Towards a High-speed Router-based Anomaly/Intrusion Detection System

Zhichun Li, Yan Gao, Yan Chen
Northwestern University {lizc,yga751,ychen}@cs.northwestern.edu

I. Introduction

Traffic anomalies and attacks are commonplace in today's networks, and identifying them rapidly and accurately is critical for large networks. With the rapid growth of network bandwidth and fast emergence of new attacks/worms, existing network intrusion detection systems (IDS) are insufficient for the following two reasons.

First, they are mostly host-based or located on low-end routers, and not scalable to high-speed networks. However, it is crucial to identify fast propagation of worms in their early phases, which can only possibly be achieved by detection at high speed edge/backbone routers instead of at end hosts. Unfortunately, the existing schemes are not scalable to the link speeds and number of flows for high-speed networks. According to a recent research agenda [7] by DARPA, detection on edge networks is particularly critical, powerful and efficient (without deploying IDSs on all the edge hosts).

Second, most of the existing approaches are signature based, which cannot detect unknown attacks. Statistical IDSs are therefore proposed to detect anomalies. Current systems are mostly designed to detect based on the overall traffic, thus they tend to be inaccurate or cannot find real attack flows for mitigation even when spotting anomalies. For instance, the state-of-the-art stateless router-based SYN flooding detection techniques, Change Point Monitoring (CPM), use the statistical behavior of SYN-FIN, or SYN-SYN/ACK packet pairs based on overall traffic for detection [5]. It detects well with pure SYN flooding traffic. However, it suffers high false positives when there is port scan traffic, because port scans can also cause lots of SYN without SYN/ACK or FIN.

There are also a few flow-level detection schemes. However most of them are not scalable to high speed networks. For instance, Threshold Random Walk (TRW) can accurately and quickly detect port scans [2], but it is vulnerable to DoS attacks with randomly spoofed IP addresses, which will cause memory overflow because it needs to store the status of *every* source IP address. It was recently improved by limiting its memory consumption with approximate caches (AC) [6]. However, now spoofed DoS attacks will cause collisions in AC, and make the real port scans undetected. Recently, PCF was recently proposed for scalable network detection [3]. But even when attacks are detected, the attacker or victim is still unknown, making mitigation impossible.

II. HRAID system design

To address these challenges, we propose a new paradigm called High-speed Router-based Anomaly and Intrusion Detection system, HRAID, a flow-level anomaly detection system, which can not only be used on high speed networks, but also detection different types of attacks rapidly and accurately. Two performance features are strongly desirable for practical heavy change detection systems: 1) a small amount of memory usage (to be implemented in SRAM); 2) a small number of memory accesses per packet. Leveraging our recent work on data streaming computation and in particular, sketches [4], we can satisfy both requirements.

Approaches	Spoofed DoS	Non-spoofed DoS	HScan	VScan
HRAID	Yes	Yes	Yes	Yes
TRW(AC)	No	No	Yes	No
CPM	Yes, but high FP with scans		No	No

TABLE I
FUNCTIONALITY COMPARISON OF FOUR APPROACHES. HSCAN STANDS FOR HORIZONTAL SCAN, AND VSCAN, VERTICAL SCAN

As the first step towards this ambitious goal, our threat model includes various port scans (horizontal scan, vertical scan) and TCP SYN flooding (DoS attacks). Our goal is to identify them simultaneously in near real-time on high speed networks, and to obtain the attacks' key characteristics for mitigating the attacks. Our approach is suitable to be deployed to edge network routers and even backbone routers. In Table I, we compared our approach with others.

We analyze the attributes in TCP/IP headers and select a small set of metrics for flow-level sketch-based traffic monitoring. Time series analysis (TSA) algorithms like exponentially-weighted moving average (EWMA) and Holt-winter are applied to detect anomalies and intrusions. To further distinguish different type of attacks we proposed the two-dimensional sketches (2D sketches). Both analytical and empirical results show the effectiveness of the 2D sketches. The HRAID system architecture is shown in Figure 1.

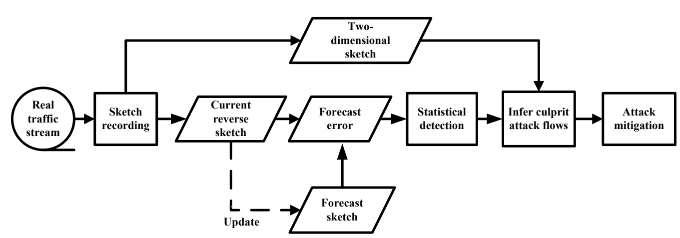

Fig. 1. HRAID System architecture.

Reversible sketch [4] is a compact data streaming data structure, which can record network traffic in terms of {`Key, Value`} pairs, aggregate the traffic, and output a set of heavy keys in the traffic accurately. Using sketches, we can archive extremely fast recording speed for network traffic. Our prototype single FPGA board implementation achieve a throughput of over 16 Gbps for 40-byte-packet streams (the worst case). Another good property of reversible sketch is its linearity. Thus in addition to supporting various time-series forecasting, it enables aggregation both in temporal and spatial domains, *e.g.*, for distributed attack detection.

A forecast TSA model can help remove noises from detections. As shown in Figure 1, by subtracting the forecast sketch from the current one, we obtain the forecast error sketches. Intuitively, a large forecast error implies there is an anomaly, thus the forecast error is the key metric for detection in our system. Moreover, we adopt the 2D sketches to further distinguish different types of attacks, such as SYN Flooding vs. vertical scan, and SYN flooding vs. horizontal scan. Finally, we use the key characteristics of the culprit flows revealed by the reversible sketches to mitigate the attacks.

For **reversible sketches based flow level detection**. We denote the key of the sketch as `K`, the feature value recorded in the sketch as `V`, and the reversible sketch as `RS(K,V)`. We also denote the destination IP address as `DIP`, the source IP address as `SIP` and the destination port as `Dport`. We only consider the attacks in TCP traffic. Our detection has the following three steps: First, we use `RS((DIP, Dport), SYN−SYN/ACK)` to detect SYN flooding attacks because they usually target a certain service as characterized by the `Dport` on a small set of machine(s). We use the similar metric SYN-SYN/ACK as [5], [3]. Secondly, we use `RS((SIP, DIP), SYN−SYN/ACK)` to detect any intruder tries attack to a particular IP address. They can be non-spoofed SYN flooding attacks or vertical scans. Thirdly, we use `RS((SIP, Dport), SYN−SYN/ACK)` to detect any source IP which causes a large number of uncompleted SYN connections to a particular destination port. They can be non-spoofed SYN flooding or horizontal scan.

After this step, we still cannot fully classify three types of attacks. The problem is due to traffic anomalies often have complicated structures: they are often multidimensional. For example, if the port distribution of an attack is unknown, it is very hard to distinguish non-IP-Spoofing SYN flooding attacks from vertical scans because both of them will observe a single (or a small number) of source IPs send a large number of un-responded SYN packets to the destination. So it is essential to know the port distribution given a specific source IP and destination IP pair ($\{$`SIP,DIP`$\}$). Put it more general, we need to know the conditional distributions of one dimension given the other.

To address this challenge, we design a novel **two-dimensional k-ary sketch (2D sketch)**. Instead of using H independent one-dimensional hash tables, we use H independent two-dimensional hash tables (matrices). For each two-dimensional hash matrix, we hash two groups of fields to its x-dimension or y-dimension. Considering the previous example of separating SYN flooding attacks and vertical scans, the x-dim is the combination of source IP and destination IP $\{$`SIP,DIP`$\}$, and the y-dim is the destination port `Dport`, as shown in Figure 2. Similarly, we can update all the H matrices in the recording stage.

In the detection stage, when we find an attack by using reversible sketch or other methods (*i.e.*, the $\{$`SIP,DIP`$\}$ is known), we use the column of buckets in the hash matrix selected by the $\{$`SIP,DIP`$\}$ to infer the distribution of `Dport` and pinpoint the type of attack, *e.g.*, a SYN flooding or a vertical scan.

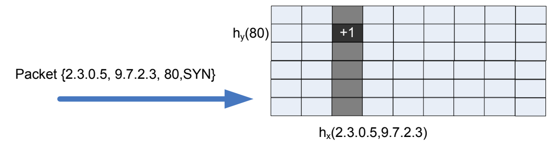

Fig. 2. An example of UPDATE operation for 2D sketch

We use *Generalized Increment-Decrement Counter Model* proposed in [3] to obtain the bound of the false positives and negatives for 2D sketches. For the detailed proof please refer to our technique report [1].

III. Preliminary evaluation

Using reversible sketches and 2D sketches only consumes a quite small amount of memory and can be easily implemented in hardware for very high speed networks, but they may introduce some false positives (FP) and negatives (FN) compared with the same detection algorithm implemented using accurate flow table to hold per-flow information (we call it non-sketch method).

In our evaluation, we used the 3-day traces collected from the gateway of Northwestern University in 2005. It consists of 536M Netflow records (3.48TB traffic). In the speed test, we found we could use a 1.6GHz AMD Opteron machine to record 536M items with one reversible sketch in 46.2 seconds, *i.e.*, 11.6M insertions/second. For 2D sketches we can archive even higher performance. In HRAID we used 3 reversible sketches and 2 2D sketches which totally only consume 9MB memory for the recording stage. And we need only 15 memory accesses for each reversible sketch and 3 memory accesses for each 2D sketch per packet. When comparing with non-sketch version, the FN is less than 0.52% and FP is less than 2.35%. Moreover, we further compared our method with other approaches and results show that we can detect most attacks detected by others and can avoid obvious false positives introduce by other approaches. Also, we manually validate the results and find what we detected are almost real attacks.

References

[1] Gao, Y., Li, Z., and Chen, Y. Towards a high-speed router-based anomaly/intrusion detection system. `http://list.cs.northwestern.edu/graid.html` (poster available).

[2] Jung, J., Paxson, V., et al. Fast portscan detection using sequential hypothesis testing. In *Proceedings of the IEEE Symposium on Security and Privacy* (2004).

[3] Kompella, R. R., et al. On scalable attack detection in the network. In *Proc. of ACM SIGCOMM IMC* (2004).

[4] Schweller, R., Gupta, A., Parsons, E., and Chen, Y. Reversible sketches for efficient and accurate change detection over network data streams. In *Proc. of ACM SIGCOMM IMC* (2004).

[5] Wang, H., Zhang, D., and Shin, K. G. Detecting SYN flooding attacks. In *Proc. of IEEE INFOCOM* (2002).

[6] Weaver, N., et al. Very fast containment of scanning worms. In *USENIX Security Symposium* (2004).

[7] Weaver, N., Paxson, V., Staniford, S., and Cunningham, R. Large scale malicious code: A research agenda. Tech. Rep. DARPA-sponsored report, DARPA, 2003.