LA-UR -86- 4113

LA-UR--86-4113

DE87 003720

TITLE: LAVA: A CONCEPTUAL FRAMEWORK FOR AUTOMATED RISK ANALYSIS

AUTHOR(S): S. T. Smith, J. J. Lim, J. R. Phillips, R. M. Tisinger,
D. C. Brown, and P. D. FitzGerald

SUBMITTED TO: 1986 Annual Meeting of the Society for Risk Analysis,
Boston, Massachusetts, November 9-12, 1986

DISCLAIMER

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

Los Alamos

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

MASTER

LAVA: A CONCEPTUAL FRAMEWORK FOR AUTOMATED RISK ANALY?IS

S. T. Smith,* J. J. Lim,** J. R. Phillips,*
R. M. Tisinger,* D. C. Brown,† and P. D. FitzGerald†

ABSTRACT

At Los Alamos National Laboratory, we have developed an original methodology for performing risk analyses on subject systems characterized by a general set of asset categories, a general spectrum of threats, a definable system-specific set of safeguards protecting the assets from the threats, and a general set of outcomes resulting from threats exploiting weaknesses in the safeguards system. The Los Alamos Vulnerability and Risk Assessment Methodology (LAVA) models complex systems having large amounts of "soft" information about both the system itself and occurrences related to the system. Its structure lends itself well to automation on a portable computer, making it possible to analyze numerous similar but geographically separated installations consistently and in as much depth as the subject system warrants. LAVA is based on hierarchical systems theory, event trees, fuzzy sets, natural-language processing, decision theory, and utility theory. LAVA's framework is a hierarchical set of fuzzy event trees that relate the results of several embedded (or sub-) analyses: a vulnerability assessment providing information about the presence and efficacy of system safeguards, a threat analysis providing information about static (background) and dynamic (changing) threat components coupled with an analysis of asset "attractiveness" to the dynamic threat, and a consequence analysis providing information about the outcome spectrum's severity measures and impact values. Each sub-analysis can be simplified or made complex, depending on the sensitivity and relative worth of the subject system. Personnel at the subject site see only an interactive questionnaire eliciting from them data about the presence and quality of the safeguards, the potential consequences of a successful threat, and the target organization's preference structure--the technical expertise is built into the model (and the computer code) itself. LAVA yields quantitative <u>and</u> qualitative insights: a pair of values (monetary and linguistic) express loss exposure for each threat/asset/safeguards-function/outcome quadruple. Using LAVA, we have modeled our widely used computer security application as well as LAVA/CS systems for physical protection, transborder data flow, contract awards, and property management. It is presently being applied for modeling risk management in embedded systems, survivability systems, and weapons systems security. LAVA is especially effective in modeling subject systems that include a large human component.

* Los Alamos National Laboratory, MS-E541, Los Alamos, NM 87544.

** Lim and Orzechowski Associates, Consultants.

† U. S. Government.

I. INTRODUCTION

LAVA (Los Alamos Vulnerability and Risk Assessment Methodology) is an original approach to risk management developed at the Los Alamos National Laboratory. It is a systematic methodology for assessing vulnerabilities and risks in complex safeguards and security systems. The associated LAVA software[1] is part of a research effort to provide tools to identify vulnerabilities and risks in large, complex systems whose modeling is generally intractable by other methods. LAVA is being implemented as a set of computer programs that run on a widely used class of personal computers; these programs operate on a variety of application system models and are executed in a team environment. Each LAVA application's implementation is designated as LAVA/XX, where XX is an identifier associated with the specific application.

Users are not required to be expert risk analysts to use LAVA/XX--that mathematical and analytical expertise already exists as a part of the methodology's software system. 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 automated questionnaires[2] administered to evaluation teams whose members have diverse backgrounds and responsibilities. LAVA/XX generates both general reports for management and detailed reports for operations staff from information obtained in the questionnaires.

The subject systems to which the LAVA methodology can be applied are massive, complicated systems characterized by a large human component, by large bodies of imprecise data (very little "hard" information and enormous quantities of "soft" information), and by often indeterminate events (events that may or may not have happened, or, if they happened, may not have been detected). The outcomes resulting from threats exploiting system vulnerabilities are often of a catastrophic nature, defying quantification in ordinary terms.

The methodology makes use of hierarchical multilevel systems theory, event-tree-like analysis, fuzzy set theory, decision theory, utility theory, knowledge-based expert-system theory, and natural language processing. The methodology gives both qualitative and quantitative insights into the vulnerabilities in the system of safeguards, yields an accurate picture of the state of the subject organization's safeguards system, and produces both qualitative and quantitative expressions for the system's loss exposure (risk).

The LAVA methodology has been applied successfully to modeling vulnerabilities and risks in computer-security systems[3-6], plant control-room operation[7], computer security for nuclear safeguards[8-9], contract control systems, transborder data-flow systems[10], and Government-property control systems. It is presently being applied for modeling risk management in embedded systems, in survivability systems, and in weapons systems security.

II. DEFINITIONS OF TERMS

Within the context of LAVA, some terms have a specialized meaning. These include, in alphabetical order,

ASSET - an item or category of items having some intrinsic value to the subject organization; assets are acted upon by threats, leading to outcomes.

HARD DATA - data or information that can be represented easily in quantitative terms, like a valve failure rate or the yearly average number of times that a known event occurs.

IMPACT - the consequence, cost, or effect upon the subject organization of an outcome of severity "X" resulting from a threat successfully exploiting a safeguards function vulnerability for a particular asset; the impact is given as a pair of values, one monetary (economic) and one non-monetary (linguistic).

LOSS EXPOSURE - the risk to the subject organization of a threat successfully exploiting a safeguards function vulnerability to produce an outcome of some calculated degree of severity; given as a pair of values, one monetary and one non-monetary.

OUTCOME - the (usually undesirable) event that occurs when a threat successfully exploits a vulnerability in a safeguards function for a particular asset. (Often wrongly confused with threat, for example, "the threat of total destruction;" total destruction is actually an OUTCOME. See next page for definition of threat.)

OUTCOME SEVERITY - a measure of how successful the threat is for this specific outcome or how much damage occurs. Outcome severity is a function of the relative weakness of a safeguards function and the relative strength of the threat.

SAFEGUARDS - policies, procedures, physical or logical devices, and so forth designed to prevent the undesirable outcomes by protecting the assets from the threats.

SAFEGUARDS FUNCTION - the functional representation of the protective mechanism that a safeguard or set of safeguards is intended to achieve.

SOFT DATA - data or information that is difficult or impossible to quantify, like the value of a human life or of a highly classified document, or whether a subtle event has occurred or has been detected.

SUBJECT SYSTEM - the system or universe upon which the assessment is being performed (the SUBJECT ORGANIZATION is responsible for the subject system).

THREAT - a person, force of nature, thing, or idea (or category of same) posing some danger or menace to the assets; a threat is an active force (or actor) and is not to be confused with OUTCOME.

VULNERABILITY - a weakness or flaw in a safeguards function, such as in a security system or procedural system, that can be exploited by a threat to cause harm to an asset or set of assets.

III. LAVA'S TECHNICAL BASIS

The LAVA methodology is a structured, modular, systems approach to risk management. For any specific application, this approach consists of four phases: modeling the system, gathering the information necessary for LAVA's analysis, determining the subject organization's loss exposures from potentially three separate analyses, and "solving" the problem of risk management in terms of action.

A. The modeling phase

For a subject system, LAVA evaluates system vulnerabilities, analyzes the consequences of vulnerability exploitation, and calculates the set of loss exposures of the subject organization resulting from the consequence set. There are well-defined steps that we must take when modeling an application for LAVA. First we identify the subject system, delineate the characteristics of that system, and specify the scope of the analysis. Next we define the system assets, the set of undesirable outcomes, and the threats to the assets. Then we consider all the ways that the threats can interact with the assets so we can understand which of the outcomes might result from these interactions and what safeguards functions must be in place to protect the assets from the threats. We then determine what factors might affect the outcome's severity and what subfunctions will determine the performance level of the safeguards functions. We then design interactive questionnaires to model the specifics of the subject system.

Assets are items of value to the organization that must be safeguarded against harm or compromise. Assets include (but are not limited to) real property, equipment of all kinds, documents, personnel, information, reputation, the ability to do business, and anything else of value, including the organization itself. Assets possess properties affecting their value, such as sensitivity, criticality, compromisability, theftworthiness, timeliness, and so forth. In modeling assets for a LAVA application, we create asset categories--categories that treat similar assets as one, such as human-readable information (information in the form of reports, letters, computer-terminal displays, plots, or other human-readable form) or machine-readable information

(information stored in coded form in a computer or word processor or on storage media for such machines, requiring a machine to translate the coded information into a form that can be understood by a human).

Threats are active forces posing some danger or menace to the assets. Threats can be people (insiders or outsiders), forces of nature, things, or ideas that can cause the assets harm. Again, in modeling a LAVA application's threat, we use broad threat categories, such as natural hazards, on-site humans, and off-site humans. Threats can be treated as always lurking in the background and not changing very much (the static or background threat) or, in the case of very sensitive applications in which the subject organization has access to certain intelligence information, threats can be treated as having two parts--a static component and a dynamic component changing with time.[11] We will discuss the concept of dynamic threat further in the analysis section.

We define an asset space, $\underline{A}$, and a threat space, $\underline{T}$, to describe the organization's assets and the threats to them, having in mind as we do this a specific set of undesirable outcomes. The asset and threat spaces comprise the specific categories modeling a specific application system:

$$\underline{A} = \{a_1, a_2, \ldots, a_i\}$$

$$\underline{T} = \{t_1, t_2, \ldots, t_j\} .$$

We consider the threats and assets in [threat, asset] pairs,

$$\underline{T}\times\underline{A} = \{t_1a_1, t_1a_2, \ldots, t_1a_i, t_2a_1, \ldots, t_ja_i\} ,$$

so that we can break down in a systematic way the kinds of threat-asset interactions that are possible for the application system. The kinds of possible interactions when coupled with the set of undesirable outcomes, represented as

$$\underline{O} = \{o_1, o_2, \ldots, o_q\} ,$$

determine what the safeguards function set should be.

The relationship of threats, assets, and safeguards functions is illustrated in the hierarchical multilevel disaggregation structures[12] (Figs. 1 and 2) taken from LAVA's application to computer security. In all LAVA models, there is a separate hierarchy for each threat; the top level is the threat, the second level contains the assets, and the bottom level lists the safeguards functions for each threat-asset pair. In the computer-security application, there are three threat categories: natural or random hazards, on-site humans,

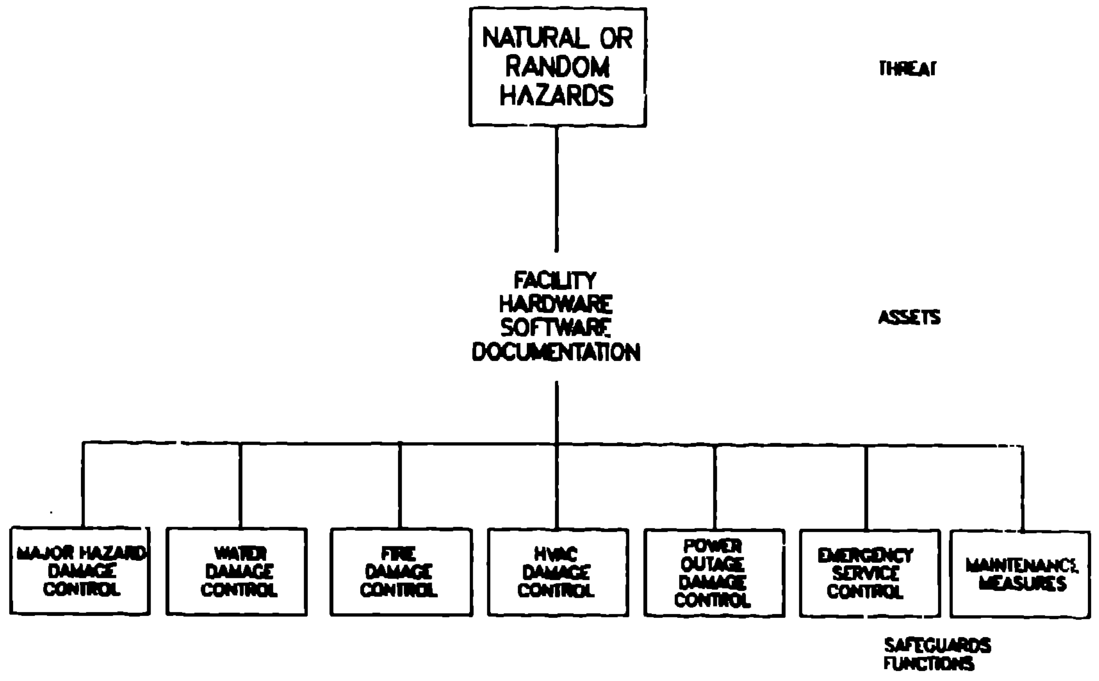

Fig. 1
Risk assessment structure for natural hazards threat.

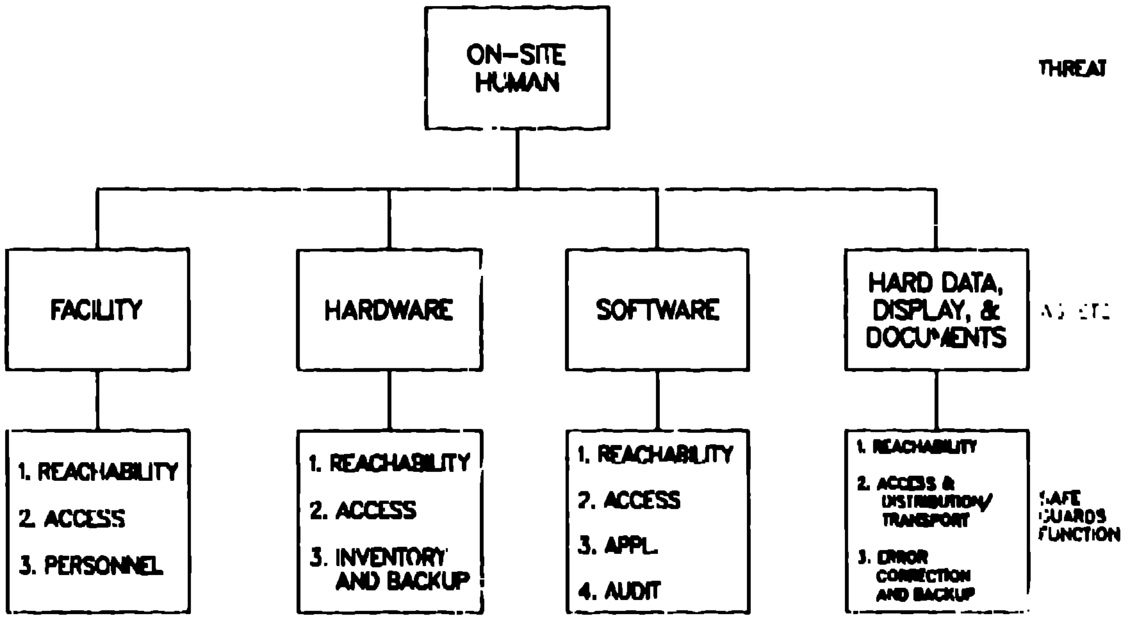

Fig. 2
Risk assessment structure for on-site human adversary.

and off-site humans. There are four asset categories: the facility or physical plant, computer-related hardware, machine-readable information, and human-readable information. In this example application, there are six outcomes making up the set of undesirable outcomes that the total set of safeguards functions are designed to prevent: unauthorized access or use, modification or tampering, damage or destruction, disclosure, theft, and denial of use. Figure 1 is the hierarchy for the natural hazard threat; the assets are taken as a single target because the natural hazards cannot differentiate among the assets in terms of attractiveness. Figure 2 is the hierarchy for the on-site human threat; here the threat agent is able to differentiate among the assets and, hence, is able to mount a selective attack.

A safeguards system comprises a set of mechanisms protecting the assets from the threats. Such a system includes physical measures such as guards, fences, and dogs; mechanical things such as locks and security cradles; procedural controls such as rules, guidelines, and standard operating procedures (SOPs); electronic devices such as monitors, sensors, alarms, and closed-circuit television (CCTV); and technical measures such as passwords, intelligent cards, shielding, and so forth. Safeguards can be thought of as specific controls and countermeasures relative to each [threat, asset, safeguards function, outcome] quadruple.

Safeguards functions perform the objective of protecting the assets from the threats. We define the safeguards function space relative to the [threat, asset, outcome] triples wherein each [threat, asset] pair has a set of safeguards functions whose purpose is to thwart the threat-asset interactions and prevent the undesirable outcomes from occurring. This relationship can be represented as

$$\underline{F}(\underline{T}x\underline{A}) = \{F_{ij1}, F_{ij2}, \ldots, F_{ijk}\} ,$$

where each of the k safeguards functions for the ith asset and jth threat is performed by a set of safeguards subfunctions that can be represented as

$$F_{ijk} <=> U_m \, (f_{ijkm}) ,$$

in which the symbol <=> means "is defined to be" and the symbol U_m is the union over m.

Each safeguards subfunction is made up of elements that determine performance adequacy and can be represented as

$$f_{ijkm} = \Sigma_n \, (e_{ijkmn}) ,$$

and each element is further composed of attributes, α_{ijkmnp}, determining element completeness and additional information, ι_{ijkmnp}, about the elements and the attributes, represented as

$$e_{ijkmn} = \sum_p (\alpha_{ijkmnp}, \iota_{ijkmnp}) \cdot$$

The elements and attributes are, in fact, specific safeguards or countermeasures.

An outcome possibility matrix[13] relates the [threat, asset] pairs and the outcome set. One can think of the outcome possibility matrix as a fuzzy matrix. The values in the matrix represent the degree of possibility that the outcome in question could occur as a result of an interaction of the specific [threat, asset] pair. One can obtain the values from a fairly complex analysis, or one can simply assume that the outcomes are all either possible (a possibility degree of unity) or not possible for the potential threat-asset interactions (a possibility degree of zero). For most applications, the simpler case is more than adequate. Figure 3 shows the outcome possibility matrix for the computer-security application example we have been using throughout this paper.

	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
On-site Human – Facility	1	1	1	1	1	1
On-site Human – Hardware	1	1	1	1	1	1
On-site Human – Software	1	1	1	1	1	1
On-site Human – Documents/ Displays	1	1	1	1	1	1

Fig. 3
Outcome possibility matrix.

Each outcome for each [threat, asset, safeguards function] triple can occur in varying degrees of severity. Outcome severity is a function of the relative weakness of the safeguards function (a performance measure of the safeguards function) and the relative strength of the threat. In the case where the dynamic threat is not analyzed separately, the threat strengths are assumed to be equal and unity. Outcome severity is expressed as a degree of membership in the fuzzy set of outcome severity, but it can also be translated into a linguistic descriptor to be used with the non-monetary measures for impact and loss exposure.

We also define a set of impacts to measure the effect that an outcome of some specific degree of severity would have upon the subject organization. where the outcome resulted from some threat-asset-vulnerability interaction. Impacts are given as pairs of values, one expressed in monetary terms for those components of impact that can be measured in financial terms, and one expressed in non-monetary or linguistic terms for those components of impact that are better measured another way. The impact set is represented as

$$\underline{I}_{ijkq} = \{I(M)_{ijkq}, I(L)_{ijkq}\} ,$$

which defines the consequences for the [threat, asset, safeguards function, outcome] quadruples.

Risk, or potential loss exposure, results from the interaction of three components: the threat component, measuring the relative potential strength of threat agents in producing a specific outcome by exploiting a safeguards function vulnerability; the vulnerability component, measuring the relative potential weaknesses in the safeguards functions; and the consequence component, measuring the relative potential costs of specific outcomes. A pair of loss exposures is given for each [threat, asset, safeguards function, outcome] quadruple, and is expressed

$$\underline{R}_{ijkq} = \{R(M)_{ijkq}, R(L)_{ijkq}\} ,$$

where

$$R(M)_{ijkq} = f\ (V_{ijk}, S_{ijk}, O_{ijkq}, I(M)_{ijkq})$$

$$R(L)_{ijkq} = f\ (V_{ijk}, S_{ijk}, O_{ijkq}, I(L)_{ijkq}) ,$$

in which V_{ijk} is a measure of the relative weakness of the kth safeguards function for the interaction of the ith asset and the jth threat; S_{ijk} is a measure of the relative strength of the jth threat against the kth safeguards function for the ith asset; O_{ijkq} is (usually) 1 or 0, indicating that the qth outcome can occur for the interaction of the ith asset, jth threat, and kth safeguards function; and $I(M)_{ijkq}$ and $I(L)_{ijkq}$ are the monetary and non-monetary impact measures for the [threat, asset, safeguards function, outcome] quadruple.

The risk pairs can be aggregated to whatever level is desired. More aggregation provides a "bottom line" for upper management, along with a coarse indication of where the bottom line came from. However, less aggregation provides more detailed and specific information to those whose job it is to improve security and the overall risk posture. In general, we think that more aggregation has a tendency to lose information and to smear out important results; hence, we think that less aggregation is better for most purposes.

B. The information-gathering phase

The information-gathering phase acquires the data for LAVA to operate upon by means of automated, interactive questionnaires. Information is collected about the organization's mission, its assets and the potential threats to its assets, its environment, the safeguards (or control) system it uses to protect its assets from the threats, and its value and preference structures. Interaction with the questionnaires is accomplished in a team environment: the safeguards vulnerability questionnaire is executed by management, operations, and security personnel; the dynamic threat questionnaire is executed by a group of people having access to the appropriate information; and the impact questionnaire is executed by a team made up of high-level management and operations personnel.

The subject organization's mission determines the necessary security level the safeguards system must achieve for adequate protection of the assets. This is essentially an assessment of the mission's sensitivity, its criticality, and its integrity requirements.

Many factors contribute to the organization's environment. Those having the most effect upon the analysis are geographical location, community environment, physical environment, and procedural environment. Geographical location determines, to a large extent, the potential for catastrophic natural events such as earthquakes, volcanic eruptions, hurricanes, tornadoes, floods from swollen rivers or burst dams, and so forth; it also indicates nearness to population centers, major transportation hubs, a ready source of spare parts, and so forth. Community environment describes the social, political, and intellectual climate in which the organization finds itself, including such factors as the presence of organized crime, political dissent, universities, and social and moral arbitrators. Physical environment describes the campus of the organization--the land, fences, buildings, and so forth. Procedural environment details the philosophy, policy, and procedures set forth by the organization's management.

The organization's value structure determines what the effect a successful attack might have upon the organization. Consider, for example, a widely used computer, Brand X, whose cost is roughly half a million dollars. If the organization's sole computing power is this single computer, and if the organization depends heavily upon the computer for carrying out its daily business, the destruction of the machine would have a catastrophic effect upon the organization. If, on the other hand, the organization has a great many of these computers as well as several more powerful supercomputers and depends only marginally upon any one of the Brand X computers, the destruction of one of them, while unpleasant, would be far from catastrophic and indeed may be only inconsequential or a nuisance.

C. The analysis phase

The analysis phase consists of potentially three separate analyses--two required analyses and an optional analysis. The first analysis uses information gathered about the subject organization's safeguards system to assess system vulnerabilities. The optional second analysis, if it is needed, uses information gathered about the dynamic component of the threat to assess current threat strength. The third analysis uses information gathered about the subject organization's values, preferences, and philosophies to evaluate the potential consequences of successful attacks. From these analyses come a pair (monetary and non-monetary) of loss exposures for each [threat, asset, safeguards function, outcome] quadruple.

1. From the vulnerability analysis, a value is calculated for each safeguards subfunction that represents the subfunction's degree of membership in the fuzzy set of safeguards-function effectiveness; its fuzzy complement[14] represents the vulnerability (or relative weakness) of the subfunction. (Recall that the degree of membership in a fuzzy set lies in the interval [0,1], where a zero degree of membership implies no membership and a value of unity for degree of membership implies full or complete membership; a set consisting only of these extremes represents a special case of fuzzy set, which is an ordinary set.)

Each safeguards function is represented as a fuzzy tree, whose branches represent the subfunctions for this particular function. For example, if the safeguards function is "Fire Damage Control," then it follows that appropriate subfunctions would be "Fire Prevention Controls," "Fire Detection Controls," "Fire Emergency Administration Measures," and "Fire Damage Mitigation Controls." This example is depict 1 in Figure 4. Note that these trees are not probabilistic trees; the values that eventually appear on the branches are not probabilities but instead are degrees of membership in the fuzzy set of subfunction control performance. The subfunction values can be combined to produce a value for the safeguards function by taking their fuzzy union[14], essentially equivalent to the maximum of the subfunction values.

A desirable set of safeguards that will accomplish the safeguards functions' objectives is modeled as an automated interactive questionnaire that elicits specific information about the presence and quality of the safeguards existing at the subject site. Each element e_{ijkmn}, attribute a_{ijkmnp}, or added information contribution is represented as a separate question. Just as the subfunctions are more or less equal within a specific safeguards function, each element question within a specific subfunction is more or less equal Further, each of the attributes within the element they modify is more or less equal. Any element safeguard can help to accomplish the objectives of more than one safeguards subfunction (a complex relational database keeps track of the numerous interrelationships), but each attribute is linked to a specific element.

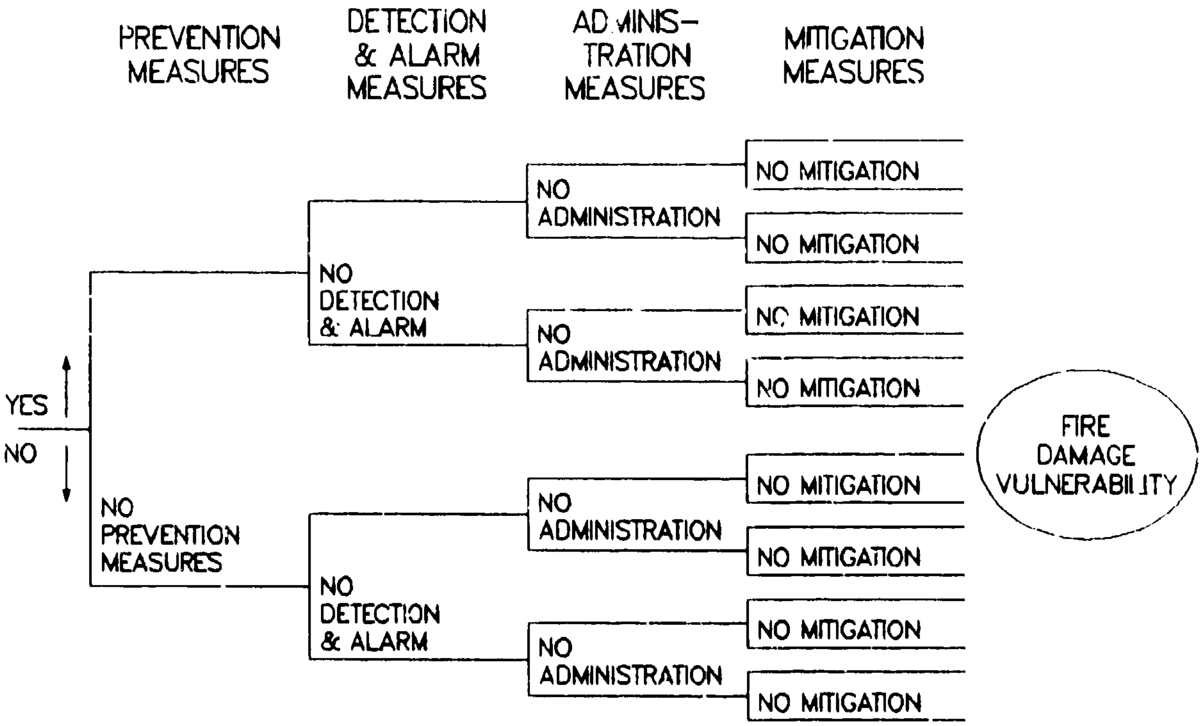

Fig. 4.
Fuzzy tree for Fire Damage Control Safeguards Function.

In scoring the safeguards system questionnaire, each safeguards element has a maximum potential degree of vulnerability equal to 1. If the element is one whose mere presence implies adequate performance, the vulnerability degree is 0 if the element is present and 1 if it is not; such an element does not have attributes. If the element has attributes, this means that not only should the safeguard implied by the element question be present but also the safeguard has associated criteria to determine element quality and performance adequacy (and hence degree of vulnerability). Each of the performance criteria (attributes) are of about equal importance and have a maximum vulnerability value of 1/p, where p is the number of attributes for the particular element under consideration. In this case, the element vulnerability value is given as

$$v(e_{ijkmn}) = \Sigma_p\, v(\alpha_{ijkmnp})$$

where the $v(\alpha_{ijkmnp})$ is either zero if the criterion expressed in the question is satisfied or 1/p if it is not.

Each safeguards element (and hence its associated attributes) can contribute to more than one safeguards subfunction. We use a relational database to keep track of which subfunctions are affected by each element, and after the safeguards questionnaire has been answered completely, the database is used to assist in the subfunction tabulations. The degree of vulnerability for each subfunction is the sum of the vulnerability values for all the elements that contribute to the subfunction, normalized to unity, or

$$v_{ijkm} = 1/n\ [\Sigma_n\ v(e_{ijkmn})]$$

where n is the number of safeguards elements contributing to the mth subfunction.

The vulnerability values for each of the subfunctions are interesting in themselves, but they also can be aggregated into vulnerability measures for the safeguards functions by taking their fuzzy union, obtained as the maximum of the values for the subfunctions that make up a safeguards function. This is the measure carried through into the loss exposure calculations.

The vulnerability component measures the relative weaknesses in the safeguards functions with respect to the spectrum of threat-asset pairs. In the vulnerability analysis, we assume that the threat is static (the dynamic component is zero), that all attacks are equally likely to occur, and that the consequences are extreme. The resulting vulnerability measure can later be made more realistic--it can be reduced or increased by including the dynamic threat and the "real" consequence measures when the risk measures are calculated. LAVA defines the vulnerability measure, V_i, as the complement of the membership function in the fuzzy set of "complete safeguards-function effectiveness," S_i (or $V_i = 1 - S_i$, where both V_i and S_i lie between zero and unity).

The LAVA software system includes an automated report generator that produces a report of the vulnerability analysis. The report includes rankings of the vulnerabilities by safeguards function and subfunction, a breakdown of vulnerabilities by threat category, and bar charts and scatter diagrams illustrating these. It includes a detailed report listing the missing or inadequate safeguards for each of the safeguards subfunctions, providing information to those whose mission is to reduce vulnerabilities.

2. A dynamic threat[11] analysis can be performed 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 takes into account possible threat agents and their potential attack goals. The magnitude of the threat is determined from the motivation, capabilities, and opportunities of the various threat agents with respect to the target(s) of the attack.

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) 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 threat agents in each of these categories all act for different reasons, and 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),
2) sabotage,
3) theft, embezzlement, fraud,
4) damage or destruction,
5) extortion,
6) disrupting business or mission,
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 threat component measures the relative strength of identifiable threat agents in terms of motivation, opportunity, and capability with respect to the spectrum of assets, the corresponding safeguards functions, and the set of possible outcomes. 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 tree structure for the dynamic threat analysis.

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 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 (motivation, capability, and opportunity) in the threat

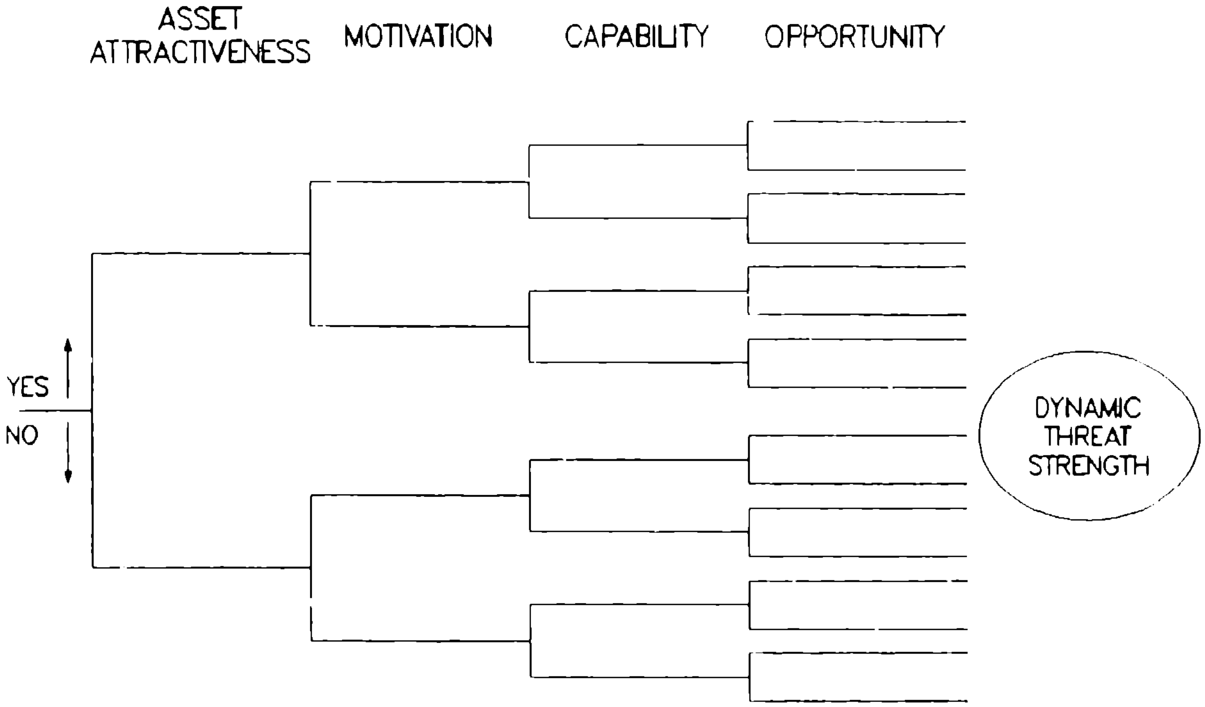

Fig. 5.
Assessment structure for .dyn. .ic threat.

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

3. Evaluating the consequences, or impacts, is the third analysis in LAVA. The object of this analysis is to determine the effect that an outcome of a successful attack would have upon the subject organization.

The consequence (or impact) component measures the potential costs (both monetary and non-monetary) of a threat successfully exploiting a safeguards function's vulnerability with respect to the severity of the exploitation's outcome. The outcome severity metric results from combining the results of the previous two analyses, the relative weakness of the safeguards function and the relative strength of the dynamic threat.

LAVA's consequence measures are given as pairs of monetary and nonmonetary descriptors: the monetary descriptor $I(M)_{ijkq}$ is used when the consequences can be given in terms of monetary cost (dollars, pesos, francs, pounds, kroner, etc.), and the non-monetary descriptor $I(L)_{ijkq}$ is used when the consequences can be given only in terms of intangible cost (such as "catastrophic" as a linguistic or non-monetary measure of irreparable reputation loss). The consequence values are obtained from another interactive questionnaire. Figures 6 and 7 illustrate the monetary and non-monetary consequence tree structure.

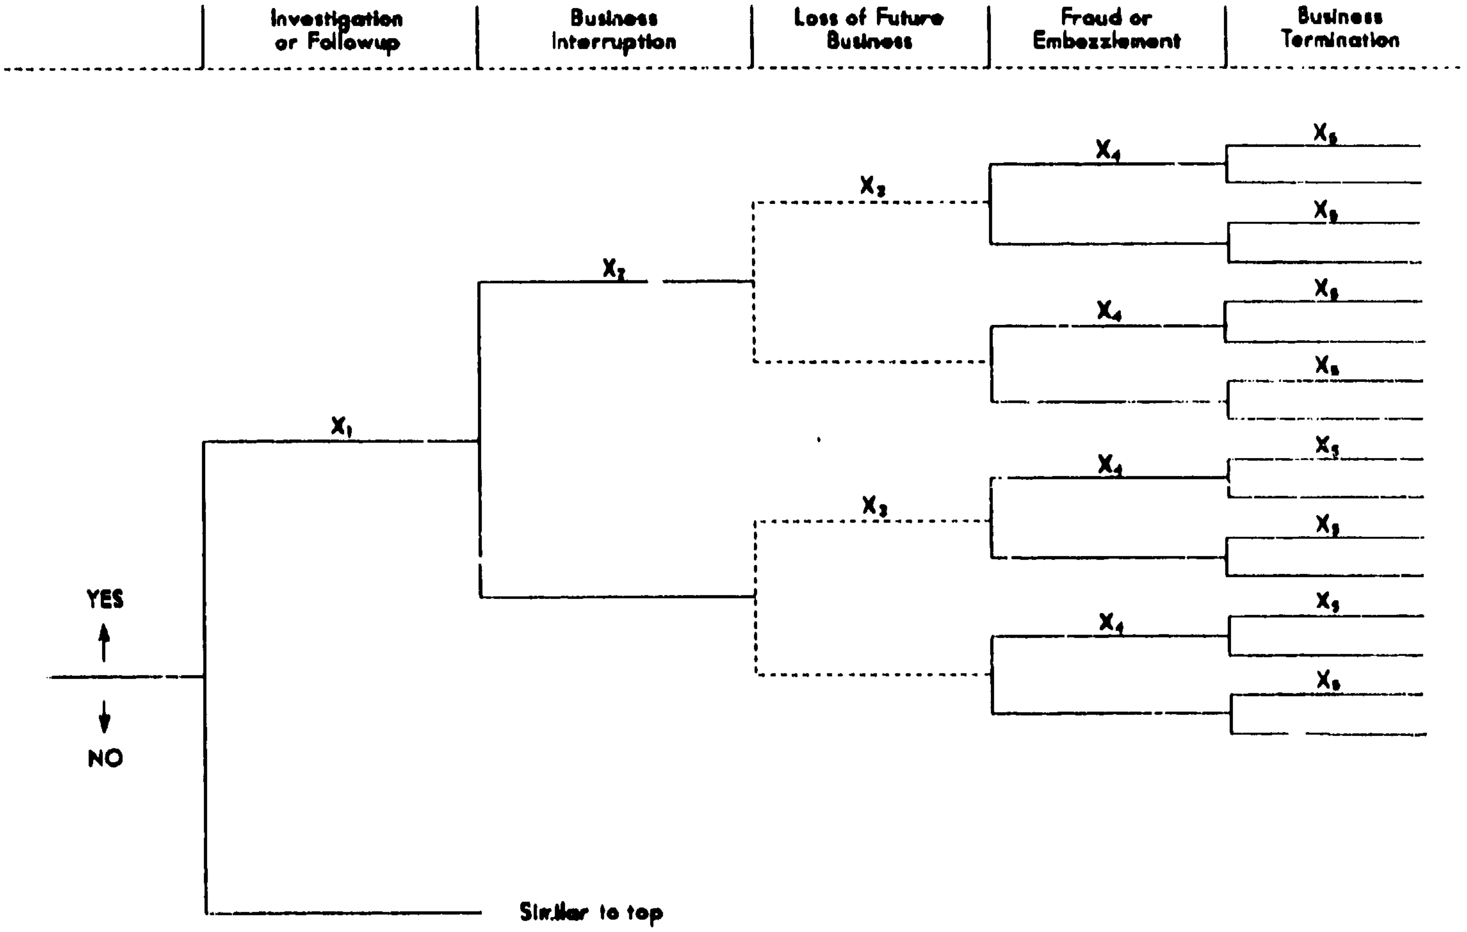

Let I_M represent the Monetary Impact Factor.
Then I_M is quantified by:

$$I_M = X_1 + X_2 + X_3 + X_4 + X_5$$

Fig. 6.
Monetary consequence tree.

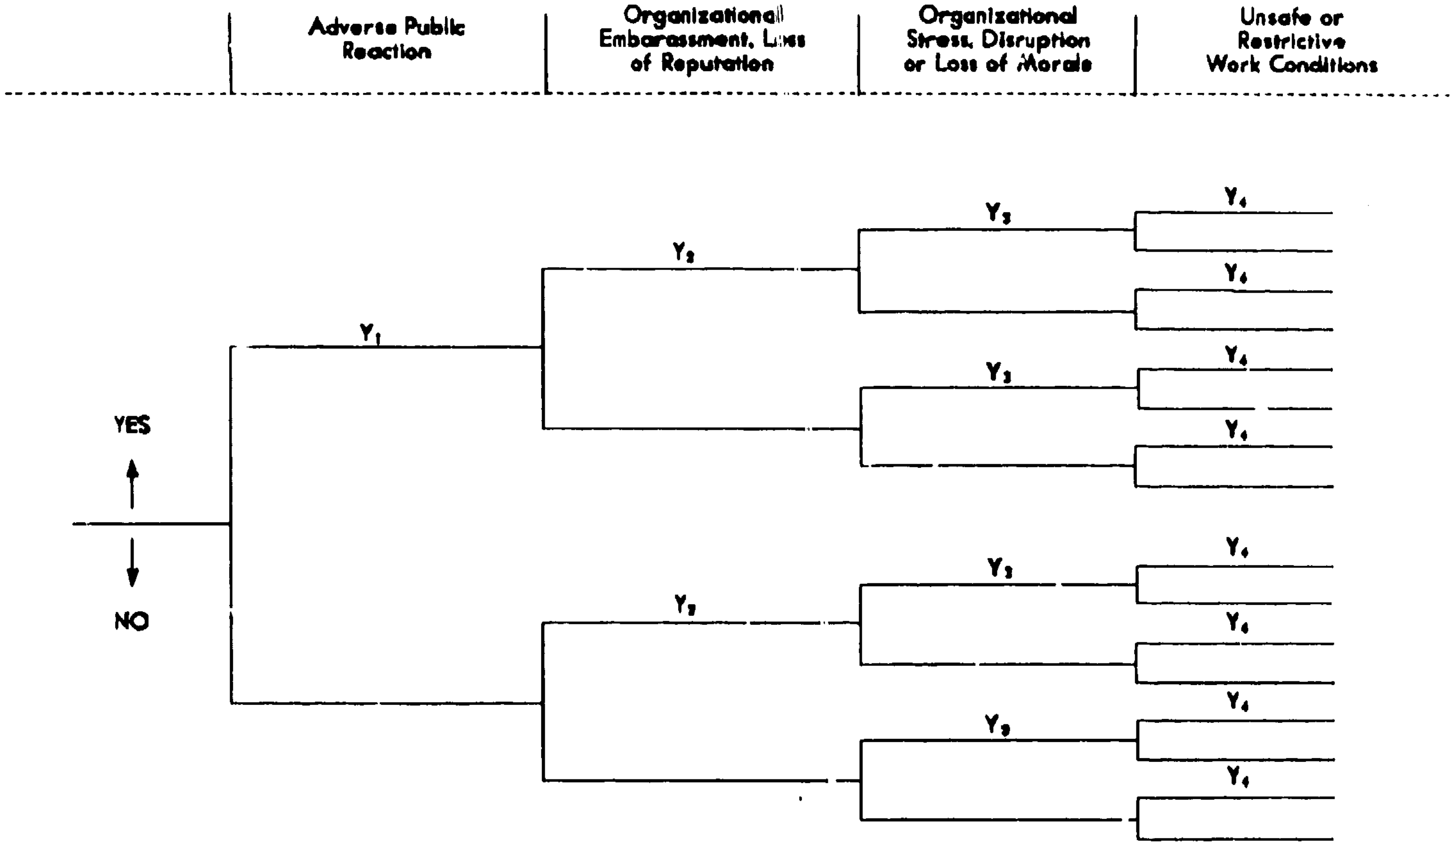
Adverse Public Reaction
Organizational Embarrassment, Loss of Reputation
Organizational Stress, Disruption or Loss of Morale
Unsafe or Restrictive Work Conditions
YES
NO
Y_1
Y_2
Y_3
Y_4

Let I_{NM} represent the Non-monetary Impact Factor.
Then I_{NM} is quantified by:

$I_{NM} = \max\ (Y_i)$, where $i = 1, 2, 3, 4$.

Fig. 7.
Non-monetary consequence tree.

The monetary consequence values come about as the sum of the costs of investigation and followup, business interruption, future business loss, fraud or embezzlement, and business termination. Each of these contributors is calculated from questions about specific costs, such as costs for reprimand or legal action; costs associated with lost time, interim operation, and full replacement, repair, and recovery; costs for employee replacement and training; direct losses from damage and destruction, waste, fraud, or embezzlement; and so forth.

The non-monetary (linguistic) consequence values are expressed as the maximum of the set of values derived for the cost of adverse public reaction, the cost of organizational disruption and/or loss of morale, the cost of embarrassment and/or potential reputation loss, and the intangible cost associated with unsafe or restrictive working conditions. The linguistic values may include, for example, such measures as "catastrophic," "nuisance," and so forth.

D. Solving the risk management problem

The best analysis does no good if appropriate action is not taken to improve the organization's risk posture. A thorough cost/benefit analysis provides a basis for optimally selecting a safeguards set to be improved or added to the existing safeguards system. The analysis not only should include how much the proposed action path will cost and how much it will reduce loss exposure, but also should include well-thought-out implementation plans and the effect that the implementation will have on the organization in terms of morale, training (or re-training), throughput reduction, and so forth.

Our plans for LAVA's future include automating the cost/benefit analysis and the optimal safeguards selection process. However, at this time, these features are left to the subject organization's analysts and management.

IV. CONCLUSIONS

We have discussed LAVA, a methodology for evaluating vulnerabilities, threats, and risks in large, complex systems that often are too intractable to be handled adequately and reasonably by other methods. LAVA's measures for loss exposure (risk) for each [threat, asset, safeguards function, outcome] quadruple are functions of the relative weaknesses of the safeguards function (vulnerability), the relative strength of the threat, and the consequences of the outcome of a successful attack.

The baseline vulnerability measures are derived with the assumption that the background (static) threats are the only operable threats (that is, their relative strength is 1). In evaluating the severity of the outcome of a successful attack when a vulnerability exists in a safeguards function, the effect of the vulnerability can be increased by the "real" threat measure if the information necessary to the threat analysis is accessible and available.

Our automated procedure can be used to elicit and analyze information to determine vulnerabilities and risks inherent in many application systems, such as computer systems, material control and accounting systems, physical protection systems, security systems, plant process-control systems, and a host of others. The subject system is modeled as interactive questionnaires that include inherent *a priori* decisions about the desirability of or aversion to applications-specific features of the system, about the subject system's threat environment, and about the value structure of the subject organization. Event-tree-like structures are used for organizing the analyses. Performance measures and associated decisions are evaluated in both qualitative and quantitative terms, giving greater insights into system performance than a strictly quantitative analysis can.

LAVA's approach is technically sound, accurate, interactive, secure, and portable. The accuracy is derived from exhaustive and comprehensive questions provided by the developer of a specific application. The interactive conversational questionnaire in natural language is straightforward and easy to use. The expertise and rigor built into the model, as well as the model's logical structure, makes decision-making simple. The functional structure of the model clearly indicates what safeguards are missing and provides a rational for selecting the safeguards to add. The LAVA software system is compatible with standard IBM-PC software, making the system portable. Because it is interactive and self-contained, the organization using this methodology and the associated software does not require the services of outside consultants, hence adding another layer of security to the entire safeguards system.

V. REFERENCES

1. 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).

2. S. Sudman and N. M. Bradburn, Asking Questions: A Practical Guide to Questionnaire Design (Jossey-Bass, Inc., San Francisco, CA, 1982).

3. S. T. Smith and J. J. Lim, "An Automated Method for Analyzing Computer Security Risk," Proc. 7th DOE Computer Security Group Conference, New Orleans, LA, April 17-19, 1984, Los Alamos National Laboratory document LA-UR-84-438.

4. S. T. Smith and J. J. Lim, "An Automated Method for Assessing the Effectiveness of Computer Security Safeguards," Proc. IFIPS Second International Congress on Computer Security, Toronto, Canada, September 10-12, 1984, Los Alamos National Laboratory document LA-UR-84-2744.

5. S. T. Smith and J. J. Lim, "An Automated Interactive Expert System for Evaluating the Effectiveness of Computer Security Measures," Proc. 7th Department of Defense/National Bureau of Standards Computer Security Conference, Gaithersburg, MD, September 24-26, 1984, Los Alamos National Laboratory document LA-UR-84-1580.

6. S. T. Smith and J. J. Lim, "Framework for Generating Expert Systems to Perform Computer Security Risk Analysis," Proc. First Annual Armed Forces Communications and Electronics Association Symposium and Exposition on Physical and Electronics Security, Philadelphia, PA, August 19-21, 1985, Los Alamos National Laboratory document LA-UR-85-1933.

7. S. T. Smith and J. J. Lim, "Assessment of Computer Security Effectiveness for Safe Plant Operation," Proc. American Nuclear Society Annual Meeting, New Orleans, LA, June 3-8, 1984) Los Alamos National Laboratory document LA-UR-84-99.

8. S. T. Smith and J. J. Lim, "An Automated Procedure for Performing Computer Security Risk Analysis," Proceedings Sixth Annual Symposium on Safeguards and Nuclear Material Management (European Safeguards Research and Development Association, Joint Research Centre, Ispra, Italy, 1984), ESARDA 17, pp. 527-530.

9. S. T. Smith, D. C. Brown, T. H. Erkkila, P. D. FitzGerald, J. J. Lim, L. Masasgli, J. R. Phillips, and R. M. Tisinger, "LAVA - A Conceptual Framework for Automated Risk Assessment," Nucl. Mater. Manage. XV (Proceedings Issue), 256-259 (1986).

10. S. T. Smith, J. J. Lim, and J. Lobel, "Application of Risk Assessment Methodology to Transborder Data Flow," Handbook on the International Information Economy, Transnational Data Report, Springfield, VA LA-UR-85-2416 (November 1985).

11. S. T. Smith, J. R. Phillips, D. C. Brown, and P. D. FitzGerald, "Assessing the Threat Component for the LAVA Risk Management Methodology," Proc. 9th DOE Computer Security Group Conference, Las Vegas, NV, May 6-8, 1986), Los Alamos National Laboratory document LA-UR-86-741.

12. M. D. Mesarovic, D. Macks, and Y. Takahara, Theory of Hierarchical Multilevel Systems (Academic Press, New York and London, 1970).

13. K. J. Schmucker, Fuzzy Sets, Natural Language Computations, and Risk Analysis (Computer Science Press, Rockville, MD, 1984).

14. L. A. Zadeh, K.-S. Fu, K. Tanaka, and M. Shimura (Eds.), Fuzzy Sets and Tneir Applications to Cognitive and Decision Processes (Academic Press, New York, 1975), pp. 1-39.