

An Accurate Sampling Scheme for Detecting SYN Flooding Attacks and Portscans

Maciej Korczyński* Lucjan Janowski† and Andrzej Duda*

*Grenoble Informatics Laboratory, Grenoble Institute of Technology, Grenoble, France

†AGH University of Science and Technology, Cracow, Poland

Abstract—In this paper, we propose an accurate sampling scheme for defeating SYN flooding attacks as well as TCP portscan activity. The scheme examines TCP segments to find at least one of multiple ACK segments coming from the server to validate legitimate connections. The method achieves good detection performance with false positive rate close to zero even for very low sampling rates. Our trace-based simulations show that the effectiveness of the proposed scheme only relies on the sampling rate regardless on the sampling method.

I. INTRODUCTION

Distributed Denial-of-Service (DDoS) attacks and portscan activity have strong influence on the Internet security. Their major cause is TCP SYN flooding that consists in sending many SYN segments from a large number of compromised computers. It prevents victim machines or even whole sub-networks from offering a service to their legitimate users. A portscan activity is usually a precursor of an intrusion attempt—a compromised computer sends multiple SYN segments to probe other hosts for open ports to gain control over more computers that become potential attackers. Both types of traffic exploit the inherent asymmetry in the TCP three-way handshake mechanism and the fact that the victim cannot authenticate TCP SYN segments it receives. As a result, malicious packets can easily reach the victim without its approval.

Among various defense mechanisms, SYN flooding detection mechanisms placed in border routers have received much attention in the recent literature [1], [2], [3], [4], [5]. All these methods take advantage of the relations between TCP control segments responsible for connection establishment and release. However, they all may fail when routers sample traffic by inspecting only some packets and drawing conclusions about the whole behavior of the system. As a consequence of packet sampling, detecting DDoS attacks and portscans becomes even more difficult.

In this paper, we propose a novel and accurate sampling detection scheme of high-volume malicious traffic composed of SYN flooding attacks and low-volume portscan activity. The scheme examines TCP segments to find at least one of multiple ACK segments coming from the server. In this case, it concludes that the connection was successfully established so its opening SYN segment was not a part of a SYN flooding attack or portscan activity. This principle is particularly suitable when routers sample packets with very low rates. We combine the proposed method with a rate

limiting scheme that controls traffic rates and compare with three other representative detection methods. We show that our method achieves a high attack detection rate (true positives). In comparison to existing methods, we significantly reduce the false positive rate, i.e., when legitimate packets are classified as malicious ones.

We also study the impact of three basic packet sampling techniques proposed by PSAMP IETF working group [6] on our detection scheme. The results reveal that even the simplest and the most commonly used sampling technique—*systematic* sampling also known as *deterministic* sampling [7], performs fairly well under low sampling rates when combined with our detection and rate limiting method. Unlike some other proposals that used network simulations or experiments on obsolete data sets with outdated background traffic, we validate our scheme on two recent data sets of network traces captured during real network attacks.

II. ISSUES IN TRAFFIC ANALYSIS

A. Analyzing TCP Connections

Analyzing TCP connections is one of the most important issues to address in the case of SYN flooding attacks and portscans. To open a connection, a client sends an initial SYN segment. Upon its reception, the server allocates some resources in the backlog queue and replies with a SYN/ACK segment. Finally, the client returns an ACK segment (further called *Client ACK*) to complete the three-way handshake. Then, communication goes on until the client or the server sends a segment with the FIN flag set, a RST segment, or the connection times out. The potential for exploiting this behavior for denial of service lies in the early allocation of the server resources. During a TCP SYN flooding attack, the attacker generates multiple SYN requests without sending the Client ACK to complete the connection establishment. The requests can quickly exhaust the server memory so it cannot accept more incoming connection requests.

SYN scanning is fairly similar to TCP SYN flooding attacks: an attacking computer tries to identify vulnerable hosts by sending multiple TCP SYN segments. If a port is open, the server responds with a SYN-ACK segment, the port scanner completes the three-way handshake and immediately closes the connection with a RST segment.

Several authors proposed interesting detection methods that can operate in border routers to detect attacks and block them near their sources [1], [2], [3], [4], [5]. They take advantage

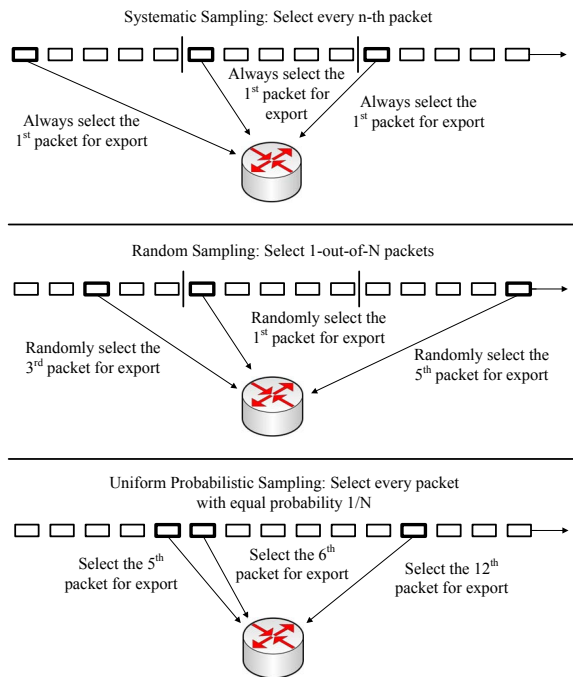


Figure 1. Principle of three sampling methods.

of the relationships between the TCP control segments: the appearance of a SYN segment implies further SYN/ACK, Client ACK, and FIN or RST segments.

However, if we want to apply sampling at border routers, considering only a small part of packets may degrade the detection capacity of all existing methods because of a fairly low probability of sampling essential control segments. As a consequence, the existing methods result in poor detection performance especially with respect to the false positive rate.

B. Sampling Techniques

We consider three basic and most commonly used count-based sampling techniques: systematic, random 1-out-of-N, and uniform probabilistic sampling (cf. Figure 1) proposed by the PSAMP IETF working group [6] and thoroughly investigated in the literature [8]. They present the advantage of simple implementation with low CPU and memory requirements.

Systematic sampling takes every N -th packet, whereas random 1-out-of-N sampling randomly chooses one packet in every bucket of size N . Finally, uniform probabilistic sampling analyzes every packet with the same small probability. Systematic sampling, also known as deterministic sampling, is usually used in current network devices, one example being the Cisco Netflow protocol [7]. Some previous work addressed the problem of how sampling techniques influence the anomaly detection process [9]. The authors examined various kinds of sampling methods with respect to volume and scanning anomalies and they concluded that packet sampling can introduce a fundamental bias by changing traffic features. Other authors considered the impact of sampling methods on various detection metrics by examining traces with TCP SYN flooding attacks [8]. Their results reveal that systematic sampling does

not perform well under low sampling rates when the detection process depends on specific packet characteristics like TCP flags.

III. PRINCIPLES OF THE DETECTION SCHEME

To overcome the limitations of methods that match pairs of TCP control segments, we propose a novel method not limited to the analysis of the three-way handshake or connection termination. To detect legitimate established connections, we take advantage of the fact that all segments originated from the server with the ACK flag set on and the SYN flag set off indicate a successfully established connection. In this case, the probability that the sampled packet contains one of multiple ACK segments coming from the server is greatly increased.

This approach decreases the false positive rate and does not influence the true positive rate, because in the case of SYN flooding attacks as well as portscans, there are almost no corresponding ACK segments coming from the server. Finally, it is impossible for the attacker to avoid detection by spoofing control segments.

The proposed scheme is placed in a border router that monitors packets generated in the controlled part of the network (e.g. an Intranet or an enterprise LAN) to confine the possible malicious activity close to the source of an attack. It is composed of three modules: the first one validates outgoing TCP segments, the second one processes corresponding control segments, while the third one changes the packet filter list if needed.

We combine the method with a rate limiting scheme. If the traffic rate is less than or equal to a predefined rate for a given IP address, it is allowed to pass the filter of outgoing traffic, whereas traffic that exceeds the rate is dropped or delayed.

We provide a detailed description of the proposed defense scheme below.

A. TCP History Check

For each sampled packet, we extract its source and destination IP addresses, and other information such as timestamps, sequence numbers, and ACK sequence numbers. When the router samples any outgoing TCP SYN segment, the module checks if a timeout has elapsed. Depending on the result, it either resets the source and destination IP lists and allows the segment to pass or it increases request counter R_{src} corresponding to the particular source IP address by a positive integer. If there are more unacknowledged SYN segments originating from the specific source IP address and $R_{src} > R_{src}^{max}$, then this module decides that the segments are parts of portscan activity and inserts the source IP address in the filter blacklist. Moreover, the module increases request counter R_{dst} by a positive integer for a particular destination IP address. If $R_{dst} > R_{dst}^{max}$, it means that there is an excessive number of connections to the destination address. Then, as this behavior may indicate host scan activity or a SYN flooding attack, the module updates the filter blacklist to block packets that follow.

B. TCP Validation Check

The goal of this module is to overcome the problem of losing some useful information because of sampling. It analyzes TCP control segments to determine whether the three-way handshake was successfully completed. Any incoming segment from the server side with the ACK flag set and SYN flag disabled indicates that the particular connection has been successfully established. In this case, the module decreases the R_{src} (R_{dst}) counter, because the connection becomes legitimate. Consequently, the requirement $R_{src} > R_{src}^{max}$ ($R_{dst} > R_{dst}^{max}$) might not be valid any more, so the module will eventually update the packet filter blacklist to permit further outgoing TCP requests from/to the specified IP address.

C. Filtering

This module applies all changes to the Access Control List (ACL) in the border router so that it will discard all malicious segments.

IV. EVALUATION RESULTS

To evaluate the method and compare it with the previous work, we have developed a prototype in the Matlab environment. We use the open source TracesPlay program [10] to read traces and to directly put the required data into Matlab.

A. Dataset Description

We have validated our scheme by means of trace-driven simulations on two data sets: the first traces were gathered on an operational university campus network at the National Technical University of Athens (NTUA) with an average traffic of 70-80 Mbits/sec and 20000 packets/sec. It contains a Distributed Denial of Service attack (TCP SYN flooding attack) captured on May 21, 2003 against a single host inside the NTUA campus. The second set has been collected on the link connecting an operational university campus network at the AGH University of Science and Technology in Cracow with a limit of 45 Mbits/s for incoming and 22 Mbits/s for outgoing traffic. In the evaluation presented in this paper, we have used set of three traces collected on March 24, 2010 containing host scans and port scans originated from the campus network as well as a packet trace without malicious activity. Moreover, traces contain rich background traffic including recent p2p applications as well as standard services like web, ftp, or mail.

B. Criteria of Detection Performance

We consider two meaningful metrics to evaluate the performance of detection methods: the true positive (TP) and false positive (FP) rates. Such rates are usually presented as the Receiver Operating Characteristics (ROC) curve by plotting the TP rate as a function of the FP rate. As attack detection is a Boolean action, the ROC curve is useful for network operators, because it indicates how to find the right tradeoff between the FP and TP rates. However, in our evaluation, we have separated both values and presented them as a function of the sampling rate, because the evaluation is also based on a trace without malicious activity.

C. Comparing with Existing Detection Schemes

To evaluate our scheme, we have compared it with other three representative detection schemes that leverage TCP relationships: SYN-SYN/ACK, SYN-FIN, and SYN-Client ACK. The key feature of schemes based on matching SYN-SYN/ACK and SYN-Client ACK pairs is the need of finding the corresponding SYN/ACK or Client ACK segment after the first SYN segment. The time interval between them is the RTT (Round Trip Time), usually less than 500ms for more than 90% of connections. Therefore, the methods have to inspect all control segments during at least this interval to conclude that the connection was successfully established. The detection methods based on matching SYN-FIN (or RST) pairs, simply waits for the corresponding FIN (or RST) segment.

D. Calibration Process

We had to face the problem of setting the right rate limiting thresholds, i.e., maximum values of the request counters corresponding to a regular traffic pattern. We have calibrated them for every examined trace in order to achieve a high TP rate with no FPs regardless of the detection method when we analyze all packets. The calibrated values reflect the relation between outgoing SYNs per destination and per source, and corresponding control packets (SYN/ACK, Client ACK, FIN). We have empirically found that the rate limiting thresholds expressed in packets are directly proportional to the sampling rate, which alleviates the problem of losing potentially useful data during the sampling process.

E. Influence of the Sampling Process on Different Detection Schemes

In our experiments, we have evaluated the influence of uniform probabilistic sampling on the proposed method and compared it with other three schemes. We have decided to choose this particular sampling method, because it is claimed to be more effective in the process of packet selection compared to systematic sampling, e.g. when there are periodic data patterns [8].

We have repeated all simulations to obtain 95% confidence intervals computed according to the bootstrap method [11].

As shown in Figure 2, all methods present approximately the same high TP and very low FP rate in case of TCP SYN flooding attacks. Similar results for all four methods are due to setting high rate thresholds corresponding to regular traffic for this particular trace. As we can observe, TP curves of all detecting schemes are similar until 0.08% when sampling process increases randomness in the results. As far as the FP rate is concerned, we can see that SYN-SYN/ACK, SYN-FIN, SYN-Client ACK methods deviate from our scheme, but differences are insignificant. For sampling rates lower than 0.08%, FPs are completely eliminated by the sampling process itself.

Our experiments considered host as well as network scans. The results however demonstrate similar behavior due to a similar attacking rate. Therefore, due to the space limitation in this paper, we only present the corresponding results for

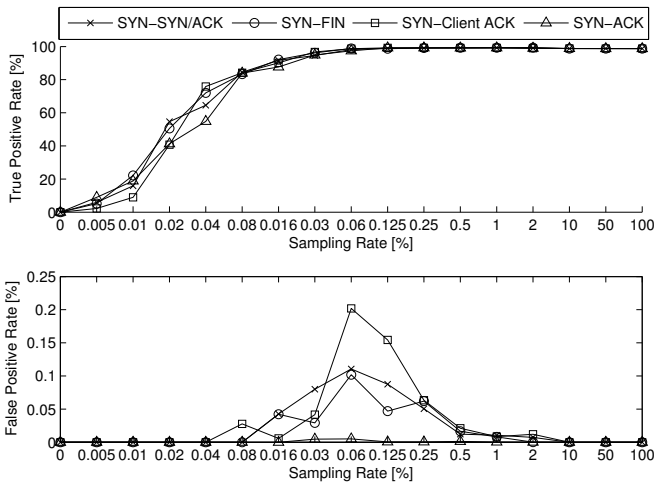


Figure 2. Comparison of the TP and TP rate for schemes based on analyzing SYN-SYN/ACK, SYN-FIN, SYN-Client ACK, and SYN-ACK segments in the case of TCP SYN flooding.

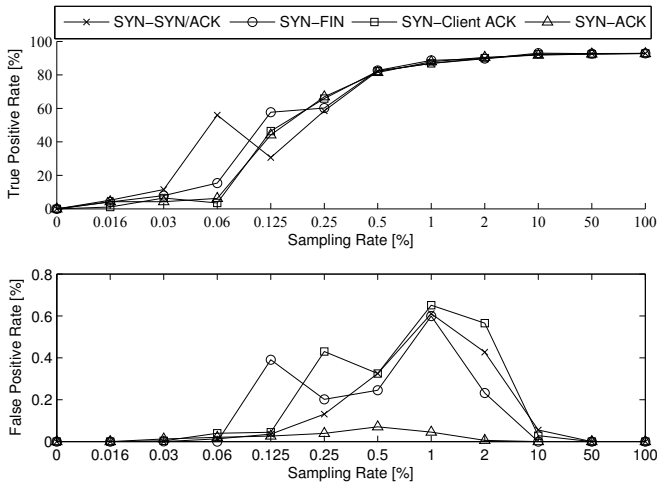


Figure 3. Comparison of the TP and FP rate for schemes based on analyzing SYN-SYN/ACK, SYN-FIN, SYN-Client ACK, and SYN-ACK segments for host scan activity.

host scan activity (cf. Figure 3). Again, all four detecting schemes show nearly the same behavior until 0.5% when the sampling process introduces randomness in the results. It is slightly higher in comparison with the case of the TCP SYN flooding attack due to a lower attacking ratio of host scans. Here, we can observe that SYN-SYN/ACK, SYN-FIN, SYN-Client ACK methods result in some FPs that do not however exceed 0.8%.

Figure 4 presents the corresponding results for the packet trace without malicious activity. Obviously, only the FP rate is presented and the results differ from the previous ones. In the following trace, rate limiting thresholds corresponding to legitimate traffic were established lower than in the previous cases. The results reveal that only our detection method based on sampling ACK segments coming from the server instead of looking for a single control segment to detect legitimate

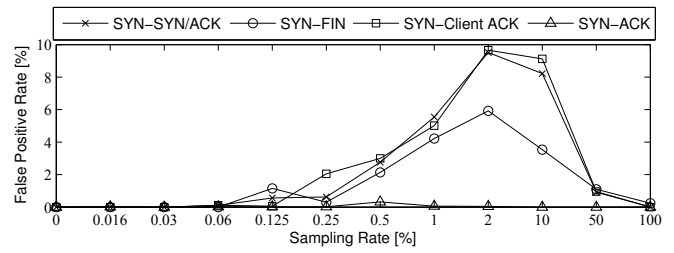


Figure 4. Comparison of the FP rate for schemes based on analyzing SYN-SYN/ACK, SYN-FIN, SYN-Client ACK, and SYN-ACK segments in for the trace without malicious activity.

established connections works well. When we reduce sampling rate to 2%, we can observe that detection schemes based on analyzing control segments (SYN/ACK, Client ACK and FIN) result in 6% up to 10% of false positives while our detection method still presents almost the same behavior as the unsampled one.

To conclude, the evaluation shows that our method can significantly reduce the false positive rate in comparison with other three methods and still accurately detect both high-volume DDoS attacks as well as low-volume scanning activity even if the method does not analyze all packets. Good performance comes from two features: first, at the instant of an attack, there are almost no ACK segments coming from the victim side so we obtain the high TP rate. Second, the method benefits from the novel approach to identify legitimate TCP connection request by sampling one of many ACK segments sent by the server to a legitimate client.

F. Impact of Sampling Techniques on the Proposed Scheme

Let us now examine the detection performance of our scheme under three representative sampling techniques described in Section II-B.

The results presented in Figure 5 correspond to a high-volume TCP SYN flood attack under the sampling process. We can observe that all three curves are similar regardless of the sampling method even at the sampling rate as low as 0.08%.

The results corresponding to low-volume portscan activity like host scans (cf. Figure 6) and network scans (cf. Figure 7) show similar behavior. Again, our detecting scheme shows nearly the same performance for all examined sampling methods until 0.125% in the case of host scans and 2% in the case of network scans when the sampling process itself introduces randomness in the results.

For the packet trace related to the traffic without network attacks, the FP rate does not exceed 0.3% in the worst case at the sampling rate 0.5% irrespective of the sampling method.

To sum up, the evaluation process indicates fairly similar behavior regardless of the sampling technique when we benefit from our novel detection scheme combined with the rate limiting scheme. This is due to the fact that this scheme is based on detecting segments with the ACK flag set. Such segments seem to appear evenly in the traffic thus resulting in

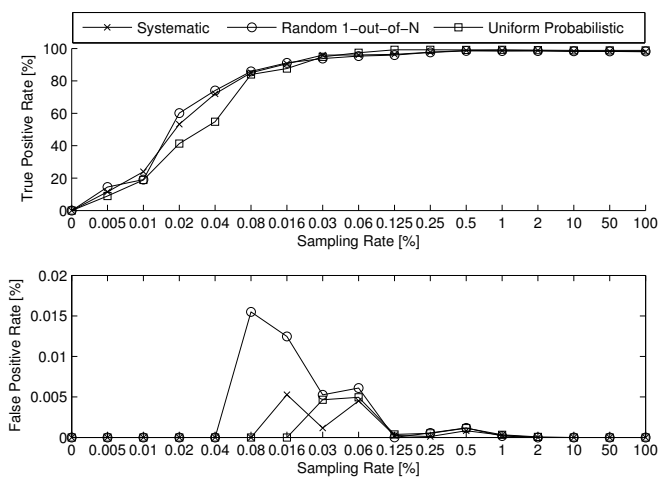


Figure 5. TP and FP rate—the influence of sampling methods on the proposed scheme based on analyzing SYN-ACK segments for the TCP SYN flooding.

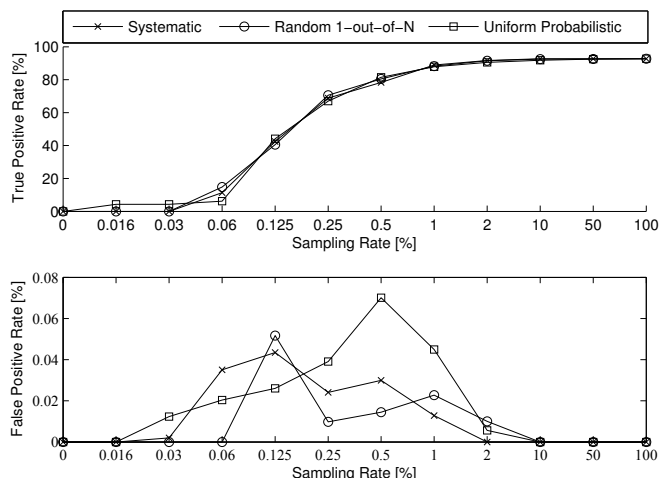


Figure 6. TP and FP rate—the influence of sampling methods on the proposed scheme based on analyzing SYN-ACK segments for host scan activity.

good performance under sampling even if systematic sampling is used.

V. CONCLUSION

We have proposed a novel scheme for detecting TCP SYN flooding attacks and portscans that offers good performance in the case of sampling. The scheme considers TCP connections as legitimate if it samples one of multiple ACK segments (with disabled SYN flag) coming from the server. This differs from existing methods based on pair matching of control segments SYN/ACK, FIN (RST) or Client ACK etc. Our trace-based simulations show that unlike other techniques, the proposed method significantly decreases the false positive rate under a sampling process. Moreover, the results reveal that our method alleviates the problem of losing some information

when systematic sampling is used. The effectiveness of the

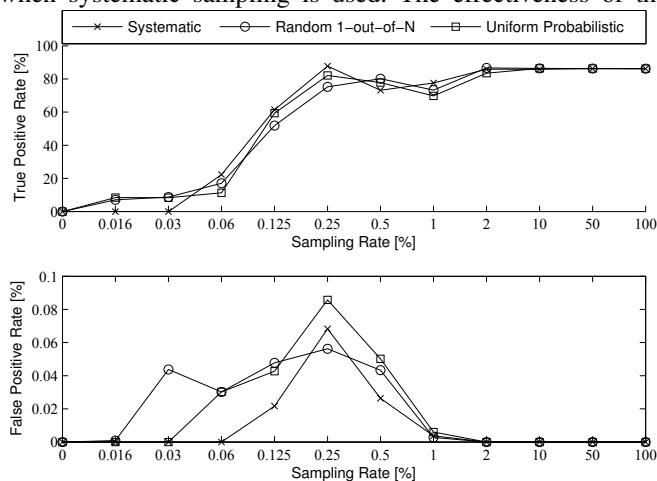


Figure 7. TP and FP rate—the influence of sampling methods on the proposed scheme based on analyzing SYN-ACK segments for network scan activity.

presented method only relies on the sampling rate and not on the type of a sampling method.

ACKNOWLEDGMENTS

We would like to thank Symeon Papavassiliou and Georgios Androulidakis for traces from the university campus network at the National Technical University of Athens. This work was partially supported by the EC FP7 project INDECT under contract 218086 and by COST action TMA IC0703 (647/N-COST/2010/0).

REFERENCES

- [1] H. Wang, D. Zhang, and K. G. Shin, "Detecting SYN Flooding Attacks," in *Proceedings of IEEE INFOCOM'2002*, June 2002, pp. 1530–1539.
- [2] H. Wang, D. Zhang, and K. Shin, "Change-Point Monitoring for the Detection of DoS Attacks," in *IEEE Transactions on Dependable and Secure Computing*. IEEE Computer Society Press, October 2004.
- [3] W. Chen and D. Y. Yeung, "Defending against TCP SYN flooding attacks under different types of IP spoofing," in *Fifth International Conference on Networking (ICN)*, 2006.
- [4] R. Kompella, S. Singh, and G. Varghese, "On scalable attack detection in the network," in *IEEE/ACM Transactions on Networking*, vol. 15, no. 1, San Francisco, CA, USA, February 2007, pp. 14–25.
- [5] C. Sun, C. Hu, Y. Zhou, X. Xiao, and B. Liu, "A More Accurate Scheme to Detect SYN Flood Attacks," in *IEEE INFOCOM Workshops*, April 2009, pp. 1–2.
- [6] "Packet Sampling (PSAMP) IETF Working Group Charter," <http://www.ietf.org/html.charters/psamp-charter.html>.
- [7] "NetFlow Services Solutions Guide," <http://www.cisco.com/>.
- [8] G. Androulidakis, V. Chatzigiannakis, S. Papavassiliou, M. Grammatikou, and V. Maglaris, "Understanding and Evaluating the Impact of Sampling on Anomaly Detection Techniques," in *IEEE MILCOM Military Communications Conference*, October 2006, pp. 1–7.
- [9] J. Mai, C.-N. Chuah, A. Sridharan, T. Ye, and H. Zang, "Is Sampled Data Sufficient for Anomaly Detection?" in *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, October 2006.
- [10] "TracesPlay," <http://tracesplay.sourceforge.net/>.
- [11] B. Efron and R. Tibshirani, "Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy," in *Statistical Science*, vol. 1, no. 1, February 1986, pp. 54–75.