Available online at www.sciencedirect.com
SCIENCE DIRECT®
Statistical Methodology 3 (2006) 252–293
Statistical Methodology
Official Journal of the International Indian Statistical Association
www.elsevier.com/locate/stamet

Detection of intrusions in information systems by sequential change-point methods

Alexander G. Tartakovsky[a,*], Boris L. Rozovskii[a], Rudolf B. Blažek[a,b], Hongjoong Kim[c]

[a] *Department of Mathematics and the Center for Applied Mathematical Sciences, University of Southern California, Kaprielian Hall, KAP 108, 3620 S. Vermont Avenue, Los Angeles, CA 90089-2532, USA*
[b] *Advanced Science and Novel Technology, 2790 Skypark Dr. Ste 104, Torrance, CA 90505-5300, USA*[1]
[c] *Department of Mathematics, Korea University, 1, Anam-dong, Sungbuk-ku, Seoul, 136-701, Republic of Korea*

Received 7 January 2005; received in revised form 3 May 2005; accepted 6 May 2005

Abstract

Sequential multi-chart detection procedures for detecting changes in multichannel sensor systems are developed. In the case of complete information on pre-change and post-change distributions, the detection algorithm represents a likelihood ratio-based multichannel generalization of Page's cumulative sum (CUSUM) test that is applied to general stochastic models that may include correlated and nonstationary observations. There are many potential application areas where it is necessary to consider multichannel generalizations and general statistical models. In this paper our main motivation for doing so is network security: rapid anomaly detection for an early detection of attacks in computer networks that lead to changes in network traffic. Moreover, this kind of application encourages the development of a nonparametric multichannel detection test that does not use exact pre-change (legitimate) and post-change (attack) traffic models. The proposed nonparametric method can be effectively applied to detect a wide variety of attacks such as denial-of-service attacks, worm-based attacks, port-scanning, and man-in-the-middle attacks. In addition, we propose a multichannel CUSUM procedure that is based on binary quantized data; this procedure turns out to be more efficient than the previous two algorithms in certain scenarios. All proposed detection algorithms are based on the change-point detection theory. They utilize the thresholding of test statistics to achieve a fixed rate of false alarms, while allowing changes in statistical models to be detected "as soon as possible". Theoretical frameworks for the performance analysis of detection procedures, as well

* Corresponding author. Tel.: +1 213 740 2450; fax: +1 213 740 2424.
E-mail addresses: tartakov@math.usc.edu (A.G. Tartakovsky), rozovski@math.usc.edu (B.L. Rozovskii), blazek@math.usc.edu (R.B. Blažek), hongjoong@korea.ac.kr (H. Kim).
[1] From September 2003.

1572-3127/$ - see front matter
doi:10.1016/j.stamet.2005.05.003

as results of Monte Carlo simulations for a Poisson example and results of detecting real flooding attacks, are presented.

Keywords: Change-point detection; Sequential tests; Multichannel information systems; Rapid detection; Page's test; Cumulative sum; Intrusion detection; Denial of service

1. Introduction

The study in this paper is motivated by intrusion detection applications. During the past few years, considerable interest in the field of defense against cyber-terrorism in general, and network intrusion detection in particular, has been induced by a series of *external* and *internal* attacks on important corporate and governmental networks, server clusters, and other network resources. This interest has been further fueled by external distributed denial-of-service (DOS) attacks against several well known Internet servers such as Yahoo, Amazon, eBay, and E*Trade [18]. Other examples are Internet-wide worm attacks and stealthy attacks by intruders posing as regular users. Software allowing hackers to initiate many varieties of external (e.g. DOS and worm) and internal (e.g. Address Resolution Protocol Men-in-the-Middle (ARP MiM)) attacks is becoming more and more available and easy to use. The aforementioned attacks represent serious threats for information systems, causing failures and interrupting services to users [14]. As a result, the defense against these serious threats is rapidly gaining importance.

A number of defense methods have been proposed for dealing with DOS attacks. For example, SYN cookies, SYN defender, SYN cache, and SYN proxying can be used to counter TCP (Transmission Control Protocol) SYN flooding attacks [2,7,26,29]. However, these defense tools are installed at the firewall of the victim server or inside the victim server and do not indicate the sources of the SYN flooding. Therefore, these tools have to use the expensive IP traceback to locate the flooding sources.

There are a wide variety of other intrusion detection methods proposed in the literature, which are mostly based on artificial intelligence techniques, including expert systems, neural networks, data mining, pattern matching, and many others. See [9] for a detailed overview. Existing intrusion detection systems (IDSs) can be classified as either Signature Detection Systems or Anomaly Detection Systems [9,22]. Signature-based detection systems detect attacks by comparing the observed patterns of the network traffic with known attack signatures from a database [15]. If the true attack belongs to the class of attacks listed in the database, then it can be successfully detected and, moreover, identified/classified. Anomaly-based IDSs compare the parameters of the observed traffic with legitimate network traffic. The attack is declared once a deviation from legitimate traffic is observed. Both classes of systems have certain pros and cons [22].

The approach undertaken in the present paper belongs to the class of anomaly IDS. Typically network intrusions (e.g. DOS, worm, port-scanning, and ARP MiM attacks) occur at unknown points in time and lead to changes in the statistical properties of certain observables. It is therefore intuitively appealing to formulate the problem of detecting attacks as a quickest change-point detection (CPD) problem: to detect changes in statistical models as rapidly as possible (i.e. with minimal average delays) while maintaining the false alarm rate (FAR) at a given level. See [1,5, 24,25,27,28,30–45] for relevant results of CPD theory.

In the standard formulation of the CPD problem, there is a sequence of observations whose distribution changes abruptly at some unknown instant; the goal is to detect this change as soon as possible, subject to false alarm constraints [1,27,31,34,36]. The problem of detecting an abrupt