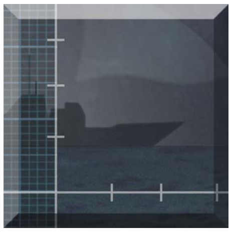

Enhancing Realism in Computer Simulations: Environmental Effects

Fred C. Newman, Albert C. Biondo, Michael D. Mandelberg, Cathy C. Matthews, and J. Ross Rottier

Environmental effects play an essential role in many DoD advanced computer simulations. Without adequately representing the environment and its effects on sensors, weapons, platforms, and people, simulation outcomes are suspect and ensuing conclusions unreliable. Accurate physics-based models of environmental effects, however, are often computationally intensive and require more computer speed and power than can feasibly be made available to the simulation. APL and NAVAIR Orlando have teamed to develop a methodology using cluster analysis that provides a pragmatic and general solution to this dilemma. The methodology, called Model-response Investigation and Visualization, has been successfully applied in the Navy Fleet Battle Experiment "Hotel" and the Navy War College Global 2001 War Game, and is currently being applied to the Joint Warfare System.

INTRODUCTION

The DoD is developing and implementing computer simulations of complex situations involving combinations of people, systems, and the natural environment to support training, analysis, and acquisition. The effects of the natural environment on systems and people are significant factors impacting the outcome of situations represented by many of these simulations. For example, ocean properties affect underwater acoustic sensor performance; properties of the atmosphere affect radar and infrared sensor performance; and terrain characteristics affect vehicle mobility, performance of local wireless communications, etc. Correspondingly, the simulations themselves must adequately and appropriately represent all environmental effects that significantly impact the processes being modeled as a necessary, if not sufficient, condition for the resulting conclusions and recommendations to be accurate and complete. We have chosen the terms "adequate" and "appropriate" because a given simulation may not require that every nuance of a particular environmental effect be represented to the greatest achievable accuracy in order for that effect to fulfill its role in the scheme of modeled processes composing the simulation.

Detailed physics-based models of environmental effects are generally computationally complex. Since simulations often include significant numbers of entities (i.e., sensors, weapons, platforms, people) simultaneously experiencing some effect(s) of the natural environment, and since the changing locations of these entities are continually modulating the details of the respective environmental effects, the corresponding models must be run and rerun rapidly and simultaneously to keep pace with the evolving sequence of events. However, many simulations must run either in "real time" to allow human-in-the-loop interaction (e.g., for training or mission rehearsal) or much faster than real time to permit the rapid repetition of runs and accumulation of data and results for analysis. In most cases, the available computer power is not sufficient to support the rapid and simultaneous run of several to numerous complex physics-based environmental effects models required to keep pace with real-time or faster simulations. Hence the current problem becomes how to achieve sufficiently accurate and appropriate representation of important environmental effects while maintaining the required speed of simulation.

Calculation speed is improving with the advent of faster processors and models employing faster and more efficient algorithms.[1] Significant modeling advances typically occur at intervals of a few years. But even though available computer power continues to increase, the processing power required to generate environmental effects model calculations to keep up with real-time and faster simulations must be at least an order of magnitude greater than the current processor rates (1–2 GHz). In all probability, however, other simulation requirements will continue to increase as computer technology advances and as users and developers recognize the potential for enhancing other aspects of simulation capability. The net result is likely to be that only a fraction of the increasing computer processing power will be available for the real-time calculation of environmental effects.

This article describes a practical and cost-effective solution to the problem. As part of the Defense Advanced Research Projects Agency (DARPA) Advanced Simulation Technology Thrust, JHU/APL and NAVAIR Orlando teamed to develop a methodology that has come to be known, somewhat whimsically, as Model-response Investigation and Visualization (MIV). MIV uses cluster analysis to identify, for a given model, a relatively small set of model calculations that can be used to approximate, to a specified degree of accuracy, every model calculation that would otherwise be required during the course of a given simulation. We also review other approaches to this and related problems, describe the MIV approach in modest detail, review our experience with applications of MIV thus far, and offer a perspective on the potential value of MIV to Navy and Joint training.

RELATED WORK

Over the past several years, much emphasis has been placed on finding ways to accelerate acoustic model predictions for use in tactical decision aids, mission planning and analysis, training, and simulations. Some of the approaches have focused on extensions to the processing hardware (e.g., using multiple fast processors in parallel to share the computational load[2]). Others have focused on algorithm efficiency improvements or on alternative, simpler approximations to the models being used (e.g., ASTRAL[3]). However, three specific efforts have gained considerable interest for their approaches to the problem.

Investigators at the Applied Physics Laboratory of the University of Washington (APL/UW) have concentrated their efforts on developing a neural network–based "approximation" to a complex acoustic model (the Navy's Comprehensive Acoustic Simulation System/Gaussian Ray Bundle [CASS/GRAB]) for use in trainers and real-time predictive systems.[4] Using the computational efficiency and speed that neural networks afford, these investigators are attempting to produce faster and more robust estimates of active sonar performance in various acoustic environments. The neural network "estimators" they are using accept 24 parametric inputs (e.g., sensor depth, acoustic frequency, sound velocity versus depth, bottom type, bottom depth, etc.) for a specific region. The neural network is then "trained" on collected or modeled transmission loss and reverberation data, with weights for the parameters adjusted according to the sensitivity of the output to the particular input. This results in a clever technique to train the neural network without unnecessarily adjusting and compensating for all of the weights associated with the parameters.

Initial results of the APL/UW investigation have been encouraging, and this continues to be an area of active inquiry as part of the Environmentally Adaptive Sonar Technologies (EAST) Program sponsored by the Office of Naval Research. However, the technique is currently limited to very benign acoustic environments (i.e., those that do not change appreciably over a specific geographical region), as their input parameters allow for a single sound velocity versus depth and three bottom depths to characterize a region. Presumably, adapting the APL/UW technique to more range-dependent environments would require more input parameters, resulting in increased data requirements and (possibly) appreciably longer times for the training phases. Further work in this area will be monitored for significant progress. (A similar approach to "simulating" the acoustic environment was attempted here at JHU/APL[5] with similar results. But the in-house technique was clearly intended as a proof of concept and currently lacks the accuracy and robustness to capture the salient characteristics of other than very benign acoustic environments.)

Another approach to providing accurate transmission loss estimates efficiently and rapidly is being pursued by researchers at the Naval Research Laboratory (NRL), Stennis Space Center, Mississippi.[6] Recognizing that sonar tactical decision aids and mission planning tools require acoustic predictions over large geographical areas, these researchers are attempting to provide the needed estimates by specifying a nonuniform grid, over the region of interest, which will optimize the computations necessary to characterize the acoustic conditions within an area. The grid is specified by the expected variability and/or complexity of the acoustic environment, being finer in more complex regions and coarser in more benign regions. The variability of the acoustic conditions is "anticipated" by monitoring the transmission loss computations at nearby grid locations and examining the similarities of these adjacent cells. The degree of similarity determines whether the grid point in question is specified as a required computation point or if its associated transmission loss can be approximated by the loss at nearby grid points.

The nonuniform grid technique allows for the optimum allocation of available computational assets to those regions that most require them. Furthermore, requiring computations in the more highly variable portions of a specific region also ensures a higher degree of accuracy in the resulting transmission loss characterization for that region. NRL investigators have found this to be true in the synthetic and real-world test cases that they have investigated.

Unfortunately, the number of computations required to both generate the nonuniform grid and provide the full transmission loss characterization for the area of interest cannot be known *a priori*. Indeed, in highly variable, complex acoustic environments (such as might be found in most littoral areas), the nonuniform grid might actually "degenerate" into a uniform one to achieve a reasonable accuracy criterion for the region. In such cases, the computational load will be significant as it approaches (or reaches) the "compute each time" condition. NRL is continuing this investigation, hoping to employ genetic algorithm techniques to determine the optimum grid specification and its attendant computational requirements.

Still another approach to the transmission loss estimation problem has been promoted by investigators at the NRL in Washington, DC.[7] This method uses a fast and efficient algorithm (FeyRay) to generate acoustic transmission loss and other quantities in near–real time. FeyRay was developed with the speed, fidelity, and implementation requirements of sonar trainers and simulators in mind. It is a broadband, range-dependent, point-to-point propagation model optimized for computational efficiency. In this NRL approach, FeyRay is an embedded component of the Acoustic Transmission Loss Server (ATLoS), which is used to provide transmission loss estimates to target-sensor pairs as required during a given simulation. (ATLoS was developed as a range-dependent follow-on to the Personal Computer Shallow Water Acoustic Toolkit [PCSWAT], a ray-theory–based propagation model fashioned primarily for mine warfare applications.) Environmental information is retrieved from dynamic and static databases and provided directly to FeyRay to perform its calculations. The results are then "served" to the requesting target-sensor pairs.

In the recent Fleet Battle Experiment-Juliet (FBE-J), ATLoS was able to provide some transmission loss estimates as requested. However, the demand nominally exceeded 50 requests per minute, a rate faster than it could support, although an algorithm was employed to ensure that transmission loss values were provided for each contact-sensor pair as frequently as possible. Nevertheless, NRL expects that the FeyRay developer may be able to improve the computational efficiency of the model even further, if not relegate the task to several processors. In addition, NRL investigators feel that they may be able to streamline the environmental data retrieval and setup for FeyRay, either through direct-access files or specialized data structures.

Each of the above approaches has demonstrated some limitations with respect to the complexity of the environment that can be addressed or the response time needed to obtain the transmission loss values in order to support the required pace of the simulation. The MIV methodology, however, accommodates environments of arbitrary complexity and provides instantaneous approximations to the required transmission loss values. The methodology can also be used to gain insight into the response of a model to the environment of interest (e.g., regions of greater sensitivity to input parameters vice regions of less sensitivity) to enhance model deployment strategies in support of computer simulations and training exercises. This enhanced insight can also serve to minimize the effort required to evaluate differences between (or among) candidate models for simulations and training systems.

MIV METHODOLOGY

The application of the MIV methodology begins with an examination of the particular environment to be used in the simulation. The environment contained within specified time, latitude, and longitude intervals (or otherwise-specified geographic limits) is sectioned into regions that are roughly uniform in those properties that impact the particular environmental effect of interest such as acoustic propagation in the ocean. These regions may be as small or as large as conditions dictate and, in the case of acoustic propagation, are based on ocean bathymetry, bottom material characteristics, and the shape of the sound speed profile between the surface and the bottom of the ocean. For our first applications of the MIV methodology, we developed these regions subjectively, guided by our understanding of the impact of these properties on acoustic propagation. We are now testing an automated technique for creating these regions. Significantly, the effective application of the MIV methodology is not sensitive to the details of the region boundaries, so development of the regions, while requiring careful thought, is not an unduly exacting process.

Following the sectioning of the battlespace environment, a variety of representative contact-sensor pair locations, judged to be typical of the entire range of propagation conditions expected to be encountered during the simulation, are selected within and among the regions. Each contact-sensor pair location is chosen to be separated by a distance marginally greater than the maximum expected detection range for the entire battlespace because the resulting transmission loss calculations include all ranges from the source out to the maximum expected detection range. The selected acoustic propagation model is then run for many combinations of contact depth, receiver depth, and frequency between each of the selected contact-sensor pair locations as described below. The resulting set of model runs constitutes a database of model output from which to select a subset of model runs that will serve as a compact representation of the entire spectrum of propagation behavior to be encountered over the simulation battlespace.

In addition to sectioning the environment into regions, consideration must be given to the ranges (or intervals) of relevant system parameter values required to ensure provision for all possible circumstances likely to be encountered during the simulation. For acoustic sensors, the system parameter values include the range of acoustic frequencies that will be encountered during the simulation as well as the range of sensor depths and contact depths expected as the sensor and contact platforms maneuver about the battlespace. A multiparameter space of model input values is defined by the collection of these intervals, and an appropriate sampling of each parameter over its respective interval is required so that the desired environmental effects model (e.g., acoustic propagation model) can be run for every combination of parameter values along each propagation path specified in the regional map. The sampling intervals need not be constant—it is more economical to sample more densely where it is known that the model is likely to be more sensitive to variation of that parameter and to sample less densely where the model is likely to be less sensitive.

Next, the environmental effects model is run for all combinations of system parameter values for each selected contact-sensor propagation path. The number of sample values for each parameter and the number of propagation paths may cause the total number of model runs to be in the thousands or tens of thousands. Automating the creation of model input files, including extraction of the environmental data along the propagation path, greatly reduces the labor and the required computer time per run. Some environmental effects models run sufficiently fast on modern desktop computers or workstations that all the runs for a single propagation path may be accomplished in a few hours, thereby allowing all of the calculations to be done in a few days. For the Joint Warfare System (JWARS), which is discussed in the section below, we calculate 10,400 model runs for a single contact-sensor pair location in approximately 2 h on a 2.4-GHz dual-processor PC. If an environmental effects model takes a significant fraction of an hour on a desktop computer or workstation, however, access to a high-performance computer would be a great advantage. The entire set of model runs for all contact-sensor pair locations is collected into a database for analysis.

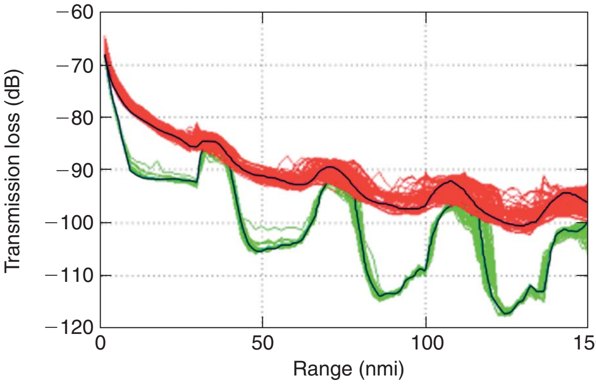

Figure 1. Two example clusters of acoustic transmission loss curves. The representative curve for each cluster is shown in black.

Cluster analysis is applied to the database of model runs to identify a relatively small subset of runs that constitutes a good approximation, or representation, of the entire spectrum of model behavior encountered over the simulation battlespace. It is this subset that becomes the library of environmental effects model calculations in the simulation.

Cluster analysis is an objective technique for identifying and characterizing distinct categories, or clusters, of objects with common properties from a set with a variety, possibly a virtual continuum, of characteristics.[8] One object from each cluster can then be identified as most similar to all other objects in that cluster and therefore as the best single approximation to, or representative of, all other objects in that cluster. Figure 1 is a plot of all transmission loss curves in each of two example clusters. The black line within each cluster is the selected representative curve for that cluster. This is the basis for reducing a very large data set of model runs down to a very manageable subset consisting of those model runs that have been identified as the best representatives in their respective clusters. This subset becomes the library of model runs specifically tailored to the simulation.

To determine quantitatively how the objects differ, a measure of the difference between two given objects must be defined. For acoustic propagation loss models, Biondo et al.[9] define this difference as

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

where

d_{rs} = standardized m-space root-mean-square (RMS) difference between transmission loss curves,
c_{ri} = transmission loss of curve r at the ith range,
c_{si} = transmission loss of curve s at the ith range, and
m = number of range points in each curve.

A small value of d_{rs} shows that curves are similar and a large value indicates dissimilarity. RMS differences are calculated between all curves. The "diameter" of a given cluster can then be defined as the maximum value of the set of RMS differences between all pairs of transmission loss curves in that cluster.

Ideally, the RMS difference is calculated between all pairs of transmission loss curves. However, the number of such curves can be very high, resulting in an unnecessarily large number of RMS difference calculations. Instead, a random sample is selected from the calculated transmission loss curves and the clustering is applied to the random sample. Once the clustering is complete and the representative transmission loss curves have been selected, each of the curves that were not in the random sample is assigned to the cluster for which the RMS difference between the curve to be assigned and the representative transmission loss curve from the cluster is the smallest.

Testing has shown that the accuracy of the final clustering results is similar to clustering without drawing a random sample. One measure of the accuracy is the overall average RMS difference between the representative curves and the other curves in their respective clusters. That is, the mean RMS difference is calculated between each representative curve and all other curves in the corresponding cluster. Then the average of all mean RMS differences is calculated, resulting in a single number representing the accuracy of the representative curves for a particular number of clusters. Figure 2 shows the dependence of the overall average RMS difference on the number of clusters. This type of curve can serve as a useful guide to the appropriate number of clusters for a given application.

Once the clustering has been completed, the transmission loss library and directory are generated. The library consists of the set of representative transmission loss curves and the directory consists of the information required to determine the appropriate representative curve, given a simulation request for transmission loss for a particular source-receiver geometry and frequency. Specifically, the directory is a table of parameter interval combinations together with the representative curve corresponding to each combination. Thus to obtain the required transmission loss value, the intervals containing the frequency of interest, latitude, longitude, and depth of both the contact and the sensor are located in the directory, the corresponding representative transmission loss curve is identified, and the contact-sensor separation distance is used to select the correct value from the transmission loss versus range curve.

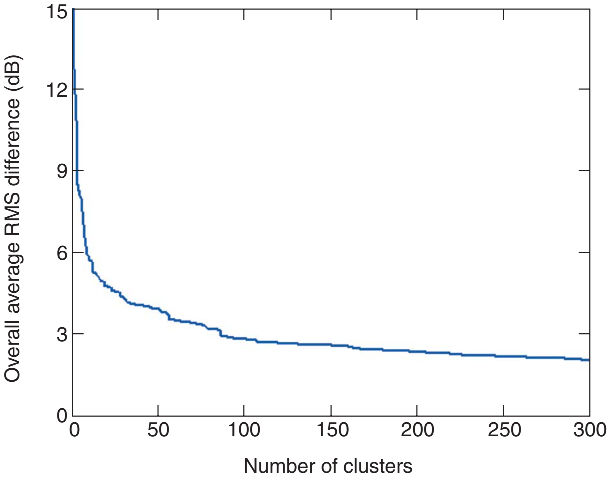

Figure 2. Dependence of representative transmission loss curve accuracy on the number of clusters used to characterize acoustic model response to the sound speed profile environment.

An interesting and valuable empirical property of the cluster analysis results is that the mapping of sample runs to representative curves is insensitive to minor changes in the environment. If the input parameter values (except for environment) for each representative curve are saved and used with the new environment parameter values to calculate a new set of "representative" curves, the resulting curve set will be a very good approximation to the set of representative curves that would result from cluster analysis applied to a set of sample runs for the new environment. This property allows extremely rapid (less than 1 h) generation of a new set of representative curves each time a new environmental data set becomes available and makes development of a new directory unnecessary. For underwater acoustics, the ocean bottom topography and bottom material composition remain constant, and only the sound velocity field varies as ocean circulation and weather affect the upper ocean temperature and salinity fields.

MIV APPLICATIONS

In July 2001, MIV was applied to the anti-submarine warfare (ASW) simulation component of the Joint Synthetic Forces (JSAF) simulation. The JSAF simulation was operated by the Naval Warfare Development Command for the Navy War College–sponsored Global 2001 War Game. The war game "playbox" ocean areas included the Persian Gulf and the Gulf of Oman during a 3-week period in March 2001. The NRL (Stennis) supplied daily Modular Ocean Data Assimilation System (MODAS) 1/8° × 1/8° gridded analyses of the ocean volume temperature and salinity from which we calculated the sound speed profiles on the same grid. We sectioned the ocean sound speed field into feature areas (Fig. 3) as described above, identified appropriate contact-sensor pair locations, and calculated acoustic transmission loss for all combinations of 20 contact depths, 20 sensor depths, and 26 frequencies ranging from 20 Hz to 10 kHz. Next, we used the Navy standard acoustic transmission loss model ASTRAL, the fastest Navy standard model available at that time, to calculate the transmission loss for each of these combinations and for each pair of contact-sensor locations. We then grouped the resulting calculations into 200 clusters and selected a representative transmission loss from each cluster. These 200 transmission loss curves and the corresponding directory were the data that JSAF applied in the sonar equation to determine whether a sensor had detected a contact. Using this approach, JSAF was able to satisfy calls for transmission loss at rates ranging from several to several tens of requests per second.

The effectiveness of the MIV methodology was dramatically demonstrated during the war game when a friendly submarine was detected and sunk by a threat platform. The commanding officer of the friendly submarine challenged the ability (based on 1 of the 200 representative ASTRAL calculations) of the threat platform to detect the submarine at that range in that environment. To adjudicate this challenge, war game personnel performed the exact ASTRAL calculation for the friendly contact depth, threat sensor depth, contact frequency, locations of the friendly submarine and the threat, and the range between them. The resulting exact transmission loss calculation confirmed the MIV-based detection. This work is described in detail in Biondo et al.[9]

We are currently applying the MIV methodology to the ASW simulation in JWARS. JWARS is a "constructive" simulation that includes all aspects of theater-scale engagement and runs as much as a thousand times faster than real time. Before running a constructive simulation, all information specifying the scenario to be run is entered. The simulation is then run to completion with no further human interaction, although certain information may be monitored as the simulation proceeds. Often a slight change is then made to the simulation input information, such as a change in the value of the random number generator seed, and the simulation is run again. This process is repeated many times to accumulate a data set of results that can be analyzed statistically. JWARS will be employed as a force analysis and course-of-action analysis tool.

Before applying the MIV methodology, JWARS used a "sweep-width" approximation that summed up all the disjoint areas where a contact was detectable by a particular sensor. The disjoint areas were then equated to a singular circular area, centered on the sensor, whose radius resulted in the area equivalent of the sum of the disjoint areas. While the sweep-width approximation required only swift table look-up operations to adjudicate possible contact detections, an important drawback of this approximation was that it greatly reduced most detection ranges, which had important and often misleading tactical consequences. The new MIV-based ASW simulation will require table lookups combined with an arithmetic equation (the sonar equation) evaluation to determine whether a sensor can detect a particular contact. Although this approach requires a few more simple computer operations for each adjudication, little or no impact is anticipated on JWARS run time, but the detailed effects of the ocean environment on sensor performance will be much more accurately and completely represented.

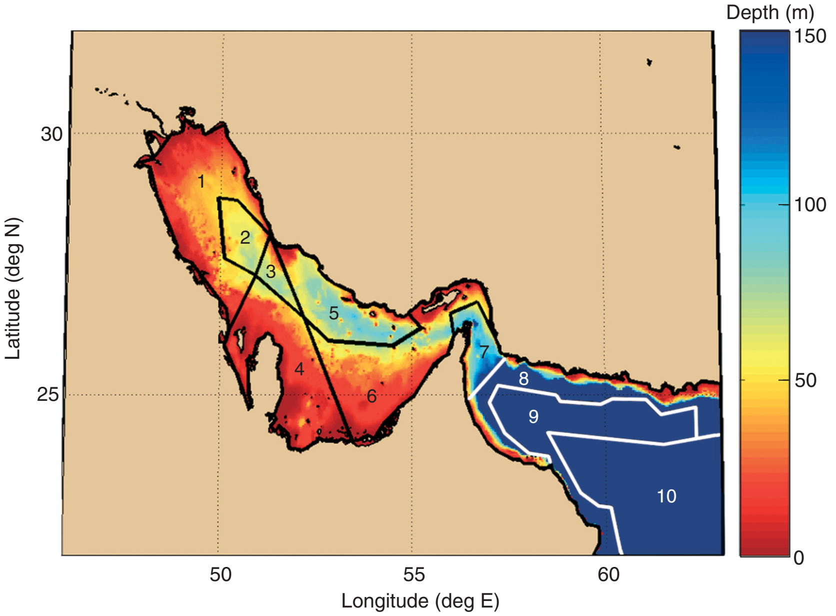

Figure 3. Feature areas showing classification of the Persian Gulf and Gulf of Oman into regions of similar acoustic propagation environments. All depths greater than 150 m are shown in blue at the top of the scale.

Preliminary Application to Radar

In addition to sonar applications, a preliminary application to radar has been examined. For this investigation, the Tropospheric Electromagnetic (EM) Parabolic

Equation Routine (TEMPER)[10] model was used to calculate the pattern propagation factor (F) for a standard atmosphere at frequencies of 1, 3, and 10 GHz. (F is the ratio of the predicted electric field amplitude, including antenna beam pattern and refractive effects, to the amplitude associated with free-space spreading.) TEMPER was chosen because it is the model accepted by the Navy's Aegis Program and is incorporated in the Shipboard Environmental Assessment/Weapon System Performance (SEAWASP) tactical decision aid for predicting radar performance. Constant altitude cuts were made through the two-dimensional TEMPER F^2 field to produce F^2 curves at altitude increments of 1 m. The F^2 curves were then subjected to MIV cluster analysis. The results are documented below.

One can calculate F^2 results according to a variety of parameters including source altitude, receiver altitude, source frequency, pointing (or elevation) angle, antenna pattern, refractivity condition, and region in the problem space (e.g., beyond the horizon or within a surface-based duct). However, the scope of this initial investigation was limited to simply determining whether cluster analysis has potential application to EM propagation. Therefore a single set of F^2 curves at one frequency (1 GHz) and one source height (18 m) was analyzed for a single sector antenna pattern (omnidirectional over a sector of ±90°) and a fixed antenna angle of 0°. The analysis included 1024 receiver (or target) altitudes (from 1 to 1000 m) over a range of 200 km.

This analysis used one-way F^2 results expressed in decibels as $20 \log_{10} F$. The first case considered was propagation over a smooth ocean with a standard atmosphere. The standard atmosphere is the simplest possible case that can occur approximately in the real atmosphere in well-mixed layers.

An F^2 plot for EM frequencies of 1 GHz is shown in Fig. 4. This TEMPER propagation run was made using the standard atmosphere with an antenna pattern of ±90° and a source at 18 m. TEMPER's wide propagator option was used to allow more accurate calculations at the higher angles. A 200-km range was chosen to avoid truncating the main beam of propagation at any altitude within the selected altitude interval (0–1000 m). The lobes are due to multipath interference between the direct path and the reflected ocean path, and occur at a higher spatial rate for the higher frequencies. Note the greater variation in F^2 at altitudes near the ocean surface than at the top of the problem space. Since we have considered only a standard atmosphere, the variation near the surface is due to the source-target geometry and the greater influence of the Earth's horizon at the lower altitudes.

Figure 5 shows an overplot of F^2 at 1 GHz for altitudes of 500 and 550 m. The axis scales have been expanded to focus on the shorter-range oscillations in F^2. Note how the F^2 maxima lobes are similar in shape for the two curves, but the phasing changes with altitude. (This similarity with altered phases is potentially exploitable for clustering purposes but that is left for future analysis.) It is clear from Figs. 4 and 5 that the character of the curve changes with altitude, and in fact the RMS difference between curves increases with altitude separation between curves.

The preliminary results of this investigation are very encouraging. Figure 6 displays the EM accuracy curve. It is the same type of accuracy curve used in the MIV underwater acoustic analysis (Fig. 2) and its shape and magnitude are similar to the underwater acoustics curve.[9] The 20-cluster case shows a mean RMS difference of about 2 dB. Figure 7 shows a plot of the contents of five sequential clusters (sequential in altitude bin) selected from the 20-cluster case. Clearly, the clusters vary uniformly with altitude. Figure 8 shows how the contents of each cluster (1 through 20) are located in altitude for the 20-cluster case. Note the monotonic behavior with altitude. In this simple example, the only input parameter varied was the receiver (or target) altitude. The clusters indicated in Fig. 8 are entirely contiguous in altitude (no cluster has curves from another altitude bin that are within the 2-dB criterion for the 20-cluster case). Also note that the number

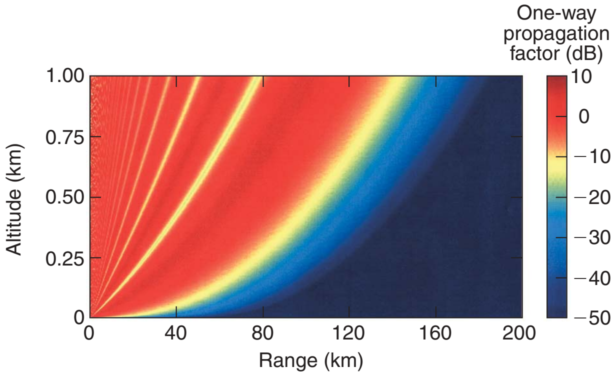

Figure 4. Field plot of the one-way propagation factor F^2 computed using the TEMPER model at 1 GHz and an antenna pattern of ±90°.

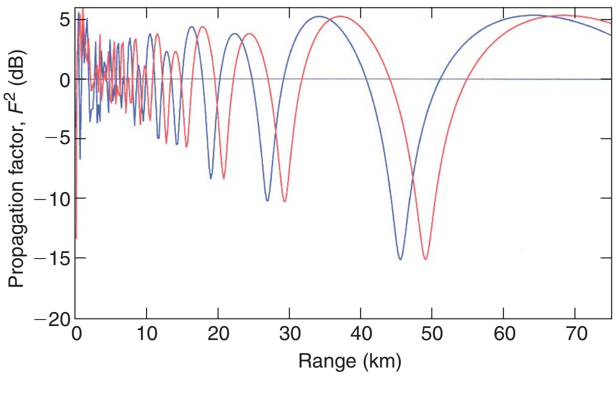

Figure 5. Dependence of F^2 on range for altitudes of 500 m (blue) and 550 m (red).

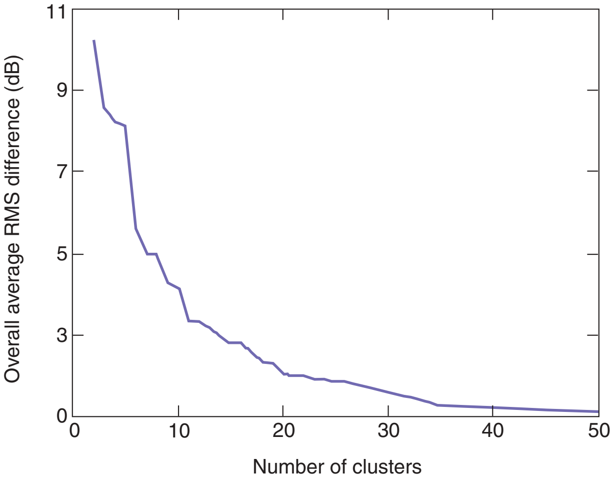

Figure 6. Dependence of representative EM curve accuracy on the number of clusters used to characterize TEMPER response to the given refractivity profile, antenna height, and antenna pattern.

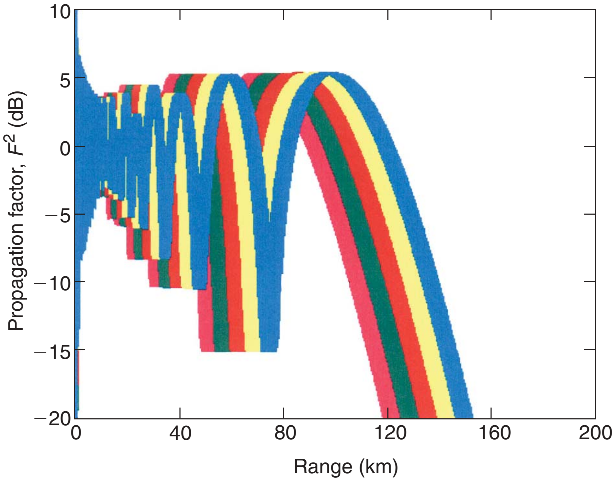

Figure 7. Five clusters of F^2 curves from adjacent altitude bins. Each cluster appears to be a colored band because the curves change gradually in shape with altitude, thereby creating a continuum when all curves in a given cluster are plotted together.

of F^2 curves per cluster increases with altitude and is reduced near the ocean surface where the variability is the greatest.

Although this analysis varied only one input parameter (receiver or target altitude) and used a simplified profile, the results indicate that the cluster analysis techniques developed for underwater acoustics were successful in a limited application. The clusters for a single frequency and single source height mapped to contiguous altitude bins. The cluster accuracy curve indicates behavior similar to the underwater acoustics case, with a reasonable number of clusters to achieve an acceptable level of accuracy. The MIV cluster analysis therefore shows promise for application to EM propagation. The next steps in this investigation would be to

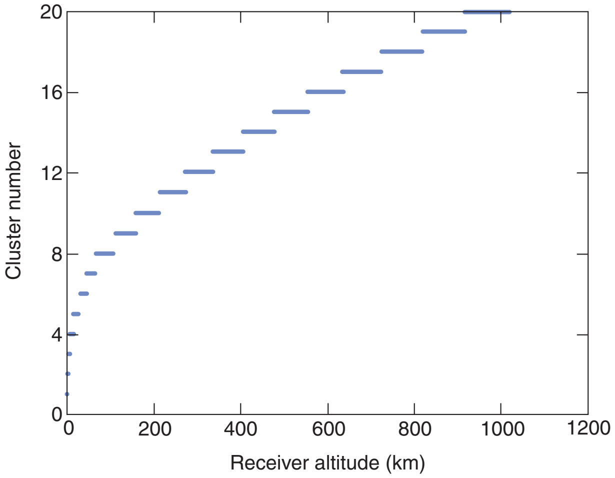

Figure 8. Graph showing how the number of F^2 curves in each cluster increases from cluster 1 through cluster 20. Each F^2 curve is represented by a dot at its corresponding altitude, but the closely overlapping dots result in line segments that increase in length with increasing altitude. The increasing line length is evidence of the increasing number of F^2 curves as altitude and cluster number increase. The apparent slight overlap in altitude of consecutive line segments is due only to the size of the dot used in the plot.

(1) include additional frequencies and source heights for the standard atmospheric profile, (2) add realistic range-independent atmospheric layers that vary with altitude, including evaporative ducts, surface-based ducts, elevated ducts, and subrefractive layers, (3) add range dependence via variation of refractivity profiles and topography with range, and (4) relate cluster variation to real weather events.

Potential Application to Navy and Joint Training Systems

There are many different kinds of training systems, each with unique requirements on simulation design and system capabilities. This section is written largely from the ASW training perspective but is transferable to many other domains. ASW training includes "real-time" virtual "man-in-the-loop," tactical, and strategic training. For real-time virtual man-in-the-loop training systems, the simulation must represent environmental effects in the training scenario to a degree that results in important nuances of system performance being reflected in the sensor displays. For tactical training systems, which might not require the sensor operator level of detail, an overall result that identifies "contact gain/contact loss/bearing to contact/range to contact" information may be adequate. In strategic training—a command level of training above both tactical and man-in-the-loop—the scenario is much larger and involves many more military assets. Strategic training simulations may be required to run faster than real-time military scenarios in order to support course-of-action analysis or to accumulate results for statistical analysis. These training events may not address environmental effects based on a representation of the natural environment simply because of the limited system computational assets. Instead, somewhat arbitrary criteria may be used to determine whether sensor contacts are gained or lost.

The differences in the environmental effects models for man-in-the-loop and tactical levels of simulation can be used to illustrate some differences in system requirements. For example, the virtual man-in-the-loop system may require the acoustic arrival angle, phase, frequency, and signal excess for each propagation path, whereas for the tactical system, signal excess may be adequate since the operator displays may not need to be simulated. In both cases, however, the respective environmental model is required to capture the appropriate effects of the natural environment.

In today's tactical environment, which has changed from deep water to littoral operational conditions, it is even more imperative to teach, train, and practice the impacts of varying and complex environmental conditions. All advantages should be optimized to achieve mission success. Understanding the impact of the environment is crucial to this success. By correctly training environmental impacts, the true cause and effect of the interaction of contacts, environment, and sensors can be fully analyzed and explained, and positive training can occur.

In each of these types of training systems, the physical environmental model calculations are computationally intensive. This condition requires assumptions, shortcuts, and simplifying approaches to be invoked in order to provide an environment for training. The objective is to accurately capture environmental impacts based on the need of the training system and on in-depth analysis. This analysis allows the system designer to ensure that critical events of tactical significance are captured to optimize training while also meeting reasonable computational (system) requirements. Historically, no tools have been available to quantify the massive amount of data analysis and assessments required to determine these design criteria. MIV is an approach that allows this analysis to be manageable *and* quantifiable. This enables the system designer to play "what-if" approaches to achieve an optimum design without sacrificing the fidelity of the environmental representation. MIV also allows for a number of efficiencies that prove advantageous to the training system designer. *In particular, MIV results highlight where the propagation effects are most sensitive to environmental conditions and operational parameters* (e.g., sensor depth, etc.). This enables the training system designer to

- Focus calculations on the "boundaries" of high-sensitivity areas. This results in more efficient use of system resources and realistic representation of the natural environment to support large, complex training scenarios, including highly variable and complex areas. In addition, it allows what-if tactical evaluation before, during, and after training sessions.
- Focus attention on explaining "cause and effect" conditions to optimize the use of military assets.
- Provide better tools (on affordable computers) to explore various tactical options.
- Facilitate assessment of the degree of consistency (or "fair fight") with which different training systems represent important effects of the environment.

While these points include the computational efficiency that supports the detailed realism of simulated environmental effects, they also include capabilities that MIV supports for providing the insight required by both instructor and trainee regarding tactically important effects of the environment. For most military personnel, training time is at a premium. The objective of a training session is not solely to make the training as realistic as possible, but also to take maximum advantage of the time involved for all participants. Therefore an instructor or on-site training lead is responsible for crafting a scenario to get the maximum training capability from each session. This requires insight into both environmental conditions and the effects of the environment on the contacts and sensors. Such insight advances the trainee's knowledge of how the systems should optimally be employed for a given scenario. To best plan a scenario, then, multiple transmission loss conditions should be investigated by the instructor throughout the gaming areas for cause and effect situations.

Training typically includes a "pre-briefing" session, the actual training session, and a debrief or post-evaluation session. The pre-briefing session describes the conditions into which the training crews will be immersed. It provides intelligence similar to that given before actual deployment. In addition to a description of the operational environment, the critical impacts of the physical environment should be presented to the crew. Tactical decision aids are also used to assess the environment, and their products should be assessed as well when developing the tactical strategy.

During the actual training session, the instructor needs to assess current and future conditions so that crew members can later be debriefed on their operational employment skills, tactics, and overall strategy. Also, a scenario may not play out as initially intended by the instructor. Crew decision making may cause the scenario to move into unintended situations. To intervene and best plan the remainder of the scenario, the instructor must understand the new environmental cause and effect situations.

Finally, the debrief or post-evaluation session includes a full disclosure of environmental impacts, allowing the trainee to gain insight into the "invisible" physics of the tactical environment and to assess whether there were better tactics or whether advantages were overlooked. This reinforces an understanding of the effect of the environment on mission success/failure.

For all of the phases of training outlined above, the common principle is the criticality of understanding and demonstrating the training mission, predominately in terms of the capabilities of the tactical resources within a defined environment. The success or failure of a mission may strongly depend on the conditions of the environment and on knowing the areas where the combination of tactical capabilities and the environment change (e.g., a contact of interest may be lost). These usually occur at the boundaries of relatively uniform conditions. This description of the requirements for the virtual training system is, in principle, very similar to the requirements for other types of training systems such as onboard or "organic" training systems.

MIV allows environmental boundaries to be captured and the use of system computational resources to be optimized. It is also an affordable approach to providing insight into the impact of the environment on a military mission. "Rules of thumb" typically used in deep-water operations are no longer applicable to the shallow-water areas where the military may be expected to deploy. The better we understand the environments in these areas, the greater our chances of success.

SUMMARY

Environmental effects are likely to become increasingly recognized as essential to the validity of advanced DoD simulations. This will occur as simulation users and developers realize that simulations at all levels of warfare (tactical, operational, and strategic) require that systems' behaviors include the effects of noise, clutter, clouds, precipitation, inhomogeneities in propagation media, etc., in order to have the most realistic impact on the evolution of the scenario being examined. Although computer speed and power continue to advance steadily, modeling of environmental effects is only one of several competing aspects of advancing simulation technology that will draw heavily on increasing computer capability. For this reason, there will continue to be, for the foreseeable future, a need for methods that enable the computer-efficient representation of complex environmental effects in DoD simulations.

REFERENCES

[1] Weinberg, H., Deavenport, R. L., McCarthy, E. H., and Anderson, C. M., *Comprehensive Acoustic System Simulation (CASS) Reference Guide*, Technical Memorandum 01-016, NUWC-NPT, Norfolk, VA (Mar 2001).

[2] Beatty, W. F., "The Design of a High-Speed Transmission Loss Server and Its Tactical Implications," *U.S. Navy J. Underwater Acoustics* **50**(4), 1233–1247 (Oct 2000).

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

[4] Jung, J-B., *Neural Network Ensonification Emulation: Training and Application*, Ph.D. dissertation, University of Washington, Seattle (2001).

[5] Chin, D. C., and Biondo, A. C., "A Non-Linear Image Inerpolator for a Time-Varying Simulation System," in *Proc. 35th Annual Conf. on Information Sciences and Systems* **1**, pp. 301–317 (21 Mar 2001).

[6] Rike, E. R., and DelBalzo, D. R., "Environmentally Sensitive Sparse Gridding for Efficient Transmission Loss Calculations," in *Proc. IEEE/MTS Oceans '02 Conf.*, Biloxi, MS, pp. 1422–1429 (Oct 2002).

[7] Collins, J., Foreman, T. L., and Speicher, D., "A New Design for a Sonar Environmental Effects Server," *Simulation Interoperability Workshop*, CD-ROM, Paper 02F-SIW-114 (Fall 2002).

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

[9] Biondo, A. C., Mandelberg, M. D., Newman, F. C., and Mathews, C., "Enhanced Representation of Environmental Effects," *Simulation Interoperability Workshop*, CD-ROM, Paper 00S-SIW-121 (Spring 2000).

[10] Newkirk, M. H., Gehman, J. Z., and Dockery, G. D., "Advances in Calculating Electromagnetic Field Propagation Near the Earth's Surface," *Johns Hopkins APL Tech. Dig.* **22**(4), 462–472 (2001).

ACKNOWLEDGMENTS: Funding for this work has been provided by the DARPA Advanced Simulation Technology Thrust program, the Defense Modeling and Simulation Office, and the Navy Modeling and Simulation Management Office. It is currently funded by the Navy Modeling and Simulation Office. The developer of the FeyRay acoustic propagation model is Terry Foreman, Anteon Corp., Arlington, Virginia. MODAS data for the Persian Gulf and the Gulf of Oman were graciously provided by Dan Fox, NRL, Stennis Space Center, Mississippi.

THE AUTHORS

FRED C. NEWMAN is a physicist on the APL Principal Professional Staff. He received a B.A. from Williams College (1964), an M.A. from Williams College (1966), and a Ph.D. from the University of North Carolina, Chapel Hill (1971), all in physics. He was a research associate in physical oceanography at MIT, a National Research Council resident research associate in physical oceanography and acoustic remote sensing for NOAA, and principal investigator for projects related to acoustic and nonacoustic detection of submarines at Science Applications International Corp. before joining APL. He is now assistant supervisor of the Oceanic, Atmospheric, and Environmental Sciences Group of NSTD. Dr. Newman has done research in techniques for including environmental effects in advanced computer simulations, effects of the ocean environment on acoustic and nonacoustic submarine detectability, physical oceanography, and solid-state physics. He is a member of AGU and the Simulation Interoperability Standards Organization. His e-mail address is fred.newman@jhuapl.edu.

ALBERT C. BIONDO is a Senior Staff physicist in the Acoustics and Electromagnetics Group of NSTD. 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. Mr. Biondo has been engaged in numerous technical evaluation and oversight activities involving acoustic sensor development, acoustic modeling and simulation, and tactical decision aid development and 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, IEEE/Simulation Interoperability Standards Organization, and Phi Kappa Phi. His e-mail is al.biondo@jhuapl.edu.

MICHAEL D. MANDELBERG is a member of the Senior Professional Staff of APL's National Security 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 modeling and simulation, physical oceanography, and acoustical oceanography. Mr. Mandelberg is a member of the Simulation Interoperability Standards Organization and the Marine Technology Society. His e-mail address is mike.mandelberg@jhuapl.edu.

CATHY C. MATTHEWS is an electronics engineer at NAVAIR Orlando, with 19 years of experience in military and training systems. She received a B.S. in acoustical engineering from Florida Atlantic University and an M.S. in computer engineering from the University of Central Florida. Ms. Matthews is currently Branch Manager for the acquisition of surface platform training systems. She has previously worked R&D programs with DARPA, ONR, and DMSO. Ms. Matthews has also participated in the development of training systems for all the Navy communities. She has diligently worked toward transition of R&D products to acquisition programs. Her e-mail address is cathy.matthews@navy.mil.

J. ROSS ROTTIER is senior field test engineer in the Theater Systems Development Group at APL, with a B.S. in mathematics (Virginia Tech, 1977) and an M.S. in applied mathematics (JHU, 1988). He worked at Columbia University from 1977 to 1980 as a marine geophysics research assistant and for the Coastal Engineering Research Center from 1980 to 1983 as a coastal oceanographer. He joined APL's Submarine Technology Department in 1983. His work has included environmental effects on underwater acoustics, submarine tactical oceanography, radar tactical decision aid development, and laboratory research on internal waves. His current work includes air-sea interaction, environmental effects on radar propagation, and field testing of sensor systems. Mr. Rottier is a member of the Oceanography Society, AGU, and American Meteorological Society. His e-mail is ross.rottier@jhuapl.edu.