Baseline Profile Stability for Network Anomaly Detection*

Yoohwan Kim[1], Ju-Yeon Jo[2], and Kyunghee Kim Suh[3]
(Corresponding author: Yoohwan Kim)

School of Computer Science, University of Nevada, 4505 Maryland Parkway, Las Vegas, NV 89154-4019[1]
School of Informatics, University of Nevada, 4505 Maryland Parkway, Las Vegas, Las Vegas, NV 89154-4005[2]
American Institutes for Research, Sacramento, CA 95833[3]

(Selected paper from ITNG 2006)

Abstract

Network attacks are commonplace in the Internet. One of the defense mechanisms against the network attacks is using a baseline profile established during normal operation to detect the traffic that deviates from the baseline profile. However, this approach works only if there is a stable base profile representing the legitimate network traffic. Although there has been some preliminary research, the details of profiling, such as the profile format, its size and the traffic stability by site or time, have not been widely available. In this study, we analyze actual traffic traces from two Internet traffic archives and verify the traffic stability by various aspects. The analysis shows that there are significant differences in the traffic patterns among different sites. In addition, there are some differences between different time of day or different days, even within a site, suggesting that different profiles are needed for different times. The result of this study can be used practically to anomaly-based IDS for determining the stability of the traffic for a particular site, and the number of required traffic profiles based on the traffic patterns.

Keywords: Denial-of-Service Attack, Internet traffic profile, network security

1 Introduction

Network attacks are commonplace in the Internet nowadays. Especially the Distributed Denial-of-Service (DDoS) became a great threat [16]. In DDoS attack, a large number of attack packets are sent to the victim network to exhaust the victim network resources. One of the defense mechanisms against such attacks is filtering the packets that deviate from the normal traffic [1, 4, 13]. The baseline traffic profile is collected during the normal operation. This method has been used for anomaly-based IDS [12, 13] and DDoS attack filtering schemes [5, 6, 8, 9, 10, 15]. The Intrusion Prevention Systems (IPS) based on the traffic anomaly detection are called rate-based IPSs and some commercial devices have been developed [2, 3, 18, 19].

One may argue that it is relatively straightforward for a sophisticated attacker to learn the approximated distribution of some attributes, e.g. protocol-type, TCP-flag pattern and packet-size, based on publicly available data on Internet traffic characteristics. Thus the attacker may be able to generate the traffic pattern accordingly to circumvent the baseline profile-based detection schemes. However, distributions of the other attributes, such as TTL, source IP-prefixes, or server-port distribution, are expected to be site-dependent and thus more difficult for an outside attacker to learn such information. For instance, it is quite difficult for an outsider to determine the joint-distribution of source-IP-prefix and the TTL value for a given site. As long as there exists profiling information which is known only to the site/network-operator but not to the attacker, our scheme can use this information as the basis to differentiate among attacking and legitimate packets.

The greatest challenge in the baseline profile-based schemes is the validity of the traffic stability. It has been known that there is a distinct traffic pattern in terms of packet attribute value distribution for a particular time and/or day for a given subnet [7, 11]. In general, the nominal traffic profile is believed to be a function of time which exhibits periodic, time-of-day, day-of-the-week variations as well as long-term trend changes. However, the baseline profiling mechanisms, the stability of periodic traffic patterns and per-site traffic pattern differences have not been adequately studied yet. In this research, we study whether there are unique traffic characteristics for different sites and different times using recent packet trace data.

*A preliminary version of this work appeared in proceedings of international conference on Information Technology: New Generations (ITNG 2006).

Table 1: An example of a base profile

TTL Value	Period 1	Period 2	Period 3	Period 4	Profile
1	0.5%	0.8%	**1.1%**	0.3%	1.1%
2	0.7%	0.5%	0.6%	**0.8%**	0.8%
2	3.0%	**3.5%**	2.4%	2.9%	3.5%
...	...	...	...	...	...
255	**1.3%**	1.2%	0.9%	1.2%	1.3%

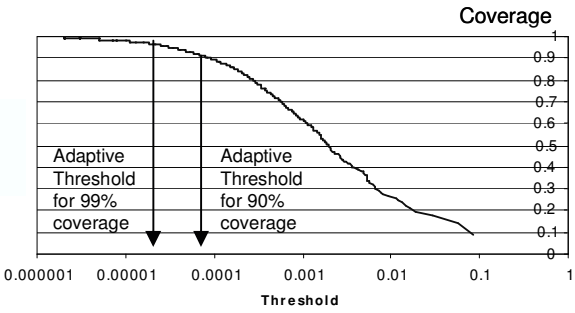

Figure 1: Adaptive thresholding

2 Collecting Baseline Profile

For different application or products, different set of base profiles can be collected and the profile format may also vary. In this research, we consider a set of marginal and joint distributions of various packet attributes. Candidate packet attributes considered to be useful for traffic profiling include: marginal distributions of the fractions of packets having various (1) IP protocol-type values, (2) packet size, (3) server port numbers, i.e., the smaller of the source port number and the destination port number, (4) source/ destination IP prefixes, (5) Time-to-Live (TTL) values, (6) IP/TCP header lengths, and (7) TCP flag patterns. It is worthwhile to employ the joint distribution of the fraction of packets having various combinations, such as (8) packet-size and protocol-type, (9) server port number and protocol-type, as well as (10) source IP prefix and TTL values, etc. Other useful candidates are the fractions of packets which (11) use IP fragmentation and (12) bear incorrect IP/TCP/UDP checksums.

During the baseline profiling period, the number of packets with each attribute value is counted and the corresponding ratio is calculated. However, if the profile is created only once from the entire traffic, temporally localized traffic characteristics may be misrepresented. To avoid this situation, the ratios of attribute values are measured over multiple periods, and one value representing all the periods is selected. Specifically, to accommodate an occasional surge of particular attribute values in legitimate traffic, the highest ratio among the periodic ratios is selected. Table 1 illustrates this process with an example of TTL values. The boldface values are the highest ratios observed among the periodic values, which are then stored in the profile.

Due to the number of attributes to be incorporated in the profile and the large number of possible values of each attribute, especially for the joint attributes, an efficient data structure is required to implement the profile. Towards this end we propose to use *iceberg-style* profiles where only the most frequently occurring items are stored. Two approaches are possible for selecting iceberg items, i.e., by static threshold and by adaptive threshold. In the static threshold approach, the profile only includes those entries which appear more frequently than a preset percentage threshold, say $x\%$. For entries which are absent from the iceberg-style profiles, we use the upper bound, *i.e.,* $x\%$ as their relative frequency. In the adaptive threshold approach, as described in Figure 1, the most frequently appearing attribute values that constitute a preset coverage of the traffic, *e.g., 95%*, are selected first. The corresponding iceberg threshold value $x\%$ is determined separately for each marginal/ joint distribution so that $y\%$ of the overall entries observed in the baseline trace are covered by the iceberg histograms. The adaptive threshold is also used for the threshold of the absent items.

Table 2: Profile storage requirements for different iceberg selection methods

Threshold Type	Storage Requirements (Kbytes)	Relative Storage Requirements (Kbytes)
Static threshold	13.6	1.0
Adaptive 90% threshold	76.0	5.6
Adaptive 95% threshold	127.8	9.4
Adaptive 99% threshold	288.3	21.2

With such iceberg-style profiles, the nominal profile can be kept to a manageable size. Typical storage requirements for storing six single attributes and two joint attributes are summarized in Table 2 for different threshold methods. For the static threshold, we used 0.01, 0.001 and 0.0001 respectively for single attribute, two-dimensional and three-dimensional joint attributes.

3 Traffic Profile Stability

3.1 Stability within a Trace

To validate our claim of the relatively "invariant" nature of the distribution of the above packet attributes, we have conducted extensive statistical analysis on real-life Internet traces collected from the traffic archive of the WIDE-project [14]. Figure 2 (a)-(d) show the time variation of the distribution of various packet attributes values observed from a moderately loaded wide area network link. For each attribute, the relative frequency of

its values is computed every 10 minutes for the period between May 10, 1999 8:00pm and May 11, 2:00pm for a total of 108 non-overlapping periods. Figure 2 (a) shows the time-variation of the distribution of TTL values. In particular, the ends of the error-bar correspond to the maximum and minimum fraction observed for the given TTL value over the aforementioned 18-hour interval and the black-dot represents the average. The corresponding time-varying distributions for protocol-type, packet-size, TCP-flag pattern, server port number (smaller of the source port number and the destination port number) and 16-bit source IP prefix are shown in Figure 2 (b)-(f) respectively.

Notice from Figure 2 that while the ratio of an attribute value does vary over the 18-hour period, it varies usually within a relatively small range. As the ratio is concentrated in a smaller range, the more stable the attribute value is. To measure the stability within one profile, we use the following metric.

$$S_P = \frac{\Sigma_{\mu_i > f} \sigma_i / \mu_i}{N}$$

where σ_i is the standard deviation and μ_i is the average for each attribute value over t periods (i is for attribute value), and N is number of attribute values where $\mu_i > f$ (f is the threshold value for choosing the iceberg values)

For example, if S_P is 0.5, most of the attribute values vary within 50% of the average (i.e. $\sigma = 0.5 * \mu$). The lower $S_P$is, the more stable the attribute is. Here are the actual values of S_P with TTL.

- When covering 100%, $S_P = 1.39$;
- For the icebergs when $f = 0.001, S_P = 0.58$;
- For the icebergs when $f = 0.01, S_P = 0.42$.

In other words, for the iceberg items that occupy more than 0.01 of the total packets, their fraction varies only within 42% of the average value.

3.2 Stability among Different Times and Sites

To further verify traffic profile stability with more recent data, we conducted a stability analysis with the packet trace data available from NLANR packet trace archives [17]. All trace data were collected for 90 seconds from 17 sites within the United States with the link speed ranging from OC-3 to OC-48. We randomly selected the four sites in Table 3 and total of 49 trace files were downloaded for analysis. Table 4 shows general statistics for 10 selected traces.

A quick examination of Table 4 reveals that each site has a distinct traffic composition. Especially, we observed that the traffic in AIX is mostly GRE rather than TCP or UDP. For each trace, a profile was created using a 99% adaptive coverage method over a series of 10-second windows. We employed only one joint attribute composed

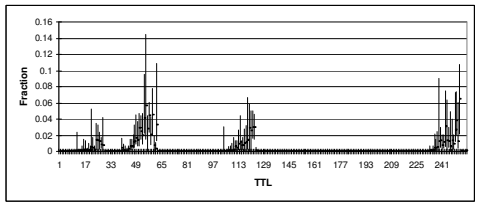

(a) Time variation of TTL value distribution

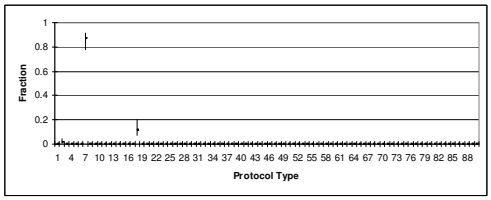

(b) Time variation of protocol type distribution

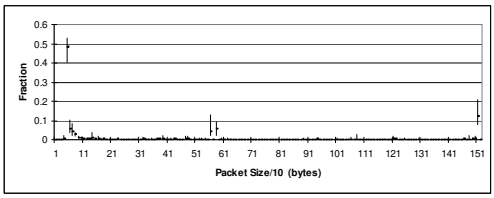

(c) Time variation of packet-size distribution

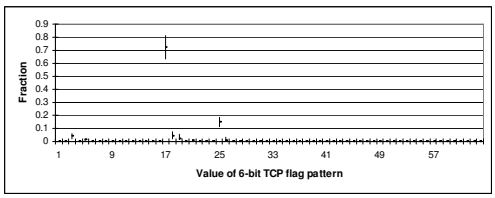

(d) Time variation of 6-bit TCP flag pattern distribution (e.g. SYN = 000010 = 2 ; ACK = 010000 = 16)

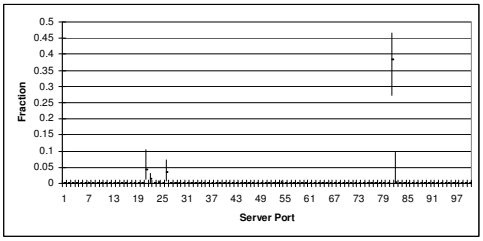

(e) Server port distribution

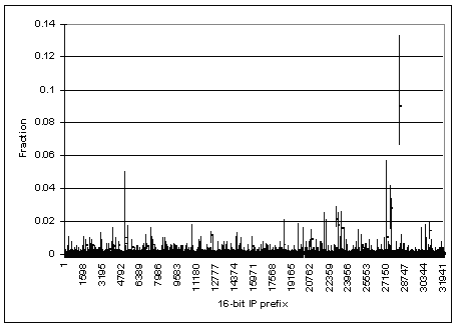

(f) Source IP prefix distribution

Figure 2: Time variation of packet attribute values distribution

Table 4: Statistics of some downloaded traces

File Name	Date	Time (24H)	File Size (MB)	Total # packets (1000)	Ave. # packets (1000)	Bandwidth (Mbps)	Traffic composition		
							TCP	UDP	ICMP
MEM-1127750202-1.tsh	9/26/05 Mon	9:12	8.4	197.352	2.381	11.889	0.926	0.060	0.008
MEM-1127791306-1.tsh	9/26/05 Mon	21:10	7.3	171.343	2.066	13.457	0.899	0.086	0.009
MEM-1124810381-1.tsh	8/23/05 Tue	8:28	9.9	231.021	2.787	14.246	0.939	0.053	0.005
MEM-1125415231-1.tsh	8/30/05 Tue	8:35	6.5	151.524	1.827	7.609	0.834	0.147	0.012
AIX-1127750202-1.tsh	9/26/05 Mon	9:31	0.5	12.116	0.149	0.417	0.000	0.000	0.000
AIX-1127835595-1.tsh	9/27/05 Tue	8:42	0.5	12.397	0.158	0.267	0.002	0.000	0.001
AMP-1127747110-1.tsh	9/26/05 Mon	8:13	31.0	739.100	9.157	50.311	0.850	0.140	0.003
AMP-1127836180-1.tsh	9/27/05 Tue	8:58	21.0	496.516	6.140	36.308	0.922	0.075	0.002
PSC-1127747111-1.tsh	9/26/05 Mon	8:38	163.0	3799.190	47.051	276.604	0.849	0.132	0.018
PSC-1127836180-1.tsh	9/27/05 Tue	9:37	136.0	3171.690	39.570	233.482	0.864	0.113	0.022

Table 3: Trace data sites from NLANR

Site	Location	Link Speed	# of traces analyzed
AIX	NASA Ames to MAE-West	OC12 (655 Mbps)	28
AMP	AMPATH, Miami, Florida	OC12 (655 Mbps)	7
MEM	University of Memphis	OC3 (155 Mbps)	7
PSC	Pittsburgh Supercomputing Center	OC48 (2.5 Gbps)	7

of Protocol type, server port, packet size because joint attributes are more unique per site than single attributes.

For an objective comparison of two profiles, we define the stability metric S_C as follows:

$$S_C = C \times D.$$

C indicates how many items are common to both profiles:

$$C = \frac{(\text{\# of common items in both profiles} = n}{(\text{\# of total items in both profiles})}.$$

D indicates how closely these common items are related. For example, the two profiles may have 3% vs. 1.7%, or 3.2% vs. 2.9% for a given attribute value. Obviously the latter shows a stronger resemblance. For an item i that is in both profiles, the comparison ratio is defined as the smaller ($R_{small}(i)$) of two values divided by the larger ($R_{large}(i)$) of two values. D is defined as the average of the comparison ratios. When the comparison ratio is too small, e.g., below 0.01, we consider it 0.

$$D = \frac{\Sigma_{i=1}^{n} \frac{R_{small}(i)}{R_{large}(i)}}{n}.$$

S_C varies between 0 and 1, from no stability to perfect stability. When two profiles are exactly the same ($C = 1$ and $D = 1$), S_C becomes 1. When there are no common items ($C = 0$) or the comparison ratios are zero for all of the common items ($D = 0$), S_C becomes zero. We believe that this S_C metric is more accurate than the usual correlation coefficient due to many absent attribute values in the iceberg style profile. For a better comparison of the stability at low S_C values, we define S_L as a log version of S_C:

$$S_L = \log_{10} 10C \times \log_{10} 10D, S_L = 0, \text{ if } C \leq 0.1 \text{ or } D \leq 0.1.$$

For stability analysis, one reference profile was compared with other profiles and the S_L value is calculated. The selected reference profile is from the MEM trace of Tuesday, September 27, 2005, 8:49 a.m. We investigate the following questions.

- Are the profiles similar for the 10-second windows within a 90-second trace?
- Are the profiles similar between mornings and evenings at the same site?
- Are the profiles similar over multiple weeks at the same time of a specific day? (e.g., 8:00 p.m. every Tuesday)
- Are the profiles different at different sites at the same date and time?

The analysis results are shown in Figure 3. In Figure 3(a), the profile of the first 10-second window is compared with other 10-second window profiles. It indicates that there is strong correlation among the 10-second windows, thus validating the S_L metric. Figure 3(b) compares seven profiles from September 26, 2005 to October 2, 2005 at approximately 9:00 a.m. (morning) and 8:00 p.m. (evening) each day. It indicates that there is moderate correlation among the daily profiles, although weaker than within the same trace. It should be noted that when the Tuesday profile is compared with itself, the $S_L = 1$. It also shows that there is a higher correlation for the same time of day (approximately 9:00 a.m.) than at a different

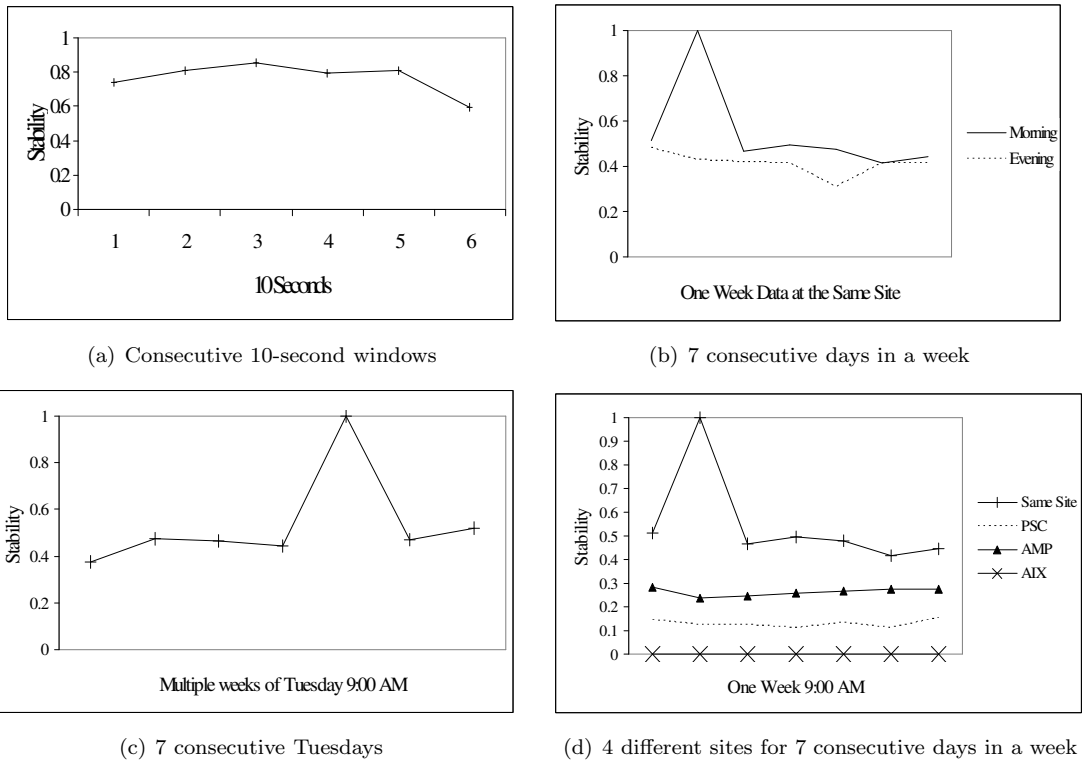

(a) Consecutive 10-second windows

(b) 7 consecutive days in a week

(c) 7 consecutive Tuesdays

(d) 4 different sites for 7 consecutive days in a week

Figure 3: Stability Comparison

time of day (approximately 8:00 p.m.). Figure 3(c) compares the profiles for seven Tuesdays at the same time of day (approximately 9:00AM) from August 23, 2005 to October 11, 2005. Although it spans seven weeks, it still shows a similar correlation to the short-term profiles of 9:00 a.m. as in Figure 3(a). These seven Tuesday morning profiles are slightly closer than the evening profiles in Figure 3(b). In all cases within the same site, the S_L is generally above 0.4. However in Figure 3(d), the S_L is much lower when compared with other sites, showing much weaker correlation.

In summary, we observe that traffic profiles are most similar among on the same day at the same time, even over multiple weeks. A traffic profile is still very similar for a different time or day within a site, although stability is slightly lower than the same time of day. On the other hand, there are considerable differences among different sites, so it is necessary to keep separate profiles for each site. By tracking S, we can determine the stability of the profiles for different times or days and can also decide how many profiles are needed. Unless there is a significant difference between profiles, we may use one uniform profile to minimize the maintenance effort.

4 Discussions

One challenging issue of the anomaly-based IDS is the need for a clean baseline profile as in other profile-based systems. This may not be easy today because various attack traffic is already prevalent in the Internet and a quiet attack-free period may be hard to find. Especially the Distributed denial-of-service attack is very common, and as a result, the constructed baseline profile may be biased by the DDoS traffic. A cleaner profile can be made one of two ways. First, the packet trace data can be analyzed to identify legitimate flows that show proper two-way communication behavior. The packets from the legitimate flows are used for constructing the profile. Secondly, we can use a packet filtering algorithm [8] with a generic profile to create a cleaner packet trace, and use the new trace to create the cleaner base profile. The generic profile reflects overall Internet traffic characteristics, e.g., TCP vs. UDP ratio, common packet size, common TCP flags, etc. Our preliminary research shows that this two-step profiling is very effective to filter generic attacks.

5 Conclusions

One of the defense mechanisms against the network attacks is using a baseline profile established during normal network operation, and detecting the traffic that deviates from the baseline profile. This is a common approach in many Intrusion Detection Systems and some Distributed Denial-of-Service (DDoS) attack defense mechanisms. However, this approach works only if there is a stable baseline profile representing the legitimate network traffic. In this study, we analyzed actual traffic traces from two Internet traffic archives and verified the traffic

stability in several aspects. Our analysis shows that there is stable and distinct traffic patterns for different sites, and different time/day for a particular site. When there is a significant difference per day or time, using establishing multiple profiles is recommended. This research guides how to check whether a particular site has meaningful traffic stability, how to measure the stability within a site, and how to decide the number of required traffic profiles. This research uses sample traces collected at the NLANR archive in September 2005. We plan to analyze more packet trace data to confirm the findings in this paper for wider range of sites and to investigate long-term traffic pattern stability.

References

[1] J. D. Brutlag, "Aberrant behavior detection in Time Series for Network Monitoring," in *the 14th USENIX Conference*, pp. 139-146, Dec. 2000.

[2] *Captus Networks.* (http://www.captusnetworks.com)

[3] *DeepNines Technologies.* (http://www.deepnines.com)

[4] C. Estan, S. Savage, and G. Varghese, "Automatically inferring patterns of resource consumption in network traffic," in *Proceedings of 2003 ACM SIGCOMM*, pp 137-148, 2003.

[5] L. Feinstein, D. Schnackenberg, R. Balupari, and D. Kindred, "Statistical approaches to DDoS attack detection and response," in *Proceedings DARPA Information Survivability Conference and Exposition*, vol. 1, pp. 303-314, Washington, DC, Apr. 2003.

[6] S. Jin, and D. S. Yeung, "A Covariance analysis model for DDoS attack detection," in *Proceedings of 2004 IEEE ICC*, pp. 1882-1886, 2004.

[7] J. Jung, B. Krishnamurthy, and M. Rabinovich, "Flash crowds and denial of service attacks: Characterization and implications for CDNs and web sites," in *Proceedings of the International World Wide Web Conference*, pp. 252-262, May 2002.

[8] Y. Kim, W. C. Lau, M. C. Chuah, and H. J. Chao, "PacketScore: statistics-based overload control against distributed denial-of-service attacks," in *Proceedings of IEEE INFOCOM*, pp. 2594-2604, Mar. 2004.

[9] J. Li, and C. Manikopoulos, "Early statistical anomaly intrusion detection of DOS attacks using MIB traffic parameters," in *Proceedings of 2003 IEEE Workshop on Information Assurance*, pp. 53-59, June 2003.

[10] Q. Li, E. C. Chang, and M. C. Chan, "On the effectiveness of DDoS attacks on statistical filtering," in *Proceedings of IEEE INFOCOM*, pp. 1373-1383, 2005.

[11] D. Liu and F. Huebner, "Application profiling of IP traffic," in *Proceedings of the 27th Annual IEEE Conference on Local Computer Networks (LCN)*, pp. 220-229, 2002.

[12] M. Mahoney and P. K. Chan, "Learning nonstationary models of normal network traffic for detecting novel attacks," in *Proceedings of ACM 2002 SIGKDD*, pp 376-385, 2002.

[13] D. Marchette, "A Statistical method for profiling network traffic," in *the 1st USENIX Workshop on Intrusion Detection and Network Monitoring*, pp. 119-128, Apr. 1999.

[14] *MAWI Working Group Traffic Archive.* (http://tracer.csl.sony.co.jp/mawi)

[15] J. Mirkovic, G. Prier, and P. Reiher, "Attacking DDoS at the source," in *Proceedings of 10th IEEE International Conference on Network Protocols*, pp. 312-321, Nov. 2002.

[16] D. Moore, G. M. Voelker, and S. Savage, "Inferring internet denial-of-service activity," in *Proceedings of 10th USENIX Security Symposium*, pp. 9-22, Aug. 2001.

[17] *NLANR PMA Internet Packet Trace Data Archive.* (http://pma.nlanr.net/Traces)

[18] *Top Layer Networks.* (http://www.toplayer.com)

[19] *Vsecure Technologies.* (http://www.v-secure.com)

Yoohwan Kim is an Assistant Professor of Computer Science at University of Nevada, Las Vegas (UNLV). He received his Master's and Ph.D. degrees in Computer Engineering from Case Western Reserve University, Cleveland, Ohio, in 1994 and 2004 respectively, and his Bachelor degree in Economics from Seoul National University, Korea in 1989. Before joining UNLV in 2004, he worked in software and communication networking industry for several years. He was a Member of Technical Staff at Lucent Technologies, Whippany, New Jersey, developing software for wireless networking equipment between 1997 and 1999. In 2000, he co-founded and managed a New Jersey-based software company that developed technologies for delivering and customizing video advertising over the Internet. His current research interests include network security, Internet traffic analysis, software architecture, and real-time embedded software design

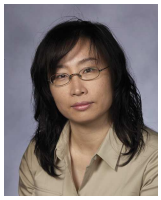

Ju-Yeon Jo received PhD degree in computer science from Case Western Reserve University, Cleveland, Ohio. She is an assistant professor of school of informatics at the University of Nevada, Las Vegas where she joined in August 2006. From 2003 to 2006, she was an assistant professor of computer science department at California State University, Sacramento. Prior to that she spent several years in communication networking and software industry. She was a member of technical staff at Lucent Technologies, Bell Labs, in Homdel, New Jersey, and a software engineer

at Coree Networks, a New Jersey based start-up company. Her current research interests include information security, network security, networking protocol design and performance analysis, and Internet traffic characterization.

Kyunghee Suh is working at American Institues for Research as a Research Scientist.She is working on all aspects of psychometric work on the development of large scale assessment such as High School Exit Exam item analysis, item calibrations, equating design, quality control of data, item banking and score reports to meet the technical standards of the assessment. Before joining AIR, Dr. Suh worked on various researches in education field-focus on constructing of statistical research design, large-scale statistical data analysis, item analysis, and item calibrations-as statistical consultant at University of Northern Colorado. Dr. Suh is familiar with statistical analysis software, educational research design, complex data analysis and data manipulation in psychometric aspect, and scoring of performance assessment. Dr. Suh specializes in statistical analysis and educational measurement. Dr. Suh received her Ph.D in Applied Statistics and Research Methods, with a Modified Rater Agreement Index in educational measurement, from University of Northern Colorado