Using Genetic Algorithms to Optimize Bathymetric Sampling for Predictive Model Input

DINESH MANIAN

Moffatt and Nichol Engineers, New York City, New York

JAMES M. KAIHATU AND EMILY M. ZECHMAN*

Zachry Department of Civil Engineering, Texas A&M University, College Station, Texas

(Manuscript received 29 March 2011, in final form 23 September 2011)

ABSTRACT

This paper describes the use of an optimization method to effectively reduce the required bathymetric sampling for forcing a numerical forecast model by using the model's sensitivity to this input. A genetic algorithm is developed to gradually evolve the survey path for a ship, autonomous underwater vehicle (AUV), or other measurement platform to an optimum, with the resulting effect of the corresponding measured bathymetry on the model used as a metric. Starting from an initial simulated set of possible random or heuristic sampling paths over the given bathymetry using certain constraints like limited length of track, the algorithm can be used to arrive at the path that would provide the best possible input to the model under those constraints. This suitability is tested by a comparison of the model results obtained by using these new simulated observations, with the results obtained using the most recent and complete bathymetric data available. Two test study areas were considered, and the algorithm was found to consistently converge to a sampling pattern that best captured the bathymetric variability critical to the model prediction.

1. Introduction

a. Environmental variability at different scales

Ocean wave prediction models that operate on global scales have reached sufficient maturity to allow for routine operational prediction. Various government agencies worldwide such as the National Oceanic and Atmospheric Administration (NOAA) and the European Centre for Medium-Range Weather Forecasts (ECMWF) offer daily predictions of wave properties (heights, periods, and directions) around the globe, usually using various combinations of wave models such as Wave Model (WAM; Komen et al. 1994) or Wavewatch-III (Tolman 1991) with global wind models. Closer to shore, however, the efficacy of forecasting becomes less clear primarily because of the highly dynamic environment present in the nearshore areas. This dynamism is caused by breaking waves and hydrodynamic currents as well as more infrequent events such as surge due to storms, hurricanes, and tsunamis. The underlying bathymetry in the nearshore areas has a strong effect on these processes, particularly the occurrence of wave breaking. Information on the nearshore bathymetry, therefore, is a primary consideration for effective modeling of waves and currents in this area.

b. Model sensitivity to bathymetric input

Predictive modeling of the coastal ocean has also been on track toward routine operational use; these nearshore models are of significant utility for many operations where predictive knowledge of waves and currents in a region near the coast is required. Among the inputs required for these models are estimates of wind, incident wave conditions at open boundaries, and sufficient bathymetric data at the modeled locations. An example of the development and use of a nearshore modeling system is discussed by Allard et al. (2008).

* Current affiliation: Department of Civil, Construction and Environmental Engineering, North Carolina State University, Raleigh, North Carolina.

Corresponding author address: James Kaihatu, Zachry Department of Civil Engineering, Texas A&M University, 3136 TAMU, College Station, TX 77843-3136.
E-mail: jkaihatu@civil.tamu.edu

DOI: 10.1175/JTECH-D-11-00051.1

However, prior to a modeling exercise, the degree to which one can consider the bathymetry sufficient for a particular model is not clear. Phase-resolving wave models, for example, simulate the evolution of surface waves on a subwavelength scale, and thus may need higher-resolution bathymetry (capturing bathymetric variations at subwavelength scale) than would be required by a phase-averaged wave model. This potential constraint is also affected by the degree to which high-resolution coastal bathymetry databases exist; significant knowledge gaps exist in the bathymetric record in many areas in the coastal and nearshore ocean. Furthermore, bathymetric records required for modeling are often created by melding and interpolating survey measurements of varying accuracy taken at different times, the results of which are dependent on the interpolation method used or the quality of the data (Plant et al. 2002). Thus, additional field surveys may be required to supplement existing databases, particularly if the areas in question were undersampled or have undergone an anticipated degree of bathymetric change. Under these circumstances, where timely information is required and highly valued, determination of the amount of information absolutely necessary for useful model prediction becomes important. Other constraints affecting the efficacy of measurement pertain to the platforms (boat, jet ski, etc.) and their power requirements. For example, autonomous underwater vehicles (AUVs) have been shown to be a tenable platform for bathymetric data collection in such situations (Bourgeois et al. 2005); however, they generally have a fixed operational time because of battery life. Thus, for optimal management of resources and data, techniques that make use of existing datasets to optimize the application of such platforms toward a particular end use would be beneficial.

There has been some related prior work on the sensitivity of bathymetric information on the results of predictive numerical models. Kaihatu and O'Reilly (2002) investigated the effect of updated bathymetric surveys of Scripps Canyon, California (CA) on the results of the Simulating Waves Nearshore (SWAN; Booij et al. 1999). Plant et al. (2002) developed an interpolation technique using Gaussian-shaped smoothing windows for bathymetric data processing, which allowed for control of the spatial scale of various bottom features in the interpolated result. The resulting technique was also useful in determining the relative amount of smoothing any small-scale features in the measured bathymetry undergo when interpolated to a model grid. The effect of this smoothing on wave and wave-averaged hydrodynamic models was analyzed by Plant et al. (2009). They determined that the wave and wave-averaged hydrodynamic models had different sensitivities to changes in the bathymetry. Use of these aspects of the sensitivity to bathymetric variability scales would thus be useful for designing survey routing and sampling such that the collected bathymetry would have the most influence upon the response of the model.

c. *Nearshore numerical wave model*

The model under consideration here is SWAN (Booij et al. 1999), which has been used for routine forecasting for regional and coastal domains. It is the default wave driver for the Delft3D model (Lesser et al. 2004)—a hydrodynamic software package widely used in many engineering and defense applications for simulations of waves, flow, sediment transport, and morphology. This study explores the behavior of the SWAN model and its sensitivity to bathymetric input and develops a method to optimize the required data collection.

To develop a feasible way to determine an optimal sampling for use as a model input to Delft3D, the spatial variation in the sensitivity of model-predicted wave heights and currents to bathymetry particular to a given study area needs to be examined. The goal of this study is to identify the critical bottom features that must be captured in sampling and to obtain an interpolated bathymetry that enables the model in generating the most accurate results of wave heights and currents. The use of a global optimization scheme for this purpose is proposed in section 2, which compares different approaches to the sampling problem. Computer-generated random survey paths are evaluated for their ability to capture the areas critical to the model function, and a genetic algorithm is developed and applied to evolve the fittest possible path. The sampled bathymetry along this path is interpolated to the model grid using inverse distance-weighted triangular interpolation. The fitness is ascertained by evaluating the differences between the model results produced from the path-interpolated bathymetry and those from the best available bathymetry. A general background of the concept of genetic algorithms is laid out in section 2, and section 3 describes the procedure followed to adapt and apply this technique to the present problem. Sections 4 and 5 present the results and conclusions, respectively, along with possible improvements and alternatives to the present approach.

The work presented here establishes a new methodology based on genetic algorithms (GAs) to identify optimal bathymetric survey paths for fitting predictive nearshore models. Specifically, new operators are necessary to facilitate the use of a genetic-algorithm-based search for bathymetric survey paths. Three new problem representations, or encodings, are developed and tested here to enable an efficient search for a survey path that will maximize the quality of new data. The specific application in this study requires a priori knowledge of the high-resolution bathymetry, since the model results over this bathymetry are being used as the comparison metric. However, in the present context, the GA-driven survey path is optimized in the sense that the bathymetry so measured exerts the greatest influence on the wave field in a particular region. Thus, should resampling of the bathymetry (often performed to help improve simulations and forecasts) become necessary, this algorithm can be implemented and extended to advise where the available measurement resources should be concentrated. Additionally, if high-resolution measurements of the free surface properties are available, the algorithm may be altered to provide optimal estimates of the bathymetry with no further a priori knowledge.

The development outlined here is not intended for general utility at this stage. The optimum path calculations are sensitive to the incident wave environment, as will be shown. To increase the applicability of the algorithm would involve calculation of a large number of optimum paths, corresponding to the relevant wave conditions for a given area. Some aspect of this is envisioned for future development using the present work as a basis.

2. Genetic algorithms

a. Approaches to the bathymetric sampling problem

If we assume that the objective of the bathymetric sampling is defined as the generation of the most cost effective input for the SWAN model, one possible method of determining the sampling strategy could be by studying the theoretical sensitivity of the predicted wave heights and currents at the required location(s) to the bathymetry. This would generally be done by performing a perturbation expansion on one of the dependent variables and one of the inputs; the input perturbations represent input errors, and the resulting analysis would determine the growth or decay of the model errors as a function of the input errors. As an example, this was performed by Chen and Svendsen (2003) to evaluate the sensitivity of modeled nearshore currents on errors along the lateral boundaries. If one focused on the sensitivity to bottom depth, this sensitivity would vary in space, and a high sensitivity at any given set of locations would indicate that greater importance must be given to capturing those locations during sampling. However, the ubiquity of depth-dependent terms in the governing equations of wave and hydrodynamic models, and the nonlinear dependence of predicted waves and currents on depth, precludes this in anything other than the simplest conditions. Therefore, any effort to quantify the effect of the input local bathymetry on the model results would have to use nonlinear or stochastic methods. Rather than an analytic approach, these methods make use of random simulations from a physics-based model to deduce the appropriate sensitivities.

b. Genetic algorithms

We wish to investigate the tenability of GAs for this application. GAs are widely used to solve problems in optimization and machine learning. A GA-based approach evolves a set of candidate solutions for an optimization model using operators inspired by natural selection and genetic variation.

The stochastic approach of GA-based approaches overcomes some limitations of traditional methods of optimization, especially for multimodal problems, which are characterized by more than one solution that satisfy system objectives to a similar degree. While traditional methods such as gradient descent may converge to locally optimal solution, a GA samples a better representation of the entire search space, increasing the probability of convergence to a global extremum. The idea of evolutionary computation, of which genetic algorithms are a subclass, came about as a result of efforts to emulate nature to solve optimization problems (Holland 1975). In nature, there exist variations within species that are brought about by reproduction. Of the many offspring produced by the individuals of a population, only a few survive to adulthood and reproduce in turn. This is known as "natural selection," where the survival of an offspring is dependent on its suitability and adaptability to its environmental conditions. Similarly, a genetic algorithm maintains a population of possible solutions to the given problem, to which the concept of "survival of the fittest" is applied. Reproduction is achieved by the "crossover" of parent chromosomes, and the mutation operator is responsible for introducing diversity by randomly changing some genes. Holland (1975) was the first to successfully model the mixing of chromosomes that occurs in nature to create a functioning genetic algorithm, and Goldberg (1989) formally developed the GA methodology in the context of solving optimization models. The following section further describes this method, and brings out the meaning of a few of these technical terms in the context of the optimization problem.

c. Method outline

A basic genetic algorithm has the following steps (Goldberg 1989):

1) problem representation,
2) problem initialization,
3) fitness evaluation,
4) selection of individuals to produce new offspring,

5) reproduction of new individuals using crossover and mutation,
6) reinsertion of new solutions to the population using elitism, and
7) repeat steps 3 to 6 till convergence.

The following is a brief general description of each of these steps.

1) Problem representation

A set of candidate solutions to the problem as maintained by the GA is referred to as a population. Each individual of the population is called a chromosome, which is composed of a string of genes. Genes can be binary, integer, or real-valued numbers.

2) Population initialization

To produce the initial population, random number generators are typically used to produce the gene values for each individual to ensure sufficient spread over the entire search space. The size of the population depends on the nature of the problem or (more specifically) the nature of the search space. A larger population size requires a larger number of computations to produce offspring, but allows greater diversity of solution values in the population, which can contribute to efficient convergence to global optima.

3) Fitness evaluation

The suitability of a particular solution is characterized by computing the ''fitness value.'' This fitness corresponding to the specified optimization criterion is calculated for each individual, and the objective of the algorithm is to maximize this in subsequent generations.

4) Selection of individuals to produce new offspring

In the genetic algorithm framework, individuals are selected to survive to the next generation based on fitness values, which ensures that individuals with higher fitness values have a greater effect on the properties of the next generation of solutions.

5) Reproduction of new individuals using crossover and mutation

New individuals are produced from the selected pairs in each generation using the genetic operations of crossover and mutation. The crossover operation combines the characteristics of the two parent chromosomes to form new offspring. The probability of a crossover operation is specified by the crossover rate. Similarly, the probability of mutation of a gene is specified by the mutation rate, and mutation is executed by inserting a new randomly created gene in place of the selected gene.

6) Reinsertion of new solutions into population using elitism

At the end of each generation, the best offspring are inserted into the population in place of the least fit individuals, which are discarded. Elitism ensures that the best-performing individuals are always retained and can increase the speed of convergence.

3. Methodology

Genetic algorithms have been used as a global optimization scheme in a variety of applications, including oceanography. They were first used in the area of oceanographic experiment design by Barth (1992), who considered a time-dependent design problem for an idealized experiment. Baehr et al. (2004) first employed it for predeployment array design in optimizing an observing system for the North Atlantic meridional overturning circulation. They compared this technique with both the simulated annealing method and a heuristic approach, and found genetic algorithms to be a significantly faster method than simulated annealing (Barth and Wunsch 1990) and at the same time more successful in finding the optimum solution than the heuristic approach.

In terms of nonsynoptic measurements using moving platforms, Bellingham and Wilcox (1996) described a method to determine the optimum resolution and extent of survey that would minimize the energy cost and the total error for oceanographic AUV deployment. This was an attempt to accommodate the temporal variation of the ocean environment during AUV deployments using the physical limitations of the vehicle as a constraint. However, no investigations of methods to increase the efficiency of temporally static, spatially nonuniform oceanographic surveys could be found in the scientific literature, particularly with the overall goal of input into predictive models. The parameter to be optimized in such surveys would be the path of motion of the survey vehicle, rather than a distribution of single-point measurement sensors. The problem thus becomes one of path planning to minimize the total survey error at an acceptable energy cost. The problem of path planning in general and for AUVs in particular has been quite extensively studied before in the context of ocean navigation with a given current field at minimum energy cost (Alvarez et al. 2004; Fox et al. 1999). However, the survey path planning problem has somewhat conflicting requirements of maximizing survey extent to achieve minimum survey error, and doing so with minimum path distance, which calls for a different GA scheme.

a. Objective

The broad objective of this study is to estimate the required amount of bathymetric information for the hydrodynamic model and the sampling strategy needed to minimize this amount. This is done by setting ourselves two goals.

The first goal is to estimate the spatial variation in the optimum cross-shore and longshore resolution requirements of bathymetry for the model for a given study area. This is hence referred to as the ''optimum resolution problem.'' This is equivalent to the problem of designing the spatial distribution of parallel longshore or cross-shore bathymetry survey tracks for best possible bathymetric input to the model.

The next goal is to design an optimum continuous path for a bathymetric survey vessel such as an AUV, focusing on the utility of the thus-sampled, model–grid interpolated data as a model input for Delft3D (hence referred to as the ''survey design problem''). Two different encoding schemes for the survey design problem are developed and evaluated. Under the first scheme (scheme 1), the problem is somewhat simplified by constraining the number of degrees of freedom of the path such that it can only take perpendicular turns, and is of the form of a line-by-line sweep of the coverage area as shown in Fig. 1. The sweep direction shown in the figure is long shore. This approach ensures at the outset a more or less uniform coverage of the area of interest. However, in this case, the basic structure of the path to be followed is predetermined. With the second approach (scheme 2), the GA is freer to decide the structure of the path to be followed by using a strategy that only sets an upper limit of the length of the path, and allows it to move in any of eight possible directions (sectors of 45° each around a circle starting from due north). Starting from an initial randomly chosen point, the step length is chosen as the minimum unit of distance that must be traveled in a straight line before changing direction. Thus, the path has to pass through adjacent points on a grid of resolution equal to the step length. Though a smaller step length might be desirable in order to make the algorithm select and converge on the best possible path (i.e., one that would lead to an interpolated bathymetry and model results with smallest deviation to those using the high-resolution bathymetry), the step length must be set high enough so that there is no excessive concentration of measurements or looping of the path in a small area at the expense of wider coverage of the complete area of interest. For each of the encoding schemes, we make no a priori assumptions concerning the endpoints of the AUV's path. Constraints on the AUV pathlength involve assumptions concerning power consumption. In a field situation, however, the AUV path endpoints can become important because of their proximity (or lack thereof) to land or infrastructure (boats, batteries, etc.). Incorporation of these endpoint effects can take the form of additional constraints for the GA optimization; this is not pursued here.

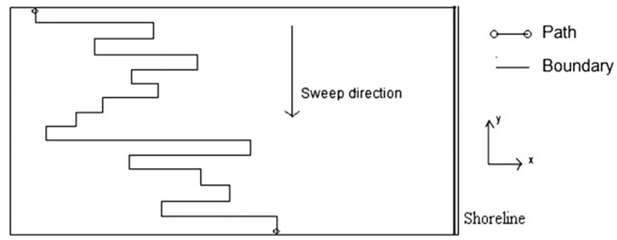

FIG. 1. A typical path for bathymetric sampling as modeled in the study (scheme 1).

b. Application of the GA

The general blueprint of the GA followed in both these problems mentioned in the objective remains the same as described in the previous chapter. However, the problem representation schemes and implementations vary. The following subsections provide a detailed description of the implementation of these steps for the stated problems in the previous section.

1) OPTIMUM SPATIAL RESOLUTION PROBLEM

(i) Encoding scheme

The ''encoding scheme,'' which describes how each individual solution is encoded as a chromosome, needs to be carefully selected, as it determines the effectiveness of the genetic operations performed on the chromosomes. In this case, each individual solution was encoded as a string of a fixed number of real-valued numbers (genes) that correspond to the y coordinates of the longshore tracks, or x coordinates of the cross-shore tracks. The positive x direction was assumed to point toward the shoreline. The number of genes in the encoded solution varied according to the number of tracks desired.

(ii) Initial population generation

The individuals of the initial population were randomly generated to simulate a real numbered value (x or y coordinate in the given domain) corresponding to each gene. Preliminary studies tested different population sizes to determine the effect on the performance of the GA. The population size leading to the best convergence characteristics was identified for use in further analysis.

(iii) Fitness calculation

The wave model was first run with the best available, highest-resolved bathymetry and the results thus obtained were used as the standard of comparison (a so-called "golden standard") for evaluating the fitness of the subsampling schemes. The inverse of the mean absolute difference between the golden standard results and the individual results over the entire area of interest was then defined as the fitness parameter of that individual.

(iv) Selection and replacement

There are various methods of selection documented in the literature (Goldberg and Deb 1991). The "Roulette selection" method was implemented, in which the probability of a member being selected to reproduce is directly proportional to its calculated fitness value relative to the other members. Elitism was executed at every generation.

(v) Genetic operations

Single-point crossover was used, which means that the parent chromosomes were truncated at a single randomly selected gene before being recombined to form new offspring. A 100% crossover rate was used so that this operation is always performed to create new individuals. A mutation operation was used in order to introduce new properties (random genes) to the offspring solution at a specified mutation rate, which was selected for best convergence on the basis of the results of the application of the GA to a test problem (not shown).

2) SURVEY DESIGN PROBLEM: SCHEME 1

(i) Encoding scheme

In this case, each individual is a random path of fixed length. Since spatial location is the primary property characteristic of these various alternate solutions, the encoding scheme must be such that each gene has a uniquely identifiable geographical area associated with it. Only then can crossover and other genetic operations be used to create progressively better solutions. Therefore, the path (chromosome) was divided into sections (genes) of equal length, and the study area divided along the y axis (see Fig. 1) at equal distances, so that each cross-shore strip of area had just one associated gene.

The structure of the solution is constrained to represent a path in the form of a line-by-line sweep, where the distance between sweeps in the y direction is fixed. For this representation, each gene is encoded as a string of numbers representing the starting x coordinates and the direction of the lines associated with a gene. To ensure that the length of the section of path corresponding to each gene is a constant l, each section was divided into n subsections. These subsections consisted of parallel tracks oriented perpendicular to the sweep direction (see Fig. 1), and which were created by selecting $n - 1$ random points between 0 and l, where n is the number of equidistant parallel tracks (either long shore or cross-shore) associated with that gene. This produced a set of n real numbers d_i, obeying the following constraint:

$$\sum_{1}^{n} d_i = l. \tag{1}$$

The x coordinate of the starting point of the section was set equal to the x coordinate of the end point of the previous adjacent section, and the x coordinate of the start of each new subsection was decided by

$$x_{j,i} = x_{j,i-1} + k_{i-1}d_{i-1}, \tag{2}$$

where $k_i = \pm 1$ is a random number, j is the gene index, and i is the allele index.

The next step was to check whether the path thus described obeys the constraints of the given boundary and thus remains in the domain. If not, the current sequence was discarded and a new one created until the constraints were met. The number of such iterations required was found not to be large enough to significantly increase computational costs. Thus the encoded solution was of the form

$$[(x_{1,1}, \ldots, x_{1,n}), (x_{2,1}, \ldots, x_{2,n}), \ldots, (x_{m,n}, \ldots, x_{m,n})],$$

where m is the number of genes and n the number of alleles of a gene.

(ii) Initial population generation

The genes of all the chromosomes of the initial population were randomly initialized and then checked to ensure that each chromosome satisfied the set constraints of the domain boundaries and the constant length.

(iii) Fitness calculation

The fitness parameter was calculated in the exact same manner as in the "optimum spatial resolution problem" (section 3a). Also, the Roulette selection method was similarly used to select parents for reproduction.

(iv) Genetic operations

The crossover rate used here was 100%, and a single-point crossover at a randomly selected rth gene was performed. However, a simple crossover would not produce a feasible offspring solution as the resultant path would most likely be discontinuous. Therefore, while the remaining genes of either parent were carried over to the offspring during the crossover in a typical fashion, a random mutated gene was inserted in place of the rth gene to create a feasible individual, enforce conditions of continuity, and maintain the pathlength.

To further ensure that the length of the section of path corresponding to each gene is constant at l, and at the same time continuous with the adjacent sections, $n - 2$ subsections of the form $[(0, l_1), (l_1, l_2), \ldots, (l_{n-3}, l_{n-2})]$ were created by selecting $n - 2$ random points between 0 and l. The first $n - 2$ alleles were then encoded and tested as described in the encoding scheme. The $(n - 1)$th allele was chosen to satisfy

$$d_{n-1} + d_n l- = \sum_{i}^{(n-2)} d_i, \quad \text{and}$$
$$k_{n-1}d_{n-1} + k_n d_n = x_{j,n} - x_{j,n-2}. \tag{3}$$

In the absence of real positive solutions for d_{n-1} and d_{n-2}, the procedure was repeated until one was found. The new chromosome was then inserted into the population in place of the least fit member.

3) SURVEY DESIGN PROBLEM: SCHEME 2

The following steps in this scheme differ from those in scheme 1.

(i) Encoding scheme

A grid of a specified step length was defined. The start point of the vehicle path was picked at random from the nodes of the defined grid. The integer index numbers of this grid point constituted the first two genes of the candidate solution. One of the adjacent grid points was picked as the next decision point on the path. The relative position of these grid points to the current location was encoded according to the numbering convention shown in Fig. 2, where the center represents the current location. This string of relative directions encoded as integers formed the remaining genes of the individual.

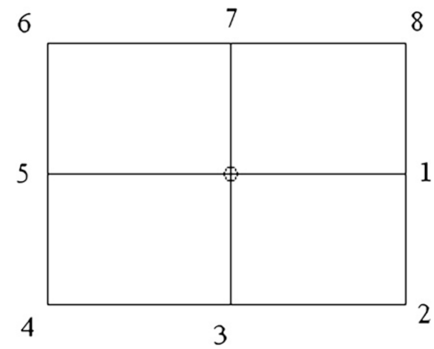

FIG. 2. Relative direction encoding scheme.

(ii) Genetic operations

A single-point crossover was used at a rate of 100%, and mutation of individual genes carried out at different rates. Thus in this case, the geographic location information of one of the two parent paths was lost, and only its structure passed on to the offspring.

4. Results

a. Study area descriptions

We chose two field sites to test the GA implementation for bathymetric surveys. A geographic reference map of the sites is shown in Fig. 3. For each test site, we used measured bathymetry and wave forcing measured at a nearby wave buoy. Each site has unique bathymetric characteristics that test the robustness of the GA implementation.

The first site is the area offshore of La Jolla, CA. This site is marked by two large undersea canyons (Scripps Canyon and La Jolla Canyon) and relatively smooth, featureless bathymetry near shore of the canyon heads (Fig. 4a). This area was the site of the Nearshore Canyon Experiment (NCEX) in 2003 (Elgar 2003). The steep topography at the canyons could be expected to cause significant changes in wave energy in the long shore.

The incident wave forcing for this site was obtained from a wave buoy operated by the Coastal Data Information Program (CDIP) at Scripps Institution of Oceanography (Coastal Data Information Program 2009). This buoy (station identifier: 100) is presently located at 32°55.84′N, 117°23.54′W (approximately 12 km offshore) at a water depth of approximately 550 m. The incident wave condition used had a significant wave height of 1.16 m, a peak period of 9.09 s, and a peak approach direction of 10° south of due east. To avoid multiple iterative model runs over a large domain, we employed a nested model approach. The SWAN model was first run over a larger domain (shown in Fig. 4a), and the resulting wave spectra along the offshore boundary of an approximately 3×3 km^2 grid (shown in Fig. 4b) was written out to be used as the boundary condition for the SWAN runs over this smaller area of interest. Since the focus is on the bathymetry—a static quantity—the SWAN model was run as a time-stationary model; thus these boundary spectra remain constant over the GA iterations, and there is no consideration of computational time step. A computational grid resolution of 40 m both in the cross-shore and long shore was used for the smaller domain. This resolution was used to ensure sufficient numerical accuracy for this analysis.

A second study site was selected off the coast of Camp Lejeune, North Carolina (NC). This area of about 22×18 km^2 has a higher degree of small-scale irregularity in the bathymetry (Fig. 5) relative to the La Jolla region. Therefore, one might expect a significant difference in

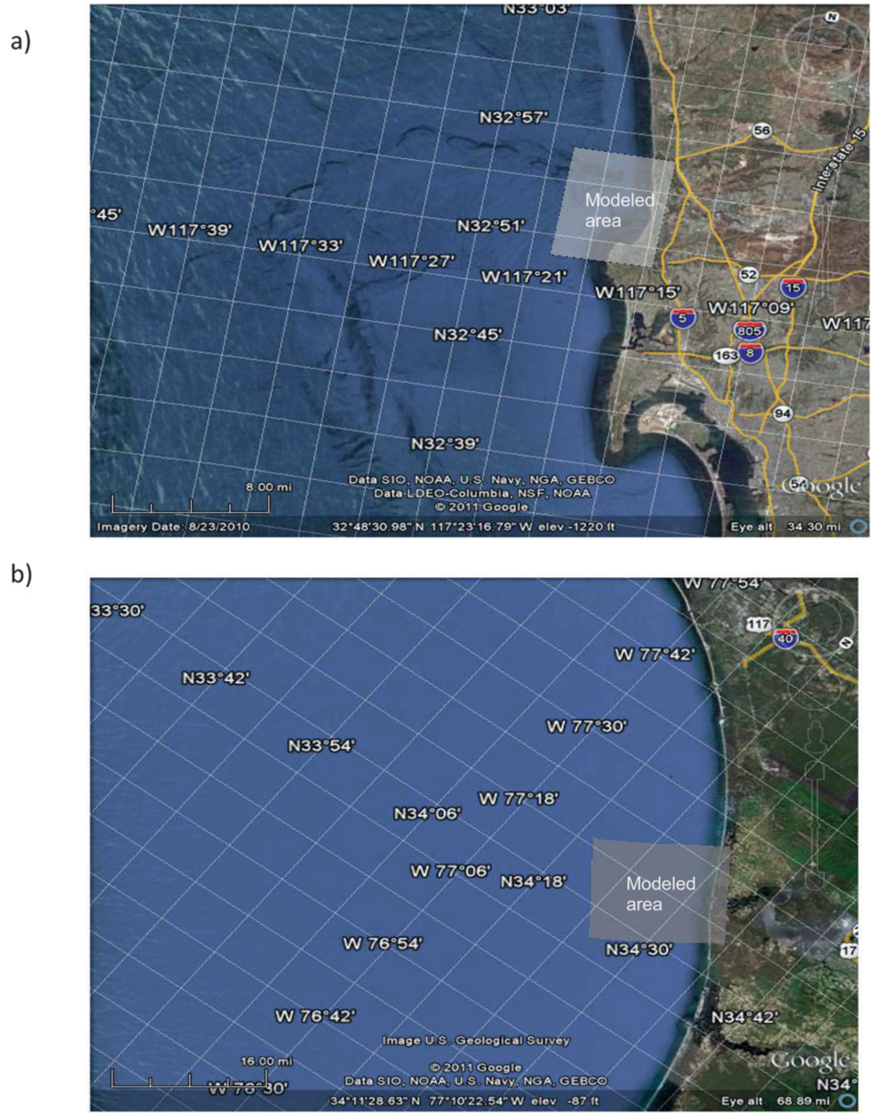

FIG. 3. Geographic reference showing the study area (shaded gray) at (a) the site of NCEX, La Jolla, CA, and (b) Camp Lejeune, NC.

the product bathymetry and thereby model results depending on where the sampling is done. This might provide a greater incentive for performing the optimization and would be a test of the usefulness of the method. The golden standard model wave heights corresponding to the best bathymetric input are also shown in Fig. 5. This served as a benchmark for comparing the model results from sparsely sampled input. The typical wave boundary conditions for this region were obtained from a National Data Buoy Center (NDBC) buoy located at 34.476°N, 77.280°W (station 41035, Onslow Bay, NC). The wave condition used for this had a significant wave height of 0.79 m, a peak period of 5 s, and was oriented normal to the shoreline.

b. *Results for optimum spatial resolution*

The performance of the method was measured in terms of the relative percentage mean absolute difference between the computed wave heights (with subsampled bathymetry as input) over the entire study area and the expected wave heights—predicted using all available bathymetric data as input. Figures 6a–c show the results for the case when the longshore and cross-shore sampling were optimized to a minimum cost of 5.5% and 3% error, respectively.

As can be seen from the figures, the resultant sampling scheme is denser near the south canyon, near which the wave-height gradient is also at its greatest (Fig. 6b).

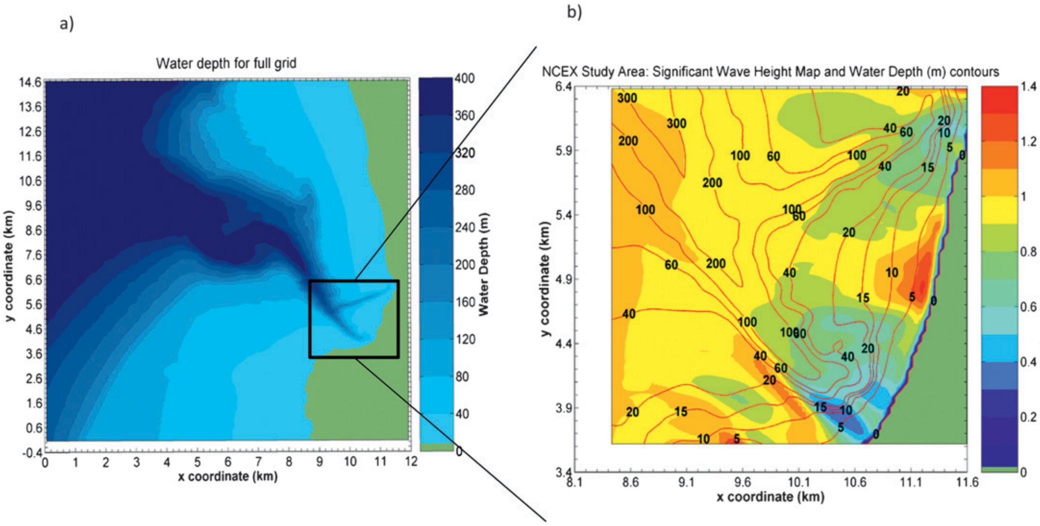

FIG. 4. The site of NCEX, La Jolla, CA: (a) large offshore wave model domain, and (b) bathymetry and modeled significant wave heights over nearshore area of interest.

Also, the area closer to the shore seems to require greater sampling (Fig. 6b), which is consistent with what one might expect given the bathymetric complexity in this area. Though the average discrepancy in the newly computed wave height over the entire study area was used as an indicator of the performance of the method, the performance at any given location might be very different.

c. Results for survey design: Scheme 1

Figure 7a shows a 27-km length of survey track, optimized to serve the model to produce the best possible wave heights over the selected study area. The corresponding grid-interpolated bathymetry and the resultant wave heights from the model are shown in Figs. 7b and 7c, respectively. The performance metric used was the same as described in the previous section, and Fig. 7d shows the convergence of the GA scheme. The computation time was about 12 s per iteration of the GA on a 2.66-GHz dual core processor of 2 GB memory. Increasing the extent of sampling has a definite but relatively small effect in improving the performance of the algorithm.

For different specified lengths of path and given input wave spectrum, the optimized paths always converged to a pattern that sampled the deep northwest canyon and the trend line lying in the northwest–southeast direction, so that the south canyon was also mostly captured. As mentioned above, the peak wave approach direction on the input wave spectrum was about 10° south of east. On shifting the peak wave direction in the input wave spectrum more toward north, significant changes in the optimized survey path were observed as the trend line of the path oriented itself roughly parallel to the wave direction, resulting in a trend line more along the northeast–southwest axis. This is probably indicative of increased wave-height variability near the northern part of the coastline, which would require a better representation of the north canyon for the model to be able to reproduce correctly. Figures 8a and 8b, respectively, show the wave-height fields that would result from changing the peak direction of the input wave spectrum and the effect of this change on the optimized survey tracks produced by the GA.

d. Results for survey design: Scheme 2

In this scheme, there are fewer constraints on the structure of the desired path. In the study over the La Jolla

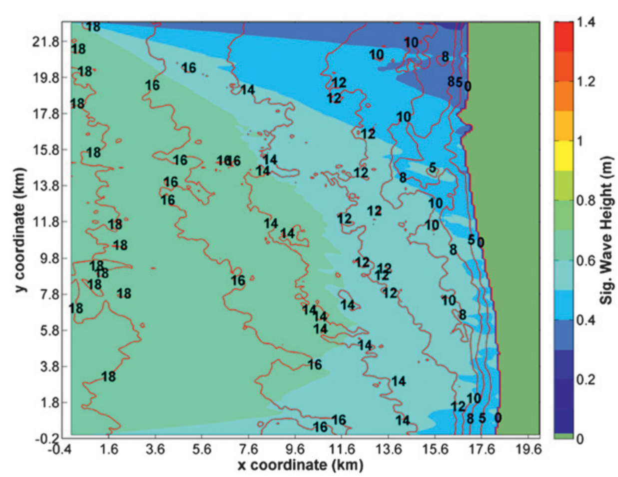

FIG. 5. Camp Lejeune, NC: bathymetric field contours and color map of modeled significant wave heights.

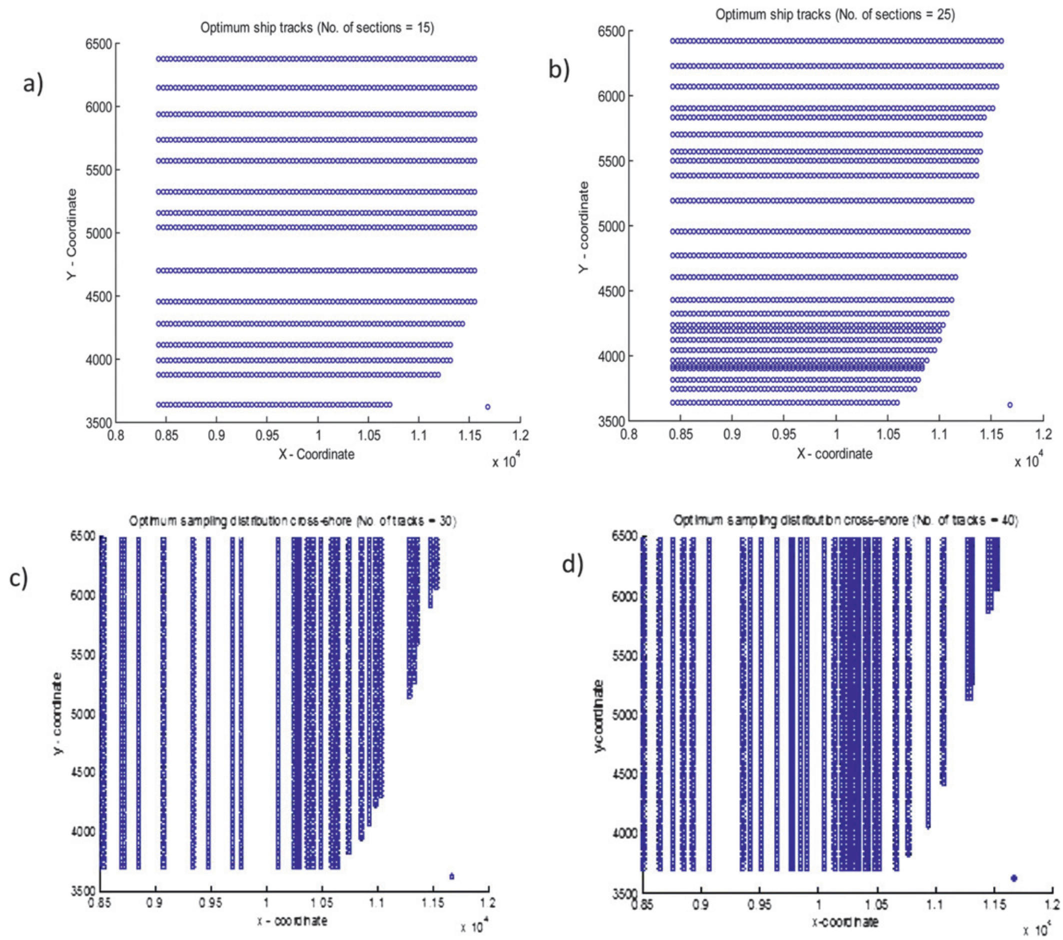

FIG. 6. Variation in wave model sensitivity to variations in the bathymetry, NCEX case: (a) varying longshore resolution, (b) varying longshore resolution with increased sampling, (c) varying cross-shore resolution, and (d) varying cross-shore resolution with increased sampling.

region, a relatively low maximum survey pathlength of 8 km was specified. An optimized solution path after 15 generations is shown in Fig. 9a, while the derived bathymetry and wave heights and convergence are shown in Figs. 9b, 9c, and 9d, respectively. The computation time for this study region was of the order of 7 s per iteration of the GA. A step length (see section 3b) of 400 m was used here, which means every decision point on the path was followed by a straight line section of at least 400 m. This was chosen such that excessive local looping that would impede possible wider coverage is avoided, and at the same time enough decision points are allowed to give room for the solution path to evolve.

As model wave height is the only criterion considered for optimization of the bottom sampling, the possibility of many distinct bottom configurations producing similar looking wave-height fields arises, especially so when the sampling extent is small. Also, since the cost function used considers only the spatially averaged wave-height errors with respect to the golden standard (see Fig. 4b), there could be considerable difference in the actual spatial variation of wave heights as seen in Fig. 9c. As seen in the figure, the optimum path sampled the south canyon well, while the north canyon was almost completely ignored. The effect of this is evident in the wave-height fields produced, which show a better match with the golden standard in the southern part of the modeled area.

The results of the study over the Camp Lejeune, NC bathymetry—the optimized path and the derived bathymetry and model wave heights—are shown in Figs. 10a,b. The upper limit used for the length of survey was 55 km on an area of about 18 km × 20 km. This length was chosen to be small relative to the bathymetric complexity and the domain size in order to provide the algorithm with a particularly strenuous test. A step length of 10 times the computational grid resolution of 92.5 m was used. In contrast with the study over the La Jolla

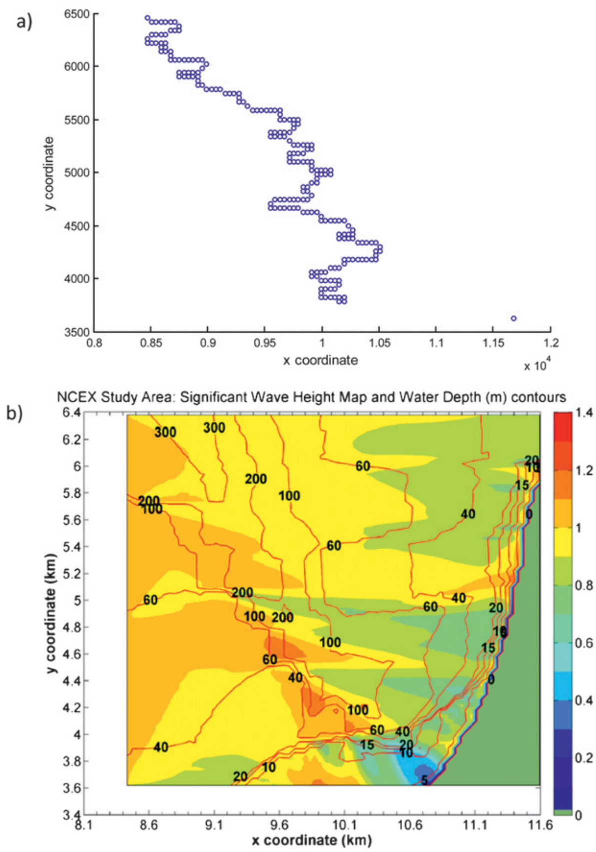

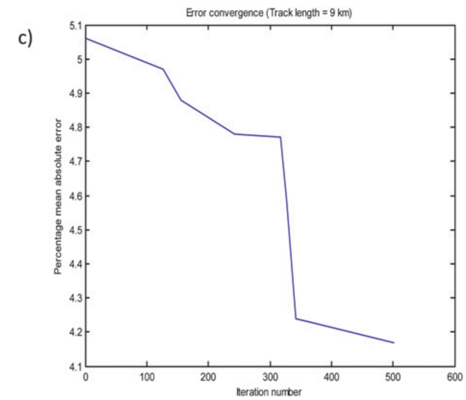

FIG. 7. NCEX case: (a) an optimized survey track length of 9 km using scheme 1, (b) bathymetry interpolated from sampling along optimum track and corresponding model wave-height results, and (c) convergence plot for track length of 9 km.

region, the spatially averaged relative error in modeled wave height for such a small length of survey path was found to be relatively very high—of the order of 23%. Moreover, because of the large size of the study area, the computation time required was of the order of 2 min per iteration.

5. Conclusions

a. Overview

Bathymetric information for a particular region is a vital input to coastal wave and hydrodynamic predictive models. Direct measurement of bathymetry, however, is laborious under the most favorable of circumstances. If we take the view of the bathymetric field as an input to the model, it is sensible to determine the spatial scales of bathymetric variability to which the models will respond, since most predictive nearshore models are averaged over some length scale. This also has the effect of optimizing a bathymetric survey, since only the bathymetry most influential to the model response would be measured. The genetic algorithm (GA) approach developed here attempts to make more efficient use of bathymetric measurement platforms for the specific task of providing better hydrodynamic model input. As this is an initial proof-of-concept study, we make no attempt at creating a utility with general applicability, but instead explore the use of GA in optimizing surveys with selected wave conditions. A set of encodings for this type of problem are explored here to allow higher degrees of freedom in the search for an optimal path. Small inefficiencies in the paths that are identified are noted; for example, a path that is identified using these encodings can loop back or cross itself. The GA-based search provides a flexible framework for investigating new representations to both improve convergence and ensure that the entire decision space is explored.

To test the algorithm, two study areas were considered, and the GA was then applied to determine the amount and locations of bathymetric data that have the greatest impact on model response. It was also shown that this result changed as the incident wave condition changed. In a practical application, a range of incident wave conditions prevalent for a given region can be run with the GA to ascertain all regions in the domain exerting high influence on the model response.

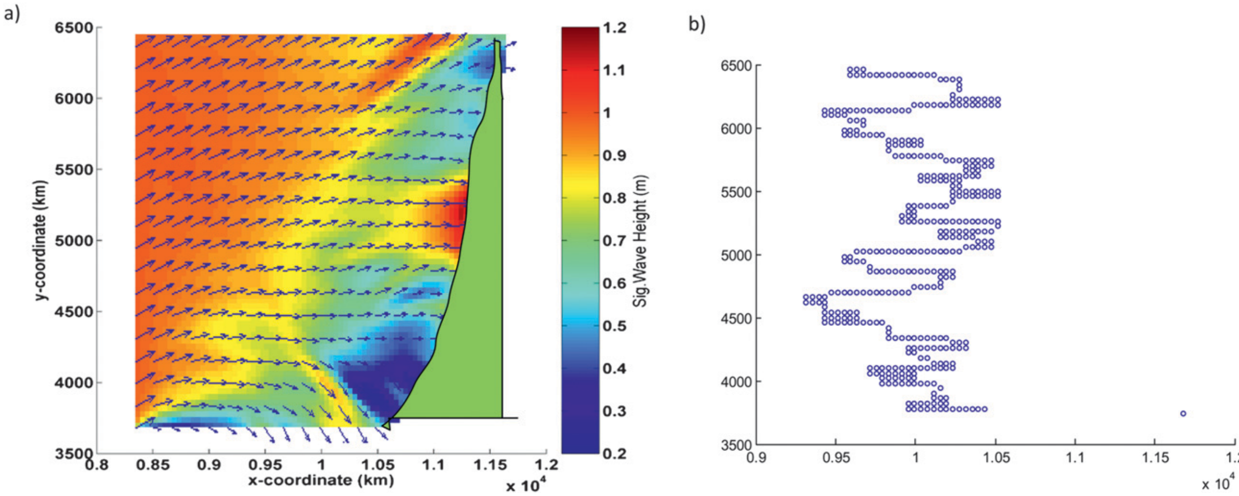

FIG. 8. Rotating original wave directional spectrum 30° toward the north: (a) wave heights and (b) optimal path.

The development detailed herein is an initial application of GA toward addressing the issue of model sensitivity to bathymetric variability and its potential influence on sampling strategies. For this study we required a standard to which ensuing model fields over iterative bathymetry interpolations would be compared, which required knowing the entire high-resolution bathymetric field in advance. However, if wave-height data were available at

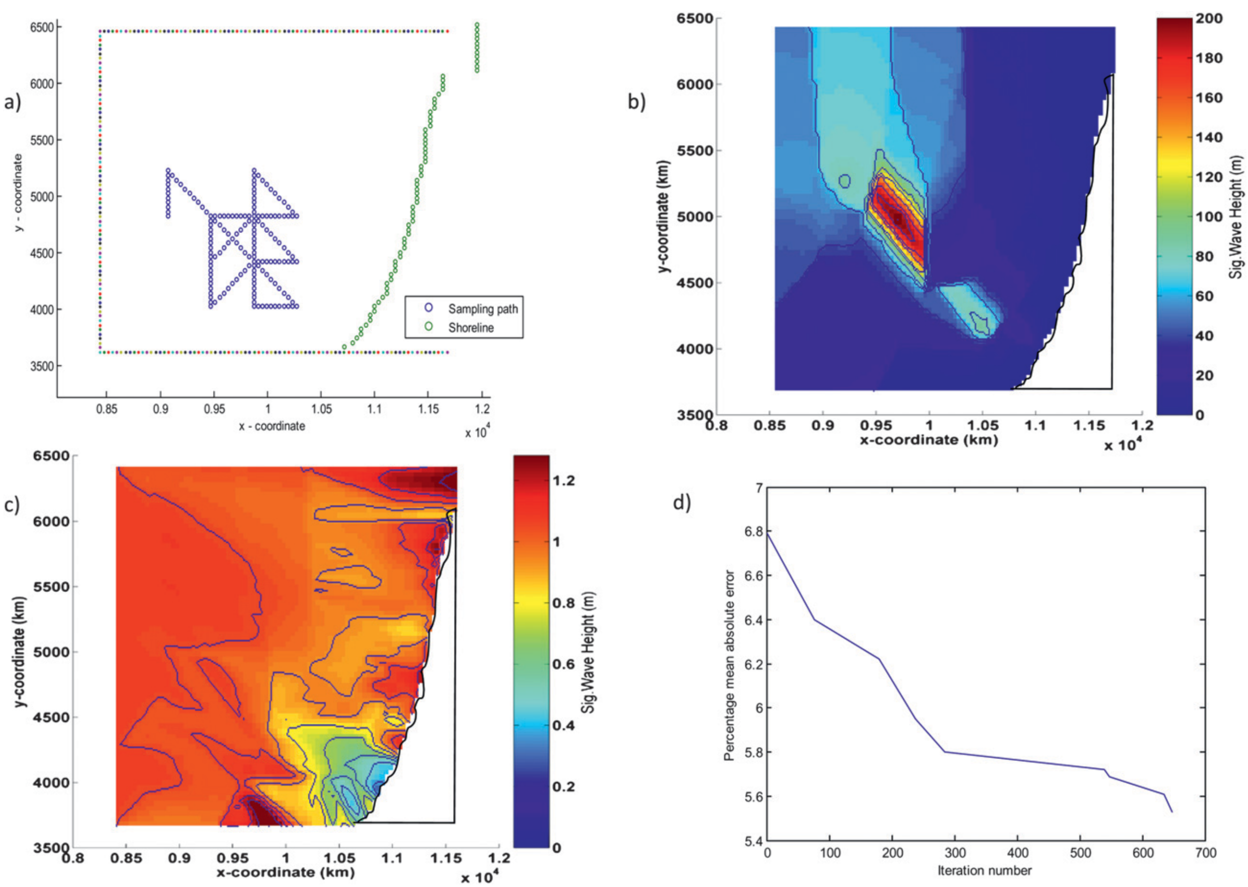

FIG. 9. NCEX case: (a) optimized survey path for specified maximum survey length of 8 km using scheme 2, (b) bathymetry interpolated from sampling along optimum path, (c) corresponding modeled wave heights, and (d) convergence plot for track length of 8 km.

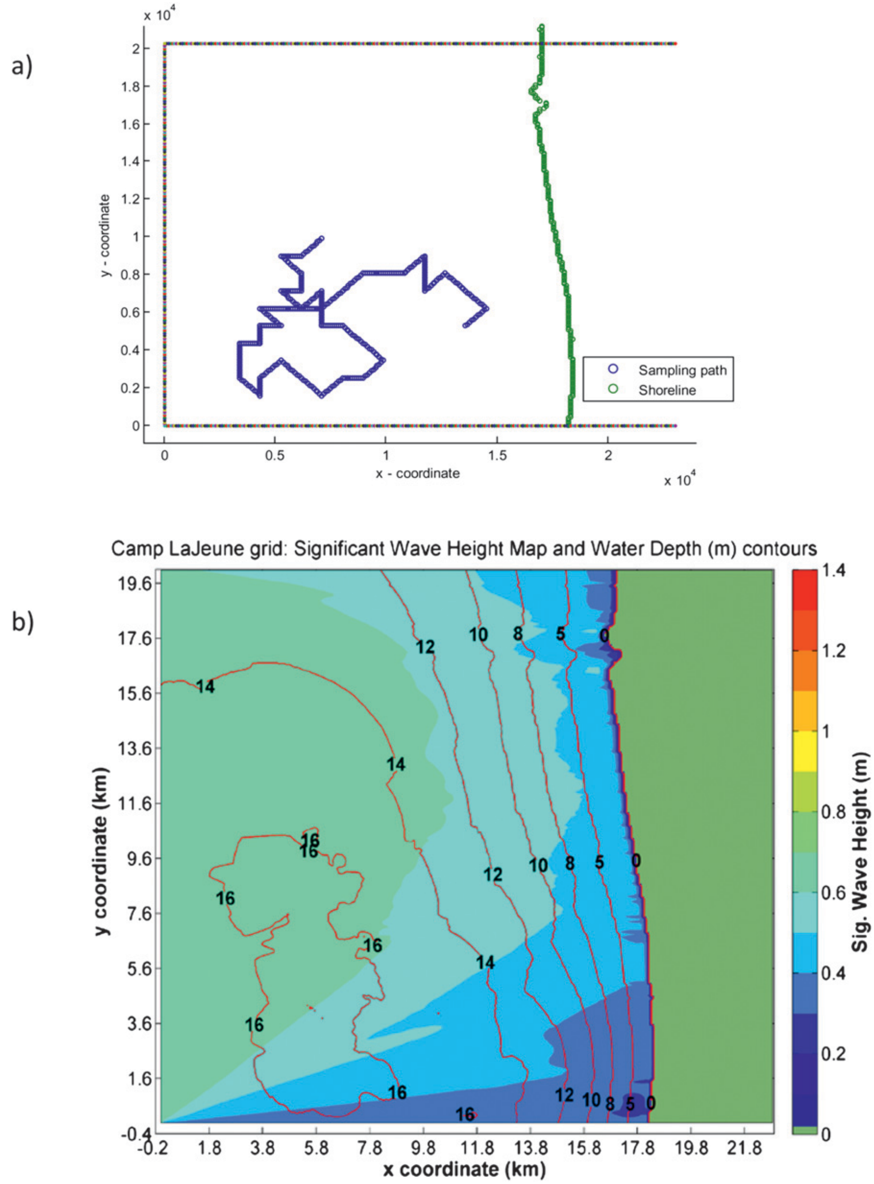

FIG. 10. Camp Lejeune case: (a) optimized survey path for maximum survey length of 55 km, using scheme 2; (b) modeled wave heights shown over bathymetry interpolated from sampling along path.

locations distributed throughout the domain, these could be used to serve the same purpose. One extension of this study would be to use data from the NCEX study and assess the robustness of the algorithm; this will be pursued in a future study.

b. Further improvements

One target of improvement is the interpolation scheme used to arrive at the model input bathymetry from the sampled one. The chosen scheme can have an impact on the solution path and needs to be carefully selected. However, the more important requirement in this case is that a consistent form of interpolation be used to evaluate all the different sampling strategies. Inverse-distance-weighted triangular interpolation was used for the purpose of this study; however, using some of the techniques of Plant et al. (2009) could potentially help improve the interpolation process used in this study.

Improvements can also be made in the error evaluation. The objective of the optimization schemes used here was to minimize the spatial average of the error in wave height, which was effectively used as a cost function. This, however, may not always be a good measure of performance for the model, and was used in this study only as a rough indicator. The objective function can then be easily redefined to be location specific, or to suit a different output parameter, or be more representative of the spatial variance of the error.

One potential issue in the practical applicability of the proposed methods is the possible redundancy in data collection that would occur by not integrating and acknowledging the acceptable data from previous surveys. To this end, a future study could employ the sediment and morphology modules of Delft3D to predict the expected changes in bathymetry over the given area during the interim period between surveys with the help of an older bathymetric dataset and long-term model forcing. This would make it possible to concentrate the search for an optimal solution path in areas that are classified as having high expected bathymetric change. Another possible extension to this work involves calculating optimum paths for various relevant incident wave conditions, then combining the ensemble of paths to determine a mean optimum path for the general wave environment for an area. This would offer some general guidance for surveys for different areas.

Acknowledgments. This work was sponsored by the Coastal Geosciences Program, Office of Naval Research, through Grant N00014-07-1-0169.

REFERENCES

Allard, R., J. Dykes, Y. L. Hsu, J. Kaihatu, and D. Conley, 2008: A real-time nearshore wave and current prediction system. *J. Mar. Syst.,* **69,** 37–58.

Alvarez, A., A. Caiti, and R. Onken, 2004: Evolutionary path planning for autonomous underwater vehicles in a variable ocean. *IEEE J. Oceanic Eng.,* **29,** 418–429.

Baehr, J., J. Hirschi, J. O. Beismann, and J. Marotzke, 2004: Monitoring the meridional overturning circulation in the North Atlantic: A model-based array design study. *J. Mar. Res.,* **62,** 283–312.

Barth, N. H., 1992: Oceanographic experiment design II: Genetic algorithms. *J. Atmos. Oceanic Technol.,* **9,** 434–443.

——, and C. Wunsch, 1990: Oceanographic experiment design by simulated annealing. *J. Phys. Oceanogr.,* **20,** 1249–1263.

Bellingham, J. G., and J. S. Wilcox, 1996: Optimizing AUV oceanographic surveys. *Proc. Symp. on AUV '96,* Monterey, CA, IEEE, 391–398.

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

Bourgeois, B. S., K. T. Holland, and J. M. Kaihatu, 2005. Field experiment report for the AUV-Fed Nearshore Nowcast Project, La Jolla, CA, REMUS bathymetry survey, Oct. 13-18, 2003. Naval Research Laboratory Formal Rep. NRL/FR/ 7440—05-10095, 31 pp.

Chen, Q. J., and I. A. Svendsen, 2003: Effects of cross-shore boundary condition errors in nearshore circulation modeling. *Coastal Eng.,* **48,** 243–256.

Coastal Data Information Program, cited 2009: The Coastal Data Information Program. [Available online at http://cdip.ucsd.edu/.]

Elgar, S., cited 2003: Nearshore Canyon Experiment (NCEX). [Available online at http://science.whoi.edu/users/elgar/NCEX/ncex.html.]

Fox, R., A. Garcia, and M. L. Nelson, 1999: A three dimensional path planning algorithm for autonomous vehicles. *Proc. 11th Int. Symp. on Unmanned Untethered Submersible Technology,* Durham, NH, Autonomous Undersea Subsystems Institute, 546–556.

Goldberg, D. E., 1989: *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., 432 pp.

——, and K. Deb, 1991: A comparative analysis of selection schemes used in Genetic Algorithms. *Foundations of Genetic Algorithms,* G. J. E. Rawlins, Ed., FOGA, Vol. 1, Morgan Kaufmann Publishers, 69–93.

Holland, J. H., 1975: *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 183 pp.

Kaihatu, J. M., and W. C. O'Reilly, 2002: Model predictions and sensitivity analysis of nearshore processes near complex bathymetry. *Proc. Seventh Int. Workshop on Wave Hindcasting and Forecasting,* Banff, Alberta, Canada, Environment Canada and U.S. Army Corps of Engineers, 353–360.

Komen, G. J., L. Cavaleri, M. Donelan, K. Hasselmann, S. Hasselmann, and P. A. E. M. Janssen, 1994: *Dynamics and Modeling of Ocean Waves*. Cambridge University Press, 556 pp.

Lesser, G. R., J. A. Roelvink, J. A. T. M. van Kester, and G. S. Stelling, 2004: Development and validation of a three dimensional morphological model. *Coastal Eng.,* **51,** 883–915.

Plant, N. G., K. T. Holland, and J. A. Puleo, 2002: Analysis of the scale of errors in nearshore bathymetric data. *Mar. Geol.,* **191,** 71–86.

——, K. L. Edwards, J. M. Kaihatu, J. Veeramony, L. Hsu, and K. T. Holland, 2009: The effect of bathymetric filtering on nearshore process model results. *Coastal Eng.,* **56,** 484–493.

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.