SANDIA REPORT
SAND2008-5730
Unclassified Unlimited Release
Printed September 2008

Capturing the Uncertainty in Adversary Attack Simulations

John Darby, Bruce Berry, and Traci Brooks

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.osti.gov/bridge

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online order:	http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2008-5730
Unclassified Unlimited Release
Printed September 2008

Capturing the Uncertainty in Adversary Attack Simulations

John Darby, Bruce Berry, and Traci Brooks
Security Systems Analysis and Security Risk Assessment Departments

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1455

Abstract

The “probability of interruption”, P_I, for a specific attack scenario is typically evaluated using conservative, point-estimate values. The result is significant expenditure of resources – both hardware and response force personnel – to address such scenarios. Also, less attention is paid to balance protection across the entire range of possible scenarios.

This work provides a comprehensive uncertainty technique to evaluate uncertainty, resulting in a more realistic evaluation of P_I, thereby requiring fewer resources to address scenarios and allowing resources to be used across more scenarios. For a given set of adversary resources, two types of uncertainty are associated with P_I for a scenario: (1) *aleatory* (random) uncertainty for detection probabilities and time delays and (2) *epistemic* (state of knowledge) uncertainty for the adversary resources applied during an attack. Adversary resources consist of *attributes* (such as equipment and training) and *knowledge* about the security system; to date, most evaluations have assumed an adversary with very high resources, adding to the conservatism in the evaluation of P_I. The aleatory uncertainty in P_I is addressed by assigning probability distributions to detection probabilities and time delays. A numerical sampling technique is used to evaluate P_I, addressing the repeated variable dependence in the equation for P_I. The epistemic uncertainty for adversary resources is considered using plausibility from the belief/plausibility measure of uncertainty. The uncertainty in adversary resources is epistemic (state of knowledge), not aleatory, so the belief/plausibility measure of uncertainty is appropriate to weight the likelihood of various adversary resources. An evaluation was performed on the fidelity of the data available for use in the technique. Insufficient data are available for a “look up” of a probability distribution for every detection and delay element for every set of adversary resources. However, sufficient data exist such that the data combined with expert judgment can provide probability distributions for a given set of adversary resources. The uncertainty in adversary resources can only be addressed by expert judgment. This report develops a process for applying expert judgment to selected scenarios of concern to address uncertainty, thereby providing a better, less conservative, evaluation for $P_{I.}$ The capabilities needed to support the steps in that process are discussed.

Acknowledgments

Brian Rutherford, Calvin Jaeger, Andrew Walter, and Carla Ulibarri of Sandia National Laboratories provided peer review of this report. Ron Moya of Sandia National Laboratories provided suggestions for the process proposed in this report. Mark Snell of Sandia National Laboratories provided comments and suggestions for improving the report. Diane Ross of Sandia National Laboratories prepared the report for publication.

Contents

Figures

Tables

Executive Summary

For a specific attack scenario, the traditional measure of the effectiveness of the detection and delay elements of a physical security system is the probability that the security system detects the adversary in time for interdiction by the response force. This measure is denoted as P_I: the "probability of interruption." This measure is typically evaluated using conservative, point-estimate values for the detection and time-delay elements of the security system, and the conservatism in the individual elements and in the overall P_I for a scenario is not evaluated. The result is significant expenditure of resources – hardware and response force personnel – to address such scenarios. Also, less attention is paid to balance protection across the entire range of possible scenarios. The subject of this report is the presentation of a comprehensive uncertainty technique to evaluate uncertainty resulting in a more realistic evaluation of P_I, thereby requiring fewer resources to address scenarios and allowing resources to be used across more scenarios.

The Laboratory Directed Research and Development (LDRD) program at Sandia National Laboratories (SNL) sponsored this project on developing a comprehensive uncertainty technique for application to the derivation of P_I. The work was performed between June 30 and September 18, 2008 at SNL in Albuquerque, NM.

This report addresses the following topics:

- Development of the mathematical technique to include uncertainty in evaluation of P_I.
- Discussion of the fidelity of the data available to use in the approach.
- A process for implementing the approach using the available data.

Two types of uncertainty are associated with P_I for a scenario:

- For a given set of adversary resources, there is *aleatory* (random) uncertainty for detection probabilities and time delays.
- *Epistemic* (state of knowledge) uncertainty addresses the adversary resources that will be brought to bear during an attack.

Adversary resources consist of *attributes* (such as equipment and training) and *knowledge* about the security system; to date, most evaluations have assumed an adversary with very high resources, thereby adding to the conservatism in the evaluation of P_I.

In the past, the following main approaches have been used by security analysts in the DOE community to evaluate P_I for a given set of adversary resources:

- the Systematic Analysis of Vulnerability to Intrusion (SAVI)
- the Analytic System and Software for Evaluating Safeguards and Security (ASSESS)
- the Adversary Time-Line Analysis System (ATLAS)

Each of these tools use point estimates for all variables. The Estimate of Adversary Interruption (EASI) technique treats the time variables as random variables, each with a normal distribution; all the probability variables are evaluated as point estimates. Our technique expands these approaches to include uncertainty for all the detection and delay elements, but incorporates the point estimate and the EASI approaches as special (degenerate) cases. None of these earlier approaches addresses uncertainty in adversary resources; our technique addresses this uncertainty.

This mathematical technique treats all variables – time and probability – as uncertain. For a given set of adversary resources, a probability distribution is assigned to each variable to reflect the aleatory (random) uncertainty in the performance of the adversary in defeating the security system elements. For example, in penetrating a reinforced wall barrier, random effects are associated with the location of rebar, the distribution of rubble, and human factors for adversary actions that provide uncertainty for the time to penetrate the wall even if the set of adversary resources is fixed.

This work describes a process to use available data for the technique developed in this report. The steps are:

1. Specify the fuzzy sets for adversary resources, using expert opinion.
2. Select scenarios of concern to be evaluated.
3. Generate probability distributions for detection and delay elements for each set of adversary resources for each scenario.
4. Evaluate P_I for each set of adversary resources.
5. Assign evidence over sets of adversary resources.
6. Evaluate P_I for weighted adversary resources.

By considering both the random uncertainty associated with a specific set of adversary resources, and the state of knowledge uncertainty in the adversary resources, we produce a more realistic, less conservative estimate for P_I.

The application of this approach requires significant effort, including assembling a team of experts, gathering many sources of data, performing additional tests to gather new data in selected areas, and elicitation of expert opinion. Experts in security system elements, human factors, statistics, and threat assessment are required.

This report makes the following recommendations:

- The comprehensive uncertainty approach should be applied only to scenarios where conservative point-estimate values result in prohibitively high costs of hardware and/or response force size to address that scenario. Thus, a screening process can evaluate scenarios and retain only those of concern for subsequent detailed evaluation using this approach. This screening process can be the one typically used where scenarios are identified – using such tools as table-top exercises or path-finding tools such as ATLAS – and evaluated using conservative point estimates.
- A pilot application of this approach to a specific set of scenarios should be performed. This will determine the usefulness, cost, and time required to apply the approach.

Extensions to the technique are suggested. Here, we focused on the uncertainty in P_I. Overall effectiveness, P_E, is the product of P_I and P_N, where P_N is the probability of neutralization of the adversary by the response force given interruption. Uncertainty in neutralization, P_N, could be addressed. Other possible extensions to the technique are also discussed.

Acronyms

ASSESS	Analytic System and Software for Evaluation Safeguards and Security
ATLAS	Adversary Time-Line Analysis System
CCDF	Complementary Cumulative Distribution Function
CDP	Critical Detection Point
DBT	Design Basis Threat
EASI	Estimate of Adversary Interruption
LDRD	Laboratory Directed Research and Development
SAVI	Systematic Analysis of Vulnerability to Intrusion
SNL	Sandia National Laboratories

This page intentionally left blank.

1 Introduction

1.1 Background

The Laboratory Directed Research and Development (LDRD) program at Sandia National Laboratories (SNL) sponsored this work, which is an investigation of a technique to evaluate uncertainty in adversary attacks against physical protection systems, resulting in a more realistic evaluation of the probability of interruption (P_I). Better definition of P_I requires fewer resources to address adversary scenarios and allow resources to be used across more scenarios. The work was performed between June 30 and September 18, 2008 at SNL in Albuquerque, NM.

1.2 Probability of Effectiveness and Probability of Interruption

The effectiveness of a physical security system for a specific adversary attack path is evaluated as the probability that the security system defeats an adversary, given an attack along the path.[1] [Vulnerability Assessment] This measure is denoted as P_E: the "probability of effectiveness." P_E includes the probability that the response force neutralizes the adversary, given that the adversary is detected in time for the response force to interdict the adversary.[2] The probability that the adversary can be detected in time for response is referred to as the "probability of interruption," denoted as P_I. P_I is also referred to as the "probability of timely detection." Here we focus on the evaluation of P_I.

1.3 Defining the Path and its Relationship to Detection

A path consists of a number of layers where the adversary may be detected and/or delayed by elements of the security system. Typical layers include:

- the protected area boundary,
- the outer surfaces of buildings within the protected area,
- areas internal to buildings, and
- targets within the areas.

At each layer, there is a probability of detection and a time delay; in general, time delay may occur both before and after detection. Let "i" denote a layer. Let P_i denote the probability of detection at the i^{th} layer. Let T_{i1} denote the time delay before detection at the i^{th} layer, and let T_{i2} denote the time delay after detection at the layer.[3] Let P_C denote the probability that a detection event is correctly assessed and communicated to the response force. Let T_R denote the time required for the response force to respond to detection and interdict the adversary. For a given set of security system elements, P_i, T_{1i}, and T_{2i} depend on the adversary; P_C and T_R depend on the defender.

The general equation for P_I is:

[1] We are evaluating a specific scenario. Scenarios are selected using such techniques as expert judgment and table-top exercises, or path analysis with tools such as ATLAS. [ATLAS]

[2] P_E is the probability of effectiveness for a scenario: the probability that the security system defeats an adversary given an attack. $P_E = P_I \times P_N$. P_I is the probability of interruption of the adversary attack path: the probability that the security detects the attack in time for the response force to interdict the adversary. P_N is the probability of neutralization: the probability that the response force neutralizes the adversary given interruption of the attack path.

[3] Time delays include time for the adversary to move from layer to layer, time for setup, etc., as well as the time to actually penetrate a delay element.

$$P_I = P_C \sum_{i=1}^{k} P_i \prod_{j=1}^{i-1} (1 - P_j) P(T_{i2} + \sum_{m=i+1}^{k} (T_{m1} + T_{m2}) - T_R > 0) \quad (1)$$

where there are k layers in the path.[4]

At layer i,

$$P_i \prod_{j=1}^{i-1} (1 - P_j) \quad (2)$$

is the probability of detection: more precisely, the probability of detection at layer i and no detection at any layer prior to i.

At layer i,

$$P(T_{i2} + \sum_{m=i+1}^{k} (T_{m1} + T_{m2}) - T_R > 0) \quad (3)$$

is the probability that the time remaining after detection is sufficient for the response force to interdict the adversary; it is the probability that the sum of the time delays from layer i into the final layer k – following detection at layer i – exceeds the response force time. Note that only delay *after* detection is counted for layer i, since at the layer of interest delay *prior* to detection is not counted.

Table 1-1 summarizes the nomenclature for the variables in Equation 1.

Table 1-1. Nomenclature for Variables in Equation 1

Variable	Meaning of Variable
P_I	Probability of interruption for a specific scenario
P_C	Probability that an event that is detected is correctly assessed and communicated to the response force
P_i	Probability that adversary action at layer "i" is detected
T_{i1}	Time delay at layer "i" before detection
T_{i2}	Time delay at layer "i" after detection
T_R	Time for response given detection and assessment
$P(T_{i2} + \sum_{m=i+1}^{k} (T_{m1} + T_{m2}) - T_R > 0)$	Probability that the time delay following detection at layer "i" exceeds the guard response time given "k" total layers

[4] In general T_R and P_C may depend on the layer, and T_R and P_C may be further developed as combinations of other variables. Extension of the approach in this report to include such considerations is straightforward; here we are focusing on uncertainty for the adversary and these considerations address the defender. Section 8 discusses expansion of the approach to include aspects associated with the defender.

2 Prior Evaluations

This section summarizes two simplifications of equation 1 that have been used by various evaluation tools.

2.1 Point-Estimate Evaluation

In the past, Equation 1 has been solved using point estimates for the variables. Such tools as SAVI, ASSESS, and ATLAS use this approach, where paths are found using graph theory algorithms and each path is evaluated using Equation 1 with point estimates for probabilities and times. [SAVI] [ASSESS] [ATLAS]

Using point estimates, if there is a layer in the path – call it c – that satisfies the following:

$$\begin{gathered} P(T_{c2} + \sum_{m=c+1}^{k} (T_{m1} + T_{m2}) - T_R > 0) = 1 \; and \\ P(T_{(c+1)2} + \sum_{m=c+2}^{k} (T_{m1} + T_{m2}) - T_R \leq 0) = 0 \end{gathered} \tag{4}$$

then layer c is called the Critical Detection Point (CDP). There will be a CDP if the time delay remaining after detection at any layer in the path exceeds the response time. Detection after layer c is of no use as there is insufficient time left for response. If there is no CDP, P_I is zero since there is insufficient delay for the response force to interrupt the adversary regardless of where or how well the adversary may be detected. If there is a CDP, then Equation 1 simplifies to:

$$P_I = P_C \sum_{i=1}^{c} P_i \prod_{j=1}^{i-1} (1 - P_j) \tag{5}$$

If there is no CDP, P_I is zero.

Typically, conservative values are used for the point estimates of detection and time delays for adversary actions in Equation 1, so that the point estimate for P_I may in reality be conservatively low. That is, low values for probabilities of detection and delay times result in a possible under-prediction of P_I.[5]

If the uncertainty in the variables is accurately known and represented using probability distributions, conservative point estimate values will correspond to high percentiles of the distributions instead of means and therefore result in a relatively unlikely point estimate for P_I.

[5] A point estimate alone provides no indication of uncertainty. Such a point estimate could be conservative or non-conservative. The point estimates used by SNL are known to be conservative, and it is assumed in the rest of the report that the point estimate values are conservative.

2.2 The EASI Evaluation

In reality, all the variables in Equation 1 have uncertainty. The EASI technique treats all the time variables in Equation 1 as random variables, each with a normal distribution. [EASI] All the probability variables in Equation 1 are evaluated as point estimates. EASI evaluates Equation 1 by analytical convolution of the probability distributions for times.

In this report, convolution means the evaluation of combinations of uncertainty distributions of random variables. Analytical convolution means non-numerical convolution as opposed to a numerical sampling technique. A random variable is a mapping from a sample space to the real number space that provides a probability distribution over real numbers.

Using probability distributions for times, there is no CDP because Equation 3 can be greater than zero for more than one layer due to "overlap" of the probability distributions for the time variables. For the point estimate model discussed in Section 2.1, at each layer Equation 3 is either zero or one. Using probability distributions for times, Equation 3 can be any value within [0,1] for each layer.

Note that Equation 3 always is a point value, not a distribution, even if the time variables in Equation 3 are random variables with probability distributions. Since Equation 3 is a point value, and since EASI uses point values for all the probabilities in Equation 1, P_I as calculated by EASI is a point value.

3 Consideration of Aleatory Uncertainty for All Variables

Aleatory (random) uncertainty is present in the performance of a given security system element for a given set of adversary resources. For example, in penetrating a reinforced wall barrier, random effects are associated with the location of rebar, the distribution of rubble, and human factors for adversary actions that provide uncertainty for the time to penetrate the wall even if the set of adversary resources is fixed.

It is not possible to correctly evaluate Equation 1 by analytically treating all variables as distinct random variables because of the "repeated variable dependence" in the equation. For example, assuming three layers, Equation 1 is of the form:

$$P_I = P_1 * X + P_2(1-P_1) * Y + P_3(1-P_2)(1-P_1) * Z \tag{6}$$

where X, Y, and Z represent other variables in each term. In Equation 6, P_1 is a repeated variable – as is P_2 – and analytical convolution will overestimate uncertainty as it will treat P_1 and $(1–P_1)$ as independent variables, which they are not. This problem can be evaluated using a numerical sampling technique to convolute probability distributions for the variables in Equation 1 instead of using an analytical technique for the convolution. With a sampling technique, the repeated variable dependence is correctly considered by using the sample value for P_1 to evaluate $(1-P_1)$, thereby explicitly accounting for the dependence.

Also, in general the probability and time variables in Equation 1 do not have a normal distribution. A given variable may be better represented using another distribution, such as:

lognormal, uniform, triangular, beta, etc. In fact a normal distribution does not have the correct range for either a probability or a time, since the range for the normal distribution is [–∞,∞] but probability is restricted to [0,1] and time is restricted to [0,∞]. (In application, if the "tails" of the normal distribution are small outside the actual range of the variable, the effect is small.)

Equation 1 can be evaluated by assigning probability distributions to each variable in the equation, and using a sampling technique to evaluate the equation.[6] The sampling technique performs numerical convolution of algebraic combinations of random variables.[7]

3.1 Definition of Probability

We are using the name "probability" in two different ways. For the probability variables in Equation 1, probability is used in the objective (classical or frequency) sense; it is the number of times an event occurs divided by the number of trials in the limit as the number of trials is infinite. The uncertainty in any variable is represented by probability used in the subjective (state of knowledge) sense. Therefore, the probability distribution for a probability variable is a subjective probability of an objective probability. This confusing use of nomenclature is discussed in a 1981 paper. [Kaplan] Kaplan points out that we use the name "probability" to refer to two different concepts, and he recommends using the name "frequency" for objective probability and reserving the name "probability" to mean subjective probability. He uses the nomenclature "probability of frequency" to denote uncertainty in a random variable. The Kaplan nomenclature can confuse engineers because the frequency is not a physical rate, but the dimensionless classical definition of probability. Also, both uses of probability satisfy the Kolmogorov axioms for a probability space, and in this sense probability is the correct term for both concepts. In this report, for probability variables in Equation 1, we mean objective probability; for the uncertainty of any variable in Equation 1 – be it probability or time – we use probability in the subjective sense.

Note that each term in Equation 1 considered as an event in a sample space is mutually exclusive since detection at a given layer considers no detection at all for the prior layers as discussed earlier.

3.2 Crystal Ball Software Tool

For this study, the Crystal Ball software tool was used to evaluate Equation 1. [Crystal Ball] Crystal Ball is on overlay on the Excel spreadsheet software package that allows probability distributions to be assigned to variables, and evaluates algebraic combinations of these random variables expressed as equations in Excel. Different types of standard probability distributions can be assigned to each variable, or a custom probability distribution can be specified for any variable. Convolution is performed by the user specifying either a Monte Carlo or a Latin Hypercube numerical sampling technique.

For example, let A and B be two variables in Excel, and let "C = A + B" be an equation in Excel. Crystal Ball allows the assignment of probability distributions to A and B, and generates a

[6] The selection of the correct probability distribution is based on the nature of the problem, the data available, and experience.

[7] Others have addressed the consideration of aleatory uncertainty in detection and delay variables; for example, the SASRAP and Nextgen projects. [Snell Communication] [Nextgen] [SASRAP]

probability distribution for C using sampling. For our purposes, we will assign probability distributions to all the variables in Equation 1 and generate a probability distribution for P_I.

In contrast to EASI, our result for P_I will be a probability distribution, not a point value. But as previously discussed, Equation 3 is always a single value, not a distribution.

4 Consideration of Epistemic Uncertainty for Adversary Resources

Section 3 discusses solving Equation 1 using the probability measure of uncertainty. Probability addresses random (stochastic) uncertainty, more precisely called aleatory uncertainty. Probability has difficulty dealing with state-of-knowledge uncertainty, more precisely called epistemic uncertainty.[8]

For example, for a fair coin the uncertainty is aleatory: the probability of heads is ½ and the probability of tails is ½.

However, if we do not know the coin is fair, the coin may be biased for heads or may even be two-headed. Our uncertainty for the coin is not random at all – the coin is either fair or not – we just do not know. Our uncertainty is epistemic, or state of knowledge. In the limit where we have no information about the coin, the case of total ignorance, all we can state is that the probability of either heads or tails is somewhere in [0,1]. Belief/plausibility is a measure of uncertainty that is a generalization of the probability measure of uncertainty that can capture epistemic uncertainty. [Theory of Evidence] Belief and plausibility form lower and upper bounds on probability, respectively. For total ignorance the belief/plausibility for heads – or tails – is 0/1.

4.1 Defining Variables

A variable may be difficult to describe numerically; a purely linguistic description is more appropriate for variables that have an unknown numeric scale. For example, for the variable "Health" the linguistic bins "Poor", "Moderate", and "Excellent" have more meaning than an arbitrary numeric scale, since the scale is unknown. Does "Health" have the numerical range [0,3] or $[1,10^6]$ or something else? The problem of unknown scale is made worse when different variables are combined, leading to a result that is dependent on the arbitrary numeric scale that is used. For such situations, it is more appropriate to reason on the words themselves instead of forcing the use of an arbitrary numeric scale. Linguistic bins for a variable are fuzzy sets, and combinations of variables with fuzzy sets can be accomplished using approximate reasoning.

We have applied the belief/plausibility measure of uncertainty to variables described using purely linguistic fuzzy sets. [Terrorist Risk] A computer tool called LinguisticBelief© has been written to evaluate uncertainty using approximate reasoning for combinations of variables represented as purely linguistic fuzzy sets using the belief/plausibility measure of uncertainty. [LinguisticBelief]

[8] The SASRAP project addressed epistemic uncertainty, but not with belief/plausibility. [Snell Communication] [SASRAP]

4.2 Adversary Resources and Creating Fuzzy Sets

Our uncertainty in adversary resources is epistemic, not aleatory, in that the level of adversary resources is not random but fixed and known to the adversary that decides to attack, but unknown to us, the defender. Resources include both attributes and knowledge. Attributes include equipment, weapons, number of adversaries, level of training, etc. Knowledge is the information about the target and the security system known to the adversary.

P_I is dependent on adversary resources. Most prior evaluations assume that the adversary has extensive resources and evaluate P_I accordingly. We denote this adversary by the linguistic fuzzy set "omniscient" or all-knowing.[9] Such an adversary may have significant knowledge supplied by an "insider" as well as significant attributes.

We will consider two other fuzzy sets for the adversary: "Expected" and "Poor". Thus our linguistic bins for adversary resources are the fuzzy sets {Omniscient, Expected, Poor}. These fuzzy sets form our sample space for the adversary. In application, each set of adversary resources needs to be defined; that is, the fuzzy sets over adversary resources need to be defined.

Using expert opinion for our epistemic uncertainty for adversary resources we can assign evidence to families of fuzzy sets. Figure 4-1 is an example of such an assignment.

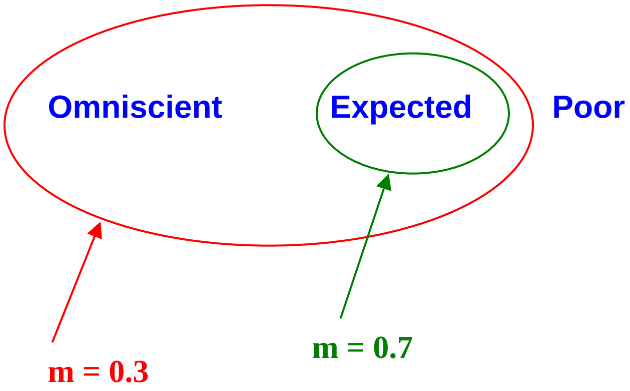

Figure 4-1. Example Evidence

Subsets of the sample space with evidence are called focal elements. Evidence is denoted by "m." In Figure 4-1 we have two focal elements: {Omniscient, Expected} with evidence 0.3 and {Expected} with evidence 0.7. For any subset A, the belief and plausibility of A can be evaluated from the focal elements as follows:

[9] We do not assume that the adversary is omnipotent, or all powerful, as an all-powerful adversary has unlimited resources and P_I would always be zero for such an adversary.

$$Bel(A) = \sum_{B|B \subseteq A} m(B)$$

$$Pl(A) = \sum_{B|A \cap B \neq 0} m(B) \qquad (7)$$

where B is a focal element. If all the focal elements are singletons – that is, each subset B with evidence has only one element – both belief and plausibility are the same, the probability.

For the example in Figure 4-1, the belief/plausibility for each fuzzy set is:

- "Omniscient" has Belief 0 and Plausibility 0.3
- "Expected" has Belief 0.7 and Plausibility 1.0
- "Poor" has Belief 0 and Plausibility 0.

4.3 Using PoolEvidence Software to Compile Results

A set of experts may not agree on the evidence over the adversary fuzzy sets; the PoolEvidence© software can pool the results from different experts to produce a set of pooled focal elements. [Qualitative Uncertainty] PoolEvidence is a utility for LinguisticBelief.[10] For example, assume four different experts assign evidence to our set of adversaries. Figure 4-2 is the model for this situation in PoolEvidence. Figure 4-3 shows the evidence assigned by each expert and the overall pooled evidence.

[10] The development of variables, fuzzy sets, combinations of variables, approximate reasoning rules, and the assignment of evidence is an art that requires training of experts and the use of formal expert elicitation techniques. LinguisticBelief is a tool that captures and processes the information so produced. Here, we are reasoning on one variable, Adversary Resources, and the belief/plausibility for that single variable can be easily calculated using Equation 7. To evaluate combinations of variables, LinguisticBelief can be used.

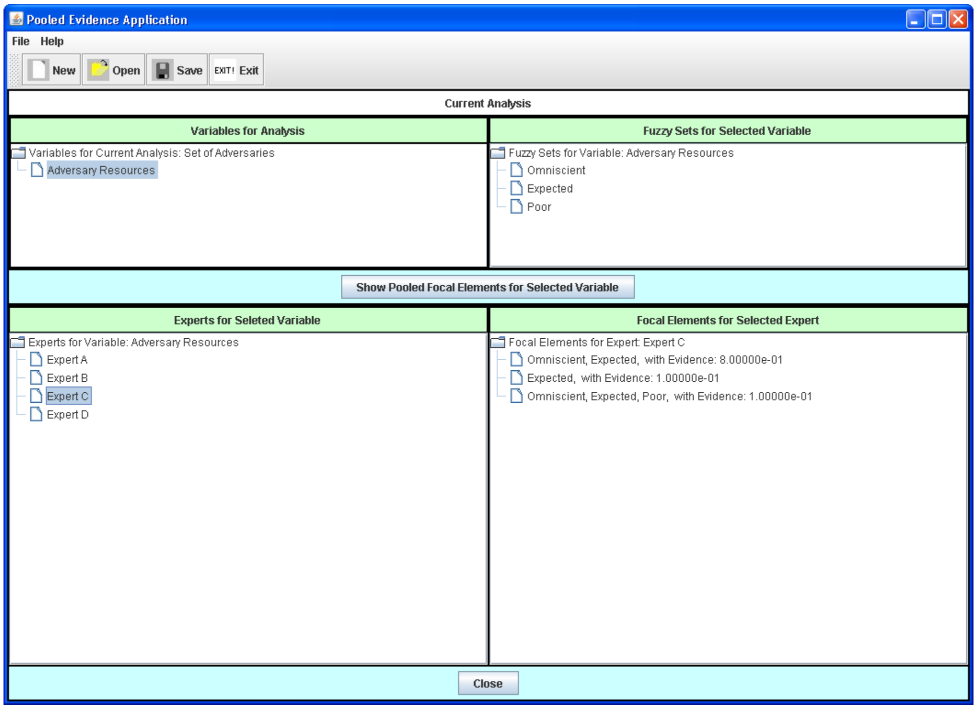
Pooled Evidence Application
File
Help
New
Open
Save
EXIT! Exit
Current Analysis
Variables for Analysis
Fuzzy Sets for Selected Variable
Show Pooled Focal Elements for Selected Variable
Experts for Seleted Variable
Focal Elements for Selected Expert
Close

Variables for Current Analysis: Set of Adversaries
- Adversary Resources

Fuzzy Sets for Variable: Adversary Resources
- Omniscient
- Expected
- Poor

Experts for Variable: Adversary Resources
- Expert A
- Expert B
- Expert C
- Expert D

Focal Elements for Expert: Expert C
- Omniscient, Expected, with Evidence: 8.00000e-01
- Expected, with Evidence: 1.00000e-01
- Omniscient, Expected, Poor, with Evidence: 1.00000e-01

Figure 4-2. Example of Evidence from Numerous Experts in PoolEvidence

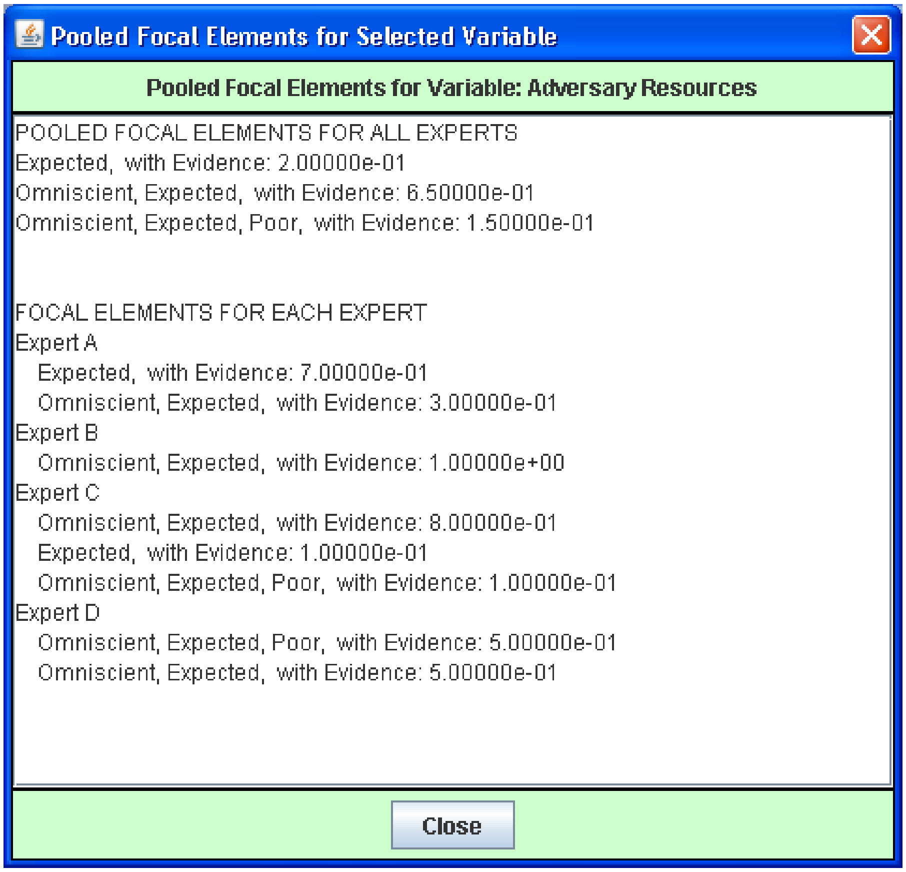
Pooled Focal Elements for Selected Variable
Pooled Focal Elements for Variable: Adversary Resources
Close

```
POOLED FOCAL ELEMENTS FOR ALL EXPERTS
Expected,  with Evidence: 2.00000e-01
Omniscient, Expected,  with Evidence: 6.50000e-01
Omniscient, Expected, Poor,  with Evidence: 1.50000e-01


FOCAL ELEMENTS FOR EACH EXPERT
Expert A
   Expected,  with Evidence: 7.00000e-01
   Omniscient, Expected,  with Evidence: 3.00000e-01
Expert B
   Omniscient, Expected,  with Evidence: 1.00000e+00
Expert C
   Omniscient, Expected,  with Evidence: 8.00000e-01
   Expected,  with Evidence: 1.00000e-01
   Omniscient, Expected, Poor,  with Evidence: 1.00000e-01
Expert D
   Omniscient, Expected, Poor,  with Evidence: 5.00000e-01
   Omniscient, Expected,  with Evidence: 5.00000e-01
```

Figure 4-3. Example of Pooled Evidence

4.4 Probability Distributions for the Variables

The probability distributions for the variables in Equation 1 are dependent on the adversary resources, and thus the final probability distribution for P_I is dependent on the adversary resources. The detection probabilities and time delays for the Omniscient adversary are worse – lower probability of detection and shorter time delay – than are the detection probabilities and time delays for the Expected adversary. Similarly, the detection probabilities and time delays for the Expected adversary are worse than for the Poor adversary. We denote the conditionality of P_I on adversary resources as $P_I|$ adversary, and we have three such conditional probabilities: $P_I|$ Omniscient, $P_I|$ Expected, and $P_I|$ Poor.

We evaluate Equation 1 for each of these adversaries. We – the Defender – use plausibility for the weighting, since we wish to conservatively evaluate the scenarios of concern, and plausibility is an upper bound for P_I.[11]

The overall P_I is evaluated by weighting each conditional P_I by plausibility. Specifically:

$$P_{I\ weighted} = \frac{Plaus(Omniscient) \bullet P_I \mid Omniscient + Plaus(Expected) \bullet P_I \mid Expected + Plaus(Poor) \bullet P_I \mid Poor}{Plaus(Omniscient) + Plaus(Expected) + Plaus(Poor)} \tag{8}$$

where "Plaus" denotes plausibility. For example, P_I|Expected is the probability distribution for P_I evaluated using Equation 1 for the Expected adversary. Plaus(Expected) is the plausibility that the adversary is the Expected adversary evaluated from the assignment of evidence over the set of adversaries as discussed previously.

In summary, the technique is to assign probability distributions for all the variables in equation 1 for each of three sets of adversary resources, evaluate Equation 1 for each of these three cases, then weight the three cases using plausibility calculated using Equation 8.

This technique has been implemented in Crystal Ball, and example results are provided in Section 5. Section 6 discusses the difficulties in generating both (a) the probability distributions from available data, and (b) the evidence using expert opinion to calculate the plausibility for each set of adversary resources.

5 Example Results

The prior evaluations discussed in Section 2 are special cases of the general evaluation technique discussed in Sections 3 and 4.

For only one adversary – say Omniscient – and using point estimates for detection probabilities, the solution of Equation 1 is the same as results obtained using EASI, where each time variable is assigned a normal distribution. This was verified by solving Equation 1 in the Crystal Ball

[11] The Adversary may weight by belief to select a scenario with a high lower bound for P_I. One of the references discusses this in more detail. [Terrorist Risk]

framework for the example problem #1 in the EASI report; the Crystal Ball solution was the same as the EASI solution. [EASI]

The following sections provide an example of the approach. First, P_I is evaluated using conservative point estimate values. Then, P_I is evaluated considering uncertainty using the approach previously described. This illustrates how conservative the point estimate for P_I may be.

5.1 Example with Conservative Point Estimates

Consider an example with seven layers. Figure 5-1 conceptually shows the detection probabilities and time delays for the seven layers along the attack path for this example using the nomenclature of Table 1-1. Dummy data are used in this example.

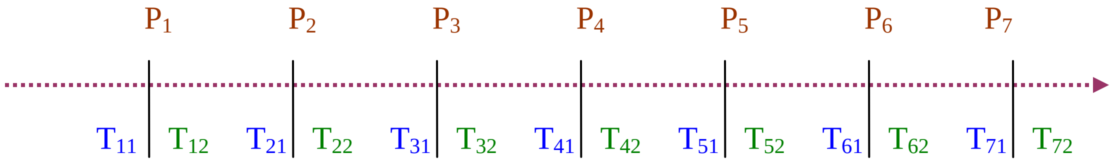

Figure 5-1. Example Attack Path

For this example, T_R is 3.00 minutes, and Pc is 0.9. For the Omniscient Adversary, assume the conservative point estimate values in Table 5-1.

Table 5-1. Point Estimate Values for the Omniscient Adversary

Layer	Time Delay Before Detection at Layer, Minutes	Probability of Detection at Layer	Time Delay After Detection at Layer, Minutes
Layer 1	T_{11} is 0.15	P_1 is 0.10	T_{12} is 0.00
Layer 2	T_{21} is 0.10	P_2 is 0.40	T_{22} is 0.04
Layer 3	T_{31} is 0.30	P_3 is a 0.75	T_{32} is 0.10
Layer 4	T_{41} is 0.40	P_4 is 0.08	T_{42} is 0.00
Layer 5	T_{51} is 1.00	P_5 is 0.30	T_{52} is 0.30
Layer 6	T_{61} is 0.15	P_6 is 0.00	T_{62} is 0.00
Layer 7	T_{71} is 0.70	P_7 is 0.00	T_{72} is 0.00

Using these point values in Equation 4, the CDP is at layer 1, and using Equation 5, P_I is only 0.09 for the Omniscient Adversary. We do not know the conservatism in this point estimate; but, if conservative estimates are used for each detection probability and delay time, the conservatism in the overall P_I can be quite large, as indicated in Section 5.2.[12]

[12] For example, if the conservative point estimate for each of two time delays has a 10% chance of being too high, the sum of the times has less than a 10% chance of being too high.

5.2 Example Considering Uncertainty

The example was evaluated using probability distributions for the security system elements instead of the conservative point estimate values used in the previous section. Crystal Ball was used for the evaluation; Latin Hypercube sampling was selected with 10,000 trials.

All of the three adversaries discussed in Section 4 were considered. For each adversary, the response time T_R and the probability of communication P_C have probability distributions as follows:

- T_R is lognormal with a mean of three minutes and a standard deviation of one minute.
- P_C is uniform with a minimum of 0.87 and a maximum of 1.0.

For the Omniscient Adversary, assume the probability distributions in Table 5-2.

Table 5-2. Probability Distributions for the Omniscient Adversary

Layer	Time Delay Before Detection at Layer, Minutes	Probability of Detection at Layer	Time Delay after Detection at Layer, minutes
Layer 1	T_{11} is triangular with minimum 0.10, likeliest 0.25, and maximum 0.40	P_1 is uniform with minimum 0.05 and maximum 0.20	T_{12} is 0.00
Layer 2	T_{21} is triangular with minimum 0.10, likeliest 0.17, and maximum 0.30	P_2 is triangular with minimum 0.36, likeliest 0.40, and maximum 0.44	T_{22} is lognormal with mean 0.03 and standard deviation 0.03
Layer 3	T_{31} is triangular with minimum 0.20, likeliest 0.50, and maximum 0.80	P_3 is a beta distribution with minimum 0.72, maximum 0.88, alpha 2, and beta 3	T_{32} is triangular with minimum 0.10, likeliest 0.25, and maximum 0.60
Layer 4	T_{41} is lognormal with mean 0.50 and standard deviation 0.20	P_4 is uniform with minimum 0.00 and maximum 0.20	T_{42} is 0.00
Layer 5	T_{51} is lognormal with mean 1.5 and standard deviation 0.70	P_5 is triangular with minimum 0.25, likeliest 0.50, and maximum 0.75	T_{52} is uniform with minimum 0.25 and maximum 0.75
Layer 6	T_{61} is uniform with minimum 0.08 and maximum	P_6 is 0.00	T_{62} is 0.00
Layer 7	T_{71} is triangular with minimum 0.5, likeliest 0.75, and maximum 0.85	P_7 is 0.00	T_{72} is 0.00

For the Expected Adversary, assume the probability distributions in Table 5-3.

Table 5-3. Probability Distributions for the Expected Adversary

Layer	Time Delay Before Detection at Layer, Minutes	Probability of Detection at Layer	Time Delay After Detection at Layer, Minutes
Layer 1	T_{11} is triangular with minimum 0.30, likeliest 0.42, and maximum 0.60	P_1 is uniform with minimum 0.10 and maximum 0.30	T_{12} is 0.00
Layer 2	T_{21} is triangular with minimum 0.20, likeliest 0.33, and maximum 0.56	P_2 is triangular with minimum 0.50, likeliest 0.70, and maximum 0.80	T_{22} is lognormal with mean 0.17 and standard deviation 0.04
Layer 3	T_{31} is triangular with minimum 0.40, likeliest 0.75, and maximum 1.20	P_3 is a beta distribution with minimum 0.81, maximum 1.00, alpha 2, and beta 3	T_{32} is triangular with minimum 0.12, likeliest 0.33, and maximum 0.70
Layer 4	T_{41} is lognormal with mean 1.00 and standard deviation 0.55	P_4 is uniform with minimum 0.00 and maximum 0.20	T_{42} is 0.00
Layer 5	T_{51} is lognormal with mean 2.17 and standard deviation 1.40	P_5 is triangular with minimum 0.30, likeliest 0.80, and maximum 0.90	T_{52} is uniform with minimum 0.25 and maximum 0.75
Layer 6	T_{61} is uniform with minimum 0.16 and maximum 0.66	P_6 is 0.00	T_{62} is 0.00
Layer 7	T_{71} is triangular with minimum 1.00, likeliest 1.50, and maximum 2.30	P_7 is 0.00	T_{72} is 0.00

For the Poor Adversary, assume the probability distributions in Table 5-4.

Table 5-4. Probability Distributions for the Poor Adversary

Layer	Time Delay Before Detection at Layer, Minutes	Probability of Detection at Layer	Time Delay After Detection at Layer, Minutes
Layer 1	T_{11} is triangular with minimum 0.70, likeliest 0.83, and maximum 1.20	P_1 is uniform with minimum 0.45 and maximum 0.55	T_{12} is 0.00
Layer 2	T_{21} is triangular with minimum 0.30, likeliest 0.67, and maximum 0.92	P_2 is triangular with minimum 0.65, likeliest 0.80, and maximum 0.94	T_{22} is lognormal with mean 0.33 and standard deviation 0.10
Layer 3	T_{31} is triangular with minimum 0.93, likeliest 1.50, and maximum 2.1	P_3 is a beta distribution with minimum 0.92, maximum 1.00, alpha 5, and beta 2	T_{32} is triangular with minimum 0.36, likeliest 0.50, and maximum 1.23
Layer 4	T_{41} is lognormal with mean 2.00 and standard deviation 0.87	P_4 is uniform with minimum 0.45 and maximum 0.55	T_{42} is 0.00
Layer 5	T_{51} is lognormal with mean 5.12 and standard deviation 1.23	P_5 is triangular with minimum 0.86, likeliest 0.90, and maximum 1.00	T_{52} is uniform with minimum 0.60 and maximum 0.74
Layer 6	T_{61} is uniform with minimum 0.50 and maximum 0.84	P_6 is 0.00	T_{62} is 0.00
Layer 7	T_{71} is triangular with minimum 3.00, likeliest 3.33, and maximum 6.39	P_7 is 0.00	T_{72} is 0.00

The plausibility for each adversary is assumed to be that as discussed for Figure 4-1, specifically:

- Omniscient has Plausibility 0.3
- Expected has Plausibility 1.0
- Poor has Plausibility 0.

P_I is a random variable over [0,1] calculated by convoluting the probability distributions for its constituent variables. The Complementary Cumulative Distribution Function (CCDF) for P_I is a graph of the probability that P_I exceeds any value in [0,1].[13]

The results for the example follow. Figure 5-2 is the CCDF for P_I for the Omniscient Adversary.

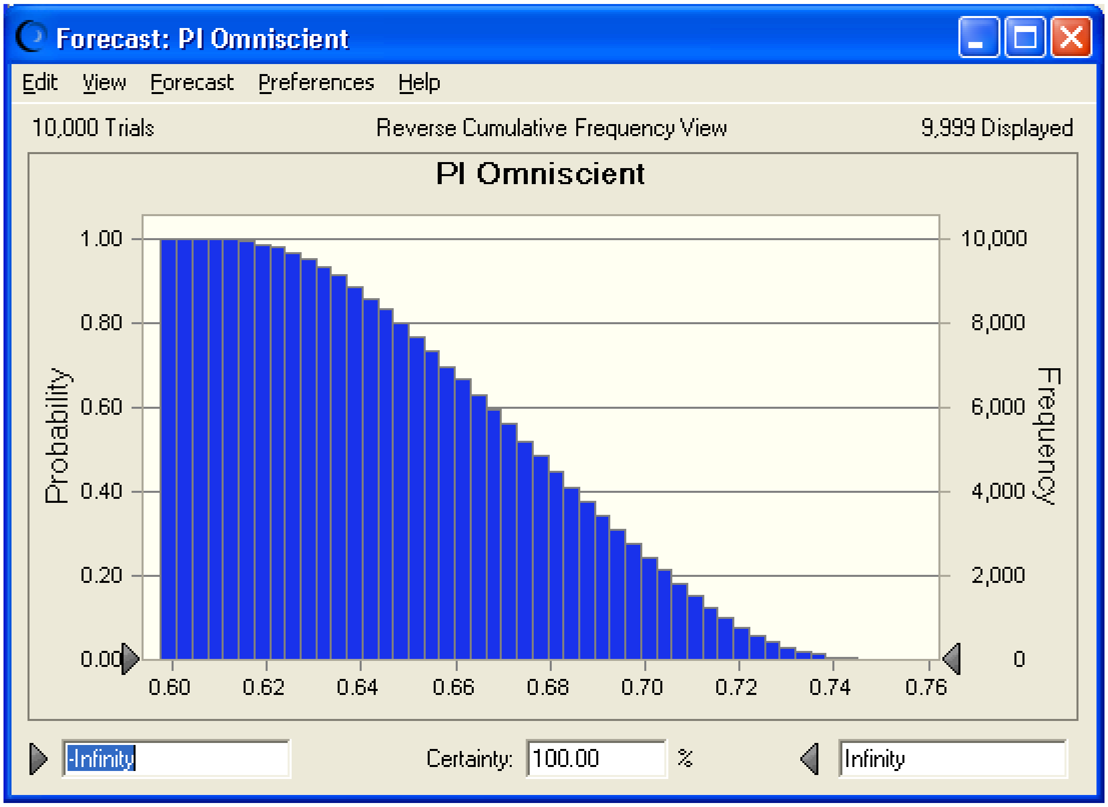

10,000 Trials Reverse Cumulative Frequency View 9,999 Displayed

-Infinity Certainty: 100.00 % Infinity

Figure 5-2. P_I CCDF for the Omniscient Adversary

Based on the statistics calculated by Crystal Ball, P_I for the Omniscient Adversary has a mean of 0.68. However the use of a mean as the point estimate is misleading in that the uncertainty is not captured; different estimates for P_I can have the same mean but have very different uncertainty. A better point estimate is the probability that P_I exceeds some value, such as 0.85. The CCDF can be used to provide this value. The probability that P_I for the Omniscient Adversary exceeds 0.85 is essentially 0.

[13] For a random variable X, let x denote a specific value of X. The CCDF over x is the probability that X is greater than x.

The use of probability distributions instead of conservative point estimates provides a less conservative estimate of P_I. For example, using the conservative point estimates from Section 5.1, the conservative point estimate for P_I for the Omniscient Adversary is 0.09. Using distributions for the detection and delay variables, the mean value of P_I is 0.68, considerably higher than the 0.09 point estimate. Also, Figure 5-2 indicates that it is essentially certain that P_I will exceed 0.60. So, the point estimate of Section 5.1 is so conservative that it is not useful.

Figure 5-3 is the CCDF for P_I for the Expected Adversary.

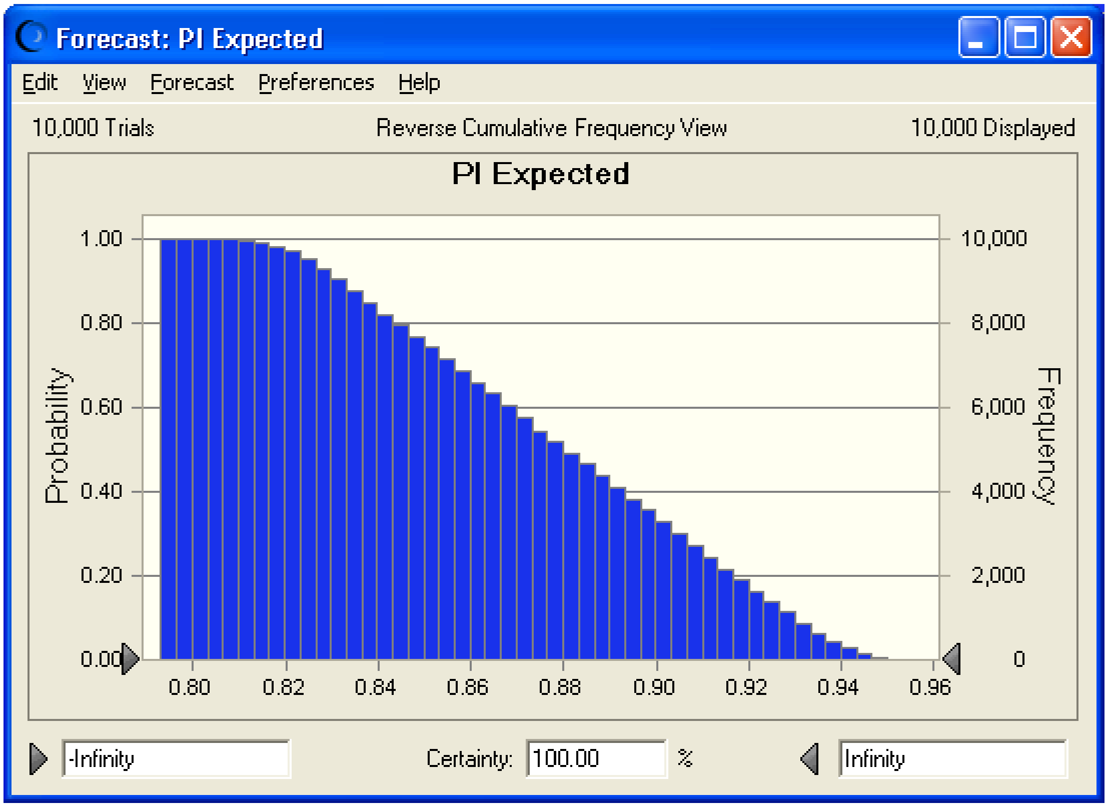

10,000 Trials Reverse Cumulative Frequency View 10,000 Displayed

-Infinity Certainty: 100.00 % Infinity

Figure 5-3. P_I CCDF for the Expected Adversary

Based on the statistics calculated by Crystal Ball, P_I for the Expected Adversary has a mean of 0.88. However, the use of a mean as the point estimate is misleading in that the uncertainty is not captured; different estimates for P_I can have the same mean but have very different uncertainty. A better point estimate is the probability that P_I exceeds some value, such as 0.85. The CCDF can be used to provide this value. The probability that P_I for the Expected Adversary exceeds 0.85 is about 0.70.

Figure 5-4 is the CCDF for P_I for the Poor Adversary.

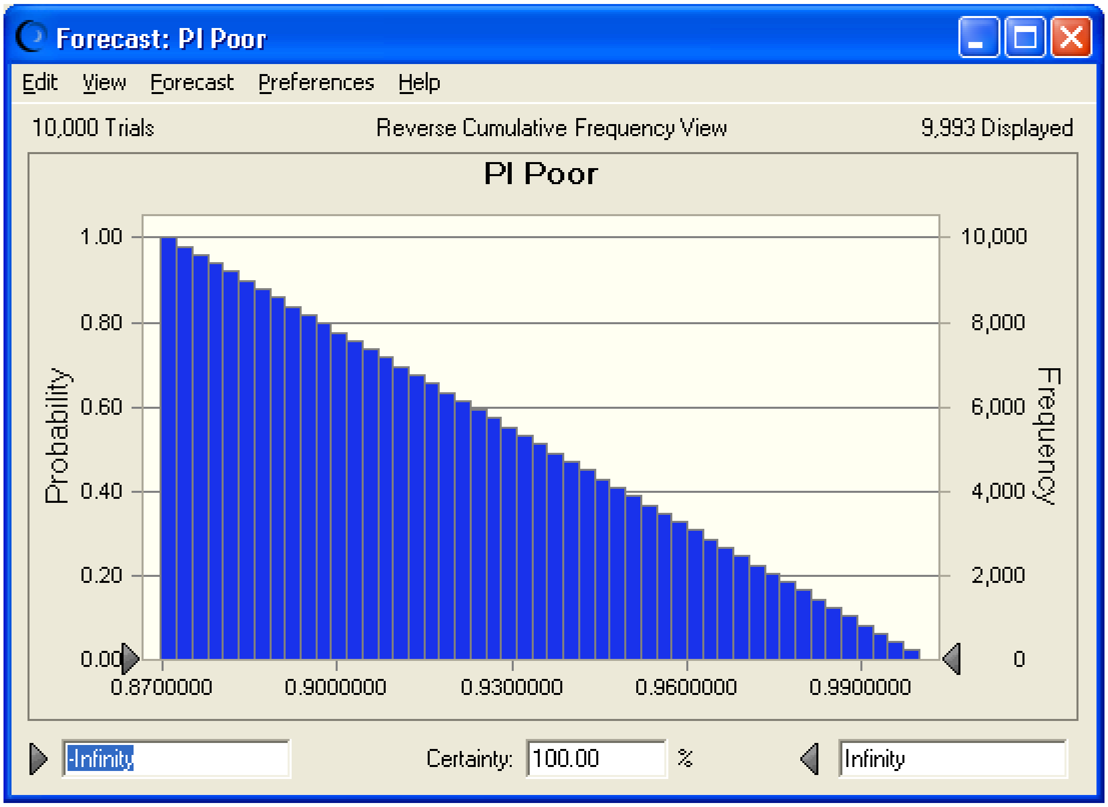

10,000 Trials

9,993 Displayed

-Infinity

Certainty: 100.00 %

Infinity

Figure 5-4. P_I CCDF for the Poor Adversary

Based on the statistics calculated by Crystal Ball, P_I for the Poor Adversary has a mean of 0.93. However, the use of a mean as the point estimate is misleading in that the uncertainty is not captured; different estimates for P_I can have the same mean but have very different uncertainty. A better point estimate is the probability that P_I exceeds some value, such as 0.85. The CCDF can be used to provide this value. The probability that P_I for the Poor Adversary exceeds 0.85 is essentially 1.0.

The weighted result $P_{I\ weighted}$ presented as a CCDF is provided in Figure 5-5.

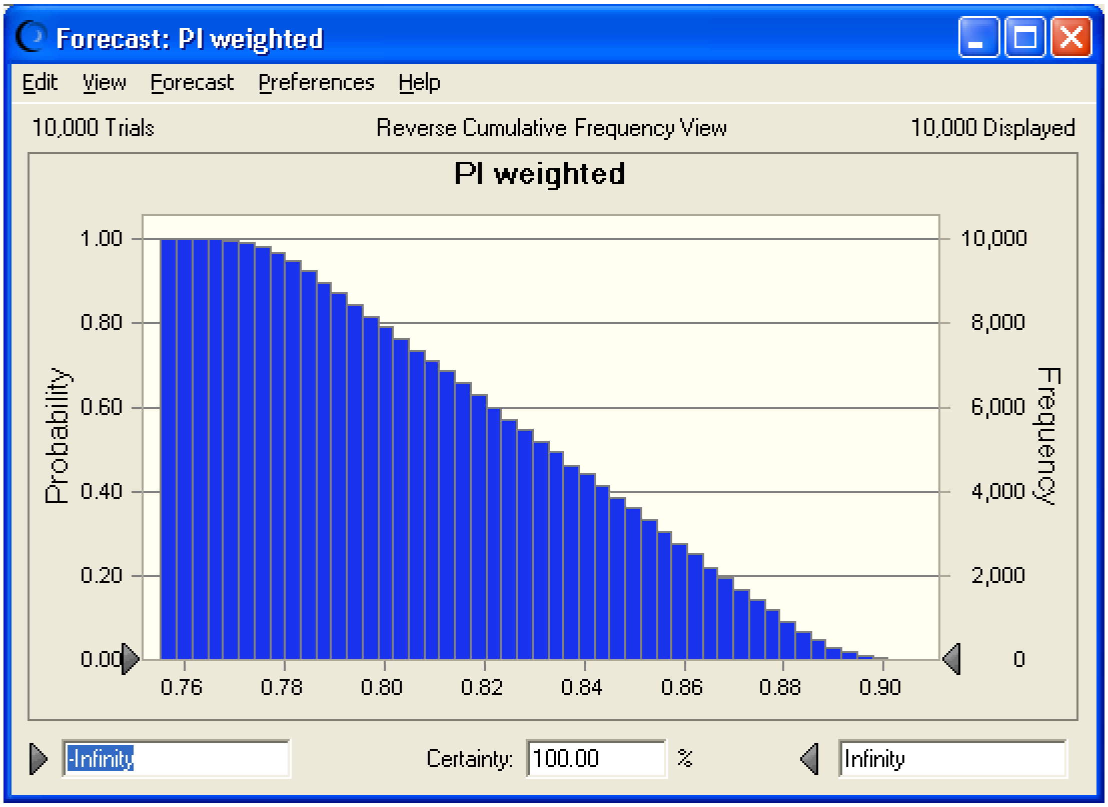

Figure 5-5. P_I CCDF Weighted for All Adversaries

Based on the statistics calculated by Crystal Ball, P_I for the Weighted Adversary has a mean of 0.83. Note that $P_{I\ weighted}$ is considerably higher than P_I|Omniscient, due to the consideration of P_I|Expected. (In this example, P_I| Poor has zero plausibility and has no effect on $P_{I\ weighted}$.) Specifically, $P_{I\ weighted}$ has a mean of 0.83 while P_I|Omniscient has a mean of 0.68.

However, the use of a mean as the point estimate is misleading in that the uncertainty is not captured; different estimates for P_I can have the same mean but have very different uncertainty. A better point estimate is the probability that P_I exceeds some value, such as 0.85. The CCDF can be used to provide this value. The probability that P_I for the Weighted Adversary exceeds 0.85 is about 0.38. Note that the probability that P_I exceeds 0.85 for the Weighted Adversary is 0.38, while the probability that P_I exceeds 0.85 for the Omniscient Adversary is essentially 0.

This example used dummy data for illustrative purposes, but it illustrates how the technique works in application.[14]

This example evaluation supports the following conclusions. Evaluations that use conservative point estimates for detection and delay values to evaluate P_I produce a point estimate for P_I that

[14] Of course, if very conservative probability distributions are used, the answer considering uncertainty will be too conservative. So, the probability distributions have to reflect the actual uncertainty to the extent possible without being overly conservative.

has a low (perhaps extremely low) probability of being that low. The use of probability distributions for a given adversary indicates the uncertainty in P_I.
Although the consideration of aleatory uncertainty for the Omniscient Adversary produces a less conservative estimate of P_I than that produced using conservative point estimates, evaluations that only consider the Omniscient Adversary produce a conservative estimate for P_I. The consideration of the plausibility of different adversaries indicates the uncertainty in P_I based on the resources that may be brought to bear by the adversary, and provides a more realistic, less conservative estimate for P_I.

6 Performance Data Sources and Application to the Technique

A great deal of data has been collected over the years for the performance of detection and delay elements of a security system. The references provided in this report summarize some of these data sources. Due to classification constraints, this report will not discuss actual data in any detail.

The data sources were reviewed to determine their fidelity for implementing the technique developed in this report. There are more data on the uncertainty in delay times than on the uncertainty in detection probabilities. For both time delays and detection probabilities it is not possible to have a comprehensive "look up" database with probability distributions for every element for every set of adversary resources. It is also concluded that cost precludes performing enough tests to generate a comprehensive database that includes every element, due to the extremely large number of elements and the variation in adversary resources used to defeat the elements.

However, in many cases the data are sufficient so that expert judgment can generate approximate, conservative probability distributions for some of the security elements for a given set of adversary resources. Where the data are not sufficient, conservative point estimates can be used, or more specific testing can be performed to allow generation of the probability distributions.

The data sources also include the adversary resources used to defeat the security element, thereby providing information for the consideration of different sets of adversary resources. Section 4 discussed using the fuzzy sets Omniscient, Expected, and Poor for adversary resources.

7 Process Steps

This section proposes a process to use available data for the technique developed in this report. The steps are:

1. Specify the fuzzy sets for adversary resources, using expert opinion.
2. Select scenarios of concern to be evaluated.
3. Generate probability distributions for detection and delay elements for each set of adversary resources for each scenario.
4. Evaluate P_I for each set of adversary resources.
5. Assign evidence over sets of adversary resources.
6. Evaluate P_I for weighted adversary resources.

7.1 Overview of Process

This section describes each of the steps to implement the comprehensive uncertainty approach.

Step 1: Use expert judgment to define the fuzzy sets for adversary resources, specifically the types of attributes and knowledge (including insider supplied information). This requires consideration of the Design Basis Threat (DBT), knowledge of the fidelity of the available data, and expertise in assessing threats.

Step 2: Select scenarios of concern to be evaluated. Each scenario is defined to the level where the detection and delay elements encountered by an adversary for that scenario can be identified. The specific detection and delay elements are specified sufficiently such that an evaluation of detection probabilities and time delays can be performed.

Step 3: For each scenario, generate probability distributions for each detection and delay element for each set of adversary resources. The probability distributions are generated based on expert judgment using the data available. The following areas of expertise are needed:

- Expertise in time delays for the elements
- Expertise in detection probabilities for the elements
- Expertise in human factors associated with the complexity of defeating the elements

For a given scenario, for a given set of adversary resources, we can segregate each element into one of two bins:

1. We have sufficient data that with expert judgment we can generate a conservative probability distribution for that element, or

2. We have insufficient data to generate a probability distribution for that element.

Note that a given element may fall into bin #1 for one set of adversary resources and bin #2 for another set of adversary resources, since the available data may focus on a limited set of adversary resources. For elements in bin #2, we have the following choices:

a. Perform additional tests to allow the generation of a probability distribution using expert judgment and generate a probability distribution, or

b. Assign a conservative point-estimate value.

Step 4: Evaluate P_I using equation 1 for each set of adversary resources for each scenario of concern.

Step 5: For each scenario, assign degrees of evidence over the set of adversary resources using expert judgment. The following areas of expertise are needed:

- Expertise in intelligence associated with adversaries gathering and using attributes.

- Expertise in intelligence associated with adversaries gathering and using information.
- Expertise in human factors associated with adversary decisions for gathering and using resources.

Step 6: Evaluate $P_{I\ weighted}$ using Equation 8 for each scenario. If $P_{I\ weighted}$ is high and if the scenario has some elements that were modeled using conservative point estimates instead of probability distributions, consider performing more tests of these critical elements to generate probability distributions, then re-evaluate $P_{I\ weighted}$ with the newer, less conservative data for these elements.

7.2 Applying the Comprehensive Uncertainty Approach

To apply this approach, this significant effort requires assembling a team of experts, gathering many sources of data, and eliciting much expert opinion. Experts in security system elements, human factors, statistics, and threat assessment are required. NUREG-1563 provides useful guidance for expert elicitation. [NUREG-1563]

Due to the effort involved, it is recommended this approach be applied only to scenarios of concern where conservative point-estimate values result in prohibitively high costs of hardware and/or response force size to address those scenarios. Thus, a screening process is needed to first evaluate scenarios using conservative point estimates; only those of concern are retained for subsequent detailed evaluation using this approach. This screening process can be the one typically used where scenarios are identified – using such tools as tabletop exercises or path-finding tools such as ASSESS or ATLAS – and evaluated using conservative point estimates.

7.3 Capabilities Required for Process

This process (described in Section 7.1) requires the following capabilities:

- a team of experts in intelligence
- a team of experts in security system element detection probabilities and time delays
- a formal procedure for expert elicitation
- a database for security system elements
- testing capabilities for security elements
- a few experts for applying the mathematics of the technique to evaluate weighted P_I.

Figure 7-1 summarizes how these capabilities integrate into the process steps.

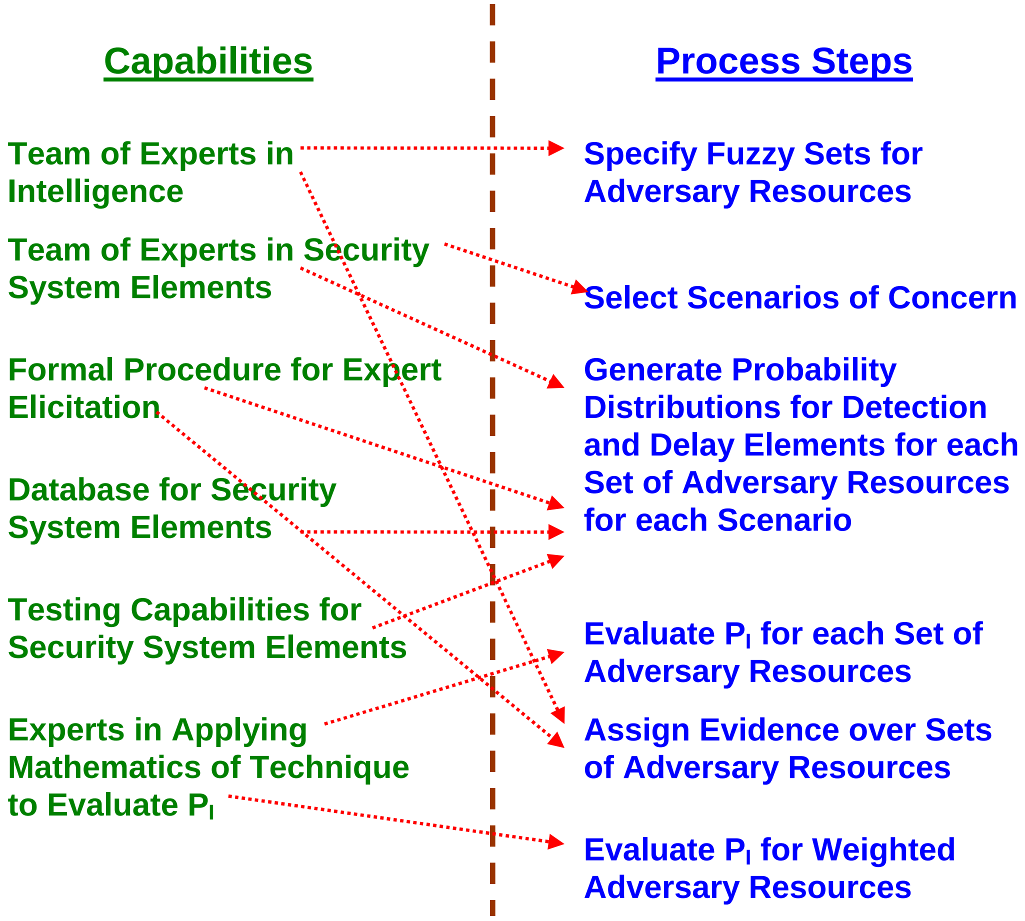
Capabilities
Process Steps
Team of Experts in Intelligence
Team of Experts in Security System Elements
Formal Procedure for Expert Elicitation
Database for Security System Elements
Testing Capabilities for Security System Elements
Experts in Applying Mathematics of Technique to Evaluate P_I
Specify Fuzzy Sets for Adversary Resources
Select Scenarios of Concern
Generate Probability Distributions for Detection and Delay Elements for each Set of Adversary Resources for each Scenario
Evaluate P_I for each Set of Adversary Resources
Assign Evidence over Sets of Adversary Resources
Evaluate P_I for Weighted Adversary Resources

Figure 7-1. Integration of Capabilities into Process Steps

It is recommended that a pilot application of this approach to a specific set of scenarios be performed. This will determine the usefulness, cost, and time required to apply the approach.

8 Extensions of the Technique

The mathematical approach developed in this report addresses uncertainty, with the focus on the Adversary. For example, the approach considers different sets of adversary resources as described in Section 5.

Uncertainties are associated with the Defender as well. Equation 1 considers uncertainty in the response time, T_R, and in the probability of communication, P_C, but only at a high level. T_R should be treated as dependent on the layer and should be segregated into constituent factors such as time for detection, time for assessment, time for communication, and time for response after communication. Similarly, P_C should be segregated into constituent factors. The incorporation of these details into Equation 1 is straightforward as they merely add more variables into the equation. These details could easily be addressed in future work.

Equation 1 uses a set of values for each security system layer, specifically, $\{T_{i1}, P_i, T_{i2}\}$ using the nomenclature of Table 1-1. At each layer, the adversary may attempt to defeat the layer using force, stealth, or deceit tactics, and the set of values for that element depends on the tactics. For example, stealth typically has a lower P_i than does force, but stealth typically has higher T_{i1} and T_{i2} than does force. In practice, the set $\{T_{i1}, P_i, T_{i2}\}$ uses the minimum times and minimum detection probabilities regardless of the tactics, and therefore the evaluation of P_I assumes that the adversary optimizes the tactics for defeating each element. In reality, the adversary does not know when to switch tactics and our evaluation of P_I is conservative. This conservatism should be addressed in future work, using some of the ideas that have been proposed for this issue. [Utility Theory and Path Timeline] [Snell Communication]

Here, we focused on the uncertainty in P_I. Overall effectiveness, P_E, is the product of P_I and P_N, where P_N is the probability of neutralization of the adversary by the response force given interruption. Uncertainty in neutralization, P_N, should be addressed.

The Design Analysis and Neutralization Technique Evaluation (DANTE) simulation framework addresses the uncertainty in P_N. [DANTE] The approach developed in this effort could perhaps be integrated with the simulation of P_N from DANTE.

9 Conclusions and Recommendations

For a specific attack scenario, the traditional measure of the effectiveness of the detection and delay elements of a physical security system is the probability that the security system detects the adversary in time for interdiction by the response force. This measure is denoted as P_I: the “probability of interruption.” This measure is typically evaluated using conservative, point-estimate values for the detection and time delay elements of the security system; the conservatism in the individual elements and in the overall P_I for a scenario is not evaluated. The result is significant expenditure of resources – hardware and response force personnel – to address such scenarios. Also, less attention is paid to balance protection across the entire range of possible scenarios. Two types of uncertainty are important:

1. ***Aleatory*** (random) uncertainty for detection probabilities and time delays for a given set of adversary resources.

2. ***Epistemic*** (state of knowledge) uncertainty for the adversary resources.

For a given set of adversary resources, there is aleatory uncertainty for detection probabilities and time delays. Also, there is epistemic uncertainty as to the adversary resources that will be brought to bear during an attack. Adversary resources consist of both attributes – such as equipment and training – and knowledge about the security system; to date, most evaluations have assumed an adversary with very high resources, thereby adding to the conservatism in the evaluation of P_I.

This work provides a mathematical technique to include both types of uncertainty to provide a more realistic evaluation of P_I.

Aleatory uncertainty for a given adversary is considered using probability distributions instead of conservative point-estimate values for element performance for a specific set of adversary resources. Epistemic uncertainty for the adversary resources is considered using plausibility from the belief/plausibility measure of uncertainty.

Sources of data were reviewed. More data exist on the uncertainty in delay times than on the uncertainty in detection probabilities. It is not possible to have a comprehensive “look up” database with probability distributions for every element for every set of adversary resources. Also, cost precludes performing enough tests to generate such a comprehensive database.

However, in many cases the data are sufficient so that expert judgment can generate approximate, conservative probability distributions for some of the security elements for a given set of adversary resources. Where the data are not sufficient, conservative point estimates can be used, or more specific testing can be performed to allow generation of the probability distributions.

The data sources also address the adversary resources used to defeat the security elements, thereby providing information for the consideration of different sets of adversary resources.

Application requires significant effort, including assembling a team of experts, gathering many sources of data, and eliciting much expert opinion. Experts in security system elements, human factors, statistics, and threat assessment are required.

This approach is recommended for application only to scenarios of concern where conservative point-estimate values result in prohibitively high costs of hardware and/or response force size to address that scenario. Thus, a screening process is used to evaluate scenarios and to retain only scenarios of concern. This screening process can be the one typically used where scenarios are identified – using such tools as tabletop exercises or path-finding tools such as ASSESS or ATLAS – and evaluated using conservative point estimates.

This report describes a process for implementing the approach, and identifies the capabilities needed to support the steps in that process.

It is recommended that a pilot application of this approach to a specific set of scenarios be performed. The pilot will determine the usefulness, cost, and time required to apply the approach.

References

[Access Delay Vol I] “Department of Energy Office of Safeguards and Security Technology Transfer Manual: Access Delay Volume I,” UCNI, SAND2001-2168, August 2001.

[ASSESS] Cousins, T.D., R.A. Al-Ayat, and J.C. Matter, “An overview of ASSESS - Analytic System and Software for Evaluating Safeguards and Security,” INMM 30th Annual Meeting, Proceedings 1989.

[ATLAS] Bhardwaj, M.K., “ATLAS Overview Presentation,” SAND2006-7380C, presented at Mod/Sim Working Group, San Antonio, TX, December 2006.

[Crystal Ball] Crystal Ball software, Version 7, Decisioneering, Inc.

[DANTE] private communication from Karen Page, Sandia National Laboratories, to John Darby, Sandia National Laboratories, August 11, 2008.

[EASI] H.R. Bennett, “User’s Guide for Evaluating Physical Security Capabilities of Nuclear facilities by the EASI Method,” SAND77-0082, NUREG-0814, June 1977.

[Exterior Detection] “Department of Energy Office of Safeguards and Security Technology Transfer Manual: Exterior Intrusion Detection,” UCNI, SAND99-2391, September 1999.

[Interior Detection] “Department of Energy Office of Safeguards and Security Technology Transfer Manual: Interior Intrusion Detection,” UCNI, SAND99-2388, Sept. 1999.

[Kaplan] Kaplan, S., and Garrick, B.J., 1981, “On the Quantitative Definition of Risk,” *Risk Analysis*, Vol. 1 No. 1, pp. 11-27.

[LinguisticBelief] Darby, J., “LinguisticBelief: A Java Application for Linguistic Evaluation using Belief, Fuzzy Sets, and Approximate Reasoning,” SAND2007-1299, March 2007.

[Nextgen] Jordan, S.E., “Next Generation Security Simulation,” SAND2002-1952C, presented at Institute of Nuclear Materials Management (INMM) 43rd Annual Meeting, June 2002, Orlando, FL.

[NUREG-1563] “Branch Technical Position on the Use of Expert Elicitation in the High-Level Radioactive Waste Program,” NUREG-1563, United States Nuclear Regulator Commission, November 1996

[Qualitative Uncertainty] Darby, J., “Uncertainty for Qualitative Variables,” SAND2007-6684P, Nov. 12, 2007.

[SASRAP] Snell, M., “Probabilistic Security Assessments: How They Differ from Safety Assessments,” SAND2002-0402C, presented at 6th International Conference on Probabilistic Safety Assessment and Management, San Juan, Puerto Rico, June 2002.

[SAVI] Matter, J.C., “SAVI: A PC-Based Vulnerability Assessment Program,” SAND88-1279, July 1988.

[Snell Communication] E-mails from Mark Snell, Sandia National Laboratories, to John Darby, Sandia National Laboratories, June 26, 2008.

[Terrorist Risk] Darby, J., “Evaluation of Risk from Acts of Terrorism: The Adversary/Defender Model using Belief and Fuzzy Sets,” SAND2006-5777, September 2006.

[Theory of Evidence] Shafer, S., *A Mathematical Theory of Evidence*, Princeton University Press 1976.

[Utility Theory and Path Timeline] Snell, M., “Utility Theory and Path Timeline Models,” SAND2008-1070A, presented at Institute of Nuclear Materials Management (INMM) 49th Annual Meeting, July 2008, Nashville TN.

[Vulnerability Assessment] “Technology Transfer Manual: Vulnerability Assessment” (UCNI), SAND2005-3929P, July 2005.

Distribution

Internal to Sandia:

1	MS0757	John Darby	06414 (electronic copy)
1	MS0759	Bruce Berry	06417 (electronic copy)
1	MS0759	Traci Brooks	06417 (electronic copy)
1	MS0757	Carla Ulibarri	06414 (electronic copy)
1	MS0791	Greg Wyss	06414 (electronic copy)
1	MS0791	Mark Snell	06754 (electronic copy)
1	MS0123	Donna L. Chavez	01011 (electronic copy)
1	MS0899	Technical Library	9536 (electronic copy)

Sandia National Laboratories