A Passive Approach to Wireless NIC Identification

Cherita Corbett
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332
Email: cherita@gatech.edu

Raheem Beyah
Department of Computer Science
Georgia State University
Atlanta, Georgia 30303
Email: rbeyah@cs.gsu.edu

John Copeland
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332
Email: jcopeland@ece.gatech.edu

Abstract—IEEE 802.11 wireless networks are plagued with problems of unauthorized access. Left undetected, unauthorized access is the precursor to additional mischief. Current approaches to detecting intruders are invasive or can be evaded by stealthy attackers. We propose the use of spectral analysis to identify a type of wireless network interface card. This mechanism can be applied to support the detection of unauthorized systems that use wireless network interface cards that are different from that of a legitimate system. The approach is passive and works in the presence of encrypted traffic.

Keywords-network security; wireless security; host identification; network management

I. INTRODUCTION

The short comings surrounding the security of 802.11 wireless local area networks (LANs) have been well documented [1, 2]. Expectedly, there are automated tools [3] that exploit these flaws to passively and actively attack wireless LANs to gain unauthorized access. Unauthorized access could manifest itself as medium access control (MAC) spoofing, as a man-in-the-middle attack, as a session-hijack or as fraudulent use of the intended services offered by the wireless network (i.e., Internet, corporate services, etc.).

Reacting to the security weakness, the IEEE 802.11 standards committee sought to provide additional security features with the 802.11i specification. To be effective, 802.11i requires new hardware and must be commonly applied to all systems on the wireless network. In general, prevention mechanisms are only effective on systems that are owned, managed, and controlled by the network administrator. This method of defense cannot be enforced on rogue systems for which an administrator has no knowledge or control. Also, users may not apply any security measures.

Despite new security enhancements, the risk of intrusion is still a legitimate concern because security policies may be circumvented, cost prohibitive, or not practiced at all. As a result, a method of detecting unauthorized access is necessary. If a breach goes undetected, sensitive information can be stolen, network resources abused, or more sophisticated attacks can be launched targeting legitimate wireless stations or the Internet. Detecting unauthorized access affords an opportunity to respond to the intrusion and curtail the potential damage to preserve the privacy and integrity of the network.

Current intrusion detection strategies [4, 5] seek to address this problem by monitoring the wireless LAN for a sequence of events that exhibit odd behavior or violate security policy, such as jumps in sequence numbers or a station operating on a prohibited channel. This approach can be evaded by stealthy intrusions that do not use brute force to gain access or do not behave out of profile. For example, a hacker may have obtained a user name and password from an authorized user via phishing techniques that lure an unsuspecting victim to divulge personal information. In such a case, the attacker appears to have legitimate access and does not exhibit alarming behavior because it had the proper credentials.

To address the problem of unauthorized access, we propose a technique to identify the type of network interface card (NIC) used by a wireless host. Establishing an identity for valid types of NICs can help detect intruding systems that have a different type of wireless NIC. Though an attacker can use the same NIC, this technique adds another barrier for the attacker[1]. We show that differences in the composition of a wireless NIC influence data transmission patterns in a manner that is observable through traffic analysis. We extract subtle differences in the temporal behavior of a wireless stream to produce a spectral profile as the identification of the NIC.

The remainder of the paper is organized as follows. In section II, we present existing approaches to device identification. Sections III and IV discuss rate switching as a viable attribute for identifying different types of wireless NICs. In section V, we present our approach to passive NIC identification. We conduct an experimental evaluation and present the results in Section VI. In Section VII, we give the conclusion and discuss future work.

II. RELATED WORK

WiMetrics [6] is a commercially available monitoring and intrusion protection system. It implements an identity profiling process that can preauthorize a user through a registration process or authorize on the fly by probing the wireless device to derive an identity profile based on the response. Probing wireless stations is intrusive and as the number of clients increases, the already constrained network becomes burdened with additional traffic imposed by the system. This approach has other drawbacks including the administrative overhead of the preauthorization process. In addition, a hacker could elude the system by crafting responses to the probe request to

[1] As with any security system, if an attacker knows the technique/algorithm, the system can be evaded.

impersonate the identity of a legitimate user, reducing the effectiveness of this scheme.

IPass Inc. developed DeviceID [7], a software-based authentication technology. DeviceID creates a digital fingerprint using random segments of serial numbers for different hardware components within the device. It consists of two components, server and client software. The server encrypts and inventories the digital fingerprint in a database. The client resides on all end-point devices to establish secure sockets layer (SSL) connections for secure transmission of the device's fingerprint required for hardware authentication. This approach is intrusive and suffers from administrative overhead involved in distributing the client software and updating the database every time a hardware component changes in the device. Further, this approach generates traffic, placing additional strain on the wireless link.

Radio frequency (RF) fingerprinting captures the unique characteristics of the RF energy of a transceiver. When a radio transmitter is placed in transmit mode, a transient is generated by the frequency synthesizer whose function it is to generate the carrier frequency used for transmission. It has been determined that the turn-on transients generated are distinct enough that positive identification of the transmitter is possible. This technology was originally used in the cellular industry to identify fraudulent clones [8]. Researchers at Carleton University [9] have extended this approach to control access amongst BlueTooth wireless devices with future plans of including 802.11 transceivers. To implement this technology in a wireless LAN, special equipment for processing RF signals would be required at each access point. The cost of new equipment can become prohibitive especially for large networks with many access points. This was not of significant concern to the cellular industry because each tower services thousands of subscribers dissipating the cost of the equipment.

Kohno et al. [10] demonstrate a method for remotely fingerprinting a physical device by exploiting the implementation of the TCP protocol stack. When the TCP timestamp option is enabled, outgoing TCP packets reveal information about the sender's internal clock. The authors' technique exploits microscopic deviations in the clock skews to derive a clock cycle pattern as the identity for a device. For machines that do not enable the timestamp option by default, such as those running Windows 2000 and Windows XP, this approach becomes an active one. In such a case, the active fingerprinting technique initiates a connection and tricks the fingerprintee into using the timestamp option. The active approach must violate the TCP specification in order to execute the trick. The drawback to the active technique is that it is detectable to the fingerprinted device. Furthermore, the entire approach only applies to TCP traffic and can be evaded by spoofing the TCP timestamp field or setting it to an arbitrary value.

III. RATE SWITCHING

The objective of our work is to identify different types of wireless NICs. A wireless NIC is installed into a host to carry out the physical transmission of a packet over radio frequency based on the 802.11 standard. The 802.11 standard specifies the services a wireless NIC must provide to carry out data transmission. The standard does not, however, dictate how several of the services are to be implemented. For example, it is left up to the vendor to establish a technique for adjusting transmission rates, reserving the link, polling for packets to conserve power, and probing the network for connectivity. The implementation of a card wields a certain behavior on the wireless traffic stream that can be exploited to distinguish between NICs. We focused on the implementation of the rate switching algorithm as an opportunity for distinguishing between cards manufactured by different vendors.

The 802.11 PHY has multiple data transfer rate capabilities that allow an opportunity for the 802.11 MAC to perform dynamic rate switching with the objective of improving performance. For example, 802.11b supports data transfer rates of 1, 2, 5.5, and 11 Mbps. Each rate corresponds to a different PHY modulation scheme with its own trade-off between data throughput and operating range. It is the responsibility of the rate switching algorithm to select the proper rate (modulation scheme) per packet that gives maximum throughput for diverse link conditions. Implementation of the rate switching algorithm is vendor specific because it is beyond the scope of the 802.11 standard. Operation of the algorithm dictates the transmission rate of a frame, number of frames transmitted at the selected rate, how often to change rates, and the order of transmission rates selected. This will directly impact the behavior of the wireless stream. The duration of the frame transmission, arrival rate of frames, inter-frame delay, and other observable traffic characteristics differ among different implementations of the rate switching algorithm. The actual rate switching algorithm implemented in a card is unknown, as it is considered sensitive proprietary property. However, [11] discusses algorithms that are speculated to be used in current products.

IV. EMPIRICIAL ANALYSIS OF RATE SWITCHING

We have pinpointed rate switching as an opportunity for distinguishing between NICs produced by different vendors. Before developing our approach, we conducted an empirical analysis to characterize the rate switching phenomenon.

A. Experimental Setup

Analysis was conducted at a local hotspot on the campus of Georgia Institute of Technology. Over the course of 7 days we captured all traffic on the wireless network. We used a Toshiba laptop with a Linksys WPC11 wireless card to collect traffic. We put the wireless card into monitor mode using the *wlanctl-ng* utility and stored the captured traffic using Ethereal. With the card in monitor mode we were able to detect the transmission rate associated with each packet collected, while Ethereal appended a timestamp to each packet. We used the timing information and transmission rate to generate statistics. Traffic was collected for a total of 13.3 hours over the course of 7 days. During our observation period, there were a total of 61 wireless clients that visited the hotspot.

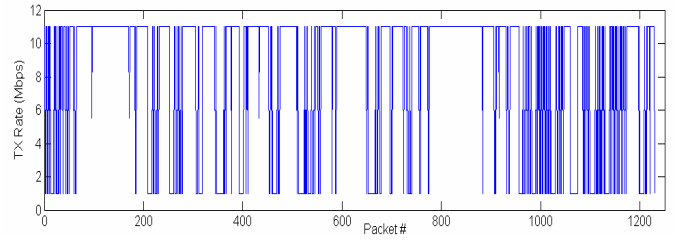

Figure 1. Host at local hotspot invoking rate switching.

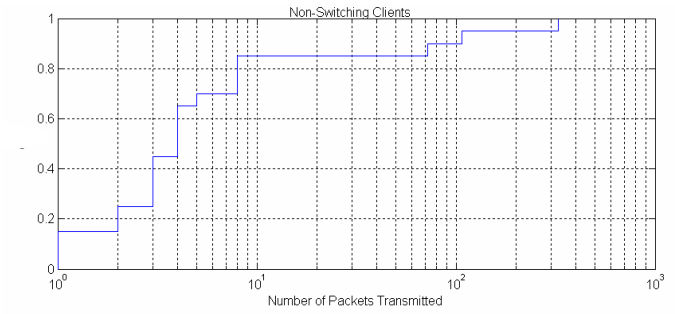

Figure 2. CDF of number of packets transmitted by non-switching clients.

B. *Results*

The results of our analysis show that rate switching is common at the hotspot. While this is definitely true for the hotspot we monitored, it is likely that RF interference occurs at most hotspots. Therefore, rate switching is likely a widespread, common phenomenon. Fig. 1 shows the transmission rate of each data frame transmitted by one of the clients at the local hotspot. This particular client switched transmission rates 279 times. Overall, Fig. 3(a) shows how often rate switching occurred for all wireless clients over the entire observation period. Fig. 3(b) shows that 67% of the clients performed rate switching, while 33% did not switch rates. Out of the clients that did not perform rate switching, 85% sent less than 9 packets (Fig. 2). If we exclude the non-switching clients that sent less than 9 packets (assuming that these clients were never properly authenticated to the network), then the percent of clients that perform rate switching becomes 92% (Fig. 3(c)).

Examining only the wireless clients that applied rate switching, Fig. 4 shows that 90% transmitted more than 37 packets and 88% were connected to hotspot more than 2 minutes. Also, 85% of these clients switched rates within the first 3 minutes of their connection.

We conclude that rate switching is a phenomenon that occurs. Our results show that the longer a wireless client is connected to the network and the more packets it transmits, the more likely rate switching is to occur. Therefore, rate switching is a viable attribute within the wireless NIC for distinguishing between cards.

V. AN APPROACH TO NIC IDENTIFICATION

We aim to identify wireless NICs without injecting traffic or requiring client software. To do so, we focus on the implementation of the rate switching algorithm as an opportunity for distinguishing between cards manufactured by

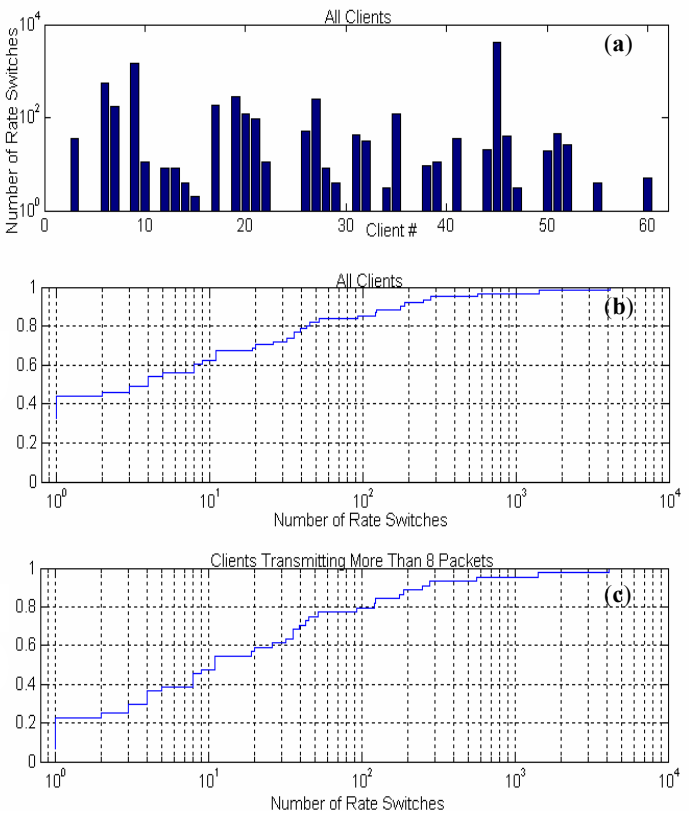

Figure 3. (a) Number of rate switches per client, (b) CDF of rate switches for all clients, (c) CDF of rate switches excluding non-switching clients that sent less than 8 packets.

different vendors. Given a trace of wireless traffic in which the rate switching algorithm has been invoked, we apply Fourier analysis to create a spectrum. The spectral features capture the temporal behavior of the trace influenced by the rate switching algorithm. We distinguish between NICs by comparing their spectral profiles.

A. *Encoding Trace into Signal*

Before we can apply spectral analysis, we must encode the traffic capture into a signal, a time series of events. Even with an encrypted payload, the 802.11 header offers a rich source of information, such as packet size, duration of payload, packet type, and retransmitted packets. Additionally, the traffic-capturing utility records a time stamp with each packet as it is collected. This information can be used to construct various types of signals.

Our goal is to construct a packet arrival timeseries from the wireless traffic stream. First we extract only data frames destined to the AP as indicated by the Subtype and ToDS fields within the header. This step isolates the actual data transmission frames from other overhead communication to minimize the amount of noise in the frequency domain. Next, we use a sampling bin of s seconds and define the arrival process, x(t), as the number of data packets that arrive in the bin[t, t+s]. Given a traffic capture T seconds long, we will have N=T/s samples. The maximum frequency that can be represented is 1/2s hertz.

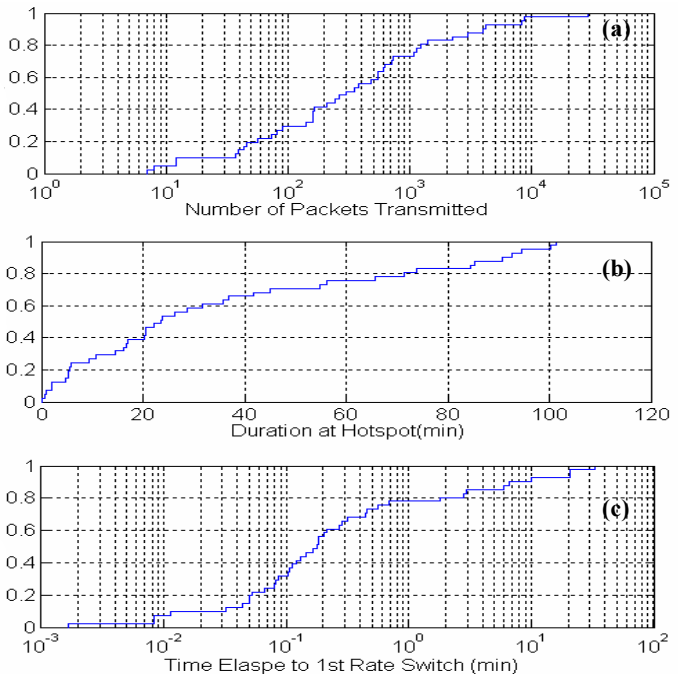

Figure 4. Rate switching clients – (a) CDF of number of packets transmitted, (b) CDF of duration at hotspot, (c) CDF of time elapsed to first rate switch.

B. *Signal Processing*

Signal processing has been previously applied to network research to detect and classify Denial of Service attacks [12], map the flow of wireless traffic, and to extract information about protocol behavior in encrypted wireless traffic [13]. We apply signal processing to wireless streams to analyze the timing of packet arrival, particularly while the rate switching algorithm is invoked to distinguish between NICs with presumably different algorithms.

Given a signal x(t) constructed using the process discussed above, we estimate the power spectral density (PSD) by computing the magnitude squared of the discrete-time Fourier transform of the sampled signal [14]. The power spectral density, $\hat{P}_{xx}$, of a process of length L is given in (1):

$$\hat{P}_{xx}(f) = \frac{|X_L(f)|^2}{f_s L} \quad (1)$$

Where the discrete-time Fourier transform, X_L is given in (2):

$$X_L(f) = \sum_{n=0}^{L-1} x_L[n] e^{-2\pi jfn/f_s} \quad (2)$$

VI. EXPERIMENTAL EVALUATION

During our experiments the only architectural differences are the NIC. We assume an environment in which the senders have a similar host configuration and similar end-user applications. Environments where these assumptions are applicable include warehouses, hotels, and restaurants [15]. In these environments, identical mobile handheld devices are used to do simple operations like inventory and tracking to speed up common functions.

A. *Setup*

The equipment used for the experiments consisted of a 1GHz Toshiba laptop, a Linksys 802.11b wireless router, and a Dell desktop. We tested three NICs: D-Link DWL-650, Linksys WPC11, and Lucent/Orinoco Gold. A fourth card was used to capture traffic in monitor mode using *tcpdump*. We also used the *iwconfig* utility to record the transmission rate of the card on one second intervals. For each experiment the laptop sends data over the wireless link through the wireless router to the desktop, which connects via a wired connection. The test was repeated ten times for each card.

We imposed a traffic load of 2.4Mbps. The load was light enough to not stress the host, but heavy enough to cause the NIC to be the bottleneck in the wireless system so that it is the primary factor influencing the behavior of the traffic. We used the *sock* [16] program to establish a user datagram protocol (UDP) connection carrying constant bit rate (CBR) traffic generating a 1470-byte packet every 5 milliseconds.

In a real environment, rate switching occurs due to changes in channel conditions caused by noise. This noise may be caused by the network contention, interference from neighboring networks operating on same channel, mobility of a wireless client, non-802.11 devices operating in the same frequency range, etc. During our experimental evaluation, we want to control the invoking of the rate switching algorithm. To emulate a real environment during our controlled experiments, we used the microwave as an artificial noise source to alter the condition of the wireless link. A microwave is capable of causing interference with the radio waves of an 802.11 network because it operates in the 2.4 GHz band. Since the energy of the microwave is normally shielded from the outside, a small 6 inch covered wire was inserted in the door with a portion of the wire hanging on the outside. In our experiments we started streaming data for 60 seconds, and then turned on the microwave causing an instant pulse of noise. After 60 seconds, the microwave was turned off and data continued streaming for another 60 seconds. We were able to trigger rate switching as seen in Fig. 5.

B. *Results*

We analyzed the captures using the encoding process and the power spectral density estimation as discussed above. We partitioned the analysis into three parts: the interval before injecting noise, the interval with noise, and the interval after injecting noise. During our analysis we used a sampling bin of 2 ms, which represents a frequency range up to 250 Hz. We also removed the mean from the sampled signal prior to calculating the discrete-time Fourier transform to remove the DC bias from the signal. If not removed it gives rise to significant power at 0 Hz.

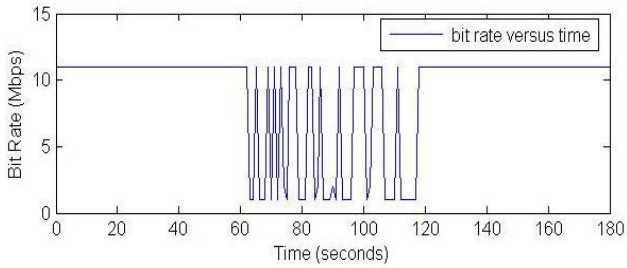

Figure. 5. Linksys card entering and exiting rate switching.

Our technique produces a spectrum that captures a view of the dynamic properties of a wireless stream. Because we filtered out all other frame types during the encoding process to only consider data frames, the spectrum captures the power or strength of the data transmission rate contained at a particular frequency. A similar analysis could be done to isolate the spectral content of management or control frames within a wireless stream or for a specific type of frame (i.e., retransmitted frames). Since the rate switching algorithm is only applied to data frames, we only encoded data frames.

Spectral analysis of the traffic trace prior to injecting noise generates a similar PSD for all three card types as illustrated by Fig. 6. The PSD also reveals that power is concentrated at discrete frequency points. Each card has the most prominent peaks at 100Hz and 200Hz. Confirming that the NICs behaved similarly in transmitting data frames when there is no rate switching (i.e., when the condition of the link is perceived as good).

During the interval in which noise was injected into the wireless environment, each card generated a distinctive PSD as shown in Fig. 7. In contrast to the noiseless period, we also observed that prominent peaks are scattered throughout the frequency range, especially at the lower frequencies. For example, in Fig. 7(a) the Lucent NIC still has prominent peaks at 100Hz and 200Hz (that were seen before injecting noise), but new distinctive peaks are found at the lower frequencies ranges (0-10Hz and 50-60Hz). The spread of prominent peaks throughout the frequency range indicates that a host is transmitting data frames at several different rates. This type of behavior is expected while a NIC is executing its rate switching algorithm. To address the stability of the PSD, we repeated experiments with the same card. The PSD for each card, respectively, is similar. A comparison between Fig. 7(a) and Fig. 8 illustrates similarity among repeated experiments for the Lucent NIC.

We also observed a distinction between NICs when comparing the normalized cumulative sum (NCS) of the spectrum during the noisy period. The slope of the NCS of the Linksys NIC is almost linear with a modest variation in the slope over the range of 90Hz to 130Hz (and 190-210Hz), indicating that the power spreads (somewhat evenly) across this frequency range. DLink shows a concentration of power around 50Hz indicated by the strong rising slope in the NCS. The strong rising slope of the NCS for Lucent indicates a concentration of power around 10Hz (and 100 Hz).

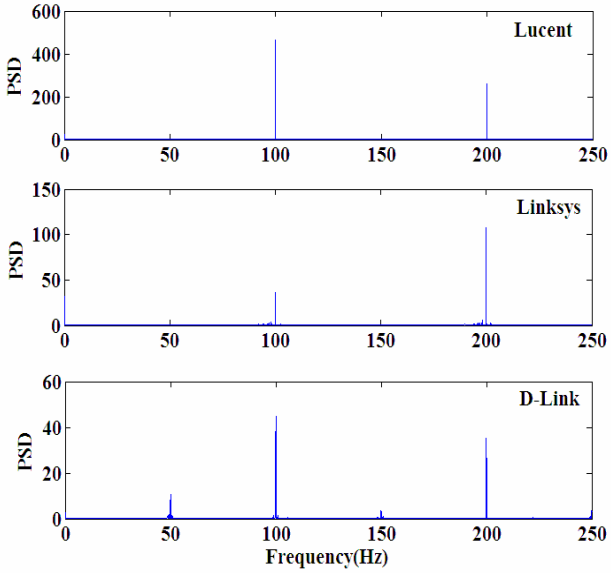

Figure. 6. PSD prior to injecting noise.

To numerically compare the spectrum between NICs, we locate the frequency points that exhibit the greatest amount of power. These key frequency points estimate the most prevalent sending rates of the NIC during rate switching. For our evaluation, we chose the top 50 frequency points to constitute a frequency set, $F = \{f_1, f_2, \ldots f_{50}\}$, as the spectral profile of a NIC. Table 1 displays the distribution of the set F for each NIC. By examining the range where the majority (54% or more) of the set F is located, we can distinguish between NICs. A Lucent NIC can be identified by 54% of frequencies concentrated between 0-10Hz. This indicates that a Lucent NIC most often sends data frames at a rate of 100ms during rate switching. The concentration of F for the Linksys card is over a broader range: 56% between 80-130Hz. A Linksys NIC most often attempts to send data frames between 7.7ms and 12.5ms during rate switching. Whereas, the DLink NIC most often sends data frames between 17ms and 25ms indicated by a concentration of 54% of the set F between 40-60Hz. The selection of 50 frequency points as a spectral profile proved to be adequate, because the distribution of the set F coincides with the observations made using NCS.

VII. CONCLUSION AND FUTURE WORK

In this paper we presented an approach for identifying a type of wireless NIC. We identified the rate switching algorithm as an attribute of a wireless NIC as an opportunity for distinguishing between NICs manufactured by different vendors. We used signal processing to extract markedly different traffic characteristics between different NICs during rate switching.

An important aspect we need to explore is the stability of the spectral profile. In the future, we plan to apply our technique to track hosts in a real network. We will also consider other factors like the setting of the NIC configuration parameters (i.e., RTS threshold, maximum retries, etc.). If the settings

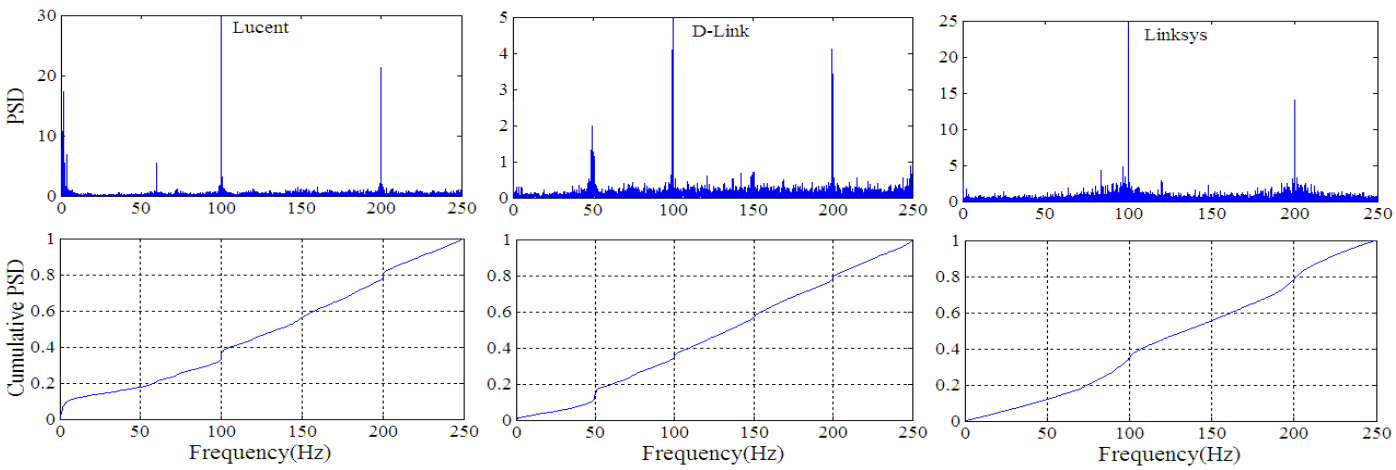

Figure 7. PSD and cumulative PSD-(a) left two Lucent, (b) middle two D-Link, (c) right two Linksys.

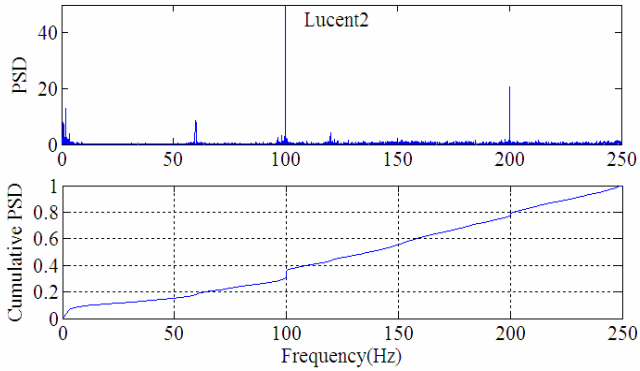

Figure. 8. PSD and cumulative PSD for repeated Lucent experiment.

change the spectrum, then they would become an essential part of the spectral profile. This could help to distinguish between wireless NICs manufactured by the same vendor. We also need to consider heterogeneous host systems. During our experiments we only used a single type of host. Hosts with different capacities may also affect the spectral profile. In this case we may need to broaden or restrict our characterization. The impact of these elements on the power spectrum may also require a different technique than establishing a set F of top frequencies, in order to have a more robust comparison.

We only examined the packet arrival rate for data frames. We plan to examine other time attributes of a stream, such as the inter-packet delay and variations in inter-packet delay.

TABLE I. DISTRIBUTION OF 50 DOMINATING FREQUENCIES

Frequency Range (Hz)	Lucent	D-Link	Linksys
0-10	54%	8%	2%
40-50	-	28%	-
50-60	4%	26%	-
80-90	-	-	10%
90-100	12%	8%	34%
100-110	10%	4%	8%
120-130	-	-	4%
140-150	-	2%	2%
150-160	-	6%	-
190-200	20%	4%	16%
200-210	-	12%	24%
220-230	-	-	-
230-240	-	2%	-

REFERENCES

[1] Nikita Borisov, Ian Golberg, and David Wagner, "Intercepting mobile communications: The insecurity of 802.11," MOBICOM 2001

[2] William A. Arbaugh, Narendar Shankar, and Y.C. Justin Wan, "Your 802.11 wireless network has no clothes," http://www.cs.umd.edu/~waa/wireless.pdf

[3] The Definitive Guide To Wireless WarX'ing, teknik.ekitap.gen.tr/TDGTW-WarXing.html

[4] AirDefense, www.airdefense.net

[5] Joshua Wright, "Detecting Wireless LAN MAC Address Spoofing," home.jwu.edu/jwright/

[6] www.wimetrics.com

[7] www.ipass.com/services/services_deviceid.html

[8] www.decodesystems.com/mt/97dec/

[9] Jeyanthi Hall, Michel Barbeau, and Evangelos Kranakis, "Detection of Transient in Radio Frequency Fingerprinting using Signal Phase," Internet and Information Technology (CIIT), St. Thomas, US Virgin Islands, November 2004.

[10] Tadayoshi Kohno, Andre Briodo, KC Claffy, "Remote physical device fingerprinting," IEEE Transactions on Dependable and Secure Computing, vol. 2, no. 2, pp. 93-108, April-June 2005.

[11] M.Lacage, M. Manshaei, and T. Turletti, "IEEE 802.11 Rate Adaption: A Practical Approach," ACM/IEEE MSWIM, Venice, Italy, October 2004.

[12] Chen-Mou Cheng, H.T. Kung, and Koan-Sin Tan, "Use of spectral analysis in defense against DoS attacks," in Proceedings of the IEEE GLOBECOM, Taipei, Taiwan, 2002.

[13] Craig Partridge et al., "Using Signal Processing to Analyze Wireless Data Traffic," *ACM Workshop on Wireless Security (WiSe),* Atlanta, GA, USA, September 28, 2002.

[14] Oppenheim, A.V., and R.W. Schafer, *Discrete-Time Signal Processing*, Prentice-Hall, 1989, pp. 730-742.

[15] Jim Geier, *Wireless LANs*, 2nd Edition. Indianapolis, IN: SAMs, 2001.

[16] W. Richard Stevens, *Unix Network Programming, Volume 1*. Upper Saddle River, NJ: Prentice Hall PTR 1998

[17] Alefiya Hussain, John Heidemann, Christos Papadopoulos, "Identification of repeated attacks using network traffic forensics," Technical Report ISI-TR-2003-577b, USC/Information Sciences Institute, August, 2003.