

CONF-9606221-7

SAND96-1541C

## INFORMATION SYSTEMS VULNERABILITY: A SYSTEMS ANALYSIS PERSPECTIVE\*

Gregory D. Wyss, Sharon L. Daniel, Heather K. Schriner  
Risk Assessment & Systems Modeling Dept.

and Timothy R. Gaylor  
Data Transport & Network Design Dept.

Sandia National Laboratories  
Albuquerque, NM 87185-0747  
Phone: (505) 844-5893 Fax: (505) 844-3321  
E-mail: gdwyss@sandia.gov

RECEIVED

JUL 25 1996

OSTI

### ABSTRACT

Vulnerability analyses for information systems are complicated because the systems are often geographically distributed. Sandia National Laboratories has assembled an interdisciplinary team to explore the applicability of probabilistic logic modeling (PLM) techniques (including vulnerability and vital area analysis) to examine the risks associated with networked information systems. We have found that the reliability and failure modes of many network technologies can be effectively assessed using fault trees and other PLM methods. The results of these models are compatible with an expanded set of vital area analysis techniques that can model both physical locations and virtual (logical) locations to identify both categories of vital areas simultaneously. These results can also be used with optimization techniques to direct the analyst toward the most cost-effective security solution.

### I. BACKGROUND

Information systems security methods have advanced considerably over the last decade, yet many field implementations of systems security are still based on a "checklist mentality" that believes that following a set of documented "best security practices" will guarantee security for *any* information system. In blindly following a checklist, an information systems manager may fail to

recognize special features of the facility that will render a typical "best practice" ineffective. In contrast, physical security for high value sites has historically been designed based on vulnerability analyses and vital area analyses. Vulnerability analyses seek to identify all sequences of events that can place a system in an undesired state. They also seek to identify which of these events could be caused by the deliberate action of a saboteur. These analyses often use probabilistic logic models (PLMs) to develop the most complete lists of vulnerabilities possible. Vital area analyses associate the identified vulnerabilities with specific locations in order to obtain the list of areas that would have to be accessed by a saboteur in order to accomplish an attack on the system. From this list it is a simple mathematical task to generate the list of locations that must be protected in order to prevent an attack from being successful.

While PLMs have been commonly used in many industries, their use in the telecommunications industry has been fairly limited. The complex topologies of communications networks, the time-dependent interactions between network elements,<sup>1</sup> and the geographically distributed nature of many information systems have made it difficult to model these systems with the fault tree, event tree, influence diagram, and reliability block diagram modeling techniques that have proven so successful in other industries. An interdisciplinary team at Sandia National Laboratories has developed a number of specialized modeling techniques that are specifically designed to enable efficient modeling of networks, and network services, for vulnerability analyses. The results of these models are compatible with an expanded set of vital area analysis techniques that can model both physical locations and virtual (logical) locations to identify both categories of vital areas simultaneously. These results can also be used with optimization techniques

\* This work was performed under the Laboratory-Directed Research and Development Program at Sandia National Laboratories. Sandia National Laboratories is operated by Sandia Corporation, a Lockheed Martin Company, for the U.S. Department of Energy under contract DE-AC04-94AL85000.

#### **DISCLAIMER**

**Portions of this document may be illegible  
in electronic image products. Images are  
produced from the best available original  
document.**

to direct the analyst toward the most cost-effective security solution. The resulting vulnerability models can provide valuable quantitative decision support during both the design and operational phases of an information system.

## II. BENEFITS OF PROBABILISTIC LOGIC MODELS

PLMs have been used by a number of different disciplines including quantitative reliability analysis (QRA), probabilistic risk analysis (PRA), and probabilistic safety analysis (PSA). Regardless of the discipline, the reasons for developing a PLM are the same:

- to identify an exhaustive list of the modes by which a system can fail,
- to find an approximate frequency with which we might expect to observe failures, and
- to determine a rank ordering of the components in the system by their "importance" to the proper function of the system.

The "importance" of a component can be defined in a number of ways, but is often thought of as answering one of the following questions:

- How sensitive is the overall system reliability to changes in the reliability of an individual component?
- If the reliability of this component is allowed to decrease (say, by using components of lesser quality), how much will this effect overall system reliability?
- If money is invested to increase the reliability of this component, how much will this effect overall system reliability?

Clearly the answers to these questions cut to the heart of the question of how data networks are designed and managed. For example, a PLM analysis might show that a particular network hub or concentrator does not contribute significantly to the unreliability of the system, but that it would become a significant contributor if its reliability were allowed to deteriorate. The analysis might also show that, while a particular router seems to be a major contributor to system unreliability, the expense that would be incurred to replace it might be more effectively spent pursuing a number of smaller and less expensive upgrades. It might also show the opposite. PLM results should not be used as the exclusive basis for design and upgrade decisions because such decisions have intangible aspects that must also be considered. However, PLM results do provide *quantitative* answers to network reliability questions, and these quantitative answers can be used as a legitimate

benchmark to get past the "gut feeling" that unfortunately forms the basis for many network design and upgrade decisions. It has also been demonstrated that PLM results are well suited for use in discrete optimization algorithms such as genetic optimization.<sup>2,3</sup>

With all of the good decision support information that comes from PLM models, it is sometimes tempting for the uninitiated to view PLM as some sort of a "silver bullet" that makes traditional forms of network analysis such as dynamic simulation obsolete. This is most certainly *not* the case. PLM and dynamic simulation should be viewed as complementary tools which, when used together, provide a more complete view of network performance than either can provide by itself. For example, dynamic simulations are often very computationally expensive, so it will not be possible to run a simulation for each network variation that might be of interest. The results of a PLM analysis can provide insights to help direct the simulation analyst to the most important variations so that they can get the most valuable information for the computational effort expended. On the other hand, the results of direct simulation analyses will help a PLM analyst to be sure that they have properly established important success criteria within their model. PLM provides a global view of the network and quantitatively leads a designer to options for its betterment, while direct simulation provides detailed information about critical situations within a particular network configuration. Clearly both perspectives are necessary for a complete understanding of the network.

At this point someone usually comments, "You speak of quantitative results, but I have no data. Surely the value of your results cannot be any better than the quality of your data, so how can this be of any benefit to me?" That statement is true *if* you are seeking to predict the absolute reliability of the system (e.g., mean time between failures). However, the most useful information from a PLM is usually the rank ordering of components by importance. An accurate rank ordering can be achieved even with relatively little measured reliability data. An analyst can often state with relatively high confidence that component A is "somewhat more likely to fail" than component B, or that a router with internal redundancy would be expected to be "much more reliable" than a workstation. The analyst can create groups of components and failure modes such that all elements in the group have similar failure rates, and then rank these groups to obtain a reasonably accurate set of *relative* reliability data. The rank-ordered results from a PLM are accurate even with only relative data. Thus, it is possible to obtain some of the most useful results from a PLM even in the absence of a great deal of measured reliability data.

### III. SYSTEM VULNERABILITY ANALYSIS

A vulnerability analysis seeks to identify how an information system can be forced into an undesired state. This undesired state may consist of an unintended disclosure of sensitive information, improper alteration of either information or system configuration, or a denial of system services (e.g., destroying network connectivity or denying access to particular information or information processing capabilities). The undesired state may be achieved due to conditions within individual information processing entities, network failures, or combinations thereof. Thus, we must examine each of these areas if we are to obtain a complete picture of system vulnerabilities.

Recall that our objective is to identify all combinations of events that can place an information system in an undesired state. Each individual combination of events that is *sufficient* to place the system in an undesired state is called a *cut set*. If each event in the cut set is also *necessary* in order for system failure to be achieved, then the cut set is said to be *minimal* (its failures are both necessary and sufficient to cause system failure). In other words, a cut set is non-minimal if the undesired system state can be achieved with the occurrence of some subset of the events in the cut set. Each of the modeling methods described below produce cut sets as results. Regardless of how it is generated, the complete list of minimal cut sets theoretically represents all of the possible ways that primary events can combine to cause system failure. Practically speaking, there are often far too many minimal cut sets for an analyst to readily examine, so the cut sets are ranked by size and/or probability, and those cut sets with the lowest rank are eliminated.

The complete list of cut sets contains a great deal of information about the system being modeled. Quantifying this list provides the overall probability of system failure. A ranking of the cut sets by probability shows the most likely failure scenarios for the system. A designer can use this information to design system improvements that remove the most likely failure scenarios. However, there is much more information buried in this list of cut sets. A simple mathematical transformation of the cut sets provides the "importance measures" described previously. The partial derivative of this list with respect to each primary event shows how quickly the reliability of a system will change given variations in reliability of each individual component. Setting a primary event's failure probability to 1.0 shows how the reliability of the system will be affected if a very low quality component is used for this function. Finally, setting a primary event's failure probability to 0.0 shows the maximum reliability improvement that could be obtained by "fixing" this component. If this value is large, then it may be appropriate to invest money to improve this

component. Thus, the list of cut sets is the key that unlocks all of the other valuable information that can be found through fault tree analysis (FTA).<sup>4</sup>

#### A. Information Processor Models

An information processor can be vulnerable to two different classes of events: static events, and dynamic events. Static events are defined as events that need not occur in a particular order to cause problems for the information processor. Examples of static events include power failures, cooling system failures, building fires, hardware failures within the processor (such as disk failures), and certain operator actions. Static events that can cause failure of the information processor are adequately modeled using fault tree analysis techniques. Cut sets are produced when the fault tree model is solved.

Dynamic events, however, are defined as events which only impact the system if they occur in a particular order or within a constrained time window. Examples of dynamic events include processor saturation, interrupt-driven operating system conflicts, and certain operator actions. Traditional fault trees do not adequately capture the time-dependent nature of dynamic events.<sup>5,6</sup> Thus, dynamic tools such as influence diagrams or event trees must be used to identify dynamic events that can cause failure of the information processor. While these types of models may not directly produce cut sets when solved, techniques such as variable transformation and event pairing can be used to transform the results into cut sets. In addition, since the presence of static events can influence the effectiveness of individual dynamic events, fault trees can be used to support the dynamic event analysis techniques.

While the techniques for developing dynamic event models still require considerable PLM expertise, Sandia has demonstrated a method by which fault tree analysis techniques can be made accessible to analysts who are at most casual fault tree analysts. Our objective was to develop a methodology that would allow an analyst to construct a fault tree model by simply "plugging together" model elements that represent easily identified generic components within the information processor. Under this "plug-and-play" modeling technique, an expert constructs a library of generic fault tree "modules" to represent the failure modes of typical generic components.<sup>7,8</sup> A casual analyst can then "plug" these modules together to quickly form a complete fault tree model for an information processor. There are a number of advantages to this approach. By creating fault trees for each generic component which can be combined together to model an overall processor, initial fault tree models can be constructed quickly and efficiently. In addition, new equipment configurations can also be easily modeled.

## B. Network Device Models

There are two types of network devices that must be modeled: passive and active devices. A passive network device transmits network traffic without assessing or transforming its contents. The most obvious example of a passive network device is network cabling, although other devices such as line amplifiers also fit into this category. Passive devices are typically the most reliable parts of a network, possessing only very simple failure modes (and, hence, fault trees). It is, however, important that passive devices not be neglected in the vulnerability model because they can represent single points of network failure, and they often travel through locations that are beyond our surveillance or control.

An active network device is in reality an information processor, albeit typically an information processor with limited capabilities. Thus, its vulnerability models are in most respects similar to, or a subset of, those that are described previously for information processors. However, since these devices are dedicated to network operations, there can be important differences related to dealing with multiple sources of network traffic simultaneously and the collisions that this traffic can cause.<sup>9</sup>

## C. Network Architecture Models

The previous two sections have dealt with fault conditions in individual components. However, the failure of any individual component may not by itself place the information system in an undesired state if there is sufficient redundancy. Therefore, it is important to model how the components are interconnected to form a network, and how the connectivity of the network can be destroyed, in order to understand how the overall information system can be adversely affected.<sup>10</sup>

Traditional computer data networks are constructed hierarchically — the network address space and/or the physical structure of the network architecture enforces a hierarchy. Many current generation telephone voice/data networks as well as Asynchronous Transfer Mode (ATM) data networks are often deployed in “flat” topologies — they are “arbitrarily interconnected” and have no enforced hierarchy. These differences in network topology require different cut set analysis methods.

**1. Hierarchical Networks.** In a hierarchical network there are usually only a few paths from one node to another. Because such networks contain few redundancies, they are less expensive to construct and easier to manage than their non-hierarchical cousins. In addition, some non-hierarchical networks exhibit characteristics that make them *behave* almost hierarchically. For example, although “911”

emergency services are provided on the non-hierarchical public telephone network, these services often behave hierarchically, with the top of the hierarchy being the public service answering point. While it is convenient to assume a rigid hierarchy within the network, the method can accommodate a limited number of “cross-cuts” through the hierarchy without becoming overly burdensome.

FTA works well to provide cut sets for hierarchical networks. Just as fault tree modules can be developed for individual components, it is also possible to develop fault tree modules for particular classes of network architectures (e.g., ethernet, token ring, and FDDI sub-network architectures) that are compatible with the “plug-and-play” fault tree methodology described previously. This enables a person with network experience but little FTA experience to successfully model most hierarchical network architectures.

Qualitatively, the user of device A will perceive that the network has failed whenever they cannot communicate with any needed device B either within their own workgroup (local connectivity) or in other parts of the network (global connectivity). The user will also perceive that the network has failed if a needed network service is unavailable for a significant amount of time. These qualitative observations by users can be used as the basis for the definition of a successful network. The top event in the “plug-and-play” fault tree then models success by the logical condition that every node must be able to communicate with and through the top of the network hierarchy (connectivity), and that network services are available.

The development of the “plug-and-play” fault tree begins by picking the component at the highest point in the hierarchy and developing the fault tree top event as described above. We then “reach out” from this component toward the bottom of the hierarchy by attaching the generic fault tree modules for each component and sub-network architecture found along the way. Thus, any components or sub-network architectures that are directly connected to the top of the network hierarchy are modeled by substituting or “plugging in” the appropriate generic fault tree module into the appropriate branch of the component’s fault tree. These newly modeled components are then examined to determine the components and architectures that are attached to them. As each new component or architecture is identified, its generic fault tree module is “plugged into” the appropriate branches of the emerging fault tree “stem”. This process continues until the entire network has been modeled. Once the entire network is included in the fault tree model, any remaining unused fault tree module branches are simply trimmed off because they represent network attachment options that were not exercised in the current network architecture. At

At this point the fault tree connectivity model is complete and can be analyzed for cut sets and component importance using existing risk analysis software.<sup>11,12</sup>

Note that the fault tree development process can be broken off before *all* elements are incorporated in the model if the analyst is interested in modeling the characteristics of only a specific portion of the network (say, the network backbone). The analyst can then extend this fault tree model to successively lower levels in the hierarchy without any loss of information by simply reviving any branches that may have been "trimmed" and continuing to apply the "plug-and-play" methodology as described previously. The fault tree paradigm naturally supports this concept of a high-level "quick look" followed by iterative model refinement. Since the model can be evaluated at any level of detail, it can provide a relatively inexpensive method for investigating high-level questions about the network. It also provides a cost effective way to play "What if?" games on early network designs as the network designer experiments with different ways to provide maximum reliability for the user community.

The fault tree model can be extended to model both network connectivity and network services if the fault tree top event is modified to reflect the following success criteria: the network-based information system is successful only if global connectivity is maintained *and* servers are available to provide all necessary network services (thus, from the user's perspective, the network-based information system fails if either the server *or* network connectivity fails). In a network where a single server provides all network services, the applicable fault tree model simply consists of a logical OR condition of the availability of the server machine and the network connectivity model we developed previously. For more advanced networks with multiple and possibly redundant servers, the single server in the OR condition would be replaced by a logical model (likely a small fault tree) that examines the combinations of server machines that must be functional in order for all network services to be available. This fault tree is usually easy to construct given the network specifications.

**2. Non-Hierarchical Networks.** Non-hierarchical networks have no enforced physical or logical hierarchy. They are often designed with a high degree of redundancy, so there may be many paths from one node to another. This makes these networks well-suited for use in areas where a high degree of reliability is important. This redundancy is, however, expensive to install and can be difficult to manage.

Fault tree models become extremely difficult to construct for non-hierarchical networks for two reasons: the presence of a large number of "cross-cuts" makes the construction of

an individual fault tree extremely difficult, and modeling global connectivity ("everyone can talk to everyone") can require the construction of many fault trees because of the absence of a defined hierarchy. Previous approaches to modeling non-hierarchical networks have focused on path set theory.<sup>13,14</sup> It is, however, very computationally expensive to obtain cut sets from path sets (path sets cannot provide the importance information that can be derived from cut sets), and the global connectivity condition can only be modeled by considering every possible pairwise combination of network endpoints — also very computationally expensive.

Because of the limitations of these commonly used techniques, Sandia National Laboratories has developed an efficient search algorithm to find the global connectivity cut sets for arbitrarily interconnected networks.<sup>15</sup> This method determines cut sets based on the network connectivity diagram, so there is no need to construct and maintain a separate reliability model. The algorithm takes advantage of a number of architectural and mathematical properties to reduce the computational effort required to obtain global connectivity cut sets for these networks. These cut sets can then be mathematically combined into the OR condition described in the previous section to obtain cut sets that consider network services.

#### IV. VITAL AREA ANALYSIS FOR INFORMATION SYSTEMS

To this point our analysis has focused only on events that must occur or equipment that must fail in order for our network-based information system to be placed into an undesired state. We have obtained cut sets, and each cut set represents one "scenario" — a set of conditions that must occur in order to achieve the undesired state. These conditions are both necessary and sufficient — in other words, if they all occur, then the undesired state is achieved. However, if even one of the conditions is not realized, then the undesired state is prevented. To a potential saboteur, each cut set represents a set of tasks that would have to be performed in order to do damage to this information system. A rudimentary security analysis, then, would examine the list of cut sets to determine which of these scenarios would be easiest for an adversary to accomplish, and consider what countermeasures might be employed to thwart the attack (*i.e.*, prevent any one of the events in the cut set from occurring). Since this approach considers one cut set at a time, it results in a piecemeal approach to security. However, with some additional processing of the cut sets, one can take a more systematic approach to the security problem.<sup>16</sup>

Since we live in a physical world, every action has to take place in a physical location. The tasks a saboteur would

have to accomplish are no different. To accomplish a task (as defined by an event in a cut set), the saboteur would have to gain access to a location where that task could be carried out. An event such as "remove the computer's hard disk" can only be accomplished in one location, while an event such as "cut the communications cable" can be accomplished in any of several locations (e.g., unplug it at either end, cut it in the wiring closet, or damage its conduit as it runs between buildings). Therefore, the first task in extending a risk and reliability analysis to be a vital area analysis consists of determining the complete set of locations from which each event can be carried out. As this is an information system, we must be careful to consider both physical locations (e.g., Room 222) and virtual locations (e.g., the Internet) in our assessment.

The next task is to combine the lists of locations with the list of cut sets. This is accomplished as a mathematical transformation of the cut sets by substituting for each event in each cut set the list of locations from which that event can be accomplished.<sup>17</sup> This will provide us with "location cut sets." A location cut set says, in effect, that our information system can be forced into an undesired state if a saboteur can gain access to all of the locations found in this location cut set. However, if we can prevent him from gaining access to even one of these locations, then we can prevent the saboteur from exploiting the scenario represented by this cut set. Since a single event may be related to more than one location, a system cut set that contains such an event will become more than one location cut set.

At this point, the group of location cut sets likely contains many redundancies that must be removed. These redundancies fall into two classes: redundancies within a particular cut set, and redundancies between cut sets. A redundancy within a cut set occurs when two or more tasks in the original system cut set (scenario) can be accomplished in the same location (say, Room 222). The location cut set for this system cut set would contain multiple instances of "Room 222" even though the saboteur would likely be able to accomplish all of these tasks during a single visit to that location. Otherwise, if he can gain access to that room once, there is no reason to believe that he could not do so more than once. In either case, the multiple instances of "Room 222" in the location cut set might give us a false sense of security because, at first glance, it looks like the saboteur must gain access to more locations than would actually be required to accomplish his intentions. Therefore, no location is allowed to appear in any location cut set more than once. Additional instances are simply removed from the cut set through the application of the laws of Boolean algebra.

The other form of redundancy involves the comparing of two location cut sets. Consider a situation in which one sabotage scenario requires access to Room 222, Room 187, and the Network Control Center, while another can be accomplished simply by visiting Rooms 222 and 187. Clearly, if we can prevent the second of these scenarios (by denying access to either Room 222 or Room 187), then the first is also prevented. Thus, that location cut set is redundant and can be removed from further consideration. Redundant cut sets are also mathematically removed through the application of the laws of Boolean algebra. While these substitution and mathematical reduction steps may seem complex, they are all performed using readily available risk analysis software. It is critical that these steps be performed because they allow us to get rid of "red herrings" and focus on those combinations of locations that are truly important to the security of our information system.

Each of the cut sets that remain at this point represents one minimal set of locations that a saboteur would have to access in order to force our information system into an undesired state. There are several ways to evaluate this information to help formulate a protection strategy. First, it may be that some of the locations that are contained in the cut sets are beyond your control. Examples of such locations include access from public networks (dial-up access, the Internet, etc.), network cables that pass through public right-of-way, and power supplied by public utilities. One should assume that a saboteur will be able to access these locations at will and, thus, that they are beyond the reach of our protective actions. Mathematically, this assumption is equivalent to assigning these event locations a value of "Always True". Events that are "Always True" can be deleted from cut sets because they are redundant. There may in fact be some location cut sets in which all of the events fall into this category. Under this mathematical operation, this cut set degenerates to the simple condition "Always True," and indicates that there is at least one scenario that an adversary can exploit without ever entering a physical or virtual location that is under our control. Such a scenario may call for a fundamental redesign of the system to incorporate additional redundancy or moving assets from public areas to controlled areas. This operation can also be performed using readily available risk analysis software. The cut sets that remain after this operation is performed represent the answer to the question, "Which locations *that are under my control* does a saboteur have to gain access to in order to force my information system into an undesired state?"

A "perfect" security program would protect our information system against any known threat scenario. While we know that perfect protection is not possible, a worthy goal might be to assure that no known threat

scenario is left totally unprotected, or, if that is not possible, to at least identify those scenarios that are left unprotected so that the system owners can consciously decide to accept the risks that are involved. In our model, the remaining location cut sets provide a list of locations that would have to be accessed to exploit a known scenario. A scenario is thwarted if we break it at any one point. Thus, if we can design a protection method that provides some assurance of breaking every cut set, then we are approaching our security system objective. This can be done in an *ad hoc* manner through examination of the cut sets. For example, if a particular location is present in a large number of cut sets, then protecting that location breaks all of those cut sets and provides some protection against all of those scenarios. It is also important to look at the "small" cut sets (those location cut sets that contain only one or a few locations) because these represent scenarios that may be particularly easy for a saboteur to carry out (he may be able to do everything from one place). As we provide protection to more and more areas, the number of unbroken cut sets is reduced. The remaining unbroken cut sets represent the sources of residual risk for the information system.

There is a mathematical method that can be used to help identify the most appropriate locations to protect. If we analyze the mathematical "dual" of the location cut sets, we obtain a list of "protection sets." Each protection set consists of a list of locations, and has the property that if every location on the list is protected, then every location cut set (threat scenario) is broken. A typical installation may have many protection sets because there may be several different combinations of locations that will achieve the same end (breaking every cut set). One could then examine each protection set to determine the cost of protecting all of the locations that it contains. This provides a way to prioritize candidate protection schemes on a cost basis. This should not be the only basis for the final security implementation decision, however, as there are always ease of use, functionality, and other intangible factors that must factor into the ultimate design of any security system.

While the mathematical duality approach described above is theoretically very appealing, it should be noted that the actual determination of duality is computationally very challenging. This operation can be performed using only a few of the available risk analysis software packages. However, if the dual can be obtained, it provides an ideal way to begin optimizing the protection strategy. The cut sets can also provide clues to help in this process, but they do not directly provide complete lists of locations to be protected. The cut sets can, however, be used as input for discrete optimization techniques such as genetic algorithms in order to obtain similar classes of results.

## V. APPLICATIONS

Vital area analyses have been performed using these techniques for a variety of facilities (including weapons-related facilities and nuclear power plants) since the late 1970's. An important feature of these vital area analysis techniques is that they follow directly from risk and reliability analyses, so the same models can be used for both results. Variations on these same techniques can be used to assess the susceptibility of systems to non-sabotage location-related threats such as fire and flood damage. As was noted previously, however, the principal strengths of these methods are in assessing static event models. While some modeling of dynamic events is possible, it is still the subject of considerable ongoing research.

Sandia has successfully used these techniques to assess the reliability of and/or the risks associated with a wide variety of information systems, including enhanced 911 emergency services architectures, public telephone common channel signaling networks, non-hierarchical data networks (LAN/MAN/WAN environments), and high-speed ATM data networks. Vulnerability models are direct descendants of these risk and reliability analyses, and their results can provide valuable decision support during both the design and operational phases of an information system.

## VI. SUMMARY

This paper has presented the results from an interdisciplinary team that was formed at Sandia National Laboratories to explore the applicability of PLM techniques to information systems. We have demonstrated that many aspects of information systems can be modeled using a "plug-and-play" fault tree analysis technique as well as other PLM techniques. We have also demonstrated that the types of results that can be obtained from PLMs can be of great practical value to network designers as well as security analysts through the use of vital area analysis techniques. These PLM techniques are not intended to replace current network analysis methods, but to supplement them. They provide additional tools for the network designers' workbench to enhance their depth of understanding so that they can design more optimal and more secure network systems.

## REFERENCES

1. M.O. Ball, "Computational Complexity of Network Reliability Analysis: An Overview," *IEEE Trans. Reliability*, Vol. R-35, No. 3, pp. 230-239, August 1986.

2. D.W. Coit & A.E. Smith, "Use of a Genetic Algorithm to Optimize a Combinatorial Reliability Design Problem," Proceedings of the Third IIE Research Conference, 467-472, 1994.
3. L. Painton & J. Campbell, "Genetic Algorithms in the Optimization of System Reliability," *IEEE Transactions on Reliability*, Special Issue on Design, 44(2), 172-178, 1995.
4. N.H. Roberts, W.E. Vesely, D.F. Haasl, and F.F. Goldberg, "Fault Tree Handbook," NUREG-0492, U.S. Nuclear Regulatory Commission, Washington, D.C., January 1981.
5. J.B. Dugan, S.J. Bavuso & M.A. Boyd, "Dynamic Fault Tree Models for Fault Tolerant Computer Systems," *IEEE Transactions on Reliability*, 41(3), 363-377, 1992.
6. L.L. Pullum & J.B. Dugan, "Fault Tree Models for the Analysis of Complex Computer Systems," Proceedings of the 1996 Reliability and Maintainability Symposium, 1996.
7. G.B. Varnado, W.H. Horton, and P.R. Lobner, "Modular Fault Tree Analysis Procedures Guide," SAND83-0963, NUREG/CR-3268, Prepared by Sandia National Laboratories for the U.S. Nuclear Regulatory Commission, Washington, D.C., August 1983.
8. T.L. Zimmerman, N.L. Graves, A.C. Payne, and D.W. Whitehead, "Microcomputer Applications of and Modifications to the Modular Fault Trees," SAND89-1887, NUREG/CR-4838, Prepared by Sandia National Laboratories for the U.S. Nuclear Regulatory Commission, Washington, D.C., June 1990.
9. "Novell's Guide to Netware LAN Analysis," 2nd Edition, Novell Press, 1993.
10. A.S. Tanenbaum, "Computer Networks," 2nd Edition, Prentice Hall, New York, 1991.
11. B. Bingham, J. Hutchison, and B. Lopez, "SEATREE Version 2.62 User Manual," prepared for Sandia National Laboratories by Science and Engineering Associates, Albuquerque, New Mexico, 1994.
12. K.M. Hays, G.D. Wyss, and S.L. Daniel, "A User's Guide to SABLE," Sandia National Laboratories, Albuquerque, New Mexico, 1996.
13. T. Polif and A. Sathyaranayana, "Efficient Algorithms for Reliability Analysis of Planar Networks — A Survey," *IEEE Trans. Reliability*, Vol. R-35, No. 3, pp. 252-259, August 1986.
14. R.K. Wood, "Factoring Algorithms for Computing *K*-Terminal Network Reliability," *IEEE Trans. Reliability*, Vol. R-35, No. 3, pp. 269-278, August 1986.
15. G.D. Wyss, H.K. Schriner & T.R. Gaylor, "Risk and Reliability Assessment For Telecommunications Networks," American Nuclear Society International Topical Meeting on Probabilistic Safety Assessment "PSA '96," Park City, Utah, 1996.
16. D.D. Boozer *et. al.*, "Safeguards System Effectiveness Modeling," SAND76-0428, Sandia National Laboratories, Albuquerque, New Mexico, September 1976.
17. D.W. Stack and M.S. Hill, "A SETS User's Manual for Vital Area Analysis," SAND83-0074, NUREG/CR-3134, Prepared by Sandia National Laboratories for the U.S. Nuclear Regulatory Commission, Washington, D.C., April 1984.

## **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.