are the two "shortcuts" described in Section 2.

5.1.1. Forcing

For forcing on the west boundary, we use the CDIP/071/067 spectral time series described in Section 4.3.1. For forcing on the south boundary, we use analyses for the NCEP ENP WW3 implementation corresponding to location 46047; this is identical to what we used for forcing the realtime system at this boundary. For wind forcing, we use the NCEP GFS winds, identical to wind forcing of the realtime system.

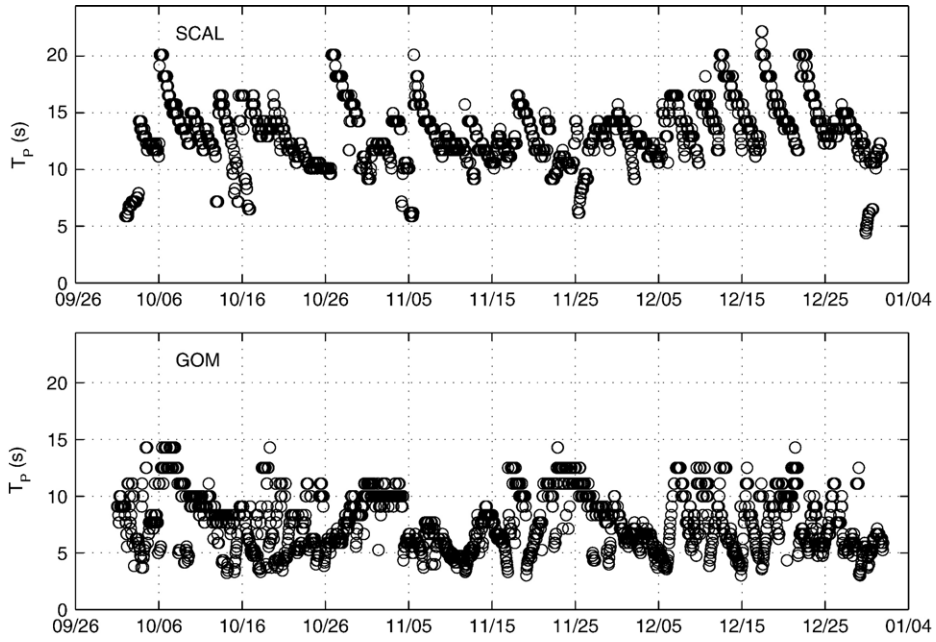

Fig. 4. Peak period of boundary forcing for the idealized cases. Southern California Bight case (upper panel) and Gulf of Maine case (lower panel).

5.1.2. Model settings

We perform the hindcast with the outer grid SC1 at two different resolutions and two different computation methods (stationary and nonstationary), so there are four hindcast simulations, denoted as

- "STAT LR", with the SC1 grid calculated with stationary computations, at relatively low resolution
- "STAT HR", with the SC1 grid calculated with stationary computations, at relatively high resolution
- "NONS LR", with the SC1 grid calculated with nonstationary (time-stepping) computations, at relatively low resolution
- "NONS HR", with the SC1 grid calculated with nonstationary (time-stepping) computations, at relatively high resolution

For these four separate hindcasts, only the operation of the outer SWAN nest (SC1) is different. So, the resolution of the "SC2" nest for the "STAT LR" hindcast is the same as the resolution used for the "SC2" nest for the "STAT HR" hindcast. Similarly, all "SC2" and "SC3" nests are calculated using stationary computations. Specific settings are as follows:

- In all cases 36 directional bins ($\Delta\theta = 10°$), and 34 frequencies are used, with logarithmic spacing from 0.0418 to 1.00 Hz. (note that the lower frequency is changed from the realtime system).
- In all cases, the dissipation settings are identical to those used for the realtime system.
- In the "low resolution" SC1 simulations, $\Delta x = \Delta y = 1/20°$. This is identical to the geographic resolution of a NAVO WAM implementation for the Southern California Bight.
- In the "high resolution" SC1 simulation, $\Delta x = \Delta y = 1'$.
- In the SC2 nests, $\Delta x = 0.3'$ (longitude) and $\Delta y = 0.4'$ (latitude).
- In the SC3 nests, $\Delta x = 1.5''$ (longitude) and $\Delta y = 2.25''$ (latitude). (unchanged from the realtime system.)
- For all stationary computations (which includes all SC2 and SC3 computations), computation at every time interval used a

Table 3
Mean error statistics for the hindcast simulations, with the "NONS, HR" hindcast as "ground truth"

Model	All locations		Open areas		North Shore		SC2 and SC3	
	Bias (m)	RMSE (m)	Bias (m)	RMSE (m)	Bias (m)	RMSE (m)	Bias (m)	RMSE (m)
STAT LR	−0.07	0.16	−0.04	0.14	−0.12	0.20	−0.04	0.13
STAT HR	−0.02	0.14	−0.01	0.14	−0.04	0.16	+0.00	0.12
NONS LR	−0.03	0.04	−0.01	0.02	−0.03	0.06	−0.04	0.05

Table 4
Time period of comparisons to observations

Location	Begin time	End time
NDBC/46047	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/067	18:00 14-Oct-2003	09:00 15-Dec-2003
NDBC/46086	18:00 10-Nov-2003	21:00 15-Dec-2003
NDBC/46023	18:00 14-Oct-2003	21:00 15-Dec-2003
NDBC/46063	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/071	18:00 14-Oct-2003	21:00 15-Dec-2003
NDBC/46054	18:00 14-Oct-2003	21:00 15-Dec-2003
NDBC/46053	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/107	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/111	18:00 14-Oct-2003	21:00 15-Dec-2003
NDBC/46025	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/102	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/028	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/092	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/DPT	15:00 15-Oct-2003	21:00 15-Dec-2003
CDIP/OSO	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/TPO	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/PLJ	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/TPI	18:00 14-Oct-2003	21:00 15-Dec-2003
CDIP/SCP	18:00 14-Oct-2003	21:00 15-Dec-2003

Table 5
Mean error statistics for the hindcast simulations

Model	All locations		Open areas		North Shore		SC2 and SC3	
	Bias (m)	RMSE (m)	Bias (m)	RMSE (m)	Bias (m)	RMSE (m)	Bias (m)	RMSE (m)
STAT LR	−0.04	0.28	−0.12	0.31	−0.06	0.26	+0.06	0.26
STAT HR	+0.01	0.28	−0.09	0.31	+0.02	0.25	+0.11	0.27
NONS LR	−0.00	0.24	−0.09	0.27	+0.03	0.23	+0.07	0.22
NONS HR	+0.03	0.24	−0.08	0.26	+0.06	0.22	+0.10	0.22

low energy wave condition as the first guess. This leads to slower convergence (increased computation time), but allows us to reproduce computations for specific time periods precisely (useful for detailed investigations) and prevents problems with "drift" in solution that may occur when large numbers of stationary computations are performed in sequence, with each using the prior computation as the first guess (see Section 6, Discussion).

- For all stationary computations, default numerical settings are used.
- For nonstationary computations (SC1 only), the default second order propagation scheme is used, with a "garden sprinkler correction wave age" of 2.0 h (see Holthuijsen et al., 2003).
- Time step sizes for the nonstationary (SC1) cases: "NONS HR": a time step of 2.5 min is used, dictated by the garden sprinkler correction scheme; "NONS LR": a time step of 6.0 min is used.
- Wind forcing, wind sea growth, and four-wave nonlinear interactions are not included in the SC3 computations.

5.2. Model sensitivity to computational shortcuts

In this section, model-based sensitivity to the two computational shortcuts is estimated by use of the simulation with nonstationary, high resolution outer nest "NONS HR" as "ground truth". Bias and RMS error computed in this manner are given in Table 3. The results suggest the following:

1) Both computational shortcuts lead to slight negative bias. A decrease in energy reaching the nearshore areas is not an expected side effect of the stationary assumption. Differences in numerics (e.g. diffusion) may be the cause, but the variability in the offshore wave climate would make bias associated with numerics less likely.
2) Biases from the two shortcuts combine in a clearly nonlinear fashion. Combined, the effect is significant in some places (−12 cm bias for the locations on the north shore of the Bight).
3) As expected, the stationary assumption has a greater impact on RMS error than bias: RMS error is greater than 12 cm at all four location groupings. This is presumably due to phase shift in swell time series (i.e. error in arrival time with the stationary assumption). RMS error associated with coarse resolution is not large (less than 7 cm at all four location groupings).

Table 6
Bias (m) for each location and each hindcast simulation

Location	STAT LR	STAT HR	NONS LR	NONS HR
NDBC/46047	−0.21	−0.20	−0.19	−0.19
CDIP/067	−0.03	−0.02	−0.01	−0.00
NDBC/46086	−0.14	−0.01	−0.10	−0.02
NDBC/46023	−0.07	−0.07	−0.06	−0.05
NDBC/46063	−0.16	−0.14	−0.13	−0.13
CDIP/071	−0.11	−0.10	−0.08	−0.08
NDBC/46054	−0.11	−0.11	−0.08	−0.08
NDBC/46053	−0.23	−0.09	−0.06	−0.02
CDIP/107	−0.23	−0.09	−0.11	−0.03
CDIP/111	0.01	0.02	0.12	0.08
NDBC/46025	0.00	0.03	0.08	0.08
CDIP/102	0.03	0.20	0.08	0.18
CDIP/028	0.07	0.10	0.11	0.12
CDIP/092	−0.06	−0.01	−0.02	0.02
CDIP/DPT	−0.02	0.07	−0.02	0.07
CDIP/OSO	0.04	0.11	0.07	0.09
CDIP/TPO	0.09	0.11	0.09	0.12
CDIP/PLJ	0.06	0.07	0.05	0.08
CDIP/TPI	0.04	0.12	0.06	0.10
CDIP/SCP	0.17	0.16	0.16	0.16

Table 7
Root-mean-square error (m) for each location and each hindcast simulation

Location	STAT LR	STAT HR	NONS LR	NONS HR
NDBC/46047	0.55	0.55	0.45	0.45
CDIP/067	0.35	0.35	0.30	0.30
NDBC/46086	0.39	0.37	0.28	0.25
NDBC/46023	0.26	0.26	0.25	0.25
NDBC/46063	0.28	0.27	0.26	0.25
CDIP/071	0.14	0.13	0.13	0.13
NDBC/46054	0.24	0.24	0.21	0.20
NDBC/46053	0.33	0.25	0.22	0.21
CDIP/107	0.31	0.22	0.23	0.20
CDIP/111	0.26	0.25	0.25	0.23
NDBC/46025	0.25	0.25	0.24	0.24
CDIP/102	0.21	0.30	0.20	0.26
CDIP/028	0.23	0.24	0.22	0.21
CDIP/092	0.26	0.25	0.23	0.22
CDIP/DPT	0.29	0.28	0.28	0.26
CDIP/OSO	0.23	0.25	0.19	0.20
CDIP/TPO	0.26	0.26	0.21	0.21
CDIP/PLJ	0.26	0.27	0.21	0.21
CDIP/TPI	0.23	0.26	0.18	0.19
CDIP/SCP	0.29	0.28	0.24	0.24

Table 8
Bias (m) for "partial wave height" for four frequency bands, for the "NONS, HR" hindcast

	Frequency band			
Locat.	0.04 to 0.08 Hz	0.08 to 0.12 Hz	0.12 to 0.15 Hz	0.15 to 0.35 Hz
B46023	−0.15	−0.03	+0.03	+0.12
B46047	−0.23	−0.12	−0.03	+0.04
B46086	−0.06	+0.02	+0.00	+0.04
OSO	+0.06	+0.06	−0.02	+0.05
TPO	+0.03	+0.10	+0.02	+0.07
TPI	+0.03	+0.07	+0.01	+0.07

Table 9
Root-mean-square error (m) for "partial wave height" for four frequency bands, for the "NONS, HR" hindcast

	Frequency band			
Locat.	0.04 to 0.08 Hz	0.08 to 0.12 Hz	0.12 to 0.15 Hz	0.15 to 0.35 Hz
B46023	0.25	0.16	0.11	0.18
B46047	0.38	0.29	0.18	0.22
B46086	0.25	0.20	0.14	0.18
OSO	0.11	0.11	0.09	0.17
TPO	0.12	0.14	0.10	0.19
TPI	0.09	0.11	0.09	0.18

5.3. Comparison of hindcasts to observations

Table 4 lists the time period used for error statistic calculations for the comparison of the hindcasts to observations. Tables 5–9 summarize the results from the hindcast simulations. The geographic grouping and averaging is the same as described in Section 3. Fig. 5 shows the geographical distribution of error of the NONS HR case. Tables 8 and 9 indicate the bias calculated from time series of "partial wave height" calculated for four frequency bands, at six locations for the NONS HR hindcast. The "partial wave height" is calculated from the variance (i.e. energy) of the wave spectrum over a frequency range defined by lower and upper bounds f_1 and f_2: $H_{m0,\mathrm{partial}} = 4\sqrt{\upsilon_{\mathrm{partial}}}$ and $\upsilon_{\mathrm{partial}} = \int_{f_1}^{f_2} F(f)\mathrm{d}f$, the "partial variance".

From these results, we can make the following observations:

1) For the hindcasts, bias tends to be negative at the open locations and positive at the SC2/3 locations. One might speculate that some problem with the SWAN implementation is allowing too much energy through the islands, reaching the NCEX area. (This speculation turns out to be incorrect, see below.)
2) The bias patterns observed in the hindcast system are generally quite different from those observed in the realtime system. For example, at the offshore locations, the realtime model has a positive bias and the hindcast models have negative bias; this is directly attributable to differences in bias in wave forcing at the west boundary.
3) For the open locations, RMS error is much lower in the hindcast (NONS HR) simulation than in the realtime system. This probably reflects a benefit from using measured buoy spectra for forcing on the west boundary.
4) For the "north shore" and "SC2/3" locations, results from the hindcast (NONS HR) are more energetic than results from the realtime systems. Since bias tends to be positive at these locations, this means that bias is worse in the hindcasts than in the blindfold realtime system, which is not expected.
5) Use of nonstationary computations significantly improved RMS error. This is expected, since arrival times will be more accurately predicted without the stationary assumption, and duration-limited windsea generation may be more accurate.
6) Use of higher resolution in the outer grid (SC1) does not have a significant effect on RMSE. Thus, from these statistics, one can conclude that other types of errors, such as boundary forcing errors, need to be reduced in order to see practical benefit from high geographic resolution.
7) At the lowest frequency band the hindcast bias is negative at open locations (as high as—23 cm at buoy 46047) and positive on the east shore of the Bight (3 to 6 cm). This clearly suggests that swell energy transmittance through the islands is overpredicted.

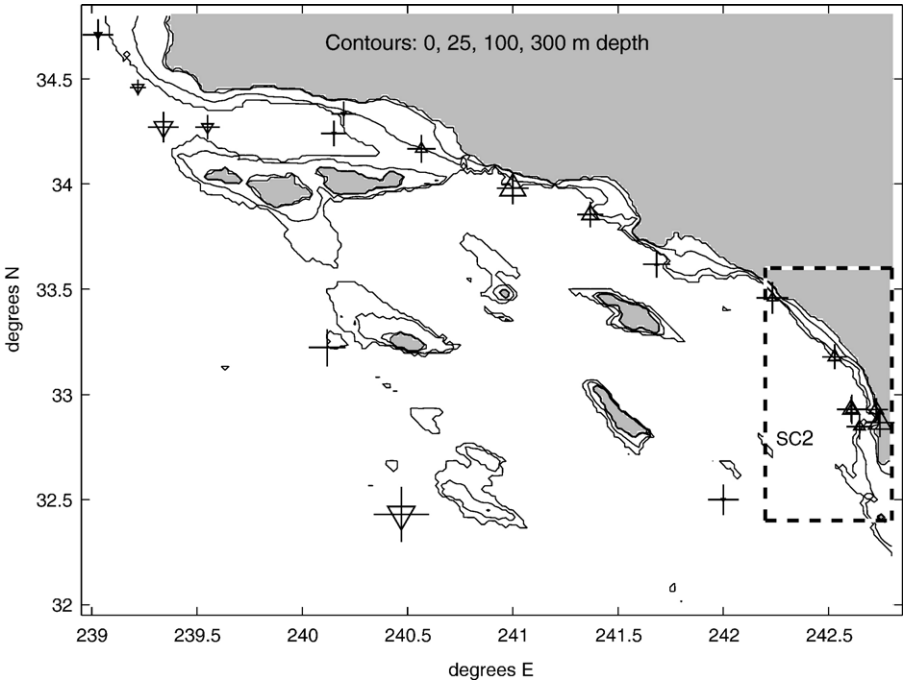

Fig. 5. Bias and root-mean-square error for the "NONS HR" case. The size of symbols indicates magnitude of bias (triangles) and root-mean-square error (plusses). The orientation of the triangles indicates sign of bias. [For the numeric values, see Tables 6 and 7.]

8) In the highest frequency band, the hindcast bias is positive at all six locations for which this calculation was made. Prior experience with this model in similar wind speed regimes—see discussion of the whitecapping term in Section 3.1—suggests that this bias is unlikely to be due to wave model physics. Taken together, this suggests that, in the atmospheric forcing, the wind speed is overpredicted for the low to moderate wind speed events.

From careful study of specific cases, we can make the following additional conclusions:

1) As mentioned in (1) above, bias patterns suggest a scenario in which swell forcing of the SWAN models is too low and the amount of swell energy getting through islands is too high. This is an oversimplification however. In fact, there exists at least one case during the hindcast where too much energy is getting through because swell from southwest is overpredicted (compensated at the Open Areas by an underprediction of swell from the northwest). To put this another way: the geography of the Bight is such that energy from the southwest will tend to reach the NCEX (SC3) area, whereas energy from the northwest tends to be blocked more before reaching this area. If a model uses forcing which overpredicts swells from the southwest and underpredicts swells from the northwest, then the total energy at the boundary may be well predicted (due to balancing of errors), while the total energy that the NCEX area will be overpredicted. This is demonstrated in Fig. 6. Here, the upper panel shows a time series of total wave height in the boundary forcing (input to SC1); based on this, the forcing appears fairly accurate. The center panel shows the same time series, except only the low frequency energy from the southwest is included in the integration to calculate wave height; in this comparison, the forcing is too high. The lower panel shows the wave height prediction near the NCEX location (output from SC3); this shows a clear overprediction, which is at least partially attributable to the overprediction in the forcing (center panel).

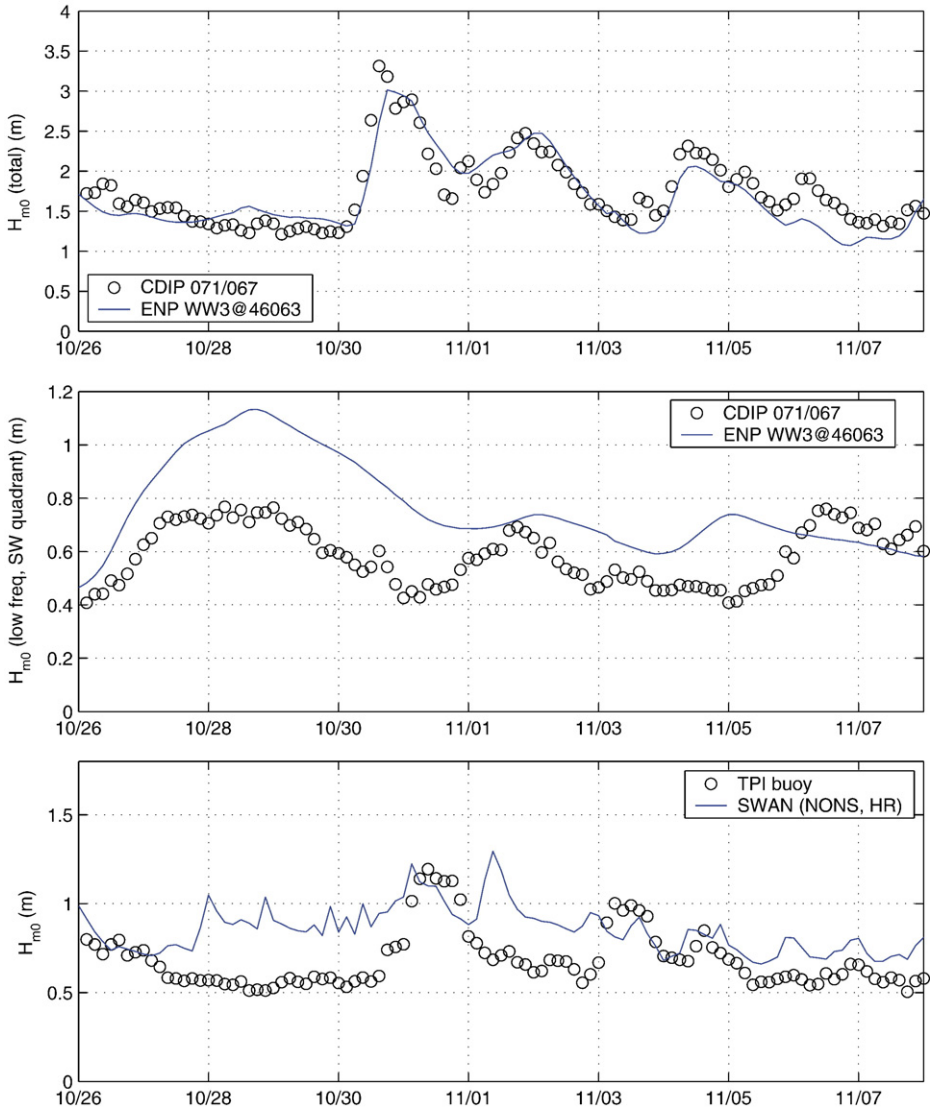

Fig. 6. Wave height time series for 26 October–7 November. See text for explanation.

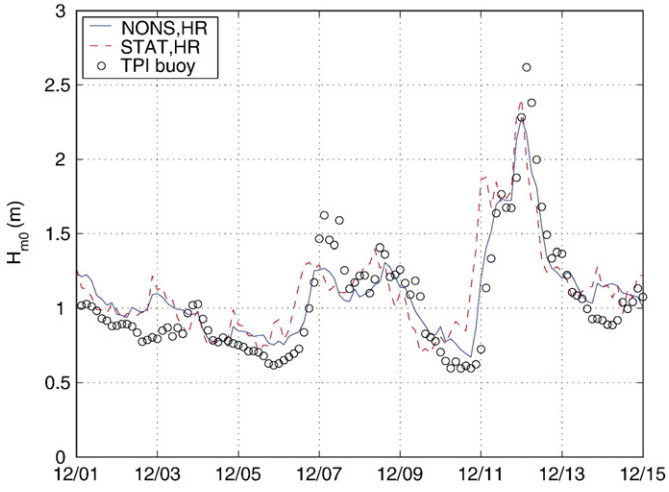

Fig. 7. Wave height time series at TPI buoy, for 1 December–16 December.

2) Use of higher resolution in the outer grid SC1 (for the purpose of better representing sheltering by islands) increases the energy level at the SC2/3 locations. This is probably due to the following: when a neighboring grid cell is a land cell, this tends to block energy traveling parallel to the coast; with a coarser grid, there are fewer wet grid points across a constriction, so blocking by neighboring land points is increased.
3) By inspection of time series, we can confirm that the lower RMS error with nonstationary computations (see (6) above) is due to better predictions of arrival/departure times of swell events. An example of this is shown in Fig. 7 (a time series from the SC3 grid, near the NCEX area).

6. Discussion

6.1. Output interval

SC2 and SC3 grids were run on workstation in serial mode, with a very large time interval between computations (24 h intervals between stationary computations, 12 h effective output interval due to staggering). Of course, this interval could have been more frequent, had we run these nests using OpenMP, with more processors. The large interval has no effect on error metrics, but when output is plotted as a time series, it would cause the SWAN output to appear excessively smooth relative to that from a model with higher temporal resolution.

6.2. Blending model and buoy spectra for southern boundary forcing

For the southern boundary of the hindcast region, non-directional buoy data (46047) were available, but not directional buoy data. A strategy of blending model and buoy spectra for boundary forcing was tested. This worked as follows: 1) at each time interval and at frequency in the model (WW3 ENP), a normalized directional spectrum $D(\theta)$ was calculated (this function integrates to unity), 2) Non-directional buoy data $E(f)$ within 1 h of the time interval were time-averaged, 3) The spectra were interpolated across frequencies to put them on similar frequency grids (38 frequencies, linearly spaced from 0.03 to 0.4 Hz), and 4) The dimensional forcing spectrum was calculated as $S(f,\theta)=D(\theta)E(f)$. However, this approach was prone to a specific problem: if a frequency band is energetic in the buoy measurement (say strong swell from northwest), but not energetic in the WW3 spectra (for instance, if the latter contains weak swell from northwest and weak swell from southwest), the result of combination is a spectrum with fairly strong swell from northwest and southwest. The swell from southwest results in overprediction of energy at the sheltered, coastal locations. For practical purposes, this approach is less accurate than using unmodified WW3 spectra for forcing (as was done for the southern boundary in the hindcasts presented above). We still feel that this method (blended boundary forcing) holds promise, but we strongly recommend quality control procedures to ensure that the blending is only performed at times at which the measured $E(f)$ is similar to the modeled (WW3) $E(f)$: at other times, the methodology should default to simple modeled (WW3) two-dimensional spectra.

6.3. Sensitivity to directional characteristics of boundary forcing

Based on specific studies of the hindcast results, it is apparent that the directional characteristics of boundary forcing play a dominant role in the predictions of energy levels at the NCEX area. Unfortunately, the spectra used for forcing on the southern boundary have the usual limitations of a global wave model, and the spectra used for forcing on the western boundary are subject to the limitations of what a buoy measures (a truncated Fourier series describing the directional distribution at each frequency) and the Maximum Likelihood Method. Problems with wave direction are likely to produce random errors in predictions at the NCEX site, whereas problems with consistent overprediction of directional spreading or smoothness of peaks could conceivably lead to bias at the NCEX area (too much or too little energy propagating past the islands). The ability of MLM to correctly reproduce multiple swells arriving from different directions at the same frequency is suspect. Also, finite directional resolution will tend to make accurate propagation through narrow channels difficult for a model.

6.4. Wave prediction error due to errors in bathymetry

One probable cause for error is the uncertainty in the bathymetry, particularly over the canyon and particularly for northwesterly waves. The bathymetric database used is comprised of data ranging from National Ocean Survey (NOS) data, to more recent ship surveys over the canyon, to nearshore surveys from airborne lidar from the US Geological Survey (USGS) measurement system. Each has different coverage and quality. Kaihatu and O'Reilly (2002) performed some sensitivity studies for model runs over various bathymetric databases and demonstrated significant sensitivity of the nearshore waveheights on the details of the canyon bathymetry. Additionally, Long et al. (2004) showed that modeled nearshore wave and circulation fields were strongly dependent on the details of the canyon, particularly the

crenellations of the depth contours. The biases seen in Table 2 are generally lower than the above estimates, however, and may indicate that the wave model is less sensitive to bathymetry errors than other types of errors, at least over the shelf. We do not specifically investigate sensitivity to bathymetric errors herein; to do this properly, it would be necessary to adjust the bathymetry while keeping the resolution fixed.

6.5. Underconvergence

While performing stationary, hindcast simulations, with forcing identical to the forcing of the stationary realtime system (not presented here), it was found that positive bias occurs which did not occur with the realtime system. By default, SWAN stationary computations use the prior stationary computation as a "first guess" for the iterative solution procedure. With many stationary computations in sequence (as with the hindcasts), there is a subtle increase in energy. This tendency was overridden in the hindcasts presented by initializing each stationary computation with a low energy condition as the "first guess", as mentioned above in Section 5. The problem with under-convergence was confirmed by running two hindcasts, identical except for method of initializing computations (at each time interval): dramatic difference in bias occurs. Note that we have not proven that hindcast results are fully converged; it is entirely possible that our method of initialization leads to energy levels that are artificially slightly low. The reader is referred to Zijlema and van der Westhuysen (2005) for further reading.

6.6. Dissipation

At specific times during the hindcasts (e.g. 25 November, 2003), there is significant local wind sea generated inside, and just west of, the SC1 grid. Energy from this wind sea is well predicted outside the islands of the Bight, but is overpredicted at the NCEX region. Unfortunately, it is very difficult to determine whether this overprediction is due to a) not enough energy being blocked by bathymetry/topography or b) not enough dissipation of these relatively short waves as they propagate from west to east across the grid (i.e. deficiency in the S_{ds} term of SWAN).

6.7. More comprehensive metrics

It is obviously desirable to evaluate model performance based on metrics other than total energy (wave height). In fact, for both the post-NCEX realtime system validation and the hindcast validations, peak period and mean wave direction is also included. However, due to the very large quantity of measurement locations and the duration of the time series, it was not possible to perform more than cursory inspection of these comparisons. Validation of directional spreading for long time series is possible, but is difficult to reduce to average quantities. Directional spreading is not very meaningful in cases where distinct wave components from multiple directions are integrated together, a probable occurrence with the wave climate of this region.

6.8. Other forcing sets

There was some interest in using operational Navy products to force the hindcasts (rather than NCEP products). The regional ("EPAC" or "East Pacific") wind and wave products were assembled for this purpose. The NCEP ("GFS") and FNMOC wind products ("COAMPS") were compared directly to winds measured at buoy 46047. This comparison suggested a possible slight advantage with the NCEP wind product, though the metrics for the two products were too close to be conclusive: one product reproduced some wind events better; the other reproduced other events better. NCEP GFS had lower RMS error; COAMPS had lower bias (COAMPS bias was positive; GFS bias negative). Preliminary hindcasts (in stationary computation mode) were performed for the outer SWAN grid (SC1) with various forcing combinations (FNMOC waves with NCEP winds, FNMOC waves with FNMOC winds, etc.) and comparisons were made to measurements at the "open locations". FNMOC boundary forcing had the advantage of being non-uniform along the boundaries (described at 1° resolution). However, the comparisons to data indicated a moderate advantage to using the NCEP forcing. This is the primary reason why NCEP forcing is used for the hindcasts presented in this paper (except at the western boundary, where buoy spectra were used). However, these preliminary hindcasts were subject to problems that were addressed in the hindcasts presented here (such as the under-convergence issue: the preliminary hindcasts were not hotstarted [re-initialized with low sea state] prior to each computation). Thus, these simulations would need to be repeated to confirm an actual advantage to the NCEP forcing.

6.9. Refraction computations at coarse resolution

SWAN has known problems calculating refraction in cases where waves turn a large amount (e.g. 50°) when propagating from one grid cell to another. In the Southern California Bight case, this is especially noticeable in the lee of the shoals east of 46047, where aphysical increase in wave height is predicted by SWAN. The refraction issue was confirmed to be the culprit: the high wave heights do not occur if either a) high geographic resolution is used, or b) refraction is disabled. Of course, neither is a good solution for a wave model (one is too expensive and the other removes physics). A test case was created which covers only the vicinity of the shoals, with geographic resolution equivalent to that used in the SC1 grid. The test case was run with a number of refraction limiters ("CDLIM", see SWAN manual, Holthuijsen et al., 2003), and a limiter of 1.25 was chosen as best replicating the results obtained with high geographic resolution ("ground truth"). This limiter was used in some of the preliminary hindcasts, but was not used in the hindcasts presented, since it was felt that more study of the practical effect of this limiter is required.

6.10. Model handoff (WAM/WW3 to SWAN)

The outer grid (SC1) could have been computed with a large scale wave model such as WW3 (and probably WAM4). We

expect that WW3 would be a bit more efficient than non-stationary SWAN at this geographic resolution (1′–2′). Note that in our hindcast nonstationary hindcasts, SWAN is a conditionally stable model, since the conditionally stable garden sprinkler correction is employed. For the high resolution model, the time step size (2.5 min) was dictated by this scheme. Due to this choice during hindcast design, SWAN in this case has no real computational advantage over WW3. The difference in efficiency is not great, however, so model choice at this scale can be governed by other concerns (user-friendliness, convenience, familiarity, ease of nest-to-nest communications).

7. Conclusions

The following conclusions can be made from this study:

- For modeling wave propagation at sub-regional scale in areas where sheltering effects are important, the accuracy of directional distribution of boundary forcing is critical. It is not enough to evaluate the accuracy of boundary forcing simply by comparing significant waveheight. The forcing may be consistently underpredicting swells from one direction and overpredicting swells from another direction: in this case, wave height may be accurate at offshore locations, but may be strongly biased (high or low) in nearshore locations, depending on the tendency of the local geography to block swells from one direction or another. This demonstrates that time invested in getting as accurate deep water directional spectra as possible is time well spent.
- Using stationary computations for an area the size of the Southern California Bight will lead to a moderate increase in root-mean-square error, primarily due to phase error: the aphysical instantaneous travel time of swells across the model grid, which occurs when the stationary assumption is used.
- Using coarse geographic resolution (e.g. $\Delta x = \Delta y = 1/20°$) in a sub-regional scale area where sheltering effects are important (such as the Southern California Bight) might be expected to carry penalties. In the model-to-model comparisons here, we found modest sensitivity to resolution (up to 6 cm RMS error). In terms of agreement with observations, there is little or no improvement derived from higher resolution. This may be due to resolution-related error being masked by errors related to boundary forcing or wind forcing.
- With the SWAN model stationary computations, extreme care must be taken with convergence criteria, especially for simulations with a long series of stationary computations. Failure to do this may lead to subtle but significant errors.
- With new parallel computing (e.g. OpenMP) features, SWAN is now a viable option for operational high-resolution nonstationary wave predictions at sub-regional scale.

Acknowledgements

This work was funded by the Naval Research Laboratory Core Program (project "Littoral Environment Nowcasting Systems"). The insightful recommendations by the anonymous reviewers are greatly appreciated. This is contribution NRL/JA/7320/05/5190.

References

Booij, N., Holthuijsen, L.H., 1987. Propagation of ocean waves in discrete spectral wave models. J. Comput. Phys. 68, 307–326.

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

Campbell, T.J., Cazes, J., Rogers, W.E., 2002. Implementation of an important wave model on parallel architectures. Proceedings of the Oceans 2002 MTS/IEEE Conference. October 29–31, 2002, Biloxi, Mississippi, pp. 1235–1239.

Günther, H., Hasselmann, S., Janssen, P.A.E.M., 1992. The WAM model Cycle 4 (revised version). Deutsch. Klim. Rechenzentrum, Techn. Rep. No. 4, Hamburg, Germany.

Hodur, R.M., 1997. The Naval Research Laboratory's Coupled Ocean/Atmospheric Mesoscale Prediction System (COAMPS). Mon. Weather Rev. 125, 1414–1430.

Hodur, R.M., Hong, X., Doyle, J.D., Pullen, J., Cummings, J., Martin, P., Rennick, M.A., 2002. The Coupled Ocean/Atmospheric Mesoscale Prediction System (COAMPS). Oceanography 15 (1), 88–98.

Holthuijsen, L.H., Booij, N., Ris, R., Haagsma, I., Kieftenburg, A., Kriezi, E., Zijlema, M., 2003. SWAN Cycle-III version 40.20 User Manual. 124 pp.[1]

Janssen, P.A.E.M., 1989. Wave-induced stress and the drag of air flow over sea waves. J. Phys. Oceanogr. 19, 745–754.

Kaihatu, J.M., O'Reilly, W.C., 2002. Model predictions and sensitivity analysis of nearshore processes near complex bathymetry. Proceedings of the 7th International Workshop on Wave Hindcasting and Forecasting, Banff, Alberta, pp. 385–396.

Komen, G.J., Cavaleri, L., Donelan, M., Hasselmann, K., Hasselmann, S., Janssen, P.A.E.M., 1994. Dynamics and Modelling of Ocean Waves. Cambridge Univ. Press. 532 pp.

Long, J.W., Özkan-Haller, H.T., Holman, R.A., 2004. Modeling and understanding remotely forced rip current systems at the Nearshore Canyon Experiment (NCEX). EOS Trans. AGU 85 (47) (Fall Meet. Suppl. Abstract OS13C-06).

O'Reilly, W.C., Guza, R.T., 1998. Assimilating coastal wave observations in regional swell predictions. Part I: Inverse methods. J. Phys. Oceanogr. 28, 679–691.

Rogers, W.E., Kaihatu, J.M., Petit, H.A.H., Booij, N., Holthuijsen, L.H., 2002. Diffusion reduction in an arbitrary scale third generation wind wave model. Ocean Eng. 29, 1357–1390.

Rogers, W.E., Hwang, P.A., Wang, D.W., 2003. Investigation of wave growth and decay in the SWAN Model: three regional-scale applications. J. Phys. Oceanogr. 33, 366–389.

Tolman, H.L., 1991. A third generation model for wind waves on slowly varying, unsteady and inhomogeneous depths and currents. J. Phys. Oceanogr. 21, 782–797.

Tolman, H.L., 2002a. User Manual and System Documentation of WAVEWATCH-III Version 2.22. NCEP Technical Note. 133 pp.[2]

Tolman, H.L., 2002b. Validation of WAVEWATCH III version 1.15 for a global domain. NCEP Technical Note. 33 pp.

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

Zijlema, M., van der Westhuysen, A.J., 2005. On convergence behaviour and numerical accuracy in stationary SWAN simulations of nearshore wind wave spectra. Coast. Eng. 52, 237–256.

[1] The reader may also refer to the current version of this manual, available from the TuDelft website at http://www.fluidmechanics.tudelft.nl/swan/default.htm.

[2] The NCEP Technical Notes are not formally published, but electronic versions are available for download from NCEP.