02S-SIW-079

Implications To The Federation Object Model On Assessing Uncertainty In Simulation Results

Stephen Kasputis
Stan Grigsby
Donna W. Blake
VisiTech, Ltd.
535A East Braddock Road
Alexandria, VA 22314-5884
703-535-6640
kasputis@visitech.com
grigsby@visitech.com
blake@visitech.com

Keywords:
Uncertainties, Federation Object Model (FOM), Federation Development Process (FEDEP)

ABSTRACT: *Simulations are used to support analyses from engineering design decisions to tactics developments, to long-term DoD-wide investment strategies. Rarely, if ever, do these analyses consider the uncertainty associated with the simulation results upon which they are based. Estimating that uncertainty is a complex matter. Some of the aspects upon which that uncertainty depends are static, such as the accuracy of environmental data used as inputs. Several of the aspects, however, are dynamic and depend not only on the scenario, but also on the topology of the processing and communications resources used for the simulation execution. Additionally, the modeling of each action, reaction, and interaction in the battlespace has an effect on the uncertainty. Therefore, the uncertainty associated with the state of the simulation is very much a function of time. To be effective, determination of uncertainty needs to become part of federation execution. This paper uses the context of the federation development process to identify the possible sources of uncertainty associated with establishing and executing a simulation. Once these sources are identified, potential implications for the federation object model are assessed.*

1. Defining Uncertainty

In characterizing or measuring uncertainty, the first thing that must be done is to define what is meant by uncertainty. There are several possible interpretations. For example, Bayesian analysis defines uncertain variables in much the same way as statisticians define random variables. Fuzzy logic deals with a different type of uncertainty in the definition of fuzzy sets [1]. For example, there may be a need to define something as "new" or "old." In non-fuzzy sets, new might be defined as less than one year old. In that case, on the366th day, items would transition from new to old in one day. This seems rather counter to logic. Fuzzy logic handles this by making the transition boundary less precise. For example, they could define a linear probability of something being new as being 1.0 at one year and zero at 2 years. At 18 months, something would have a 50% probability of being defined as new. There is thus uncertainty associated with the definition of an item 18 months old.

Perhaps the most straightforward and intuitive definition of uncertainty is one that captures the low order characteristics of the statistics of the parameter of interest. An example of how a statement like this would appear for a specific probability P is: $P = 0.89 \pm 0.04$ with a confidence interval of 90%. This is interpreted as being 90% confident that the true value of the quantity estimated by P it will fall between 0.85 and 0.93.

The definition of uncertainty can also be broadened. It can include, for example, identification of aspects that are the biggest contributors. Further consideration may also be desired to include what would need to be done to reduce the uncertainty associated with or introduced by any or all of the contributing aspects. Thus, it is important to understand what interpretation and aspects of uncertainty are of interest before addressing how to estimate it.

2. Sources of Uncertainty

To identify the sources of uncertainty in simulations, we analyze the process of creating and executing one. Each step of this process is analyzed to identify possible sources of uncertainty associated with it. It should be noted that the process discussed below is not the Federation Development Process (FEDEP). Addressing uncertainty in the context of the FEDEP is discussed in Section 5. Also, the following discussion is presented in the context of physical modeling. The same issues apply to behavioral modeling and are generally compounded by their complexity and our lack of understanding of some behaviors.

2.1 Stating the Problem

The first step is to understand the problem. All of the aspects that influence the situation or process must be identified. Along with these, their dependencies and correlations must also be identified and defined. Any aspect, dependency, or correlation that is not identified will contribute to uncertainty in the final modeling. These types of omission fall into the category that is sometimes known as "unknown unknowns." The only way to identify this uncertainty is through statistical comparison with real world observations.

2.2 Expression of the Problem as Algorithms

The next step is to express the physics as algorithms or rule sets. There are three potential issues in expressing the physical world in algorithms. The simplest of these is not using all the known aspects in the algorithms. That is, the aspects that have little effect are abstracted away. A second issue is the possibility of needing to make simplifying assumptions about the real world to be able to derive algorithms. An example would be replacing the sine of an angle with the value of the angle in radians in deriving an equation for harmonic motion. This example demonstrates two characteristics of simplifications that contribute to uncertainty. The first is that the simplification is an approximation. The second characteristic is that approximations often restrict use of the algorithm to a specific regime. In the example provided, the regime is to small angles where the sine of an angle is approximately equal to the angle.

The third issue in expressing the physics in algorithms comes from an incomplete understanding of the physics of the problem. If the phenomenon to be modeled is not completely understood, the best available algorithms will simply provide an approximation of that phenomenon. Even if a physical understanding is fairly complete at one level of detail or resolution, it may not be at the level which it is needed for the simulation. For example, modeling the stresses on a single fiber in a composite material may be well understood. How to use this model to describe the response of an aircraft wing made of the composite material when subjected to anti-aircraft fire, however, may not be well understood at all.

2.3 Evaluation of the Algorithms

There are also issues associated with the evaluation of algorithms on a computer that contribute to uncertainty. The most of obvious of these is round off error. While the absolute value of round off error may be small for any given calculation, complex scenarios require many hundreds or thousands of calculations that all build upon each other. The growth of the potential error and associated uncertainty can become significant in such cases.

Another issue in evaluating the physics expressed in algorithms is that many mathematical expressions cannot be solved exactly; higher order partial differential equations or an N-body problem for example. In such instances, numerical techniques are typically used to determine an approximate solution or evaluation of the algorithms.

A closely related issue in evaluating the representation of the real world through algorithms is that of representing continuous phenomena discretely. Consider, for example the modeling of energy propagation. The best models calculate propagation loss at discrete intervals between source and receiver. Representing propagation in this manner implies that the medium through which the energy is propagating is piece-wise constant or at best linear between the points at which propagation loss calculations are made. The error that results from the discrete representation of the continuous world provides yet another source of uncertainty in modeling.

2.4 Definition of Modeled Conditions

Closely related to the topic of evaluation of the algorithms is definition of the inputs for them. Inputs come with some degree of uncertainty whether it is defined or not. For the modeling of military operations, these inputs are generally things like the environmental description, and force locations. For any input, there is a limit to the accuracy with which any given input value can be measured. The difference between the measured or observed value and the true value adds error and uncertainty to any use of those inputs.

For input parameters that are continuous in either time or space, resolution or sampling frequency is another consideration. This issue is closely tied to the discrete calculation of continuous phenomena such as propagation as discussed above. Ideally, the sampling frequency for input parameters matches the frequency at which the algorithm is evaluated. This is seldom the case, however, for very real and practical reasons. Typically the values of the inputs are under sampled with respect to the resolution at which the algorithms will be evaluated. Additional input values are therefore generated through an interpolation or modeling scheme. Again, any difference in the generated input values and the true values adds error and uncertainty to the results.

One additional consideration with respect to the sampling of continuous phenomena is that the act of sampling alters the media and therefore the value that is being measured. The measured values are therefore different than they would be in the absence of sampling. This can add to the error of our interpolated values as those interpolations are based upon values that are not truly representative of the medium at the sampled points.

2.5 Variability of Model Performance

The performance of many models and simulations, and, therefore, their associated uncertainty can depend on the conditions under which they are operated. That is, there are some key factors that affect the performance of the simulation. This is especially true for distributed simulations. Examples of such factors are processor or network utilization, or the number of tracks held by an entity. These factors will vary by simulation and method of modeling. But it is generally the case that there are operating conditions that affect a simulations performance and associated uncertainty. It thus becomes critical in accounting for uncertainty in a federation's execution to not only identify these key operating conditions, but also to monitor and report on them.

2.6 Complexity of Federation Execution

As complex as quantifying uncertainty from all sources associated with any specific model may be, tracking uncertainty through a federation's execution can be much more difficult. Such execution requires the interaction of several models. Characterization of the uncertainty associated with any given simulation run requires that not only must the interaction between models be characterized, so also must the sequence of execution. The sequence, including iterations, is highly scenario dependent. It is possible, therefore to have low uncertainty of simulation results under one scenario, and significantly higher under another.

3. Uniqueness and Dynamic Calculation

For simulations that are non-deterministic, distributed or designed to run under different operating conditions and scenarios, last two factors discussed above introduced the dynamic nature to uncertainty in federation execution. They make the uncertainty of federation's results unique to that specific execution run and necessitate the consideration of dynamically tracking and reporting uncertainty during runtime.

To see this need consider the following. For simulations that can be characterized as above, different runs starting from the same initial conditions will likely evolve slightly differently. The nature, extent, results, a number of number of interactions will differ from run to run. Because of this the conditions under which these interactions will be simulated will also differ to some extent. As the purpose of federations is to model these interactions, it can be argued that the sequence or nature of federation execution is unique for each simulation run. This means the path of uncertainty propagation is unique for each simulation run. It follows, therefore, that the uncertainty associated with the simulation run results is also unique and suggests the need to dynamically compute it.

Consider now a federation whose execution is completely deterministic. The results from multiple runs with given identical initial conditions will be identical as will the associated uncertainties. To be of much practical use, the federation should be capable of executing with different initial conditions or scenarios. Each of these different scenarios, however, will evolve differently and thus require a unique determination of uncertainty. While this determination could be done after execution with no loss of accuracy, it would be more efficient to do so during simulation runtime.

4. Methods for determining uncertainty

There are several methods for the determination of uncertainty in federations. Three of these are briefly discussed as examples and to provide some idea of the possibilities and complexities in determining uncertainty. The first is a standard statistical method for combining uncertainties. The second is an adaptation of finite statistical element techniques. The third is the use of Bayesian networks. Other possible techniques include application of fuzzy logic and differential inclusions.

4.1 Standard Statistical Approach

The standard statistical approach of combining uncertainties is effective when multiple factors affect one aspect. Consider an aspect that is a function of several variables: $Y = f(X_1, X_2, \ldots X_n)$. If we have expressed the uncertainty of each variable as a standard uncertainty $u(X_i)$, we can calculate the standard uncertainty of the aspect of interest, $u(Y)$, as

$$u^2(Y) = \sum_{i=1}^{N} \left(\frac{\partial f}{\partial X_i}\right)^2 u^2(X_i) + 2\sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \frac{\partial f}{\partial X_i}\frac{\partial f}{\partial X_j} u(X_i, X_j)$$

(Equation 4.1)

where $u(X_i,X_j)$ is the covariance associated with X_i and X_j. The standard uncertainty $u(X_i)$ is defined as the positive square root of the estimated variance $u^2(X_i)$[2].

This technique can be used to determine uncertainty at any point in the execution of a federation. It can also support sensitivity analyses to identify the major contributors to uncertainty for the problem as a whole or any aspect or subset of the problem. Consider, for example, the problem of ocean acoustic propagation pictured in Figure 4.1.

This technique could be used to define the uncertainty associated with the receive signal as a function of the all the aspect listed. It could also provide the uncertainty associated with any intermediate aspects such as the sound speed profile (SSP), the amount of scattering, or the estimated noise level.

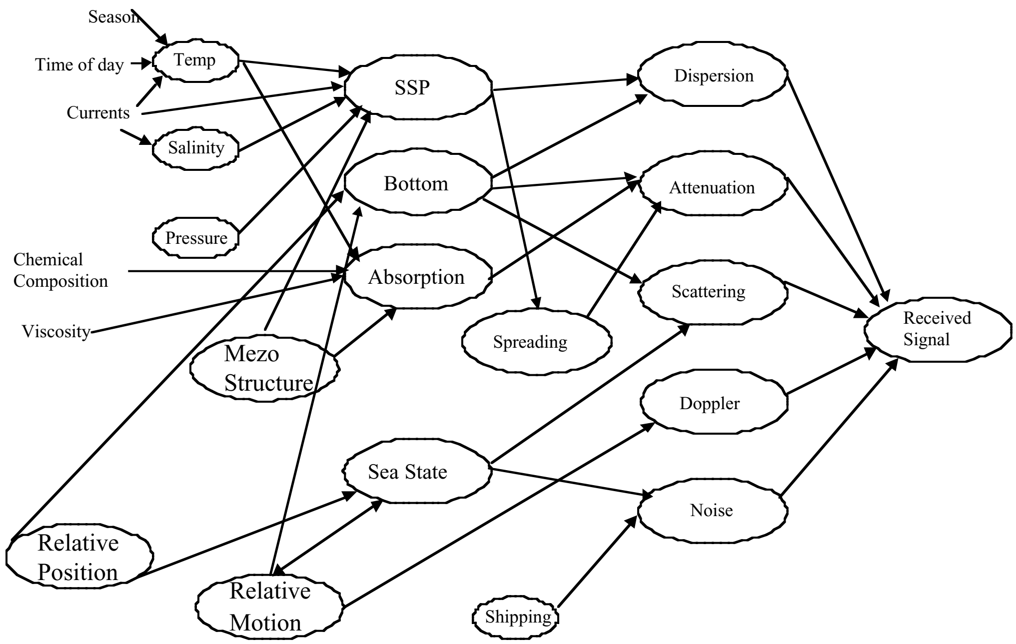

Figure 4.1. Relationship diagram for ocean acoustic propagation

This technique also supports sensitivity analyses that can identify the largest contributors to uncertainty or those to which the uncertainty of interest is most sensitive[1]. The standard uncertainties of any of the aspects can be artificially controlled. This facilitates such studies by allowing for the standard design and execution of experiments in which some aspects are held constant and others varied in a precise and predetermined manner to assess their effect on the downstream uncertainty.

A brief analysis of equation 4.1 reveals something interesting. The value of a combined uncertainty depends on the covariance associated with the factors that affect it. Some covariance values can be dependent upon conditions or scenario. For example, the correlation between temperature and absolute humidity may be different in the desert than in the Caribbean. This could mean the uncertainty determined under one set of conditions may not be very representative of the uncertainty under another set of conditions. The immediate caution is that one should not assume the uncertainty under one scenario should be approximately equal to that under another scenario. This would seem to dictate the calculation of uncertainty should, as a minimum, be done over a number of different scenarios and conditions. Employment of this statistical approach basically requires the repeated evaluation of equation (1). This is an intensive effort. To assess uncertainty under different scenarios requires that the intensive uncertainty calculations be done from start to finish for each scenario since it explicitly uses the covariance values in each step in the propagation of uncertainty.

4.2 Statistical Finite Element Approach

The specific approach identified, that of stochastic adaptive refinement, addresses the representation theory for random processes. It is an extension of the basic ideas of the deterministic finite element method to accommodate

[1] These are not necessarily the same thing. Considering the case represented in Figure 1 as an example, the uncertainty in the received signal may be most sensitive to the uncertainty in the noise. That is, changes in uncertainty in noise result the largest relative changes in received signal. Yet the situation may be such that the uncertainty in noise is low while that in scattering is high and thus the uncertainty in scattering contributes a greater absolute uncertainty to that of the received signal.

random functions. The mathematical problem addressed is that of representing a random process by a denumerable set of random variables, thereby discretizing the process. In a more applied sense, it describes random processes in such a manner that they can be implemented in a finite element formulation of the physical problem.

Following the formalization by Ghanem and Spanos [3], a continuous random process is formally defined as an indexed set of random variables, the index belonging to some continuous uncountable set. The process can be approximated as closely as desired by restricting the index to a set dense in the indexing set. A random process is then represented by its values at a discrete set of points in its domain of definition. The random processes involved are substituted for by random variables that are so chosen as to coincide with some local average of the process over each element.

This approach is roughly analogous to taking the Fourier transform of a function. In mathematical notation, the first steps of this approach can be described as follows: The random process $\omega(x,\theta)$ is expanded in terms of a denumerable set of orthogonal random variables in the form

$$\omega(\theta, x) = \sum_{i=1}^{\infty} \mu_i(\theta) g_i(x) \qquad \text{(Equation 4.2)}$$

where $\{\mu_i(\theta)\}$ is a set of orthogonal random variables and $g_i(x)$ are deterministic functions, which can be related to the covariance kernel of $\omega(x,\theta)$. Since this equation constitutes a representation of the random process in terms of a denumerable set of random variables, it may be regarded as an abstract discretization of the random process. To provide further insight into this approach, this equation can be viewed as a representation of the process $\omega(x,\theta)$ as a curve in the appropriate Hilbert space. The random process $\omega(x,\theta)$ is expressed as a direct sum of orthogonal projections in this Hilbert space whereby the magnitudes of the projections on successive basis vectors are proportional to the corresponding eigenvalues of the covariance function associated with the process. A final note on the potential applicability of this approach is that these concepts can be generalized to allow for the representation of nonlinear functionals.

In the standard statistical approach described previously, changing the scenario may have required a new complete tracing of uncertainty through the data flow architecture, with the multiple evaluation of equation (1), to ensure an accurate estimation of uncertainty. Such is not the case with the approach described here. The expansion of the random functions need only be done once. Once this transformation is made, determination of the uncertainty for any scenario or initial conditions is simply a case of solving a matrix equation.

This method does suffer from two current shortcomings. First is that application of this method to real world problems is still very much in the research stage. There is no streamlined process to mathematically defining and transforming the real world random processes. As a result, any such application may be require considerable overhead time from the few experts who exist. The second possible shortcoming is that this method is not at all intuitive. While it holds probably the greatest potential for efficient determination of uncertainty associated with federation execution, it would be difficult for most decision makers to understand the process by which that uncertainty value was derived. It could, therefore, be difficult for them gain confidence in the results produced by this method.

4.3 Bayesian Networks

A Bayesian network is a graphical model for the probabilistic relationship among a set of variables [4]. The Bayesian network has become a popular representation for encoding uncertain expert knowledge in expert systems. Additionally, methods have been developed and continue to evolve that can learn Bayesian networks for observed data.

Additional properties of Bayesian networks make them potentially useful. The first is that they can easily handle incomplete data sets. For example, consider a case where two input variables are strongly anti-correlated. Other methods for extracting knowledge from data can handle such correlation only if all inputs are measured in every case. If one of the inputs is not observed, these other methods will likely produce inaccurate predictions since they do not encode the correlation between input variables. Bayesian networks offer a natural way to do this.

Another property of Bayesian networks is that they allow one to learn about causal relationships. This is useful when trying to gain an understanding about a problem domain. It also allows one to make predictions in the presence of interventions. That is, they can address questions of cause and effect even when no experiments about the specific effect have been run.

A third property of Bayesian networks is that, in conjunction with Bayesian statistical techniques, they facilitate the combination of the domain knowledge of subject matter experts and data. Bayesian networks have a have a causal semantics that makes the encoding of causal prior knowledge particularly straightforward. Also, these networks encode the strength of causal relationships with probabilities. Therefore, prior knowledge and data can be effectively combined with mature techniques from Bayesian statistics.

Lastly, Bayesian methods, in conjunction with Bayesian networks, offer an efficient approach for avoiding the over fitting of data. That is, there is no need to hold out some of the available data for testing; all available data can be used for defining the network.

There are some cautions with the use of Bayesian Networks. The networks are essentially networks of conditional probability relationships. Each node provides the probability of some condition being true. For problems such as ship self defense with many degrees of freedom, the networks either must be excessively large, or specifics of the problem must be aggregated or abstracted away. The abstractions must be done with great care if the system is still to meet the intended use.

For even simple problems, construction of the network is often rather complicated. Deriving the proper construct so as to ensure casual relationship between parent and children nodes, especially for reactive systems such as in ship self defense, can become a significantly labor intensive problem. Additionally, construction of Bayesian Networks requires defining conditional probability tables for each node. The number of entries needed in these tables for any node is geometrically dependent upon the number of nodes that affect it. For problems involving complex interactions where many factors can affect one node, creation of these tables can become somewhat problematic. This can be mitigated somewhat with careful construction of the network, but this again complicates that effort.

A final caution on Bayesian Networks is that they don't allow for explicit identification of dependencies between children of a common parent. Consider a network in which the parent node is the probability that a radar is radiating. Two children nodes could be the probability of detecting an incoming aircraft, and the probability that the incoming aircraft counter-detects being irradiated. These networks do not provide for identification of any correlation between these two children nodes. Such correlations could provide valuable information for insight into both engineering decisions and development of tactics, techniques, and procedures.

5. Incorporation of uncertainty considerations into the Federation Development Process

If characterization of uncertainty in federation execution is desired, uncertainty considerations must become integral to the Federation Development Process (FEDEP). Activities that must be considered during the FEDEP are outlined below.

5.1 Define Federation Objectives

During this phase, the uncertainty requirements must be identified. These requirements fall into two different classifications. The first identifies the accuracy and confidence limits on the objective data or information provided by the federation. It may be important for an analysis, for example, that a federation provides probability of kill values within $\pm$ 3% under certain conditions.

The second classification of uncertainty related requirements is when some measures or form of uncertainty itself as one of the pieces of objective information desired from the federation. An example of this could be the uncertainty of enemy troop movements associated with different sensor placements in order to evaluate optimum resource employment. Once the requirements are understood, the methods of measuring or determining the needed aspects of uncertainty that best meet these requirements need to be identified. At this point, this identification can be as general as the standard statistical method or Bayesian networks. The general types of data needed to support these methods can then be specified, such as low order statistical moments.

5.2 Develop Federation Conceptual Model

During this step of the FEDEP, all the sources of uncertainty associated with the development and execution of the federation must be identified. Additionally, each source must be classified as significant or insignificant. Insignificant sources are ones that make very little contribution to uncertainty measures of interest and can be extracted away. Care must be taken to ensure that such a source is not significant under any possible conditions of the federation execution.

For significant sources, details of the method of quantifying uncertainty from each must be identified. This detail must be to the level of specific algorithms or exact processes. These details will depend on the place in the conceptual model execution of the factor with which the uncertainty is associated. These factors can be either fundamental inputs or can use other factors as inputs. For example, environmental inputs may be fundamental inputs. As such, the uncertainty associated with them would be that of the measuring technique or accuracy of historical information. Assessing the uncertainty of modeling of the final intercept stage of an incoming air target may depend on the uncertainties of the detection, track generation, and terminal guidance systems as well as environmental factors. In that case, use of Equation 4.1 would be appropriate if the standard statistical method was being employed. To ensure the proper identification and tracing of all the potential dependencies within a federation, a structured approach to conceptual modeling with suitable automated support tools, such as described by Grigsby and Blake [5] is highly recommended.

Since the methods for determining uncertainty have been detailed, the specifics of the data requirements of these methods and also be identified. This data will be either a characterization of fundamental inputs, or the output of another federate. Since such characterization data identifies information that must be passed to and between federates, it should be part of the Federation Object Model (FOM). It is during this phase of the FEDEP, therefore, that the initial FOM requirements to support tracking uncertainty are established.

There are additional aspects that must be identified for the FOM. Because of the need for dynamic determination of uncertainty as discussed earlier, there must be some tracking and communication of the state or value of the operating conditions that govern a simulation's performance. While these conditions are not generally thought of as appropriate for a FOM, there is no other formalism within the High Level Architecture to account for these. Given this and the critical nature of these conditions, it is recommended that operating conditions identified as critical during this phase be included in the FOM.[2]

One additional task can be accomplished during this phase. Once the details for determining uncertainty have been established, the conceptual model of the federation can be used to establish an uncertainty budget, if required. This would be needed if the uncertainty requirement established earlier included accuracy minimums and confidence interval limits for any of the objective data. The identified method for determining uncertainty can be inverted within the context of the federation conceptual model to establish uncertainty limits for each step in federation execution, up to and including fundamental inputs. Because of the probable furcating nature in this backward propagation of uncertainty, subject matter expert opinion may play an important role in the allocation of the uncertainty budget between multiple factors that all affect another factor. It is also probable that many factors are inputs to more than one factor in the forward propagation mode. For such factors, the proper bookkeeping of the conceptual model, such as that offered by current software development tools, is critical in ensuring that the most stringent budget limits are identified. If an uncertainty budget is established, and additional consideration for the FOM could be the dynamic value of uncertainty with respect to the budget. This may be needed if it is desired to have federation execution change in some way, such as executing at a different resolution, depending upon the trend of uncertainty during runtime.

[2] To facilitate this aspect of FOM development, the content of the conceptual model should be expanded to identify all conditions that can or will vary during the lifetime of the federation, including computing hardware and network topology. For example, if a federate is to be capable of running on different computers with different processor speeds, this must be reflected in the conceptual model.

5.3 Design Federation

During the Federation Design phase, the requirements established in the previous phases are used to supplement other requirements to establish either selection criteria for use of existing federates, or design criteria for the development of new federates. These criteria need to ensure that uncertainty is accounted for by the federates, that they determine it in the desired manner, and that they adhere to the FOM for sending and receiving the appropriate uncertainty data. The importance of the criteria addressing uncertainty must be properly weighted with respect to criteria addressing other requirements areas for overall federation performance optimization.

5.4 Develop Federation

During the Federation Development phase, the federation and FOM are finalized. Generally, some of the federate design or selection criteria are eased as compromises are made with other criteria or program constraints on cost and schedule.

6. Summary

Sound engineering and operations analysis practice dictates accounting for the accuracy limits or uncertainty of the data used in the analysis. For data supplied through simulations, this means understanding the uncertainty associated with simulation results. To properly address uncertainty in federation results, it must be considered from the earliest stages of federation development. The requirements addressing it must be on par with those of other performance areas such as execution time, processor, storage, and bandwidth utilization, or operator loading. Development of a complete conceptual model is essential to the identification of all the interactions and dependencies upon which uncertainty propagation during federation execution will depend. Specific data requirements need to be added to the FOM. As a minimum, these include data specific to the method used to track uncertainty, such as lower order moments for the standard statistical method, and information on operating conditions that affect the uncertainty associated with model results. Inclusion of information on operating conditions such as hardware utilization in the FOM is a new concept, but essential if uncertainty is to be properly estimated.

7. References

[1] Bauer, P., Nouak, S., and Winkler, R., "A Brief Course in Fuzzy Logic and Fuzzy Controls," Version 1.2, http://www.flll.uni-ac.at/pdw/fuzzy, December 1996.

[2] Taylor, B. N. and Kuyatt, C. E., "Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results," NIST Technical Note 1297, September 1994.

[3] Ghanem, R. G., and Spanos, P. D., Stochastical Finite Elements: A Spectral Approach, Springer-Verlag, New York, 1991.

[4] Hackerman, D., "A Tutorial on Learning with Bayesian Networks," Microsoft Research Technical Report MSR-TR-95-06, March 1995 (Revised November 1996)

[5] Grigsby, S. and Blake, D. W. "Assessing Uncertainties in the Probability of Raid Annihilation," SISO paper 01F-SIW-077, September 2001.

Author Biographies

STEPHEN KASPUTIS is a Vice President with VisiTech, Ltd., in Alexandria, VA. He had held numerous positions in the Undersea Surveillance Program Office including Technical Director of the Fixed Distributed System. He has been the systems engineer for numerous simulation efforts. He is currently the systems engineer for the modification of JSAF for a prototype Marine Corps training system and directing development of advanced validation techniques. He has a BS in physics from Penn State, an MS in engineering acoustics from The Naval Postgraduate School, and a doctorate in acoustics from The Catholic University of America.

STAN GRIGSBY is a Senior Environmental Systems Engineer with VisiTech, Ltd., in Alexandria, VA. He supports the MARVEDS and P_{RA} programs. He served as the Meteorology Officer on the USS Tarawa and developed and managed environmental effects programs for the Navy High Energy Laser Project. He served as a program manager in the Strategic Defense Initiative Organization. Currently his work is focused on the application of systems engineering practices to the evaluation of environmental effects on Navy systems. He has a BS in physics and an MS in meteorology.

DONNA W. BLAKE is a Senior Scientist with VisiTech, Ltd. in Alexandria, VA, supporting the MARVEDS and P_{RA} programs. She is a former Chief of the Office of the M&S Ocean Executive Agent and has performed research in both ocean and atmosphere modeling at several universities and Navy laboratories. She has served as a program manager for ocean sciences at NASA and for atmospheric sciences at NSF. She has a BA in physics and an MS in astro-geophysics, both from the University of Colorado, and a doctorate in geosciences from the University of Chicago.