Rogue Access Point Detection using Temporal Traffic Characteristics

Raheem Beyah, Shantanu Kangude, George Yu, Brian Strickland, and John Copeland
Communications Systems Center
School of Electrical and Computer Engineering
Georgia Institute of Technology

Abstract-**As the cost of 802.11 hardware continues to fall, the appeal of inserting unauthorized wireless access into enterprise networks grows. These rogue access points (APs) expose the enterprise network to a barrage of security vulnerabilities in that they are typically connected to a network port behind the firewall. Most of the current approaches to detecting rogue APs are rudimentary and are easily evaded by hackers. We propose the use of temporal traffic characteristics to detect rogue APs at a central location. This detection is independent of the wireless technology (802.11a, 802.11b, or 802.11g), is scalable, does not posses the inefficiencies of the current solutions, and is independent of the signal range of the rogue APs.**

I. INTRODUCTION

As users realize the benefits of wireless networking at home, they begin to desire the same flexibility in the workplace. Rather than waiting for their IT organizations to install a wireless network, users are taking matters into their own hands. Employees are deploying rogue APs and building large grassroots wireless networks without the knowledge or consent of their IT departments. These rogue APs represent a serious breach of network security. They are typically connected to a network port behind the corporate firewall. Additionally, employees rarely enable even the most basic security settings on rogue APs, making it easy for unauthorized outsiders to use the AP and eavesdrop on network traffic.

Corporate network administrators are not the only ones who are, or should be, concerned about the rogue AP problem. Universities' support staffs are already having a difficult time trying to manage the security of the PCs on the network. Under pressure from students, staff, and administration, the universities' networking staffs have deployed wireless networking across campus with minimum security measures. Some enable the wired equivalency protocol (WEP) and perform some degree of application-level authentication before allowing nodes to become associated with the network. While this is a good start, due to various factors, including cost constraints, many universities do not have specific wireless intrusion detection systems, nor do they have any method of preventing students, staff, or faculty from installing their own AP. This rogue AP may allow unauthorized pleasant or malicious users onto the network. Further, the network administrator will have difficulty tracking down the assailant.

Similarly, a growing number of hotels now offer broadband access in upgraded and even regular rooms. Many of the services are provided by a third-party who accepts payment in exchange for daily internet access for one machine. This is loosely enforced by assigning a temporary IP address to the requesting machine and storing that machine's medium access control (MAC) address. Once the time expires, the IP expires and the communication is blocked until the fee is paid again. One cannot share the internet access with another machine in the room because the IP is linked to one MAC address. This control is easily circumvented so that multiple users can share the access. The room's users can simply use a router that has MAC address spoofing and network address translation (NAT) features. If a user wants to share this one connection with everyone on the hall, he merely has to use a wireless router with the same features.

In general, every organization that has a network should have some form of rogue AP detection, *especially organizations that do not have wireless networks*. These organizations do not expect, or think to consider, any sort of malicious wireless activity because they have not deployed a wireless network. This thinking, can result in two undesirable outcomes: 1) an employee installs a rogue AP and a malicious user stumbles upon a wide-open invitation to the corporate network as they "war drive"; or 2) a hacker installs an AP out of site on a live port (e.g., hotel lobby, hardware store, hospital, government building, etc.) and has a gateway to the network from the parking lot or, using signal boosting antennas, even farther away.

To these authors' knowledge, the rogue AP detection problem has been overlooked by academic researchers. Most solutions have been quick fixes by wireless local area network (WLAN) security vendors. We illustrate, by empirical analysis, a novel approach to rogue AP detection using temporal traffic characteristics. The remainder of the paper is organized as follows: In section II we discuss current approaches. We discuss the background of our scheme in Section III. In Section IV we describe our experimental setup. Section V gives the results and performance analysis of our scheme. In Section VI we give the conclusion and we conclude this paper with a discussion on future work in Section VII.

II. CURRENT APPROACHES

A. Wireless Approaches

Most of the current approaches for detecting rogue APs are rudimentary and easily evaded by hackers. Some organizations have equipped IT personnel with wireless packet analyzer tools (e.g., sniffers) on laptops and handheld

 0-7803-8794-5/04/$20.00

devices (e.g., AirMagnet [4] and NetStumbler [5]), forcing IT personnel to walk the halls of the enterprise or campus searching for rogue APs. This method is generally ineffective because manual scans are time-consuming and expensive – and, therefore, are conducted infrequently. Also, with 802.11 hardware operating at separate frequencies (802.11a - 5Ghz and 802.11b - 2.4Ghz), IT personnel must upgrade their detection devices to accommodate multiple frequencies. Moreover, scans are easy to elude, since a rogue AP can easily be unplugged when the scan takes place.

Most vendors today go a step further. Rather than relying on an employee equipped with a scanner, they enable IT to initiate an enterprise-wide scan from a central location. This is possible by using separate hardware devices [2][3][7][8] (e.g., sensors) or using APs to detect beacons from surrounding APs [2], and transmitting this information back to a central management platform containing the wireless network policy for analysis [1]. This method becomes costly, considering that one must place sensors or APs throughout the entire enterprise to monitor the air waves. This technique is also completely impractical for the networks that do not have wireless APs. Much like the drawback of the "walking the halls" solution, each sensor/AP must operate at both frequencies to be completely effective. Moreover, with sensors deployed throughout the network, one still may not be able to detect the rogue AP. The clever employee could have used a directional antenna, or reduced the signal strength to cover the small range within his/her office. Another drawback of wireless-based solutions is that they will falsely report the wireless network in the coffee house next door as a rogue.

B. Hybrid Wireless and Wired Approach

Taking a step in the right direction, Wavelink [2] combined the previously mentioned techniques for detecting Rogue APs with listening at network layers 2 and 3 and querying switches and routers to determine what devices are connected to them, thus, attempting to provide a hybrid wired and wireless approach to detecting Rogue APs. This fails for the same reasons that the wired-only solutions discussed in detail below fail.

C. Wired Approaches

Cisco offers a more complete, scalable, and comprehensive approach using a suite of tools [9] that are not limited by signal range. They attempt to detect APs by querying routers and switches for company MAC address assignments (i.e., if the MAC address belongs to Linksys, the MAC address cannot belong to a PC and becomes suspicious). This fails because MAC addresses can be spoofed or cloned easily by an AP. Another approach in the suite is the use of httpd query to communicate with the web server residing on the AP. This is a good approach, but the node must already be suspected as being an AP (maybe using one of the aforementioned methods), or every node on the network must be queried. This approach assumes that the wireless router responds to httpd queries. Additionally, this invasive approach is considered active, adding significant unwanted traffic on the network and can also alert an advanced rogue AP user of a scan for the AP. The suite also has an application which allows the viewing of html code generated when configuring AP settings. Though this approach will work in theory, the window of opportunity is limited since this data is only transmitted when the AP's configuration is updated. Additionally, as signature-based IDSs can attest, reassembling application-level data becomes more difficult and impractical as network speeds increase.

Another LAN only approach is presented by Wimetrics [4]. Their product has a LAN only approach, but is ambiguous with details. The basic premise of their work is that they probe the network to identify the profile of a wireless AP. While the details were unclear, Wimetrics' general approach proves not scalable since it requires a PC to sit on each segment of the network. Their approach unjustly assumes that the network is a shared network. As discussed in the previous section, APs can be configured to ignore network queries.

III. BACKGROUND ON OUR SCHEME

The primary goal of our research is to detect rogue APs from a central location (a switch that supports a subnet) with the detection independent of the wireless technology. We show a scalable solution, thus not attempting to reassemble data before analysis. Also, this solution will function independently of the signal range of the rogue APs. Our research involves comparing traffic characteristics of flows from different sources in a LAN segment and detecting traffic coming from a wireless AP.

Our scheme starts with the hypothesis that a wireless link in a network path of multiple links would cause a more random and temporally different spreading of packets, as compared to a path that has only wired links. Consider Figure 1 on the following page. The objective is to differentiate the scenario shown in Figure 1, in which a switch port is connected to a network segment that has no wireless links, from the scenario shown in Figure 2, in which a switch port is connected to a segment with at least one wireless link. The assumption is that a majority of ports in a switch are connected to network segments that have only wired links. The processing and decision making are performed at the switch with the input as the link layer traffic traversing, in both directions, a switch port. The number of hops between the switch and end point will most likely affect the temporal characteristics of traffic as observed at the switch. Queuing and congestion tend to mask the temporal shaping of traffic through end points. However, we consider scenarios that involve network segments with, at most, 2 links from the detecting switch. Such scenarios are commonly observed in most Ethernet local networks. The reliability of wired links makes the temporal characteristics of traffic, in a path, to be

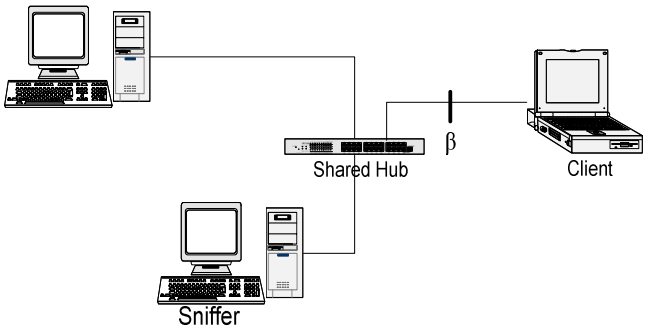

Figure 1. Configuration with wired link.

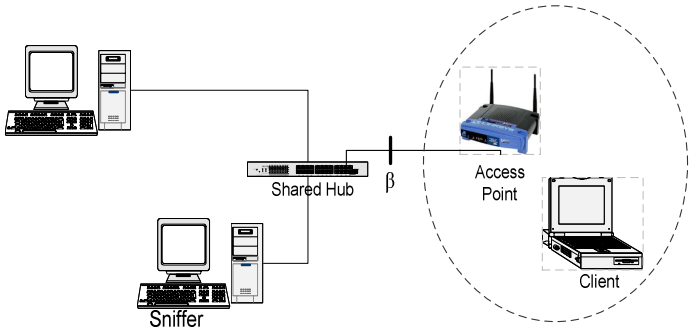

Figure 2. Configuration with wireless link.

shaped mostly due to higher layer (e.g., TCP) mechanics, with a relatively simple shaping from the link layer. A wireless link however, shapes traffic differently. Due to variations in channel conditions, wireless link capacity varies and random delays are introduced. The difference in link speed between wired and wireless links also shapes the characteristics significantly. Accordingly, our detection scheme is based on the premise that if traffic at a switch port is observed in both directions over time, and input-response correlated, different patterns may be observed for segments with and without wireless links.

The input-response correlation involves estimating which part of traffic is in response to which impulse (or input). Thus, for every response quanta of traffic from an end point in the segment, temporal characteristics may be analyzed by classifying mean, variance, and other frequency response characteristics of inter-packet spacing. For all switch ports, such analysis over time is compactly stored in a number of state variables. As time progresses, a rogue AP (or a number of them) are detected as the difference in state variables between ports crosses a threshold. The aim of this research is to experiment and derive such state representation and its derivation from the observed temporal characteristics of traffic. In this paper, we present the observed differences in inter-packet spacing in wired and rogue AP scenarios.

IV. EXPERIMENTAL SETUP

To test the inter-packet spacing theory in a controlled environment, we built an experimental testbed. The base scenario is depicted in Figure 1, where the only traffic on the network was the traffic generated from our client. The first scenario shows a PC server machine, a laptop acting as a client connected by a 100Mbps shared hub (the hub was only used to create a scenario where we could sniff the traffic on the link). An additional PC was used to observe traffic traversing the link. The network sniffing software used was Ethereal [10].

We conducted experiments using FTP traffic for ten different file sizes from 10MB to 100MB using increments of 10MB. Each test was done ten times and similar results were generated. The client machine (laptop) initiated the transfer and downloaded a file from the server. For this experiment, we label the path taken by the data traffic from the server to the client as the forward path. The data was captured to be analyzed offline. For the wireless experiments, we extended the network with a wireless router as shown in Figure 2.

V. PERFORMANCE ANALYSIS

The results presented in this section follow the same general pattern. Each figure shows a cumulative distribution function (CDF) of the inter-packet spacing of the wired and wireless traffic. As such, the y-axes of each graph peaks at one, representing the cumulative maximum value, while the x-axes are displayed in logarithmic form measuring time in seconds. The results are shown on the following page.

A. Rogue AP Detection Performed at Immediate Switch

Figure 3 shows the inter-packet spacing of the forward path. That is, the variation in spacing of the packets sent from server to client on the forward path. As expected, due to the wireless channel variability and the inferior link capacity compared to a wired link, the inter-packet spacing for the traffic traversing the wired link is noticeably less than that of the wireless link. Specifically shown in Figure 3, 80% of the inter-packet spacing for the wired link was less than 1ms, while around 90% of the inter-packet spacing values were greater than 1ms. The data shown in Figures 3 & 4 assumes that the monitoring for a rogue AP would be performed in the immediate switch connecting the rogue AP. This one hop from switch to desktop configuration is extremely common. Of course, the switch would be located in a wiring closet located several meters away from the offending host.

Figure 4 shows the inter-packet spacing for the reverse path. Though not as significant a difference as in the forward path case, a clear distinction between the inter-packet spacing on the wired and wireless link is given. The reverse path consists primarily of small packets (ACKS) which are less frequent and more asynchronous, thus loosing some of the desired characteristics. Therefore, the remaining of the CDFs depict the traffic on the forward path due to its inherent ability to present the traffic distinctions.

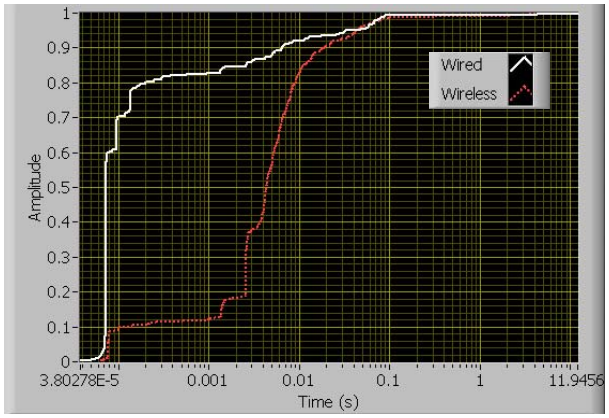

Figure 3. Inter-packet spacing of wired and wireless traffic on the forward path.

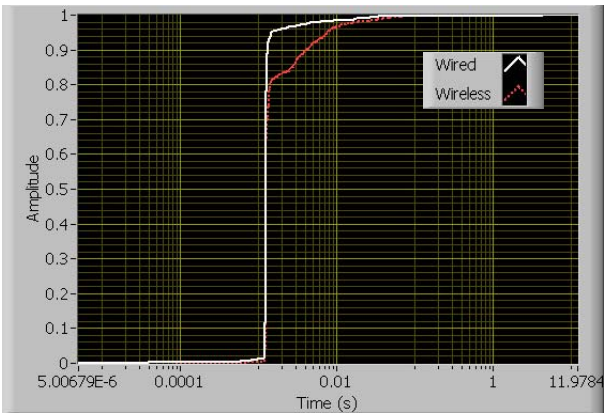

Figure 4. Inter-packet spacing of wired and wireless traffic on the reverse path.

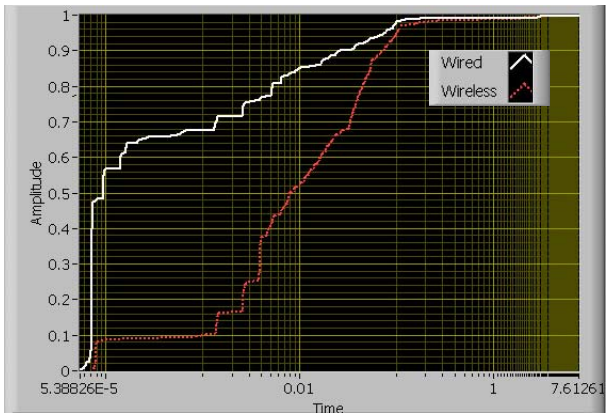

Figure 5. Inter-packet spacing of wired and wireless traffic on the forward path with TCP cross traffic.

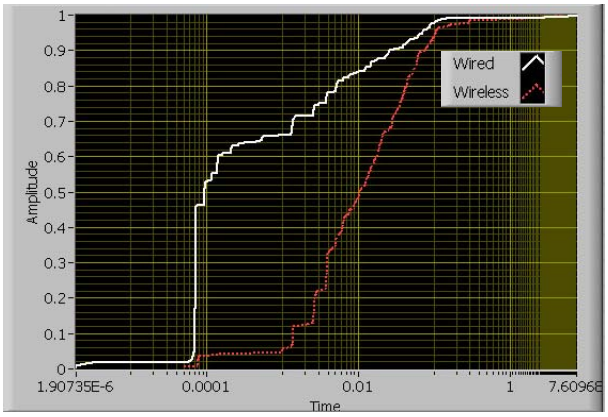

Figure 6. Inter-packet spacing of wired and wireless traffic on the forward path with UDP constant bit rate cross traffic at 500Kb/s.

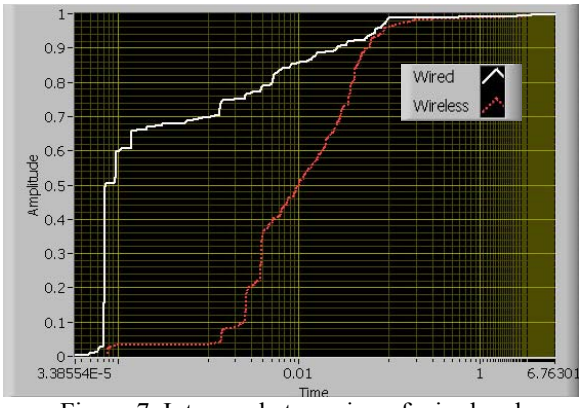

Figure 7. Inter-packet spacing of wired and wireless traffic on the forward path with UDP constant bit rate cross traffic at 10Mb/s.

B. Rogue AP Detection Performed One Hop Away

Aggregating the responsibility of monitoring rogue APs on several segments into a central location is an obvious need. We attempt to take a step in that direction by testing our approach one hop downstream from the monitored node. Specifically, we test to see if the inter-packet spacing characteristic will hold, once the flows traverse a loaded switch and once potentially variable queuing delays are introduced. To more accurately model a congested intermediate switch, we generated different levels of cross traffic at the intermediate switch. The purpose of the cross traffic is to place a load on the intermediate switch, not to inject packets into the experiment flow stream. Thus, the testbeds shown in Figures 1 and 2 were expanded to include an additional router with two machines connected. The additional components were injected at the point "β*"* shown in Figures 1 and *2*.

Figure 5 shows that the basic premise of using inter-packet spacing to detect unwanted wireless traffic holds with an additional hop placed in the network. The intermediate node is lightly loaded with TCP cross traffic traversing the switch simultaneous with the experiment traffic.

Figures 6 & 7 illustrate the resilience of the schemes and show that inter-packet characteristics can be preserved even when traversing an intermediate switch which was mildly to heavily loaded with aggressive UDP traffic.

VI. CONCLUSION

In this paper, we presented a technique to detect rogue APs from a central location using temporal traffic characteristics. The technique is novel because it presents a scalable LAN-only solution that is independent of the wireless technology. Also, this solution will function independent of the signal range of the rogue APs.

VII. FUTURE WORK

Our proposed scheme requires that the rogue AP monitoring be performed in the switch where the rogue AP is immediately connected or one immediate hop downstream. This becomes challenging and costly when one considers large-scale networks (e.g., campus networks). We are currently investigating techniques for performing the analysis further upstream. Specifically, we plan to determine the relationship between the number of hops upstream and the probability of accurate detection. We expect this will prove more efficient. We will also run actual experiments with different data models on a large campus network and perform trace-driven simulations to characterize interesting scenarios.

Additionally, we plan to create the appropriate function that will take the observed statistics as input and generate the probability of Rogue AP Index (RAI) parameter. Once the statistics and probability rogue function has been determined and deemed acceptable, we will extend our analysis to support non-traditional (other than first in first out (FIFO)) queuing at switches (i.e., queues that have priority associated with them - an example would be an intranet that supported voice over IP (VoIP), where voice traffic has higher priority than data traffic in the LAN).

Our final area of interest deals with the ability to perform the rogue AP identification in an automated fashion. Specifically, this functionality should be a utility that can run on a switch. Our current scheme uses a visual approach that would prove challenging for a computer system to analyze. We are looking at an approach where we can use the area under the curves of the traffic shown in Figures 3-7 to automate the analysis. As observed in the figures, the area under the wired curves is significantly greater than that under the wireless curves. Thus, by computing and comparing the areas, we can potentially perform this rogue AP identification, using inter-packet spacing, without human intervention.

REFERENCES

[1] www.airwave.com/airwave_rogue_detection.pdf
[2] www.wavelink.com/downloads/pdf/wlmobilemanager_wp_rogueap.pdf
[3] www.highwalltech.com/products.cfm?menu=hwsent&page=hwsent
[4] www.wimetrics.com/WAPD.htm
[5] www.airmagnet.com
[6] www.netstumbler.com
[7] www.computerworld.com/mobiletopics/mobile/story/0,10801,72065,00.html
[8] www.airdefense.net
[9] winfingerprint.sourceforge.net/presentations/APTools.ppt
[10] www.ethereal.com