Dempster-Shafer Theory and Network Intrusion Detection Systems

M. Esmaili[1]

Intrusion Detection Systems (IDS) were previously made by hand. These systems have difficulty in classifying intruders successfully and require a significant amount of computational overhead, causing problems in the creation of robust real-time IDS. Artificial Intelligence (AI) techniques can reduce the human effort required to build these systems and can improve their performance. AI has recently been used in Intrusion Detection (ID) for anomaly detection, data reduction and induction, or discovery, of rules explaining audit data [1]. The use of expert system technology allows certain intrusion scenarios to be specified much more easily and naturally than in the cases where other technologies are being used. However, expert system technology alone provides no support for developing models of intrusive behavior and encourages the development of ad hoc rules. This paper proposes the application of evidential reasoning for dealing with uncertainty in IDS. It is shown that how through dealing with uncertainty the system is allowed to detect the abnormality in the user behavior more efficiently.

INTRODUCTION

ID is the identification of attempted or ongoing attacks on a computer system or network. Issues in ID research include data collection, data reduction, behavior classification, reporting and response. Although there are many significant open problems in ID research, the focus here is on behavior classification. Classification is the process of identifying attackers and intruders. AI techniques have been used in many IDS to perform these important tasks.

There has been a deepening concern regarding the problem of ever increasing intrusions into computers and computer networks on the internet and other networks. Intruders can be characterized as "joy riders" with no malicious intent, as thieves aiming to appropriate resources of the computer system or those controlled by the system, or as terrorists aiming to destroy or incapacitate the system. Intruders are also sometimes known as "hackers".

The study of providing security in computer networks is a rapidly growing area of interest because the network is the medium over which most attacks or intrusions on computer systems are launched. One approach to solving this problem is the intrusion detection concept whose basic premise is that not only abandoning the existing and huge infrastructure of possibly-insecure computer and network systems is impossible, but also replacing them by totally secure systems may not be feasible or cost effective.

The importance of securing the data and

1. Center for Computer Security Research, University of Wollongong, Wollongong, NSW 2522, Australia.

Scientia Iranica, Vol. 3, No. 4, © Sharif University of Technology, January 1997.

information maintained by a corporation has become a driving force in the development of numerous systems that perform computer security audit trail analyses [2,3]. These systems are generally classified as "intrusion detection" systems. The primary purpose of an intrusion detection system is to expose computer security threats in a timely manner.

BACKGROUND

Intrusión detection and network security are becoming increasingly more important in today's computer-dominated society. As more and more sensitive information is being stored on computer systems and transferred over computer networks, more and more crackers are attempting to attack these systems to steal, destroy or corrupt that information. While most computer systems attempt to prevent unauthorized use by some kind of access control mechanism such as passwords, encryption and digital signatures, there are several factors that make it very difficult to keep these crackers from eventually gaining entry into a system [4–8]. Most computer systems have some kind of security flaw that may allow outsiders (or legitimate users) to gain unauthorized access to sensitive information. In most cases, it is not practical to replace such a flawed system with a new, more secure system. It is also very difficult, if not impossible, to develop a completely-secure system. Even a supposedly secure system can still be vulnerable to insiders misusing their privileges, or it can be compromised by improper operating practices. While many existing systems may be designed to prevent specific types of attack other methods to gain unauthorized access may still be possible. Due to the tremendous investment already made in the existing infrastructure of open (and possibly insecure) communication networks, it is infeasible to deploy new, secure and possibly closed networks. Since the event of an attack should be considered inevitable, there is an obvious need for mechanisms that can detect outsiders attempting to gain entry into a system, detect insiders misusing their system privileges and monitor the networks connecting all of these systems together.

The goal of any intrusion detection system must be to aid the System Security Officer (SSO) in the detection of penetration and abuse. The system should provide the knowledge of an "expert" security officer. This is a minimum standard of performance for an intrusion detection system. Humans generally do not do a very good job of audit trail analysis, since the volume of audit record data generated makes this a difficult and time consuming job. The set of penetrations or abuses detected by a security officer with the aid of the automated system should be a superset that would have been detected by the security officer unaided.

IDS require that basic security mechanisms are in place which enforce authorization controls over system, data and other resource access on computer or network and that an audit trail be available to record a variety of computer usage activity. IDS attempt to identify security threats through the analysis of these computer security audit trails. IDS are based on the principle that an attack on a computer system (or network) will be noticeably different from normal system (or network) activity [9]. A system intruder (possibly masquerading as a legitimate user) is very likely to exhibit a pattern of behavior different from the normal behavior of a legitimate user. The job of IDS is to detect these abnormal patterns by analyzing numerous sources of information that are provided by the existing systems. The two major methods used to detect intrusions are statistical analysis and rule-based expert systems analysis [2,3]. The statistical method attempts to define normal (expected) behavior, while the expert system defines proper behavior. The expert system also searches for breaches of policy. If the IDS notices a possible attack using either of these methods, then the SSO is notified. The SSO may then take some action against the aggressor.

Statistical Analysis

One means of detecting anomalous behavior is to monitor statistical measures of user activities

on the system. A popular way to monitor statistical measures is to keep profiles of legitimate user activity [10–15]. These profiles may include such items as login times, CPU usage, favorite editor and compiler, disk usage, number of printed pages per session, session length, error rate, etc [10]. The IDS will then use these profiles to compare current user activity with past user activity. Whenever a current user's activity pattern fails outside certain predefined thresholds, the behavior is considered anomalous. Legitimate behavior flagged as intrusive is defined to be a false alarm. A major problem with the statistical model is determining exactly what activities and statistical measures provide the highest detection rate and lowest false alarm rate for a particular system.

It may also be the case that a particular activity may not be threatening by itself, but when aggregated with other activities, it may constitute an attack. These statistical profiles must be adaptive, i.e., they must be updated regularly, since users may be constantly changing their behavior.

Knowledge-Base Analysis

Another means of detecting possible attacks on a computer system is by using a knowledge-based expert system analysis method [8,11,14, 16]. Codification and reapplication of knowledge under similar circumstances are the basis of an expert system. This knowledge is encoded in the form of facts (assertions about the state of a problem solution) and heuristics (rules which govern the transformation of the solution state). The expert system analyzes the audit trail records and tries to determine attacks based on the information contained in the rule base. The expert system is able to pose sophisticated queries to the knowledge-base to answer conditional questions based on sets of events. The main problem with this method is determining exactly what kinds of attacks can be detected using the method. As an example, the knowledge-base may contain rules describing known attack methods and signatures, known system flaws, expected system behavior and the site-specific security policy.

OBJECTIVE

Most of the current IDS are built on the concept of detecting anomalous behavior of users with respect to observed behavioral norms. This approach may be seen as an unsupervised learning scheme for behavioral patterns with a subsequent pattern recognition approach to determine whether observed behavior falls inside or outside the pattern. In effect, a model of a user's behavior is generated based on observations, but it is difficult to relate the model to specific (and specially proscribed) activities. Thus, validation of the behavior of IDS' statistical algorithms may prove to be difficult.

As mentioned earlier, some IDS include an expert system component that attempts to encode known system vulnerabilities and attack scenarios in its rule base. The IDS raises an alarm if observed activity matches any of its encoded rules. However, expert system technology provides no support for developing models of intrusive behavior and encourages the development of ad hoc rules.

In this paper, the idea is to extend the IDS paradigm to include specific models of proscribed activities. These models would imply certain activities with certain observables which could then be monitored. This would allow an active search for intruders by looking for activities which would be consistent with a hypothesized intrusion scenario. But the evidence can not always be matched perfectly to a hypothesized intrusion. Therefore, a determination of the likelihood of a hypothesized intrusion would be made based on the combination of evidence for and against it. The security of such an explicit model should be easier to validate. However, the system must be able to deal with information that can be uncertain.

Various numerical calculi have been proposed as methods to represent and propagate uncertainty in a system. Among the more

prominent calculi are probabilistic (in particular Bayesian) methods, the evidence theory of Dempster-Shafer, fuzzy set theory and the MYCIN and EMYCIN calculi [17]. In this paper, as considered the application of evidential reasoning in computer intrusion detection is investigated.

THE DEMPSTER-SHAFER THEORY

In the 1960s, A. Dempster laid the foundation for a new mathematical theory of uncertainty. In the 1970s, this theory was extended by G. Shafer to what is now known as the Dempster-Shafer theory. This theory may be viewed as a generalization of probability theory. Contrary to the subjective Bayesian method and the certainty factor model [17], the Dempster-Shafer theory has not been specially developed for reasoning with uncertainty in expert systems. Only at the beginning of the 1980s did it become apparent that this theory might be suitable for such a purpose. However, this theory cannot be applied in an expert system without modification. Moreover, this theory in its original form has an exponential computational complexity. For rendering it useful in the context of expert systems, Lucas and Van Der Gaag proposed several modifications of this theory [18].

The Probability Assignment

As mentioned earlier, the Dempster-Shafer theory may be viewed as a generalization of probability theory. The development of this theory has been motivated by the observation that probability theory is not able to distinguish between uncertainty and ignorance, owing to incompleteness of information. In probability theory, probabilities have to be associated with individual atomic hypotheses. Only if these probabilities are known, does the computation of other probabilities of interest become possible. In the Dempster-Shafer theory, however, it is possible to associate measures of uncertainty with sets of hypotheses, interpreted as disjoints, instead of with individual hypotheses only. This, nevertheless, renders it possible to make statements concerning the uncertainty of other sets of hypotheses. Note that, in this way, the theory is able to distinguish between uncertainty and ignorance.

The strategy followed in the Dempster-Shafer theory for dealing with uncertainty roughly amounts to starting with an initial set of hypotheses. Then, each piece of evidence is associated with a measure of uncertainty with certain subsets of the original set of hypotheses. This continues until measures of uncertainty may be associated with all possible subsets on account of the combined evidence. The initial set of all hypotheses in the problem domain is called the frame of discernment. In such a frame of discernment, the individual hypotheses are assumed to be disjoint. The distribution of a unit of belief over a frame of discernment is called a mass distribution [19]. A mass distribution, m_Θ, is a mapping from the subsets of a frame of discernment, Θ, into the unit interval. The impact of a piece of evidence (body of evidence) on the confidence or belief in certain subsets of a given frame of discernment is described by means of a function which is defined as follows [18].

Definition 1

Let Θ be a frame of discernment. If a number $m_\Theta(x)$ is associated with each subset $x \subseteq \Theta$ such that:

1. $m_\Theta(x) \geq 0$,
2. $m_\Theta(\emptyset) = 0$,
3. $\sum_{x \subseteq \Theta} m_\Theta(x) = 1$,

then m_Θ is called a basic probability assignment (or mass distribution) on Θ. For each subset $x \subseteq \Theta$, the number $m_\Theta(x)$ is called the basic probability number of x.□

There are two other notions which should be defined.

Definition 2

Let Θ be a frame of discernment and let m_Θ be a mass distribution on Θ. A set $x \subseteq \Theta$ is called a focal element in m_Θ if $m_\Theta(x) > 0$. The core

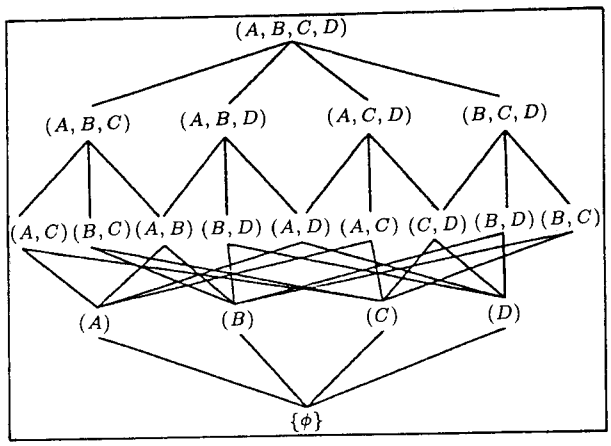

Figure 1. Lattice of all possible subsets of the universe $\Theta = \{A, B, C, D\}$.

of m_Θ, denoted by $\kappa(m)$, is the set of all focal elements of m_Θ.□

Notice the similarity between a basic probability assignment (mass distribution) and a probability function. A probability function associates each element in Θ with a number from the interval $[0, 1]$ such that the sum of these numbers equal 1. Figure 1 shows the lattice of all possible subsets for a typical set Θ. A mass distribution (basic probability) associates a number in the interval $[0, 1]$ with each element in 2^Θ such that, again, the sum of the numbers equal 1.

$$m_\Theta : \quad 2^\Theta \longmapsto [0, 1] .$$

A probability number $m_\Theta(x)$ expresses the confidence or belief assigned to precisely the set x. It does not express any belief in subset of x. It will be evident, however, that the total confidence in x is not dependent on the confidence assigned to subsets of x. For a given basic probability assignment, Lucas et al. define a function describing the cumulative belief in a set of hypotheses [18].

Definition 3

Let Θ be a frame of discernment and m_Θ be a mass distribution on Θ. Then, the belief function (or credibility function) corresponding to m_Θ is the function Bel:$2^\Theta \longmapsto [0, 1]$ defined by:

$$\text{Bel}(x) = \sum_{y \subseteq x} m_\Theta(y) ,$$

for each $x \subseteq \Theta$.□

Several properties of this belief function can easily be proved:

1. $\text{Bel}(\Theta) = 1$ since $\sum_{y \subseteq \Theta} m_\Theta(y) = 1$.
2. For each $x \subseteq \Theta$ containing exactly one element, $\text{Bel}(x) = m_\Theta(x)$.
3. For each $x \subseteq \Theta$, $\text{Bel}(x) + \text{Bel}(\bar{x}) \leq 1$ is obtained, since:

$$\begin{aligned} \text{Bel}(\Theta) &= \text{Bel}(x \cup \bar{x}) \\ &= \text{Bel}(x) + \text{Bel}(\bar{x}) + \sum_{\substack{x \cap y \neq \emptyset \\ \bar{x} \cap y \neq \emptyset}} m_\Theta(y) = 1. \end{aligned}$$

Furthermore, the inequality $\text{Bel}(x) + \text{Bel}(y) \leq \text{Bel}(x \cup y)$ holds for each $x, y \in \Theta$.

Some special belief functions follow. Recall that a basic probability assignment (mass distribution) describing lack of evidence had the following form:

$$m_\Theta(x) = \begin{cases} 1 & \text{if } x = \Theta \\ 0 & \text{otherwise} . \end{cases}$$

The belief function corresponding to such an assignment has been given a special name [18].

Definition 4

Let Θ be a frame of discernment and m_Θ be a mass distribution, such that $\kappa(m_\Theta) = \{\Theta\}$. The belief function corresponding to m_Θ is called a vacuous belief function.□

The following definition from [18] concerns functions corresponding with mass distribution of the form:

$$m_\Theta(x) = \begin{cases} 1 - C_1 & \text{if } x = \Theta \\ C_1 & \text{if } x = A \\ 0 & \text{otherwise} , \end{cases}$$

where $A \subset \Theta$ and $0 < C_1 < 1$ is a constant.

Definition 5

Let Θ be a frame of discernment and m_Θ be a mass distribution, such that $\kappa(m_\Theta) = \{A, \Theta\}$ for a certain $A \subset \Theta$. The belief function corresponding to m_Θ is called a simple support function.□

A belief function provides a lower bound for each set x to the 'actual' belief in x. It is also possible that belief has been assigned to a set w such that $x \subseteq w$. Therefore, in addition to the belief function, the Dempster-Shafer theory defines another function corresponding with a basic probability assignment (mass distribution).

Definition 6

Let Θ be a frame of discernment and m_Θ be a mass distribution on Θ. Then, the plausibility function corresponding to m_Θ is the function Pl: $2^\Theta \longmapsto [0,1]$ defined by:

$$\mathrm{Pl}(x) = \sum_{x \cap y \neq \emptyset} m_\Theta(w) \ ,$$

for each $x \subseteq \Theta$.□

A function value $\mathrm{Pl}(x)$ indicates the total confidence not assigned to $\bar{x}$, so $\mathrm{Pl}(x)$ provides an upperbound to the 'real' confidence in x. It can be shown that, for a given basic probability assignment m_Θ, the property:

$$\mathrm{Pl}(x) = 1 - \mathrm{Bel}(\bar{x}) \ ,$$

for each value $x \subseteq \Theta$, holds for the belief function Bel and the plausibility function Pl corresponding to m_Θ. The difference $\mathrm{Pl}(x) - \mathrm{Bel}(x)$ indicates the confidence in the sets w for which $x \subseteq w$ and, therefore, expresses the uncertainty with respect to x.

Definition 7

Let Θ be a frame of discernment and m_Θ be a mass distribution on Θ. Let Bel be the belief function corresponding to m_Θ and Pl be the plausibility function corresponding to m_Θ. For each $x \subseteq \Theta$, the closed interval [Bel(x),Pl(x)] is called the belief interval of x.□

The lower bound of a belief interval indicates the degree to which the evidence supports the hypothesis, while the upper bound indicates the degree to which the evidence fails to refute the hypothesis, i.e., the degree to which it remains plausible.

Example 1

Let Θ be a frame of discernment and $x \subseteq \Theta$. Now, consider a basic probability m_Θ on Θ and its corresponding functions Bel and Pl.

- If $[\mathrm{Bel}(x), \mathrm{Pl}(x)] = [0,1]$, then no information concerning x is available.
- If $[\mathrm{Bel}(x), \mathrm{Pl}(x)] = [0,0]$, then x has been completely denied by m_Θ.
- If $[\mathrm{Bel}(x), \mathrm{Pl}(x)] = [0,0.8]$, then there is some evidence against x.
- If $[\mathrm{Bel}(x), \mathrm{Pl}(x)] = [1,1]$, then x has been completely confirmed by m_Θ.
- If $[\mathrm{Bel}(x), \mathrm{Pl}(x)] = [0.3,1]$, then there is some evidence in favor of the hypothesis x.
- If $[\mathrm{Bel}(x), \mathrm{Pl}(x)] = [0.15,0.75]$, then there is some evidence in favor of as well as against x.□

If $\mathrm{Pl}(x) - \mathrm{Bel}(x) = 0$ for each $x \subseteq \Theta$, then Dempster-Shafer Theory is the same as the conventional probability theory. In such a case, the belief function is called a Bayesian belief function. This notion is defined more formally in the following definition from [18].

Definition 8

Let Θ be a frame of discernment and m_Θ be a mass distribution such that the core of m_Θ consists only of singleton sets. The belief function corresponding to m_Θ is then called a Bayesian belief function.□

Dempster's Rule of Combination

The Dempster-Shafer theory provides a function for computing a new basic probability assignment from two pieces of evidence and their associated basic probability assignment, describing the combined influence of the pieces of evidence. This function is known as the Dempster's rule of combination. The more formal definition is as follows [18].

Definition 9

Let Θ be a frame of discernment and m^1_Θ and m^2_Θ be basic probability assignments on Θ. Then $m^1_\Theta \oplus m^2_\Theta$ is a function $m^1_\Theta \oplus m^2_\Theta : 2^\Theta \longmapsto [0,1]$, such that:

1. $m^1_\Theta \oplus m^2_\Theta(\emptyset) = 0$ and
2. $m^1_\Theta \oplus m^2_\Theta(x) = \dfrac{\sum_{y \cap z = x} m^1_\Theta(y) \cdot m^2_\Theta(z)}{\sum_{y \cap z \neq \emptyset} m^1_\Theta(y) \cdot m^2_\Theta(z)}$ for all $x \neq \emptyset$.

$\text{Bel}_1 \oplus \text{Bel}_2$ is the function $\text{Bel}_1 \oplus \text{Bel}_2 : 2^\Theta \longmapsto [0,1]$ defined by:

$$\text{Bel}_1 \oplus \text{Bel}_2(x) = \sum_{y \subseteq x} m^1_\Theta \oplus m^2_\Theta(y). \qquad (1)$$

Evidential Reasoning

The goal of evidential reasoning is to assess the effect of all available pieces of evidence upon a hypothesis by making use of domain-specific knowledge. The first step in applying evidential reasoning to a given problem is to delimit a propositional space of possible situations. Within the theory of belief functions, this propositional space is called the frame of discernment. A frame of discernment delimits a set of possible situations, exactly one of which is true at any one time. Once a frame of discernment has been established, propositional statements can be represented by subsets of elements from the frame corresponding to those situations for which the statements are true. Bodies of evidence are expressed as probabilistic opinions about the partial truth or falsity of propositional statements whose granularity is appropriate to the variable evidence.

Evidential reasoning provides a number of formal operations for assigning evidence [19], including:

1. Fusion – to determine a consensus from several bodies of evidence obtained from independent sources. Fusion is accomplished through the Dempster's rule of combination (Equation 1):

$$m^3_\Theta(A_h) = \frac{1}{1-k} \sum_{A_i \cap A_j = A_h} m^1_\Theta(A_i) m^2_\Theta(A_j),$$

$$k = \sum_{A_i \cap A_j = \phi} m^1_\Theta(A_i) m^2_\Theta(A_j). \qquad (2)$$

The Dempster's Rule is both commutative and associative (meaning evidence can be fused in any order) and has the effect of focusing belief on those propositions that are held in common.

2. Translation – to determine the impact of a body of evidence upon elements of a related frame of discernment. The translation of a BOE from frame Θ_A to frame Θ_B, using the compatibility relation $\Theta_{A,B}$, is defined by:

$$m_{\Theta_B}(B_j) = \sum_{\substack{C_{A \mapsto B}(A_k) = B_j \\ A_k \subseteq \Theta_A, B_j \subseteq \Theta_B}} m_{\Theta_A}(A_k), \qquad (3)$$

where, $C_{A \mapsto B}(A_k) = \{b_j \mid (a_i, b_j) \in \Theta_{A,B}, a_i \in A_k\}$.

3. Projection – to determine the impact of a body of evidence at some future (or past) point in time. The projection operation is defined exactly as translation, where the frames are taken to be one time-unit apart.
4. Discounting – to adjust a body of evidence to account for the credibility of its source. Discounting is defined as:

$$m^{discounted}_\Theta(A_j) = \begin{cases} \alpha \cdot m_\Theta(A_j), & A_j \neq \Theta \\ 1 - \alpha + \alpha \cdot m_\Theta(\Theta), & \text{otherwise} ; \end{cases} \qquad (4)$$

where α is the assessed credibility of the original BOE ($0 \leq \alpha \leq 1$).

Independent opinions are expressed by multiple bodies of evidence. Dependent opinions can be represented either as a single body of evidence, or as a network structure that shows the inter-relationships of several BOEs. The evidential reasoning approach focuses on a body of evidence, which describes a meaningful collection of interrelated beliefs as the primitive representation. In contrast, all other such technologies focus on individual propositions.

Analysis Using an Example

To illustrate the evidential reasoning method described above in an intrusion detection system, the following example is used:

> "A user successfully logs in from a remote host after trying several bad passwords and usernames. The user enters several wrong command names and arguments and tries to look at some directories and files entry to which is denied. The user also employs commands such as 'finger' several times to find out about other system users and activities. The user also copies the `/bin/csh` file into `/usr/spool/mail/root` where the root's mail directory resides and makes it a setuid file by `chmod 4755 /usr/spool/mail/root` command. After a few minutes, the user leaves. Who was this? Could it be an intruder or just an inexperienced user who was experimenting with the system?"

In evidential reasoning the first step is to construct the sets of possibilities (the frame of discernment) for each unknown. For example, the user could either be an intruder or not; these possibilities can be represented in the `Intruder?` frame:

$$\{Yes,\ No\}.$$

Other frames could also be constructed; `Location` will be included for the user's location containing the possibilities:

$$\{Local,\ Remote\}.$$

Two types of location for a user are distinguished — local (i.e., physically at the keyboard) and remote. Because the majority of intruders do not have direct physical access to the locally connected terminals, a local keyboard is considered to indicate normal use and not an intruder. Most intrusions originate from remote internet sites. However, because an intruder can jump from host to host, intrusive behavior is also likely to appear originating from local hosts. Thus, activity originating from any location other than the local keyboard is considered equally indicative of intrusive behavior, so only the single category 'remote' will be used for this. For a remote user, it cannot be distinguished whether the user is an intruder, based on this dimension of behavior alone.

An intruder is expected to be somewhat paranoid, therefore a frame, `Fear`, is included to capture the paranoia level:

$$\{Paranoid,\ Calm\}.$$

A paranoid intruder (one who is afraid of being caught) will probably have very short sessions (eg., lasting under two minutes), because the longer the session the greater the risk of discovery. A paranoid intruder will also commonly check to see who is logged in and what they are doing. Thus, for example, in Unix an ordinate number of 'who', 'ps' and 'finger' commands can be expected to indicate a paranoid intruder. User sessions can be characterized as having a high degree of this sort of activity if two or more such commands are used. Therefore, short sessions and two or more "surveillance" commands are considered to be strong indicators of fear.

An intruder may also be unfamiliar with the system, so another frame, `familiarity`, will be defined to contain:

$$\{Familiar,\ Unfamiliar\}.$$

A person who is unfamiliar with the computer system under attack is likely to have a relatively large number of invalid commands, resulting from attempts to execute commands that are not recognized by the system. Such a person is also likely to have a relatively large number of errors resulting from invalid command usage, for example, too few arguments or invalid parameters. But this alone cannot be a good measure to condemn a user to be an intruder, since the user might be inexperienced. This frame should be looked at in conjunction with other frames. A relatively large number of file permission errors, resulting from attempts to read, write, or execute files or directories when permission is denied, is also indicative of a person unfamiliar with the computer system under attack. Therefore, relatively large number of these types of errors are considered

to be strong indicators of unfamiliarity with the system. Conversely, low error rates for all of these categories of error strongly suggest a normal, nonintrusive user.

Another frame can be constructed for the actions which raise the suspicion level, such as copying a file from `/bin` directory or trying to access somebody else's mail file, or etc. These actions can be represented in the `Actions` frame:

$$\{Malicious,\ Normal\}.$$

Authentication errors result from the use of an invalid username or password during login. A high rate of authentication errors (greater than three failed login attempts for a given username within a certain time period) is considered to be strongly suggestive of an intrusion attempt.

Once the frames are defined, the next step is to construct the compatibility relations that define the domain-specific relationships between the frames. A connection between two propositions A_1 and B_1 indicates that they may co-occur (in other words, $(A_1, B_1) \in \Theta_{A,B}$).

Figure 2 shows the frames and compatibility relations used in determining whether the user is an intruder.

Once the frames and compatibility relations have been established, the evidence can be analyzed. The goal of the analysis is to establish a line of reasoning from the evidence to determine belief in a hypothesis, in this case that the user is an intruder.

The first step is to assess each piece of evidence relative to an appropriate frame of discernment. Each piece of evidence is represented as a mass distribution, which distributes a unit of belief over subsets of the frame. For example, the fact that the user logged in from a remote host is pertinent to the `Location` frame and 1.0 is attributed to remote, to indicate the complete certainty on this point.

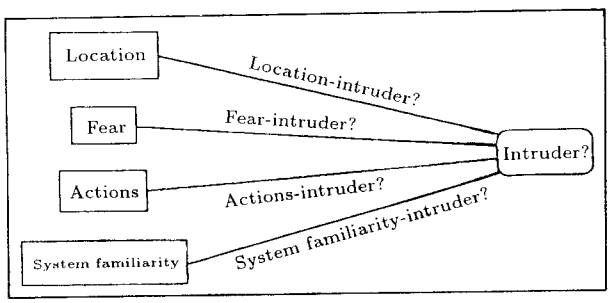

Figure 2. Frames and compatibility relations.

The fact that the user had a high number of authentication errors leads to the belief that the user may be an intruder. Based on this, a likelihood of 0.75 is assigned to this possibility.

The high number of command usage and file permission errors gives information about `Familiarity`. Based on the number and types of errors, a belief of 0.7 is assigned to the possibility, Unfamiliar; the remaining 0.3 is assigned to Familiar.

The user tried some commands, at least two of which can be interpreted as malicious intent, that might give information about `Actions`. Based on this belief, a likelihood of 0.8 is given to the possibility that the user's actions have been malicious.

The last piece of evidence, that the user employed several "surveillance" commands and had a short session, gives information about `Fear`. It might be assessed as 0.75 support for the user being paranoid and 0.25 for the user being calm. This is usual behavior for that user (perhaps the user is a system administrator).

In this example, beliefs about paranoia levels, system familiarity, actions and, authentication errors are drawn, directly from interpretations over various types of audit data. These processes can also be represented directly in evidential reasoning, at the cost of some additional complexity. In practice, reasoning processes will be required to include more extensive analysis of this sort.

Evidence from these sources will provide the inputs to the analysis. Many of these determinations are judgments that may not be of equal validity. In order to be able to weigh them differently, a means for discounting the impact of the evidence through the discounting operation will be provided. This will allow change in their relative weights.

The final step is to construct the actual analysis of the evidence as shown in Figure 3

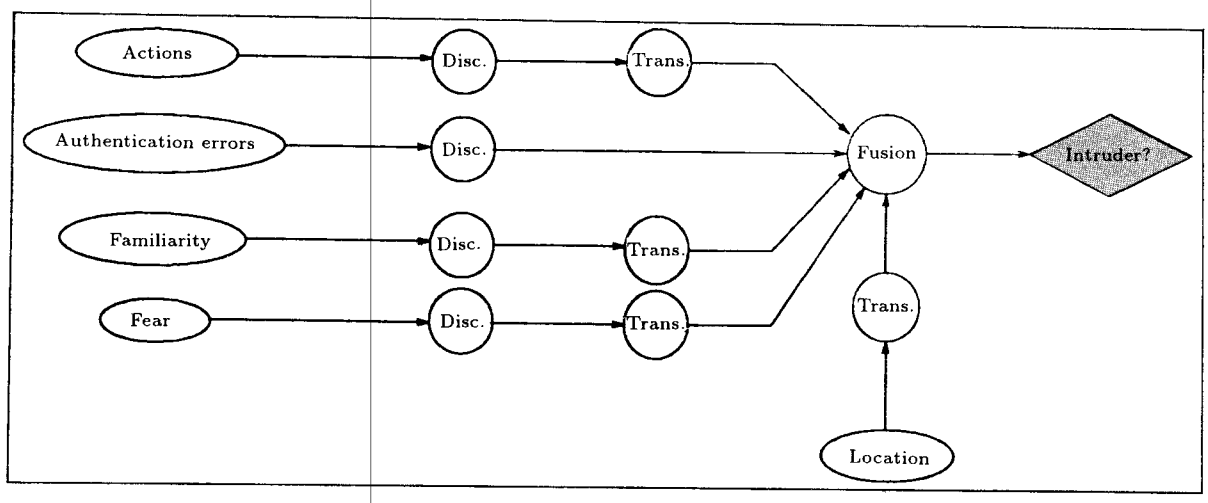

Figure 3. Frames and compatibility relations.

to determine its impact upon the question at hand. In this case the question of whether the user is an intruder can be answered by an assessment of belief over elements in the `Intruder?` frame. Evidential operations can be used to derive a body of evidence providing beliefs about whether the user is an intruder.

In the analysis shown in Figure 3, all sources except `Location` source are discounted. The Authorization Errors source provides information directly about the likelihood of an intruder, but the others must all be translated to the `Intruder?` frame. After translation, these independent BOEs are represented relative to a common frame and can be combined using the fusion operation (i.e., the Dempster's Rule). Fusing the mass distributions yields a mass distribution relative to the `Intruder?` frame, from which conclusions, as to whether the user is an intruder, can be drawn.

CONCLUSION

In this paper an attempt is made to demonstrate the applications of AI techniques, specifically Expert Systems in IDS. It is also shown how by using evidential reasoning, the system is allowed to detect abnormality in the user behavior more efficiently. The use of Expert System technology allows certain intrusion scenarios to be specified much more easily and naturally than is the case using other technologies. However, expert system technology provides no support for developing models of intrusive behavior and encourages the development of ad hoc rules.

ACKNOWLEDGMENT

The authors wish to thank Dr. Muthukumar Balachandran and Prof. Svein Knapskog for their invaluable discussions and feedback.

REFERENCES

1. Frank, J. "Artificial intelligence and intrusion detection: Current and future directions", in *Proceedings of 17th National Computer Security Conference*, **1**, Baltimore, Meryland, pp 22–33 (Oct. 1994).

2. Esmaili, M., Safavi-Naini, R. and Pieprzyk, J. "Computer intrusion detection: A comparative survey", Tech. Rep. TR-95-07, Department of Computer Science, University of Wollongong, Australia (Aug. 1995).

3. Esmaili, M., Safavi-Naini, R. and Pieprzyk, J. "Intrusion detection: A survey", in *Proceedings of Twelfth International Conference on Computer Communication ICCC'95*, **1**, Seoul, Korea, pp 409– 414 (Aug. 1995), Sponsored by International Council for Computer Communication.

4. Bauer, D.S. and Koblentz, M.E. "NIDX - an expert system for real-time network intrusion detection", in *Proceedings of the IEEE*

Computer Networking Symposium, pp 98–106 (1988).

5. Brignone, A. "Fuzzy Sets: An answer to the evaluation of security systems?", in *Proceedings of Fourth IFIP TCII International Conference on Comp. Sec. (IFIP/Sec'86)*, Monte Carlo, Monaco, pp 143–151 (Dec. 1986).
6. Debar, H. and Dorizzi, B. "An application of a recurrent network to an intrusion detection system", in *Proceedings of International Joint Conference on Neural Networks*, pp II478–II483 (June 1992).
7. Debar, H., Becker, M. and Siboni, D. "A neural network component for an intrusion detection system", in *Proceedings of the 1992 IEEE Computer Society Symposium on Research in Security and Privacy*, pp 240–250 (May 1992).
8. Lunt, T.F. "IDES: An intelligent system for detecting intruders", in *Proceedings of the Symposium: Computer Security, Threat and Countermeasures*, Rome, Italy (Nov. 1990).
9. Anderson, D. et al. "Next generation intrusion detection expert system (NIDES): User manual for security officer user interface (SOUI)", Technical Report, SRI International (March 1993).
10. Denning, D.E. "An intrusion-detection model", in *Proceedings of the 1986 IEEE Symposium on Security and Privacy*, pp 118–131 (April 1986).
11. Lunt, T.F. and Jagannathan, R. "A prototype real-time intrusion-detection expert system", in *Proceedings of the 1988 IEEE Symposium on Security and Privacy* (Apr. 1988).
12. Lunt, T.F. "Real-time intrusion detection", in *Proceedings of COMPCON Spring '89*, San Francisco, California (Feb. 1989).
13. Lunt, T.F. "Using statistics to track intruders", in *Proceedings of the Joint Statistical Meetings of the American Statistical Association* (August 1990).
14. Lunt, T.F. et al. "A real-time intrusion-detection expert system (IDES)", Final Technical Report, SRI International (Feb. 1992).
15. Lunt, T.F. "A survey of intrusion detection techniques", *Computers & Security*, **12**(4), pp 405–418 (1993).
16. Porras, P.A. and Kemmerer, R.A. "Penetration state transition analysis a rule-based intrusion detection approach", in *Proceedings of the Eight Annual Computer Security Applications*, IEEE Comp. Soc. Press, pp 220–229 (Dec. 1992).
17. Henkind, S.J. and Harrison, M.C. "An analysis of four uncertainty calculi", *IEEE Transactions on Systems, Man and Cybernetics*, **18**, pp 700–714 (Sept./Oct. 1988).
18. Lucas, P. and Van Der Gaag, L. *Principles of Expert Systems*, Addison-Wesley Publishing Company (1991).
19. Lowrance, J.D. and Garvey, T.D. "A framework for evidential-reasoning systems", in *Readings in Uncertain Reasoning*, G. Shafer and J. Pearl, Eds., Morgan Kaufmann Publishers, Inc., San Mateo, California, USA, pp 611–618 (1990).