An Application of Bayesian Networks to Antiterrorism Risk Management for Military Planners

Linwood D. Hudson
Digital Sandbox, Inc.
11710 Plaza America Dr.,
Suite 2000
Reston, VA 20190
lhudson@dsbox.com

Bryan S. Ware
Digital Sandbox, Inc.
11710 Plaza America Dr.,
Suite 2000
Reston, VA 20190
bware@dsbox.com

Kathryn Blackmond Laskey
Dept. of Systems
Engineering
and Operations Research
George Mason University
Fairfax, VA 22030-4444
klaskey@gmu.edu

Suzanne M. Mahoney
IET, Inc.
1911 North Fort Myer Dr.,
Suite 600
Arlington, VA 22209
suzanne@iet.com

Abstract

Recent events underscore the need for effective tools for managing the risks posed by terrorists. Assessing the threat of terrorist attack requires combining information from multiple disparate sources, most of which involve intrinsic and irreducible uncertainties. This paper describes Site Profiler® Installation Security Planner, a tool initially built to assist antiterrorism planners at military installations to draw inferences about the risk of terrorist attack. Site Profiler applies knowledge-based Bayesian network construction to allow users to manage a portfolio of hundreds of threat/asset pairs. The constructed networks combine evidence from analytic models, simulations, historical data, and user judgments. Site Profiler was constructed using our generic application development environment that combines a dynamically generated object model, a Bayesian inference engine, a graphical editor for defining the object model, and persistent storage for a knowledge base of Bayesian network fragment objects. Site Profiler's human-computer interaction system is tailored to mathematically unsophisticated users. Future extensions to Site Profiler will use data warehousing to allow analysis and validation of the network's ability to predict the most effective antiterrorism risk management solutions.

1 BACKGROUND

The U.S. military defines antiterrorism as the defensive posture taken against terrorist threats. Antiterrorism includes fostering awareness of potential threats, deterring aggressors, developing security measures, planning for future events, interdicting an event in process, and ultimately mitigating and managing the consequences of an event. These activities are undertaken at the installation or unit level throughout the Department of Defense. One key element of an effective antiterrorist strategy is evaluating individual military bases for terrorist risk. Site Profiler® allows military planners to develop a customized, base-specific risk assessment that combines information from many different sources. Site Profiler Installation Security Planner (ISP) was licensed by the US Department of Defense in 1999 to develop an enterprise-wide anti-terrorism risk management system.

The ISP program was initiated in response to the bombing of US Air Force servicemen in Khobar Towers, Saudi Arabia and the bombings of the US Embassies in Africa. These events and their ensuing investigations revealed that the US had inadequate methods for assessing terrorist risks and planning for future terrorist events. Recent events continue to highlight the difficulties of antiterrorism planning. Site Profiler ISP was completed in August 2001. It provides a means to assess terrorist risks, manage these risks, and develop standardized antiterrorism plans.

The next section describes the uncertainties involved in assessing and managing antiterrorism risks, and the weaknesses in current methodologies. Section 3 describes our modeling and knowledge engineering approach. Section 4 describes the software architecture for Site

Profiler and our generic Bayesian application development environment.

2 THE ANTITERRORISM RISK MANAGEMENT CHALLENGE

Terrorist activity is increasing globally and the targeting of US individuals, at home or abroad, has risen sharply since the early 90's. In particular, US military forces abroad are increasingly targets of terrorist attacks. Terrorists present an asymmetric threat that military training and planning doctrine are not well suited for. Military forces are organized and trained to fight clearly defined enemies in definitive engagements. The terrorist exploits this posture by attacking when least expected, using unconventional means, against a force that is often ill-prepared.

The risk of terrorist attack is always present for US forces abroad and at home. Vulnerabilities will always exist that could be exploited by an enemy. The challenge is to prioritize these risks by identifying exploitable vulnerabilities and the likelihood that these vulnerabilities will be targeted. However these risks involve highly uncertain and subjective assessments to be made of terrorist intent, capabilities, targeting preferences, and other features indicative of the likelihood and severity of a terrorist incident.

Risk management requires that we understand the likelihood of an event and the consequences if that event were to occur so that mitigation efforts can be optimally employed. This assessment requires integrating disparate data sources that are almost impossible for one person to grasp. Information about terrorist intent and targeting preferences, usually the province of intelligence staff, is largely subjective and highly uncertain. Understanding of vulnerabilities and mitigation options, typically areas for physical security specialists, are often based on experience or "best judgment." Estimating the consequences of an attack requires sophisticated models that are only usable by engineers and scientists. The antiterrorism planner at each military installation is responsible for assimilating all of this information for all of his installation assets and managing a dynamic risk portfolio of potentially hundreds of threat-asset pairs. Although a limited number of experts may be able to understand and manage a given risk, no human can manage all of the components of hundreds of risks simultaneously.

2.1 Existing Approaches to Antiterrorism Risk Management

Historically, antiterrorism planners have employed manual procedures derived from conventional military doctrine or Special Forces targeting criteria. These standard operating procedures are documented in military manuals and supported by paper and pencil tools. Existing methods are typically not specific to a given threat or situation, analytically dubious, extremely simplistic, and procedural rather than knowledge based. A brief description of some of the existing methods follows.

2.1.1 DSHARPP

DSHARPP is an acronym that represents a so-called "stubby pencil" process for risk assessment. First, the installation planner develops a list of potential terrorist targets on his installation. Then, for each of the terms in the DSHARPP acronym; Demography, Susceptibility, History, Accessibility, Recognizability, Proximity, Population; he awards from 1 to 5 points for the term. Summing up these scores yields a maximum of 35 points, the worst case, or a minimum of 7 points (Air Force Instruction 31-210). These scores are used to sort the list of potential targets. The installation commander then determines which of these targets he will address with additional security measures, and applies any measures he feels are appropriate. None of these terms or scores is adjusted based on threat, type of target, or any other special considerations. Several other similar approaches can be found in the military literature.

2.1.2 The FPCON System

The FPCON System is a risk management approach that is based on compliance with a set of prescribed standards (Joint Pub 3-07.2, 1998). These standards, referred to as Force Protection Condition Measures, must be implemented at every military installation. Five FPCON levels (Normal, Alpha, Bravo, Charlie, and Delta) represent an increasing level of terrorist threat, as determined by military intelligence. As the FPCON level increases from Normal to Alpha and so on, the installation employs the prescriptive FPCON Measures. Each level involves ten measures like Measure 5, "Limit access points for vehicles and personnel commensurate with a reasonable flow of traffic". These measures were developed by military experts as prudent measures to undertake under rising threat conditions. FPCON Measures represent practical expert judgment, but are not specific to a given threat, situation, or local environment. Though these measures make it easy to develop consistent plans, they have been shown to be inadequate (USS Cole Commission Report, 2001).

2.1.3 Expert Assessment Teams

In addition to the two methods cited above, both of which are required of all military installations, a third process is used—expert assessment. The DoD employs dozens of

teams composed of eight to ten expert assessors. These teams are composed of individual experts in terrorist options, structural engineering, chemical weapons, law enforcement, and disaster response. These experts spend no more than one week at the installation. During that time they interview and observe, pick three to five likely targets, three to five potential threats, and discuss the risks and mitigation options in a final report and briefing to the installation commander. The installation commander is then responsible for addressing these findings.

Since the findings of the expert teams are based solely on their judgment and experience, the analytical integrity of their assessments is subjective and their results are not repeatable. First, each expert and team of experts tends to view the situation differently. Second, the time that they have to conduct their assessment limits them to a small number and depth of considerations. Third, the situations on military installations are constantly changing, but assessment visits occur infrequently. Finally, because their findings are essentially opinions, they can easily be countered by someone else's opinion. Expert assessment is a powerful tool, but only if it can be structured to produce repeatable, high confidence results.

2.2 The Site Profiler Approach to Antiterrorism Risk Management

Site Profiler allows antiterrorism planners to analyze and manage a large portfolio of risks simultaneously. Knowledge about installations, assets, and terrorist risks is encoded in a knowledge base of Bayesian network fragments (Laskey and Mahoney, 1997) that can be dynamically combined at run-time into a Bayesian network for assessing risks specific to a given installation and situation. This is accomplished via an object-oriented database and an architecture that can supply the network nodes with data from disparate sources. These disparate sources include the planner's own subjective and objective assessments, historical database information, analytic model results, and simulation results that are integrated into various nodes on the Bayesian network. This network is dynamically constructed by the software and automatically solved and presented to the user for each combination of asset and threat that the user has described.

Because our users are not analysts, the system was designed to be understandable to users not accustomed to thinking about probability. Analyzing the consequence component of risk required the inclusion of results from physics-based weapon models that planners would be unfamiliar with. Our architecture was thus developed to provide intuitive input and feedback interfaces to the users and to have direct interfaces to multiple simulations and models that could provide data directly to the network without user intervention.

The system was designed to support the thought process of military planners, while maintaining computational efficiency and analytical integrity. The fragments in the risk influence Bayesian network were designed to match the user's domain concepts, ensuring a scalable, modular, and maintainable model. For example, we created a fragment that consisted of all the network nodes that pertained exclusively to the concept of an asset. We then provided interfaces for the user to characterize each of the assets on his installation, as seen in Figure 1. Given a characterization of the assets on an installation and a set of threats to be considered, the software automatically creates and evaluates a Risk Influence Network (RIN) for each asset/threat pair.

Figure 1: The interview-style interface allows the user to describe his assets

The model also included relational information regarding threat/asset pairs. Because the number of risks scales roughly as the number of assets times the number of threats, manual entry of relational information for each threat/asset combination was clearly infeasible. We were able to model the relational aspects so that all relational information could be calculated from a complete characterization of threats and assets. Many of these calculations were computed from simulations external to the Bayesian network. For example, we used a simulation to determine the accessibility of an asset to a threat, and a physics-based model to determine the consequences of a given explosive against a structure or a chemical weapon against a group of people. These, and other calculations, are automatically requested by the architecture and used to populate Bayesian nodes when the RIN is solved.

Because our users are not analysts, they needed a tool they could learn simply, use effectively, and trust. In particular, we wanted to avoid a "black box" into which the user feeds information and out of which an answer magically appears. Users invoke the model via a natural

and understandable interface to describe their assets, specify characteristics of their installation, and select threats to consider. The system constructs RINs for each threat/asset combination, runs offline simulations and database queries as needed, applies evidence, and computes risks which are presented back to the user in tables formatted for understandability. At this high level view, a utility function is used to reduce the probability distributions to single indicators like High, Medium, or Low, as shown in Figure 2. Users are then able to "drill-down" into the components of the risk by clicking on rows in the risk table and walking down the Bayesian network. This ultimately takes them to leaf nodes at which the information may have come directly from a question they answered or from the results of a model calculation. We present the user with graphical views of his risk and with the probability distributions of each node. They can then adjust inputs as necessary, but can feel confident that they understand the underlying components of a given risk score.

Threat	Target	Plausible	Attractive	Susceptible	Consequences	Risk	Conf	Vectors	Model
Panel Truck Bomb	Operations	High	Moderate	High	Critical	1	■	↘	◉
Panel Truck Bomb	Operations	High	Moderate	High	Critical	2	■	↘	◉
Vehicle Bomb	Operations	High	Moderate	High	Critical	3	■		
Hoaxes	Operations	Very High	Moderate	High	Moderate	4	■		
Panel Truck Bomb	Dignitaries	High	Moderate	Moderate	Critical	5	■		
Vehicle Bomb	Dignitaries	High	Moderate	Moderate	Critical	6	■		
Panel Truck Bomb	Warehouse	High	Low	Low	Critical	7	■		
Vehicle Bomb	Warehouse	High	Low	Very Low	Moderate	8	■		

Figure 2: The Risk Table allows the user to view and sort by key network nodes

The structure of the network and domain fragments facilitates the risk management process. Users can easily see the threat that is most plausible, or the asset that has high consequences, or a common element among many risk scenarios. Countermeasures, procedures, and other adjustments can be applied to the installation baseline to address issues identified in the risk influence network.

3 THE DEVELOPMENT PROCESS

Our initial research into the domain of antiterrorism risk management identified a broad consensus among both experts and policy makers that a new approach to antiterrorism risk management was needed. The new approach needed to be more sophisticated, analytically defensible, customizable, and easily modifiable than the approaches in current use. Our goal was therefore not to improve a current method, but to develop a truly revolutionary approach.

3.1 Knowledge Representation

Because of the diverse types of information that we needed to collect and the complex interrelations between the many factors that affect risk, we immediately realized

that a straightforward algorithmic approach would be inadequate. An approach was needed that allowed disparate types of data to be combined in a coherent, analytically defensible, and understandable manner.

Our research identified uncertainty as fundamental. There is uncertainty in the identities of the terrorists, uncertainty in their capabilities, uncertainty in what makes an asset attractive, uncertainty in the most likely methods of attack, uncertainty in the consequences of an attack, and uncertainty in how these factors combine to affect risk. After analyzing several approaches to reasoning under uncertainty, we determined that Bayesian networks would provide the capabilities necessary for Site Profiler. A Bayesian network could be used to model the components that affect risk and how they interact. Because of the need to represent and combine repeatable sub-structures, it became clear that an object-oriented representation (Koller and Pfeffer, 1997; Laskey and Mahoney, 1997) combined with knowledge-Based model construction (Wellman, et al, 1992; Mahoney and Laskey, 1998) was necessary. We found that experts had little difficulty understanding and suggesting improvements to Bayesian network fragments we presented to them. We could even present the constructed networks to users to give them a clearer picture of the factors influencing risk. The Bayesian network representation allowed us to combine evidence from disparate sources, such as from users, historical databases, simulation, and analytic models.

3.2 Engineering the Network

The heart of the Site Profiler risk methodology is the Risk Influence Network (RIN). The RIN is a 146-node Bayesian network that solves for the relative risk of an attack against a particular asset by a particular threat. The nodes of the RIN contain information about the installation as a whole, the asset, the threat (tactic, weapon system, and terrorist organization), the asset-threat target pairing, and the attack event.

Following a network engineering process (Mahoney and Laskey, 1996), we iteratively moved from initial concepts and definitions to a set of reusable network fragments that could be combined into an asset-threat specific RIN. The effort proceeded in six stages: initial concept, formal definition and analysis, subsection review with experts, scenario elicitation and revision, implementation, and operational revisions.

3.2.1 Initial Concept

We began by collecting all of the various pieces of data needed to support the system. This data fell cleanly into two distinct categories: physical data and domain data. Physical data includes information necessary to describe the state of a physical object, such as position, size, shape,

and weight. Domain data represents the more abstract concepts that do not necessarily relate to any physical structure, such as attractiveness, risk, and plausibility. When combined, these two types of data form a complete model of the terrorist realm.

Once we had gathered all of the necessary data for the system, we grouped it together in the form of data objects. Each object represents a particular collection of data that defines a concept, such as a car, building, asset, or threat. In Site Profiler, there are seven objects that are used to construct the RIN: installation, asset, threat, weapon system, terrorist organization, target, and attack. Our core knowledge representation consists of a set of Bayesian network fragments expressing information about attributes of and relationships among these objects.

Working with a combination of existing documents and experts, we drew the initial graph for the RIN. Nodes in the network included both evidence nodes and measures of aspects of the risk. We drew the arcs in an inferential direction from evidence to inferred measure, and then developed initial definitions for the nodes. Because the RIN was a new concept to our experts, the precise definitions for many of the nodes (e.g. Accessibility, Recoverability) remained unclear.

3.2.2   Formal definition and analysis

When we formally defined the nodes and their states, we required examples for each state. These examples helped us decide how many states were appropriate for a measurement node and made it much easier to communicate the concepts to the experts who later reviewed the network. Concurrently, we identified inferentially interesting network fragments of five to a dozen nodes, revised their structure and populated their conditional probability tables with "rough guess" values based on information we had obtained from domain experts and literature.

3.2.3   Subsection review with experts

We reviewed subsections of the RIN with three different groups: threat experts, damage experts, and accessibility experts. Each review took two days. Most of the effort was spent communicating and revising the terminology. We used Netica™ to display the fragments, one fragment at a time. Rather than explicitly asking for probability distributions we elicited relative strengths of influence, entered appropriate distributions and displayed the inferential results to the experts for their feedback. We found that this process of developing an initial model based on our understanding of the domain and obtaining review and feedback from experts was an efficient and effective approach for rapid knowledge engineering. There is a common view in the literature that elicitation of

structure is relatively straightforward relative to the difficult problem of eliciting probabilities (Druzdzel and van der Gaag, 2000). In our experience, the most difficult and time consuming part of the process was establishing a common understanding of terminology and definitions. We circumvented the difficult issue of directly assessing probabilities in favor of an approach of developing an initial model based on a review of the literature, reviewing the model with experts, and asking experts for relative strengths of inference rather than for probabilities. Based on results of our evaluations, this approach was successful.

3.2.4   Scenario elicitation and revision

In additional sessions we elicited scenarios from a cross-section of experts and entered the data into the RIN. While these scenarios showed that the RIN 'worked', they tended to be exceptional (e.g. attacks against the Pentagon).

3.2.5   Implementation

At this point, we implemented the RIN in software. This involved importing the structure of the network into the software architecture and running scenarios to ensure that the integration between the network and the rest of the system was successful. From this point on, all testing regarding the RIN occurred in conjunction with the entire Site Profiler system.

3.2.6   Operational revisions

The initial validation of the RIN by experts showed us that the network basically works in the sense that it appears to order asset-threat pairs sensibly for the small set of scenarios we evaluated during initial validation of the model. However, there is as yet insufficient data to provide a definitive evaluation of the quality of the solutions the system provides. As the system is fielded, data will be generated from military sites around the globe. We plan to use this data to validate and calibrate the RIN.

Questions that need to be answered are: Does the RIN provide enough separation among a commander's asset-threat pairs, that the commander can make a decision about spending his force protection budget? For a given security expert at a given site are the RIN results repeatable? Would two security experts using Site Profiler produce the same ordering of asset-threat pairs? Are their measures essentially the same? Can relative risk measures for one site be compared with those for another?

4 SOFTWARE IMPLEMENTATION

To meet the many needs of Site Profiler, we developed an object-oriented database architecture with native support for Bayesian networks. In our architecture, an object contains a set of attribute types, with one of these types being a Bayesian network value. This value represents the current probability distribution of a node in the network. These types of attributes reside on domain objects that are considered during the evaluation of the network. For instance, the accessibility of an Asset is considered as an influencer in the RIN, so the Asset domain object contains Bayesian attributes.

4.1 Bayesian Attributes and Objects

Bayesian attributes are attributes on domain objects that store the belief values for nodes in the network. These values represent either evidence entered into the network, or propagated belief generated by queries against the network. For each Bayesian attribute, there is associated with it a Bayesian object. A Bayesian object contains the data necessary to represent a network node, such as the states of the node, the probability distribution, and the parents of the node.

Bayesian attributes and objects work together to empower the Site Profiler RIN. Bayesian objects are the building blocks of the RIN, in that they define the structure and behavior of the network. Although designed to be generic, our Bayesian objects are optimized for use with IET's Universal Bayesian Network Solutions Engine (UBNSE), which is the Bayesian inference engine used by Site Profiler. Bayesian attributes provide a snapshot in time of the network, and they allow for the exposure of the state of the network to the users.

When a domain object containing Bayesian attributes comes into existence, the Bayesian objects associated with the attributes are also created. These collections of Bayesian nodes, or network fragments, remain with the object during its lifecycle, and are applied to the RIN as a group. When one domain becomes associated with another, such as when an Asset and Threat form a Target, the fragments also associate with one another, based on the parent/child relationships identified by the network structure, to construct an instance of the RIN as show in Figure 3.


Figure 3 - Fragments associated with domain objects combine to form the RIN

4.2 RIN Structure

Site Profiler domain objects combine in fundamental relationships to describe risk. Assets and Threats combine to form Targets, and Targets can be addressed with Countermeasures that can act on the Asset, the Threat, or both. When Targets are created from Threat-Asset pairs, an instance of the RIN is created. The RIN is composed of network fragments from the Asset, the Threat, and other domain objects.

For Assets, these fragments contain influencers, or nodes, that seek to describe how *critical* the Asset is to your mission, how *desirable* it is to an enemy, and how soft or *accessible* the Asset is. For Threats, the fragment describes how *plausible the tactic and weapon* are, the likely *intent* of an actor to target you and the asset types he's most likely to target. These risk elements combine to contribute to the key risk nodes associated with a Target – *Likelihood of Event*, *Susceptibility* of an Asset to the event, the *Consequences* of the event, and ultimately, the *Risk* of the event. The influence nodes and the fragments themselves represent the critical elements of risk for each threat-asset pair, and are attributes of the domain objects. Countermeasures, in their most raw form, essentially counter any of the positive influencers of risk.

4.3 Dynamic Object Model and User Interface

Site Profiler needed the ability for software administrators and maintainers to modify the interface contents and object model without changing source code. This was intended to provide customers with the flexibility to tailor the system to meet their needs after software delivery. The ramifications of this were that the data used by the system could not be defined in application code itself, but instead had to be accessible outside of the system.

To address this challenge, we designed a database that allows us to store the structure of our object model, complete with all of its necessary data and associations.

We also structured our user interface components as objects and stored them in the database. We then designed our software to interpret the database in order to construct the object model and user interface dynamically.

Using a graphical editor, the user interface and object model, including the Bayesian network nodes, can be modified by software administrators and maintainers. Interface screens can be added or modified, and then linked up to attributes of the object model. These screens can then be used to apply data to the model, or to present data to the users. This rapid application creation capability provides immense flexibility and scalability.

4.4 Evidence from Other Modules

Along with the database, RIN, and user interface modules of Site Profiler, we developed a 3D modeling environment for building a site in 3D, an "intelligent terrorist" module that attempts to infiltrate the site in order to identify physical vulnerabilities, and analytic models for simulating weapons effects. These three modules provide evidence that can improve the user's understanding of their risk. Integrating these results into the RIN was another requirement of Site Profiler.

Evidence in the Site Profiler architecture is supplied through the Plug-in Interface, which is an application programmer interface (API) for accessing various data sources. As shown in Figure 4, this allows the RIN to fuse information from the graphical user interface, models and simulations, an historical database, a corporate information system, or a real-time information source. This interface allows the RIN to consider new and existing evidence sources for evaluating risk contributors.
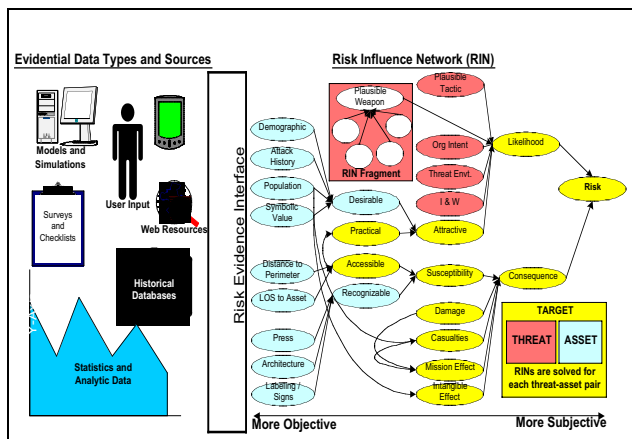

Figure 4 - The RIN uses evidence from a wide range of sources to evaluate risks

4.5 Confidence Measures

In order to differentiate between the various types of RIN evidence, we developed a confidence model that recognizes the difference between subjective user evidence and objective analytical/historical evidence. Using credibility nodes, we apply softer evidence to the RIN when the data is subjective in nature. When applying analytic data, however, we only soften the evidence if the analytic model itself is less credible than other models. This not only allows for greater confidence in analytic versus user evidence, but also for recognition of the levels of fidelity in analytic models. This approach also works for our historical data, in that we vary the credibility of the data depending upon the reliability of its source.

4.6 Support Tools

In addition to the graphical editor mentioned above, we developed several other tools for use in creating, maintaining, and analyzing the RIN. During the data elicitation and knowledge engineering phases of the project, we used the Netica modeling tool for constructing and refining the RIN. Once the RIN had been completed using Netica, we developed a tool for importing the Netica file format directly into our object model. This capability proves to be invaluable because it allows us to model and make changes to the network in an environment specifically suited for that task, and to then import the results into our model. Additionally, changes based on expert feedback can be easily integrated into the RIN.

To facilitate the inclusion of historical data into the RIN, we devised a method of importing default evidence values from an Excel spreadsheet into our object model. This allows us to compile and manage the data externally from the system, and to then quickly integrate it into our model without having to modify the objects individually using the editor. Given that the information that we based our historical data on tend to fluctuate with new findings and reports, this tool proves very handy.

We then developed two other tools for use in analyzing the RIN values and structure. The first tool provides an export capability from the RIN to a color-coded Excel worksheet. For each instance of the RIN inside of our object model, we export the node names, state names, and probability distribution to a worksheet page. This provides a snapshot in time of all nodes of the RIN for analysis.

Figure 5 - The RIN Viewer allows users to mine through the network in detail

The second tool, shown in Figure 5, is referred to as the RIN Viewer. The RIN Viewer is an embeddable user interface element that graphically represents the structure and state of the nodes of the RIN. Each node is color coded to distinguish the fragment type that it belongs to, and contains an icon to distinguish the type of the node (user input, database lookup, propagated, or analytically set) and the current state of the node. The level at which the RIN Viewer mines down through the network is adjustable, and can be set to display either a small fragment of information or the entire network.

5   DISCUSSION

Site Profiler provides a key element of an overall strategy for antiterrorist risk management. A knowledge based Bayesian network construction module forms a key component of a decision support system for assessing terrorist threats against military installations. Site Profiler provides a complete picture of risk to military planners in a way that is coherent, repeatable, and efficient. It allows planners to take knowledge that they already have and augment it with information that was previously unavailable or inaccessible to them. By embedding the RIN inside of Site Profiler's intuitive interface, we allow planners to perform complex data analysis without requiring them to be experts in antiterrorism risk management.

The preliminary validation of the Site Profiler RIN is encouraging. Current plans for data capture and warehousing will facilitate the long-term validation efforts needed to ensure that the system is successful. Site Profiler's generic environment for developing Bayesian applications has already proven to be extremely useful. We recently used it to develop a separate application for assessing the risks to seaports from drug trafficking, and other applications are on the horizon. Our generic environment provides the ability to rapidly develop and deploy decision support systems employing knowledge based Bayesian network constructions across a wide range of application domains.

Dedication

This paper is dedicated to the memory of journalist Danny Pearl, brutally murdered in Pakistan in February 2002, and to the pioneering research of his father Judea Pearl, inventor of the Bayesian network representation language and computational architecture. Danny Pearl's spirit will live on in the work of those who apply his father's research to protecting the open society for which he gave his life.

References

Air Force Instruction 31-210, The Air Force Antiterrorism/Force Protection (AT/FP) Program.

Druzdzel, M. and van der Gaag, L. Building Probabilistic Networks: Where do the Numbers Come From – A Guide to the Literature, Guest Editors' Introduction, *IEEE Transactions in Knowledge and Data Engineering*, vol. 12, pp. 481 – 486, 2000.

Joint Pub 3-07.2; Joint Tactics, Techniques, and Procedures for Antiterrorism; March, 1998

Koller, D. and A. Pfeffer (1997) Object-Oriented Bayesian Networks In Geiger, D. and Shenoy, P. (eds) *Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference*, San Francisco, CA: Morgan Kaufmann.

Laskey, K. B. and S. M. Mahoney (1997) Network Fragments: Representing Knowledge for Constructing Probabilistic Models. In Geiger, D. and Shenoy, P. (eds) *Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference*, San Francisco, CA: Morgan Kaufmann.

Mahoney, S.M. and Laskey, K.B. (1998) Constructing Situation Specific Networks. In Cooper, G. and Moral, S. (eds) *Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference*, San Francisco, CA: Morgan Kaufmann

Mahoney, S. M. and K. B. Laskey (1996) Network Engineering for Complex Belief Networks, Laskey, K.B. and Mahoney, S.M. Network Engineering for Agile Belief Network Models. *IEEE Transactions in Knowledge and Data Engineering*, 2000.

USS Cole Commission Report, January 9, 2001

Wellman, M.P., J.S. Breese, and R.P. Goldman From knowledge bases to decision models. *The Knoweldge Engineering Review*, 7(1):35-53. November, 1992.