Journal of Network and Computer Applications 66 (2016) 166–179

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

journal homepage: www.elsevier.com/locate/jnca

A likelihood ratio anomaly detector for identifying within-perimeter computer network attacks

Justin Grana [a,**], David Wolpert [c,*], Joshua Neil [d,*], Dongping Xie [a,*], Tanmoy Bhattacharya [b,*], Russell Bent [b,*]

[a] American University, Economics Department, 4400 Massachusetts Ave NW, Washington, DC 20016, United States
[b] Los Alamos Laboratory, PO Box 1663, MS B264, Los Alamos, NM, United States
[c] Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe, NM 87501, United States
[d] Ernst and Young, 370 17th St, Denver, CO 80202, United States

ARTICLE INFO

Article history:
Received 21 October 2015
Received in revised form
14 January 2016
Accepted 8 March 2016
Available online 11 March 2016

Keywords:
Anomaly detection
Computer network defense
Cyber security
Likelihood ratio detection
ROC analysis
Model misspecification

ABSTRACT

The rapid detection of attackers within firewalls of enterprise computer networks is of paramount importance. Anomaly detectors address this problem by quantifying deviations from baseline statistical models of normal network behavior and signaling an intrusion when the observed data deviates significantly from the baseline model. However, many anomaly detectors do not take into account plausible attacker behavior. As a result, anomaly detectors are prone to a large number of false positives due to unusual but benign activity. This paper first introduces a stochastic model of attacker behavior which is motivated by real world attacker traversal. Then, we develop a likelihood *ratio* detector that compares the probability of observed network behavior under normal conditions against the case when an attacker has possibly compromised a subset of hosts within the network. Since the likelihood ratio detector requires integrating over the time each host becomes compromised, we illustrate how to use Monte Carlo methods to compute the requisite integral. We then present Receiver Operating Characteristic (ROC) curves for various network parameterizations that show for any rate of true positives, the rate of false positives for the likelihood ratio detector is no higher than that of a simple anomaly detector and is often lower. We conclude by demonstrating the superiority of the proposed likelihood ratio detector when the network topologies and parameterizations are extracted from real-world networks.

1. Introduction

Many existing systems designed to detect intrusions into computer networks monitor data streams only at the perimeter of the network. In addition, many network intrusion detection systems, such as snort (Roesch et al., 1999), are signature based, meaning that every communication entering or leaving the network is examined for matches to a database of signatures, or indicators of compromise. At this point, the long list of breaches to corporate networks (Krebs, 2014, 2015) speaks loudly to the insufficiency of these methods. Attackers are able to innovate rapidly in order to avoid signature schemes, and penetrate these perimeter systems seemingly at will. Therefore, there is a pressing need to identify attackers *within* network perimeters, and to do so using behavioral methods rather than signatures.

Anomaly detectors—a model-based approach—show promise in detecting within-perimeter attacks. In general, anomaly detectors quantify "normal" network behavior, and when observed behavior significantly deviates from the baseline model, an intrusion is signaled. As a simple example, consider an anomaly detector that models a computer network as a directed graph where nodes are users within a network and edges represent a communication channel between users. The detector is then calibrated such that it specifies the average rate of packet transfer along each edge. When the observed rate of packet transfers is sufficiently different from the calibrated rate of packet transfers, the detector signals an intrusion.

In practice, many reported anomalies end up being false, reflecting behavior that is unusual but benign. This is due in part to an incomplete specification of normal network behavior in the null hypothesis as well as the difficulty in modeling and predicting the behavior of humans that interact over the network. There are at least two approaches in addressing this issue. The first is to

* Corresponding authors.
** Principal author.
E-mail addresses: jg3705a@student.american.edu (J. Grana), david.h.wolpert@gmail.com (D. Wolpert), joshua.neil@ey.com (J. Neil), xdp668@gmail.com (D. Xie), tanmoy@santafe.edu (T. Bhattacharya), rbent@lanl.gov (R. Bent).

http://dx.doi.org/10.1016/j.jnca.2016.03.008
1084-8045/

improve the specification of the network under normal conditions.[1] The second is to develop a model of attacker behavior and compare the probability of the observed behavior under the hypothesis that the network has been compromised against the hypothesis that the network is functioning under normal conditions. With an accurately specified attacker model, such an approach would rule out benign but unusual activity as being malicious since it is not consistent with attacker behavior. Our work in this paper takes the second approach. More explicitly, to our knowledge this paper is the first to incorporate an exact parametric specification of attacker behavior into a likelihood ratio detector for identifying malicious traversal activity within a network perimeter.

The challenge is in how to model the behavior of a network that has been penetrated without pre-supposing attacker methods, since these methods evolve rapidly. To see how this might be done, consider a common attack conducted on an enterprise network. First, a Phishing email or set of emails, containing either a malicious attachment or a link to an internet host serving malware, is sent to the target network. Click rates on Phishing emails, even after enterprise training is conducted, can be as high as 50% (Kumaraguru et al., 2009), providing a high-confidence intrusion vector.

At this point, the firewall is penetrated and the attacker has control of an initial host in the target network. The attack is far from complete since the initially compromised host is not the primary target of the attacker. Instead, the attacker seeks to penetrate the network and access key servers. However, since credentials are typically required to access these servers, the attacker undergoes a process known as *lateral movement* to move among hosts collecting these credentials (Kent and Liebrock, 2013). This means that there is a definite sequence in the movement of the attacker across the network, from computers with low value (for any of the goals of inserting malware, extracting data, or stealing credentials) to computers with higher value, such as data servers and active directories. This will be true *no matter what precise methods the attacker uses* to achieve that movement. As a result, the attacker's traversal will leave a trace of increasing network traffic going from low value computers to progressively higher value ones. Therefore, an increase in network traffic along paths from low value to high value nodes in a network can be used as the basis of a model of network behavior once it has been penetrated.

The approach in this paper is parallel to that of Jiang et al. (2014a) in that we first propose a model of attacker behavior and novel detection criteria based on a likelihood ratio. For various network parameterizations, we simulate network activity in both the normal and compromised state. We then employ receiver operating characteristic (ROC) curves to show that the proposed likelihood ratio detector outperforms a simple anomaly detector that does not exploit information regarding the traversing nature of an attack. In addition we develop the Monte Carlo techniques used to approximate the relevant integrals in computing our proposed likelihood ratio. Finally, we extract topologies and parameter data from real-world networks and then simulate attacker behavior. The results show that in real-world networks, our proposed likelihood ratio detector is superior to the simple anomaly detector.

[1] This work is similar to reducing prediction error of network traffic. See Jiang et al. (2015a, 2014b) and Jiang et al. (2011), for work that focuses on improving the modeling and prediction of normal network behavior.

2. Background

Model-based anomaly detection proceeds by modeling and estimating the parameters, $\hat{\theta}$, of a computer network under normal conditions. Next, given a dataset $\mathbf{D}$ under question, the likelihood of the parameters given the data can be evaluated: $\mathcal{L}(\hat{\boldsymbol{\theta}}|\mathbf{D})$. A generalized likelihood ratio test (GLRT) can then be used to infer whether a more likely alternative parameterization is present given data $\mathbf{D}$

$$GLRT = \frac{\mathcal{L}(\hat{\boldsymbol{\theta}}|\mathbf{D})}{\sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}}\mathcal{L}(\boldsymbol{\theta}|\mathbf{D})}$$

where $\boldsymbol{\Theta}$ is an alternative parameter space. Typically, we choose what data $\mathbf{D}$ to collect in order to facilitate statistical discovery of security breaches. For example, the network model under normal conditions might be a graph connecting computers (nodes or hosts) with edges representing parameterized time-series of traffic. The data collected would then be communications between nodes. When the observed communication pattern is different from the parameterized time-series, the anomaly detector would sound an alarm. Additionally, since attacks typically cover multiple nodes and edges, subgraphs can be used to group data from multiple nodes and edges into $\mathbf{D}$ for increased detection power. Such graph based methods include Borgwardt et al. (2006), Eberle et al. (2010), Neil et al. (2013); Staniford-Chen et al. (1996), and Djidjev et al. (2011)).

If we know that the attacker behaves according to a specified alternative parameter vector, say $\boldsymbol{\theta}_A$, then the uniformly most powerful test for rejecting the null hypothesis that no attack is present is a likelihood ratio test where $\boldsymbol{\theta}_A$ is used in the denominator. That is, if we know the attacker is behaving according to $\boldsymbol{\theta}_A$, the power of the test is maximized when using the test statistic

$$\overline{GLRT} = \frac{\mathcal{L}(\hat{\boldsymbol{\theta}}|\mathbf{D})}{\mathcal{L}(\boldsymbol{\theta}_A|\mathbf{D})}. \tag{1}$$

However, the set of alternatives $\boldsymbol{\Theta}$ is typically under-specified. In other words, anomaly detectors do not specify exact attacker behavior but simply restrict the parameter space of alternatives. A representative example of such a detector is the Modeled Attack Detector (MAD) given in Thatte et al. (2008). In their work, the authors consider the rate of incoming traffic in order to detect a Distributed Denial of Service (DDoS) attack. They assume that under normal conditions, the number of incoming connections can be modeled by a Poisson distribution with average rate of messages per unit time of λ_B. The authors treat λ_B as a known and calibrated parameter. Therefore, given a sequence of incoming connections (i.e. one unit of network traffic) $\mathbf{D} = \{d_1, d_2 \ldots d_N\}$ per unit time interval, the probability of observing $\mathbf{D}$ under the hypothesis that H_0=*no attack is taking place* is given by

$$P(\mathbf{D}|H_0) = \prod_{i=1}^{N} \frac{e^{-\lambda_B}\lambda_B^{d_i}}{d_i!}. \tag{2}$$

The authors assume that under a DDoS attack the network receives additional malicious connections at fixed, deterministic time intervals but at an unknown rate. If the rate was known, the probability of an observed sequence under the hypothesis that H_1=*DDoS attack is occurring* is given by

$$P(\mathbf{D}|H_1) = \prod_{i=1}^{N} \frac{\lambda_B^{d_i-\lambda_m} e^{-\lambda_B}}{(d_i-\lambda_m)!} \tag{3}$$

where λ_m is the rate at which the network receives malicious connections. In reality, λ_m—the rate under the alternative hypothesis—is unknown so a simple likelihood ratio test is