Protecting Against Cyber Threats in Networked Information Systems

L. Ertoz [a,b], A. Lazarevic[a,b♣], E. Eilertson [a,b], Pang-Ning Tan [a,b],
Paul Dokas [a,] , V. Kumar[a,b], Jaideep Srivastava[a,b]
[a]Dept. of Computer Science & Electrical Engineering, Minneapolis, University of Minnesota
[b]Army High Performance Computing Research Center, Minneapolis, MN 55455

ABSTRACT

This paper provides an overview of our efforts in detecting cyber attacks in networked information systems. Traditional signature based techniques for detecting cyber attacks can only detect previously known intrusions and are useless against novel attacks and emerging threats. Our current research at the University of Minnesota is focused on developing data mining techniques to automatically detect attacks against computer networks and systems. This research is being conducted as a part of MINDS (Minnesota Intrusion Detection System) project at the University of Minnesota. Experimental results on live network traffic at the University of Minnesota show that the new techniques show great promise in detecting novel intrusions. In particular, during the past few months our techniques have been successful in automatically identifying several novel intrusions that could not be detected using state-of-the-art tools such as SNORT.

Keywords: Data mining, cyber threat analysis, network intrusion detection, learning from rare classes, anomaly / outlier detection, characterization.

1. INTRODUCTION

As the cost of information processing and Internet accessibility falls, military and government organizations are becoming increasingly vulnerable to potential cyber threats such as network intrusions. According to a recent survey by CERT/CC (Computer Emergency Response Team/Coordination Center), the rate of cyber attacks has been dramatically increasing every year in recent times (see Figure 1). Although existing security policies and mechanisms (e.g. firewalls) provide a practical protection against such cyber threats, they are not perfect and usually have inevitable vulnerabilities. In fact, military and law enforcement authorities report that every month, assailants make thousands of unauthorized attempts to gain access to these systems, amounting to a nearly continuous assault [1]. In addition, there is an increasing awareness that cyber strategies can be a major force multiplier and equalizer. Therefore, mechanisms for detecting cyber attacks are of great interest for national defense and security.

Traditional methods for cyber threat detection are based on extensive knowledge of signatures of known attacks that are provided by human experts. The signature database has to be manually revised for each new type of intrusion that is discovered. A significant limitation of signature-based methods is that they cannot detect emerging cyber threats. In addition, once a new attack is discovered and its signature developed, often there is a substantial latency in its deployment across networks. These limitations have led to an increasing interest in intrusion detection techniques based upon data mining.

The Data Mining represents a non-trivial interactive process that involves search for structure, patterns or models in large and typically heterogeneous and multi-dimensional data. The field spans several research areas such as statistics, pattern recognition, data visualization, databases, and high performance/distributed computing. Extraction of useful information from network traffic data by human analysts is tedious and time consuming due to its high volume, dimensionality and heterogeneity. Consequently, much of the output of network traffic monitoring is simply stored away on disks and is never analyzed at all. Given its success in commercial applications, data mining holds great promise for

♣ aleks@cs.umn.edu; phone: 1-612-626-8096; fax: 1-612-626-1596

the development of tools for gaining fundamental insights into the network traffic data, thereby allowing system administrators and network engineers to automatically detect emerging cyber attacks [2, 3, 4, 5, 6].

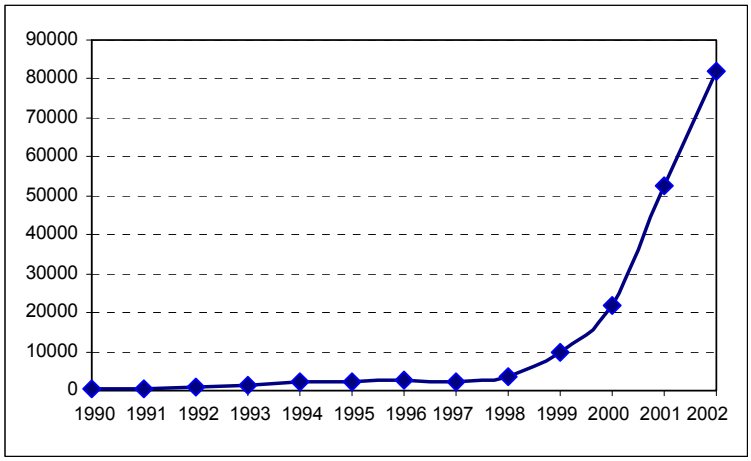

Figure 1. Incidents that are reported to Computer Emergency Response Team/Coordination Center (CERT/CC) during the last decade shows the growth rate of cyber threats

Data mining based intrusion detection techniques generally fall into one of two categories; misuse detection and anomaly detection. In misuse detection, each instance in a data set is labeled as 'normal' or 'intrusion' and a learning algorithm is trained over the labeled data. These techniques are able to automatically retrain intrusion detection models on different input data that include new types of attacks, as long as they have been labeled appropriately. Unlike signature-based intrusion detection systems, models of misuse are created automatically, and can be more sophisticated and precise than manually created signatures. A key advantage of misuse detection techniques is their high degree of accuracy in detecting known attacks and their variations. Their obvious drawback is the inability to detect attacks whose instances have not yet been observed. Anomaly detection, on the other hand, builds models of normal behavior, and automatically detects any deviation from it, flagging the latter as suspect. Anomaly detection techniques thus identify new types of intrusions as deviations from normal usage [7, 8]. While an extremely powerful and novel tool, a potential drawback of these techniques is the rate of false alarms. This can happen primarily because previously unseen (yet legitimate) system behaviors may also be recognized as anomalies, and hence flagged as potential intrusions.

Our research group at the University of Minnesota is developing high performance data mining algorithms and tools that will provide support required to analyze the massive data sets generated by various processes that monitor computing and information systems. This research is being conducted as a part of MINDS (Minnesota Intrusion Detection System) project at the University of Minnesota that is developing a suite of data mining techniques to automatically detect attacks against computer networks and systems [9]. This presentation will provide an overview of the MINDS system and results of its applications to the real network data at the University of Minnesota.

2 THE MINDS SYSTEM OVERVIEW

The *MINDS* (Minnesota INtrusion Detection System) is a data mining based system for detecting network intrusions. Figure 2 shows the architecture of the MINDS system.. There are three integral parts of the research within the *MINDS* project:

- Known attack detection module (Learning from rare class) – building rare class prediction models
- Anomaly detection algorithms
- Summarization of attacks using association pattern analysis

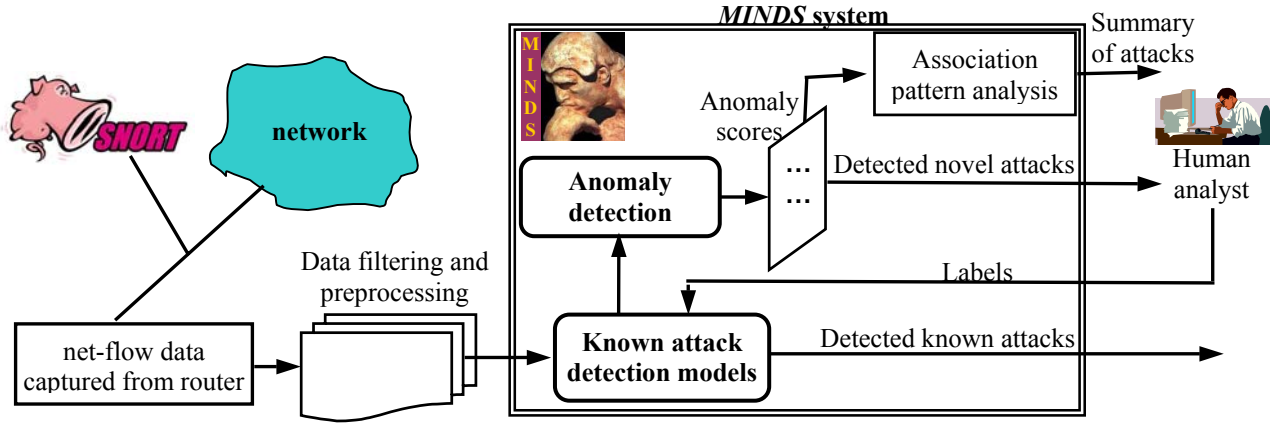

Figure 2. Architecture of MINDS system

At present, a prototype of the MINDS system is being used by the University of Minnesota (UM) network security analysts, as follows. The Traffic from the University of Minnesota network is monitored using SNORT, an open source network intrusion detection system based on signatures, as well as using MINDS. As the first step in MINDS, the net flow tools are used to collect the network traffic data from CISCO routers. This data is filtered to remove network connections not interesting for analysis and preprocessed to collect basic features (such as source and destination IP addresses, source and destination ports, protocols) and to extract some derived features (such as number of flows to unique destination IP addresses inside the network within the last T seconds from the same source). Such created data is fed into the MINDS system. The known attack detection module detects network connections of attacks for which the signatures are available. The remaining connections are fed to the anomaly detection module, which assigns a score reflecting how anomalous the connection is compared to the normal network traffic. Connections that are highly anomalous are analyzed by the UM network security analysts to determine if they are truly intrusions, new normal behavior or just false alarms. Highly anomalous connections are further analyzed and summarized using association pattern analysis. Such association rules provide a concise characterization of the detected anomalies, and are helpful in creating new signatures and models for emerging attacks.

However, straightforward analysis of network data using data mining techniques may be challenging for several reasons. *First*, data generated from network traffic monitoring tends to have very high volume, dimensionality and heterogeneity, since tens of millions network connections, representing a variety of applications (ssh, http, sftp, etc.), are available on a daily basis for commercial and academic network sites. *Second*, in addition to the fact that the network intrusions are very rare they also represent sequences of events, and therefore data mining algorithms need not only to deal with highly skewed class distribution but also to learn from data streams. *Third*, it is problematic to get high-quality data for evaluation due to privacy and competitive issues, since many organizations are not willing to share their data with other institutions. Even if real life data were available, labeling network connections as normal or intrusive may require enormous time for many human experts. *Fourth*, the constant change of the network traffic can not only introduces new types of intrusions but can also change the aspects of the "normal" behavior, thus making detecting novel attacks and their evaluation even more difficult. *Finally*, there is a significant problem due to intermixing the normal network behavior with the intrusive behavior. For example, connections associated with a scanning attack are intermixed with normal network connections.

The MINDS system attempts to successfully overcome these challenges by developing novel data mining algorithms that are appropriate for learning from network traffic data. The following sections provide the overview of these developments for all three MINDS modules.

Known Attack Detection Module - Learning From Rare Class
The existing approaches for building predictive models from rare classes, such as decision trees, neural networks, support vector machines, are applicable provided labeled training data (normal and abnormal users' or applications' behavior) are available. However, for network traffic data, where intrusions correspond to only small percentage of data examples, these standard data mining techniques offer very poor results. Therefore, we have developed several novel classification algorithms designed especially for learning from the rare classes, e.g. rule based systems (PN rule [13], CREDOS [14]), ensemble-based methods (Rare-Boost [11, 12], SMOTEBoost [15]), and modification of standard association-based classification algorithm.

Anomaly Detection
Anomaly detection is a key element of intrusion detection in which perturbations of normal behavior suggest the presence of intentionally or unintentionally induced attacks, faults, defects, etc. Most anomaly detection approaches attempt to build some kind of a model over the normal data and then check to see how well new data fits into that model. The MINDS system currently uses several outlier detection algorithms as well as on unsupervised support vector machine algorithms for detecting network intrusions. An outlier may be defined as a data point which is very different from the rest of the data, based on some measure. In the MINDS system, we utilize outlier detection schemes that are based on computing the full dimensional distances of the points from one another [18, 19] as well as on computing the densities of local neighborhoods [20].

Association Pattern Analysis for Summarization of Attacks
As previously noted, *MINDS* uses the anomaly scores of the connections to determine whether a connection belongs to the normal or attack class. In the MINDS system, we choose connections that have a prespecified percentage of top anomaly scores to be the anomaly class and the prespecified percentage of bottom anomaly scores to be the normal class. Connections with intermediate anomaly scores are ignored. Next, the association pattern generator is applied to each class and the patterns are ranked according to the various measures described in [18]. The extracted patterns can be used to create summaries and profiles for normal and anomaly connections. Once the profiles for the attack class have been established, a follow-up analysis is often performed to study the nature of the anomalous connections. A typical follow-up analysis involves connecting to the suspected computer at the specific port and examining the returned information. Another possibility of analyzing the suspected computer is to start capturing packets on that machine at the particular port and to investigate the contents of the packets.

3. EXPERIMENTAL RESULTS

The University of Minnesota network security analyst has been using MINDS in a production mode for the past several months and is routinely detecting novel intrusions that could not be identified using state-of-the-art signature-based tools such as SNORT [19]. Many of the attacks detected by MINDS, have already been on the CERT/CC list of recent advisories and incident notes. Some of these are attacks for which SNORT did not have a signature. Although these were targeted to multiple machines, the rate of attacks were too slow to be detected by SNORT's scan detector. As an example, on August 13th, 2002, the top ranked anomalous connections corresponded to scanning of Microsoft-DS service on port 445/TCP. SNORT could not detect this attack, since the port scanning was slow (a signature for this attack was added to SNORT in September 2002). The second highest ranked block of connections on the same day corresponded to scanning for an Oracle server. This scan attack was embedded within much larger Web scan, and the SNORT alerts generated by the Web scan were overwhelming, making it hard for the analyst to identify the Oracle scan from the SNORT alarm trace. MINDS routinely detects variations of known worms that cannot be identified by SNORT, as SNORT does not yet have signatures for their variations. For example, the most common version of the Slapper worm uses port 2002 for communication, but some variations use other ports. Connections corresponding to many these variations are often ranked highly by MINDS anomaly detection algorithms. Finally, MINDS anomaly detection scheme is much more successful than SNORT in detecting policy violations (e.g. rogue and unauthorized services), since it looks for unusual network behavior. SNORT may detect these policy violations only if it has a rule for each of these specific activities. Since the number and variety of these activities can be very large and unknown, it is not practical to incorporate them into SNORT. For example, on August 8th and 10th, 2002, MINDS detected a machine running a Microsoft PPTP VPN server, and another one running a FTP server, on non-standard ports, which were policy violations. In the absence of manual screening of incoming network connections,

it is not possible to provide any estimate of MINDS's detection rate (i.e. the fraction of attacks that are identified by MINDS as anomalous). However, nearly all connections that are ranked highly by our anomaly detection algorithms are found to be interesting and anomalous by the network security analyst on our team. MINDS also uses association pattern analysis to summarize highly scored connections. Given the very high volume of connections observed per unit time, such characterization of highly anomalous connections is essential in enabling a security analyst to understand emerging threats, as well in creating new attack signatures. For example, the following set of rules was identified by MINDS association pattern scheme:

```
If (SrcIP=IP1, DstIP = xxx, DstPort=80, Protocol=TCP, Flag=SYN, NoPackets: 3, NoBytes:120…180) ⇒ attack
If (SrcIP=IP1, DstIP = IP2, DstPort=80, Protocol=TCP, Flag=SYN, NoPackets: 3, NoBytes:120…180) ⇒ attack
```

At first glance, the first rule indicates a Web scan since it appears mostly in the anomaly class with a fixed source IP address but not with a fixed destination IP address. However, the second rule suggests that an attack was later launched to one of the specific machines since the pattern originates from the same source IP address but has a specific destination IP address and covers only anomalous connections. Further analysis confirms that a scan has been performed from the source IP1, followed by an attack on a specific machine that was previously identified by the attacker to be vulnerable.

4. CONCLUSIONS

Our continuing objective is to develop an overall framework for defending against attacks and threats to computer systems. Data generated from network traffic monitoring tends to have very high volume, dimensionality and heterogeneity, making the performance of serial data mining algorithms unacceptable for on-line analysis. In addition, cyber attacks may be launched from several different locations and targeted to many different destinations, thus creating a need to analyze network data from several networks in order to detect these distributed attacks. Therefore, development of new classification and anomaly detection algorithms that can take advantage of high performance computers and be computationally tractable for on-line and distributed intrusion detection is a key component of this project. According to our preliminary results on real network data, there is a small overlap of our anomaly detection algorithms with the SNORT intrusion detection system, which implies that they could be combined in order to increase coverage. We are also developing a visualization tool to aid the analyst in better understanding anomalous/suspicious behavior detected using our techniques.

ACKOWLEDGEMENTS

This work was supported by Army High Performance Computing Research Center contract number DAAD19-01-2-0014. The content of the work does not necessarily reflect the position or policy of the government and no official endorsement should be inferred. Access to computing facilities was provided by the AHPCRC and the Minnesota Supercomputing Institute.

REFERENCES

1. F. Vizard, Waging War.com: A Hacker Attack Against NATO Spawns a War in Cyberspace, *Popular Science*, p. 80, July 1999.
2. W. Lee, S. J. Stolfo, Data Mining Approaches for Intrusion Detection, *Proceedings of the 1998 USENIX Security Symposium*, 1998.
3. E. Bloedorn, et al., Data Mining for Network Intrusion Detection: How to Get Started, *MITRE Technical Report*, August 2001.
4. J. Luo, Integrating Fuzzy Logic With Data Mining Methods for Intrusion Detection, *Master's thesis, Department of Computer Science, Mississippi State University*, 1999.

5. D. Barbara, N. Wu, S. Jajodia, Detecting Novel Network Intrusions Using Bayes Estimators, *Proceedings of the First SIAM Conference on Data Mining*, Chicago, IL, 2001.
6. S. Manganaris, M. Christensen, D. Serkle, and K. Hermix, A Data Mining Analysis of RTID Alarms, *Proceedings of the 2nd International Workshop on Recent Advances in Intrusion Detection* (*RAID* 99), West Lafayette, IN, September 1999.
7. D.E. Denning, An Intrusion Detection Model, *IEEE Transactions on Software Engineering*, SE-13:222-232, 1987.
8. H.S. Javitz, and A. Valdes, The NIDES Statistical Component: Description and Justification, *Technical Report, Computer Science Laboratory, SRI International*, 1993.
9. MINDS, Minnesota Intrusion Detection System, www.cs.umn.edu/~aleks/MINDS.
10. M. Joshi, R. Agarwal, V. Kumar, PNrule, Mining Needles in a Haystack: Classifying Rare Classes via Two-Phase Rule Induction, *Proceedings of ACM SIGMOD Conference on Management of Data*, May 2001.
11. M. Joshi, V. Kumar, CREDOS: Classification using Ripple Down Structure (A Case for Rare Classes), *in review*.
12. M. Joshi, V. Kumar, R. Agarwal, Evaluating Boosting Algorithms to Classify Rare Classes: Comparison and Improvements, *First IEEE International Conference on Data Mining*, San Jose, CA, 2001.
13. M. Joshi, R. Agarwal, V. Kumar, Predicting Rare Classes: Can Boosting Make Any Weak Learner Strong?, *Proceedings of Eight ACM Conference ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Edmonton, Canada, 2002.
14. A. Lazarevic, N. Chawla, L. Hall, K. Bowyer, SMOTE-Boost: Improving the Prediction of Minority Class in Boosting, *AHPCRC Technical Report*, 2002.
15. C. C. Aggarwal, P. Yu, Outlier Detection for High Dimensional Data, *Proceedings of the ACM SIGMOD Conference*, 2001.
16. E. Knorr, R. Ng, Algorithms for Mining Distance-based Outliers in Large Data Sets, *Proceedings of the VLDB Conference*, 1998.
17. M. M. Breunig, H.-P. Kriegel, R. T. Ng, J. Sander, LOF: Identifying Density-Based Local Outliers, *Proceedings of the ACM SIGMOD Conference*, 2000.
18. L. Ertoz, E. Eilertson, A. Lazarevic, P. Tan, J. Srivastava, V. Kumar, V., P. Dokas: The MINDS - Minnesota Intrusion Detection System, *in review*.
19. SNORT Intrusion Detection System. www.snort.org.