

# Information-Theoretic Detection of Masquerade Mimicry Attacks

Juan E. Tapiador and John A. Clark  
Department of Computer Science, University of York  
Heslington, York, YO10 5DD, UK  
{jet, jac}@cs.york.ac.uk

**Abstract**—In a masquerade attack, an adversary who has stolen a legitimate user’s credentials attempts to impersonate him to carry out malicious actions. Automatic detection of such attacks is often undertaken constructing models of normal behaviour of each user and then measuring significant departures from them. One potential vulnerability of this approach is that anomaly detection algorithms are generally susceptible of being deceived. In this paper, we first investigate how a resourceful masquerader can successfully evade detection while still accomplishing his goals. We then propose an algorithm based on the Kullback-Leibler divergence which attempts to identify if a sufficiently anomalous attack is present within an apparently normal request. Our experimental results indicate that the proposed scheme achieves considerably better detection quality than adversarial-unaware approaches.

**Index Terms**—Anomaly detection; insider threats; masqueraders; mimicry attacks; Kullback-Leibler divergence.

## I. INTRODUCTION

It has been long acknowledged that one of the worst threats in computer security is that posed by internal users who misuse their privileges for malicious purposes. Such actions could potentially result in enormous damages for the organisation, arguably far greater than those expected from external adversaries. Classical access control models can partially alleviate the risks associated with internal security issues, but the reality of many systems is much more complex [20]: specifying good security policies is very hard; policies are frequently and purposely bypassed to get the job done; sharing information among different organisations is too often necessary and current security models are very poor at controlling the potential repercussions of wrong-sharing; etc. As a consequence, it has been recognised that access control systems are necessary measures, but clearly insufficient to deal with all the complexities posed by insider attacks. Research on this area has been in place for the last 20 years and, to some extent, has proliferated lately (see e.g. [33], [3], [2], [7] for a few examples of recently reported research initiatives).

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

One traditional way of classifying insiders is as *traitors* and *masqueraders* [35]. A traitor is a user who already enjoys some privileges within the system and whose purposes will affect negatively the security properties of the organisation’s information and systems. A masquerader, on the contrary, is an (often external) attacker who succeeds in obtaining a legitimate user’s credentials and attempts to use the stolen identity to carry out malicious actions (e.g. credit card fraudsters).

Virtually all existing masquerade detection approaches rely upon one key observation: “*behaviour is not something that can be easily stolen*” [35]. Profiling users behaviours could therefore establish models of normalcy such that deviations from them would presumably indicate the presence of an impersonation attempt. The idea of using anomalies as proxies for attacks has been extensively studied in various security domains and, albeit generally useful, is not free from drawbacks and controversies [38]. Beyond that, there are inherent limitations in using an anomaly detection algorithm as the basis for masquerade detection. Firstly, profiles are ultimately derived from data provided by the user, who might well be in the business of forcing the learning process to build something undesirable (e.g. a model of normalcy such that future misbehaviours will not be identified). Some works (e.g. [22], [6]) have already pointed out that the data used to train a security application could be actively manipulated by an adversary. When applied to such adversarial domains, learning algorithms should be conveniently adapted, but research in this area is still scarce. A second threat stems from the fact that knowledge of some details about the detection process facilitates evasion. Furthermore, in general it is reasonable to assume that such information is public, as it can generally be possible for an adversary to obtain it by careful experimentation with the system [27].

In this paper we investigate some of these threats in the context of masquerade detection. More specifically, we make the following contributions:

- We show that an adversary with sufficient knowledge about the internals of the employed algorithms can successfully evade the masquerade detection process by means of carefully constructed attacks that are not identified as anomalous. Such *mimicry* attacks have been demonstrated before in other contexts (e.g. intrusion detection and spam filtering). We implement and empirically evaluate their effectiveness against some widely

used masquerade detection algorithms.

- We propose a defence mechanism based upon the idea of separating, in a probabilistic sense, the attack from the padding sequence in a block of data. The proposed algorithm makes use of the Kullback-Leibler divergence and does not rely on any assumptions about the attack (other than, once isolated, it is anomalous). We empirically demonstrate the improvement achieved through this method in terms of detection quality.

The remainder of this paper is organised as follows. In Section II we discuss previous works on masquerade detection and mimicry attacks. In Section III we introduce mimicry attacks in the context of a masquerade detection scenario. We describe various methods for generating such attacks and empirically evaluate their success in evading detection. In Section IV we describe and evaluate the PPI algorithm. The results obtained over a dataset containing normal samples, as well as mimicry and non-mimicry masquerade attacks, are shown in Section V. Finally, Section VI concludes the paper.

## II. BACKGROUND

### A. Masquerade Detection

Schonlau *et al.* presented in [37] a dataset<sup>1</sup> for the evaluation of different masquerade detection methods. The dataset consists of sequences of truncated UNIX commands corresponding to the normal activity of 70 users and collected over a period of several months. Fifty users were chosen as intrusion targets and the remaining 20 as masqueraders. The dataset provides 15000 commands for each one of the target users, organised as follows. The first 5000 commands contain actions actually executed by the user. The remaining 10000 commands are split in blocks of 100 commands, each block being either self (i.e., corresponding to the user) or non-self. Non-self blocks represent masquerade attempts and contain commands issued by a randomly chosen masquerader user.

As introduced in [37], the main task for a masquerade detection algorithm is to accurately identify non-self blocks as anomalous (and, therefore, implicitly mark them as masquerade attempts), while correctly classifying the self blocks as belonging to the user. The work in [37] explores the performance of six different machine learning algorithms for this task in the so-called SEA configuration: each user's first 5000 commands are used for training and the remaining 10000 commands for testing on a per-block basis.

A series of papers by Maxion *et al.* improved on the results reported in [37] and provided further analysis of the masquerade detection problem. In [30] it is shown how the naïve Bayes classifier achieves much better performance than previously proposed schemes. The paper also provides an excellent articulation of why some users are more difficult to attack than others and introduces a new experimental setting called 1v49, as opposed to the original SEA experiment described in [37]. In this setting, the first 5000 commands of each user are employed to build the normalcy model. In testing mode,

self data is extracted from the remaining 10000 commands of the user, and the other 49 users' first 5000 commands as non-self data. The 1v49 experiment is arguably a better way of evaluation the performance of detection algorithms. We refer the reader to [30] for a rationale behind this claim.

Further works explored the consequences of using datasets enriched with information other than commands [31], as well as the effects of applying privacy-preserving sanitisation strategies over the data [23]. Wang and Stolfo argued in [42] that detection methods based on one-class training (i.e., relying only on self data) are more appropriate for a real-world setting. They showed that both naïve Bayes and Support Vector Machine (SVM) algorithms attain similar results in a one-class configuration than by using two-class data.

Work on masquerade detection has proliferated over the last decade, especially concerning the study of different detection strategies. Some of the proposals include information-theoretic approaches [1], [10], hidden Markov models [34], or sequence- and text-mining [32], [26], [5], [16] schemes, among others. Despite the diversity of principles behind these methods, the results reported show that their effectiveness is similar.

### B. One-Class Naïve Bayes Masquerade Detection

The naïve Bayes (NB) classifier [18] is a supervised learning algorithm which has been used in a wide range of applications. NB is often a very attractive solution because of its simplicity, efficiency and excellent performance. It uses the Bayes rule to estimate the probability that an instance  $x = (x_1, \dots, x_m)$  belongs to class  $y$  as

$$P(y|x) = \frac{P(y)}{P(x)} P(x|y) = \frac{P(y)}{P(x)} \prod_{i=1}^m P(x_i|y) \quad (1)$$

so the class with highest  $P(y|x)$  is predicted. (Note that  $P(x)$  is independent of the class and therefore can be omitted.) The naïvety comes from the assumption that in the underlying probabilistic model all the features are independent, and hence  $P(x|y) = \prod_{i=1}^m P(x_i|y)$ .

NB has been used before in the context of masquerade detection [30], [42], particularly using Schonlau *et al.*'s dataset. In the multinomial model (or bag-of-words approach), every block of commands  $B$  to be classified is represented by a vector of attributes  $[n_1(B), \dots, n_m(B)]$ , where  $n_i(B)$  is the number of times command  $c_i$  appears in the block. The probability  $P(y|B)$  given by (1) can be then computed as

$$P(y|B) = P(y) \prod_{i=1}^m P(c_i|y)^{n_i(B)} \quad (2)$$

The probabilities  $P(c_i|y)$  are derived from a training set consisting of labelled instances for all possible classes (i.e., from each user's first 5000 commands), and the priors  $P(y)$  are often ignored. In order to control the sensitivity to previously unseen commands, it is convenient to ensure that all commands appear with non-zero probability even if some of them are not present at all in the training set. This can be

<sup>1</sup>Publicly available at <http://www.schonlau.net>.

achieved by using an additive smoothing over the estimated probabilities

$$P(c_i|y) = \frac{\sum_{B \in \mathcal{T}(y)} n_i(B) + \alpha}{|B| \cdot |\mathcal{T}(y)| + \alpha \cdot m} \quad (3)$$

where  $\mathcal{T}(y)$  is the training set for class  $y$  and  $\alpha$  the smoothing parameter.

For convenience, in this work we will use minus the logarithm of (2) rather than the probability itself as basic indicator of the nature of a block (again, ignoring the priors):

$$\text{score}(B) = -\log P(y|B) = -\sum_{i=1}^m n_i(B) \log P(c_i|y) \quad (4)$$

The result can be seen as an anomaly score: the higher its value, the more anomalous the block is, and vice versa.

Following [42], in a one-class setting the training set for each user consists exclusively of data corresponding to self activities. Since a profile of non-self behaviour is not required, the detection is performed by simply comparing the probability of a block being self (or, equivalently, the anomaly score) to a threshold. Such a threshold can be adjusted to control the false and true positive rates, and the resulting ROC (Receiver Operating Characteristic) curve provides a way of measuring the detection quality.

### C. Mimicry Attacks

The basic idea behind mimicry attacks is to evade an anomaly detector by altering the attack to make it look normal. Successful evasion is attained when the (modified) characteristics of the block of data fit the normal profile used by the detector, while simultaneously preserving the intended attack semantics. Introducing such transformations generally requires the attacker to know both the detection algorithm and the model of normalcy in use.

Early works on mimicry attacks mainly targeted host-based IDSs, in particular systems based on the analysis of system call sequences as introduced by Forrest *et al.* [13], [14], [19], [45]. Wagner and Soto [41] and Tan *et al.* [39], [40] developed various strategies for generating mimicry attacks against such detectors. Subsequent works (e.g., [15], [17], [24], [21] among others) further explored this idea, mainly focusing on the problem of how to generate a mimicry sequence that evades detection and achieves the attacker goals. The task is generally computationally hard, and techniques drawn from domains such as model checking, code analysis, or genetic programming have proven useful.

Similar ideas have also been investigated in the area of network-based IDS, where detection is accomplished by analysing payload features such as byte distributions or, more generally,  $n$ -gram or more complex models (such as, e.g., [43], [44], [25], [28], [29], [8], [9]). Fogla *et al.* introduced in [11], [12] polymorphic blending attacks, where the main idea is to generate each attack instance in such a way that its statistics match the profile of normalcy used by an anomaly detector. Such attacks would therefore be able to evade both signature- and anomaly-based IDSs. (Again, it is shown that the problem

of generating such instances is NP-complete, though some heuristic techniques are of help.)

## III. MASQUERADE MIMICRY ATTACKS

In this section we introduce mimicry attacks in the context of a masquerade detection problem. We consider an adversary who intends to launch an attack consisting of a sequence of actions or commands. We make three fundamental assumptions about this process:

- 1) *Perfect knowledge*: The adversary knows perfectly the detection algorithm being used and all the relevant parameters, as well as the model of normalcy for the user whose system account is impersonating. Alternatively, the adversary could be the user himself attempting to launch an attack without being spotted by the anomaly detector.
- 2) *Non-poisoned detector*: The detector has been trained with attack-free data, so we do not consider the possibility of frog-boiling attacks (e.g. [4]) or other forms of evasion based on training the detection algorithm with carefully crafted data.
- 3) *Attack padding*: The attack sequence must be executed within a block, though not necessarily in a contiguous way. Thus, the adversary could insert padding commands at any point of the attack sequence. We do not put any restriction on the type, length, position, or number of padding sequences (other than both attack and padding must add up to a block size).

### A. Notation

We will denote sequences (or blocks) of commands by capital letters, in particular  $A$  for attacks,  $P$  for padding, and  $B$  for entire blocks. The symbol  $|\cdot|$  denotes the length of a sequence. Sequences will be treated as arrays, so  $S(i)$  denotes the  $i$ -th command in the sequence. The probability density function of a sequence will be specified by a calligraphic font, e.g.,  $\mathcal{A}$ ,  $\mathcal{P}$ ,  $\mathcal{B}$ , etc. Thus,  $S(c_i)$  will denote the frequency of command  $c_i$  in sequence  $S$ .

### B. Evading OCNB

Consider an attack consisting of  $|A| \leq |B|$  commands, so the number of padding commands the adversary must generate is  $|B| - |A|$ . We assume that the attack sequence will contribute significantly to identify the block as anomalous. For example, in the case of a detector based on the OCNB classifier described above, this translates into a very low probability induced by the commands comprising the attack. In this case, the optimal padding strategy for the attacker consists of filling the block with the command  $c_{max} = \arg \max_{c_i} \mathcal{M}(c_i)$ ,  $\mathcal{M}$  being the model of normalcy, as this will cause the maximum possible increment in the probability of the block being classified as normal given the attack. Despite being optimal against OCNB, we will not consider such a strategy here since the results might not be generally useful for different detection algorithms. We shall instead look into the more general strategy of producing a padding sequence such that

the histogram of the resulting block (attack plus padding) is statistically indistinguishable from that observed during training. Such attacks would be presumably effective against a wider range of masquerade detection algorithms.

1) *Attack Generation*: We will assume that the distinguishability metric we attempt to minimise is  $\sum_{c_i} |\mathcal{B}(c_i) - \mathcal{M}(c_i)|$ , where  $\mathcal{B}$  and  $\mathcal{M}$  are the histogram of the block and the normalcy model, respectively, and the sum is taken over the available set of commands. We will also restrict ourselves to the case where the attack sequence is immutable, i.e. no command in it can be deleted or replaced by other. In this case, it is not difficult to see that the optimal strategy for generating the padding sequence consists of:

- (i) Compute the difference histogram  $\mathcal{D}(c_i) = \mathcal{M}(c_i) - \mathcal{A}(c_i)$ .
- (ii) Add to the padding sequence  $|B| \cdot \mathcal{D}(c_m)$  instances of the command  $c_m = \arg \max_{c_i} \mathcal{D}(c_i)$ .
- (iii) Set  $\mathcal{D}(c_m) = 0$  and repeat step (ii) until no more padding is needed.

Alternatively, a suboptimal (but certainly much faster) strategy consists of generating the padding by just sampling from the model of normalcy  $\mathcal{M}$ . (The procedure is straightforward once the inverse of cumulative distribution,  $F_{\mathcal{M}}^{-1}$ , is computed.)

To build the final block of commands, we first select  $|A|$  different random positions of the block and place one attack command in each of them, respecting the original order in the attack sequence. The remaining empty positions are then filled up with the padding commands previously generated in no particular order.

2) *Results*: In order to quantify the performance of such attacks, we have conducted the following experiment using the Schonlau *et al.*'s dataset. Given a user  $u$ , we first repeat the 1v49 experiment and record the raw scores issued by OCNB. We then plot the distribution of the scores for both self and non-self blocks. This serves to illustrate visually the discriminative capability of the classifier: the higher the overlapping between both distributions, the lower the detection quality. As an example, Fig. 1 shows the distribution of the scores given by OCNB to user 2's self and non-self blocks (first two boxplots).

"Attacks" are generated by randomly choosing a sequence of  $|A|$  commands from a block belonging to the training dataset of a user other than  $u^2$ . This sequence is then placed into an empty block, and the remaining  $100 - |A|$  positions are filled with a padding sequence obtained by following the optimal strategy described above. The score for the block as given by OCNB is computed and the procedure is repeated 10000 times for randomly generated attacks. The ten rightmost boxplots in Fig. 1 show the score distribution for attacks of length 10, 20, ..., 100. It is observed that the bulks of the self

<sup>2</sup>Note that such sequences are not by any means *actual* attacks. However, our emphasis here is not on the consequences of the adversary's actions in a real setting, but rather on the assumption that attacks are anomalous events which nonetheless might be conveniently camouflaged to avoid detection. For this purpose, the methodology followed should do as far as the detection of such *concealed anomalies* is concerned.




Fig. 1. Distribution of OCNB scores for user 1 including mimicry attacks of various lengths.

TABLE I  
AVERAGE DETECTION RATE OF MIMICRY ATTACKS USING OCNB.

$ A  = 10$	$ A  = 20$	$ A  = 30$	$ A  = 40$	$ A  = 50$	$ A  = 60$
0.081	0.206	0.314	0.407	0.474	0.521

and non-self distributions are largely non-overlapping, and a threshold around 500 might serve to detect most nonself sequences with some rate of false positives and negatives. Mimicry attacks (ten rightmost plots) of low length present a score distribution below any reasonable detection threshold, thus being essentially impossible to detect. An increasing attack length generates more anomalies per block and also leaves less space available for padding, which translates into a greater score and, consequently, more chances of detection. The plots for most users are completely analogous.

In global terms, OCNB performs rather poorly in detecting this form of attacks. Table I gives the average detection rate of mimicry attacks of length up to 60 commands computed for the 50 users in the dataset. The detector for each user was tuned so as to limit the false positive rate to a maximum of 5%, and the average is computed for the 50 users. The majority of the attack blocks passed unnoticed by the detector, only approaching a detection rate higher than 50% (which is still remarkably low) when the attack sequence comprises more than half the block length.

#### IV. PROBABILISTIC PADDING IDENTIFICATION

In this section we develop an algorithm which attempts to separate the attack from the padding sequence in a given block of commands. The process will be carried out with the help of the normalcy model presumably used to generate the padding, but without any further knowledge about the attack

length (which, incidentally, could be zero). We first review some properties of the Kullback-Leibler divergence, a concept which will be central in our algorithm.

### A. Kullback-Leibler Divergence

The Kullback-Leibler (KL) divergence is a non-symmetric measure of the difference between two probability distributions. If  $P$  and  $Q$  are two discrete distributions, then the KL divergence of  $Q$  from  $P$  is defined by

$$D_{KL}(P \parallel Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)} \quad (5)$$

Note that  $D_{KL}(P \parallel Q)$  can be rewritten as

$$\begin{aligned} D_{KL}(P \parallel Q) &= - \sum_i P(i) \log Q(i) + \sum_i P(i) \log P(i) \\ &= H(P, Q) - H(P) \end{aligned} \quad (6)$$

where  $H$  denotes the entropy. Consequently,  $D_{KL}$  admits a simple interpretation as the expected number of extra bits necessary to encode samples taken from  $P$  when using a code based on  $Q$  rather than one based on  $P$ .

From a different perspective, the KL divergence can also be seen as the expected discrimination information between two hypothesis. Given a sample  $x$  and two possible hypothesis  $H_0$  and  $H_1$ ,  $D_{KL}(P(x|H_1) \parallel P(x|H_0))$  provides the mean information per sample for discriminating in favour of  $H_1$  against  $H_0$ , given that  $H_1$  is true. Or, in other words, it measures as the amount of evidence for  $H_1$  over  $H_0$  to be expected per sample.

### B. The PPI Algorithm

Based on the properties of the KL divergence, we next describe an algorithm to probabilistically identify the padding portion of a block of commands. Assume that  $A$  and  $P$  are the attack and padding portions of a block  $B$ , and assume that  $\mathcal{M}$  is the normalcy model for a given user. The algorithm relies upon two main observations: 1)  $\mathcal{A}$  is sufficiently different from  $\mathcal{M}$  (otherwise it would not be necessary to add padding); and 2)  $\mathcal{P}$  is highly similar to  $\mathcal{M}$ , as it has to compensate for the effects of  $\mathcal{A}$ . Note that the problem of extracting  $P$  from  $B$  is further complicated by the fact that we generally do not know the length of the attack.

Our approach consists of identifying subsets  $\hat{P}, \hat{A} \subseteq B$ , with  $\hat{P} \cup \hat{A} = B$  and  $\hat{P} \cap \hat{A} = \emptyset$ , such that  $D_{KL}(\hat{P} \parallel \mathcal{M})$  is very low and, simultaneously,  $D_{KL}(\hat{A} \parallel \mathcal{M})$  is very high. An exhaustive search would require to check  $2^{|B|}$  possible subsets and compute two KL divergences for each one of them, which is clearly impractical. Instead, we propose a greedy strategy where suitable candidates for  $\hat{P}$  and  $\hat{A}$  are identified in one single pass over the block.

The algorithm (detailed below) attempts to identify the portion  $\hat{P}$  of  $B$  that best fits the model. A vector  $C$  is used to indicate whether command  $B(i)$  is padding or not, so at each step such a vector partitions the block into two sequences,  $\hat{P}$  and  $\hat{A}$ . The procedure DIFFKL computes the KL divergences between each of these sequences and the model

$\mathcal{M}$ , and returns the absolute value of the difference. At each step, the PPI algorithm is governed by a simple rule: add the  $i$ -th command to the tentative padding if, by doing so, the increment of the differential KL divergence is greater than that corresponding to not adding the command. The rationale behind such a rule can be better understood by observing that

$$\begin{aligned} |D_p - D_a| &= \left| \sum_i \hat{P} \log \frac{\hat{P}}{\mathcal{M}} - \sum_i \hat{A} \log \frac{\hat{A}}{\mathcal{M}} \right| \\ &= \left| H(\hat{A}) - H(\hat{P}) + H(\hat{P} - \hat{A}, \mathcal{M}) \right| \end{aligned} \quad (7)$$

i.e., a command is accepted as belonging to padding if that translates into a higher difference of the entropies of  $\hat{P}$  and  $\hat{A}$ , plus a higher difference in the cross entropy between  $(\hat{A} - \hat{P})$  and the model  $\mathcal{M}$ . Implicit in this utility function is the idea that padding and attack have different *information content*, hence its use to identify both of them.

---

### Algorithm 1 PPI

---

**Input:** Block  $B$ , model  $\mathcal{M}$

**Output:** Boolean vector:  $C(i) = \mathbf{true}$  if  $B(i)$  is padding

- 1: Initially  $C(i) \leftarrow \mathbf{false}$  for all  $i$ .
  - 2: **for**  $i = 1$  **to**  $|B|$  **do**
  - 3:    $\bar{d} = \text{DIFFKL}(C, B, \mathcal{M})$
  - 4:    $C(i) \leftarrow \mathbf{true}$
  - 5:    $d = \text{DIFFKL}(C, B, \mathcal{M})$
  - 6:   **if**  $d \leq \bar{d}$  **then**
  - 7:      $C(i) \leftarrow \mathbf{false}$
  - 8:   **end if**
  - 9: **end for**
  - 10: **return**  $P = \text{commands } B(i) \text{ such that } C(i) \text{ is } \mathbf{true}$
- 

---

### Algorithm 2 DIFFKL

---

**Input:** Boolean vector  $C$ , block  $B$ , model  $\mathcal{M}$

**Output:** Difference of K-L divergences

- 1:  $\hat{A} \leftarrow$  PDF of those  $B(i)$  such that  $C(i)$  is **false**
  - 2:  $\hat{P} \leftarrow$  PDF of those  $B(i)$  such that  $C(i)$  is **true**
  - 3:  $D_a \leftarrow D_{KL}(\hat{A} \parallel \mathcal{M})$
  - 4:  $D_p \leftarrow D_{KL}(\hat{P} \parallel \mathcal{M})$
  - 5: **return**  $|D_p - D_a|$
- 

A simpler and more natural approach would appear to be to accept the  $i$ -th command as padding if that decreases the KL divergence between the candidate  $\hat{P}$  and  $\mathcal{M}$ . This alternative, to which we will refer as PPI KL as opposed to the previously discussed PPI DIFFKL, turns out to be less effective in practice. We next discuss some experimental results.

### C. Evaluation

In this section we report results of the evaluation of the PPI algorithm over masquerade mimicry attacks only. Next section provides details on the overall behaviour over a dataset composed of both attacks and self samples.

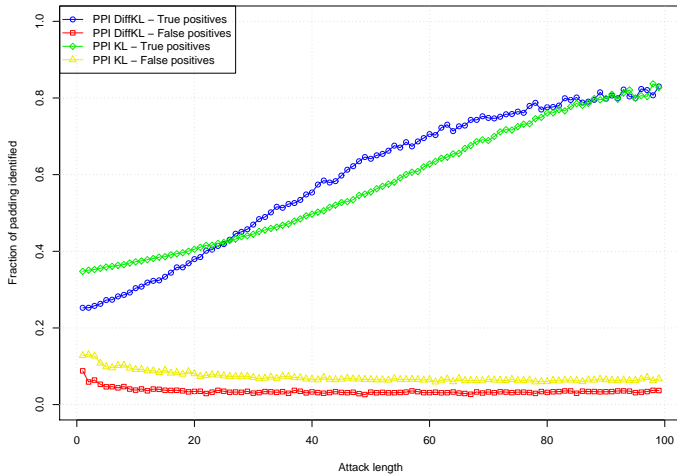


Fig. 2. (In colour in the electronic version.) Accuracy of the PPI algorithm in identifying the padding portion of attacks of various lengths.

For each possible attack length from 1 to 100, we have generated 10000 mimicry attacks following the procedure described in Section III-B. Each attack is analysed by the PPI algorithm, which returns the estimated positions of the padding. We then compute how many true positives (i.e., true padding positions correctly identified) and false positives (i.e., attack positions incorrectly identified as padding) are produced. Fig. 2 shows the figures for both PPI DIFFKL and PPI KL. PPI DIFFKL performs better in terms of FP, with a rate below 5% except for extremely short attacks. As far as TP are concerned, PPI DIFFKL outperforms PPI KL for attacks of length approximately 25 or greater. We suspect that the reason for such a behaviour is related to the fact that PPI DIFFKL makes use of both padding and attack information. While this certainly helps the algorithm to keep down the FP rate, it turns out to be a drawback when dealing with blocks when the attack portion is very short. Regarding TP, the identification rate increases with the attack length almost linearly, up to a limit of around 80%. As we will see later, even these imperfect figures will be of help to assess the likelihood of an apparently normal block containing a mimicry attack.

The algorithm is reasonably fast. In our experiments with an average laptop, the original OCNB requires (average  $\pm$  standard deviation)  $0.01 \pm 0.28$  ms per block. The inclusion of the PPI increases this time up to  $38.22 \pm 5.80$  ms. In a real-world system, these figures do not constitute a problem, especially when considering that the analysis is performed every 100 user actions.

## V. MASQUERADE MIMICRY ATTACK DETECTION

In this section we describe how the PPI algorithm can be integrated within an anomaly detector to improve the identification of mimicry attacks. Even though we will limit




Fig. 3. (In colour in the electronic version.) Distribution of scores for attack (red) and padding (blue) sections in blocks containing attacks of various lengths.

our discussion to the case of OCNB, the same principle could be extended to a wider family of detectors.

In a first experiment, we generated 10000 blocks  $B$  containing mimicry attacks and applied the PPI algorithm to each one of them. We then have computed the anomaly score, given by (4), to each one of 2 sequences (attack and padding) returned by the algorithm *separately*. The purpose of this is to measure the contribution towards the overall anomaly score of the identified padding and attack portions. (Recall that the overall score is merely the sum of these two scores.) Fig. 3 shows the distribution of anomaly scores for the attack and padding sections for attacks of various lengths. As expected, padding sequences map to very low scores (around 50) which, besides, are almost independent of the attack length. On the contrary, the attack portion generally receives a much higher score, which obviously increases with the attack length.

When applied to self blocks, the result is completely similar. Nevertheless, in this case the identified “attack” portions correspond to false negatives of the PPI algorithm. These, however, are comparatively very few, a fact that will facilitate the construction of a combined anomaly score capable of detecting mimicry attacks. The measure we propose below is not the only way of exploiting this behaviour, but in our experiments it turned out to be the best performing. The idea consists of reusing the OCNB-based anomaly score and applying it to each portion, attack and padding, separately. The overall score is then computed as a weighted combination of both scores, with a major reward put on the attack portion:

$$\text{score}(B) = - \sum_{c_i \in P} n_i(P) \log P(c_i | \text{self}) - \beta \left( \sum_{c_i \in A} n_i(A) \log P(c_i | \text{self}) \right) \quad (8)$$

with  $\beta \geq 1$ .

TABLE II  
DETECTION RATES (FP RATE 5%) USING THE ORIGINAL OCNB (NORMAL FACE) AND THE PPI-BASED OCNB (BOLD FACE).

User	1v49	$ A  = 10$	$ A  = 20$	$ A  = 30$	$ A  = 40$	$ A  = 50$	$\beta$
0	0.805 / <b>0.653</b>	0.000 / <b>0.110</b>	0.070 / <b>0.368</b>	0.237 / <b>0.554</b>	0.359 / <b>0.643</b>	0.473 / <b>0.656</b>	4.0
1	0.964 / <b>0.945</b>	0.392 / <b>0.970</b>	0.821 / <b>0.968</b>	0.937 / <b>0.974</b>	0.970 / <b>0.984</b>	0.960 / <b>0.986</b>	4.0
2	0.968 / <b>0.958</b>	0.000 / <b>0.180</b>	0.080 / <b>0.676</b>	0.311 / <b>0.914</b>	0.573 / <b>0.946</b>	0.736 / <b>0.969</b>	3.0
3	0.926 / <b>0.851</b>	0.039 / <b>0.401</b>	0.348 / <b>0.677</b>	0.576 / <b>0.790</b>	0.653 / <b>0.849</b>	0.734 / <b>0.855</b>	4.0
4	0.806 / <b>0.805</b>	0.089 / <b>0.149</b>	0.426 / <b>0.467</b>	0.599 / <b>0.619</b>	0.687 / <b>0.674</b>	0.696 / <b>0.728</b>	2.0
5	0.984 / <b>0.961</b>	0.018 / <b>0.150</b>	0.599 / <b>0.650</b>	0.872 / <b>0.897</b>	0.945 / <b>0.984</b>	0.972 / <b>0.974</b>	4.0
6	0.819 / <b>0.706</b>	0.028 / <b>0.267</b>	0.292 / <b>0.526</b>	0.434 / <b>0.648</b>	0.550 / <b>0.692</b>	0.620 / <b>0.720</b>	3.0
7	0.908 / <b>0.908</b>	0.000 / <b>0.002</b>	0.000 / <b>0.129</b>	0.005 / <b>0.438</b>	0.159 / <b>0.605</b>	0.361 / <b>0.688</b>	5.0
8	0.767 / <b>0.668</b>	0.000 / <b>0.357</b>	0.191 / <b>0.574</b>	0.374 / <b>0.703</b>	0.460 / <b>0.709</b>	0.558 / <b>0.777</b>	4.0
9	0.143 / <b>0.162</b>	0.000 / <b>0.000</b>	0.000 / <b>0.000</b>	0.000 / <b>0.010</b>	0.000 / <b>0.011</b>	0.000 / <b>0.014</b>	4.0
10	0.780 / <b>0.647</b>	0.004 / <b>0.180</b>	0.214 / <b>0.457</b>	0.394 / <b>0.553</b>	0.508 / <b>0.613</b>	0.551 / <b>0.688</b>	3.0
11	0.524 / <b>0.505</b>	0.000 / <b>0.085</b>	0.015 / <b>0.275</b>	0.080 / <b>0.387</b>	0.205 / <b>0.460</b>	0.294 / <b>0.504</b>	4.0
12	0.059 / <b>0.053</b>	0.000 / <b>0.000</b>	0.000 / <b>0.000</b>	0.000 / <b>0.000</b>	0.000 / <b>0.000</b>	0.000 / <b>0.001</b>	5.0
13	0.888 / <b>0.780</b>	0.002 / <b>0.374</b>	0.221 / <b>0.639</b>	0.461 / <b>0.775</b>	0.584 / <b>0.844</b>	0.634 / <b>0.885</b>	4.0
14	0.716 / <b>0.625</b>	0.005 / <b>0.172</b>	0.186 / <b>0.392</b>	0.365 / <b>0.547</b>	0.465 / <b>0.600</b>	0.556 / <b>0.616</b>	3.0
15	0.236 / <b>0.253</b>	0.000 / <b>0.089</b>	0.000 / <b>0.257</b>	0.000 / <b>0.429</b>	0.000 / <b>0.548</b>	0.033 / <b>0.566</b>	6.0
16	0.924 / <b>0.875</b>	0.071 / <b>0.791</b>	0.319 / <b>0.896</b>	0.508 / <b>0.935</b>	0.668 / <b>0.933</b>	0.734 / <b>0.947</b>	3.0
17	0.935 / <b>0.913</b>	0.000 / <b>0.000</b>	0.000 / <b>0.024</b>	0.000 / <b>0.231</b>	0.064 / <b>0.427</b>	0.222 / <b>0.635</b>	6.0
18	0.851 / <b>0.795</b>	0.309 / <b>0.522</b>	0.560 / <b>0.687</b>	0.638 / <b>0.754</b>	0.712 / <b>0.781</b>	0.748 / <b>0.801</b>	3.0
19	0.031 / <b>0.041</b>	0.000 / <b>0.085</b>	0.000 / <b>0.219</b>	0.000 / <b>0.282</b>	0.000 / <b>0.391</b>	0.000 / <b>0.413</b>	8.0
20	0.904 / <b>0.886</b>	0.000 / <b>0.000</b>	0.000 / <b>0.000</b>	0.000 / <b>0.001</b>	0.096 / <b>0.132</b>	0.311 / <b>0.408</b>	2.0
21	0.788 / <b>0.739</b>	0.128 / <b>0.670</b>	0.410 / <b>0.774</b>	0.557 / <b>0.781</b>	0.556 / <b>0.807</b>	0.630 / <b>0.839</b>	3.0
22	0.942 / <b>0.901</b>	0.026 / <b>0.087</b>	0.253 / <b>0.464</b>	0.441 / <b>0.609</b>	0.627 / <b>0.742</b>	0.721 / <b>0.823</b>	2.0
23	0.875 / <b>0.832</b>	0.008 / <b>0.375</b>	0.270 / <b>0.682</b>	0.501 / <b>0.774</b>	0.619 / <b>0.812</b>	0.659 / <b>0.818</b>	3.0
24	0.861 / <b>0.816</b>	0.000 / <b>0.092</b>	0.078 / <b>0.431</b>	0.297 / <b>0.606</b>	0.503 / <b>0.657</b>	0.584 / <b>0.754</b>	3.0
25	0.860 / <b>0.812</b>	0.000 / <b>0.015</b>	0.003 / <b>0.200</b>	0.085 / <b>0.488</b>	0.418 / <b>0.637</b>	0.574 / <b>0.758</b>	3.0
26	0.016 / <b>0.004</b>	0.000 / <b>0.002</b>	0.000 / <b>0.012</b>	0.000 / <b>0.046</b>	0.000 / <b>0.116</b>	0.000 / <b>0.118</b>	8.0
27	0.812 / <b>0.716</b>	0.000 / <b>0.262</b>	0.155 / <b>0.529</b>	0.377 / <b>0.646</b>	0.507 / <b>0.676</b>	0.648 / <b>0.741</b>	4.0
28	0.251 / <b>0.209</b>	0.000 / <b>0.000</b>	0.000 / <b>0.001</b>	0.000 / <b>0.001</b>	0.000 / <b>0.014</b>	0.000 / <b>0.056</b>	3.0
29	1.000 / <b>1.000</b>	1.000 / <b>1.000</b>	1.000 / <b>1.000</b>	1.000 / <b>1.000</b>	1.000 / <b>1.000</b>	1.000 / <b>1.000</b>	1.0
30	0.837 / <b>0.787</b>	0.055 / <b>0.097</b>	0.390 / <b>0.394</b>	0.538 / <b>0.612</b>	0.679 / <b>0.712</b>	0.735 / <b>0.808</b>	2.0
31	0.993 / <b>0.985</b>	0.844 / <b>0.996</b>	0.976 / <b>0.997</b>	0.988 / <b>0.999</b>	0.987 / <b>1.000</b>	0.989 / <b>1.000</b>	4.0
32	0.764 / <b>0.725</b>	0.000 / <b>0.002</b>	0.000 / <b>0.036</b>	0.000 / <b>0.266</b>	0.014 / <b>0.544</b>	0.066 / <b>0.694</b>	4.0
33	0.821 / <b>0.764</b>	0.101 / <b>0.642</b>	0.400 / <b>0.786</b>	0.585 / <b>0.817</b>	0.652 / <b>0.835</b>	0.694 / <b>0.843</b>	4.0
34	0.971 / <b>0.931</b>	0.007 / <b>0.543</b>	0.643 / <b>0.892</b>	0.864 / <b>0.954</b>	0.903 / <b>0.960</b>	0.935 / <b>0.976</b>	4.0
35	0.772 / <b>0.761</b>	0.000 / <b>0.000</b>	0.000 / <b>0.001</b>	0.000 / <b>0.001</b>	0.000 / <b>0.035</b>	0.020 / <b>0.136</b>	2.0
36	0.773 / <b>0.785</b>	0.000 / <b>0.053</b>	0.127 / <b>0.380</b>	0.372 / <b>0.539</b>	0.460 / <b>0.638</b>	0.555 / <b>0.681</b>	2.0
37	0.070 / <b>0.086</b>	0.000 / <b>0.097</b>	0.000 / <b>0.229</b>	0.000 / <b>0.370</b>	0.000 / <b>0.422</b>	0.000 / <b>0.427</b>	9.0
38	0.033 / <b>0.043</b>	0.000 / <b>0.155</b>	0.000 / <b>0.269</b>	0.000 / <b>0.339</b>	0.000 / <b>0.378</b>	0.000 / <b>0.395</b>	9.0
39	0.471 / <b>0.493</b>	0.000 / <b>0.089</b>	0.000 / <b>0.316</b>	0.000 / <b>0.489</b>	0.040 / <b>0.577</b>	0.078 / <b>0.633</b>	5.0
40	0.510 / <b>0.566</b>	0.000 / <b>0.474</b>	0.000 / <b>0.669</b>	0.002 / <b>0.786</b>	0.051 / <b>0.809</b>	0.174 / <b>0.809</b>	5.0
41	0.815 / <b>0.796</b>	0.000 / <b>0.000</b>	0.000 / <b>0.070</b>	0.054 / <b>0.243</b>	0.220 / <b>0.391</b>	0.335 / <b>0.544</b>	2.0
42	0.460 / <b>0.426</b>	0.000 / <b>0.191</b>	0.000 / <b>0.438</b>	0.009 / <b>0.633</b>	0.066 / <b>0.725</b>	0.114 / <b>0.734</b>	5.0
43	0.791 / <b>0.718</b>	0.000 / <b>0.095</b>	0.059 / <b>0.370</b>	0.218 / <b>0.593</b>	0.371 / <b>0.676</b>	0.510 / <b>0.704</b>	3.0
44	0.649 / <b>0.602</b>	0.000 / <b>0.001</b>	0.003 / <b>0.042</b>	0.102 / <b>0.210</b>	0.289 / <b>0.352</b>	0.400 / <b>0.460</b>	2.0
45	0.994 / <b>0.992</b>	0.908 / <b>0.926</b>	0.981 / <b>0.981</b>	0.989 / <b>0.988</b>	0.995 / <b>0.995</b>	0.996 / <b>0.996</b>	2.0
46	0.991 / <b>0.986</b>	0.000 / <b>0.031</b>	0.000 / <b>0.535</b>	0.290 / <b>0.928</b>	0.786 / <b>0.982</b>	0.918 / <b>0.987</b>	4.0
47	0.733 / <b>0.704</b>	0.005 / <b>0.073</b>	0.143 / <b>0.320</b>	0.311 / <b>0.477</b>	0.437 / <b>0.549</b>	0.523 / <b>0.573</b>	3.0
48	0.598 / <b>0.576</b>	0.000 / <b>0.102</b>	0.000 / <b>0.329</b>	0.045 / <b>0.476</b>	0.157 / <b>0.572</b>	0.261 / <b>0.603</b>	4.0
49	0.651 / <b>0.599</b>	0.000 / <b>0.661</b>	0.072 / <b>0.780</b>	0.284 / <b>0.802</b>	0.353 / <b>0.813</b>	0.386 / <b>0.815</b>	4.0
Avg	0.701 / <b>0.667</b>	0.081 / <b>0.253</b>	0.206 / <b>0.423</b>	0.314 / <b>0.558</b>	0.407 / <b>0.625</b>	0.474 / <b>0.671</b>	-

### A. Experimental Results

Table II summarises the behaviour of the OCNB detector based on the use of expression (8). As before, each threshold has been tuned so as to limit the false positive rate to 5%. The first column (1v49) shows the detection rate computed as per the 1v49 experiment (i.e., blocks belonging to other users are considered as masquerading attempts, but no mimicry attack is included). Note that using the PPI algorithm generally has

some impact on the detection rate of non-mimicry attacks. The reasons for this behaviour are related to the false positives generated by the identification algorithm, particularly in the case of users with similar profiles, as expression (8) tends to reduce the anomaly score of blocks coming from users with similar profiles. The overall effect, however, is very limited, and the global detection rate only degrades by less than 4% on average. The remaining columns in Table II show the fraction

of detected mimicry attacks of lengths between 10 and 50. In all cases, the inclusion of the PPI algorithm increases the rate by more than 20%. For some users the improvement is enormous; see, for example, users 8, 16, 33, 34, or 49. In other cases (e.g., users 20, 26, 35) the algorithm is of little help. We have not investigated yet the reasons for this behaviour.

## VI. CONCLUSIONS

The majority of current approaches to identify masquerade attempts ultimately rely upon an anomaly detection algorithm and, consequently, are susceptible of being evaded by a resourceful adversary. In this paper we have investigated the extent of such vulnerability in the case of a widely used algorithm (OCNB), and shown that most mimicry attacks can effectively evade detection. We have proposed a very efficient algorithm that attempts to separate the attack sequence from the padding based on the KL divergence. The resulting outcome could be used in a number of ways to improve upon current detection methods. Our experimental results show that it is possible to significantly increase the successful detection of mimicry attacks with almost no degradation in terms of false positives. Moreover, the principle behind the PPI algorithm is general and can be adapted to detectors other than OCNB.

## REFERENCES

- [1] M. Bertacchini and P.I. Fierens. "Preliminary Results on Masquerader Detection using Compression-based Similarity Metrics". *Electronic Journal of SADIO* 7(1), 2007.
- [2] B.M. Bowen, M. Ben Salem, S. Hershkop, A.D. Keromytis, and S.J. Stolfo. "Designing Host and Network Sensors to Mitigate the Insider Threat". *IEEE Security & Privacy*, pp. 22–29, Nov/Dec 2009.
- [3] D.D. Caputo, G.D. Stephens, and M.A. Maloof. "Detecting Insider Theft of Trade Secrets". *IEEE Security & Privacy*, pp. 14–21, Nov/Dec 2009.
- [4] E. Chan-Tin, D. Feldman, N. Hopper, and Y. Kim. "The Frog-Boiling Attack: Limitations of Anomaly Detection for Secure Network Coordinate Systems". In *SecureComm 2009*.
- [5] L. Chen and G. Dong. "Masquerader Detection using OCLEP: One-class Classification using Legth Statistics of Emerging Patterns". In *WAIMW 2006*, pp. 5.
- [6] N. Delvi, P. Domingos, Mausam, S. Sanghai, and D. Verma. "Adversarial Classification". In *ACM KDD 2004*, pp 98–108.
- [7] F.A. Durán, S.H. Conrad, G.N. Conrad, D.P. Duggan, and E.B. Held. "Building a System for Insider Security". *IEEE Security & Privacy*, pp. 30–38, Nov/Dec 2009.
- [8] J.M. Estevez-Tapiador, P. Garcia-Teodoro, and J.E. Diaz-Verdejo. "Stochastic Protocol Modeling for Anomaly-Based Network Intrusion Detection". In *IWIA 2003*, pp. 3-12.
- [9] J.M. Estevez-Tapiador, P. Garcia-Teodoro, and J.E. Diaz-Verdejo. "Detection of Web-based Attacks through Markovian Protocol Parsing". In *ISCC 2005*, pp. 457–462.
- [10] S. Evans, E. Eiland, S. Markham, J. Impson, and A. Lacro. "MDL-compress for Intrusion Detection: Signature Inference and Masquerade Attack". In *MILCOM 2007*, pp. 1–7.
- [11] P. Fogla, M. Sharif, R. Perdisci, O. Kolesnikov, and W. Lee. "Polymorphic Blending Attacks". In *15th USENIX Security Symposium*, 2006.
- [12] P. Fogla and W. Lee. "Evading network anomaly detection systems: formal reasoning and practical techniques". In *CCS 2006*, pp. 59–68.
- [13] S. Forrest, S.A. Hofmeyr, A. Somayaji, and T.A. Longstaff. "A Sense of Self for Unix Processes". In *IEEE Symp. Security and Privacy*, 1996.
- [14] S. Forrest, A.S. Perelson, L. Allen, and R. Cherkuri. "Self-Nonslef Discrimination in a Computer". In *IEEE Symp. Security and Privacy*, 1994.
- [15] D. Gao, M.K. Reiter, and D. Song. "On Gray-Box Program Tracking for Anomaly Detection". In *USENIX Security Symposium*, 2004.
- [16] M. GebSKI and R.K. Wong. "Intrusion Detection via Analysis and Modelling of User Commands". In *DAWAK*, LNCS Vol. 3589, pp. 388–397. Springer-Verlag, 2005.
- [17] J.T. Giffin, S. Jha, and B.P. Miller. "Automated Discovery of Mimicry Attacks". In *RAID 2006*.
- [18] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd Edition. Springer-Verlag, 2009.
- [19] S. Hofmeyr, S. Forrest, and A. Somayaji. "Intrusion Detection Using Sequences of System Calls". *J. Computer Security* 6:151–180, 1998.
- [20] M.C. Jason Program Office. "Horizontal Integration: Broader Access Models for Realizing Information Dominance". *Technical Report JSR-04-132*, The MITRE Corporation, JASON Program Office, Mclean, Virginia, Dec 2004. <http://www.fas.org/irp/agency/dod/jason/classpol.pdf>.
- [21] H.G. Kayacik, A.N. Zincir-Heywood, and M.I. Heywood. "Automatically Evading IDS Using GP Authored Attacks". In *IEEE Conf. on Computational Intelligence for Security and Defense Applications*, 2007.
- [22] M. Kearns and M. Li. "Learning in the Presence of Malicious Errors". In *Proc. ACM Symposium on Theory of Computing*, pp 267–280, 1988.
- [23] K.S. Killourhy and R.A. Maxion. "Toward Realistic and Artifact-Free Insider-Threat Data". In *ACSAC 2007*, pp. 87–96.
- [24] C. Kruegel, E. Kirda, D. Mutz, W. Robertson, and G. Vigna. "Automating Mimicry Attacks using Static Binary Analysis". In *USENIX Security Symposium*, 2005.
- [25] C. Kruegel, T. Toth, and E. Kirda. "Service Specific Anomaly Detection for Network Intrusion Detection". In *SAC 2002*, pp. 201–208.
- [26] M. Latendresse. "Masquerade Detection via Customized Grammars". In *DIMVA 2005*, LNCS Vol. 3548, pp. 141–159. Springer-Verlag, 2005.
- [27] D. Lowd and C. Meek. "Adversarial Learning". In *ACM KDD 2005*.
- [28] M. Mahoney. "Network Traffic Anomaly Detection Based on Packet Bytes". In *Proc. ACM SAC*, 2003.
- [29] M. Mahoney and P.K. Chan. "Learning Nonstationary Models of Normal Network Traffic for Detecting Novel Attacks". In *Proc. SIGKDD*, 2002.
- [30] R.A. Maxion and T.N. Townsend. "Masquerade Detection using Truncated Command Lines". In *DSN 2002*, pp. 219–228.
- [31] R.A. Maxion. "Masquerade Detection using Enriched Command Lines". In *DSN 2003*, pp. 5–14.
- [32] M. Oka, Y. Oyama, H. Abe and K. Kato. "Anomaly Detection Using Layered Networks Based on Eigen Co-occurrence Matrix". In *RAID 2004*, LNCS Vol. 3224, pp. 223–237. Springer-Verlag, 2004.
- [33] S.L. Pfleeger and S.J. Stolfo. "Addressing the Insider Threat". *IEEE Security & Privacy*, pp. 10–13, Nov/Dec 2009.
- [34] R. Posadas, J.C. Mex-Perera, R. Monroy, J.A. Nolzaco-Flores. "Hybrid Method for Detecting Masqueraders using Session Folding and Hidden Markov Models". In *Proc. 5th Mexican Intl. Conf. on Artificial Intelligence*, pp. 622–631, 2006.
- [35] M. Ben Salem, S. Hershkop, S. Stolfo. "A Survey of Insider Attack Detection Research". In *Insider Attack and Cyber Security: Beyond the Hacker*, Springer, 2008.
- [36] M. Ben Salem and S. Stolfo. "Masquerade Attack Detection Using a Search-Behavior Modeling Approach". Columbia University, Computer Science Department, Technical Report CUCS-027-09, 2009.
- [37] M. Schonlau, W. DuMouchel, W.-H. Ju, A.F. Karr, M. Theus, and Y. Vardi. "Computer Intrusion: Detecting Masquerades". *Statistical Science* 16(1):58–74, Feb 2001.
- [38] R. Sommer and V. Paxson. "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection". In *IEEE Symposium on Security and Privacy*, 2010.
- [39] K.M.C. Tan, K.S. Killourhy and R.A. Maxion. "Undermining an Anomaly-Based Intrusion Detection Systems Using Common Exploits". In *RAID 2002*.
- [40] K Tan, J. McHugh, and K.S. Killourhy. "Hiding Intrusions: From the Abnormal to the Normal and Beyond". In *Proc. 5th Information Hiding Workshop*, 2002.
- [41] D. Wagner and P. Soto. "Mimicry Attacks on Host-Based Intrusion Detection Systems". In *ACM CCS 2002*.
- [42] K. Wang and S. Stolfo. "One-class Training for Masquerade Detection". In *ICDM Workshop on Data Mining for Computer Security*, 2003.
- [43] K. Wang and S. Stolfo. "Anomalous Payload-based Network Intrusion Detection". In *RAID 2004*.
- [44] K. Wang and S. Stolfo. "Anomalous Payload-based Worm Detection and Signature Generation". In *RAID 2005*.
- [45] C. Warrender, S. Forrest, and B. Pearlmutter. "Detecting Intrusions Using System Calls: Alternative Data Models". In *IEEE Symposium on Security and Privacy*, 1999.