

Risk-Based Methodology for Scenario Tracking, Intelligence Gathering, and Analysis for Countering Terrorism

Barry M. Horowitz and Yacov Y. Haimen*

¹Department of Systems and Information Engineering, and Center for Risk Management of Engineering Systems, University of Virginia, Charlottesville, VA 22903

Received 8 September 2002; Revised 27 November 2002 and 26 January 2003; Accepted 8 March 2003
DOI 10.1002/sys.10043

ABSTRACT

Disruption of a terrorist attack depends on having information facilitating the identification and location of those involved in supporting, planning, and carrying out the attack. Such information arises from myriad sources, such as human or instrument surveillance by intelligence or law enforcement agencies, a variety of documents concerning transactions, and tips from a wide range of occasional observers. Given the enormous amount of information available, a method is needed to cull and analyze only that which is relevant to the task, confirm its validity, and eliminate the rest. The risk-based methodology for scenario tracking, intelligence gathering, and analysis for countering terrorism builds on the premise that in planning, supporting, and carrying out a terrorist plot, those involved will conduct a series of related activities for which there may be some observables and other acquirable evidence. Those activities taken together constitute a *threat scenario*. Information consistent with a realistic threat scenario may be useful in thwarting an impending attack. Information not consistent with any such scenario is irrelevant. Thus, the methodology requires a comprehensive set of realistic threat scenarios that would form a systemic process for collecting and analyzing information. It also requires a process for judging the validity and usefulness of such information. The key questions for intelligence gathering and analysis are: how to produce a compre-

*Author to whom all correspondence should be addressed (e-mail: haimes@virginia.edu).

hensive set of threat scenarios, how to winnow that set to a subset of most likely scenarios, what supplementary intelligence is worth pursuing, how to judge the relevance of available information, and how to validate and analyze the information. The methodology presented in this paper can serve as a vehicle with which to enable the intelligence community to better: (a) assess the intent and capabilities of terrorist groups, (b) develop and compare terrorist scenarios from different sources and aggregate the set that should guide decisions on intelligence collection, (c) assess the possible distributions of responsibility for intelligence gathering and analysis across various homeland security agencies at the federal, state, and local levels, and (d) establish effective collection priorities to meet the demands of counterterrorism. Some of the critical issues addressed in this paper include: (1) how to create a reasonably complete set of scenarios and filter it down to a more manageable set to establish intelligence collection priorities, (2) how to integrate the wide variety of intelligence sources associated with monitoring for terrorism and analytically account for the corresponding disparities in information reliability, and (3) how to incorporate these new methodologies into existing information management efforts related to protecting our nation's critical infrastructures. © 2003 Wiley Periodicals, Inc. Syst Eng 6: 152–169, 2003

Key words: risk analysis, scenario tracking, terrorism, hierarchical holographic modeling, Bayesian analysis, multiobjective decision trees

1. INTRODUCTION

As a part of addressing new intelligence-system needs related to antiterrorism, there is significant interest in scenario-based tracking. (See NRC [2002] for a comprehensive document on the roles of science and technology in countering terrorism, which includes a chapter on systems engineering.) In response to this need, this paper develops a holistic generation of possible terrorist scenarios which envisage potential attacks on specific targets. New issues must be addressed relative to traditional intelligence-gathering efforts. These issues relate to the design of an analytically based process for identifying the great diversity of potential attacks, and, correspondingly, a process for filtering this number down to a more manageable number of high-risk scenarios that would actually be tracked.

There are two principal motivations for focusing specific attention on scenario-based tracking. First, scenario analysis will illuminate new items of information for collection; as a result, it will increase the performance of existing approaches. Second, tracking scenarios will increase the likelihood of intercepting potential terrorist actions, which may already be in progress. The design approach developed in this paper is based on applying a legacy of methodologies that have been developed for risk assessment and risk management. (Risk is defined as a measure of the probability and severity of adverse effects [Lowrance, 1976].) For example, an imperative objective is to be able to detect beforehand that a terrorist effort is underway to poison food or a source of water. In order to accomplish this

goal, detailed scenarios of potential attacks must be developed to guide intelligence collection. These scenarios must include possible terrorist actions and behaviors during the attack-planning phases, since intelligence-based observations of such early activity would signal that an attack scenario is in progress and prevent its execution. Computer security is one area of activity that has taken this approach. Intrusion-detection systems related to computer security provide a methodology that can potentially be extended to general scenario structuring. Bace [2000] examines several approaches to misuse detection and anomaly detection. *Misuse detection* focuses on identifying known use patterns. Pouzol and Ducassé [2001] propose using the signature specification language *Sutekh* to develop misuse-detection algorithms. Valdes and Skinner [2001] present a probabilistic approach to alertness correlation. Vigna, Kemmerer, and Blix [2001] describe a new approach to support a highly configurable intrusion-detection sensing infrastructure. Currently, surveillance sensors developed on an ad hoc basis are difficult to configure, extend, and remotely control. Debar and Wespi [2001] present an algorithm to aggregate and correlate intrusion-detection alerts. Their work is intended to address the current intrusion-detection weaknesses of flooding, context, false alerts, and scalability. Cunningham and Stevenson [2001] describe their efforts in detecting attack code before the attack is actually executed. Intrusion Detection Systems (IDS) that audit data afford attackers the opportunity to exploit the delay between attack execution and detection to disable the IDS.

However, when developing terrorist scenarios, two points not addressed in the above literature quickly become apparent: (1) There are a large number of potential scenarios to track, some related to a specific target and some generally aimed, and (2) both the quality and evaluation of the scenarios depend on in-depth specialized knowledge. Therefore, a filtering process is needed for comparing and choosing the set of scenarios that are most worthwhile to track. In turn, a large team of people may be required to create and assess the huge potential-scenario set for tracking. Needless to say, this must include team members with highly specialized knowledge. The dual need for capacity and specialization leads us to delegate a major role in scenario tracking to government organizations whose everyday focus is on areas related to specific scenarios. For example, an agriculture agency can develop scenarios related to poisoning food, while a transportation agency can develop scenarios related to highway attacks. Once developed in the specialized agencies, the scenarios can be relayed to the intelligence community to be tracked effectively by an integrated central intelligence collection and analysis system.

This paper is organized as follows:

1. A scenario-tracking process, or methodology, is introduced, using diverse intelligence and knowledge sources.
2. The objectives of the methodological-based process are explained.
3. Specific technical methods are described. The integration of these methods constitutes the suggested methodology. They are: Hierarchical Holographic Modeling (HHM) for scenario development and characterizations of observable terrorist actions; Risk Filtering, Ranking, and Management (RFRM) for filtering and ranking the classes of observables; Bayesian analysis for updating the intelligence; and Multiple-Objective Decision Trees (MODT) for making decisions that build on tradeoffs between time, effort, and risk.
4. An example is presented associating these results with existing intelligence data-processing systems.

2. OVERVIEW OF THE METHODOLOGY

There is a need for the development and use of a methodological process that will enable the intelligence community to:

- (a) assess the intents and actions of terrorist networks and our associated needs for intelligence gathering;
- (b) develop and compare terrorist scenarios from different sources and aggregate the set that should guide decisions for intelligence collection;
- (c) assess the possible distributions of responsibility for intelligence gathering and analysis across various homeland security and intelligence collection agencies at the federal, state, and local levels.

Equally important, the selected scenarios can become part of a standard process for conversion into the intelligence gathering and analysis systems. This process and the associated methods and tools must be applicable to a broad and diverse set of possible threats across a wide range of potentially threatened domains. The participants selected to carry out the methodology must be willing to explore scenarios that might be considered farfetched by competing methods and evaluation teams. The professionals in the domains being evaluated for potential terrorist attacks must be able to understand the methodologies, evaluate them, and contribute to their adaptation. The methodologies also must have an analytical basis to enable comparing risks from different domains, so that the process for filtering out less important scenarios is both manageable and reasonably objective [NRC, 2002]. Finally, the methodologies should have a successful record of solving a related set of problems, so that the antiterrorist capability has as great a chance of success as possible. For determining those scenarios most likely to be worth tracking, an appropriate holistic methodology is needed. Once the universe of scenarios has been identified, a systemic process capable of filtering and ranking the set of scenario is required.

The value of intelligence depends to a large extent on its credibility and on continuous updating. For this purpose, we apply Bayesian analysis of the sensitivity and specificity of intelligence observables related to particular scenarios. Once a set of scenarios has been selected, an immediate requirement is to analyze the potentially observable actions that terrorists might take in order to prepare for and execute an attack. For this purpose, we can divide the set of potential observables into three subsets:

- (1) those that would be collected by the intelligence community,
- (2) those that would be submitted by the general public, as a result of scenario-specific government advisories or circulars, and

(3) those that would be collected via new requirements established by the appropriate governmental department, to be carried out by the industrial base under its domain of operation.

In all three cases, to determine what is worth collecting, it would be necessary to consider the importance of a particular scenario coupled with the value of a particular data item, the effort (cost and time), and the risk of collecting and processing the desired information. In addition, information that is relevant to a large-enough number of scenarios can become part of the general collection that is considered *scenario-independent*. Based on the effort and risk involved, certain information would only be collected when the value of the potential new data was sufficient in terms of making decisions. In this area, we recognize that collections from the three subsets are likely to have very different reliabilities and varying levels of criticality to actual decision-making. Therefore, Bayesian analysis is presented as the foundation for assessing the potential significance of the information.

Given the impact of extreme and catastrophic terrorist attacks, a risk-of-extremes metric is needed that supplements and complements the expected-value-of-risk metric. Finally, a decision-making mechanism is needed that can utilize the added knowledge derived from newly-discovered intelligence and make use of Bayesian analysis. In particular, the noncommensurate objectives—effort (cost and time) and risk—must be addressed in the multiobjective tradeoff analysis. Follow-up actions could vary from calling for special new

information, to calling in experts to further evaluate the data, to initiating interception of the anticipated terrorist activity.

The methodologies used for analyzing results must be consistent with the scenario-structuring and observation-collection techniques. This would require integrating the existing methodologies used by the intelligence community with other appropriate methodologies. Therefore, the efforts of the designers of scenario methodology and collection-analysis methodologies must be integrated. In addition to the efforts to apply risk-based analysis for scenario tracking, there is a need for the scenario-structuring communities to collaborate with the intelligence-analysis communities in order to create appropriately integrated data-processing support systems.

Figure 1 constitutes a graphical representation of the methodology, where collected intelligence from myriad sources is tracked and analyzed through structured scenarios. The comprehensive nature of scenario structuring necessitates a process of filtering and ranking. As new intelligence is gathered, Bayesian analysis adds value and credibility to the intelligence (through posterior probabilities). To account for extreme and catastrophic events, the expected-value metric must be supplemented by a new extreme-event metric. Finally, decisions must be made over time; what is needed is a multiobjective-based sequential decision-making process under uncertainty, where cost, time, and risk are analyzed and traded off.

The remainder of this paper elaborates on each of the above points. Section 3 discusses methodologies for

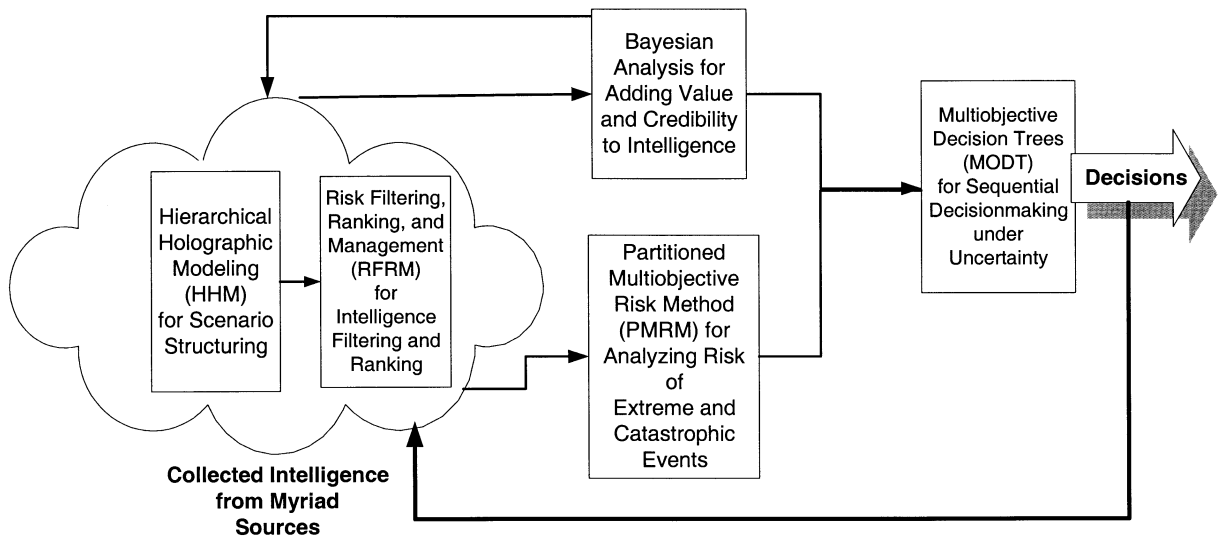


Figure 1. Graphical representation of methodology with specific methods.

scenario structuring as they relate to scenario-based tracking. Section 4 presents an example application of the methodology. Section 5 discusses integrating scenario-related collection and analysis methodologies. Conclusions are presented in Section 6.

No single model, or methodology can effectively meet the technical challenges posed by (1) anticipating and tracking terrorism through scenario generation and structuring, (2) updating and quantifying the value of intelligence through Bayesian analysis, (3) assigning priorities to the scenarios in a well-established risk-based methodology, (4) evaluating the cost-effectiveness of the entire process of intelligence gathering and analysis, and (5) tracking terrorists' attack plans. To meet these challenges, the proposed methodology builds on, modifies, and integrates several appropriate risk-based techniques.

3. THE METHODOLOGY

3.1. Scenario Generation and Structuring through Hierarchical Holographic Modeling (HHM)

It is impracticable to represent within a single model all the aspects of a large-scale system, such as the homeland's myriad vulnerabilities to terrorism networks. Several methods have been developed over the years to address the complexity of modeling large-scale systems and to offer remedies. For example, Kott, Pollack, and Krogh [1999] apply a *Situation Assessment Problem* to the military command and control process, and identify a host of gaps that must be overcome in this difficult task. Josang [2002] describes a framework for combining and assessing subjective evidence from different sources. In *Synerctics, the Development of Creative Capacity*, Gordon [1968] introduces an approach that uses metaphoric thinking as a means to solving complex problems. Other early works include a book on societal systems and complexity by Warfield [1976] and a book on large-scale systems by Sage [1977]. Other scenario-structuring methods include *Anticipated Failure Determination (AFD)* [Kaplan et al., 1999] and *TRIZ* [Altshuller, 1988, 1994; Kaplan, 1996].

Hierarchical Holographic Modeling (HHM) [Haimes, 1981, 1998] forms the basis for the proposed methodology. HHM reflects a difference in kind from other modeling schemas and emerged from a generalization of Hierarchical Overlapping Coordination (HOC). In HOC (see Haimes and Macko [1973] and Macko and Haimes [1978]), a system's single model is divided into several decompositions in response to the various aspects of the system, and these decompositions are coordinated to yield an improved solution. Although

both HOC and HHM are quantitative methods, only the conceptual aspect of HHM is deployed in this paper. A quantitative discussion on HOC and HHM can be found in *Hierarchical Multiobjective Analysis of Large-Scale Systems* [Haimes et al., 1990b]. The name Hierarchical Holographic Modeling is suggested by holography—the technique of lensless photography. The difference between conventional mathematical modeling techniques (yielding what might be termed planar models) and the HHM schema is analogous to the difference between holography and conventional photography, which captures only two-dimensional planar representations of scenes. In the abstract, a mathematical model may be viewed as a one-sided image of the real system which it portrays. For example, with single-model analysis and interpretation, it is quite impossible to identify and document the sources of risk associated with not only the multiple components of an infrastructure (e.g., transportation or hydroelectric power structure or food processing plants), but also with their welter of societal aspects (functional, temporal, geographical, economic, political, legal, environmental, sectoral, institutional, etc.). Central to the mathematical and systems basis of holographic modeling (from the perspective of theoretical constructs) is the overlapping among various holographic models with respect to the objective functions, constraints, decision variables, and input-output relationships of the basic system. Arthur D. Hall III, whose first book on systems engineering was published in 1962 [Hall, 1962], recognized the contributions of HHM in his seminal book *Metasystems Methodology* [Hall, 1989, p. 6]: “In this way,” he wrote, “history becomes one model needed to give a rounded view of our subject within the philosophy of hierarchical holographic modeling [Haimes, 1981] being used throughout this book, defined as using a family of models at several levels to seek understanding of diverse aspects of a subject, and thus comprehend the whole.”

HHM has turned out to be particularly useful in modeling large-scale, complex, and hierarchical systems, such as defense and civilian infrastructure systems. As pointed out by Kaplan, Haimes, and Garrick [2001], it can be regarded as a general method for identifying a set of risk scenarios. At the same time, we are reminded that the process of identifying the risk scenarios for a system of any kind should begin with a diagram that represents the “success” or “as-planned” scenarios.

The multiple visions and perspectives of HHM add strength to risk analysis. It has been extensively and successfully deployed to study risks for government agencies such as the President's Commission on Critical Infrastructure Protection (PCCIP), the FBI, NASA,

the Virginia Department of Transportation (VDOT), and the National Ground Intelligence Center, among others. The HHM methodology/philosophy is grounded on the premise that in the process of modeling large-scale and complex systems, more than one mathematical or conceptual model is likely to emerge. Each of these models may adopt a specific point of view, yet all may be regarded as acceptable representations of the infrastructure system. Through HHM, multiple models can be developed and coordinated to capture the essence of the many dimensions, visions, and perspectives of infrastructure systems. One example is the study conducted for the PCCIP on the U.S. water supply system. Sixteen different visions/perspectives (*head-topics*) with an additional 94 subvisions (*subtopics*) were identified as sources of risk [Haimes et al., 1998; Haimes, 1998]. These modeling perspectives are: physical (pipes, pumps, wells, aqueducts, water treatment plants, wastewater treatment plants), scope, temporal, maintenance, institutional, organizational, management, resource allocation, SCADA, system's configuration, hydrology, geography, external factors, buffer, contamination, and ground and surface water. Other applications of HHM are discussed in Haimes [1998, Chap. 3], Lambert et al. [2001], and Dombroski et al. [2002].

Perhaps one of the most valuable and critical aspects of HHM is its ability to facilitate the evaluation of the subsystem risks and their corresponding contributions to the risks in the total system. In the planning, design, or operational mode, the ability to model and quantify the risks contributed by each subsystem markedly facilitates identifying, quantifying, and evaluating risk. In particular, HHM has the ability to model the intricate relationships among the various subsystems and to account for all relevant and important elements of risk and uncertainty. This makes for a more tractable modeling process and results in a more representative and encompassing risk assessment process.

To present a holistic view of the elements that must be included in the model, the HHM approach involves organizing a team of experts with widely varied experience and knowledge bases (e.g., technologists, psychologists, political scientists, criminologists, and others). The broader the base of expertise that goes into identifying potential risk scenarios, the better is the comprehensiveness of the ensuing HHM. The result of the HHM process is a very large number of risk scenarios, hierarchically organized into sets and subsets. If done well, the set of scenarios at any level of the hierarchy would approach a "complete set." The result of the HHM effort is organized into what is called the candidate *scenario model*. The HHM approach [in conjunction with the Risk Filtering, Ranking, and Management (RFRM) method to be discussed subsequently]

then ranks the elements of the candidate scenario model, giving strong preference to those elements that are considered most important from several different areas of expertise. The result is a filtered scenario model, and the search and analysis efforts for tracking terrorists then would be designed based on this model. For this application of HHM, the elements that would be part of the scenario model would be time-variant, depending on the resulting intelligence-collection system workload. The HHM-derived model would automatically add or delete elements as a function of system workload by referring to a master model that incorporates all of the HHM-identified elements. This automatic feedback process requires practical subsystem workload measures to be defined and monitored. Research efforts are required to develop and experiment with different approaches for adjusting the models dynamically and for modifying the detailed search parameters correspondingly. Figure 1 is a representation of the methodology with the specific methods indicated.

3.1.1. Intelligence Filtering and Ranking

It is clear that the first and most important step in a quantitative risk analysis is identifying the set of risk scenarios. If the number of such scenarios is large, then (as noted above) the second step must be to filter and rank the scenarios according to their importance, as determined by their likelihoods and consequences. The need for such ranking arises in a variety of situations. For example: Thousands of military and civilian sites have been identified as contaminated with toxic substances; myriad risk scenarios are commonly identified during the development of software-intensive engineering systems; and thousands of mechanical and electronic components of the Space Shuttle are placed on a critical item list in an effort to reveal significant contributions to program risk. In all such risk-identification procedures we must then prioritize a large number of risk scenarios according to their individual contributions to the overall system risk. Dependable and efficient ranking and filtering of identified risk elements can be an important aid toward systematic risk control and reduction. Other methodologies for filtering and ranking include Morgan et al. [1999, 2000], Baron, Hershey, and Kunreuther [2000], Weblar et al. [1995], and Sokal [1974].

3.2. Risk Filtering, Ranking, and Management

Risk filtering, ranking, and management (RFRM) [Haimes et al. 2002] is a methodology developed to identify, prioritize, assess, and manage risk scenarios of a large-scale system. The RFRM is a modified and

much-improved version of the risk ranking and filtering (RRF) [Haimes, 1998] developed a decade ago for NASA for the Space Shuttle qualitative screening of scenarios and classes of scenarios. The eight phases of the RFRM methodology are:

1. Phase I, *Scenario Identification*—a hierarchical holographic model (HHM) is developed to describe the “as-planned” or “success” scenario.
2. Phase II, *Scenario Filtering*—the risk scenarios identified in Phase I are filtered according to the responsibilities and knowledge of the domain experts.
3. Phase III, *Bi-Criteria Filtering and Ranking* (based on likelihood and consequences) is applied.
4. Phase IV, *Multi-Criteria Evaluation* (based on 11 tabulated criteria of risk) is applied.
5. Phase V, *Quantitative Ranking*—the filtering and ranking of scenarios continues based on quantitative and qualitative matrix scales of likelihood and consequence, as well as ordinal-scale response to scenario resiliency, robustness, and redundancy.
6. Phase VI, *Risk Management*—intelligence collection options for dealing with the filtered scenarios are identified, and the cost and the potential for intercepting a terrorist attack for each are estimated.
7. Phase VII, *Safeguarding against Missing Critical Items*—scenarios previously filtered out in Phases II to V are reexamined and compared to the consequences, cost, and intercept potential of the selected options.
8. Phase VIII, *Operational Feedback*—experience and information gained during this application are used to refine the scenario filtering and decision processes in earlier phases.

These eight phases manifest a philosophical approach rather than a mechanical one. In this philosophy the filtering and ranking of discrete scenarios is viewed as a precursor to, rather than a substitute for, considering the totality of all risk scenarios. *Note that no system of weights in the prioritization process is used in the RFRM.* This attribute is central to the repeatability of the results. Indeed, with weights one can “cook the books” and manipulate the results.

3.3. Bayesian Analysis

The National Security Agency collects a vast number of pieces of information—about 2,000,000 data items per day—and this represents a fraction of all collections

that are part of our nation’s antiterrorism campaign. The need for giving priorities to a flood of intelligence has been discussed above. A second need is to add synergy to this vast database by “connecting the dots” of widespread intelligence collection. In the field of counterterrorism, the opportunity exists to build upon extensive experience in applying and contributing to the theory, methodology, and practice of Bayesian analysis [Kaplan, 1990, 1992; Kaplan and Garrick, 1981; Kaplan, Haimes, and Garrick, 2001]. A wide range of applications have been addressed with Bayesian analysis. For example, a Bayes-based methodology for the carcinogenicity prediction of chemicals and battery selection of assays (CPBS) was developed and successfully deployed [Pet-Edwards et al., 1985a, 1985b; Chankong et al., 1985; Haimes, 1998]. Another example is the Bayesian model that was developed for the U.S. Army Corps of Engineers for flood warning and evacuation systems [Haimes, Li, and Tulsiani, 1990c; Haimes, 1998].

The value of intelligence depends on its credibility—more specifically, on the credibility of its sources. In particular, when the public at large is encouraged to report suspicious activities, a welter of pieces of information from unknown sources may flood government agencies. To cope with this information overload, a two-step process is required: (1) a screening mechanism with which to filter and rank the incoming messages, and (2) a scientifically based process with which to assign a level of credibility to the remaining messages. In the proposed risk-based methodology, the RFRM methodology is the first step and a Bayesian-based process is deployed as the second step.

Adding credibility to the intelligence that survives the risk filtering and ranking can be related to Bayes’ Theorem through the following three phases. (1) An intelligence report whose credibility is unknown (i.e., the probability of its trustworthiness is unknown) is termed *prior probability*, and may be assigned any appropriate value, such as 0.5. (2) Through Bayes’ Theorem, an added level of confidence in the intelligence (prior probability) can be achieved by using new evidence (e.g., verification by an intelligence agency such as the FBI or the CIA). This new level of information trustworthiness is termed *posterior probability*. (3) To use Bayes’ Theorem, we need to assess the credibility of intelligence agencies or individuals that might be used to corroborate any of the myriad messages received on terrorist activities. This scientific corroboration will be based, among other criteria, on the historical (statistical) number of false-positive and false-negative assessments made by these intelligence agents or entities. A false-positive assessment means that the intelligence entity assesses an impending activity as true

when it is not true. In a false-negative assessment, the intelligence entity assesses an impending activity as false when it is true. In several fields (such as environmental and health sciences), the measures of *sensitivity* and *specificity* are commonly used to represent the false positive and false negative, respectively.

Caveat: Our Bayesian analysis is predicated on the concept of making decisions when the change in the conditional probabilities is large, and is not based on the actual values of the probabilities themselves. Furthermore, the credibility of intelligence collected by an agency largely depends on its source, location, circumstances, the agent(s) involved, and other factors. Thus, one cannot assign, for example, sensitivity or specificity to an agency per se. The best one can expect is to assign credibility to a specific intelligence report. The above caveat notwithstanding, intelligence agents believe that such a quantitative process, if adopted in the future, could markedly enhance the cost-effectiveness of intelligence collection and analysis.

Sensitivity of an intelligence source ($\alpha+$) is defined as the ratio of the number of intelligence messages indicating a threat (T), given that the threat was real, divided by the total number of messages received from this source.

Specificity of an intelligence source ($\alpha-$) is defined as the ratio of the number of intelligence messages indicating no threat (NT), given that the threat was not real, divided by the total number of messages received from this source:

$$(\alpha+) = \Pr(+ | T), \quad (\alpha-) = \Pr(- | NT).$$

Of course, we are interested in ascertaining the probability that a threat is indeed real, especially when the information is received from a heretofore unknown source of intelligence (i.e., a new source whose sensitivity and specificity are unknown). This can be achieved through Bayes' formula. Let

($\theta+$) = the probability that the threat information is real, given that the intelligence agency whose sensitivity and specificity are known to us indicate that the threat is real,

($\theta-$) = the probability that the threat information is false, given that the intelligence agency whose sensitivity and specificity are known to us indicate that the threat is false,

$$(\theta+) = \Pr(T | +) = [\Pr(T) \Pr(+ | T)] / [\Pr(T) \Pr(+ | T) + \Pr(NT) \Pr(+ | NT)],$$

$$(\theta-) = \Pr(NT | -) = [\Pr(NT) \Pr(- | NT)] / [\Pr(T) \Pr(- | T) + \Pr(NT) \Pr(- | NT)]$$

As an example of the application of Bayes' analysis, consider the following U.S. Department of Homeland Security's definition of various threat conditions: There are five *threat conditions*, each identified by a level description and corresponding color. From lowest to highest, these are: *Low = Green; Guarded = Blue; Elevated = Yellow; High = Orange; Severe = Red*. The higher the threat condition, the greater is the risk of a terrorist attack. Risk includes both the probability of an attack occurring and its potential gravity. As established above, the Bayesian estimation of the potentials for missed warnings and false detections (sensitivity and specificity) can be applied generally here. To do so: (i) Define the emergence of sets of behaviors in the scenario evolution, (ii) relate the sets of behaviors to threat conditions that are the basis for action by risk managers, and (iii) relate the threat conditions to the estimations of missed warnings (sensitivity) and false detections (specificity). Importantly, the analysis will delineate extended sensitivity and specificity factors associated with transitions among threat conditions.

In sum, intelligence that amounts to information overload with unknown credibility can be evaluated through two steps: (1) filtering and ranking using the RFRM, and (2) using Bayes' Theorem to determine its trustworthiness.

3.3.1. Assessing Risks of Extreme and Potentially Catastrophic Events

In scenario tracking for terrorism, both the probabilities and consequences are considered when evaluating the importance of any specific scenario. The most common metric used in measuring risk is the expected value of risk. However, in the face of such unforeseen calamities as the terrorist attacks of September 11, 2001, we are more willing to acknowledge the importance of studying extreme events. We are no longer asking questions about expected risk; instead, we are asking questions about expected maximum risk. Indeed, decision-makers involved in homeland security are likely to be most concerned with risks associated with a specific terrorist scenario, rather than with the likelihood of the average adverse outcomes that may result from various terrorist risk scenarios. In this sense, the expected value of risk, which has for years dominated most risk analysis in the field, is not only inadequate, but it can lead to fallacious results and interpretations. Furthermore, people in general are not risk-neutral. They are often more concerned with low-probability catastrophic events than with more frequently occurring but less severe accidents or events. In some cases, a slight increase in the cost of modifying a structure might have a very small effect on the unconditional expected risk (the commonly used business-as-usual measure of risk), but would make a

significant difference to the conditional expected catastrophic risk (to be defined subsequently). Consequently, the conditional expected catastrophic risk can be of significant value in many multiobjective risk problems.

From the perspective of public policy, imagine that decision-makers are faced with two possible scenarios. One is a catastrophic dam failure, which might cause flooding of 10^6 acres of land with associated untold human fatalities, as well as damage to the economy and the environment; however, this has a very low probability of 10^{-6} of happening. The other scenario is minor flooding of 10^2 acres of land that has a high probability of 10^{-2} of happening. Obviously, extreme and catastrophic events cannot be viewed by decision-makers in the same vein as common events. Yet this is exactly what the expected-value function would ultimately generate. Most importantly, analysts' precommensuration of these low-probability/high-damage events with high probability/low-damage events into one expectation function markedly distorts the relative importance of these events and consequences as they are viewed, assessed, and evaluated by decision-makers [Mitsopoulos, Haimes, and Li, 1991; Frohwein, Haimes, and Lambert, 2000; Haimes, 1998; Haimes et al., 1990a].

3.4. The Partitioned Multiobjective Risk Method (PMRM)

To introduce the Partitioned Multiobjective Risk Methodology (PMRM) [Asbeck and Haimes, 1984; Haimes, 1998], consider that a continuous random variable X of damages has a cumulative distribution function $P(x)$ and a probability density function $p(x)$, which are defined by the relationships $P(x) = \text{probability}[X \leq x]$. The expected value of the risk is defined as

$$E[X] = \int_0^{\infty} x p(x) dx$$

In the PMRM, the concept of the expected value of damage is extended to generate multiple *conditional expected-value functions*, each associated with a particular range of exceedance probabilities or their corresponding range of damage severities. A *conditional expectation* is defined as the expected value of a random variable, given that this value of the random variable lies within some prespecified probability range. Clearly, the values of conditional expectations are dependent on where the probability axis is partitioned. The choice of where to partition is made by the analyst/decision-maker in response to the extreme charac-

teristics of the problem. The resulting conditional expected-value functions, in conjunction with the traditional expected value, provide a family of risk measures associated with a particular policy. For example, the conditional expected value risk of extreme and catastrophic events, $\psi(\cdot)$, is defined as

$$\psi(\cdot) = E[X | \beta \leq x],$$

where β is the partitioning point on the damage axis, e.g., the number of fatalities. Note that probability density functions for terrorism scenarios can be generated on the basis of expert evidence, utilizing the fractile method or the triangular distribution [Haimes, 1998]. Furthermore, the fundamental tenet upon which the PMRM is based is that for safety-critical systems, where the consequences are so severe and unacceptable, a low probability of occurrences should not be undervalued by the averaging process of the expected value of risk metric. Thus, in scenario tracking for terrorism, intelligence collectors and analysts should not overlook the importance of the low-probability and catastrophic consequences of any given scenario. Furthermore, they should subject such scenarios to the systemic multiobjective tradeoffs among all relevant costs, benefits, and risks associated with the assessment and management of the associated risks, as advocated by the PMRM:

$$\psi(\cdot) = \frac{\int_{\beta}^{\infty} xp(x) dx}{\int_{\beta}^{\infty} p(x) dx}$$

In sum, the PMRM is an enabling methodology that supplements and complements the expected value of risk metric for extreme and catastrophic events by introducing the conditional expected value of risk. This renders the scenario tracking, intelligence gathering, and analysis for countering terrorism more realistic and effective.

3.4.1. Decisionmaking

Decision-tree analysis has emerged over the years as an effective and useful tool in decision-making. More than two decades ago, Howard Raiffa [1968] published the first comprehensive and authoritative book on decision-tree analysis. Ever since, its application to a variety of problems from numerous disciplines has grown exponentially. Advances in science and scientific approaches to problem-solving are often made on the basis of the earlier works of others. In this case, the foundation for Raiffa's contributions to decision-tree analysis can be

traced to the works of Bernoulli on utility theory (see von Neumann and Morgenstern [1953]).

3.5. Multiple-Objective Decision Tree (MODT) for Intelligence Decision-Making

The single-objective models that were advanced in the fifties, sixties, seventies, and eighties are today considered by many to be unrealistic, too restrictive, and often inadequate for most real-world complex problems. During the last decade or two, there has been a proliferation of books, articles, conferences, and courses on what has come to be known as Multiple-Criteria Decision-Making (MCDM). This is a vivid indicator of the maturation of the field of decision-making (see, for example, Chankong and Haimes [1983]; Haimes [1998]). In particular, an optimum derived from a single-objective mathematical model, including that derived from a decision tree, often may be far from representing reality, and thereby may mislead analysts as well as decision-makers. In particular, the Multiple-Objective Decision Tree (MODT) methodology presented here for scenario tracking of terrorism must analyze at least three noncommensurate objectives—cost, time, and risk. This is done by folding back Pareto optimal solutions, as opposed to a single utility function in the conventional single-objective decision tree. Given a system of multiple objectives,

$$\min_{x \in X} \{f_1(x), f_2(x), \dots, f_n(x)\},$$

where x is an N -dimensional vector of decision variables, and X is the set of all feasible solutions,

$$X = \{x \mid g_i(x) \leq 0, \quad i = 1, 2, \dots, m\},$$

a decision x^* is said to be a Pareto optimal solution to the above system of multiple objectives if and only if there does not exist another \bar{x} so that

$$f_j(\bar{x}) \leq f_j(x^*), \quad j = 1, 2, \dots, n,$$

with strict inequality holding for at least one j . (In other words, one can improve one objective function only at the expense of degrading another.)

Indeed, decision trees can better serve both analysts and decision-makers when they are extended to deal with the above multiple objectives. Furthermore, MODT also incorporates risk of extreme, rare, and catastrophic events via the Partitioned Multiobjective Risk Method (PMRM) [Asbeck and Haimes, 1984; Haimes, 1998], which also has been applied to studies for several agencies.

There is a world of difference between MODT and decision trees with multiple objectives that utilize a single utility function to convert the multiple objectives into a single utility function. In MODT, the analysis proceeds forward with all objectives addressed in their own noncommensurate units, and folding back with Pareto optimal solutions. Therefore, at the first decision node, decision-makers make explicit tradeoffs among the various objectives. Contrary to this approach, the decision-tree analysis is likely to be distorted and lose credibility with the use of weights to convert the multiple objectives into a single one at the beginning of the decision tree and/or the generation of a utility function to represent the decision-maker(s) preferences at the start. Indeed, how can any analyst solicit and obtain sufficient and credible responses from decision-makers to be able to construct a utility function that relates the risks from weapons of mass destruction to the cost of the overall protection from such attacks?

3.6. Analysis of Observable Actions

The scenario structuring described in Sections 3.1 and 3.2 must be converted into a set of observable activities that an intelligence collection operation can use as the basis for discovering terrorist scenarios that are in progress. This section presents initial ideas on how to identify the observables of a scenario. We begin by dividing a potential terrorist attack into six stages (see Table I).

These six steps, which serve as divisions of a terrorist action, can be used to consider what is observable in a scenario. To establish observables, an appropriate team of intelligence and domain experts must collaborate to hypothesize the detailed steps the terrorist would need to take. For example, the intelligence team would determine what communications and interactions might be observed. A specific scenario might include a variety of observables, depending on the specific attack methods. As a result, each terrorist act spawns multiple observable scenarios that the intelligence analysis systems must keep linked together. This process also creates a structure for dealing with when events might occur and with the timing involved in each part of a scenario. Most importantly, this set of information could be used to determine the actions to be taken before, during, or after an attack.

Returning to the HHM structure, one must begin by determining which parts of the model provide potential observation points. Note that the HHM elements (without further analysis) represent a set of unlinked vulnerabilities that a terrorist can exploit by developing a specific operational plan. The method for observing each HHM element must be determined, so that as real

Table I. Six Stages for Identifying Observable Terrorist Scenarios

Stage	Description
Intent	Earliest stage, where the terrorist develops malice and an intent to harm via a general plan of attack
Target Acquisition	At this stage, the terrorist chooses specific target(s)
Plan	The terrorist researches the target(s) and various attack options
Preparation	Full commitment stage: At this point, the wheels are in motion as the terrorist prepares to launch the attack
Execution	The attack is carried out
Grace Period	Depending on the nature of the attack, there is sometimes a time-lag between a successful attack and its impact; for example, poisoning of food does not result in harm until someone eats the food

observations are made, both the likelihood of the scenario and a corresponding operational plan can be revealed. The likelihood that a scenario is in progress does not require initiating a specific operational plan, but determining the time from execution to completion of a planned terrorist action does. Correspondingly, since it is desirable that decisions resulting from a set of observations will account for the remaining time before the actual plan is executed, this part of the process must be tightly linked to decision-making. The HHM can be used to develop the possible steps that lead to the completion of a plan once it is in progress. This, in turn, can be the basis for estimating the time it takes to complete a plan. Timing estimates can then be coupled with the observables associated with completing possible plans. This is to establish new priorities on collecting specific observations that would provide a stronger indication that a plan is indeed underway, and how much time might exist prior to its completion.

4. HHM FOOD-POISONING SCENARIO: BAYESIAN ANALYSIS

The following specific example illustrates the use of the methodology for scenario-based tracking of terrorism. This involves the possibility of a terrorist organization poisoning meat products at a slaughterhouse. For the purposes of this paper, a limited HHM analysis is presented that attempts to illuminate the various possibilities a terrorist organization might consider in developing an attack scenario (see Fig. 2). A more complete HHM, covering many times the number of elements presented here, is beyond the scope of this paper. The top boxes in the diagram—termed “head-topics”—integrate lower-level slaughterhouse characteristics that might attract the attention of a terrorist and in turn, should be of concern to the intelligence system. The analyst’s initial concerns are: (1) which slaughter-

houses might be the more likely targets and (2) what avenues of attack are most inviting. The head-topic box in Figure 2—termed *Macro*—lists three subtopic boxes. The first of these is *Capacity*. If we assume that targeting a large facility can result in impacting a larger segment of the population, bigger is better for the terrorist and constitutes a higher risk for the target. The second box is termed *Location*. Assuming that convenience of location matters, then correlating known terrorist locations with slaughterhouse locations could be important. The third box is termed *Ownership*. If the owners/operators of a slaughterhouse have relationships with terrorist organizations, this could be a significant factor in target selection. As a result, tracking the relationships of slaughterhouse owners might be a useful intelligence activity. The fourth box is termed *Customers*. If the slaughterhouse provides products to a kosher food market, for example, this might influence its selection as a target; in turn, keeping track of slaughterhouse customers becomes an item of interest. Other head-topic areas, such as *Security Measures*, *Procedures for Operations*, and *Employment Practices*, could all contribute to the terrorists’ choice of target.

In the HHM food-poisoning scenario, one possibility is to insert a poison into the cleaning fluid used on a meat-grinding machine at a slaughterhouse (see Fig. 2). The poison would not lose its ability to affect the meat even after cooking. Other scenarios would be considered in parallel. For the Bayesian analysis:

- Let T denote that a terrorist scenario is in progress, focused on meat poisoning at a slaughterhouse.
- Let \bar{T} denote that the above scenario is not in progress.
- Let M denote that a supply of machine-cleaning fluid has been reported stolen.
- Let P denote that a supply of a poison that can be applied to meat and stay lethal after cooking has been reported stolen.

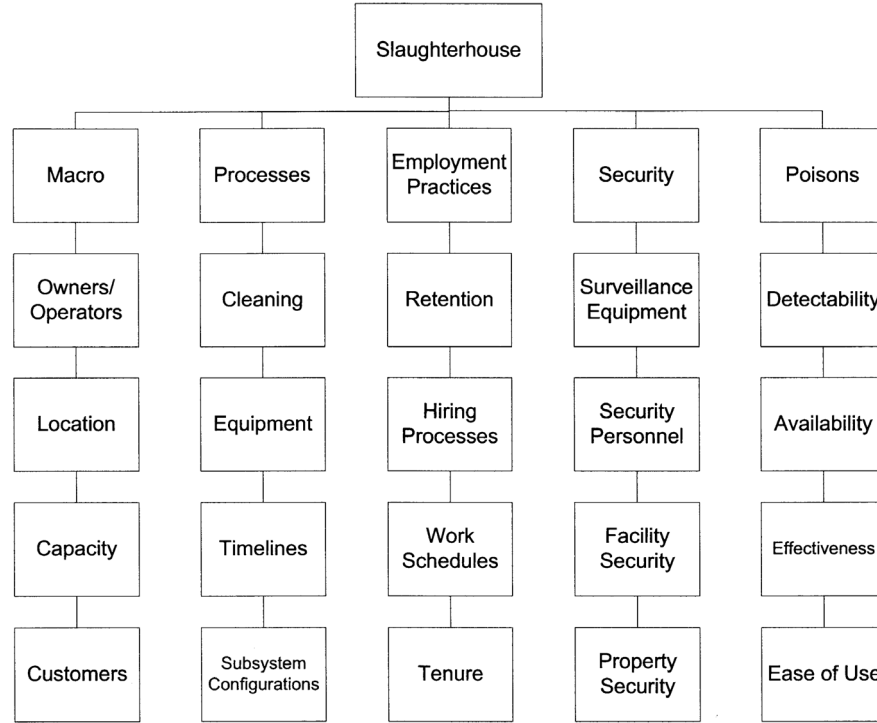


Figure 2. Simplified HHM for Assessing the risks of food poisoning by a terrorist organization.

Let S denote that the same person has been associated with both robberies.

Let A denote that the thief is known to be a member of a terrorist organization.

The following probability values pertain:

$$p(T) = 0.01,$$

$$p(M | T) = 0.01,$$

$$p(M | \bar{T}) = 0.005,$$

$$p(P | T) = 0.3,$$

$$p(P | \bar{T}) = 0.05,$$

$$p(S | T, M, P) = 0.5,$$

$$p(S | \bar{T}, M, P) = 0.1,$$

$$p(A, S | T, M, P) = 0.4,$$

$$p(A, S | \bar{T}, M, P) = 0.005.$$

Correspondingly,

$$p(\bar{A}, S | T, M, P) = 0.1,$$

$$p(\bar{A}, S | \bar{T}, M, P) = 0.095.$$

These probabilities are derived from assessments of the meat-poisoning scenario and its relationship to all other scenarios, whether or not related to terrorism. For example, when viewing all the ways that meat poisoning might be accomplished, as depicted in the HHM analysis, it can be judged that $p(M | T) = 0.01$ since there are many other ways to accomplish the attack without using or stealing cleaning fluid. Establishing a value for $p(M | \bar{T})$ would require gathering information on thefts from national crime data repositories as a basis for judging the probability. Using Bayesian analysis,

$$p(T | M) = \frac{p(T)p(M | T)}{p(M | T)p(T) + p(M | \bar{T})p(\bar{T})}.$$

Substituting the example values yields that $p(T | M) = 0.02$. The value of knowing that machine-cleaning fluid was stolen doubles the initial likelihood of a meat-poisoning scenario.

Continuing with the Bayesian analysis,

$$p(T | M, P) =$$

$$p(T | M) \left[\frac{p(P | T)}{p(P | T)p(T | M) + p(P | \bar{T})p(\bar{T} | M)} \right].$$

This equation assumes $p(P | T, M) = p(P | T)$ that since all of the scenarios for food poisoning include gaining access to a poison. Also, it is reasoned that it would be quite unlikely that the use of cleaning fluid would significantly impact the probability of stealing a poison, given that the terrorist activity is underway. Similarly, it is assumed that $p(P | \bar{T}, M) = p(P | \bar{T})$ since we assume that the thefts of poison and machine-cleaning fluid are independent situations outside of the meat-poisoning scenario. Substituting example values yields $p(T | M, P) = 0.11$. This is a factor ten times greater than the likelihood without the information, and five times greater than when only having knowledge of the stolen cleaning fluid.

Similarly, it can be shown that $p(T | M, P, S) = 0.38$ and $p(T | M, P, S, A) = 0.91$ providing factors of over three and eight of increased probability compared to knowing only about the thefts occurring. It is important to note that the information called M , P , and S may be collected by local law officials and the information called A may be collected at the federal level. In such a case, the opportunity to detect the example scenario is reduced by more than 50% if the federal and local information is not integrated (due to information security reasons, for example).

Now suppose that the information about the thief being a member of a terrorist organization is 80% likely to be true and 20% likely to be false. If \tilde{A} denotes the measurement of A , then

$$p(\tilde{A} = A | A) = 0.8$$

$$p(\tilde{A} = A | \bar{A}) = 0.2,$$

where \bar{A} denotes the thief is NOT a member of a terrorist organization.

Then,

$$p(T | M, P, S, \tilde{A}) = p(T | M, P, S, A)p(\tilde{A} = A | A) +$$

$$p(T | M, P, S, \bar{A})p(\tilde{A} = A | \bar{A}).$$

Using Bayesian analysis, we can compute that

$$p(T | M, P, S, \tilde{A}) = 0.11(0.2) + 0.91(0.8),$$

$$p(T | M, P, S, \tilde{A}) = 0.75.$$

This result is nearly a 19% reduction compared to having perfect information. Now suppose that a decision tree for the meat-poisoning scenario is structured thus: Given the various pieces of evidence, if the probability of a terrorist attack exceeds 0.80, the federal government will provide special protection to large slaughterhouses, and below 0.80 they will provide less.

The example illustrates how the quality of data can impact decision-making and must be accounted for in developing a decision tree. A well-designed decision tree will be based on a sensitivity analysis relating decisions to the range of probabilities that the detected scenario is underway. The greater the difference in the cost and risk between alternative decisions, the greater the difference in detection probability that should separate the particular choices.

Based on the results of the Bayesian analysis presented above and on the HHM model in Figure 2, assume that a decision is made to collect information related to current employees at slaughterhouses. Suppose that this reveals that a known terrorist sympathizer had been recently hired by a large slaughterhouse near a big city. This new information helps to identify the potential point of attack as well as the possible timing of the attack. This would mandate decisions on meat inspection for the slaughterhouse and city in question, as well as directed efforts to terminate the potential attack. This part of the example illustrates that decisions related to intelligence gathering are dependent on the type of observables, the cost of collection, and the risk of collection. Thus, the probability of attack serves a useful mechanism for guiding such decisions.

Another aspect of scenario tracking is working with alternative scenarios that are being assessed in parallel. As in the case of selecting widely spaced decisions for a single scenario, the analysis must gauge the relationship between the costs and risks of decisions, the consequences of the alternate scenarios, and the probabilities that each scenario is indeed in progress. Depending on the specific values yielded by the analysis, the resulting actions might relate to one or more scenarios at the same time.

A number of other important conclusions can be drawn from the example. First, intelligence collection ranges from information about terrorist relationships that might be best done at the federal and state levels, to information about the slaughterhouse and its processes that might best be done at a local level. Of course, the information must eventually be integrated in order to provide maximum value. Second, even from this example, it is clear that a filtering scheme is needed to make collection practical. This is where RFRM comes into play.

When dealing with risk analysis, the issue of probability of occurrence must be addressed. Assessing the probability and the consequences of the attack serves as critical input for guiding decisions on intelligence collection. Since we have no experience to draw upon, for analyzing risk of extreme and catastrophic events the PMRM must use information that is related to the probability of occurrence as a substitute for statistical

data. This information could be based on such factors as the terrorist organizations' skills, funding, and goals. Information on these factors is potentially available through intelligence collections, best done at the federal level. In addition, an evaluation could be made of a terrorist's likelihood of getting caught, either before or after a potential attack. Information on this can be derived, for example, from the HHM mode of Figure 2 and an analysis of the subtopics under the head topics *Security* and *Macro*. These head topics include subtopics that give an indication of the surveillance systems and possible escape routes at the slaughterhouse. Information related to these subtopics is best done at the local level.

Using the HHM and related consequences and likelihoods, the RFRM process, through filtering and ranking, identifies those risk management actions that address the most serious problems within cost, schedule, and other resource limitations. For example, requiring certain controls over employees may be an effective way to prevent attacks that depend on internal employee cooperation. Note that the HHM of Figure 2 includes an *Employee* head-topic with several subtopics that relate to risk that can be reduced, if controlled.

If we explore the time relationships for several of the subjects addressed above, the sequencing presented in Table 1—intent, target acquisition, plan, preparation, execution, and grace period—is illuminated as follows: Historical relationships between owners/operators or employees and terrorist organizations can occur well before an actual attack. In addition, the terrorist may take certain actions to infiltrate the work force before the attack. Purchasing poison and inserting it in a cleanser can occur well before the cleanser is used, depending on slaughterhouse procedures. Once the poison is applied, there is a grace period between application and arrival at the dinner table. Part of risk management can include added inspections to address this possibility. These would need to be regulated by the government, most likely at the federal level, with results of intelligence analyses utilized as the basis for increasing inspections when other indicators warrant enhanced protection measures. The Bayesian analysis and the multiobjective decision tree analysis described in Section 3 would serve to drive decisions regarding increased inspection or specific situations. Clearly it is impractical to demonstrate in this example the full deployment of MODT and the other methods discussed in this paper.

In sum, the food-poisoning example illustrates how we might deal with one scenario. Similar analyses would need to be performed on numerous other potential scenarios, which in turn would be compared for decisions to be made about the levels of risk manage-

ment needed. The example also illustrates the degree of integration that would be required to combine federal, state, and local data sources. The following section addresses the factors that influence solving the integration problem, and corresponding approaches to achieve this.

5. INTEGRATING THE METHODOLOGY WITH EXISTING INTELLIGENCE-COLLECTION SYSTEMS

For many years, the intelligence community has been supporting research and development efforts to create automation tools to help collect and analyze the large quantities of data that emerge from intelligence-collection efforts. One of the more recent and significant efforts is the Genoa Program sponsored by the Defense Advanced Research Projects Agency (DARPA). The tools deal with what to collect (e.g., which websites to monitor), how to combine data from different sources, how to present outputs for easier interpretation by analysts, how to permit teams of analysts working on separate data sources to fuse their results, how to change conditionally the data to be collected, how to adapt tools based on the experience of prior collections, and other aspects. The results of prior efforts have been applied selectively, and as a recent follow-up to the Genoa Program, DARPA has initiated an effort to integrate a significant and diverse set of collection and analysis support tools into a community-wide asset for addressing terrorism. These tools are likely to include scenario-based tracking capabilities as described in earlier sections of this paper.

The key to developing a system for scenario-based tracking of terrorists is to successfully integrate capabilities across local, state, and federal agencies. Historically, the federal government has created its own systems for intelligence collection and analysis, using substantial research and development funding focused on addressing international security risks or weapons of mass destruction. Little effort has been placed on the potential roles of states and localities in defending point targets in our own country. New processes are needed in order to manage and use a new scenario-based tracking system as described above. For example, new policies related to information security will need to be developed. Once a terrorist scenario seems to be in progress, the visibility of information related to the impacted community's information becomes a key factor. This implies a conditional set of security controls. Also, handling tips from the general public will be a sensitive matter. While it would be horrendous to have received and ignored an actual tip, it also would be

impossible to react to all, and in certain circumstances, counterproductive to overreact. A "learning system" is needed that draws on the methodology presented in this paper to start developing processes and management approaches based on partial experience. This might be accomplished by fielding a very limited representative system in the short term, one that includes all elements of an eventual real system, from scenario structuring through intelligence collection and analysis. This system could be tried in a small number of localities and states, dealing with only a few scenarios. Even within these limits, it would bring together federal, state, and local participants, and would illuminate the integration issues that will need to be resolved when the full-scale system is implemented.

In addition to management issues, an approach for computer systems integration is required to tie together scenario-related information from national, state, and local subsystems. The approach for systems integration must recognize the budget constraints of state and local governments. As a result, commercial-off-the-shelf (COTS) components must play a major role in the new system. In developing a commercially based, well-integrated, widely distributed intelligence system, there are significant technical problems. These are related to the issues of (a) combining information from diverse and separately designed databases and (b) maintaining security while controlling access to different data sets. In this regard, the internet and the move toward business-to-business e-business systems has stimulated the commercial development of certain standards as well as continuing research into more advanced standards for dealing with both of these problems. XML (Extensible Markup Language), which deals with managing data definitions and boundaries, RDF (Resource Description Framework), which deals with the relationships between different data sources, and SAML (Security Assertion Markup Language), are all commercially-driven standards for dealing with data and security that have strong potential for application in a distributed intelligence collection and analysis system. To benefit fully from the methodology presented here, more research is required on using this commercial technology to deal with the potential observables tied to various scenarios. Such research could develop rapid prototype capabilities for converting scenario results into corresponding computer inputs for guiding data integration. It would also deal with organizing the security arrangements that must go along with a broad system.

Another important element in designing such a system successfully is the computer logic that adjusts collection based on actual results, allowing for analyst-directed as well as automated adjustment. The scenario structuring methodology described in Section 3 can be

designed and should be used to serve as the basis for integrating collection management and analysis. The Bayesian analysis capabilities can drive the fusion requirements as well as serve as the method for assessing the correctness of prior collections. Bayesian analysis also can adjust correspondingly the sensitivity and specificity values used as the basis for evaluating the quality of various data sources. The scenarios derived via HHM can be converted to computer models that set the parameters that control the performance of individual collection tools. Additional work and study are needed to determine how to understand both the design details of the collection tools and the possible scenarios that could be derived via HHM. In addition, the HHM scenarios can be designed so that an automated system can integrate analyst inputs with those obtained from the domain experts involved in deriving the model. Field experience would thus be fully accounted for in the collection strategy. If the collection and analysis tools are designed on the basis of the HHM scenario structuring and Bayesian analysis tools, a consistent design will run from scenario structuring to implementation and operational use. This should create a coherent process for evaluating the quality of the system and making changes that are consistent throughout a complex, multi-participant system. It would make conducting system management (configuration, evolution, and quality assurance management, among others) a more practical activity across the anticipated large number of participants (operators, designers, analysts, modelers, and others).

6. CONCLUSIONS

The risk-based methodology for scenario tracking for terrorism introduced in this paper builds on the premise that intelligence gathering and analysis for combating terrorism constitutes a complex process. This process may be characterized as a large-scale system of systems with numerous components: dynamic and nonlinear; spatially distributed; involving multiple government and nongovernment agencies, agents, and decision-makers; agencies with different missions, resources, timetables, agendas, and cultures; and multiple constituencies. Risk of extreme and catastrophic events are of paramount importance, organizational and human errors/failures are common, and the process is fraught with multiple conflicting and competing objectives.

Clearly, no silver-bullet approach can address this complexity; neither can a single model do justice to the inherent difficulties associated with the intelligence process. Furthermore, no single methodology, including this one, can be expected to provide a unified and

comprehensive scientific basis for intelligence gathering, analysis, and decision-making. The scope and objective of this paper is to stimulate further discussion that would advance our nation's scientific-analytical capabilities, and thus contribute an important building block to the overall gigantic modeling, analysis, and decision-making efforts facing the country's myriad intelligence agencies. Cognizant of this reality, the methodology presented here builds in part on available methodologies and tools. The term *methodology* connotes here a flexible, adaptive, and repeatable process that is grounded on theoretically sound and tested methods.

Applying Bayesian analysis to the problem of terrorist scenario tracking supports determining how information collection should be distributed across organizational boundaries, as well as understanding the consequences of reducing collection due to organizational considerations. In addition, it provides a quantitative basis for helping the intelligence community understand how improvements in the quality of individual data items relate to overall system performance, based on specific scenarios of concern.

The complexity of the intelligence process calls for iterative learning, unlearning, and relearning [Toffler, 1991]. Learning activities need to be initiated as soon as possible in order to provide timely feedback to the major efforts underway for designing and implementing major components for the intelligence system to fight terrorism. Training must include technical, process management, and organizational components. It must be based on actual experience and therefore requires the fielding of early capabilities for teaching the user community. These "learning systems" must include measurement capabilities derived from measures of effectiveness established for the intelligence system. Since most operational systems are not designed for training, they frequently do not include the teaching of metrics. As a result, an emphasis on metrics is a critical feature of a "learning system" strategy. Once such a strategy for learning is established, it is likely that our government will find that it must initiate more than one philosophical/methodological approach for addressing scenario tracking and then select and integrate the best solutions based on actual experience.

Disrupting a terrorist attack depends on having information that facilitates identifying and locating those involved in supporting, planning, and carrying out the attack. Such information arises from myriad sources, such as human or instrument surveillance by intelligence or law enforcement agencies, a variety of databases and documents concerning transactions, and tips from a wide range of occasional observers. Given the enormous amount of information available, a method is

needed to cull and analyze only the data relevant to the task, confirm their validity, and eliminate the rest. Scenario structuring and tracking could similarly have a broad impact in the field of law enforcement by helping to prevent criminal activities as well as assisting in the collection and analysis of forensic evidence after the crime. The Bayesian analysis of intelligence, used here for purposes of assessing the reliability or accuracy of information, would be useful more broadly in the entire field of intelligence analysis where information from different sources must be combined and their synergistic value assessed.

ACKNOWLEDGMENTS

The authors thank Jim Lambert, Stan Kaplan, and Irwin Pikus for their assistance and contributions to this paper, Grace Zisk for her editorial assistance, and Della Dirickson for her administrative support. The authors also acknowledge the most valuable comments and suggestions made by Dr. Andrew P. Sage, the Editor-in-Chief of *Systems Engineering*, and by other reviewers.

REFERENCES

- G. Altshuller, *Creativity as an exact science*, Gordon and Breach, New York, 1988.
- G. Altshuller, *And suddenly the inventor appeared, TRIZ, the theory of inventive problem solving*, Technical Innovation Center, Worcester, MA, 1994.
- E. Asbeck and Y.Y. Haimes, The partitioned multiobjective risk method, *Large Scale Syst* 6 (1984), 13–8.
- R.G. Bace, *Intrusion detection*, Macmillan Technical Publishing, New York, 2000, pp. 79–117.
- J. Baron, J.C. Hershey, and H. Kunreuther, Determinants of priority for risk reduction: The role of worry, *Risk Anal* 20(4) (2000), 413–427.
- V. Chankong and Y.Y. Haimes, *Multiobjective decision making: theory and methodology*, Elsevier-North Holland, Amsterdam, 1983.
- V. Chankong, Y.Y. Haimes, H.S. Rosenkranz, and J. Pet-Edwards, The carcinogenicity prediction and battery selection (CPBS) method: A Bayesian approach, *Mutation Res* 153 (1985), 135–166.
- R. Cunningham and C. Stevenson, *Accurately detecting source code of attacks that increase privilege, RAID 2001, LNCS 2212*, Springer-Verlag, Berlin, 2001, pp. 104–116.
- H. Debar and A. Wespi, *Aggregation and correlation of intrusion-detection alerts, RAID 2001, LNCS 2212*, Springer-Verlag, Berlin, 2001, pp. 85–103.
- M. Dombroski, Y.Y. Haimes, J.H. Lambert, K. Schluskel, and M. Sulcoski, Risk-based methodology for support of operations other than war, *Mil Oper Res* 7(1) (2002), 19–38.
- H.I. Frohwein, Y.Y. Haimes, and J.H. Lambert, Risk of extreme events in multiobjective decision trees, Part 2. Rare events, *Risk Anal* 20(1) (2000), 125–134.

- W.J.J. Gordon, *Synergetics, the development of creative capacity*, Macmillan, New York, 1968.
- Y.Y. Haimes, Hierarchical holographic modeling, *IEEE Trans Syst Man Cybernet* 11(9) (1981), 606–617.
- Y.Y. Haimes, *Risk modeling assessment, and management*, Wiley, New York, 1998.
- Y.Y. Haimes and D. Macko, Hierarchical structures in water resources systems management, *IEEE Trans Syst Man Cybernet SMC-3*(4) (1973), 396–402.
- Y.Y. Haimes, P.O. Karlsson, J.M. Mitsiopoulos, and D. Li, “Risk management of extreme events,” Supplement 1. *Systems and control encyclopedia: Theory, technology, applications*, M. Singh (Editor), Pergamon, New York, 1990a, pp. 259–265.
- Y.Y. Haimes, K. Tarvainen, T. Shima, and J. Thadathil, Hierarchical multiobjective analysis of large-scale systems, Hemisphere, New York, 1990b.
- Y.Y. Haimes, D. Li, and V. Tulsiani, Multiobjective decision tree method, *Risk Anal* 10(1) (1990c), 111–129.
- Y.Y. Haimes, N.C. Matalas, J.H. Lambert, B.A. Jackson, and J.F.R. Fellows, Reducing the vulnerability of water supply systems to attack, *J Infrastruct Syst* 4(4) (1998), 164–177.
- Y.Y. Haimes, S. Kaplan, and J.H. Lambert, Risk filtering, ranking and management framework using hierarchical holographic modeling, *Risk Anal* 22(2) (2002), 383–397.
- A.D. Hall, III, *A methodology for systems engineering*, D. Van Nostrand, New York, 1962.
- A.D. Hall, III, *Metasystems methodology: A new synthesis and unification*, Pergamon, New York, 1989, p. 6.
- A. Josang, Subjective evidential reasoning, *Proc 9th Int Conf Inf Process Manage Uncertainty Knowledge-Based Syst (IPMU 2002)*, Annecy, France, July 1–5, 2002.
- S. Kaplan, On inclusion of precursor and near miss events in quantitative risk assessments: A Bayesian point of view and a space shuttle example, *J Reliab Eng Syst Safety* 27 (1990), 103–115.
- S. Kaplan, “Expert information” vs. “expert opinion;” another approach to the problem of eliciting/combining/using expert knowledge in PRA, *J Reliab Eng Syst Safety* 35 (1992), 61–72.
- S. Kaplan, An introduction to TRIZ: The Russian theory of inventive problem solving, Ideation International, Southfield, MI, 1996.
- S. Kaplan and B.J. Garrick, On the quantitative definition of risk, *Risk Anal* 1(1) (1981), 11–27.
- S. Kaplan, S. Vishnepolschi, B. Zlotin, and A. Zusman, New tool for failure and risk analysis, anticipatory failure determination (AFD) and the theory of scenario structuring, Ideation International, Southfield, MI, 1999.
- S. Kaplan, Y.Y. Haimes, and B.J. Garrick, Fitting hierarchical holographic modeling (HHM) into the theory of scenario structuring and a refinement to the quantitative definition of risk. *Risk Anal* 21(5) (2001), 807–819.
- A. Kott, M. Pollack, and B. Krogh, The situation assessment problem: Toward a research agenda, *AEC Proc DARPA-JFACC Symp Adv Enterprise Control*, San Diego, November 15–16, 1999, pp. 251–258.
- J.H. Lambert, Y.Y. Haimes, D. Li, R. Schooff, and V. Tulsiani, Identification, ranking, and management of risks in a major system acquisition, *Reliab Eng Syst Safety* 72(3) (2001), 315–325.
- W.W. Lowrance, *Of acceptable risk*, William Kaufmann, Los Altos, CA, 1976.
- D. Macko, and Y.Y. Haimes, Overlapping coordination of hierarchical structures, *IEEE Trans Syst Man Cybernet SMC-8*(10) (1978), 745–751.
- J. Mitsiopoulos, Y.Y. Haimes and D. Li, Approximating catastrophic risk through statistics of extremes, *Water Resources Res* 27(6) (1991), 1223–1230.
- M.G. Morgan, B. Fischhoff, L. Lave, and P. Fischbeck, “A proposal for risk ranking within federal agencies,” Comparing environmental risks: Tools for setting government priorities, J. Clarence Davies (Editor), Resources for the Future, Washington, DC, 1999.
- M.G. Morgan, H.K. Florig, M.L. DeKay, and P. Fischbeck, Categorizing risks for risk ranking, *Risk Anal* 20(1) (2000), 49.
- National Research Council (NRC) of the National Academies, *Making the nation safer: the role of science and technology in countering terrorism*, Washington, DC, 2002.
- J. Pet-Edwards, H.S. Rosenkranz, V. Chankong, and Y.Y. Haimes, Cluster analysis in predicting the carcinogenicity of chemicals using short-term assays, *Mutation Res* 153 (1985a), 167–185.
- J. Pet-Edwards, V. Chankong, H.S. Rosenkranz, and Y.Y. Haimes, Application of the carcinogenicity prediction and battery selection (CPBS) method to the Gene-Tox data base, *Mutation Res* 153 (1985b), 187–200.
- J. Pouzol and M. Ducassé, *From declarative signatures to misuse IDS*, RAID 2001, LNCS 2212, Springer-Verlag, Berlin, 2001, pp. 1–21.
- H. Raiffa, *Decision analysis: Introductory lectures on choice under uncertainty*, Addison-Wesley, Reading, MA, 1968.
- A.P. Sage, *Methodology for large scale systems*, McGraw-Hill, New York, 1977.
- R.R. Sokal, Classification: Purposes, principles, progress, prospects, *Science*, September 27, 1974.
- A. Toffler, *The third wave: the classic study of tomorrow*, Bantam Books, New York, 1991.
- A. Valdes and K. Skinner, Probabilistic alert correlation, RAID 2001, LNCS 2212, Springer-Verlag, Berlin, 2001, pp. 54–68.
- G. Vigna, R. Kemmerer, and R. Blix, Designing a web of highly-configurable intrusion detection sensors, RAID 2001, LNCS 2212, Springer-Verlag, Berlin, 2001, pp. 69–84.
- J. von Neumann and O. Morgenstern, *Theory of games and economic behavior*, 3rd edition, Princeton University Press, Princeton, NJ, 1953.
- J.N. Warfield, *Social systems-planning and complexity*, Wiley, New York, 1976.
- T. Webler, H. Rakel, O. Renn, and B. Johnson, Eliciting and classifying concerns: A methodological critique, *Risk Anal* 15(3) (1995), rs.: 421.



Barry M. Horowitz, Professor of Systems and Information Engineering at the University of Virginia, received an MSEE and a PhD from New York University in 1967 and 1969 respectively, and a BEE from the City College of New York in 1965. Dr. Horowitz joined the University of Virginia's faculty as a Professor in the Systems and Information Engineering Department in September 2001, after an industrial career involving the application of systems engineering to many large and complex systems. From 1969 through 1996 he was employed in a variety of positions at the Mitre Corporation, including the last five years as President and CEO and the three prior years as Executive Vice President and COO. During his time at Mitre he played major roles in the company's military, intelligence, and civil aviation sectors. He received the Air Force's highest award for a civilian as a result of this effort. In 1995, he authored a book entitled *Strategic Buying for the Future* that highlighted significant problems in the development of large military systems and corresponding approaches to solving these problems. In 1996, he founded Concept Five Technologies, an e-business systems development company focused on the creation and application of standards-based frameworks for the secure integration of large business-to-business e-business systems. As a result of his efforts, in 1996 Dr Horowitz was elected into the National Academy of Engineering. He is also a member of the Tau Beta Pi and Eta Kappa Nu honor societies.



Yacov Y. Haimes, Professor of Systems and Information Engineering at the University of Virginia, is the Founding Director (1987) of the University of Virginia's Center for Risk Management of Engineering Systems, and holds the Lawrence R. Quarles Professorship in the School of Engineering and Applied Science. He is a member of the Systems and Information Engineering and Civil Engineering faculties. On the faculty of Case Western Reserve University for 17 years, he was Chair of the Systems Engineering Department. During the 1977–1978 sabbatical year, he was an AAAS/AGU Congressional Science Fellow, joining the staff of the Executive Office of President Carter, and later the staff of the House Science and Technology Committee. He is the recipient of several major awards in his field, including the Norbert Weiner Award from IEEE Systems, Man & Cybernetics, the Distinguished Achievement Award from the Society for Risk Analysis, the Georg Cantor Award from the International Society on Multiple Criteria Decision Making, and the Warren A. Hall Medal from the Universities Council on Water Resources. He is a Fellow of the following professional societies: AAAS, IEEE, ASCE, IWRA, AWRA, INCOSE, and the Society for Risk Analysis (SRA). He also served as President of SRA. He has published more than 200 articles and technical papers, over 120 of which are in archival journals. He has authored/co-authored five books and edited 20 volumes. His most recent book is *Risk Modeling, Assessment, and Management*, John Wiley & Sons, 1998. He has worked extensively with the military and domestic agencies on boards and projects.