

# Physically Locating Wireless Intruders

**Frank Adelstein**

(ATC-NY, New York  
[fadelstein@atc-nycorp.com](mailto:fadelstein@atc-nycorp.com))

**Prasanth Alla**

(ATC-NY, New York  
[prasanth@atc-nycorp.com](mailto:prasanth@atc-nycorp.com))

**Rob Joyce**

(ATC-NY, New York  
[rob@atc-nycorp.com](mailto:rob@atc-nycorp.com))

**Golden G. Richard III**

(University of New Orleans, Louisiana  
[golden@cs.uno.edu](mailto:golden@cs.uno.edu))

**Abstract:** Wireless networks, specifically IEEE 802.11, are inexpensive and easy to deploy, but their signals can be detected by eavesdroppers at great distances. Even with existing and new security measures, wireless networks have a higher risk than wired nets. WIDS, Wireless Intrusion Detection System, provides an additional layer of security by combining intrusion detection with physical location determination, using directional antennas. We briefly describe WIDS and present our initial results of remote station location using inexpensive hardware.

**Key Words:** Wireless, intrusion detection system, 802.11, antenna

**Category:** D.4.6

## 1. Introduction

Wireless local area networks (WLANs) are very popular due to their availability and low price. Portable device manufacturers are already providing 802.11 wireless cards as a standard built-in networking device. Installing a base station and a wireless card in one or more devices (PC, laptop, printer, etc) gives an almost instant mobile network, which can use a high-speed, typically broadband, Internet connection. This is faster, cheaper and much more convenient than running CAT-5 cable and installing outlets, hubs and switches for a traditional Ethernet network.

The wireless medium by its very nature cannot be contained. Reliable omnidirectional communication for devices at 100m (the typical range of an 802.11 Access Point (AP)) requires a signal strength that is easily detected at greater distances. Craig Ellison's [Ellison 2001] research showed that a majority of 802.11b

wireless LANs are vulnerable. Using a laptop with a wireless card and a 14db Yagi antenna mounted on a tripod, he quickly identified 61 APs within a six-block radius in Manhattan. The shareware program NetStumbler reports detailed information about each AP. This technique, called “war-driving,” is increasingly popular, and there are sites dedicated to mapping unprotected wireless networks (e.g., <http://netstumbler.com> and <http://wardriving.com>).

The industry approach has been to layer data encryption onto the wireless signal with first 40 bit and then 128 bit encoding. The 802.11 standard specifies Wired Equivalent Privacy (WEP), a link-layer security protocol. WEP is based on the RC4 stream cipher, a symmetric cipher (the same key is used for both encryption and decryption). These security mechanisms—intended to maintain the confidentiality, integrity, and availability of wireless communications—are problematic. Several WEP flaws have been widely documented and disseminated [see Hayes 2001, Information Security 2001, and <http://www.cs.umd.edu/~waa/wireless.html>]. Each of these flaws allows passive or active attacks on wireless transmissions, by which attackers can decrypt information or inject forged information into the transmissions.

Several vendors, such as 3Com, Cisco, DLink, LinkSys, added access control lists (ACLs), implemented through MAC address filtering, to increase security. MAC address filtering amounts to allowing predetermined clients with specific hardware addresses to authenticate and associate. Unfortunately, MAC addresses can be forged and MAC address filtering is not available for ad-hoc (i.e., peer-to-peer) 802.11 networks.

The 802.11i protocol addresses most of WEP’s shortcomings; however, several problems remain. First, a large installed base of legacy systems will remain unprotected for some time. Second, flaws will always exist, due to misconfiguration and implementation bugs. And third, authentication mechanisms can be compromised by lost or stolen equipment.

Because of these ever-present risks, a layered protection mechanism is needed. WIDS, Wireless Intrusion Detection System, can provide an extra layer of protection that detects intruders. In addition, it determines the physical location of intruders, information not provided by any other means. The WIDS approach allows the presence and *location* of the intruder to be determined. This paper focuses on the experimental results obtained from our initial work on developing an early WIDS prototype.

The rest of the paper is organized as follows. Section 2 presents the problem. Section 3 describes the WIDS approach. Section 4 outlines the experiments we performed on directional location of intruders. We present the results of the experiments in Section 5 and conclusions and future work in Section 6.

## **2. Problem Statement**

The combination of inherent insecurity of wireless networks (signals radiating further than the intended coverage area) and weaknesses in the current security mechanisms make them open targets for attacks, which limits their deployment. We focus on active attacks. Because an intruder can attack from any point close enough

to an Access Point, it is a challenge to devise an effective intrusion detection system. Knowing the physical location of the attacker aids in the intrusion detection, as well as the response. Although the industry has been betting that the benefits of wireless technology outweigh the security risks, some customers, such as the military, have no alternative to very selective deployment of wireless networks and severe limits on the data they carry. Current intrusion detection mechanisms are not flexible enough to provide early detection of intruders in wireless networks. WIDS improves upon the state-of-the-art by providing earlier detection capabilities than are currently available.

### 3. WIDS Approach

WIDS comprises three tasks. First, a WIDS access point must detect a signal from a remote station. Second, based on the data in the signal (MAC address, IP header information, application data, etc.), the AP determines the remote station is an intruder. And finally, two or more APs determine the location of the intruder, using directional antennas. We describe experiments to test the location capability later in this paper.

An intrusion detection system (IDS) coordinates the physical location data, events seen on the network, and preset administrative policies.

Related work in WLAN intrusion detection has been done by Wright and Foust. Wright [Wright 2002] proposed techniques to detect war-driving programs, including NetStumbler, but focused on probe detection only, not physical location. Foust [Foust 2002] proposed a simple method of locating remote stations using signal strength but used only fixed omnidirectional antennas.

In the following sections, we describe the WIDS APs, the directional antennas, and the IDS architecture.

#### 3.1. WIDS AP

The typical WIDS installation [see Fig. 1], consists of a normal omni-directional AP located in the center of the physical facility and WIDS APs located around the perimeter and directional antennas pointing outward. (Note that there could be multiple omni, or directional, APs inside the perimeter; we do not show them in order to keep the figure uncluttered.) “Authorized” users connect to the omni AP from within the perimeter. We assume that internal security procedures handle authentication inside the perimeter. The problem is that the omni’s coverage extends *beyond* the physical perimeter.




Figure 1: WIDS access point protecting a perimeter

WIDS addresses this problem because an intruder attempting to break into the wireless network from the outside will contact the WIDS APs *before* coming within range of the omni AP. The WIDS APs will detect intrusions based on both known attack signatures (such as network probes from NetStumbler or other non-passive “war-driving” programs) and behavior based signatures (such as an internal MAC address suddenly appearing in an external location on a machine with different OS characteristics than it had previously). WIDS will detect anomalous behavior by tying the signature data into a behavior-based intrusion detection system [Hofmeyr et al 1998 and Marceau 2000].

WIDS is based on open-source access point code, HostAP [Malinen 2003], using Prism II wireless card drivers. The capabilities of WIDS include user-specifiable, trigger-based intrusion detection, allowing the user to tailor the criteria that trigger alarms. Triggers are based on both packet/frame data and historical/behavioral data—including typical signal strength levels, locations, time-to-live (TTL) values, IPIDs, etc. for particular MAC addresses. For example, we could specify that packets from MAC 00:20:E0:8C:92:88 must have a TTL of 30 and a signal strength greater than 10; any packets from this MAC address not meeting those criteria will raise an alarm. This is similar to a honeypot or honeynet (see <http://project.honeynet.org>), but allows more detailed control.

We used the publicly available software HostAP [Malinen 2003], a set of loadable kernel modules and a user-space daemon, under Linux, for interacting with and configuring the module that make an 802.11b wireless Ethernet card become an access point (instead of a remote station). The Zoom Air 4100 series cards feature both the Prism II chipset and an external antenna plug (a reverse-polarity SMA plug), two requirements for this project.

### 3.2 Directional Antenna

An intruder's location can be determined by rotating a directional antenna 360 degrees while monitoring the signal strength. Ideally, the signal would have a single, global maximum representing the direction of the intruder. Unfortunately, antennas are not ideal due to multipath reflection and other environmental scattering, resulting in a more complex signal. However, we can determine the intruder's bearing by measuring an antenna's "signature" ahead of time, and then comparing the intruder's data to the signature data. Gathering a full set of data points would take too much time and delay the response. Therefore, we may collect fewer data points for the intruder as compared to the signature. We conducted experiments to verify the accuracy of correlating these two data sets.

We used three different types of directional antennas for the project: grid array, parabolic dish, and Vagi, described below. Different antenna types have different signatures but the signature of each individual antenna type should be independent of the distance to the target (i.e., the distance will only attenuate the signal and affect its amplitude).

The *grid array*, a parabolic grid array antenna, is the largest of the antennas and has the greatest gain. Its specifications are: 8 lbs, 24 dBi gain, 10° beam angle, >28 dB (see <http://www.ydi.com/products/pt2421-pt2424.php>).

The *Vagi*, a V-shaped Yagi-style antenna, is the lightest of the group, with decent gain. Its specifications are: 1.5 lbs, 16 dBi gain, 25° beam angle and 19 dB F/B ratio (see [http://www.pacwireless.com/html/vagi\\_series.html](http://www.pacwireless.com/html/vagi_series.html)).

The *echo*, a parabolic dish antenna, is a nice compromise and has the best F/B (front-to-back) ratio. Its specifications are: 4.4 lbs, 14 dBi gain, 26° beam angle, >30 dB F/B (see [http://www.pacwireless.com/html/echo\\_series.html](http://www.pacwireless.com/html/echo_series.html)).

### 3.3 IDS Architecture

WIDS Access Points (WIDS AP) are fully functional access points. Security mechanisms can examine each incoming packet by using the logging facility in HostAP.

The system is configured as follows. The controller knows all the legitimate MAC addresses and their system profiles. Each antenna knows its scanning range, its neighbors, and their scanning ranges. Each WIDS AP has a set of triggers (security alerts) and response mechanisms for the triggers defined. [Figure 2] shows the components of the IDS. There are two general states to the IDS operation: startup and steady state. We describe each below.




Figure 2: Intrusion Detection System Overview

On startup, the controller distributes the list of legitimate nodes, profiles, and their MACs to the WIDS Access Points (WIDS APs). Then each WIDS AP initializes its security mechanisms with the coordinates of the security perimeter.

During the steady state, a WIDS AP receives an Association Request (AR) from an intruder node (outside the perimeter). It then captures the MAC address. If the MAC address is not in the access control list distributed by the controller, the information is passed on to the controller to take appropriate action. Otherwise the controller may request a history (activity log), if any, of that MAC from the WIDS AP or its neighbors. This includes the omnidirectional AP as well. The requested information includes last seen sequence numbers, system profile, and observed activity.

The controller then contacts the neighbors of the reporting WIDS AP and requests the neighbors to scan for the intruder to accurately determine its location. Based on the location and signal strength of the intruder, the controller determines who should monitor the intruder.

The controller then analyzes the frames against the history of that MAC. This can include information gathering from that node, such as an OS fingerprint. After the controller has completed processing, it may send an alert to an operator. The WIDS AP takes an appropriate response, which can include redirecting the intruder to a fake machine (a “fishbowl”).

### Response Modes

We have defined several response modes: none, locate, connect, and phalanx, each of which is described below. The *configure* and *command* messages can be used to

set the response, with the former setting the default response and the latter specifying a response for a particular MAC.

**None.** This is the most basic response mode, in which the WIDS AP does nothing. It will refuse to accept association messages from the intruder. Because the WIDS AP refuses associations, the internal omni-directional AP may accept association requests if the remote station (“intruder”) is within the field of the omni.

**Locate.** In this response mode, after one WIDS AP detects the intruder and reports its angle via an *alert* message, the controller issues a *command* message instructing the AP to send a *disassociate* message to the intruder and ignore further associate requests, and sends a *command* to its neighbor APs instructing them to track the intruder’s MAC address and report on its location. When the WIDS AP sends an *alert* message, the controller can determine the location of the intruder using triangulation.

**Connect.** The connect response mode instructs the WIDS AP to allow the intruder to associate with it and will allow it to “connect” to a network. The WIDS AP will route traffic to a fake network (a “fishbowl,” sometimes referred to as a “honeypot”) where further analyses on the intruder can be performed. The fishbowl and analyses are performed by a separate component; the WIDS AP merely routes traffic to that network.

**Phalanx.** The final response mode defined in the initial prototype version is the “phalanx” mode. This uses the “fakeAP” software which generates thousands of virtual, fake access points. By changing the MAC address, ESSID, and channel every time a beacon is broadcast, fakeAP can make it appear that there are thousands of other access points, which makes it difficult for the intruder to detect where the real access point is. In addition, the intruder must send a lot of packets to find the real access point, which makes its detection and location easier. Of course, this mode is the opposite of “run silent, run deep” and would only be appropriate to use in certain situations.

**Track.** This response mode is used to track the location of an intruder and report if the location changes. This mode is similar to the locate mode, but can have extra triggers. The Track mode can be configured so that an *alert* message is sent only if the intruder moves more than a specified distance or into or out of a specified location.

## 4. Experiments

In the experiments, we computed the remote station and compared it to the actual location. We performed some initial tests to characterize each antenna’s beam pattern. This data was then used as training information that would allow us to estimate the incident angle of an intruder by correlation. For each experiment, we set up a portable WIDS AP, using a telescope tripod mount and one of the directional antennas described above. A Linux laptop running HostAP was connected to the antenna under test. An “intruder” machine was placed a few hundred feet away and allowed to associate with the WIDS AP. The intruder was also set to ping the AP once per second, to ensure that some traffic was flowing through the wireless link.

## 4.1 Correlation

We used the `iwspy` Linux command to measure the intruder’s signal strength at the WIDS AP. `iwspy` reports both noise and signal level in dB, along with a “signal quality” function of unspecified units. As we wish to find the location of the transmitter (intruder), we only care about the signal strength: noise level and overall quality are secondary, and likely to be affected by external factors. The values reported by `iwspy` are updated only when traffic from the target is seen, thus the need for the continuous pinging. We wrote a short Perl script to collect and timestamp values from `iwspy`, sorting them by MAC address.

With the setup complete, we then noted the true angle of the intruder machine as indicated by the tripod. Working slowly and in fixed increments, we rotated the directional antenna and noted the signal strength value reported after a few seconds. Readings were taken over the entire 360° range to help in characterizing side- and back-lobes.

The tests were performed in a realistic environment in which WIDS might be deployed, specifically, an office building and adjacent parking lot with cars present. In the interest of space, we will present only the final results of our experiments.

One of the primary goals of the WIDS effort is to locate intruders using the beam patterns of the directional antennas. For each individual antenna, we first obtain (experimentally) a training beam pattern using the methods described above. The training data set consists of signal strength measurements at 1° increments over the entire 360° of antenna rotation. A reference point for 0° is chosen to be north, though any other convenient 0°-point can be used. Given a set of signal strength measurements of the intruder’s transmissions, at different angles of the WIDS AP’s antenna, we correlate the training data with these new measurements to determine the intruder’s location (angle). The computed angle is that with the best fit to the known antenna beam pattern from the training set.

Signal strength measurements from the intruder will not, in general, occur at 1° increments, nor will there be measurements for even a large subset of the 360° range, due to physical constraints limiting the antenna’s rotation. In such cases, we interpolate among these non-uniform measurements to correlate among uniformly-spaced samples. For our initial work, we use simple linear interpolation. Future work can include more sophisticated techniques, as well as estimates of the uncertainty introduced by such interpolation.

After interpolation of test samples, we then compute the correlation for each possible angular offset:

$$\rho(k) = \sum_{n=1}^N s_{train}(n+k \bmod N) s_{test}(n)$$

where  $s_{train}(n)$  is the  $n^{th}$  training sample of signal strength versus angle and  $s_{test}(n)$  is the  $n^{th}$  test sample of signal strength for this potential intruder (likely interpolated from measured data);  $N = 360$ . Both  $s_{train}$  and  $s_{test}$  are normalized by removing their mean and dividing by their magnitude before performing this calculation; the raw signal power is not as important in this application as the relative

measured powers at different angles. As training and test data samples are spaced in  $1^\circ$  increments,  $\rho(k)$  is the correlation at an offset of  $k$  degrees.

The offset with the largest correlation is that at which the training and test samples best match, thus it is our estimated bearing of the target (assuming the training data's  $0^\circ$  point corresponds to north). The entire set of correlation values  $\rho(k)$ , a vector of 360 samples, can be used to estimate uncertainty in the correlation calculation; larger values indicate more probable intruder bearings, smaller ones, less probable bearings. Antennas with narrow beam patterns will yield narrower spikes in the correlation vector, due to their increased resolving power, while broader antenna beam patterns give broader correlation vectors—i.e., less certain angle estimates.

The direct correlation technique described above requires  $O(N^3)$  operations to compute the entire  $\rho(k)$  vector. The  $O(N^3)$  technique executes in approximately 0.3 seconds on a 2.5GHz P4 CPU, more than fast enough for our purposes, though faster FFT-based techniques could be used [Oppenheim and Schaffer 1989].

## 5. Experimental Results

[Figure 3] shows Vagi antenna signal strength measurements in  $5^\circ$  increments, where the target was located about 150 feet from the directional antenna. The angles have been adjusted so that the target corresponds to a  $0^\circ$  bearing (making it easier to use the data in later correlations). The graph shows the effects of real-world obstacles, reflections, and distortions on what should be an ideal radiation pattern. These effects prevent pinpointing the exact location of the target using a naïve approach, such as maximum value.

The target was then positioned at a different angle, about 75 feet from the WIDS AP on a small hill. Only a few measurements were taken at this location, shown in [Figure 4]. These measurements simulate the availability of only a few signal strength samples in deployed versions of WIDS.




Figure 3: Vagi signal strength versus angle, June 16 2003. Target is at 0 degrees bearing (up), approximately 150 feet from the antenna.




Figure 4: Vagi signal strength versus angle, June 16 2003. Target is at 330 degrees bearing (upper-left), approximately 75 feet from the antenna.

## 5.1. Correlation Results

Despite the temporal variance, physical obstructions and reflections, and weaker than expected signal strength, the correlation algorithm described in Section 4.1 worked quite well in pinpointing the target machine's bearing. For the grid antenna, the 3°-sampled measurements experiment were used as training data. (In deployed systems, the training data would be an amalgam of measurements taken under a number of conditions.)

Each of the remaining figures shows the correlation value at each candidate angle, with 1° increments;  $N = 360$  in the above equation. As described in Section 4.1, the angle corresponding to the highest correlation value is declared to be the target angle estimate.

[Fig. 5] shows the correlation versus angle for the grid antenna measurements when compared with earlier training data. The antenna was located at a bearing of 25°; our algorithm estimated the angle to be 13°. The figure shows that 13° corresponds to the peak in signal response.




Figure 5: Correlation values of the May 9 grid antenna trial with the April 11 training data; the maximum correlation occurs at 13°.

[Fig. 6] shows the correlation with a later grid antenna trial, again with the earlier training data. The target was estimated to be at 23°, whereas it was actually at 25° (two degrees is within the likely error range of our visual angle measurements). [Fig. 7] shows a grid antenna trial in which the target was estimated to be 37°, while the target was sighted at 40°.




Figure 6: Correlation values of the first June 10 grid antenna trial with the April 11 training data; the maximum correlation occurs at 23°.




Figure 7: Correlation values of the second June 10 grid antenna trial with the April 11 training data; the maximum correlation occurs at 37°.

The Vagi antenna trials are correlated in [Fig. 8] with the first (more detailed) trial as training data, and the second trial's sporadic measurements as test data. The target was sighted at a bearing of  $330^\circ$  and the angle estimate was  $326^\circ$ .

The results of [Fig. 8] show a definitive correlation. It should be noted that even when presented with less than ideal correlation results, as in [Fig. 7], we can locate the target with a high degree of accuracy. WIDS was designed to function in a realistic, noisy environment. These results support that claim.




Figure 8: Correlation values of the June 16 Vagi antenna trials, with the first trial as training data and the second as testing; the maximum correlation occurs at  $326^\circ$ .

## 5.2. Comparison of Antennas

All three antennas, Vagi, grid array, and echo, performed reasonably close to their marketing specifications. In order to best determine the accuracy of the correlation calculations, most of our experiments used the grid array antenna (the most selective). This antenna was physically the largest and most unwieldy; in practical applications, antenna size could be an important factor in terms of robustness, visibility to potential intruders, and cost for accurate positioning.

Cable and connector losses may be problematic as well. The coax cable between the WIDS AP's wireless card and the grid antenna, while having a 10 AWG core, is nearly 15 feet in length and contains two intermediate connectors. The received signal strength measured by the intruder is significantly ( $\sim 20$  dBm) higher than the received signal strength measured by `iwspy` on the WIDS AP (the former being a

direct connection, the latter going through 15 feet of cable). Other variables, however, prevent us from drawing a direct conclusion: different software is used on the intruder to measure signal strength, and the two measurements are of symmetric—but not necessarily identical—signal strengths. In future work we will quantify the effects of cable attenuation.

An additional correlation experiment uses the Vagi antenna. While not as sharply directional, the Vagi is significantly lighter and smaller. Our correlation code does not need a sharp peak in the reception pattern in order to determine the incident angle; what matter are the fluctuations in response through the entire 360° sweep.

The three antennas were roughly equal in cost, approximately \$80 each. Performance characteristics and allowable size will be the main factors that drive the decision of what antenna is most appropriate for an installation. Most likely, no one antenna will be suitable for *all* installations. Overall, the Vagi antenna offers an attractive tradeoff between size/weight and accuracy, with angle estimates nearly as precise as the grid array. The needs of a particular installation, however, will dictate which antenna type to use.

Overall, the antennas' beam patterns match what was expected, with the grid array antenna being more directional than the other two. The patterns were not as ideal as those in the antennas' marketing literature, possibly due to interference and reflections from surrounding objects (particularly cars). To construct final training sets for each antenna in an installation, more detailed measurements under varied conditions would be required. Specific WIDS installations would benefit from training data measured on-site, where many potential obstacles and reflectors are in place.

We were surprised by the fluctuation in signal strength measurements reported by `iwspey`. These fluctuations varied from 2 to up to 8-10 dBm over a few seconds; it seems unlikely that reception conditions change that rapidly in our environment. `iwspey` simply uses values reported by the wireless card driver (Prism, in this case). Our suspicions were confirmed by measurements taken by using the target machine's "Toshiba Client Manager Link Test" software under Windows; the target reported signal strength variance of less than 5 dBm, often 0 (a constant value).

## 6. Conclusions and Future Work

WIDS provides intrusion detection and intruder location for wireless access points. This paper focused on the ability to determine the angle representing the bearing of the remote station to the directional antenna. We were able to get an accurate angle using inexpensive, readily available hardware. By using triangulation on the bearings obtained from two directional antennas at different locations, we accurately determined the location of the remote station.

WIDS includes intrusion detection capabilities in addition to intruder location. Users of WIDS will be able to detect network intruders *before* they connect to the "real" (non-WIDS) access points. In some cases, intruders will be detected when they are attempting to scan for networks (e.g., by using NetStumbler), long before they associate with the omni access point. By using a combination of signature and

behavior based techniques, WIDS can detect a “spoofed” MAC address (i.e., an intruder masquerading as a legitimate user). Since the signals in wireless networks radiate beyond the intended coverage area, intruders beyond the physical perimeter can attack the network. WIDS protects against these attacks.

Currently, WIDS comprises the directional capabilities described in this paper, as well as triangulation location, an IDS component, a protocol for WIDS AP communication, and an interface to control the system. Future work will include taking the proof-of-concept components developed and integrating them into one prototype. Additional features include adding motor controls on the directional antennas and increasing the sophistication of the IDS pattern matching capabilities.

Antenna arrays, an alternative to physically rotating directional antennas, measure the phase difference of signals incident on each element of the array. Due to the temporal resolution required to measure phase differences of 2.4 GHz signals, this is generally done with beamforming techniques. Current work in this field has focused on cellular telephony applications, and is often based on algorithms such as MUSIC [Swindlehurst and Kailath 1992]. Such measurements are impossible within the framework of HostAP and likely require radio-level access to the 802.11 hardware or custom hardware of our own.

## References

[Ellison 2001] Ellison, Craig: “Exploiting and Protecting 802.11b Wireless Networks”; Extreme Tech, 04 Sep 2001, <http://www.extremetech.com/article2/0,1558,11388,00.asp>.

[Foust 2002] Foust, Robert: “Identifying and Tracking Unauthorized 802.11 Cards and Access Points, A Practical Approach”; ;login:, 27(4), August (2002), 32 – 43.

[Hayes] Hayes, Nicki: “Wired Equivalent Privacy (WEP) – Gone in 15 Minutes!”; 2001, <http://www.wirelessdevnet.com/channels/wireless/features/newsbyte31.html>.

[Hofmeyr et al.] Hofmeyr, S. A., S. Forrest, et al. : “Intrusion detection using sequences of system calls”; Journal of Computer Security, 6 (1998), 151-180.

[Information Security] “WEP Security Goes ‘Poof’”; Information Security, 4, 9 (September 2001), 30.

[Malinen] Malinen, Jouni: “Host AP driver for Intersil Prism 2/2.5/3”; 2003, <http://hostap.epitest.fi/>.

[Marceau 2000] Marceau, C.: "Characterizing the Behavior of a Program Using Multiple-Length n-grams"; Proceedings of the New Security Paradigms Workshop, Ballycotton, Ireland (2001), 101–110.

[Oppenheim and Schafer 1989] Oppenheim, Alan V. and Ronald W. Schafer: “Discrete-Time Signal Processing”; Prentice Hall (1989).

[Swindlehurst and Kailath 1992] Swindlehurst, A. L. and T. Kailath: “A Performance Analysis of Subspace-Based Methods in the Presence of Model Errors, Part I: The MUSIC Algorithm” IEEE Transactions on Signal Processing, 40, 7 (July 1992), 1758–1774.

[Wright 2002] Wright, Joshua: "Layer 2 Analysis of WLAN Discovery Applications for Intrusion Detection"; (November 2002), <http://home.jwu.edu/jwright>.

### **Acknowledgements**

This material is based upon work supported by the Naval Surface Warfare Center, Dahlgren Division, under Contract Number N00178-03-C-2010.