A Novel Approach to Detection of Intrusions in Computer Networks via Adaptive Sequential and Batch-Sequential Change-Point Detection Methods

Alexander G. Tartakovsky, *Senior Member, IEEE*, Boris L. Rozovskii, Rudolf B. Blažek, and Hongjoong Kim

***Abstract*—Large-scale computer network attacks in their final stages can readily be identified by observing very abrupt changes in the network traffic. In the early stage of an attack, however, these changes are hard to detect and difficult to distinguish from usual traffic fluctuations. Rapid response, a minimal false-alarm rate, and the capability to detect a wide spectrum of attacks are the crucial features of intrusion detection systems. In this paper, we develop efficient adaptive sequential and batch-sequential methods for an early detection of attacks that lead to changes in network traffic, such as denial-of-service attacks, worm-based attacks, port-scanning, and man-in-the-middle attacks. These methods employ a statistical analysis of data from multiple layers of the network protocol to detect very subtle traffic changes. The algorithms are based on change-point detection theory and utilize a thresholding of test statistics to achieve a fixed rate of false alarms while allowing us to detect changes in statistical models as soon as possible. There are three attractive features of the proposed approach. First, the developed algorithms are self-learning, which enables them to adapt to various network loads and usage patterns. Secondly, they allow for the detection of attacks with a small average delay for a given false-alarm rate. Thirdly, they are computationally simple and thus can be implemented online. Theoretical frameworks for detection procedures are presented. We also give the results of the experimental study with the use of a network simulator testbed as well as real-life testing for TCP SYN flooding attacks.**

***Index Terms*—Attack detection, change point detection, denial of service, intrusion detection, man-in-the-middle, network security, network traffic, nonparametric detection, port scanning, sequential tests, service survivability, worm.**

I. Introduction

THE goal of this paper is to show that recent advances in the change-point detection theory [1], [5], [8], [12], [15], [28]–[34] can be successfully applied in the design of anomaly detection systems for the early detection of intrusions in computer networks. We show that the asymptotic theory that has been developed for change-point detection is useful in intrusion detection problems; it also allows for the development of efficient algorithms that are easily implemented and, at the same time, have certain optimality properties. While change-point detection methods have been extensively used in many branches of signal processing (such as statistical process control, target detection, and tracking), the application of these powerful methods to network security is still in its infancy.

Large-scale attacks on computer networks usually cause abrupt changes (anomalies) in the network traffic. Typical examples include denial-of-service (DOS) attacks, worm-based attacks, port-scanning, and address resolution protocol (ARP) man-in-the-middle (MIM) attacks. In this paper, we develop efficient adaptive nonparametric sequential and batch-sequential methods for an early detection of such attacks.

Existing intrusion detection systems (IDSs) can be classified as either signature detection systems or anomaly detection systems (see, e.g., [14]). Signature detection systems detect attacks by comparing the observed patterns of the network traffic with known attack templates (signatures). If the true attack belongs to the class of attacks listed in the database, then it can be successfully detected and, moreover, identified. Examples of signature-based IDSs are Snort [26] and Bro [23]. Anomaly detection systems compare the parameters of the observed traffic with "normal" network traffic. The attack is declared once a deviation from a normal traffic is observed. Examples of ad hoc anomaly IDSs are MULTOPS [11] and D-WARD [19].

The approach we undertake here belongs to the class of anomaly-based intrusion detection systems (ABIDSs), which compare the parameters of the observed traffic with normal network traffic. The idea of the approach is based on the observation that DOS, worm, port-scanning, and MIM attacks typically lead to relatively abrupt changes in statistical models of traffic compared to the traffic's "normal mode." These changes occur at unknown points in time and should be detected "as soon as possible." Therefore, the problem of detecting an attack can be formulated and solved as a change-point detection problem: to detect a change in the distribution (model) with a fixed delay (batch approach) or minimal average delay (sequential approach), while controlling the rate of false detections. (See [1], [15], and [28]–[34] for relevant results of change-point detection theory.) In addition, we combine both methods (batch and sequential) in one unit to develop a multistage (batch-sequential) detection algorithm.

Manuscript received January 6, 2005; revised Ocotber 25, 2005. This work was supported in part by the U.S. Defense Advanced Research Projects Agency under Grant N66001-00-C-8044 and the U.S. Office of Naval Research under Grant N00014-03-1-0027 at the University of Southern California, and in part by the U.S. Army under SBIR grant DAAD17-03-C0054 at Adsantec. This work was presented in part at the Second Annual IEEE SMC Information Assurance Workshop, West Point, NY, June 5–6, 2001, and at the 35th Symposium on the Interface (Interface 2003: Security and Infrastructure Protection), Salt Lake City, UT, March 12–15, 2003. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Leslie Collins.

A. G. Tartakovsky and B. L. Rozovskii are with the Department of Mathematics and the Center for Applied Mathematical Sciences, University of Southern California, Los Angeles, CA 90089-2532 USA (e-mail: rozovski@math.usc.edu; tartakov@math.usc.edu).

R. B. Blažek is with Advanced Science and Novel Technology, Rancho Palos Verdes, CA 92075 USA and the Center for Applied Mathematical Sciences, University of Southern California, Los Angeles, CA 90089-2532 USA (e-mail: blazek@math.usc.edu).

H. Kim is with the Department of Mathematics, Korea University, Seoul 136-701, Korea (e-mail: hongjoong@korea.ac.kr).

Digital Object Identifier 10.1109/TSP.2006.879308

The developed detection algorithms have several attractive features. First, they have manageable computational complexity and thus can be implemented online. Secondly, both algorithms are self-learning, which enables them to adapt to various network loads and usage patterns. Thirdly, the sequential algorithm has optimal properties among the totality of algorithms with a prespecified false-alarm rate (FAR). It is the quickest detection algorithm in the sense that it minimizes the average delay to detection of an attack for a given FAR. In addition, when augmented with the feedback about false alarms from an appropriate decision-making authority monitoring the quality of the provided service, these adaptive methods can be used to predict traffic overflows and resource hogging. This will allow the system to manage the available resources dynamically to ensure the survivability of the service.

In contrast to [9] and [35] and many other works where parametric models (including hidden Markov models) have been considered, we use a nonparametric approach that has much more robust properties.

This paper is organized as follows. In Section II, we give an overview of the contemporary quickest detection methods and relevant results in change-point detection theory. Section III discusses issues related to the measurable characteristics of the network traffic and importance of rapid detection of DOS attacks. In Section IV, we propose two nonparametric detection algorithms, fully sequential and batch-sequential, and discuss their asymptotic performance for low FAR. In Section V, we evaluate the performance of the sequential detection algorithm for detecting DOS attacks based on experiments in a network simulator testbed. The results of Monte Carlo simulations for realistic scenarios show that the performance of the sequential detection algorithm is very high. The detector allows for rapid detection of typical DOS attacks such as user datagram protocol (UDP) packet storm, Internet Control Message Protocol (ICMP) ping of death, and transmission control protocol (TCP) SYN DOS. The results are further validated in Section VI, where we test the proposed sequential nonparametric detection procedure in real environment in detecting TCP SYN flooding attacks based on data sets collected by the Massachusetts Institute of Technology (MIT) Lincoln Laboratory. The results of comparison with two other efficient detection schemes are also presented in this section.

II. Overview of Sequential and Batch Detection Methods

Network intrusions occur at unknown points in time and lead to changes in the statistical properties of certain observables. Therefore, the problem of detecting attacks can be formulated and solved as a change-point detection problem: detect changes in the distributions (models) with fixed delays (batch approach) or minimal average delays (sequential approach), while maintaining the FAR at a given level. Choosing the relevant network flow and resource usage characteristic that are to be observed represents a crucial aspect of the development of the local ABIDS. The observables used for our experiments are discussed in Section III.

There are two main approaches to detecting abrupt changes in stochastic models: the *fixed-size batch detection* (or a posteriori methods) and *sequential change-point detection* [1], [5], [28]. In the latter setting, the problem is formulated as a quickest detection problem: detect a change in the model as rapidly as possible after its occurrence, while maintaining the FAR at a given level. In what follows, we will focus on sequential and batch-sequential methods for detecting attacks. The latter ones perform fixed size processing in every stage of a multistage sequential decision-making process.

The design of the quickest (sequential) change-point detection procedures involves optimizing the tradeoff between two kinds of performance measures, one being a measure of detection delay and the other being a measure of the frequency of false alarms. There are two standard mathematical formulations for the optimum tradeoff problem. The first of these is a minimax formulation proposed by Lorden [16] and Pollak [24], in which the goal is to minimize the worst case delay subject to a lower bound on the mean time between false alarms. The second is a Bayesian formulation, proposed by Shiryaev [27], in which the change point is assumed to have a prior distribution and the goal is to minimize the expected delay subject to an upper bound on false-alarm probability. Therefore, the sequential detection methods (both Bayesian and non-Bayesian) involve two performance indexes: the *rate of false alarms* and the *detection delay*. In what follows, we will not consider the Bayesian change detection problem, as the prior distribution of the change point is usually unknown in applications of interest.

Let X_n, $n \geqslant 1$, be a sequence of observations that are being chosen for monitoring. The observed random variables $X_1, X_2, \ldots$ have a joint probability density function (pdf) $p_0(X_1, \ldots, X_n)$ (baseline distribution) until a change occurs at an unknown point in time λ, $\lambda \in \{1, 2, \ldots\}$. After the change occurs, the observations have another distribution $p_1(X_1, \ldots, X_n)$. In other words, it is assumed that $X_1, X_2, \ldots$ have the conditional pdf $p_0(X_n|X_1, \ldots, X_{n-1})$ for $n < \lambda$ and the conditional pdf $p_1(X_n|X_1, \ldots, X_{n-1})$ for $n \geqslant \lambda$, where p_0 and p_1 are prechange and postchange pdfs, respectively. Therefore, if the change occurs at time $\lambda = k$, then the conditional density of the kth observation changes from $p_0(X_k|X_1, \ldots, X_{k-1})$ to $p_1(X_k|X_1, \ldots, X_{k-1})$.

A sequential change-point detection procedure is identified with a stopping time τ for an observed sequence $\{X_n\}_{n \geqslant 1}$, i.e., the time of alarm τ, at which it is declared that a change has occurred, is a random variable depending on the observations. A good detection procedure should have a low FAR and small values of the expected detection delay, provided that there is no false alarm. To be more specific, let $\mathbf{P}_k$ and $\mathbf{E}_k$ denote the probability and the expectation that correspond to the sequence $\{X_n, n \geqslant 1\}$ when the change occurs at time $\lambda = k$. For the situation when there is no change (i.e., $\lambda = \infty$), we will use the notation $\mathbf{P}_\infty = \mathbf{P}_0$ and $\mathbf{E}_\infty = \mathbf{E}_0$, where $\mathbf{P}_0$ is a prechange distribution.

When we compute the performance of the detection procedures, we will be interested in the average detection delay (ADD) and the FAR, which are defined by

$$\mathrm{ADD}_\lambda(\tau) = \mathbf{E}_\lambda(\tau - \lambda | \tau \geqslant \lambda), \quad \mathrm{FAR}(\tau) = \frac{1}{\mathbf{E}_0 \tau}.$$

There are two major competitive sequential change-point detection algorithms: Page's cumulative sum (CUSUM) detection

procedure and the Shiryaev–Roberts–Pollak detection procedure [1], [24], [28]. Both approaches utilize the log-likelihood ratio (LLR) for the hypotheses that the change occurred at the point λ and that there is no change at all ($\lambda = \infty$), which is defined as

$$Z_{n,\lambda} = \sum_{k=\lambda}^{n} \log \frac{p_1(X_k|X_1,\ldots,X_{k-1})}{p_0(X_k|X_1,\ldots,X_{k-1})}, \quad n \geqslant \lambda.$$

Page's CUSUM procedure is motivated by a maximum likelihood argument. It is based on the comparison of the maximum LLR statistic $U_n = \max_{1\leqslant\lambda\leqslant n} Z_{n,\lambda}$ with a threshold h

$$\tau_{\mathrm{CU}}(h) = \min\{n \geqslant 1 : U_n \geqslant h\}. \tag{1}$$

Note that under the condition $h > 0$ and in the case of independent identically distributed (i.i.d.) observations when $p_0(X_n|X_1,\ldots,X_{n-1}) = p_0(X_n)$ and $p_1(X_n|X_1,\ldots,X_{n-1}) = p_1(X_n)$, the statistic U_n in (1) can be replaced by the statistic $\tilde{U}_n$ which obeys the recursion

$$\tilde{U}_n = \max\left\{0, \tilde{U}_{n-1} + \log\frac{p_1(X_n)}{p_0(X_n)}\right\} \tag{2}$$

with the initial condition $\tilde{U}_0 = 0$. This latter representation is a basis for the nonparametric detection algorithm proposed in Section IV-A.

The Shiryaev–Roberts–Pollak procedure is motivated by Bayesian, rather than maximum likelihood, considerations. Specifically, define the statistic $R_n = \sum_{\lambda=1}^{n} \exp\{Z_{n,\lambda}\}$, which can be regarded as an average likelihood ratio. The corresponding detection algorithm is identified with the stopping time

$$\tau_{\mathrm{SP}}(h) = \min\{n \geqslant 1 : \log R_n \geqslant h\}. \tag{3}$$

It is known [1], [16], [24], [28] that in the i.i.d. case both detection methods minimize the worst case average detection delay $\sup_\lambda \mathrm{ADD}_\lambda(\tau)$ among the detection algorithms for which the FAR is fixed at a given level $\overline{\mathrm{FAR}}$, i.e., $\mathrm{FAR}(\tau) \leqslant \overline{\mathrm{FAR}}$. The threshold values should be chosen from the conditions $\mathbf{E}_0\tau_{\mathrm{CU}}(h) = 1/\overline{\mathrm{FAR}}$ and $\mathbf{E}_0\tau_{\mathrm{SP}}(h) = 1/\overline{\mathrm{FAR}}$. In the preliminary engineering computations, one can set $h = \log(1/\overline{\mathrm{FAR}})$, which guarantees the inequalities $\mathrm{FAR}(\tau_{\mathrm{CU}}) \leqslant \overline{\mathrm{FAR}}$ and $\mathrm{FAR}(\tau_{\mathrm{SP}}) \leqslant \overline{\mathrm{FAR}}$.

However, the i.i.d. assumption is very restrictive for intrusion detection applications. Recent advances in general change-point detection theory (see Lai [15], Tartakovsky [31], and Tartakovsky and Veeravalli [33], [34]) allow us to conclude that the detection procedures (1) and (3) are also optimal for general statistical models when the FAR is low ($\overline{\mathrm{FAR}}$ is small). While being asymptotically optimal, the corresponding sequential procedures have manageable computational complexity. The latter features make them very attractive for intrusion detection applications where the observed data are usually correlated and nonstationary, even bursty, due to substantial temporal variability.

Another restrictive feature of the optimal detection procedures described above is that they require complete prior information regarding the prechange and postchange distributions. Parametric modifications and corresponding detection procedures that are based on the generalized likelihood ratio, the likelihood ratio mixtures, and the adaptive likelihood ratio [1], [8], [16], [24], [33] are useful for many applications but do not solve the problem when the distributions are not known. Several nonparametric procedures that have been proposed in the literature are somewhat different in nature [2], [5], [12], [18]. They usually use a sequence of statistics based on signs or ranks. For instance, in [18], a CUSUM procedure based on ranks has been proposed, and in [12], a Shiryaev–Roberts–Pollak procedure based on the sequential vectors of signs and ranks has been studied for the i.i.d. observations. Nonparametric sign-rank likelihood ratio detection methods are extremely efficient from a statistical standpoint. However, they are not quite computationally feasible for our applications.

In contrast to the sequential change-point problem in which the "homogeneity" hypothesis is tested online in the process of data acquisition, the a posteriori change-point problem is considered on the fixed-time interval $1,\ldots,n_0$. In this case, a good detection procedure is based on the comparison of the statistic $U_{n_0} = \max_{0\leqslant\lambda\leqslant n_0} Z_{\lambda,n_0}$ with a threshold h at moment n_0. The decision that a change occurred is made if $U_{n_0} \geqslant h$. The threshold h is chosen from the condition $\mathbf{P}_0(U_{n_0} \geqslant h) = \alpha$, i.e., such that the false-alarm probability is equal to a given value α. If the change occurs at the point λ, then any fixed-size (batch) method detects this change with the fixed delay $n_0 - \lambda$, which is large in all cases where n_0 is large and λ is small. The advantage of the sequential methods is obvious.

In many applications, it can be beneficial to combine both methods by grouping the data obtained in the fixed size intervals and first performing an intraprocessing of the data in these fixed-size intervals. Then, the results of this intraprocessing are further processed sequentially. The resulting procedure represents a multistage sequential procedure with batch processing within individual stages. The idea is similar to group sequential tests. The corresponding detection method will be called the *batch-sequential method*.

In Section IV-A, we propose a simple CUSUM-type nonparametric fully sequential detection algorithm and evaluate its asymptotic operating characteristic for the low FAR. The performance is evaluated under general conditions that are not confined to the restrictive i.i.d. assumption. This algorithm is extremely simple and computationally inexpensive. At the same time, it performs very well in a variety of intrusion detection scenarios, as shown in Sections V and VI. In Section IV-B, a nonparametric batch-sequential algorithm is discussed.

III. Network Traffic Flow Observables and Scope of Detection

A. Observables

While monitoring network traffic, one can observe various kinds of information related to the headers, sizes, and other characteristics of the received and transmitted packets, as well as the usage of system resources, service quality, and similar aspects associated with the utilization of the network and available resources. For example, for the purpose of detecting DOS attacks, port-scanning, and worm attacks, we observe in the transport

layer the number of TCP packets categorized by size or type (ACK, SYN, URG, etc.), the number of UDP packets and their sizes, the source and destination port for each packet, etc. In applications related to insider attacks, it is also important to observe information relevant to local-area network (LAN)-related protocols, e.g., occurrences of media access control (MAC) and Internet protocol (IP) assignment changes via ARP.

The methods described in this paper are general, and the choice of the particular set of observed characteristics depends on the application at hand. In our simulation experiments, we focus on classical DOS attacks and utilize information related to the transport layer of the network protocol and to system usage; namely, we observe the numbers of received packets categorized by size and type and monitor the size of buffers related to received and transmitted SYN packets. In particular, we consider a sequence of nonoverlapping (relatively short) time intervals; then for each packet type pt among ICMP, UDP, or TCP, we categorize the received packets by their size into $M_{\rm pt}$ size bins $A^1_{\rm pt}, A^2_{\rm pt}, \ldots, A^{M_{\rm pt}}_{\rm pt}$. In the case of the UDP and ICMP attacks, we observe the total number $N^{\rm pt}_{k,i}$ of packets of type pt with sizes in the ith bin received during the kth time interval. For the purpose of detection of the TCP SYN attack, we observe the SYN packet induced buffer size B_k at the end of the kth time interval. In the testbed, we observe the statistics $N^{\rm pt}_{k,i}$ and B_k *simultaneously*. Note that, in general, the sizes of the bins $A^i_{\rm pt}$, as well as the sizes $\Delta_k = t_k - t_{k-1}$ of the time intervals, can vary. In our simulations, we used constant size time intervals with $\Delta_k = \Delta$ and bins with unequal sizes.

B. Scope of Detection—Importance of Rapid Detection of DOS Attacks

We now argue that despite the fact that in many cases DOS attacks are obvious, there are very important scenarios when DOS attacks represent a serious threat to the service provided by the network infrastructure and there is the need for early DOS attack detection. In fact, DOS flooding attacks remain the subject of intensive research, as can be seen in [11], [13], [19], [20], [22], and [37]. We first describe DOS detection practices that are currently used in the network security community.

There are three basic ad hoc approaches that are currently used for DOS attack detection: 1) observing network performance degradation or outage; 2) monitoring link saturation or number of flows per a host and port; and 3) signature-based detection. In the first case, network centers may receive phone calls from users who are unable to access their e-mail or experience many dropped or sluggish connections. It is also common for the network centers to notice these problems themselves, e.g., by monitoring the number of dropped connections. The second approach is to monitor some reasonable network characteristic, e.g., the link saturation. One suspects that a DOS attack takes place if a customer's rented link becomes, e.g., 98% utilized. Another example is monitoring the numbers of flows at the router of a LAN and trigger an alarm if a host has too many connections on a single port. The third approach—signature-based detection—only detects selected sets of DOS attacks. Besides tremendous false-alarm rates, it suffers from frequently missed detections, especially of unknown attacks. Its nature is different than the other approaches and is not considered in this paper.

The first two approaches are intuitively attractive in that they observe very reasonable and informative network characteristics. However, their ad hoc nature results in the following most important common drawbacks.

1) It is often too late when the detection occurs (i.e., the network performance has already been degraded).
2) The FAR becomes unacceptably high when trying to increase the speed of detection.

Therefore, a rigorous framework is necessary to control the tradeoff between two requirements: a low FAR and a minimal detection delay. Our approach directly addresses these issues—it controls the average FAR at a prescribed low level while minimizing the average detection delay for the given FAR. Therefore, we do perceive our method as an important general tool for intrusion detection, including DOS attacks. In addition to detecting DOS flooding attacks (examples of which are presented in this paper), the method can be applied to the data that are commonly monitored by network experts to provide control for FAR and detection delays of the intrusion system.

Let us now focus on the scenarios where DOS attacks represent a serious threat. The first situation is related to mission-critical network services that must be protected from any kind of outages and degradation. Examples of such mission-critical services are the federal reserve, stock exchanges, and various tactical military networks. Another very important example is governmental and industrial networks that monitor and control real-time processes that are critical for the public, e.g., control of manufacturing processes in chemical industry to prevent dangerous accidents or monitoring the power grid to prevent large-scale power outages.

It is important to point out that not all DOS attacks start abruptly. Many attacks start gradually, as illustrated in Section V and described in [13], where the authors analyze the onset features of a large number of real DOS attacks. In these situations, early detection at a stage when the attack is still subtle (i.e., long before network performance degrades) is crucial and may represent the only realistic chance for effective DOS attack prevention.

The second scenario where efficient detection of DOS attacks plays a critical role is when large Internet service providers (ISPs) with huge high-speed networks provide relatively low-capacity links to many customers. This scenario includes large corporate and governmental networks that provide services to many small departments and divisions. In such a setup, a DOS attack that overwhelms the customer's link is essentially invisible to the ISP. The attack is only obvious to the customer who is defenseless. Therefore, the detection and prevention of the attack have to be performed at the level of the ISP, not at the customer's site. As a real-life example, it has been reported (see [10]) by the Gibbson Research Corporation (GRC) that it took their ISP more than two hours to react to their request to protect their link after it was attacked by a TCP SYN flood attack in January 2002. The ISP was not aware of the attack at all. The GRC engineers had to provide the ISP with detailed analysis and a detailed request for a defensive action. The prevention of the attack would have been delayed much more in the case of an ordinary, less affluent customer of the ISP.

The above example underlines the importance for ISPs to be able to detect subtle DOS attacks rapidly but with a low FAR. However, in cases with huge numbers of ISP customers, it is impossible to monitor all customer links separately. Segments of the networks or whole subnets have to be monitored in aggregate fashion where the DOS attacks are not as subtle as at the global level, but are by no means obvious. Our detection method is able to detect low-intensity attacks with reasonable detection delays and FARs and therefore promises a tool for this purpose.

IV. Detection Algorithms

Next we describe the two developed detection algorithms: the purely sequential algorithm and the batch-sequential algorithm. Both algorithms are nonparametric versions of the CUSUM-type method adapted to detect changes in multiple bins. We also present the results of asymptotic analysis and optimization of the sequential algorithm.

It is worth noting that while in the rest of this paper the observables $N^{\rm pt}_{k,i}$ are associated with DOS attacks, all or almost all of the results are valid in a more general context where the observations $X_{k,i}$, $k \geqslant 1$, $i = 1, \ldots, M$ are identified with some random variables monitored by a multichannel sensor system.

A. Sequential Detection Algorithm and Its Asymptotic Performance

1) Algorithm Description: If both the prechange and the postchange distributions are exactly known, then the optimal procedure represents a thresholding of either the CUSUM statistic or the Shiryaev–Roberts statistic [see (1)–(3)]. In our applications, however, it is very difficult, if not impossible, to build an exact model. As a result, these distributions are usually unknown. For this reason, we will use a nonparametric approach.

Let p_i and p_0 denote probabilities when the change occurs in the ith bin and when there is no change. When these probabilities are unknown, the LLRs

$$\log \frac{p_i\left(N^{\rm pt}_{k,i}|N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k-1,i}\right)}{p_0\left(N^{\rm pt}_{k,i}|N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k-1,i}\right)}, \quad i=1,\ldots,M_{\rm pt}$$

are also not known. Therefore, the LLRs should be replaced by appropriate score functions $g^{\rm pt}_i(N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k,i})$ that have a negative mean $\mathbf{E}_0 g^{\rm pt}_i(N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k,i}) < 0$ before the attack and a positive mean $\mathbf{E}_{\lambda,i} g^{\rm pt}_i(N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k,i}) > 0$, $\lambda < k$ after the attack starts. Here $\mathbf{E}_0$ and $\mathbf{E}_{\lambda,i}$ stand for the expectations when there is no attack and when the attack starts at the time λ ($\lambda = j\Delta$, $j = 1,2,\ldots$) in the ith bin. In this case, the CUSUM-type statistic

$$U^{\rm pt}_{n,i} = \max_{1\leqslant\lambda\leqslant n} \sum_{k=\lambda}^{n} g^{\rm pt}_i\left(N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k,i}\right)$$

remains close to zero or slightly negative in normal conditions, while under the attack it starts drifting upward until it crosses a threshold $h^{\rm pt}$. More specifically, the stopping time (the time of alarm) is defined as

$$\tau^{\rm pt} = \min\left\{n \geqslant 1 : \max_{1\leqslant i\leqslant M_{\rm pt}} U^{\rm pt}_{n,i} \geqslant h^{\rm pt}\right\}. \qquad (4)$$

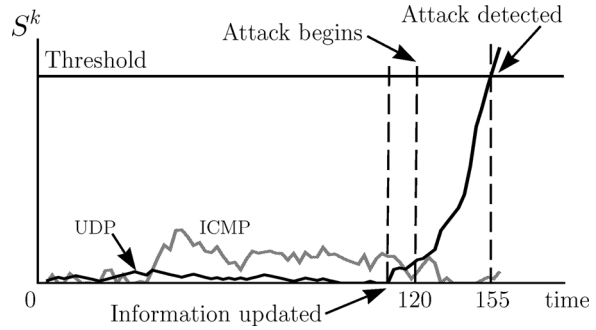

Fig. 1. An illustration of the behavior of the decision statistic for one particular run of a simulated UDP DOS attack and for ICMP normal traffic.

The score functions can be chosen in many ways. One possible solution is based on the observation that in many cases the attack leads to abrupt changes of mean values; therefore, the score functions should be sensitive to changes in mean values.

Let $\mu^{\rm pt}_i = \mathbf{E}_0 N^{\rm pt}_{k,i}$ and $\theta^{\rm pt}_i = \mathbf{E}_{1,i} N^{\rm pt}_{k,i}$ denote the prechange and postchange mean values of the corresponding packet sizes in the ith bin for the packet type pt. The value of $\mu^{\rm pt}_i$ can be estimated quite accurately in advance; hence, it is supposed to be known. Once in a while, however, it is reestimated, as discussed later (see Fig. 1). The value of $\theta^{\rm pt}_i$ is unknown and should be estimated on line. We suppose that the attack leads to a change in the mean value of the number of packets $N^{\rm pt}_{k,i}$ for some packet size bin $i = 1,\ldots,M_{\rm pt}$. In other words, the problem of detecting an attack can be regarded as the quickest change detection in the mean $\mu^{\rm pt}_i \to \theta^{\rm pt}_i$, where $\mu^{\rm pt}_i$ is known and $\theta^{\rm pt}_i$ is unknown. In the rest of this section, we suppose that $\theta^{\rm pt}_i > \mu^{\rm pt}_i$. It is worth noting that we do not assume that the attack leads only to a change of the mean. Other statistical parameters can also change along with the mean values. These parameters, however, will be treated as nuisance factors.

For the UDP and ICMP attacks, the score functions are

$$g^{\rm pt}_i\left(N^{\rm pt}_{1,i},\ldots,N^{\rm pt}_{k,i}\right) = N^{\rm pt}_{k,i} - \mu^{\rm pt}_i - \varepsilon\hat{\theta}^{\rm pt}_{k,i}$$

where ε is a tuning parameter belonging to the interval (0,1) and $\hat{\theta}^{\rm pt}_{k,i}$ is an estimate of the unknown mean $\theta^{\rm pt}_i$. The particular choice of this estimate will be discussed at the end of this section.

If the threshold $h^{\rm pt}$ is positive (which is usually the case), then the detection algorithm (4) is essentially equivalent to the detection algorithm that performs simultaneous thresholding of the statistics

$$S^{\rm pt}_{k,i} = \max\left\{0, S^{\rm pt}_{k-1,i} + N^{\rm pt}_{k,i} - \mu^{\rm pt}_i - \varepsilon\hat{\theta}^{\rm pt}_{k,i}\right\} \qquad (5)$$

with reflection from the zero barrier [compare with (2)]. If any statistic $S^{\rm pt}_{k,i}$ exceeds the threshold $h^{\rm pt}$, then an alarm message is sent to the decision making engine, i.e., the detection algorithm (4) can be rewritten in the form

$$\tau_{\rm pt} = \min\left\{k \geqslant 1 : \max_{1\leqslant i\leqslant M_{\rm pt}} S^{\rm pt}_{k,i} \geqslant h^{\rm pt}\right\}. \qquad (6)$$

In the case of the observed buffer sizes B_k (to detect a TCP SYN attack), we use a similar statistic with $N^{\rm pt}_{k,i}$ replaced by B_k. Note that the detection algorithm of (5) and (6) is sensitive to changes in the average intensity of the observed traffic. This algorithm represents a "multichannel" version of the CUSUM-type nonparametric adaptive detection procedure that is adapted to detect changes in the mean values of packet sizes in multiple size bins. In what follows, this algorithm will be referred to as the *multichannel nonparametric adaptive CUSUM* detection algorithm, and the abbreviation MNA-CUSUM will be used throughout the paper.

In contrast to the likelihood ratio-based CUSUM test, generally the MNA-CUSUM algorithm (6) is not optimal. Certain optimization is possible based on the training data, if the adaptive estimation of post-change parameters is applied. This important problem will be considered elsewhere.

To fix the FAR, the threshold $h^{\rm pt}$ is chosen from the condition $\mathbf{E}_0\tau_{\rm pt} = 1/\overline{\rm FAR}$, where $\overline{\rm FAR}$ is a given positive number that characterizes the admissible FAR. An argument provided in Section IV-A2 shows that under certain regularity conditions

$$\mathbf{E}_{\lambda,i}(\tau_{\rm pt} - \lambda | \tau_{\rm pt} \geqslant \lambda) \approx h^{\rm pt}/I^{\rm pt}_i$$

where $I^{\rm pt}_i = \mathbf{E}_i g_i(N^{\rm pt}_{k,i})$ (see also Tartakovsky *et al.* [32]). An alternative way of stabilizing the FAR is to fix the probability of a false alarm in the fixed-length time interval at a given level.

As shown in Fig. 1, the information about the patterns of regular traffic flow is updated when a statistic $S^{\rm pt}_{k,i}$ reaches and departs the zero barrier. If the decision-making engine reports that a previously issued alarm message was a false alarm, then the information about regular traffic patterns and thresholds will be updated accordingly, and the traffic monitoring starts all over again. In particular, the prechange mean values $\mu^{\rm pt}_i$ are reestimated.

We now briefly discuss a selection of the estimators. Choosing the estimators $\hat\theta^{\rm pt}_{n,i}$ is not a straightforward task. For example, the estimate

$$\hat\theta^{\rm pt}_{n,i} = \max\left\{\mu^{\rm pt}_i, n^{-1}\sum_{k=1}^{n} N^{\rm pt}_{k,i}\right\} \tag{7}$$

which is natural for $\lambda = 1$, is not a good choice for $\lambda > 1$. This estimate works well only when the attack occurs from the very beginning, but for large λ, the performance degrades dramatically. A good choice would be estimators that "forget" the past depending on the behavior of the statistic $S^{\rm pt}_{k,i}$. For example, adaptive exponentially weighted estimators perform fairly well [8]. These estimators have a similar structure to that of a sample mean (7), where the current sample size k is replaced with the adaptive number β_k; β_k is set to zero (renewed) whenever the statistic $S^{\rm pt}_{k,i}$ hits the zero level. This allows us to forget the observations that are not consistent with the traffic change. Further details are omitted and will be presented elsewhere.

2) Asymptotic Operating Characteristics: In this section, we provide an asymptotic analysis of the MNA-CUSUM detection algorithm (6). Recall that we are interested in the two performance indexes—the false-alarm rate ${\rm FAR}(\tau_{\rm pt}) = 1/\mathbf{E}_0\tau_{\rm pt}$ and the average detection delay ${\rm ADD}_{\lambda,i}(\tau_{\rm pt}) = \mathbf{E}_{\lambda,i}(\tau_{\rm pt} - \lambda | \tau_{\rm pt} \geqslant \lambda)$.

The value of $\mathbf{E}_0\tau^{\rm pt}$ can be evaluated when the estimates $\hat\theta^{\rm pt}_{k,i}$ depend only on $k-1$ previous observations at the kth stage. In other words, one has to use one-stage delayed estimates, as was suggested by Robbins and Siegmund [25] for the analysis of one-sided sequential tests. Another possibility is to replace the estimates $\hat\theta^{\rm pt}_{k,i}$ by design constants $c^{\rm pt}_i$. A reasonable choice is the minimal expected values of the postchange mean for which the attack is still detectable (see below).

It can be shown that under very general conditions that include correlated and nonstationary data

$$\mathbf{E}_0\tau_{\rm pt} \geqslant \frac{1}{M_{\rm pt}} C_1 e^{C_2 h_{\rm pt}} \tag{8}$$

where C_1 and C_2 are constants (see [5] and [32]). It is difficult to compute these constants if no assumptions are made on the models. In practice, the FAR should be evaluated experimentally for different threshold values. The important fact is that the FAR decays exponentially fast with the threshold value.

We now evaluate ${\rm ADD}_{\lambda,i}$ for large values of the threshold $h_{\rm pt}$, i.e., for low FAR. Recall that $\mathbf{P}_0$ and $\mathbf{P}_{\lambda,i}$ stand for the probabilities that correspond to the true statistical models when there is no attack, and the attack occurs at time λ in the ith size bin, respectively.

Assume that sample means converge almost surely (a.s.) to their expected values, i.e., the following conditions hold:

$$\frac{1}{n}\sum_{k=1}^{n} N^{\rm pt}_{k,i} \xrightarrow[n\to\infty]{\mathbf{P}_0-{\rm a.s.}} \mu^{\rm pt}_i, \quad \frac{1}{n}\sum_{k=\lambda}^{\lambda+n-1} N^{\rm pt}_{k,i} \xrightarrow[n\to\infty]{\mathbf{P}_{\lambda,i}-{\rm a.s.}} \theta^{\rm pt}_i$$

and moreover, that for all $\varepsilon > 0$ and $\lambda \geqslant 1$

$$\sum_{n=1}^{\infty} \mathbf{P}_0\left\{\left| n^{-1}\sum_{k=1}^{n} N^{\rm pt}_{k,i} - \mu^{\rm pt}_i \right| > \varepsilon\right\} < \infty$$

$$\sum_{n=1}^{\infty} \mathbf{P}_{\lambda,i}\left\{\left| n^{-1}\sum_{k=\lambda}^{\lambda+n-1} N^{\rm pt}_{k,i} - \theta^{\rm pt}_i \right| > \varepsilon\right\} < \infty.$$

The latter two conditions determine the rate of convergence in the strong law of large numbers. If these conditions are not satisfied, i.e., if there is no convergence of sample means to the corresponding mean values of traffic with an appropriate rate, one cannot expect "nice" properties of any detection scheme.

Finally, for the sake of simplicity, in the rest of this section we assume that in the decision statistic (5), in place of the estimates $\hat\theta^{\rm pt}_{k,i}$ the constants are used, i.e., $\varepsilon\hat\theta^{\rm pt}_{k,i} = c^{\rm pt}_i$, where $c^{\rm pt}_i$ are the design nonnegative constants.

Then arguments similar to those used in the proof of [33, Theorem 1] (see also [34]) can be applied in order to show that asymptotically as $h_{\rm pt} \to \infty$,

$${\rm ADD}_{\lambda,i}(\tau_{\rm pt}) \sim \frac{h_{\rm pt}}{\theta^{\rm pt}_i - \mu^{\rm pt}_i - c^{\rm pt}_i} \tag{9}$$

whenever $\theta^{\rm pt}_i - \mu^{\rm pt}_i - c^{\rm pt}_i > 0$. The detailed proof of the asymptotic approximation (9) is quite tedious and is given in Tartakovsky *et al.* [32].

We now discuss how the constants $c_i^{\rm pt}$ have to be selected in order to guarantee the high performance of the algorithm. The first important fact is that $\theta_i^{\rm pt} - \mu_i^{\rm pt} - c_i^{\rm pt}$ should be positive. Indeed, it can be shown that if $\theta_i^{\rm pt} - \mu_i^{\rm pt} = c_i^{\rm pt}$, then the average detection delay is on the order of the square of the threshold, in which case the ADD can be large. The situation becomes even worse when $\theta_i^{\rm pt} - \mu_i^{\rm pt} < c_i^{\rm pt}$. Therefore, a minimum requirement is $c_i^{\rm pt} < \delta_i^{\rm pt}$, where $\delta_i^{\rm pt} = \theta_i^{\rm pt} - \mu_i^{\rm pt}$.

On the other hand, we cannot choose $c_i^{\rm pt}$ too small, e.g., $c_i^{\rm pt} = 0$ is not a good choice. The reason is that the constants C_1 and C_2 in (8) and, as a result, the FAR, depend on $c_i^{\rm pt}$: the smaller the $c_i^{\rm pt}$, the bigger the FAR. Hence the threshold $h_{\rm pt} = h_{\rm pt}(\{c_i^{\rm pt}\}, \overline{\rm FAR})$ depends on the FAR constraint $\overline{\rm FAR}$ and the constants $c_i^{\rm pt}$. This means that there is a tradeoff between the growth of the numerator (the threshold) and the denominator (the value of $\delta_i^{\rm pt} - c_i^{\rm pt}$) in (9); they compensate each other, giving an optimal value of the constants provided that $c_i^{\rm pt} < \delta_i^{\rm pt}$.

Choosing the optimal values of $c_i^{\rm pt}$ is complicated by the fact that $\theta_i^{\rm pt}$, $i = 1, \ldots, M_{\rm pt}$, are unknown. The following strategy can be proposed to choose the values of the design constants. Usually one can evaluate the minimal expected values of $\theta_{i,min}^{\rm pt}$ for typical attack scenarios. These minimal values are used to obtain the trivial upper bounds $c_i^{\rm pt} < \delta_{i,min}^{\rm pt}$. Next, note that the attack should be detected within a certain prespecified time interval T or it will be missed. Thus, the approximation (9) with the minimal values $\delta_{i,min}^{\rm pt}$ can be used to obtain

$$\frac{h_{\rm pt}}{\delta_{i,min}^{\rm pt} - c_i^{\rm pt}} < T \Rightarrow c_i^{\rm pt} < \delta_{i,min}^{\rm pt} - h_{\rm pt}/T.$$

The values of $\delta_{i,min}^{\rm pt}$ and $c_i^{\rm pt}$ may change dramatically at large time scales (night, day, morning, afternoon). Therefore, the procedure requires parameter estimation at least at large scales.

B. A Batch-Sequential Algorithm

There are many reasonable detection methods currently employed in the signal-processing and networking communities that utilize batch statistics based on various kinds of important network characteristics. These batch methods utilize a single batch of data that is observed during a fixed time interval T_k. Obviously, the sensitivity of the detection method can be improved by increasing the length of the time interval. However, in this case, a subtle anomaly that is very consistent in time (over several time intervals T_k) might either be undetected or be detected too late.

As a remedy, we propose a *batch-sequential* (multistage) detection method in which both batch and sequential parts are combined in one unit. The main advantage of this approach is that the batch-sequential statistics retain enough relevant past information to detect network intrusions quickly, while maintaining the FAR below a selected level.

To illustrate our batch-sequential method, we outline a modification of a special version of the popular batch χ^2 statistics to obtain a batch-sequential detection method for detecting anomalous departures from regular network traffic patterns. More precisely, the method is designed to detect increase or decrease in the expected number of packets that are observed in all possible sets of size bins.

For the sake of simplicity, we will concentrate on the statistic

$$\chi_{{\rm pt},k}^2 = \sum_{i=1}^{M_{\rm pt}} \frac{\left(N_{k,i}^{\rm pt} - \mu_i^{\rm pt}\right)^2}{\mu_i^{\rm pt}} \tag{10}$$

which in this particular form measures the departure of the network traffic from the prechange distribution $\mathbf{P}_0$ under which the expectation $\mathbf{E}_0 N_{k,i}^{\rm pt} = \mu_i^{\rm pt}$ simultaneously for all $i = 1, \ldots, M_{\rm pt}$. In this case, the mean value has the form $\mu_i^{\rm pt} = N_k^{\rm pt} p_i^{\rm pt}$, where $p_i^{\rm pt}$ is the $\mathbf{P}_0$-probability of a packet size to fall into the ith size bin and $N_k^{\rm pt} = \sum_{i=1}^{M_{\rm pt}} N_{k,i}^{\rm pt}$. It is worth emphasizing that in the batch-sequential method, the number of packets $N_{k,i}^{\rm pt}$ is computed in the fixed-size intervals of the lengths T_k; these intervals are much bigger compared to the values of Δ_k used in the sequential method. For example, $T_k = T = m\Delta$, where m is a large number. This is equivalent to the grouping of the data obtained from the sequence of intervals $k\Delta$, $k = 1, \ldots, m$.

It is well known that if the observed packet sizes are i.i.d., and if the total number $N_k^{\rm pt}$ of packets of type pt observed during the time interval T_k is sufficiently large (and independent of the observed packet sizes), then the batch statistic (10) has asymptotically the χ^2 distribution with $M_{\rm pt} - 1$ degrees of freedom. This fact is the basis of the traditional goodness-of-fit test of the hypothesis that the distribution of the packet sizes over the bins $A_{\rm pt}^1, \ldots, A_{\rm pt}^{M_{\rm pt}}$ corresponds to regular network traffic.

The modified batch-sequential algorithm consists of the following multistage procedure. In the kth stage, we group and process the data observed during the interval T_k and form a batch statistic $\chi_{{\rm pt},k}^2$ defined by (10). We then apply our sequential approach to the sequence $\chi_{{\rm pt},1}^2, \chi_{{\rm pt},2}^2, \ldots$ to detect a change in the mean $\mu_{{\rm pt},k} = \mathbf{E}_0 \chi_{{\rm pt},k}^2$. For each packet type pt (i.e., ICMP, UDP, or TCP) at stage k, we use for thresholding CUSUM-type statistics $S_k^{\rm pt}$ that obey the recursions

$$S_k^{\rm pt} = \max\left\{0, S_{k-1}^{\rm pt} + \chi_{{\rm pt},k}^2 - \gamma_{{\rm pt},k}\right\}, \quad S_0^{\rm pt} = 0.$$

In this recursion, $\gamma_{{\rm pt},k}$ is a number that is estimated based on the recent history of the network traffic and is, in general, larger than the prechange mean $\mu_{{\rm pt},k}$. The attack is declared at the time moment $\tau_{\rm pt} = \min\{k : S_k^{\rm pt} \geqslant h_{\rm pt}\}$, where the threshold $h_{\rm pt}$ is chosen from the given FAR.

Due to space limitations, a comprehensive study of the batch-sequential algorithm will be presented elsewhere.

V. Experimental Results: DOS Attack Detection in the NS Testbed

In this section, we present the results of experimental analysis of the MNA-CUSUM detection algorithm that was described in Section IV-A. For simulations, we have used a testbed network simulator[1] NS with a network consisting of 100 nodes configured into a transit-stub topology (depicted in Fig. 2). The network contained one transit domain, four transit nodes, and 12 stub domains with 96 nodes.

[1]More information on the NS can be found at http://www.isi.edu/nsnam/ns/.

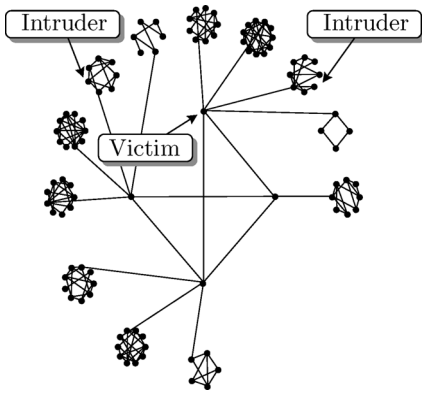

Fig. 2. Transit-stub network topology used in simulations.

Under regular conditions, the traffic consisted of approximately 5% ICMP packets, 15–20% UDP packets, and 75–80% TCP packets. The attacker's activity represented less than 1% of traffic. After a 120 s period (measured using the simulator time) of regular traffic, we initiated one of the following three kinds of DOS attacks targeted at the victim node: TCP SYN flooding, UDP packet storm, and ICMP ping flooding DOS attacks.

During a DOS attack, the attacker's traffic rapidly increased, reaching 20% of all traffic. We considered two scenarios for the attacker's traffic increase: linear and abrupt. In the former case, the level of 20% of all traffic was reached in a linear manner during a 60 s interval, while in the latter situation the traffic increased to the 20% level immediately after the beginning of the attack.

In the experiments, we simultaneously observed the statistics $N^{\mathrm{pt}}_{k,i}$ and B_k with the sampling rate 1 Hz (i.e., $\Delta_k = \Delta = 1$ sec.). As shown in Fig. 1 (see Section IV-A-1), the MNA-CUSUM detection algorithm detected the attack in its early stage. In this particular simulation of a UDP packet storm with linear traffic increase and the threshold of 84.66, the detection delay was 35 s.

The results of the experimental study for the abrupt attacks are shown in Fig. 3. Fig. 3(a) illustrates the operating characteristic (ADD versus FAR) obtained by simulations (dashed line) and by theoretical formulas (8) and (9) (solid line) for the UDP packet storm attack. Fig. 3(b) shows similar plots for the ICMP ping flooding DOS attack. The plots of theoretical estimates of ADD versus $|\log(\mathrm{FAR})|$ are straight lines with slopes that are equal to $[\theta^{\mathrm{pt}}_i - \mu^{\mathrm{pt}}_i - c^{\mathrm{pt}}_i]^{-1}$. It is seen that the experimental estimates of ADD are always bigger than the theoretical estimates. This is not surprising, since asymptotic formulas (9) ignore excesses over the thresholds of the decision statistics. In the figures, the average detection delay is measured in terms of the number of samples.

Tables I and II, respectively, summarize the operating characteristic for the UDP and ICMP attacks with the abrupt traffic increase. Tables III and IV illustrate the same performance for the linear traffic increase. In the tables, the average detection delay is shown in terms of the number of samples, the total number of observed packets, and the number of spoofed packets. The FAR is measured as the average number of false alarms per one sample. In addition, in the last column, we show the optimal values of the threshold h, which depend on the FAR.

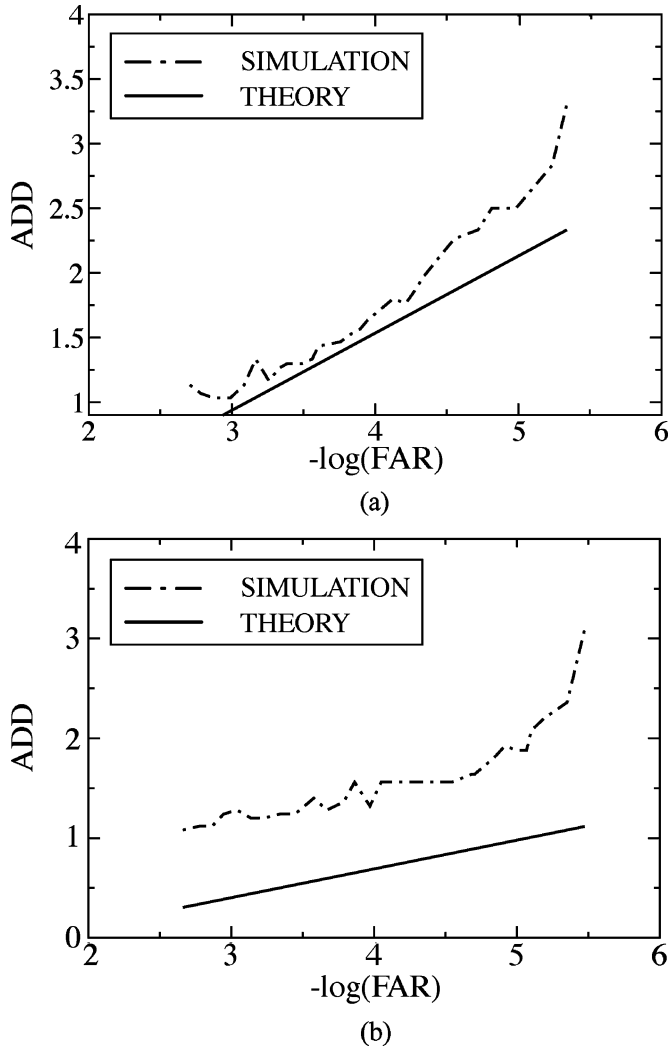

Fig. 3. Operating characteristic of the MNA-CUSUM detection algorithm for the UDP and the ICMP flooding attacks: abrupt traffic increase. (a) Operating characteristics for the UDP attack; (b) operating characteristics for the ICMP attack.

TABLE I
UDP PACKET STORM ATTACK: ABRUPT TRAFFIC INCREASE

FAR	ADD			Optimal Threshold h_{opt}
	# of samples	# of spoofed packets	# of total packets	
0.0667	1.13333	8.23	46.17	4.69072
0.0356	1.26667	9.27	52.27	6.33052
0.0193	1.63333	12.37	66.40	8.42176
0.0105	2.26667	17.90	93.87	9.94797
0.0053	2.83333	23.27	117.50	14.4556

TABLE II
ICMP FLOODING ATTACK: ABRUPT TRAFFIC INCREASE

FAR	ADD			Optimal Threshold h_{opt}
	# of samples	# of spoofed packets	# of total packets	
0.0698	1.08	8.84	42.68	2.60412
0.0354	1.24	10.12	47.24	5.16986
0.0174	1.56	13.52	60.76	4.37522
0.0091	1.64	14.12	63.40	2.87293
0.0042	3.08	27.84	121.20	14.8376

TABLE III
UDP PACKET STORM ATTACK: LINEAR TRAFFIC INCREASE

FAR	ADD			Optimal
	# of samples	# of spoofed packets	# of total packets	Threshold h_{opt}
0.0654	8.96	6.84	296.04	7.84487
0.0362	13.6	12.00	451.12	13.363
0.0197	17.56	19.00	589.84	16.2074
0.0107	31.08	49.48	1055.6	20.7214
0.0052	36.96	59.44	1105.96	34.1854

TABLE IV
ICMP FLOODING ATTACK: LINEAR TRAFFIC INCREASE

FAR	ADD			Optimal
	# of samples	# of spoofed packets	# of total packets	Threshold h_{opt}
0.0705	2.08	8.96	71.64	2.58553
0.0322	4.96	15.32	168.24	7.53183
0.0172	11.08	30.24	371.04	6.93251
0.0092	15.8	42.20	529.72	14.5461
0.0045	35.68	100.20	1202.4	29.0923

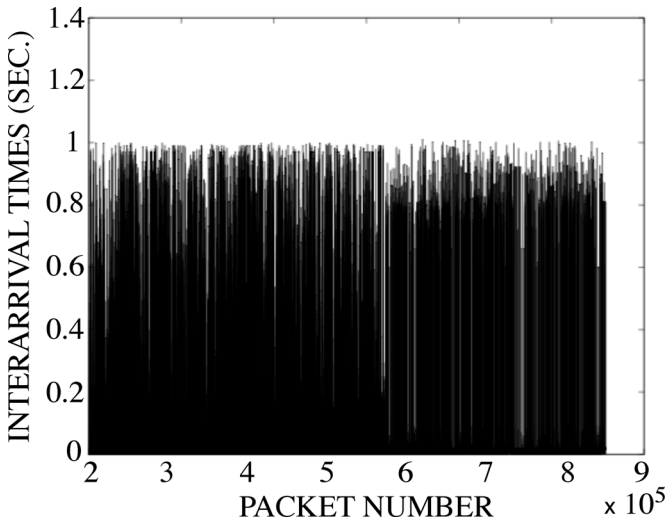

Fig. 4. Packet interarrival times for the neptune TCP SYN attack: observe that the attack is not visible to the naked eye.

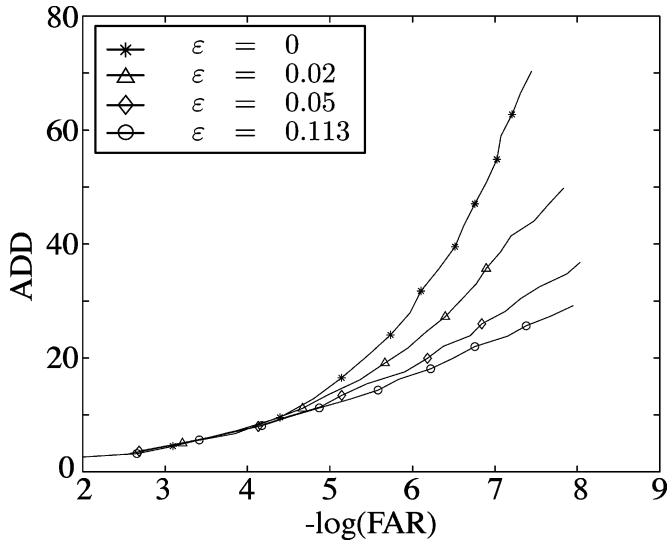

Fig. 5. Operating characteristics of the MNA-CUSUM algorithm.

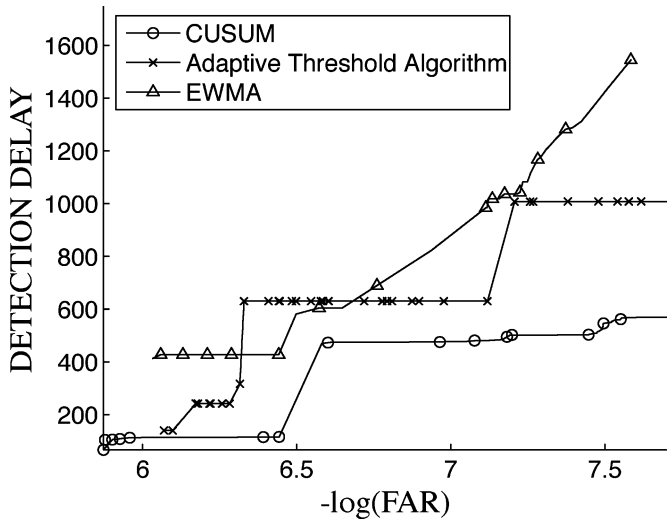

Fig. 6. Comparison of the MNA-CUSUM algorithm with the adaptive threshold and EWMA algorithms.

VI. EXPERIMENTAL RESULTS: DETECTION OF REAL TCP SYN DOS ATTACKS

We now illustrate the efficiency of the MNA-CUSUM detection algorithm using resampling techniques to obtain the operating characteristics in detecting real network intrusions, specifically the neptune TCP SYN flooding attack [7]. The data sets have been created by the MIT Lincoln Laboratory as a part of a DARPA Intrusion Detection Evaluation study in 1998. These data sets are available at http://www.ll.mit.edu/IST/ideval/.

In order to make the detection of the neptune attack more difficult, the intensity of the attack was reduced by rescaling the interpacket arrival times. The resulting data set is shown in Fig. 4. Clearly the attack cannot be seen to the naked eye.

We have observed the data with sampling period 0.1 s and used the optimized MNA-CUSUM detection algorithm with $\varepsilon_{\text{opt}} = 0.113$. Note that in this case, the pre- and postchange mean values are estimated as $\hat{\mu} = 4.394$ and $\hat{\theta} = 5.697$. The value of c in the decision statistic is equal to $c = \hat{\theta}\varepsilon_{\text{opt}} = 0.644$. The operating characteristics obtained by Monte Carlo resampling techniques for several values of ε are shown in Fig. 5. The optimal value $\varepsilon_{\text{opt}} = 0.113$ was obtained by minimizing the growth rate of curves fitted through the experimental values of $-\log(\text{FAR})$ in Fig. 5.

Fig. 6 compares the MNA-CUSUM algorithm with the adaptive threshold algorithm and exponentially weighted moving average (EWMA) algorithm. In this figure, we plot the detection delay in a particular run versus FAR. The FAR has been evaluated by the Monte Carlo experiment using resampling methods.

The adaptive threshold algorithm recently proposed in [36] triggers an alarm at time n if $\sum_{i=n-k+1}^{n} 1_{\{X_i \geqslant (\alpha+1)\bar{\mu}_{i-1}\}} \geqslant k$, where $k > 1$ is a parameter that indicates the number of successive threshold violations, $\alpha > 0$ is a parameter for the alarm threshold, and $\bar{\mu}_n$ is the mean rate estimated from measurements. It is updated by $\bar{\mu}_n = \beta\bar{\mu}_{n-1} + (1-\beta)X_n$, where β is the EWMA factor. We have used $k = 6$ and $\beta = 0.7$ and varied α from 0 to 0.5.

The EWMA algorithm is very popular in the statistical process control community. In the network security applications of interest, it is controlled by the four parameters η, α, L, and λ, where η is a smoothing constant that determines event intensity, α is the smoothed variance of the one-step-ahead prediction error, L determines the width of the EWMA control zone, and λ determines the effective memory of the EWMA filter. See [38] for further details. In our experiments we used $\eta = 0.2$, $\alpha = 0.0001$, and $L = 0.96$. The parameter λ was varied from 0 to 0.1 to obtain different FAR values.

The experiments illustrate that the MNA-CUSUM algorithm detects the TCP SYN attack much faster than both other algorithms for all FAR values. The MNA-CUSUM identified the attack at least 200 s prior to the adaptive threshold algorithm or EWMA. In some applications, this difference in detection speed may be extremely important in protecting mission critical network resources.

ACKNOWLEDGMENT

The authors would like to thank reviewers for comments that have improved this paper. The nonparametric multichannel change-point detection methods that are described in Section IV were first proposed in 1999 when the authors started to work on the DARPA AIA-SWIMM project at the University of Southern California.

REFERENCES

[1] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1993.

[2] P. K. Bhattacharya and D. Frierson, "A nonparametric control chart for detecting small disorders," *Ann. Statist.*, vol. 9, pp. 544–554, 1981.

[3] R. Blažek, H. Kim, B. Rozovskii, and A. Tartakovsky, "A novel approach to detection of 'denial-of-service' attacks via adaptive sequential and batch-sequential change-point detection methods," in *Proc. 2nd Annu. IEEE Syst., Man, Cybern. Inf. Assurance Workshop*, West Point, NY, 2001.

[4] ——, "The quickest sequential detection of intrusions in computer networks," in *Interface 2003*, Salt Lake City, UT, Mar. 12–15, 2003.

[5] B. E. Brodsky and B. S. Darkhovsky, *Nonparametric Methods in Change-Point Problems*. Dordrecht, The Netherlands: Kluwer, 1993.

[6] H. Cramér, *Mathematical Methods of Statistics*. Princeton, NJ: Princeton Univ. Press, 1946.

[7] Daemon9, Route, and Infinity, "Project Neptune," *Phrack Mag.*, vol. 7, no. 48, 1996 [Online]. Available: http://www.phrack.org/show.php?p=48&a=13, File 13 of 18

[8] V. P. Dragalin, "Adaptive procedures for detecting a change in distribution," in *Proc. 4th Wuerzburg-Umea Conf. Statist.*, 1996, pp. 87–103.

[9] F. Feather and R. Maxon, "Fault detection in an ethernet network using anomaly signature matching," in *ACM Sigcomm*, 1993, vol. 23.

[10] S. Gibson, Distributed reflection denial of service: Description and analysis of a potent, increasingly prevalent, and worrisome Internet attack Gibson Research Corp., 2002 [Online]. Available: http://www.grc.com/dos/drdos.htm

[11] T. M. Gil and M. Poletter, "MULTOPS: A data-structure for bandwidth attack detection," in *Proc. USENIX Security Symp. '01*, Washington, DC, Aug. 2001, pp. 23–38.

[12] L. Gordon and M. Pollak, "An efficient sequential nonparametric scheme for detecting a change in distribution," *Ann. Statist.*, vol. 22, pp. 763–804, 1994.

[13] A. Hussain, J. Heidemann, and C. Papadopoulos, "A framework for classifying denial of service attacks," in *Proc. Sigcomm 2003*, Karlsruhe, Germany, 2003.

[14] S. Kent, "On the trial of intrusions into information systems," *IEEE Spectrum*, vol. 37, no. 12, pp. 52–56, December 2000.

[15] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *J. Roy. Statist. Soc. B*, vol. 57, no. 4, pp. 613–658, 1995.

[16] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, pp. 1908–1987, 1971.

[17] R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling high bandwidth aggregates in the network," *ACM Comput. Commun. Rev.*, Jul. 2001.

[18] D. McDonald, "A Cusum procedure based on sequential ranks," *Naval Res. Logist.*, vol. 37, pp. 627–646, 1990.

[19] J. Mirkovic, G. Prier, and P. Reiher, "Attacking DDoS at the source," in *Proc. 10th IEEE Int. Conf. Network Protocols (ICNP)*, Paris, France, Nov. 10–12, 2002, pp. 312–321.

[20] D. Moore, G. Voelker, and S. Savage, "Inferring Internet denial of service activity," in *Proc. USENIX Security Symp.*, Washington, DC, 2001, pp. 9–22.

[21] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, pp. 100–115, 1954.

[22] C. Papadopoulos, R. Lindell, J. Mehringer, A. Hussain, and R. Govindan, "COSSACK: Coordinated suppression of simultaneous attacks," in *Proc. 3rd DARPA Information Survivability Conf. Exposition (Discex III)*, Washington, DC, 2003, vol. 1, pp. 2–13.

[23] V. Paxson, "Bro: A system for detecting network intruders in real-time," *Comput. Netw.*, vol. 31, no. 23–24, pp. 2435–2463, 1999.

[24] M. Pollak, "Optimal detection of a change in distribution," *Ann. Statist.*, vol. 13, pp. 206–227, 1985.

[25] H. Robbins and D. Siegmund, "The expected sample size of some tests of power one," *Ann. Statist.*, vol. 2, pp. 415–436, 1974.

[26] M. Roesch, "Snort: Lightweight intrusion detection for networks," in *Proc. 13th Syst. Admin. Conf. (LISA)*, 1999, pp. 229–238.

[27] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory Prob. Appl.*, vol. 8, pp. 22–46, 1963.

[28] A. G. Tartakovsky, *Sequential Methods in the Theory of Information Systems* (in Russian). Moscow, Russia: Radio i Svyaz', 1991.

[29] ——, "Efficiency of the generalized Neyman–Pearson test for detecting changes in a multichannel system," *Prob. Inf. Transmission*, vol. 28, pp. 341–350, 1992.

[30] A. G. Tartakovsky and I. A. Ivanova, "Comparison of some sequential rules for detecting changes in distributions," *Prob. Inf. Transmission*, vol. 28, pp. 117–124, 1992.

[31] A. G. Tartakovsky, "Asymptotic properties of CUSUM and Shiryaev's procedures for detecting a change in a nonhomogeneous Gaussian process," *Math. Meth. Statist.*, vol. 4, no. 4, pp. 389–404, 1995.

[32] A. G. Tartakovsky, B. L. Rozovskii, R. Blažek, and H. Kim, "Detection of intrusions in information systems by sequential change-point methods," *Statist. Methodol.*, vol. 3, no. 3, pp. 252–340, 2006.

[33] A. G. Tartakovsky and V. Veeravalli, "Change-point detection in multichannel and distributed systems with applications," in *Applications of Sequential Methodologies*, N. Mukhopadhyay, S. Datta, and S. Chattopadhyay, Eds. New York: Marcel Dekker, 2004, pp. 339–370.

[34] ——, "General asymptotic Bayesian theory of quickest change detection," *Theory Prob. Appl.*, vol. 49, no. 3, pp. 458–497, 2005.

[35] R. R. Talpade, G. Kim, and S. Khurana, "NOMAD: Traffic-based network monitoring framework for anomaly detection," in *Proc. 4th IEEE Symp. Comput. Commun.*, 1999, pp. 442–451.

[36] V. A. Siris and F. Papagalou, "Application of anomaly detection algorithms for detecting SYN flooding attacks," in *Proc. IEEE Global Telecommun. Conf. (IEEE GLOBECOM 2004)*, Dallas, TX, 2004, vol. 4, pp. 2050–2054.

[37] H. Wang, D. Zhang, and K. Shin, "Detecting SYN flooding attacks," in *Proc. IEEE Infocom*, New York, 2002, pp. 1530–1539.

[38] N. Ye, S. Vilbert, and Q. Chen, "Computer intrusion detection through EWMA for autocorrelated and uncorrelated data," *IEEE Trans. Reliab.*, vol. 52, no. 1, pp. 75–82, Mar. 2003.

Alexander G. Tartakovsky (M'01–SM'02) received the M.S. degree in electrical engineering from Moscow Aviation Institute, Russia, in 1978. He received the Ph.D. degree in statistics and information theory and the doctor of science degree in statistics and control from Moscow Institute of Physics and Technology, Russia, in 1981 and 1990, respectively.

He is the Associate Director of the Center for Applied Mathematical Sciences, University of Southern California, Los Angeles. His research interests include theoretical and applied statistics; sequential analysis; change-point detection phenomena; adaptive, minimax, and robust methods for overcoming prior uncertainty; pattern recognition; speech recognition and speaker identification; target detection and tracking; information fusion; and network security. He is the author of one book and more than 70 papers in the areas indicated above.

Dr. Tartakovsky is a member of the Institute of Mathematical Statistics, SPIE, and Information Fusion Society.

Boris L. Rozovskii is Professor of mathematics and Director of the Center for Applied Mathematical Sciences at the University of Southern California, Los Angeles. His research interests are in the areas of stochastic partial differential equations, fluid dynamics, nonlinear filtering, prediction and smoothing, inverse problems for randomly perturbed systems, target tracking, intrusion detection, and mathematical modeling of the Internet. He is the author of four books in the general area of stochastic systems and more than 100 research papers.

Prof. Rozovskii is a Fellow of Institute of Mathematical Statistics. He is a recipient of many awards.

Rudolf B. Blažek received the M.S. degree in mathematics (mathematical statistics) from Charles University, Czech Republic, in 1991 and the Ph.D. degree in statistics from Michigan State University, East Lansing, in 1998.

He is with Advanced Science and Novel Technology, Torrance, CA. He is also with the Center for Applied Mathematical Sciences, University of Southern California (USC), and is a Lecturer in the Department of Mathematics, USC. His research interests include mathematical statistics, resampling theory, α-stable processes, multiresolution analysis, statistical image analysis, and statistical methods in network security.

Dr. Blažek is a member of the American Mathematical Society and the American Statistical Association.

Hongjoong Kim received the Ph.D. degree from the Department of Applied Mathematics and Statistics, State University of New York at Stony Brook, in 2000.

He is an Assistant Professor in the Department of Mathematics at the Korea University in Seoul, Korea. His research interests include fluids in porous media, computational fluid dynamics, stochastic partial differential equations, numerical analysis, and mathematical modeling of the Internet.