Toward Enhanced Environmental Effects Representations in Advanced Computer Simulations

Fred C. Newman, Albert C. Biondo, Michael D. Mandelberg, Arthur R. Croucher II, James C. Spall, Cathy C. Matthews, and J. Thomas Warfield

As part of the Defense Advanced Research Projects Agency's Advanced Simulation Technology Thrust, the Multiresolution Interaction Validity (MIV) Project at APL is investigating methodologies to employ environmental effects models more effectively and efficiently, to aid the development of synthetic force models, and to ensure internal consistency among multiple environmental effects models. The use of cluster analysis has proven to be a powerful tool in identifying important sensitivities of environmental effects models such as ASTRAL, a Navy model for the propagation of underwater sound. Applications of this work will be found in future simulations for both training and analysis. This article describes recent accomplishments as well as an overview of the MIV Project. (Keywords: Acoustics, Cluster analysis, Environment, Interoperability, Modeling, Propagation, Simulation.)

INTRODUCTION

In advanced distributed simulations, the synthetic natural environment (SNE) provides the "playing field" on which synthetic forces will operate. SNEs locate those forces with respect to each other and mediate their interactions (e.g., terrain and atmospheric influences on intervisibility). The SNE encompasses all physical aspects of the natural environment, including the land, sea, atmosphere, and space.

Modeling the complex effects of the natural environment on warfighters and systems, however, requires prioritizations and trade-offs that can result in inconsistent or reduced-fidelity representations across networked simulations. Inconsistent representations can occur for several reasons: use of different models that are not entirely consistent with each other to represent the same environmental effects in different interacting simulations; use of different implementations of the same environmental effects model; and use of inconsistent environmental data in the environmental effects models.

Reduced-fidelity representations can result from suboptimal employment of environmental effects models. These models often involve complex algorithms to represent effects of the environment on some process such as the propagation of signal energy through the atmosphere or the ocean or the passage of a military platform over heavy seas or a muddy field. In situations where environmental effects models must be updated frequently to keep pace with simulation dynamics (e.g., movement of sources and sensors), the computational complexity of the models can require a disproportionately large share of the computer resources available to the simulation. To reduce the number of required calculations and improve the model's cost-effectiveness, estimates of model output may be employed to reduce actual model calculations.

Recognizing these problems, the Defense Advanced Research Projects Agency decided to include a task area on the natural environment in their Advanced Simulation Technology Thrust. The Multiresolution Interaction Validity (MIV) Project is part of this thrust, and was begun in mid-FY97. The project is developing prototype methodology as well as prototype guidance products to support the enhanced representation of environmental effects in networked computer simulations. The three main goals of the project are (1) development of use strategies for improving computational efficiency, (2) abstractions of a given model's response to environmental features that can guide developers of synthetic force models, and (3) consistency analysis for evaluating correlation of environmental effects across networked simulations.

MIV PROJECT OVERVIEW

Figure 1 is a conceptual model of the MIV Project. Ultimately, this methodology will be applied to multiple environmental domains, but initially it will be developed in the ocean acoustic domain. The approach begins with the selection of a specific real-world littoral environment that will be the basis for environmental inputs to an environmental effects model. The selected real-world environment is examined for critical environmental factors (e.g., fronts, eddies, surface mixed layers, seamounts, etc.) that significantly impact the environmental effects model output. This environment is then simplified into various reduced-complexity forms. The most basic is a baseline environment, which is the simplest acceptable representation of the area under study. Next, simple intermediate SNEs (ISNEs) are constructed, each consisting of a single environmental feature added to the baseline environment. Finally, progressively more complex ISNEs consisting of different assortments of multiple environmental features are considered, converging to the real-world environment.

Each of these environments becomes a case study for the environmental effects model. Other model inputs represent operational factors such as source and sensor locations and source frequency. The resulting sets of model output (or model response) are then analyzed in several exploratory ways. Cluster analysis is the main analysis method described here, although others have been investigated as well.[1] The results of these analyses form the basis for development of the MIV products.

The three MIV products are model use strategy, model-dependent environmental abstractions, and consistency analyses. The goal of model use strategy is

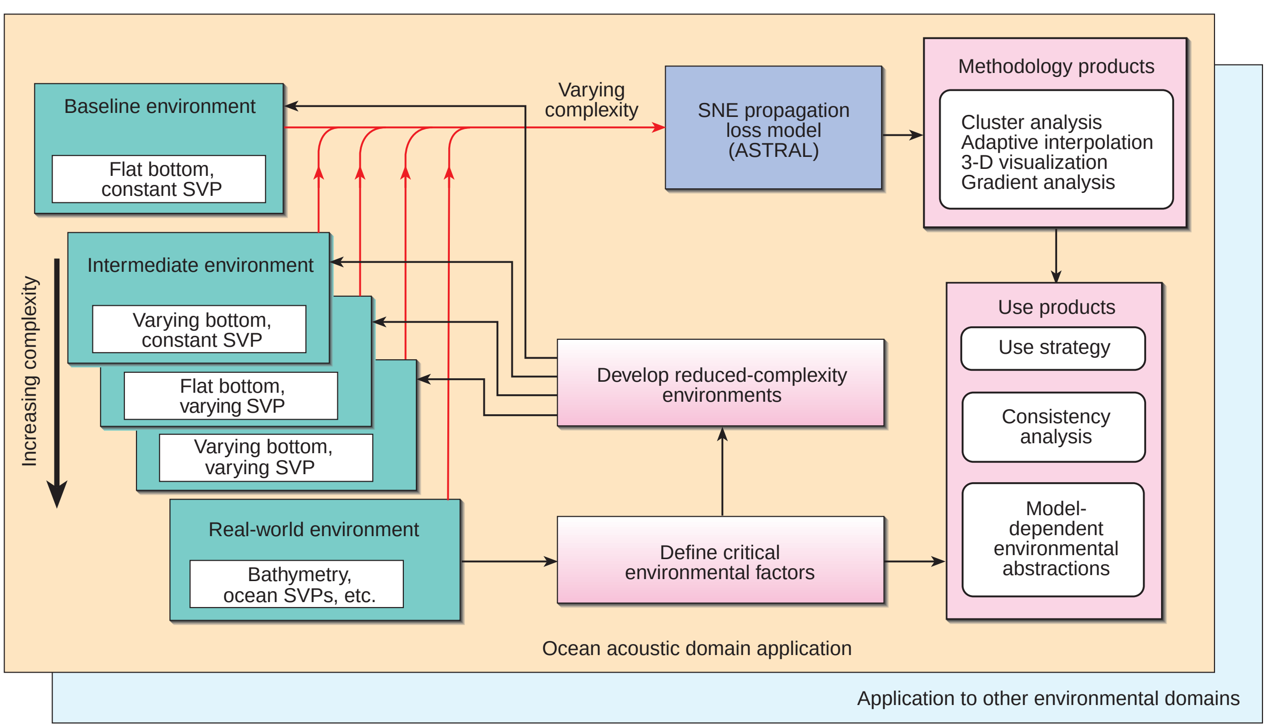

Figure 1. Conceptual model of the Multiresolution Interaction Validity (MIV) Project (SVP = sound velocity profile, SNE = synthetic natural environment).

the development of more computationally efficient representations of environmental models. A model use strategy is defined as the conditions for which an environmental effects model calculation will be performed or the conditions for which one or more methods will be exercised to estimate that model's output response.

Model abstractions developed for use by synthetic-force model developers have come to be called model-dependent environmental abstractions. They are model-dependent descriptions of the effects of a specified environmental feature on the output response of a given environmental effects model. Synthetic forces are automated models used to represent system and/or human resources in large simulations, and must be sufficiently realistic to satisfy the objectives of the simulation.

Consistency analysis assists the intercomparison of environmental effects model outputs from different models, different implementations of the same model, or the same model receiving different environmental inputs describing the same environment (e.g., varying grid sizes, feature descriptions). Model consistency within an advanced simulation is essential to the use of advanced simulations for training and analysis, yet is problematic when many participants bring their own models to the simulation.

The substrate on which these products are constructed is the representation of the SNE, including both models and data. Thus the approach to this problem begins there.

SNE REPRESENTATION

Area Selection

Areas of potential military interest were identified using guidance provided by the Program Executive Officer/Undersea Warfare Advanced Systems and Technology Office for another project,[2] as well as a knowledge base of APL's past and current work in support of the Navy and other government sponsors. The area best suited for this study was the Sea of Japan (SOJ). Types of bathymetric features found in the SOJ include a continental shelf, continental slope, abyssal plain, spreading ridge, and seamounts. Primary water column features are fronts and surface mixed layers. The SOJ study area and a typical trackline are shown in Fig. 2.

Ocean Model Selection

Princeton Ocean Model (POM) data sets for the SOJ were available from the Naval Research Laboratory, Stennis Space Center, Mississippi (personal communication, M. Carnes, May 1997). The POM is a free-surface, primitive-equation ocean model that includes a turbulence submodel and has depth scaling that is proportional to local water depth. It was developed in the late 1970s by Blumberg and Mellor,[3] with subsequent contributions from others. The model has been used for modeling of estuaries, coastal regions, and open oceans, and currently has hundreds of registered users from many different countries. The particular POM data set used in this investigation was designed to capture mesoscale features (≈40 km), which, though lacking the smallest scale features, nonetheless captures enough of the natural variability for the present study.

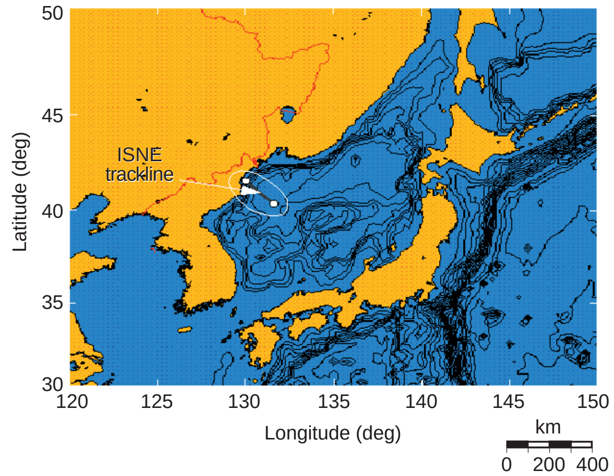

Figure 2. Sea of Japan study area, showing a typical intermediate synthetic natural environment (ISNE) trackline.

Underwater Acoustic Propagation Model

The ASTRAL Underwater Acoustic Propagation Model (Version 4.2) was selected to calculate the acoustic transmission losses (TLs) in this analysis. Requirements for the acoustic model were that it be a Navy standard model and commonly used in Navy simulations. ASTRAL is a Navy standard model and is available in the Navy's Ocean and Atmospheric Master Library.[4] ASTRAL is widely used in real-time simulations because it runs 10 to 1000 times faster than the more accurate parabolic equation models. It is a range-dependent, adiabatic, range-smoothed, mode-theory model with additional algorithms to model important acoustic features such as convergence zones and surface ducts that are not appropriately scaled for the primary algorithm.

ISNE Plan

Identification of Features

A comprehensive review of the SOJ was conducted to identify all the oceanographic and bathymetric features and to assess their acoustic significance. The results of this analysis are reported in Ref. 5. The individual features are listed in Fig. 3.

The baseline environment consists of a range-independent sound velocity profile that was taken from the Navy's General Digitized Environmental Model and is representative of summer SOJ sound velocity profile structures and a range-independent abyssal plain bottom. Using full-resolution POM data, we constructed single-feature ISNEs for each feature shown in Fig. 3, and multiple-feature ISNEs for the sequence shown in red lines. All single-feature ISNEs are built up as modifications or additions to the baseline environment.

Selected Tracklines

The POM data were reviewed to locate specific examples of the features expected in the SOJ. Once the features were found, a trackline through them was defined for the ISNE. These tracklines were the basis for propagation runs using ASTRAL. The ISNEs analyzed to date were chosen with the expectation that using simple environments would simplify the initial analysis.

The Data Set

Twenty-six standard frequencies ranging from 20 Hz to 10 kHz in an approximately logarithmic scale were selected, as shown in Table 1. In addition, 17 standard depths were used for both the source and receiver, covering a range from 50 to 1500 ft. Thus, for each ISNE, TLs were calculated for each combination of frequency, source depth, and receiver depth shown in Table 1, resulting in 7514 TL curves per trackline. ASTRAL output consists of TL as a function of range of the receiver from the source. Range increments of 0.25 nmi were used for these calculations.

On a Pentium Pro 200-MHz system, the run time for a single geometry was 0.25 s. Combining multiple frequencies and source depths into a single run eliminated redundant setup calculations and optimized disk access, reducing the run time to 0.03 s per geometry. The resulting run time for an entire ISNE was approximately 5 min.

The ASTRAL output generated for a single ISNE comprises a 10-MB, four-dimensional array (range, frequency, source depth, and receiver depth), containing 3,005,600 TL values. TL data sets of this size and dimensionality are not commonly analyzed, even by experts, making this data set a unique resource within the community.

ANALYSIS OF PROPAGATION MODEL DATA

Several analysis methods including gradient analysis, interpolation, and Markov estimation methods have been examined and reported in the MIV Project.[1] The most promising method developed so far, cluster analysis, is described in the following subsection.

Cluster Analysis

Development of use strategy, model-dependent environmental abstractions, and consistency analysis requires detailed knowledge of model sensitivity to the input parameters. All three of these products depend on the observation that there are often broad regions in the multidimensional input parameter space (source depth, receiver depth, and frequency) where the model output is relatively insensitive to small variations in some or all of the input parameter values. Within these domains, we should be able to achieve good representation of the model response without having to repeat the full set of model calculations at each point. Hence, we seek to identify and map out these regions.

Cluster analysis is a multivariate statistical procedure for detecting natural groupings in data. These groupings, or clusters, of model output data can be mapped back to input parameter intervals, and therefore identify those regions where model output is relatively insensitive to small variations in input parameter values.

This analysis uses a hierarchical divisive cluster technique based on the method devised by Macnaughton-Smith et al.[6] The method is illustrated in Fig. 4. First, all objects are in a single cluster and then

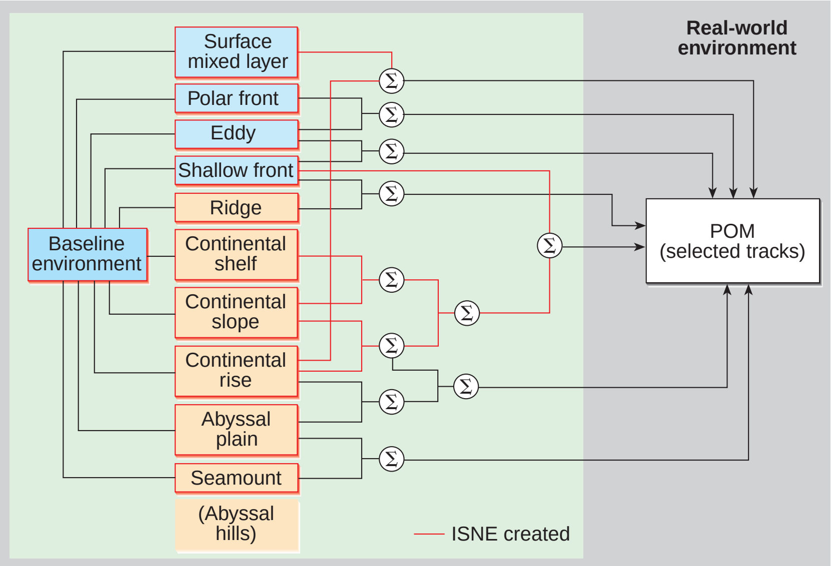

Figure 3. Sea of Japan ISNE plan (POM = Princeton Ocean Model).

Table 1. Standard values selected for the ISNE data set.

Frequencies (Hz)	Source depths (ft)	Receiver depths (ft)
20	50	50
32	100	100
40	150	150
50	200	200
64	250	250
80	300	300
100	350	350
160	400	400
200	450	450
250	500	500
320	600	600
400	700	700
500	800	800
640	900	900
800	1,000	1,000
1,000	1,250	1,250
1,250	1,500	1,500
1,600		
2,000		
2,500		
3,200		
4,000		
5,000		
6,400		
8,000		
10,000		

the data set is systematically split into smaller clusters. This process would ultimately result in each object belonging to a separate cluster, but in practice it is terminated when sufficient clustering is achieved. Examination of how clusters are split assists in determining the appropriate number of clusters to characterize the structure in a data set.

After all objects are placed in a single cluster to be divided into two smaller ones, the divisive algorithm proceeds by examining the average dissimilarity for each object in a cluster. The dissimilarity between two objects is simply the Euclidean distance between them, i.e., the geometric distance d_{rs} between any two vectors **r** and **s**, defined as

$$d_{rs} = \sqrt{\frac{1}{m}\sum_{i=1}^{m}(x_{ri} - x_{si})^2}\,, \tag{1}$$

where x_{ri} is the value of the ith element in the vector **r**, x_{si} is the value of the ith element in the vector **s**, and m is the dimension of the vectors **r** and **s**. A small distance indicates that the objects are similar, whereas

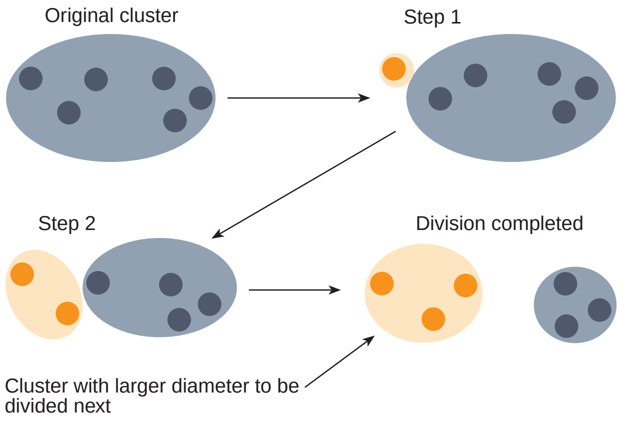

Figure 4. Diagram of the divisive hierarchy algorithm.

a large distance indicates dissimilarity. The average dissimilarity of an object is then just the ensemble average of the dissimilarity of that object from each of the other objects in the cluster. The diameter of a cluster is simply the maximum dissimilarity between any two objects in the cluster.

The object with the largest average dissimilarity is removed from the original cluster and placed into a new cluster as a starter object. The average dissimilarity is then recomputed for each of the objects remaining in the original cluster. These are rank-ordered by their average dissimilarity, and the object with the largest average dissimilarity is considered next. The average dissimilarity of this object to the current object(s) in the new cluster is also computed. If this object is "closer" to the new cluster than to the original cluster (i.e., smaller average dissimilarity relative to the new cluster than to the original cluster), then it is placed in the new cluster. The process is repeated iteratively until no further objects are "attracted" to the new cluster. Once this process is complete, the cluster with the largest diameter is selected for division next.

When the clustering is complete, representative objects are identified for each cluster. The representative object for a given cluster is the object with the smallest average dissimilarity in that cluster; note that it is an actual object from a cluster, rather than a calculated value or average of some number of objects.

Application to Use Strategy Development

The plan, then, is to use a single representative object where possible, instead of the multitude of actual objects. That this works is a major result of the project. Recall that we seek to find input parameter intervals within which the model output may be reliably estimated without calculating the actual TL. To this end, the divisive cluster analysis algorithm was applied to the TL calculated by ASTRAL in the baseline environment. The goal was to see if the clusters would form regions in the parameter space that would enable a use strategy to be defined. As part of this process, three input parameters were examined: frequency, source depth, and receiver depth. Each individual TL-versus-range calculation (for a given frequency, source depth, and receiver depth combination) was considered to be a single object.

The total number of TL curves to be clustered for the baseline environment was 7514. Since the full clustering calculation takes several days on a PC, a sampling approach was taken here. For an initial test, a random sample of 100 TL curves was chosen from the full set of 7514, and this sample was then clustered. Representative TL curves were selected from each cluster. The full clusters were then approximated by associating each unsampled curve from the full set with the nearest cluster, i.e., the cluster with the smallest Euclidean distance between its representative curve and the unsampled curve.

The TLs analyzed included all the source depths, receiver depths, and frequencies shown in Table 1. The entire data set was reduced to 20 clusters in order to determine if the clustering would yield continuous, connected volumes in the three-dimensional source depth-receiver depth-frequency space. If so, the method would hold promise. But if the clusters were to consist of the union of many small, isolated regions in input space, it would not. The clusters were indeed found to represent continuous, connected volumes in the three-dimensional space, and did not consist of disconnected points. Two of the clusters are shown in Fig. 5a, and the corresponding regions in input parameter space are shown in Fig. 5b. We concluded from these studies that cluster analysis can find volumes in the three-dimensional (source depth, receiver depth, frequency) space that exhibit similar TL characteristics.

Cluster Analysis Toolbox

To help reduce the amount of required analysis of the TL calculations, a graphical user interface (GUI) was developed using Matlab 5.2. The GUI implements the cluster algorithm and displays the results in a user-friendly format. It enables the user to display the actual TL in a variety of ways as an aid in analysis. A typical display from the toolbox is shown in Fig. 6. The representative TL curves for the 20-cluster case are shown in Fig. 6a. The size of each cluster (i.e., the number of TL curves per cluster) is shown in Fig. 6b. One can see that the number of TL curves in each cluster varies considerably, implying that a large part of the input parameter space can be represented by a small number of TL curves.

The cluster analysis toolbox enables a user to quickly examine precalculated TLs for a large range of source depths, receiver depths, and frequencies. It aids in identifying areas in the parameter space that have similar TL characteristics and in finding representative TL curves for those areas. The toolbox can also provide information on the appropriate number of clusters, based on the amount of within-cluster variation that is acceptable. Although presently configured for TL analysis, the toolbox is easily reconfigurable and extendible to other parameter spaces. This toolbox is now available for use by others.

THEORETICAL INVESTIGATIONS OF USE STRATEGIES

A parallel effort is also under way on developing an analytical basis for the efficient use of models so that run time may be dramatically reduced. Essentially, we must rapidly produce an output that closely matches the physics-based model output as the input variables change. Such speedups will be required to run many models effectively in real time, e.g., in a simulation of acoustic propagation for submarine training exercises.

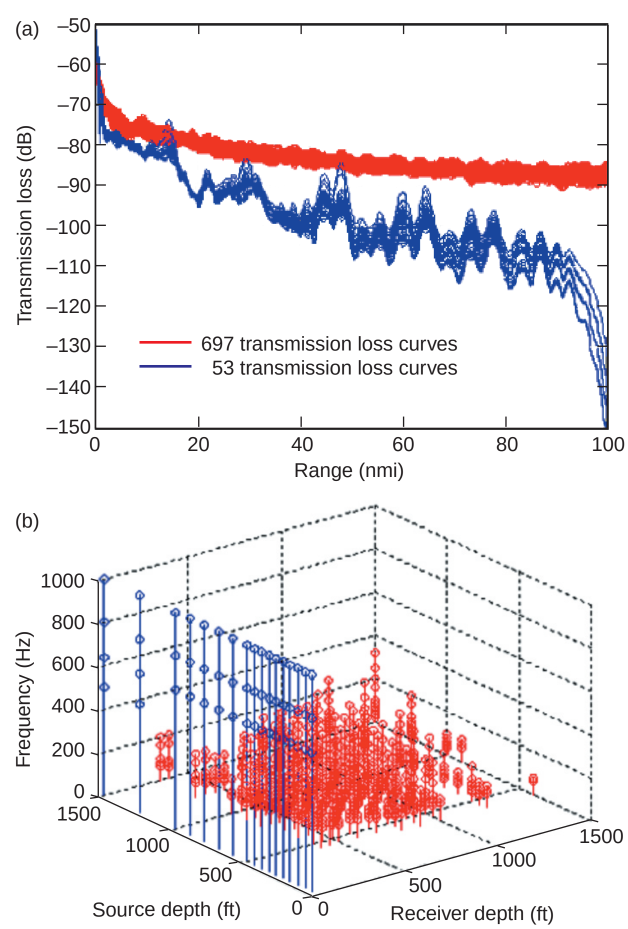

Figure 5. Plots of clusters representing continuous, connected volumes in three-dimensional space: (a) transmission loss for 2 (of 20) clusters, and (b) location in input space of the 2 clusters.

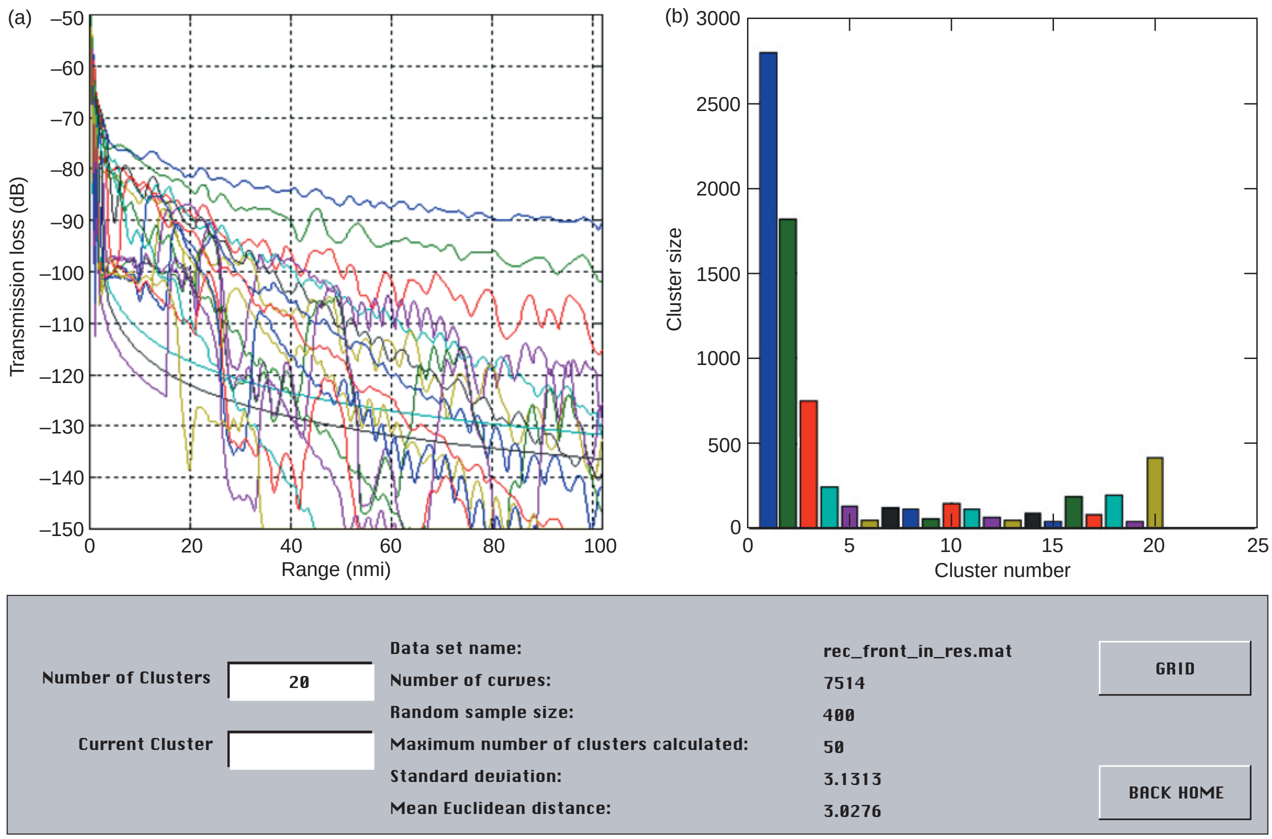

Figure 6. Example display from the APL cluster analysis toolbox showing (a) transmission loss for two clusters and (b) size of the clusters.

This task will entail intelligent use of a relatively small number of full physics-based model runs. Two of the many issues regarding these runs are how small a "seedling" set of model runs will suffice and how to disperse them throughout the parameter space. Although we are using ASTRAL as a motivating example, the basic ideas would apply to any physics-based model, including models relevant to land- and air-based combat settings.

Overview of Analytical Approach

The means to accomplish this task is the use of selected runs of the physics-based code to build an efficient interpolator. The near-term aim of this project is to develop means by which the ASTRAL code can be interpolated with respect to the underlying input variable space of relevant terms (likely to be 10 or more variables, including the relative spatial coordinates of the submarine and signal source, water salinity, ocean-bottom shape, sound-speed profile, etc.) as these relevant terms are changing. Hence, in a training exercise, for example, a relatively small number of runs of ASTRAL would be made at strategically chosen points in the multivariate input variable space, with interpolation used to "fill in" the remaining space and provide the real-time output. This interpolator will run much faster than actual ASTRAL runs, and hence can be used to provide good real-time acoustic transmission estimates.

The existing literature on multivariate interpolation methods provides little guidance on feasible methods for picking the evaluation points in a dynamic setting. The lack of guidance in this area presents a challenge to developing an overall interpolation approach.

We are currently looking at analytical ways to efficiently and effectively determine these "strategically chosen" points using advanced optimization techniques. A naive choice of evaluation points for building the interpolator will be hopeless except in the smallest problems. A complication for the optimization process of finding an intelligent grid of points for interpolation is that no loss-function gradients will be available because the code is much too complex to differentiate analytically or even numerically. Hence, the optimization method must rely entirely on simulation inputs and outputs.

Our approach is based on the well-known information matrix from estimation theory. The information matrix plays a central role in the practice and theory of statistical estimation. It provides a summary measure of the amount of information in the data relative to the parameters that are being investigated. Some of its applications are confidence-region calculations for parameter estimates, determination of inputs in experimental design, estimation of performance bounds in adaptive systems (such as a control system), and uncertainty bounds on predictions (such as with a neural network). However, calculation of the information matrix is often a difficult and lengthy task, especially with nonlinear models like neural networks.

A fast approximation to the information matrix is being developed in this project. This will aid in determining where the evaluation points should be placed to produce an accurate interpolation.

Calculation of the Information Matrix

Our work on this problem[7] consists of a resampling-based method for computing the information matrix. This method applies in problems of any level of complexity and is relatively easy to implement. Analytically, the matrix is defined as

$$F_n(\theta) = -\frac{1}{n} E\left(\frac{\partial^2 \log L}{\partial\theta \partial\theta^T} \,\Big|\, \theta \right), \tag{2}$$

where the scalar L is the statistical likelihood function (the probability density function of the output values expressed as a function of the input vector θ), n is the number of data points, and $E(\cdot)$ represents a statistical expected value with respect to all the randomness in the problem. In practice, it is common that neither the indicated second derivative matrix (i.e., the Hessian matrix) nor the expectation can be computed, or at least not easily. This is especially the case with nonlinear models such as those likely to be used as interpolators.

The essence of our method is to produce efficient, "almost unbiased" estimates of the Hessian matrix of $\log L$ at a large number of pseudo data vectors and then average these to obtain an approximation to $F_n(\theta)$. This approach is directly motivated by the definition of $F_n(\theta)$ as the mean value of the Hessian matrix of the function $\log L$. The pseudo data are generated according to a bootstrap resampling scheme treating the chosen θ as "truth." That is, each pseudo data set represents a sample of size n from the assumed distribution of data based on the unknown parameters taking on the chosen value of θ. The critical part of this conceptually simple scheme is efficient estimation of the Hessian matrix.

This method is surprisingly simple and is based on ideas from simultaneous perturbation stochastic approximation.[8] Because the simultaneous perturbation Hessian estimate is easy to compute and because the information matrix estimate is a simple average of the Hessian estimates, the approach provides a powerful and easy-to-implement means for calculating the information matrix in arbitrarily complex models or interpolators.

Two applications of the information matrix estimate having special relevance to the use-strategy part of the MIV Project are picking the input values for use in building the interpolator and producing error bounds on the prediction errors for the interpolator. Although the general approach is reasonably clear and consistent with previous work in experimental design, our work is continuing on the details of the information matrix estimate in specific cases.

PRODUCT DEVELOPMENT

Use Strategy

The previous discussion of cluster analysis suggested that alternative approximate representations of specific model responses are indeed feasible. Therefore, our first prototype use strategy will be developed for the SOJ baseline environment, based on the cluster analysis results. This strategy consists of selecting one of the representative TL calculations that describes a cluster, based on the input parameter values of the desired situation, in lieu of actually performing a specific ASTRAL calculation. The number of clusters required for a particular simulation is determined by the fidelity requirements of the simulation. A larger number of clusters produces a higher-fidelity representation of the acoustic environment than a smaller number of clusters by reducing the within-cluster variation. The degree of variability within each cluster is shown in Fig. 7 as a function of the total number of clusters. When the entire TL data set is separated into two clusters, the mean Euclidean distance between the representative TL curves and all other curves in its cluster is over 7 dB. As the number of clusters increases to 20, the mean Euclidean distance decreases to approximately 3 dB.

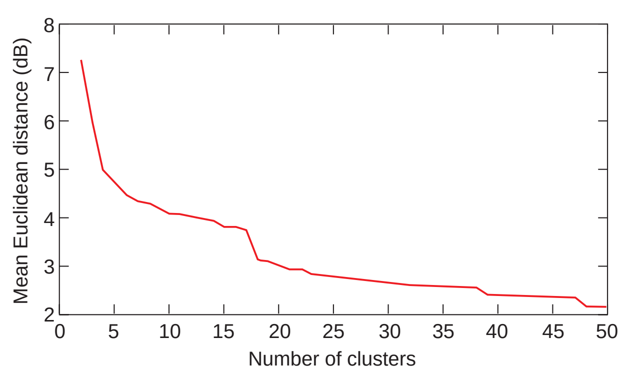

Figure 7. Accuracy curve for prototype use strategy.

Another use strategy being investigated is interpolation. The cluster analysis showed that there were large parts of the input parameter space in which the variability of TL was not great. A simple experiment was conducted to assess the feasibility of interpolation. Transmission loss curves for a constant source depth (200 ft) and for all frequencies were calculated for receiver depths every 10 ft. These model-calculated TL curves were then compared with interpolated TL curves. Transmission loss was interpolated from four receiver depths. Initially, these were approximately evenly spaced at 50, 500, 1000, and 1500 ft. The interpolated TL curves were then compared to the ASTRAL-calculated TL curves. This comparison is shown in Fig. 8a. At receiver depths greater than 500 ft and frequencies less than 6400 Hz, the interpolation errors were less than 1 dB. At depths less than 500 ft, the errors increased to over 20 dB. At frequencies higher than 6400 Hz, there were bands of low and high interpolation errors.

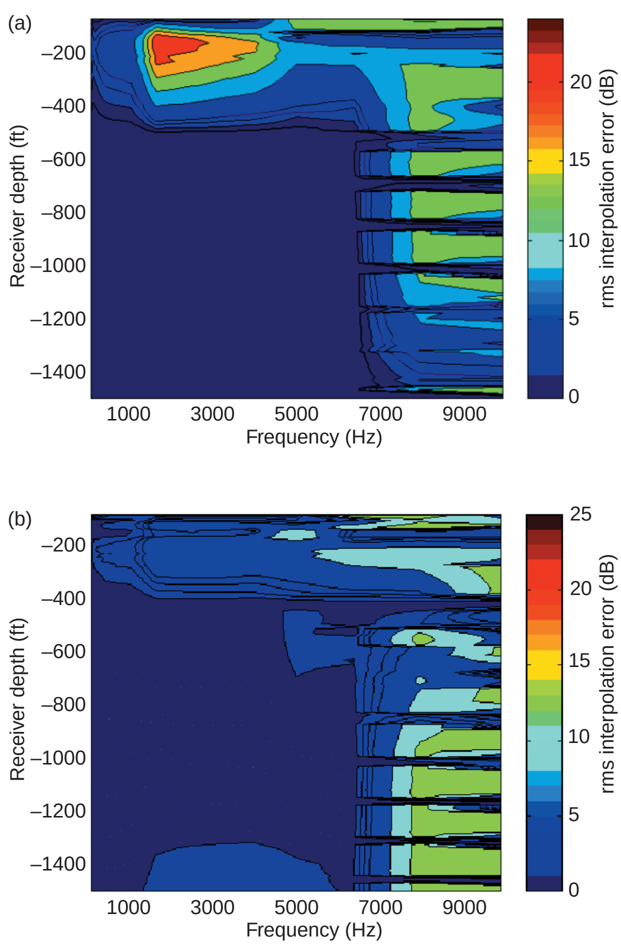

Figure 8. Accuracy contour plots for proposed interpolation use strategy at (a) receiver depths of 50, 500, 1000, and 1500 ft, and (b) adjusted receiver depths of 60, 120, 400, and 1000 ft (rms = root mean square).

The size of the individual clusters varied considerably, indicating that the interpolation could be improved by varying the depths from which to interpolate. Cluster analysis showed that the TL varied more quickly at shallow receiver depths than at deeper receiver depths. Therefore, receiver depths of 60, 120, 400, and 1000 ft were chosen to interpolate from. The differences between the interpolated and the model-calculated TL are shown in Fig. 8b. The interpolation errors at the depths shallower than 500 ft have been reduced from over 20 to 4 dB. The interpolation error for receiver depths deeper than 1300 ft have increased only slightly. This potential use strategy increased the fidelity of the TL without increasing the computational load on the simulation.

One distinct advantage of such a strategy is the order-of-magnitude decrease in storage capacity required to save the representative cluster TL sets versus storing 7514 individual TL files. In addition, the selection of a representative TL set in place of an actual calculation is based on a more nearly optimal choice than selection from a conventionally precomputed table.

Model-Dependent Environmental Abstractions

In developing model-dependent environmental abstractions, a model is exercised for a variety of preselected environmental conditions and sensor parameter inputs. The results of the model runs are then examined to determine (1) if the model response is appropriate for the inputs provided, and (2) how the model response can be characterized for the conditions presented. Initially, single-feature ISNEs were used in this study to increase the likelihood of correlating the model response to an identifiable environmental element. As our capability matured, the model was subjected to increasingly more complex ISNEs to examine the model response to multiple features. The goal is to provide compact representations of propagation that can be easily incorporated in the (relatively) simple models characteristic of synthetic forces.

The first environment investigated was the baseline environment, in which the sound speed profile was simplified to a nearly bilinear profile, and the bottom was taken to be flat and horizontal. Propagation is well understood in this environment, and we were able to provide good descriptors. Next we examined single-feature ISNEs, where we were also able to establish some simple propagation descriptors. The challenge was describing the effects of two or more environmental features, and we have succeeded in developing a corresponding ASTRAL model-dependent environmental abstraction that would provide a concise description of propagation behavior for these ISNEs over the entire operational parameter space.

Figure 9 depicts the range of characteristic TL behaviors along the trackline shown in Fig. 2. This trackline was dubbed the "front/rise/slope" ISNE since it crosses an oceanographic front, then a bathymetric rise, and finally terminates on a continental slope. The mechanisms giving rise to the front/rise/slope TL behaviors in Fig. 9 are understood,[1] and the four characteristic types of TL behavior shown there could provide a preliminary characterization of acoustic propagation in the SOJ baseline environment for synthetic force acoustic sensor models. (The black curve shown occurred in the analysis and was found to be due to a high-frequency anomaly in ASTRAL rather than to any real environmental conditions.)

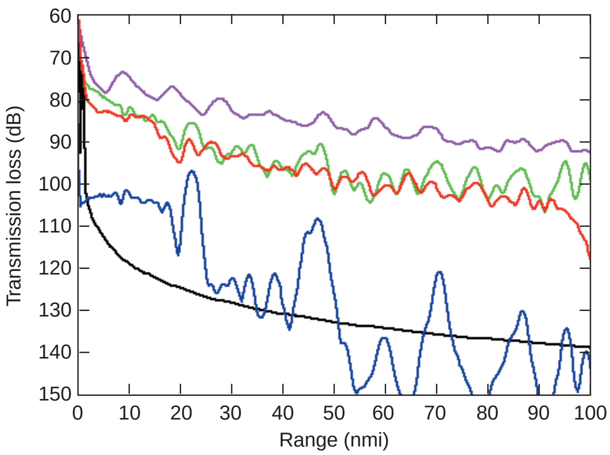

Figure 9. General TL curves that characterize the front/rise/slope ISNE. The black curve is an ASTRAL anomaly and is not due to any environmental conditions.

Consistency Analysis

In simulations, different participants may employ diverse propagation models of varying degrees of complexity and fidelity. To avoid legislating the universal use of some single model—an undesirable approach to distributed simulation—we must develop methods by which two or more propagation models, independently using environmental information (which may or may not be identical in resolution or format), can be compared with relatively little effort. The objective of consistency analysis is to help ensure that the different propagation or other environmental effects models used in the simulation are providing sufficiently consistent results to satisfy the objectives of the simulation exercise.

Two acoustic TL models were considered in this context: ASTRAL, the model chosen for the overall MIV effort, and the parabolic equation (PE) model, another Navy standard TL prediction model. The latter is a fully coherent wave-theoretic solution to the acoustic wave equation. Although sufficiently accurate to serve as a "benchmark" for other TL models, the PE model can be computationally intensive, particularly at higher acoustic frequencies (> 500 Hz). A hypothetical question was considered: Could the PE model and ASTRAL be used consistently in the same simulation?

The slope/rise/front ISNE was chosen, and identical environmental information and tracklines were used in both models. Because of the high computational burden required by the PE model at higher frequencies, we limited this analysis to frequencies of 1 kHz and less, noting that the methodology would generalize to higher frequencies as well.

We decided to use the ASTRAL output cluster analysis to determine the system parameter values (source depth, receiver depth, and frequency) associated with each representative TL curve, and then to exercise the PE model for those system parameter values and environmental inputs. This produced a set of TL curves representing the range of ASTRAL TL behaviors, and by assumption (vice by cluster analysis), a set of TL curves representing the range of PE TL behaviors.

An overplot of the 20 representative curves from the ASTRAL data set is shown in Fig. 10a, and the PE results from the same conditions are overplotted in Fig. 10b. (The PE results were range-averaged to be comparable with the naturally averaged ASTRAL results.)

Figure 10 shows encouraging similarities in TL character between the two models, although such similarities are not to be taken as proof. The difference between each pair of corresponding ASTRAL and PE curves was computed, as were the mean and standard deviation of the absolute value of those differences. The results are plotted against cluster number in Fig. 11. The clusters have been rank-ordered according to increasing level of overall TL of each representative TL curve.

Fifteen of the mean differences (75% of the clusters) are between 2 and 3 dB, and all are less than 6 dB. Some of the poorer results coincide with either high-frequency (1000 Hz: clusters 16 and 18) or shallow source or receiver depth (50 ft: clusters 17–20). The shallow source/receiver depths in this environment typically resulted in significantly higher TL, particularly on the slope (at the far end of the track), as the acoustic energy is stripped away by the lossy interaction with the shallowing bathymetry.

The small mean differences between the sets of ASTRAL and PE curves are very encouraging from the perspective of establishing consistency between the two models. A variety of TL types are included in the clusters, suggesting that this comparison method might be robust enough to serve as a quantitative measure of consistency. Of course, other factors within the simulation of interest may dictate the accuracy required, but

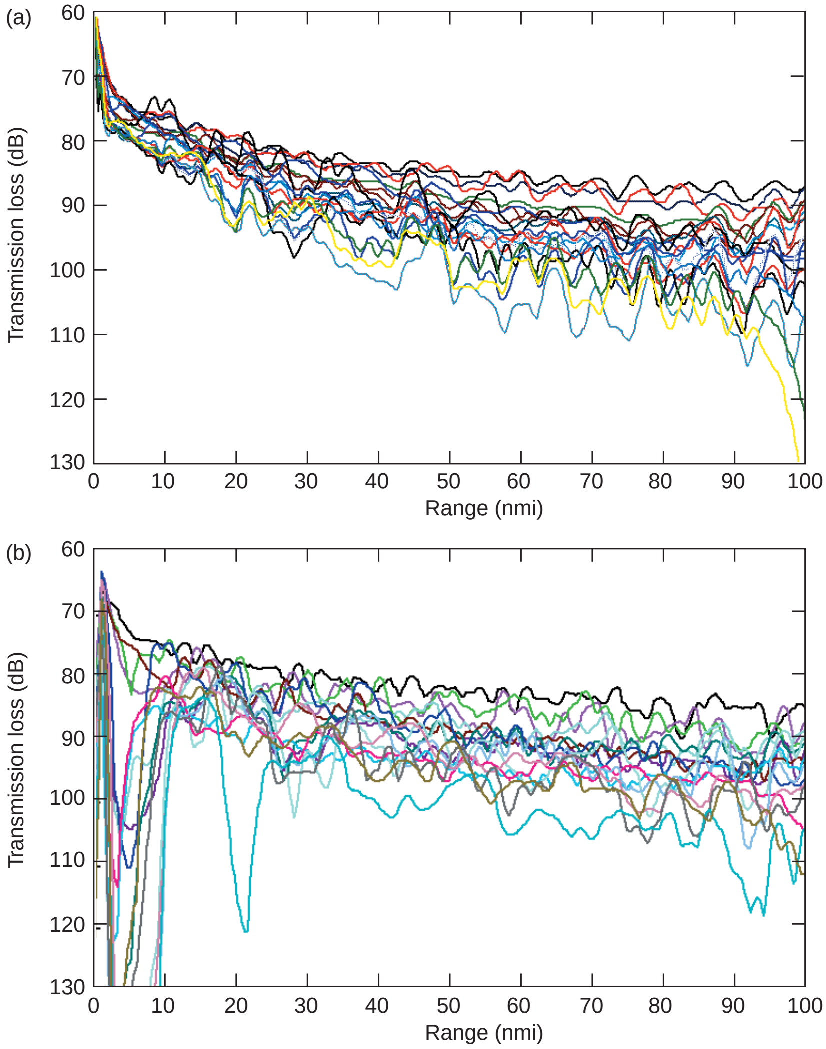

Figure 10. Overplot of (a) ASTRAL representative TL curves for all 20 clusters, and (b) 20 parabolic equation model TL curves generated using the same system parameter values.

this type of analysis would clearly support determination as to whether that accuracy condition could be met. Also, as suggested by the increasing differences for high cluster numbers (higher overall TL) in Fig. 11, inspection of the pairwise ASTRAL–PE TL curve comparisons helps to identify potential problem areas where the models do not agree as well, and provides an opportunity to consider mitigating measures.

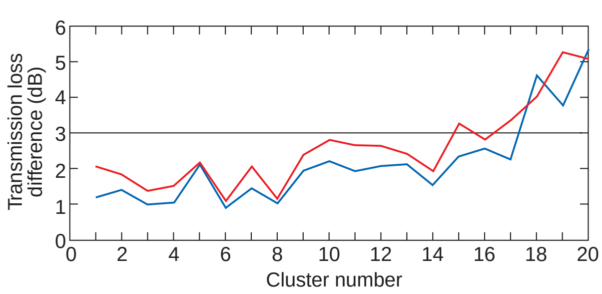

Figure 11. Mean and standard deviation (red and blue curves, respectively) of the absolute value of differences between ASTRAL TL curves and the corresponding parabolic equation model TL curves.

CONCLUSION

The promise of advanced simulation to aid in addressing a wide range of problems from training to acquisition to mission rehearsal is being fulfilled almost daily. But many problems appear as we try to push the envelope of capability beyond the state of the art. Representation of the natural environment has been a major topic in this field since the earliest days of SIMNET. However, progress is needed if we are to move forward in realizing the vision of come-as-you-are, plug-and-play simulations.

The work undertaken here addresses three of the problems facing the environmental community and offers some new directions and results in this regard. The application of cluster analysis to use strategy, consistency analysis, and model-dependent environmental abstractions is promising, and we are continuing our investigations on this topic.

REFERENCES

[1]Newman, F. C., Croucher, A. R., Biondo, A. C., Mandelberg, M. D., Spall, J. C., et al., *FY98 Progress Report: Multiresolution Interaction Validity Project*, STD-R-2774, JHU/APL, Laurel, MD (Oct 1998).

[2]Mandelberg, M. D., "*Observed Upper-Ocean Extremes in Temperature and Salinity for Selected Littoral Areas*," STD-N-1040, JHU/APL, Laurel, MD (May 1995).

[3]Blumberg, A., and Mellor, G., "A Description of a Three-Dimensional Ocean Circulation Model," in *Three-Dimensional Coastal Ocean Circulation Models, Coastal Estuarine Science, Vol. 4*, N. S. Heaps (ed.), American Geophysical Union, Washington, DC, pp. 1–16 (1987).

[4]*Software Test Description for the ASTRAL Model*, Naval Oceanographic Office, NRL OAML-STD-23E, Stennis Space Center, MS (1994).

[5]Croucher, A., *Sea of Japan Description and Environmental Feature List*, STD-L-14814, JHU/APL, Laurel, MD (14 Jul 1997).

[6]Macnaughton-Smith, P., Williams, W. T., Dale, M. B., and Mockett, L. G., "Dissimilarity Analysis: A New Technique of Hierarchical Sub-Division," *Nature* **202**, 1034–1035 (1964).

[7]Spall, J. C., "Resampling-Based Calculation of the Information Matrix in Nonlinear Statistical Models," in *Proc. 4th Joint Conf. on Information Sciences, Vol. 4*, Research Triangle Park, NC, pp. 35–39 (23–28 Oct 1998).

[8]Spall, J. C., "Multivariate Stochastic Approximation Using a Simultaneous Perturbation Gradient Approximation," *IEEE Trans. Autom. Control* **37**, 332–341 (1992).

THE AUTHORS

FRED C. NEWMAN is a physicist and a member of APL's Principal Professional Staff. He received B.A. and M.A. degrees in physics from Williams College in 1964 and 1966, respectively, and his Ph.D. in physics from the University of North Carolina, Chapel Hill, in 1971. Before coming to APL in 1986, Dr. Newman was Principal Investigator for projects related to acoustic and nonacoustic detection of submarines at Science Applications International Corp. He is now Assistant Supervisor of the Ocean Atmospheric and Environmental Sciences Group. He has conducted research into environmental effects in advanced computer simulations, effects of the ocean environment on acoustic and nonacoustic submarine detectability, physical oceanography, and solid-state physics. Dr. Newman is a member of the American Geophysical Union. His e-mail address is fred.newman@jhuapl.edu.

ALBERT C. BIONDO is a Senior Staff physicist in the Acoustics and Electromagnetics Group of APL's Submarine Technology Department. He received a B.A. in physics (cum laude) from the University of Maryland, Baltimore County, in 1976 and an M.S. in physics from the University of Maryland, College Park, in 1982. He joined APL in June 1977. Mr. Biondo's professional efforts have included investigations of underwater acoustic propagation and antisubmarine warfare systems assessment. His more recent investigations have involved increasing the use of higher-fidelity, physics-based models in high-level simulations, either directly or via abstracted responses of the models to varying environmental conditions. Mr. Biondo is a member of the American Physical Society, the Acoustical Society of America, and the honor society Phi Kappa Phi. His e-mail address is albert.biondo@jhuapl.edu.

MICHAEL D. MANDELBERG is a member of the Senior Professional Staff of APL's Submarine Technology Department. He received a B.A. in physics from Brown University in 1980 and an M.S. in applied mathematics from The Johns Hopkins University in 1990. Since joining APL in 1990, he has worked in a variety of areas including physical oceanography, oceanographic data acquisition systems, oceanographic analysis systems, and acoustical oceanography. Mr. Mandelberg is a member of the Marine Technology Society. His e-mail address is mike.mandelberg@jhuapl.edu.

ARTHUR R. CROUCHER II is a member of the Senior Professional Staff of the Ocean Atmospheric and Environmental Sciences Group of APL's Submarine Technology Department. He has been involved in numerous physical oceanography, geophysics, and analysis programming projects. His e-mail address is art.croucher@jhuapl.edu.

JAMES C. SPALL joined APL's Strategic Systems Department in 1983 after receiving an S.M. from M.I.T. and a Ph.D. from the University of Virginia. He was appointed to the Principal Professional Staff in 1991. He also teaches in the JHU Whiting School of Engineering and holds a joint appointment in the Department of Mathematical Sciences. Dr. Spall has published many articles in the areas of statistics and control and holds two U.S. patents. He received the Hart Prize in 1990, and was the Chairman of the 4th Symposium on R&D in 1997. He is an Associate Editor for the *IEEE Transactions on Automatic Control* and a Contributing Editor for the *Current Index to Statistics*. He also was the editor and co-author for the book *Bayesian Analysis of Time Series and Dynamic Models* and is the author of the forthcoming book *Introduction to Stochastic Search and Optimization* (Wiley). Dr. Spall is a senior member of IEEE, a member of the American Statistical Association, and a fellow of the engineering honor society Tau Beta Pi. His e-mail address is james.spall@jhuapl.edu.

CATHY C. MATTHEWS is an electronics engineer at the Naval Air Warfare Center Training Systems Division (NAWCTSD) with 17 years of experience in military and training systems. She received her B.S. degree in acoustical engineering from Florida Atlantic University and an M.S. degree in computer engineering from the University of Central Florida. Ms. Matthews is presently a member of the team responsible for developing ocean acoustics requirements for Navy training systems. She is also a member of the Integrated Product Team for the Generic Acoustic Stimulation System modification for air antisubmarine warfare trainers, the SEDRIS (Synthetic Environment Data Representation Interchange Specification) management team, and the Advanced Simulation Technology Thrust Program for DARPA. Her e-mail address is mathewsc@navair.navy.mil.

J. THOMAS WARFIELD has been a consultant to APL since 1995, primarily in the area of modeling and simulation. He has been a member of the MIV project since its inception, and has also contributed to numerous other APL projects. Before becoming a consultant, Dr. Warfield worked for 25 years in government service, his last two positions being Technology Manager for Ocean and Atmospheric Technology in the Office of Naval Technology and Associate Superintendent of the Acoustics Division at the Naval Research Laboratory. During the early 1960s, he was responsible for pilot performance data collection and analysis on the first NASA-sponsored real-time (7-day) pilots-in-the-loop simulations of the Apollo lunar landing mission. His e-mail address is tom.warfield@jhuapl.edu.