

Sensor Fusion for Enhancement in Intrusion Detection

Ciza Thomas¹ and Balakrishnan Narayanaswamy²

¹Professor, College of Engineering, Trivandrum

²Associate Director, Indian Institute of Science, Bangalore
India

1. Introduction

An Intrusion Detection System (IDS) gathers information from a computer or a network, and analyzes this information to identify possible security breaches against the system or the network. An observation of various IDSs available in literature shows distinct preferences for detecting a certain class of attack with improved accuracy, while performing moderately on the other classes. The availability of enormous computing power has made it possible for developing and implementing IDSs of different types on the same network. The integration of the decisions coming from different IDSs has emerged as a technique that could strengthen the final decision. Sensor fusion can be defined as the process of collecting information from multiple and possibly heterogeneous sources and combining them to obtain a more descriptive, intuitive and meaningful result (1).

An analysis of the poorly detected attacks reveals the fact that the attacks are characterized by features that do not discriminate them much. In this chapter, we prove the distinct advantages of sensor fusion over individual IDSs. All the related work in the field of sensor fusion has been carried out mainly with one of the methods like probability theory, evidence theory, voting fusion theory, fuzzy logic theory or neural network in order to aggregate information. The Bayesian theory is the classical method for statistical inference problems. The fusion rule is expressed for a system of independent learners, with the distribution of hypotheses known *a priori*. The Dempster-Shafer evidence theory is considered a generalized Bayesian theory. It does not require *a priori* knowledge or probability distribution on the possible system states like the Bayesian approach and it is mostly useful when modeling of the system is difficult or impossible (2). The improved performance of multiple IDSs using rule-based fusion and data-dependent decision fusion has been demonstrated in the work of Thomas and Balakrishnan (3).

An attempt to prove the distinct advantages of sensor fusion over individual IDSs is done in this chapter using the Chebyshev inequality. Fusion threshold bounds were derived using the principle of Chebyshev inequality at the fusion center using the false positive rates and detection rates of the IDSs. The goal was to achieve best fusion performance with the least amount of model knowledge, in a computationally inexpensive way. The anomaly-based IDSs detect anomalies beyond a set threshold level in the features it detects. Threshold bounds instead of a single threshold give more freedom in steering system properties. Any threshold

within the bounds can be chosen depending on the preferred level of trade-off between detection and false alarms.

The remaining part of the chapter is organized as follows. Section 2 discusses the related work of sensor fusion in IDS. In section 3, the modeling of the Intrusion Detection System is presented. Section 4 includes the modeling of the fusion of Intrusion Detection Systems. Section 5 contains the experimental results along with the discussions regarding the higher performance of the proposed fused IDS. Finally, the concluding comments are presented in section 6.

2. Related work

Tim Bass (4) presents a framework to improve the performance of intrusion detection systems based on data fusion. A few first steps towards developing the engineering requirements using the art and science of multi-sensor data fusion as an underlying model is provided in (4). Giacinto et al. (5) propose an approach to intrusion detection based on fusion of multiple classifiers. Didaci et al. (6) attempt the formulation of the intrusion detection problem as a pattern recognition task using data fusion approach based on multiple classifiers. Wang et al. (7) present the superiority of data fusion technology applied to intrusion detection systems. The use of data fusion in the field of DoS anomaly detection is presented by Siaterlis and Maglaris (1). The detection engine is evaluated using the real network traffic. Another work incorporating the Dempster-Shafer theory of evidence is by Hu et al. (8).

Siraj et al. (9) discuss a Decision Engine for an Intelligent Intrusion Detection System (IIDS) that fuses information from different intrusion detection sensors using an artificial intelligence technique. Thomopolous in one of his work (10), concludes that with the individual sensors being independent, the optimal decision scheme that maximizes the probability of detection at the fusion for fixed false alarm probability consists of a Neyman-Pearson test at the fusion unit and the likelihood ratio test at the sensors. The threshold based fusion of combining multiple IDSs by fixing a certain number of false alarms is discussed in the work of Thomas and Balakrishnan (11). This is a case of combining the top ranking outputs of each IDS after removing the duplicate alerts and setting the maximum acceptable false alarm rate.

The other somewhat related works albeit distantly are the alarm clustering method by Perdisci et al. (12), aggregation of alerts by Valdes et al. (13), combination of alerts into scenarios by Dain et al. (14), the alert correlation by Cuppens et al. (15), the correlation of Intrusion Symptoms with an application of chronicles by Morin et al. (16), and aggregation and correlation of intrusion-detection alerts by Debar et al. (17). In the work of Thomas and Balakrishnan (3), a sensor fusion architecture, which is data-dependent and different from the conventional fusion architecture is demonstrated. The focus of the present work is modeling the fusion of IDSs using threshold bounds in an attempt to optimize both the fusion rule as well as the sensor rules.

3. Modeling the Intrusion Detection Systems

Consider an IDS that either monitors the network traffic connection on the network or the audit trails on the host. The network traffic connection or the audit trails monitored are given as $x \in X$, where X is the entire domain of network traffic features or the audit trails respectively. The model is based on the hypothesis that the security violations can be detected by monitoring the network for traffic connections of malicious intent in the case of network-based IDS and a system's audit records for abnormal patterns of system usage in

the case of host-based IDS. The model is independent of any particular operating system, application, system vulnerability or type of intrusion, thereby providing a framework for a general-purpose IDS.

When making an attack detection, a connection pattern is given by $x_j \in \mathcal{R}^{jk}$ where j is the number of features from k consecutive samples used as input to an IDS. As seen in the DARPA dataset, for many of the features the distributions are difficult to describe parametrically as they may be multi-modal or very heavy-tailed. These highly non-Gaussian distributions has led to investigate non-parametric statistical tests as a method of intrusion detection in the initial phase of IDS development. The detection of an attack in the event x is observed as an alert. In the case of network-based IDS, the elements of x can be the fields of the network traffic like the raw IP packets or the pre-processed basic attributes like the duration of a connection, the protocol type, service etc. or specific attributes selected with domain knowledge such as the number of failed logins or whether a superuser command was attempted. In host-based IDS, x can be the sequence of system calls, sequence of user commands, connection attempts to local host, proportion of accesses in terms of TCP or UDP packets to a given port of a machine over a fixed period of time etc. Thus IDS can be defined as a function that maps the data input into a normal or an attack event either by means of absence of an alert (0) or by the presence of an alert (1) respectively and is given by:

$$IDS : X \rightarrow \{0, 1\}.$$

To detect attacks in the incoming traffic, the IDSs are typically parameterized by a threshold T . The IDS uses a theoretical basis for deciding the thresholds for analyzing the network traffic to detect intrusions. Changing this threshold allows the change in performance of the IDS. If the threshold is very low, then the IDS tends to be very aggressive in detecting the traffic for intrusions. However, there is a potentially greater chance for the detections to be irrelevant which result in large false alarms. A large value of threshold on the other hand will have an opposite effect; being a bit conservative in detecting attacks. However, some potential attacks may get missed this way. Using a 3σ based statistical analysis, the higher threshold (T_h) is set at $+3\sigma$ and the lower threshold (T_l) is set at -3σ . This is with the assumption that the traffic signals are normally distributed. In general the traffic detection with s being the sensor output is given by:

$$Sensor\ Detection = \begin{cases} attack, & T_l < s < T_h \\ normal, & s \leq T_l, s \geq T_h \end{cases}$$

The signature-based IDS functions by looking at the event feature x and checking whether it matches with any of the records in the signature database D .

$$\begin{aligned} Signature - based\ IDS : X &\rightarrow \{1\} && \forall x \in D, \\ &: X \rightarrow \{0\} && \forall x \notin D. \end{aligned}$$

Anomaly-based IDS generates alarm when the input traffic deviates from the established models or profiles P .

$$\begin{aligned} Anomaly - based\ IDS : X &\rightarrow \{1\} && \forall x \notin P, \\ &: X \rightarrow \{0\} && \forall x \in P. \end{aligned}$$

4. Modeling the fusion of Intrusion Detection Systems

Consider the case where n IDSs monitor a network for attack detection and each IDS makes a local decision s_i and these decisions are aggregated in the fusion unit f . This architecture is often referred to as the parallel decision fusion network and is shown in Figure 1. The fusion unit makes a global decision, y , about the true state of the hypothesis based on the collection of the local decisions gathered from all the sensors. The problem is casted as a binary detection




Fig. 1. Parallel Decision Fusion Network

problem with the hypothesis “Attack” or “Normal”. Every IDS participating in the fusion has its own detection rate D_i , and false positive rate F_i , due to the preferred heterogeneity of the sensors in the fusion process. Each IDS indexed i gives an alert or no-alert indicated by s_i taking a value one or zero respectively, depending on the observation x .

$$s_i = \begin{cases} 0, & \text{normal is declared to have been detected} \\ 1, & \text{attack is declared to have been detected} \end{cases}$$

The fusion center collects these local decisions s_i and forms a binomial distribution y as given by $y=s = \sum_{i=1}^n s_i$, where n is the total number of IDSs taking part in fusion.

4.1 The effect of setting threshold

To detect the attack in the incoming traffic, the IDSs are typically parameterized with a threshold, T . Changing this threshold allows the change in performance of the IDS. If the threshold is very large, some potentially dangerous attacks get missed. A small threshold on the other hand results in more detections, with a potentially greater chance that they are not relevant.

The final step in the approach towards solving of the fusion problem is taken by noticing that the decision function $f_i(\cdot)$ is characterized by the threshold T_i and the likelihood ratio (if independence is assumed). Thus the necessary condition for optimal fusion decision occurs if the thresholds (T_1, T_2, \dots, T_n) are chosen optimally. However, this does not satisfy the sufficient condition. These refer to the many local minima, each need to be checked to assure the global minimum.

The counterintuitive results at the individual sensors with the proper choice of thresholds will be advantageous in getting an optimum value for the fusion result. They are excellent paradigms for studying distributed decision architectures, to understand the impact of the limitations, and even suggest empirical experiments for IDS decisions.

The structure of the fusion rule plays a crucial role regarding the overall performance of the IDS since the fusion unit makes the final decision about the state of the environment. While a few inferior IDSs might not greatly impact the overall performance, a badly designed fusion rule can lead to a poor performance even if the local IDSs are well designed. The fusion IDS

can be optimized by searching the space of fusion rules and optimizing the local thresholds for each candidate rule. Other than for some simple cases, the complexity of such an approach is prohibitive due to exponential growth of the set of possible fusion rules with respect to the number of IDSs. Searching for the fusion rule that leads to the minimum probability of error is the main bottleneck due to discrete nature of this optimization process and the exponentially large number of fusion rules. The computation of thresholds couples the choice of the local decision rules so that the system-wide performance is optimized, rather than the performance of the individual detector.

4.2 Threshold optimization

Tenney and Sandell in their work (21) establish the optimum strategy that minimizes a global cost in the case where the *a priori* probabilities of the hypotheses, the distribution functions of the local observations, the cost functions, and the fusion rule are given. They concluded that each local detector is optimally a likelihood ratio detector but that the computation of the optimum thresholds for these local detectors is complicated due to cross coupling.

The global optimization criterion for a distributed detection system would encompass local decision statistics, local decision thresholds, the fusion center decision statistic, and the fusion center decision threshold. For each input traffic observation x , the set of n local thresholds should be optimized with respect to the probability of error. With a fusion rule given by a function f , the average probability of error at the fusion unit is given by the weighted sum of false positive and false negative errors.

$$P_e(T, f) = p * P(s = 1|Normal) + q * P(s = 0|Attack) \quad (1)$$

where p and q are the respective weights of false positive and false negative errors. Assuming independence between the local detectors, the likelihood ratio is given by:

$$\frac{P(s|Attack)}{P(s|Normal)} = \frac{P(s_1, s_2, \dots, s_N|Attack)}{P(s_1, s_2, \dots, s_N|Normal)} = \prod_{i=1}^n \frac{P(s_i|Attack)}{P(s_i|Normal)}$$

The optimum decision rule for the fusion unit follows:

$$f(s) = \log \frac{P(s|Attack)}{P(s|Normal)}$$

Depending on the value of $f(s)$ being greater than or equal to the decision threshold, T , or less than the decision threshold, T , the decision is made for the hypothesis as "Attack" or "Normal" respectively. Thus the decisions from the n detectors are coupled through a cost function. It is shown that the optimal decision is characterized by thresholds as in the decoupled case. As far as the optimum criterion is concerned, the first step is to minimize the average probability of error in equation 1. This leads to sets of simultaneous inequalities in terms of the generalized likelihood ratios at each detector, the solutions of which determine the regions of optimum detection.

4.3 Dependence on the data and the individual IDSs

Often, the data in the databases is only an approximation of the true data. When the information about the goodness of the approximation is recorded, the results obtained from

the database can be interpreted more reliably. Any database is associated with a degree of accuracy, which is denoted with a probability density function, whose mean is the value itself. In order to maximize the detection rate it is necessary to fix the false alarm rate to an acceptable value, taking into account the trade-off between the detection rate and the false alarm rate. The threshold (T) that maximizes the TP_{rate} and thus minimizes the FN_{rate} is given as:

$$FP_{rate} = P[alert|normal] = P \left[\sum_{i=1}^n w_i s_i \geq T | normal \right] = \alpha_0 \quad (2)$$

$$TP_{rate} = P[alert|attack] = P \left[\sum_{i=1}^n w_i s_i \geq T | attack \right] \quad (3)$$

The fusion of IDSs becomes meaningful only when $FP \leq FP_i \quad \forall i$ and $TP \geq TP_i \quad \forall i$; where FP and TP correspond to the false positives and the true positives of the fused IDS and FP_i and TP_i correspond to the false positives and the true positives of the individual IDS indexed i . It is required to provide low value of weight to any individual sensor that is unreliable, hence meeting the constraint on false alarm as given in equation 2. Similarly, the fusion improves the TP_{rate} as the detectors get weighted according to their performance.

4.4 Modeling the fusion IDS by defining proper threshold bounds

Every IDS participating in the fusion has its own detection rate D_i , and false positive rate F_i , due to the preferred heterogeneity of the sensors in the fusion process. Each IDS indexed i gave an alert or no-alert indicated by s_i taking a value of one or zero respectively. The fusion center collected these local decisions and formed a binomial distribution s as given by $s = \sum_{i=1}^n s_i$,

where n is the total number of IDSs taking part in the fusion.

Let D and F denote the unanimous detection rate and the false positive rate respectively. The mean and variance of s in case of attack and no-attack, are given by the following equations:

$$E[s|alert] = \sum_{i=1}^n D_i, \quad Var[s|alert] = \sum_{i=1}^n D_i(1 - D_i)$$

; in case of attack

$$E[s|alert] = \sum_{i=1}^n F_i, \quad Var[s|alert] = \sum_{i=1}^n F_i(1 - F_i)$$

; in case of no-attack

The fusion IDS is required to give a high detection rate and a low false positive rate. Hence the threshold T has to be chosen well above the mean of the false alerts and well below the mean of the true alerts. The figure 2 shows a typical case where the threshold T is chosen at the point of overlap of the two parametric curves for normal and attack traffics. Consequently, the threshold bounds are given as:

$$\sum_{i=1}^n F_i < T < \sum_{i=1}^n D_i$$




Fig. 2. Parametric curve showing the choice of threshold T

The detection rate and the false positive rate of the fusion IDS is desired to surpass the corresponding weighted averages and hence:

$$D > \frac{\sum_{i=1}^n D_i^2}{\sum_{i=1}^n D_i} \tag{4}$$

and

$$F < \frac{\sum_{i=1}^n (1 - F_i)F_i}{\sum_{i=1}^n (1 - F_i)} \tag{5}$$

Now, using simple range comparison,

$$D = Pr\{s \geq T|attack\} = Pr\{s - \sum_{i=1}^n D_i \leq (\sum_{i=1}^n D_i - T)|attack\}.$$

Using Chebyshev inequality on the random variable s , with $Mean = E[s] = \sum_{i=1}^n D_i$ and

$$Variance = Var[s] = \sum_{i=1}^n D_i(1 - D_i),$$

$$Pr\{|s - E(s)| \geq k\} \leq \frac{Var(s)}{k^2}$$

With the assumption that the threshold T is greater than the mean of normal activity,

$$Pr\{s - \sum_{i=1}^n D_i \leq (\sum_{i=1}^n D_i - T)|attack\} \geq 1 - \frac{\sum_{i=1}^n D_i(1 - D_i)}{(\sum_{i=1}^n D_i - T)^2}$$

From equation 4 it follows that $1 - \frac{\sum_{i=1}^n D_i(1 - D_i)}{(\sum_{i=1}^n D_i - T)^2} \geq \frac{\sum_{i=1}^n D_i^2}{\sum_{i=1}^n D_i}$

The upper bound of T is derived from the above equation as:

$$T \leq \sum_{i=1}^n D_i - \sqrt{\sum_{i=1}^n D_i}$$

Similarly, for the false positive rate, $F = Pr\{S \geq T \mid \text{no-attack}\}$, in order to derive the lower bound of T ,

From equation 5 it follows that $\frac{\sum_{i=1}^n F_i(1 - F_i)}{(T - \sum_{i=1}^n F_i)^2} \leq \frac{\sum_{i=1}^n F_i(1 - F_i)}{\sum_{i=1}^n (1 - F_i)}$

The lower bound of T is derived from the above equation as:

$$T \geq \sum_{i=1}^n F_i + \sqrt{\sum_{i=1}^n (1 - F_i)}$$

The threshold bounds for the fusion IDS is:

$$\left[\sum_{i=1}^n F_i + \sqrt{\sum_{i=1}^n (1 - F_i)}, \sum_{i=1}^n D_i - \sqrt{\sum_{i=1}^n D_i} \right]$$

Since the threshold T is assumed to be greater than the mean of normal activity, the upper bound of false positive rate F can be obtained from the Chebyshev inequality as:

$$F \leq \frac{\text{Var}[s]}{(T - E[s])^2} \quad (6)$$

In a statistical intrusion detection system, a false positive is caused due to the variance of network traffic during normal operations. Hence, to reduce the false positive rate, it is important to reduce the variance of the normal traffic. In the ideal case, with normal traffic the variance is zero. The equation 6 shows that as the variance of the normal traffic approaches zero, the false positive rate should also approach zero. Also, since the threshold T is assumed to be less than the mean of the intrusive activity, the lower bound of the detection rate D can be obtained from the Chebyshev inequality as:

$$D \geq 1 - \frac{\text{Var}[s]}{(E[s] - T)^2} \quad (7)$$

For an intrusive traffic, the factor $D_i(1 - D_i)$ remains almost steady and hence the variance given as:

Variance = $\sum_{i=1}^n D_i(1 - D_i)$, is an appreciable value. Since the variance of the attack traffic is

above a certain detectable minimum, from equation 7, it is seen that the correct detection rate can approach an appreciably high value. Similarly the true negatives will also approach a high value since the false positive rate is reduced with IDS fusion.

It has been proved above that with IDS fusion, the variance of the normal traffic is clearly dropping down to zero and the variance of the intrusive traffic stays above a detectable minimum. This additionally supports the proof that the fusion IDS gives better detection rate and a tremendously low false positive rate.

5. Results and discussion

5.1 Test set up

The test set up for the experimental evaluation consisted of a combination of shallow and deep sensors. Hence, for the purpose of fusion we have incorporated two sensors, one that monitors the header of the traffic packet and the other that monitors the packet content. The experiments were conducted with the simulated IDSs PHAD and ALAD (22). This choice of heterogeneous sensors in terms of their functionality is to exploit the advantages of fusion IDS (4). In addition, complementary IDSs provide versatility and similar IDSs ensure reliability. The PHAD being packet-header based and detecting one packet at a time, is totally unable to detect the slow scans. However, PHAD detects the stealthy scans much more effectively. The ALAD being content-based will complement the PHAD by detecting R2L (Remote to Local) and U2R (User to Root) attacks with appreciable efficiency.

5.2 Data set

The fusion IDS and all the IDSs that form part of the fusion IDS were separately evaluated with the same two data sets, namely 1) the real-world network traffic and 2) the DARPA 1999 data set. The real traffic within a protected University campus network was collected during the working hours of a day. This traffic of around two million packets was divided into two halves, one for training the anomaly IDSs, and the other for testing. The test data was injected with 45 HTTP attack packets using the HTTP attack traffic generator tool called libwhisker (23). The test data set was introduced with a base rate of 0.0000225, which is relatively realistic. The MIT-DARPA data set (IDEVAL 1999) (24) was used to train and test the performance of Intrusion Detection Systems. The network traffic including the entire payload of each packet was recorded in tcpdump format and provided for evaluation. The data for the weeks one and three were used for the training of the anomaly detectors and the weeks four and five were used as the test data. Each of the IDS was trained on distinct portions of the training data (ALAD on week one and PHAD on week three), which is expected to provide independence among the IDSs and also to develop diversity while being trained.

Even with the criticisms by McHugh (25) and Mahoney and Chan (26) against the DARPA dataset, the dataset was extremely useful in the IDS evaluation undertaken in this work. Since none of the IDSs perform exceptionally well on the DARPA dataset, the aim is to show that the performance improves with the proposed method. If a system is evaluated on the DARPA dataset, then it cannot claim anything more in terms of its performance on the real network traffic. Hence this dataset can be considered as the base line of any research (27). Also, even after 12 years of its generation, there are still a lot of relevant attacks in the data set for which signatures are not available in database of even the frequently updated signature based IDSs. The test data of the DARPA data set consisted of 190 instances of 57 attacks which included 37 probes, 63 DoS attacks, 53 R2L attacks, 37 U2R/Data attacks with details on attack types given in Table 1.

Attack Class	Attack Type
Probe	portsweep, ipsweep, queso, ntinfoScan, mscan, lsdomain, satan, illegal-sniffer
DoS	apache2, smurf, neptune, pod, mailbomb, back, teardrop, udpstorm, processtable, arpoison, tcpreset, crashiis, dosnuke, syslogd, land, selfping, warezclient
R2L	dict, guest, ftpwrite, xlock, xsnoop, httptunnel, framespoof, netbus, netcat, ppmacro, imap, named, ncftp, phf, sendmail, sshotrojan, snmpget
U2R/ Data	perl, xterm, eject, fdformat, ffbconfig, ps, loadmodule, casesen, nukepw, sechole, yaga, secret, ntfstdos, ppmacro, sqlattack

Table 1. Various attack types in DARPA'99 data set

The large observational data set were analyzed to find unsuspected relationships and was summarized in novel ways that were both understandable and useful for the detector evaluation. There are many types of attacks in the test set, many of them not present in the training set. Hence, the selected data also challenged the ability to detect the unknown intrusions. When a discrete IDS was applied to a test set, it yields a single confusion matrix. Thus, a discrete IDS produced only a single point in the ROC space, whereas scoring IDSs can be used with a threshold to produce different points in the ROC space.

5.3 Evaluation metrics

Let TP be the number of attacks that are correctly detected, FN be the number of attacks that are not detected, TN be the number of normal traffic packet/connections that are correctly classified, and FP be the number of normal traffic packet/connections that are incorrectly detected as attack. In the case of an IDS, there are both the security requirements and the usability requirements. The security requirement is determined by the TP_{rate} and the usability requirement is decided by the number of FPs because of the low base rate in the case of a network traffic.

The commonly used IDS evaluation metrics on a test data are the overall accuracy and F-score.

$$Overall\ Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

Overall Accuracy is not a good metric for comparison in the case of network traffic data since the true negatives abound.

Precision is a measure of what fraction of test data detected as attack are actually from the attack classes.

$$Precision = \frac{TP}{TP + FP}$$

Recall is a measure of what fraction of attack class was correctly detected.

$$Recall = \frac{TP}{TP + FN}$$

There is a trade-off between the two metrics precision and recall. As the number of detections increase by lowering of the threshold, the recall will increase, while precision is expected to decrease. The recall-precision characterization of a particular IDS is normally used to

analyze the relative and absolute performance of an IDS over a range of operating conditions. F-score, which is the harmonic mean of recall (R) and precision (P), scores the balance between precision and recall. The F-score is given by:

$$F\text{-score} = \frac{2 * P * R}{P + R}$$

The standard measures, namely precision, recall, and F-score are grounded on a probabilistic framework and hence allows one to take into account the intrinsic variability of performance estimation.

5.4 Experimental evaluation

The fusion element analyzes the IDS data coming from PHAD and ALAD distributed across the single subnet and observing the same domain. The fusion unit performed the aggregation of the IDS outputs for the purpose of identifying the attacks in the test data set. It used binary fusion by giving an output value of one or zero depending on the value of the aggregation of the various IDS decisions. The packets were identified by their timestamp on aggregation. A value of one at the output of the fusion unit indicated the record to be under attack and a zero indicated the absence of an attack.

The fusion IDS was initially evaluated with the DARPA 1999 data set. The individual IDSs chosen in this work are PHAD and ALAD, two research IDSs that are anomaly-based and having extremely low false alarm rate of the order of 0.00002. The other reason for the choice of PHAD and ALAD was that they are almost complementary in attack detection as evident from table 2 and table 3. This helps in achieving best results from the fusion process. The

Attack type	Total attacks	Attacks detected	% detection
Probe	37	22	59%
DOS	63	24	38%
R2L	53	6	11%
U2R/Data	37	2	5%
Total	190	54	28%

Table 2. Types of attacks detected by PHAD at 0.00002 FP rate (100 FPs)

Attack type	Total attacks	Attacks detected	% detection
Probe	37	6	16%
DOS	63	19	30%
R2L	53	25	47%
U2R/Data	37	10	27%
Total	190	60	32%

Table 3. Types of attacks detected by ALAD at 0.00002 FP rate (100 FPs)

analysis of PHAD and ALAD has resulted in a clear understanding of the individual IDSs expected to succeed or fail under a particular attack. On combining the two sensor alerts and removing the duplicates, an improved rate of detection is achieved as shown in table 4. The performance in terms of F-score of PHAD, ALAD and the combination of PHAD and ALAD is shown in the tables 5, 6 and 7 respectively and figure 3 for various values of false positives by setting the threshold appropriately. In our experiment we are trying to maximize the true positive rate by fixing the false positive rate at α_0 . α_0 determines the threshold T by

Attack type	Total attacks	Attacks detected	% detection
Probe	37	24	65%
DOS	63	39	62%
R2L	53	26	49%
U2R/Data	37	10	27%
Total	190	99	52%

Table 4. Types of attacks detected by the combination of ALAD and PHAD at 0.00004 FP rate (200 FPs)

FP	TP	Precision	Recall	Overall Accuracy	F-score
50	33	0.39	0.17	0.99	0.24
100	54	0.35	0.28	0.99	0.31
200	56	0.22	0.29	0.99	0.25
500	56	0.10	0.29	0.99	0.15

Table 5. F-score of PHAD for different choice of false positives

FP	TP	Precision	Recall	Overall Accuracy	F-score
50	42	0.45	0.21	0.99	0.29
100	60	0.37	0.31	0.99	0.34
200	66	0.25	0.34	0.99	0.29
500	72	0.12	0.38	0.99	0.18

Table 6. F-score of ALAD for different choice of false positives




Fig. 3. F-score of PHAD, ALAD and fusion IDS for different choices of false positives

FP	TP	Precision	Recall	Overall Accuracy	F-score
50	44	0.46	0.23	0.99	0.31
100	73	0.42	0.38	0.99	0.40
200	99	0.33	0.52	0.99	0.40
500	108	0.18	0.57	0.99	0.27

Table 7. F-score and Detection Performance for different choice of false positives for fused IDS




Fig. 4. ROC curve of PHAD, ALAD and fusion IDS

trial and error. We have noticed that within two or three trials in our case. This is done with the training data and hence it is done off line. The improved performance of the combination of the alarms from each system can be observed in table 7, corresponding to the false positives between 100 and 200, by fixing the threshold bounds appropriately. Thus the combination works best above a false positive of 100 and much below a false positive of 200. In each of the individual IDSs, the number of detections were observed at false positives of 50, 100, 200 and 500, when trained on inside week 3 and tested on weeks 4 and 5. The improved performance of fusion IDS compared to the two IDSs PHAD and ALAD is also illustrated with the ROC semilog curve shown in figure 4. The improved performance of the fusion IDS over some of the fusion alternatives using the real-world network traffic is shown in table 8 and figure 5.

6. Summary

Simple theoretical model is initially illustrated in this chapter for the purpose of showing the improved performance of fusion IDS. The detection rate and the false positive rate quantify the performance benefit obtained through the fixing of threshold bounds. Also, the more independent and distinct the attack space is for the individual IDSs, the better the fusion IDS performs.

The theoretical proof was supplemented with experimental evaluation, and the detection rates, false positive rates, and F-score were measured. In order to understand the importance

Detector/ Fusion Type	Total Attacks	TP	FP	Precision	Recall	F-score
PHAD	45	10	45	0.18	0.22	0.2
ALAD	45	18	45	0.29	0.4	0.34
OR	45	22	77	0.22	0.49	0.30
AND	45	9	29	0.24	0.2	0.22
SVM	45	21	44	0.32	0.47	0.38
ANN	45	21	61	0.26	0.47	0.28
Fusion IDS	45	22	32	0.41	0.49	0.45

Table 8. Comparison of the evaluated IDSs with various evaluation metrics using the real-world network traffic




Fig. 5. F-score of the evaluated IDSs using the real-world network traffic


of thresholding, the anomaly-based IDSs, PHAD and ALAD have been individually analyzed. Preliminary experimental results prove the correctness of the theoretical proof. The chapter demonstrates that our technique is more flexible and also outperforms other existing fusion techniques such as OR, AND, SVM, and ANN using the real-world network traffic embedded with attacks. The experimental comparison using the real-world traffic has thus confirmed the usefulness and significance of the method. The unconditional combination of alarms avoiding duplicates as shown in table 4 results in a detection rate of 52% at 200 false positives, and F-score of 0.4. The combination of highest scoring alarms as shown in table 7 using the DARPA 1999 data set results in a detection rate of 38% and threshold fixed at 100 false positives, and F-score of 0.4.

7. References

- [1] C. Siaterlis and B. Maglaris, Towards Multisensor Data Fusion for DoS detection, ACM Symposium on Applied Computing, 2004

- [2] H. Wu, M. Seigel, R. Stiefelhagen, J. Yang, Sensor Fusion using Dempster-Shafer Theory, IEEE Instrumentation and Measurement Technology Conference, 2002
- [3] C. Thomas and N. Balakrishnan, Advanced Sensor Fusion Technique for Enhanced Intrusion Detection, IEEE International Conference on Intelligence and Security Informatics, Jun. 2008
- [4] T. Bass, Multisensor Data Fusion for Next Generation Distributed Intrusion Detection Systems, IRIS National Symposium, 1999
- [5] G. Giacinto, F. Roli, and L. Didaci, Fusion of multiple classifiers for intrusion detection in computer networks, Pattern recognition letters, 2003
- [6] L. Didaci, G. Giacinto, and F. Roli, Intrusion detection in computer networks by multiple classifiers systems, International Conference on Pattern recognition, 2002
- [7] Y. Wang, H. Yang, X. Wang, and R. Zhang, Distributed intrusion detection system based on data fusion method, Intelligent control and automation, WCICA 2004
- [8] W. Hu, J. Li, and Q. Gao, Intrusion Detection Engine on Dempster-Shafer's Theory of Evidence, Proceedings of International Conference on Communications , Circuits and Systems, vol.3, pp. 1627-1631, Jun 2006
- [9] A. Siraj, R.B. Vaughn, and S.M. Bridges, Intrusion Sensor Data Fusion in an Intelligent Intrusion Detection System Architecture, Proceedings of the 37th Hawaii international Conference on System Sciences, 2004
- [10] S.C.A. Thomopolous, R. Viswanathan, and D.K. Bougoulas, Optimal distributed decision fusion, IEEE Transactions on aerospace and electronic systems, vol. 25, No. 5, Sep. 1989
- [11] C. Thomas and N. Balakrishnan, Selection of Intrusion Detection Threshold bounds for effective sensor fusion, Proceedings of SPIE International Symposium on Defense and Security, vol.6570, Apr 2007
- [12] R. Perdisci, G. Giacinto, and F. Roli, Alarm clustering for intrusion detection systems in computer networks, Engg. applications of Artificial intelligence, Elsevier publications, March 2006
- [13] A. Valdes and K. Skinner, Probabilistic alert correlation, Springer Verlag Lecture notes in Computer Science, 2001
- [14] O.M. Dain and R.K. Cunningham, Building Scenarios from a Heterogeneous Alert Stream, IEEE Workshop on Information Assurance and Security, 2001
- [15] F. Cuppens and A. Mieke, Alert correlation in a cooperative intrusion detection framework, Proceedings of the 2002 IEEE symposium on security and privacy, 2002
- [16] B. Morin and H. Debar, Correlation of Intrusion Symptoms : an Application of Chronicles, RAID 2003
- [17] H. Debar and A. Wespi, Aggregation and Correlation of Intrusion-Detection Alerts, RAID 2001
- [18] C. Thomas and N. Balakrishnan, Improvement in Attack Detection using Advances in Sensor Fusion, IEEE Transactions on Information Forensics and Security, vol.4, No.3, Sep.2009
- [19] M.V. Mahoney and P.K. Chan, Detecting Novel attacks by identifying anomalous Network Packet Headers, Florida Institute of Technology Technical Report CS-2001-2
- [20] M.V. Mahoney and P.K. Chan, Learning non stationary models of normal network traffic for detecting novel attacks, SIGKDD, 2002
- [21] R.R. Tenney and N.R. Sandell, Detection with distributed sensors, IEEE Transactions on Aerospace and Electronic Systems, vol. 17, No.4, Jul 1981

- [22] M.V. Mahoney, A Machine Learning approach to detecting attacks by identifying anomalies in network traffic, PhD Dissertation, Florida Institute of Technology, 2003
- [23] rfp@wiretrip.net/libwhisker
- [24] DARPA Intrusion Detection Evaluation Data Set, http://www.ll.mit.edu/IST/ideval/data/data_index.html
- [25] J. McHugh, Testing Intrusion Detection Systems: A Critique of the 1998 and 1999 DARPA IDS evaluations as performed by Lincoln Laboratory, ACM Transactions on Information and System Security, vol.3, No.4, Nov. 2000
- [26] M.V. Mahoney and P.K. Chan, An analysis of the 1999 DARPA /Lincoln Laboratory evaluation data for network anomaly detection, Technical Report CS-2003-02
- [27] C. Thomas and N. Balakrishnan, Usefulness of DARPA data set in Intrusion Detection System evaluation, Proceedings of SPIE International Defense and Security Symposium, 2008



Sensor Fusion - Foundation and Applications

Edited by Dr. Ciza Thomas

ISBN 978-953-307-446-7

Hard cover, 226 pages

Publisher InTech

Published online 24, June, 2011

Published in print edition June, 2011

Sensor Fusion - Foundation and Applications comprehensively covers the foundation and applications of sensor fusion. This book provides some novel ideas, theories, and solutions related to the research areas in the field of sensor fusion. The book explores some of the latest practices and research works in the area of sensor fusion. The book contains chapters with different methods of sensor fusion for different engineering as well as non-engineering applications. Advanced applications of sensor fusion in the areas of mobile robots, automatic vehicles, airborne threats, agriculture, medical field and intrusion detection are covered in this book. Sufficient evidences and analyses have been provided in the chapter to show the effectiveness of sensor fusion in various applications. This book would serve as an invaluable reference for professionals involved in various applications of sensor fusion.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Ciza Thomas and Balakrishnan Narayanaswamy (2011). Sensor Fusion for Enhancement in Intrusion Detection, Sensor Fusion - Foundation and Applications, Dr. Ciza Thomas (Ed.), ISBN: 978-953-307-446-7, InTech, Available from: <http://www.intechopen.com/books/sensor-fusion-foundation-and-applications/sensor-fusion-for-enhancement-in-intrusion-detection>

INTECH

open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2011 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](#), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.