LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

LA-UR -89-527

LA-UR--89-527

DE89 007990

TITLE RISK ASSESSMENT AND LAVA'S DYNAMIC THREAT ANALYSIS

AUTHOR(S) Suzanne T. Smith

SUBMITTED TO 12th National Computer Security Conference, Baltimore, Maryland, October 10-13, 1989

DISCLAIMER

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

RISK ASSESSMENT AND LAVA'S DYNAMIC THREAT ANALYSIS

Suzanne T. Smith
Los Alamos National Laboratory
Safeguards Systems Group, MS-E551
P. O. Box 1663
Los Alamos, New Mexico 87545

ABSTRACT

LAVA (the Los Alamos Vulnerability/Risk Assessment system) is a three-part systematic approach to risk assessment that can be used to model risk assessment for a variety of application systems such as computer security systems, communications security systems, and information security systems. The first part of LAVA is the mathematical methodology based on such disciplines as hierarchical system theory, event-tree analysis, possibility theory, and cognitive science. The second part is the general software engine, written for a large class of personal computers, that implements the mathematical risk model. The third part is the application data sets written for a specific application system. The methodology provides a framework for creating applications for the software engine to operate upon; all application-specific information is data. Using LAVA, we build knowledge-based expert systems to assess risks in application systems comprising a subject system and a safeguards system. The subject system model comprises sets of threats, assets, and undesirable outcomes; because the threat to security systems is ever-changing, LAVA provides for an analysis of the dynamic aspects of the threat spectrum. The safeguards system model comprises sets of safeguards functions for protecting the assets from the threats by preventing or ameliorating the undesirable outcomes; sets of safeguards subfunctions whose performance determine whether the function is adequate and complete; and sets of issues that appear as interactive questionnaires, whose measures (in both monetary and linguistic terms) define both the weaknesses in the safeguards system and the potential costs of an undesirable outcome occurring. The user need have no knowledge of formal risk assessment techniques - all the technical expertise and specialized knowledge are built into the software engine and the application system itself. LAVA applications include our popular computer security application and other applications for embedded systems, survivability systems, transborder data flow systems, and property control systems. LAVA application systems have been used by federal government agencies since 1984.

LAVA'S DYNAMIC THREAT ANALYSIS

Suzanne T. Smith
Los Alamos National Laboratory
Safeguards Systems Group, MS-E551
P. O. Box 1663
Los Alamos, New Mexico 87545

Introduction

LAVA (the Los Alamos Vulnerability/Risk Assessment system) is an original systematic approach to risk assessment developed at the Los Alamos National Laboratory to deal with risks inherent in massive, complicated systems. Characteristics of such systems are huge bodies of imprecise data, indeterminate (and possibly undetected) events, large quantities of subjective information, and a dearth of objective information. The impetus for developing LAVA was the existence of Federal requirements for periodic risk assessments of a variety of systems, coupled with the need for an inexpensive, reusable, automated risk assessment tool firmly rooted in science [1]. When the LAVA project began in 1983, there was no such tool [2]; LAVA was designed to fill that gap [3].

LAVA is an alternative to existing quantitative methods, providing an approach that is both objective and subjective, and producing results that are both quantitative and qualitative. In addition, LAVA could be used as a self-testing aid in preparing for inspections, as a self-evaluating device in testing compliance with the various orders and criteria that exist, and as a certification device by an inspection team.

LAVA is a three part systematic approach to risk assessment that can be used to model a variety of application systems such as computer security systems, communications security systems, information security systems, and others. The first part of LAVA is the mathematical model based classical risk assessment [4,5], hierarchical multilevel system theory [6,7], decision theory [8-11], fuzzy possibility theory [11-14], expert system theory [15,16], utility theory [17,18], and cognitive science [19,20]. (The mathematical model has been presented at other technical meetings [21-23], and generally will not be addressed in depth in this paper.) The second part is the implementation of the mathematical risk model as a general software engine, written in a commercially available programming language for a large class of personal computers. The third part is the application data sets written for a specific application system. LAVA provides a framework [24] for creating applications upon which the software engine operates; all application specific information appears as data.

We use the LAVA system to develop a hierarchical structure and sets of fuzzy analysis trees for modeling risk assessment for a variety of systems associated with computer and information security. With LAVA, we build knowledge-based expert systems to assess risks in application systems comprising a subject system and a safeguards system. The subject system model is sets of threats, assets, and undesirable outcomes; because the threat to security systems is ever-changing, LAVA provides for an analysis of the dynamic aspects of the threat spectrum--the dynamic threat analysis [25] is the subject of this paper. The safeguards system model has three parts: sets of safeguards functions for protecting the assets from the threats by preventing or ameliorating the undesirable outcomes; sets of safeguards subfunctions whose performance determines whether the function is adequate and complete; and sets of issues, appearing as interactive questionnaires, whose measures (in both monetary and linguistic terms) define both the weaknesses in the safeguards system and the potential costs of an undesirable outcome occurring.

The user need have no knowledge of formal risk assessment techniques. All the technical expertise and specialized knowledge are built into the software engine and the application system. LAVA applications include the popular computer security application [26-29] and applications for nuclear power plant control rooms [30], embedded systems, survivability systems, transborder data flow systems [31], property control systems, nuclear processing plant safeguards systems [32], and others. LAVA application systems have been in use by Federal government agencies since 1984.

LAVA Application Models

The General LAVA Application Model

Using LAVA, we build knowledge based expert systems for assessing risks in applications systems. There are two parts that define an application model. The first part is composed of the following elements: the hierarchical structure and trees that define the framework of the model the threat, asset, and outcome sets; the fuzzy outcome possibility matrix; the safeguards functions for each threat asset pair, based upon the kinds of interactions that might result in one or more of the outcomes; the safeguards subfunctions for each function; mitigating factors for outcome severity; and the contributing factors, both linguistic and monetary, to the potential cost of a successful attack. The second part is the set of questionnaires, implemented as data sets on which the general software engine operates: the vulnerability assessment questionnaire, the outcome severity mitigation questionnaire, the dynamic threat questionnaire (if applicable), and the monetary and linguistic impact (or cost) questionnaires.

The vulnerability assessment questionnaire for a given application is concatenated from a library of category questionnaires;

that come from specific security orders, inspection criteria, interviews with various experts in the field, and general good security practice. The questions themselves represent individual safeguards (called "safeguards elements") or portions of safeguards (called "safeguards attributes") that are related through a database structure to one or several of the safeguards subfunctions. The vulnerability questionnaire can comprise from a few hundred to several thousand questions, depending on the required analytical depth.

The other questionnaires are all considerably smaller than the vulnerability questionnaire. The outcome severity mitigation questionnaire inquires about the presence and estimated effectiveness of any mitigating situations that might be pertinent. If intelligence information is available and analytical detail about the dynamic threat is required, the dynamic threat questionnaire seeks information about the motivation, capability, and opportunity of the current known threat and about the attractiveness of each asset set to the threat; if such information is not available, the user estimates a relative attractiveness factor for the asset sets and whether the dynamic threat is the same as or, in varying degrees, larger or smaller than the background (static) threat. The impact questionnaires ask cost-related questions in either linguistic or monetary terms. With the exception of the intelligence-based dynamic threat questionnaire, all of the questions in these questionnaires number in the single or double digits (usually not more than a dozen or so questions).

Users are not required to be expert risk analysts to use a LAVA application--that mathematical and analytical expertise already exists as a part of the mathematical model and its general software engine. Expert knowledge about the structure and characteristics of safeguards and security systems is a part of the specific application model. The only knowledge required of users is information about that which they know best: their own facility, organization, assets, equipment, policies, procedures, and security practices. The LAVA software system elicits this information by means of the automated questionnaires administered to evaluation teams whose members have diverse backgrounds and responsibilities. LAVA generates both general reports for management and detailed reports for operations staff from information obtained from the questionnaires.

LAVA/CIS: The Computer/Information Security Model

For our computer/information security application model, LAVA/CIS, we postulate four assets: 1) the facility, including physical plant and personnel; 2) hardware, including all computing and ancillary pre- and post-processing hardware; 3) machine interpretable information, including software, input and output files, and databases; and 4) human interpretable information, including documents, screen displays, graphs, charts, film output, and so

forth. The model's threat set consists of three threats: 1) natural, random, and environmental hazards; 2) direct or onsite humans, including the authorized insider; and 3) indirect or offsite humans. Figures 1-2 show the hierarchical structures for two of the threat categories with respect to the four asset categories; included in these hierarchies, and discussed later in this paper, are representative safeguards functions and subfunctions associated with each threat-asset pair. Figure 3 shows how this relates to the entire model.

There are six undesirable outcomes considered in the computer/information security model: 1) unauthorized access or use; 2) modification or tampering; 3) damage or destruction; 4) theft; 5) unauthorized disclosure; and 6) denial of use. It is important to note that a single event can result in the simultaneous occurrence of more than one of the outcomes. Figure 4 shows the outcome possibility matrix for the threat-asset combinations; a value of zero indicates that the outcome is impossible for that threat-asset combination, and a value of unity means the outcome is possible for that threat-asset pair; greater granularity can be achieved by assigning values lying between zero and unity.

Once we have established the threat, asset, and outcome sets and the outcome possibility matrix, we then address what constitutes the ideal safeguards system for preventing the threats from attacking the assets and achieving the postulated outcomes. For this we define a set of safeguards functions for each of the distinguishable threat-asset pairs (nine T-A pairs, in this application) in such a way that the relative importance of each function within the set of functions for each T-A pair is about the same. Then, for each of the individual safeguards functions, we define a set of subfunctions that provide performance criteria for the adequacy and completeness of that safeguards function; each of the subfunctions is devised so that the relative importance of each subfunction within a specific function is about the same. Again, Figs. 1-3 show the safeguards functions and subfunctions for each distinguishable threat-asset pair.

The Dynamic Threat Analysis

Both government and corporate organizations may be the targets of a variety of hostile agents [33,34], and the intensity of the threat may change with time and circumstances. The dynamic threat strength can be analyzed if the subject system is extremely sensitive to a changing threat and if the subject organization has access to the kinds of information the analysis requires. The dynamic threat analysis takes into account possible threat agents and their potential attack goals with respect to the target(s) of the attack.

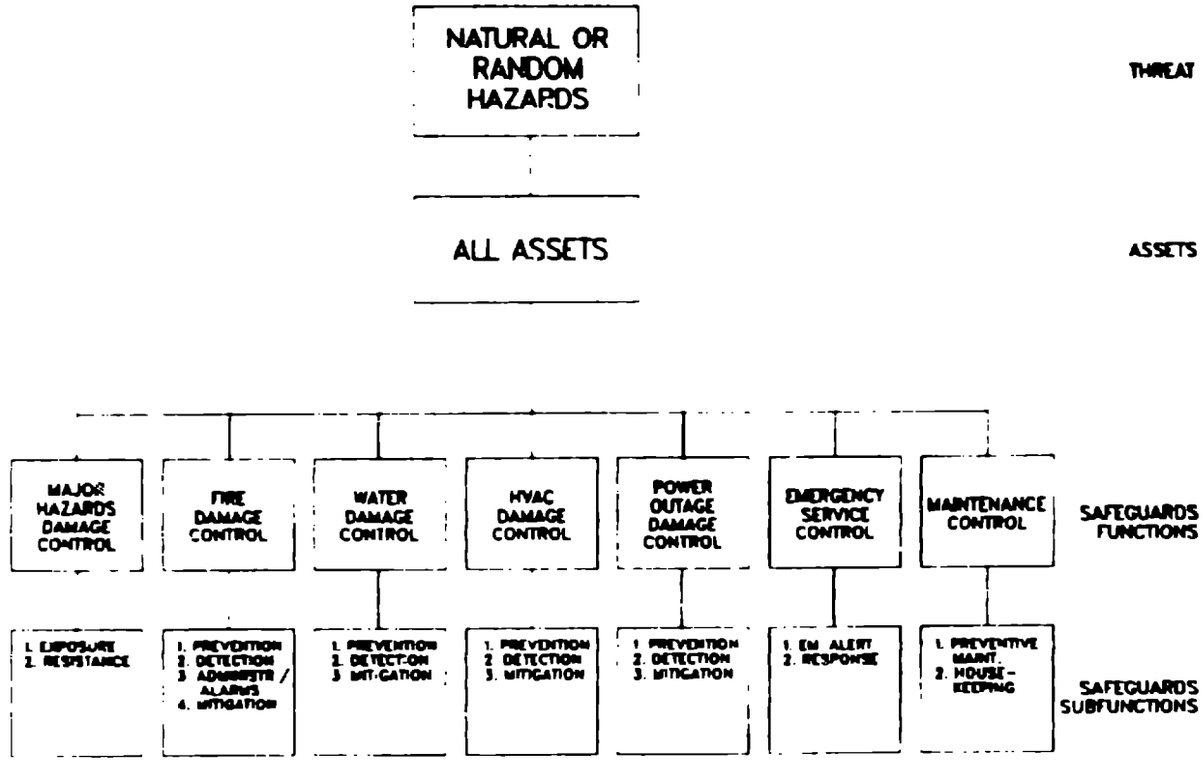

Fig. 1. Natural hazards hierarchy for computer/information security application.

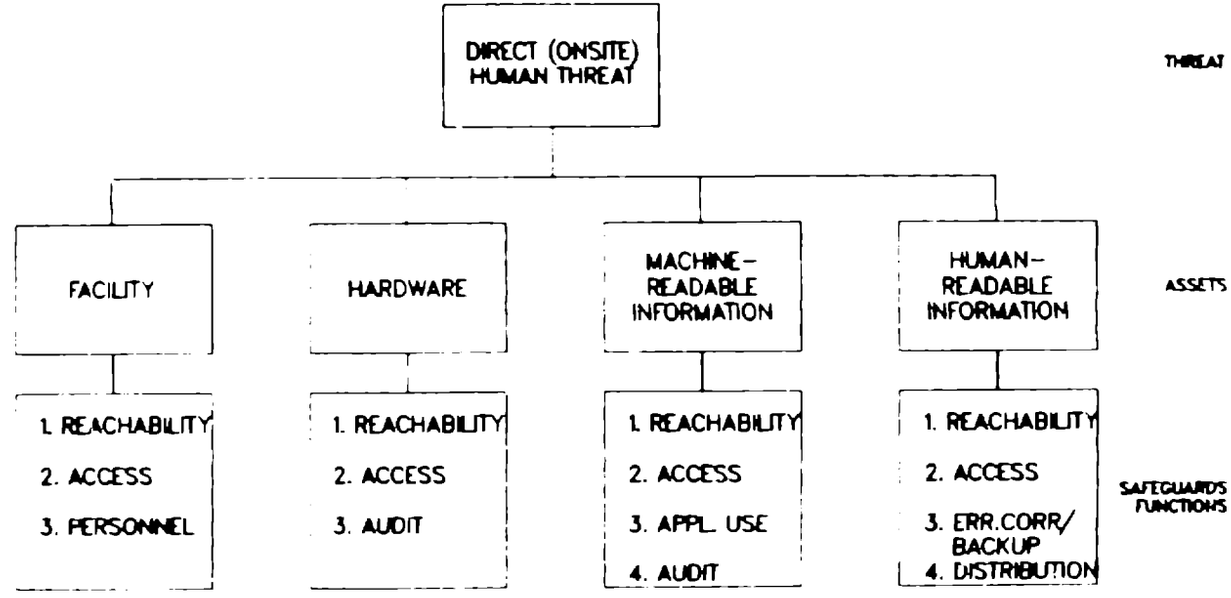

Fig. 2. Direct (onsite) human threat hierarchy for computer/information security application.

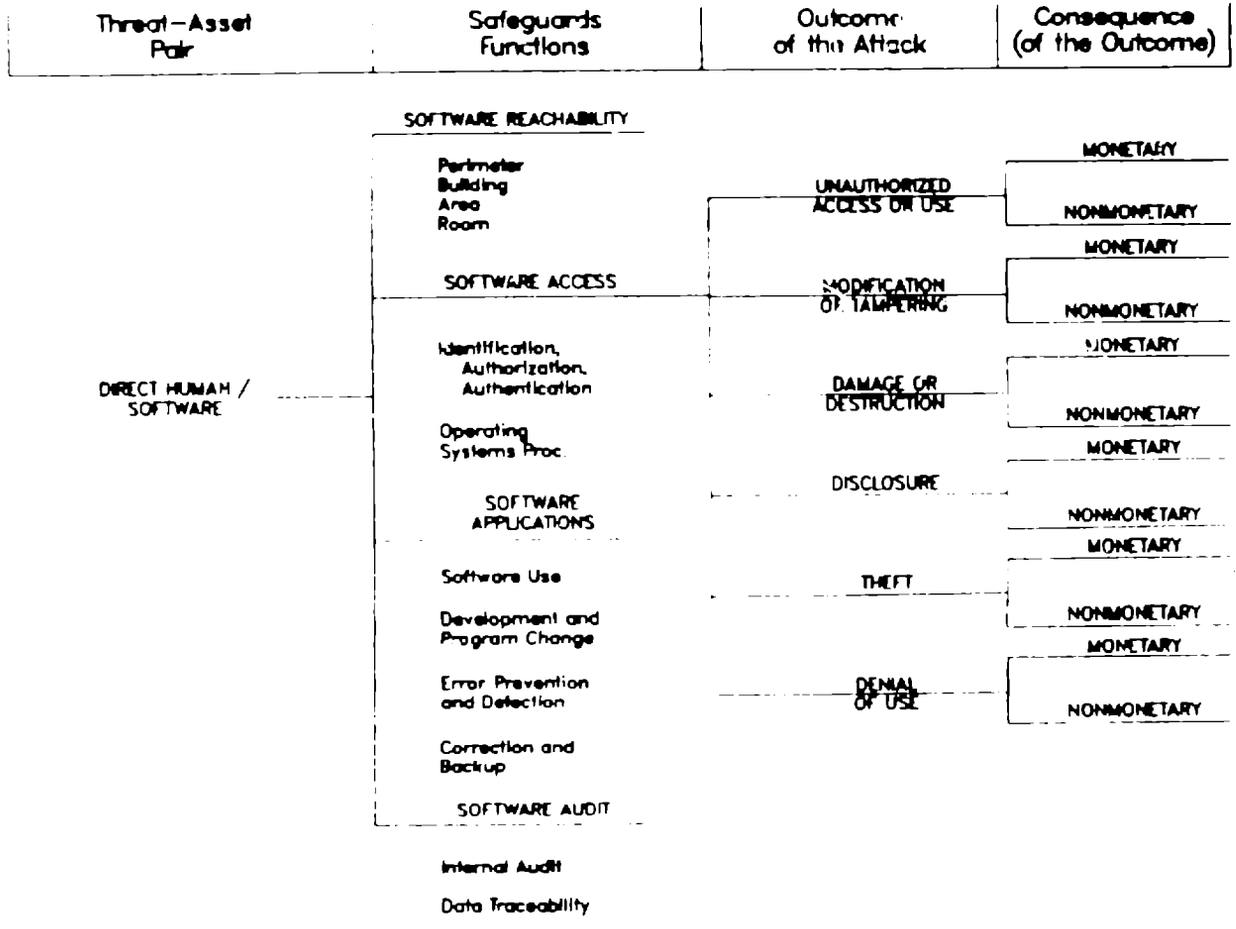

Fig. 3. Direct human/software scenario analysis tree.

	Unauthorized Access or Use	Modification or Tampering	Damage or Destruction	Disclosure	Theft	Denial of use
Natural Hazards Facility	0	1	1	0	0	1
Natural Hazards Hardware	0	1	1	0	0	1
Natural Hazards Software	0	1	1	0	0	1
Natural Hazards - Documents/ Displays	0	1	1	0	0	1
Direct Human Facility	1	1	1	1	1	1
Direct Human Hardware	1	1	1	1	1	1
Direct Human Software	1	1	1	1	1	1
Direct Human Documents, Displays	1	1	1	1	1	1

Fig. 4. Outcome possibility matrix for computer/information security application.

The threat component measures the relative strength of identifiable threat agents in terms of asset attractiveness, motivation, opportunity, and capability with respect to the spectrum of assets, the corresponding safeguards functions, and the set of possible outcomes. Asset attractiveness to the threat agent is different from asset value to the organization, reflecting the different value structure of the threat agent; it is a rough indicator of attack likelihood in that a threat agent is unlikely to mount an attack on an unattractive asset. Motivation is a measure of how much effort or what part of his resources a threat agent is willing to expend on an attack and how dedicated he is to carrying out the attack. Capability is a measure of the resources--knowledge (training), information (intelligence), funds, skills, equipment, armament, personnel--the threat agent has at his disposal. Opportunity is a measure of how easy it is for the threat agent to achieve an enabling proximity for an attack: how easy it is for him physically to reach the object of attack, how easy it is for him to attack or to access the object, how easy it is for him to travel undetected (both in the neighborhood of the object of attack and from afar to get near the object), and so forth. Opportunity is separate and different from potential site vulnerabilities. Figure 5 illustrates the analysis structure for the dynamic threat analysis.

There are several broad categories of threat agents having a variety of goals. Possible categories of threat agents might be, for example:

a) information gatherers (e.g., spies or hostile intelligence services),
b) terrorists,
c) pro- or anti-"X" radicals or extremists (where "X" could be almost anything!),
d) representatives of organized crime,
e) other criminals (non-malicious criminals and pranksters),
f) insiders (employees, contractors, etc.),
g) outsiders with access, and
h) Mother Nature.

The dynamic aspects of the natural hazards may or may not be of interest; these include both random natural hazards, such as volcanic eruptions or earthquakes, as well as the natural hazards more cyclic in nature, such as hurricanes, tornadoes, torrential rains, and the like. The human threat agents in each of these categories all act for different reasons, so they may differ widely in motivation, capability, and opportunity. Similarly, the goals of the attacks may vary, but all categories of goals may be used by all categories of threat agents. Some possible goal categories are

1) information and/or material collection (e.g., espionage or theft of nuclear materials),
2) sabotage,

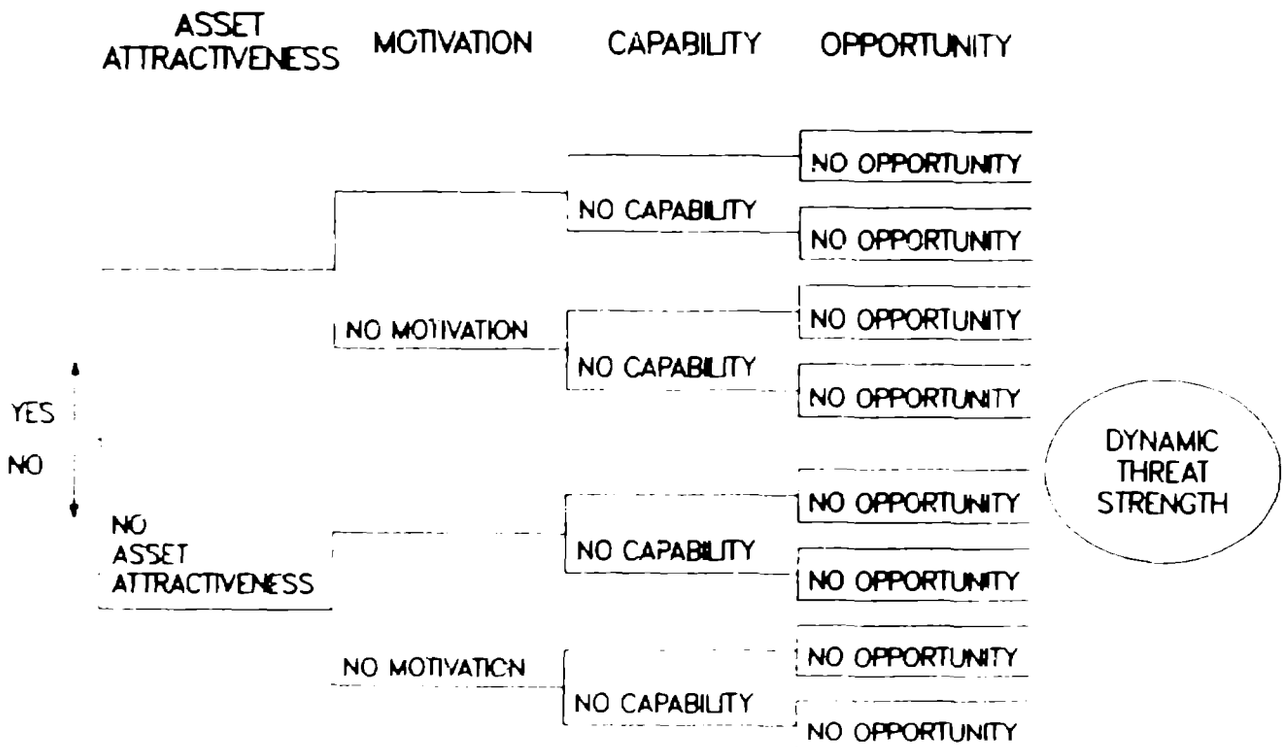

Fig. 5. Analysis structure for dynamic threat.

3) theft, embezzlement, fraud--generally for monetary gain,
4) damage or destruction,
5) extortion,
6) disrupting business or mission, and
7) surmounting an intellectual challenge.

Clearly, more than one of the categories may be the goal of a single attack, and a single attack may be perpetrated by more than one category of threat agent.

The approach to assessing the dynamic part of the threat component by considering categories of threat agents and possible categories of attack goals is parallel to the approaches used for both the vulnerability analysis and the general consequence analysis. Potential scenarios are modeled implicitly as the relationship between the threat-asset pairs and the safeguards functions in the vulnerability analysis, and as the relationship between the assets and the threat elements (asset attractiveness, motivation, capability, and opportunity) in the threat assessment. Similarly, the attack goals are modeled implicitly in the capability component of the dynamic threat measure and are approximately equivalent to the outcomes used in the consequence analysis.

An interactive questionnaire models the contributors to the dynamic threat in terms of specific threat groups. A fuzzy degree of strength is calculated for each group based on asset attractiveness, motivation, capability, and opportunity relative to a specific [threat, asset, safeguards function, outcome] quadruplet. A relational database keeps track of which threat groups can affect each quadruplet so that an overall or total value for the dynamic threat strength can be calculated for each quadruplet, which is used subsequently in the loss exposure calculations.

Conclusions

LAVA's capability to assess the dynamic aspects of the threat spectrum makes it an ideal tool for modeling applications of interest to the intelligence and military communities. It would also be highly applicable in the business community in situations ripe for industrial espionage.

Using the LAVA approach for risk assessment has benefits that do not accrue from the use of other methods. First, the automated report generators produce results that are immediately usable both to managers who must make major, far reaching decisions and to the security personnel in the field whose job it is to maintain an acceptable level of safeguards. Second, because LAVA produces both qualitative and quantitative results, users feel more comfortable with the results because they understand both the results and the information that produced those results. Third, because LAVA does not require the user to generate probabilities (often unfounded)

for its operation but instead relies on a natural-language user-friendly interface to acquire its data, users are more willing to act upon its results. Fourth, LAVA includes a way to assess the changing, or dynamic, aspects of the threat spectrum. And finally, because of the team environment in which an assessment is performed and discussions that arise among team members, using a LAVA application has proved to be an experience that both raises the security consciousness of the users and enhances the overall working environment at the facility.

References

[1] S. Katzke, "National Bureau of Standards Perspective on Risk Analysis: Past, Present, and Future," presented at the 1st Federal Risk Analysis Workshop, Montgomery, Alabama, January 1985.

[2] S. T. Smith, "A Government-Wide Overview of Risk Analysis Methodologies," presented at the 8th DOE Computer Security Group Conference, Richland, Washington, April 16-18, 1985.

[3] S. T. Smith and J. J. Lim, "An Automated Procedure for Performing Computer Security Risk Analysis," in Proceedings 6th Annual ESARDA Symposium on Safeguards and Nuclear Material Management, 1984, ESARDA 17, pp. 527-530.

[4] N. J. McCormick, Reliability and Risk Analysis: Methods and Nuclear Power Applications. New York: Academic Press, 1981.

[5] W. D. Rowe, An Anatomy of Risk. New York: John Wiley & Sons, 1977.

[6] M. D. Mesarovic, D. Macko, and Y. Takahara, Theory of Hierarchical Multilevel Systems. New York and London: Academic Press, 1970.

[7] Y. M. I. Dirickx and L. P. Jennergren, Systems Analysis by Multilevel Methods. New York: John Wiley & Sons, 1979, pp. 10-82.

[8] P. C. Fishburn, Decision and Value Theory. New York: John Wiley & Sons, 1964.

[9] R. L. Keeney and H. Raiffa, Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: John Wiley & Sons, 1976.

[10] R. Schlaifer, Analysis of Decisions Under Uncertainty. Huntington, New York: Robert E. Krieger Publishing Company, 1978.

[11] R. E. Bellman and L. A. Zadeh, "Decision-making in a Fuzzy Environment," Management Science, Vol. 17, No. 4, pp. B141-B164, December 1970.

[12] A. Kaufmann and M. M. Gupta, Introduction to Fuzzy Arithmetic: Theory and Applications. New York: Van Nostrand Reinhold Company, 1985.

[13] L. A. Zadeh, "Fuzzy Sets as a Basis for a Theory of Possibility," Fuzzy Sets and Systems, Vol. 1 pp. 3-28, 1978.

[14] C. V. Negoita, Expert Systems and Fuzzy Systems. Menlo Park, California: The Benjamin/Cummings Publishing Company, Inc., 1985, pp. 52-58, 74-88, 95-112.

[15] P. H. Winston, Artificial Intelligence. Reading, MA: Addison-Wesley, 1984, pp. 251-288.

[16] R. Jain, "A Procedure for Multiple-Aspect Decision-Making Using Fuzzy Sets," Int. J. Systems Sci., Vol. 8, No. 1, pp. 1-7, January 1977.

[17] P. J. H. Schoemaker and C. C. Waid, "An Experimental Comparison of Different Approaches to Determining Weights in Additive Utility Models," Management Science, Vol. 28, No. 2, February 1982. **PAGE NUMBER**

[18] E. M. Johnson and G. P. Huber, "The Technology of Utility Assessment," IEEE Trans. Sys., Man, Cyber., Vol. SMC-7, No. 5, pp. 311-325, May 1977.

[19] L. A. Zadeh, K.-S. Fu, K. Tanaka, and M. Shimura (Eds.), Fuzzy Sets and Their Applications to Cognitive and Decision Processes. New York: Academic Press, 1975.

[20] S. Sudman and N. M. Bradburn, Asking Questions: A Practical Guide to Questionnaire Design. San Francisco: Jossey Bass, Inc., 1982.

[21] S. T. Smith and J. J. Lim, "An Automated Interactive Expert System for Evaluating the Effectiveness of Computer Security Measures," presented at the 7th Department of Defense/ National Bureau of Standards Computer Security Conference, Gaithersburg, Maryland, September 24 26, 1984.

[22] S. T. Smith, J. R. Phillips, R. M. Tisinger, J. J. Lim, D. C. Brown, and P. D. FitzGerald, "LAVA: A Conceptual Framework for Automated Risk Analysis," presented at the 1986 Annual Meeting of the Society for Risk Analysis, Boston, November 9 12, 1986.

[23] S. T. Smith, "LAVA: An Expert System Framework for Risk Analysis", presented at the 1st International Computer Security Risk Management Model Builders Workshop, Denver, Colorado, May 24-26, 1988.

[24] S. T. Smith and J. J. Lim, "Framework for Generating Expert Systems to Perform Computer Security Risk Analysis," Proceedings First Annual Armed Forces Communications and Electronics Association Symposium and Exposition on Physical and Electronics Security, 1985, pp. 24-1 - 24-7.

[25] S. T. Smith, J. R. Phillips, D. C. Brown, and P. D. FitzGerald, "Assessing the Threat Component for the LAVA Risk Management Methodology," presented at the Ninth DOE Computer Security Group Conference, Las Vegas, Nevada, May 6-8, 1986.

[26] S. T. Smith and J. J. Lim, "An Automated Method for Analyzing Computer Security Risk," presented at the Seventh DOE Computer Security Group Conference, New Orleans, April 17-19, 1984.

[27] S. T. Smith and J. J. Lim, "An Automated Method for Assessing the Effectiveness of Computer Security Safeguards," presented at the IFIPS Second International Congress on Computer Security, Toronto, Canada, September 10-12, 1984.

[28] S. T. Smith and J. J. Lim, "LAVA: An Automated Computer Security Vulnerability Assessment Software System (Version 0.9)," Los Alamos National Laboratory document LA-UR-85-4014, December 1985.

[29] S. T. Smith et al., "LAVA for Computer Security: An Application of the Los Alamos Vulnerability Assessment Methodology," Los Alamos National Laboratory document LA-UR-86-2942, 1986.

[30] S. T. Smith and J. J. Lim, "Assessment of Computer Security Effectiveness for Safe Plant Operation," Trans. Am. Nucl. Soc., Vol. 46, pp. 525-526, June 1984.

[31] S. T. Smith, J. J. Lim, and J. Lobel, "Application of Risk Assessment Methodology to Transborder Data Flow," in Handbook on the International Information Economy. Springfield, Virginia: Transnational Data Report, November 1985.

[32] S. T. Smith and R. M. Tisinger, "Modeling Risk Assessment for Nuclear Processing Plants with LAVA," Nucl. Mater. Manage., Vol. XVII (Proceedings Issue), pp. 315-318, June 1988.

[33] N. R. Bottom, Jr., and R. R. J. Gallati, Industrial Espionage: Intelligence Techniques and Countermeasures. Boston: Butterworth Publishers, 1984.

[34] R. Fells and P. Nehemkis, Corporate Intelligence and Espionage: A Blueprint for Executive Decision Making. New York: Macmillan, 1984