Available online at www.sciencedirect.com

ScienceDirect

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

Computers & Security

A comparative analysis of detection metrics for covert timing channels

*Rennie Archibald, Dipak Ghosal**

Department of Computer Science, University of California, Davis, CA 95616, USA

ARTICLE INFO

Article history:
Received 25 July 2013
Received in revised form
14 March 2014
Accepted 20 March 2014
Available online 2 May 2014

Keywords:
Covert timing channels
Detection
Welch's *t*-test
Corrected conditional entropy
Regularity test
Shape test
Entropy test

ABSTRACT

Methods to detect covert timing channels (CTCs) can be categorized into three broad classes: shape tests which include the Kolmogorov–Smirnov (KS) test, entropy tests which include first order entropy test, corrected conditional entropy (CCE) test, and Kullback–Leibler (KL) divergence test, and regularity tests. This paper contributes towards understanding and advancing the state-of-the-art of CTC detection methods. First, we present a detailed analysis of the performance of the well-known tests that are used to detect three main types of CTCs, namely, JitterBug, model-based CTC (MB-CTC) and time-replay CTC (TR-CTC). The performance analysis is carried out in an enterprise-like setting, employing large traffic traces. The detection methods are compared with respect to their applicability, computational complexity, and the classification rates for the three types of CTCs. In addition to evaluating the existing methods, we propose a new shape test based on the Welch's *t*-test and compare its performance with existing detection methods. We show that the classification rate of Welch's *t*-test is at least at par with other existing detection methods while having a relatively lower computational cost. The results also show that the Welch's *t*-test outperforms the CCE test in detecting JitterBug, while the CCE test has a better performance in detecting the TR-CTC. Furthermore, both tests perform comparably on the MB-CTC. Finally, we study the feasibility of using a multi-feature SVM classifier to increase the classification rate. We show that by combining the Welch's *t*-test we are able to increase the classification rate of MB-CTCs from 0.67 (using a single regularity measure) to 0.94.

1. Introduction

There are two main approaches to thwart covert timing channels (CTCs) - blind disruption and detection-and-disruption. Timing disruption techniques investigated in (Kang and Moskowitz, 1993; Kang et al., 2005) propose the use of a network pump that can randomize and/or homogenize inter-packet delays (IPDs), and thereby disrupt timing channels. The network pump can be deployed at the network edge of a secure facility and (blindly) applied to all flows that are potential carriers of timing channels. While network pumps can be effective in disrupting CTCs, they may be of limited applicability for applications flows with a QoS requirement such as VoIP, streaming video, SSH, or remote desktop applications. Additionally, there is also the issue of scalability with the ever increasing data rates at the egress point of the enterprise. For the above mentioned reasons, it is

* *Corresponding author*. Tel.: +1 530 754 9251.
E-mail addresses: archibal@cs.ucdavis.edu (R. Archibald), ghosal@cs.ucdavis.edu, dghosal@ucdavis.edu (D. Ghosal).
http://dx.doi.org/10.1016/j.cose.2014.03.007

important to develop CTC detection methods that can identify network flows that are covert carriers before timing disruption techniques are applied.

Methods to detect CTCs can be categorized into three broad classes - shape tests, regularity tests, and entropy tests. Shape tests quantify the difference between IPD distributions of known (clean) overt packet streams and potentially covert packet streams. The most widely used shape test is the Kolmogorov–Smirnov (KS) test (Liu et al., 2010). Regularity tests quantify the similarities of the statistical features in contiguous subsets of IPDs (Cabuk et al., 2004). In covert traffic the statistical features in contiguous subsets of IPDs are similar particularly, when the IPDs are generated following some well defined statistical models. This is in contrast to real network traffic which do not follow well defined statistical models and may exhibit significantly different statistical features over contiguous subset of IPDs. Finally, there are entropy tests, which are a subset of shape tests, but are considered as a separate class because of the depth of prior work. Entropy tests can be used to distinguish IPDs generated by independent and identically distributed (iid) statistical models which, in general, yield larger entropy values than those obtained for legitimate overt traffic. The entropy tests that have been examined in the literature include the first-order entropy test, the corrected conditional entropy (CCE) test (Gianvecchio & Wang, Nov.–Dec. 2011), and the Kullback–Leibler (KL) divergence test (Archibald and Ghosal, 2012).

A key contribution of this paper is a comprehensive analysis of the three aforementioned classes of detection methods. Our analysis is based on real traffic traces obtained from WAND at the University of Waikato (W. N. R. Group, 2005). The analysis is based on two distinct application traffic, namely HTTP and SSH traffic as overt carriers. The HTTP and SSH flows are extracted from the captures and a subset of the flows are injected with covert messages. We have considered three different types of CTCs, namely, JitterBug (Shah et al., 2006), time replay (TR-) CTCs (Liu et al., 2009; Cabuk, 2006), and model-based (MB-) CTCs (Gianvecchio & Wang, Nov.–Dec. 2011). We use the classification rate as the metric for comparing the performance.

The results of our analysis of the existing methods demonstrate a need for a detection method that performs adequately on the JitterBug CTC and is computationally inexpensive. To address this, we propose new regularity tests and a new computationally low cost shape test. Our proposed regularity tests extend existing methods; they operates on a windowed vector of IPDs and extracts higher-moment features, in particular, the kurtosis and the skew. Additionally, we propose a new shape test based on Welch's *t*-test.

Our results show, that no single test appears to be effective in detecting all three types of CTCs. Instead, the two top performing tests, the CCE test and Welch's *t*-test, both effectively detect two of the three CTCs. We find that Welch's *t*-test outperforms all existing methods to detect the JitterBug CTC, while the CCE test outperforms Welch's *t*-test in the detection of TR-CTCs. Both tests perform comparably in detecting MB-CTCs with accuracy up to 97% depending on how the model is parameterized and how many IPDs are used for testing. The advantage of Welch's *t*-test over the CCE test for MB-CTC detection is the lower computation cost. Finally, we employ SVM to study a multi-feature regularity test and find that it compares well to the CCE test and the Welch's *t*-tests, achieving over 94% classification accuracy.

The remainder of this paper is organized into four sections. We present the background and related work in Section 2. In Section 3 we give a detailed account of the tests we introduce in this paper. Our experimental set-up and testing procedure is discussed in Section 4. A detailed analysis of the experimental results is presented in Section 5.

2. Background and related work

In this section we first discuss the network traffic based CTCs that we have considered in the evaluation of the existing detection methods that have been proposed in the literature. There are two broad classes of network traffic based CTCs. In the first type, covert data is injected by modulating the inter-packet delays (IPDs). Examples include JitterBug (Shah et al., 2006), model based CTCs (MB-CTCs) (Liu et al., 2010; Gianvecchio & Wang, Nov.–Dec. 2011; Archibald and Ghosal, 2012), and time-replay CTCs (TR-CTCs) (Cabuk, 2006; Liu et al., 2009). In the second type, some combinatorial function is applied to the IPDs to inject the covert message. This type of CTCs are also referred to as *combinatorial* CTCs, which can be based on a single flow as in (Cabuk et al., 2004) or multiple flows as in (Luo et al., 2011, 2009; Houmansadr and Borisov, 2011). In the case of a single flow, packets are transmitted in groups where the size of the group is used to encode information (Cabuk et al., 2004). A pre-determined IPD is used to demarcate groups. In the case of multiple flows, both the size of the group and the flow in which the group of packets arrive are used to encode information (Luo et al., 2011).

We study detection methods for CTCs that modulate covert data into the IPDs. Each of the CTCs used to test the detection methods fit the same theoretical framework for stego-systems defined in (van Tilborg, 2005) and then expanded upon in (Kiayias et al., 2012). A complete stego-system includes all information and algorithms needed for Alice to encode the covert message and Bob to decode the message (van Tilborg, 2005). Formally, a stego-system consists of three elements, a seed generator *SK*, a steganographic encoding algorithm *SE*, and a steganographic decoding algorithm *SD*. Thus, the stego-system is defined by the triplet (*SK*,*SE*,*SD*) (van Tilborg, 2005; Kiayias et al., 2012). We summarize prior CTCs, and map the CTC systems to the formal stego-system definition in the subsequent section.

2.1. Covert timing channels

The JitterBug system employs Telnet traffic to slowly leak information over a network (Shah et al., 2006). The JitterBug inserts an additional delay into the Telnet traffic generated by the sender. The delays imposed by the sender are limited to at most w milliseconds (ms). The sender transmits a bit 0 by adding sufficient delay to the observed IPD such that the modified IPD modulo $\lfloor w/2 \rfloor$ ms is 0. Similarly, a bit 1 is transmitted by increasing the observed IPD to so that modified IPD modulo w ms is 0 (Shah et al., 2006). The sender's codebook *SK* and receiver's decode-book *SD*, consist of the shared value w.