

Title: Long Paper: Using Visualization to Locate Rogue Access Points

Authors:

Dino Schweitzer (corresponding author)

Department of Computer Science

United States Air Force Academy CO 80840

(719) 333-3945

dino.schweitzer@usafa.af.mil

Wayne Brown

Department of Computer Science

United States Air Force Academy CO 80840

(719) 333-3590

wayne.brown@usafa.af.mil

Abstract:

Unauthorized access points on a wireless network, known as rogue access points, represent significant security vulnerabilities. Commercial and open source tools are available to detect and locate such devices. Our tool, WiVis, uses interactive visualization to locate rogue access points. Distributed sensors are “profile mapped” to determine how they receive wireless signals from the environment around them. Visual displays of the maps as contour lines are overlaid onto a schematic of the office or lab space. When sensors detect an access point, contour line intersection visually show the predicted location of the device. The tool runs on standard platforms and requires no special hardware.

Keywords:

Rogue access points, Visualization, Wireless security

Using Visualization to Locate Rogue Access Points

Dino Schweitzer
Department of Computer Science
United States Air Force Academy, CO 80840
(719) 333-3945
dino.schweitzer@usafa.af.mil

Wayne Brown
Department of Computer Science
United States Air Force Academy, CO 80840
(719) 333-3590
wayne.brown@usafa.af.mil

ABSTRACT

Unauthorized access points on a wireless network, known as rogue access points, represent significant security vulnerabilities. Commercial and open source tools are available to detect and locate such devices. Our tool, WiVis, uses interactive visualization to locate rogue access points. Distributed sensors are “profile mapped” to determine how they receive wireless signals from the environment around them. Visual displays of the maps as contour lines are overlaid onto a schematic of the office or lab space. When sensors detect an access point, contour line intersection visually show the predicted location of the device. The tool runs on standard platforms and requires no special hardware.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network monitoring - *network security*. K.6.5 [Management of Computing and Information Systems]: Security and Protection.

General Terms

Management, Security.

Keywords

Rogue access points, Visualization, Wireless security

1. INTRODUCTION

Rogue access points are a critical security risk to wireless enterprise networks. A rogue access point is one which has been installed on a secure network without the explicit permission of the appropriate network management authority. While the intent of the unauthorized access point may be malicious, it is more commonly installed by legitimate network users who are trying to extend the flexibility of their office or lab environment. The growth and availability of inexpensive wireless access points for home use has compounded this problem.

The reason that a non-malicious rogue access point represents a security risk is that untrained users are likely to misconfigure them so that they do not conform to the company’s security policy [6,7]. As such, it represents a vulnerability by opening the network to malicious wireless attackers. Additionally, rogue access points can interfere with signals from legitimate Wi-Fi installations. While many corporate IT departments have put in place administrative policies to control installation of equipment on the network, it is difficult to totally eliminate unauthorized equipment such as access points in a large corporate environment. In the case of a malicious user, it is necessary to gain access to an active network port within the facility. However, unless strong access control and search procedures are in place, getting a device into a location for installation is relatively easy. Successful defense against rogue access points must rely on detection through monitoring activities.

2. ROGUE ACCESS POINT DETECTION

Detection is often broken into two phases, discovery of the existence of an unauthorized access point followed by the determination of its location. Detecting a rogue’s existence can be accomplished with different tools and techniques. Wireless sniffing tools, such as AirMagnet or NetStumbler, can be executed on laptops or handhelds capturing information about access points within their range. Any signal detected can be compared to a list of authorized access points via its MAC address, vendor name, or security configurations. The wireless sniffing software will only detect signals within the range of the device it is running on. The monitoring device(s) may be mobile and moved around the facility on a regular basis looking for rogue signals. This may not be practical in a large facility, or one that is not easily navigated. It is also not effective against an intermittent signal that is not present at the time of monitoring.

A variation of the mobile sniffer approach is to run wireless monitoring software on distributed machines on the network and have them report back suspicious signals to a centralized monitoring server. Kismet, an open source initiative, has the capability to operate in this fashion with “drone remotes” (www.kismetwireless.net). This eliminates the need for mobile monitoring, but will only detect signals within the range of the fixed location monitors. The commercial product, AirWave uses a similar approach, but rather than running sniffer software, it runs software on the company’s existing access points listening for rogue signals and feeds the results back to a centralized console. Once again, it will only detect signals within range of the existing access points. Another variation of the distributed monitor approach was the development of customized low-cost sensors at the Air Force

Research Laboratory as part of the WIDS (Wireless Intrusion Detection System) [1]. These customized devices can be placed anywhere in the area to be monitored, detect wireless signals, and report back via the network to a centralized monitor.

A simple technique for detecting rogue access points that does not rely on monitors is to run a port scan on the network looking for Port 80 (HTTP) interfaces which includes all Web servers, some printers, and nearly all access points. Even if an access point's Port 80 interface is disabled or protected with a password, the device will usually respond to a request for some basic information that may be useful in determining its status.

3. LOCATING ROGUE ACCESS POINTS

Once the existence of a potential rogue access point has been determined, the next step is to physically locate it. Locating the source of the wireless signal requires variation in some measurable factor based on the source's physical location. For example, the time it takes for a signal to transmit from the source to the receiver will vary based on distance. In theory, a network monitor can send a request, time how long the response takes, and estimate the distance of the responder based on this time delay. The amount of time necessary to process the request would be subtracted from the total round trip time. The rogue access point would then lie on the circle with distance radius around the monitor. Given data from three monitors, the originating point can be approximated by the intersection of the three circles. Unfortunately, in practice, accurate timing is difficult, signals are inherently noisy, and distance calculations are affected by delays due to walls and environmental conditions. As such, this approach is not practical as a general solution to the problem.

Another measurable variable is the angle of arrival of the signal. While not providing distance, it can give an angle at which the signal arrives at a directional antenna. Several directional antennas appropriately positioned would be necessary to accurately locate the rogue signal source. Once again, accuracy of the measured angle and a noisy environment degrade this approach. In addition, the effect of multi-path signals in a complex environment makes accurate measurement difficult. This approach can be used with other methods to get a general direction of the rogue signal.

A third measurable component is the relative signal strength received. This is the most common measurement used for locating rogue access points. In the mobile sensor method, this can be done by physically walking around while paying attention to signal strengths detected. A good quality directional antenna can assist this method by providing a relative direction to walk in. In the case of the distributed sensor technique, whether existing machines or customized devices, mathematical triangulation can be applied if the signal is received by multiple sensors at different signal strengths. Once again, these calculations are affected by a non-uniform signal distribution due to the environment.

An interesting approach using signal strengths can be adapted from solutions in the literature to the similar problem of locating or tracking a mobile user using RF technology, also known as the localization problem. If the environment knows the physical location of a user, it allows for some interesting location-based applications. For example, sending a document to a printer can automatically determine which printer is closest based on the user's physical location. When a phone call

arrives, only the phone closest to the physical location of the recipient rings. An early example of this problem and solution was the RADAR system developed at Microsoft [3]. The problem is to determine the location of a mobile user in a known space such as an office building to allow for location-aware systems and services. RADAR assumes the mobile user has a wireless receiver, such as a laptop, and the known space has a network of access points at known locations. At any point in time, the receiver can measure the relative signal strengths from the access points, and use a form of triangulation to determine its physical location. Rather than assume ideal transmission conditions, the RADAR system initially calibrated the known space by taking several signal sample points at different physical locations and storing them in a database. When a mobile user wants to know their location, the measured signal strengths are compared to the pre-measured values in the database to determine the physical point that most closely matched. Enhancements to the basic system were later introduced to add user tracking and to simplify the environment profiling step [4]. Other researchers have taken a similar approach with various refinements such as the use of Bayesian inference, improved profiling, and probabilistic techniques [2,5,8,9].

The localization problem is essentially the rogue access point location problem in reverse. That is, rather than locating a signal receiver at an arbitrary location using known signal generators (localization), the goal is to locate a signal generator (the rogue access point) using known signal receivers (monitors running wireless sniffers). This is the approach our system, WiVis, takes using visualization.

4. WIVIS

In Spring 2006, faculty and students in the Information Security class took on a project to use visualization techniques to locate rogue access points once their presence was detected. The basic approach was to use a distributed monitoring system running wireless sniffing software. Once a rogue was detected, the sensors would feed the received signal strength back to a central server. The key to locating these devices uses the same principle as the localization problem, that is, to have the distributed sensors "profiled" ahead of time to map how they receive signals based on the environment around them. This profile map is similar to mapping the signal strength from an access point for the localization problem. The difference is that the profile map is from the point of view of the wireless signal receiver not the access point generating the signal. In the ideal case, different signal strengths would be concentric circles, or contour lines surrounding the sensor. In the real world, signal strengths are affected by several environmental factors. Another major difference from the RADAR approach is the profile maps are computed as points on a fixed grid as opposed to individual samples stored in a database. This allows the map to be displayed as a raster image. Figure 1 shows four different views of a sample profile map generated for a wireless monitor set in a classroom (4H47);. Each point in the image has an associated signal strength that represents what the sensor would receive if a signal was sent from that location. The different views shown in the figure display different techniques for viewing the signal strength data. In the upper left image a contour map is drawn of signal strengths. Every line represents a single signal strength as received by the sensor. As can be seen, the contours center

around the sensor located in 4H47. The upper right image is the same data, but uses color banding to represent ranges of data. The lower right uses a continuous color ramp to represent the same information. Finally, the lower right simply uses a random colormap to show all of the data values. This view is useful in that it is easy to quickly see all points with the same signal strength (same color).

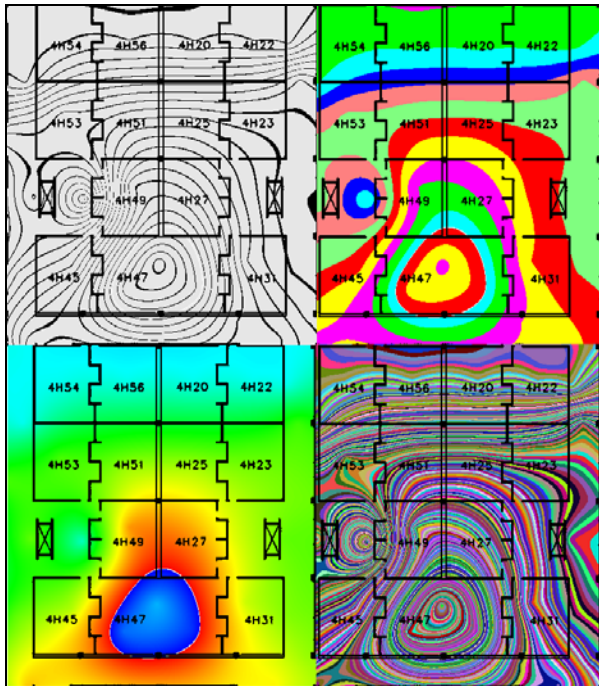


Figure 1. Four views of a profile map for a wireless network sensor: contour lines, color bands, continuous color ramp, and random colormap.

Once the profile map is created, further signals received from a different access point will (in theory) lie in the same area as that indicated by the signal strength profile map. Thus, in the upper right image in Figure 1, if the sensor received a future signal, possibly from a rogue access point, with a strength equivalent to the color red, then the location of the signal generator should physically lie on the red band of the map. To facilitate the isolation of a single signal strength for viewing, WiVis incorporates interactive sliders to allow the user to select a narrow signal range of signals to view as shown in Figure 2. The gray band represents the portion of the map containing a specified range of values. Any signal received that lies within that range should be physically located at those locations on the floorplan. The WiVis user interface allows the user to select both the median value and range of values to display. Signal ranges for multiple sensors can be viewed simultaneously as shown in Figure 3. Here two different ranges from two different sensors are overlaid on the floorplan. If sensor 1 received a signal that indicated the originating source was in the darker band, and sensor 2 received a signal from the same source (as determined from MAC address) that indicated it was in the lighter band, then the physical location of the originating source would be where the two band intersected, or in room 4F15.




Figure 2. Signal range for a single sensor.




Figure 3. Intersecting signal ranges from two sensors.

4.1 Creating Profile Maps

Profile maps are necessary for WiVis to display likely locations for a received signal. If static monitors are used, such as existing desktop machines running a wireless sniffer program, and the environment does not change significantly, the maps should remain fairly consistent. Our experience to date is that the minor changes in signal strength due to different weather conditions, such as humidity, and number of people in the area had only a minor effect on the generated profile map. For the level of accuracy necessary to detect a rogue access point, knowing the location within a given room is generally sufficient. With that assumption, profile maps only need to be created once. Note that this assumption is under scrutiny and requires further verification. For example, the impact of other devices in the vicinity and other factors needs to be further explored.

The map is essentially a two-dimensional array of data values representing signal strengths at each grid location. The grid corresponds to physical locations within the space being mapped. The process of filling in the grid is to measure some number of sample signal strengths at known points and then use data interpolation techniques to fill in the remaining grid points.

The intuitive way to collect sample points is to put a sensor at the desired location, and then walk around with an access point, shooting signals at known locations on the grid and recording the signal strength received by the sensor. To simplify and speed up the process, we did it the opposite way similar to the way profiles were created in RADAR. We placed an access point at the sensor location, and walked around with a laptop running a sniffer program. We simply recorded the signal strength as measured in dBs received at different locations on a map of the area, and translated those map coordinates into the appropriate grid locations. The rationale for using this approach is that the signal strength between a signal generator and signal receiver located at two points does not depend on which point is the generator and which is the receiver. Using this approach, we were able to quickly collect

the 70 sample points used to generate the profile map shown in Figure 1.

Once the data samples are collected, the next step is to interpolate the data to fill in the entire grid. We tried different interpolation techniques using MatLab. The floor plan in Figure 4 shows the data collected for one wireless access point. The value next to each X is the detected signal strength.




Figure 4. Sample points collected for building profile map.




Figure 5. Different interpolation schemes in Matlab: nearest, linear, cubic, V4.

These sampled values can be interpolated into discrete signal strengths for every location on the floor plan by standard interpolation techniques. We used MATLAB's `griddata()` function to investigate various interpolation methods. Figure 5 shows the four possible interpolation schemes used by the `griddata()` function. The top three techniques (nearest, linear, and cubic) are based on a Delaunay triangulation of the data. The linear and cubic methods restrict their interpolation values to points inside the convex hull of the original data points, while the nearest and v4 techniques project their interpolations to the boundaries of the image.

Nearest neighbor interpolation would only work with a large number of samples, otherwise it would not provide the expected continuity of an RF signal. Linear interpolation may be sufficient for estimation purposes, but tended to fall off sharply at the edges. Cubic interpolation had the unintended side effect of overshooting and giving negative results. The V4 method (biharmonic spline interpolation) gave the best visual results. Figure 6 shows the three dimensional surface of the signal strength as a height field over the grid.

The accuracy of the interpolation varies with the number and location of sample points. While a complete quantitative analysis has not been completed, Figure 7 gives an indication of the relative change in the profile map when interpolating 20, 40, and 60 sample points. In this example, the points were chosen randomly. In practice, care would be taken to select sample points that are representative of the environment, such as one per room.




Figure 6. 3D surface view of V4 interpolation.




Figure 7. Interpolation changes with 20, 40, and 60 sample points.

4.2 Locating Access Points with WiVis

Once the profile maps are created for all of the sensors on the system, detected signal strengths of potential rogue access points can be sent back to a central server for WiVis processing. WiVis allows the user to load in a bilevel background map of the area of interest for reference. Profile maps are then loaded for the sensors of interest. Currently, data from up to three sensors can be displayed simultaneously. Different views of the data, as shown in Figure 1, are done with colormap manipulation affecting the appropriate bits of the image. The process of locating a rogue access point is currently a manual process. The user sets the appropriate signal range for each of the sensors being displayed based on the signal strength reported from the sensor, as shown in Figure 3. The rogue access point, in theory, lies in the intersection of the contours. Future plans include automating this feature of displaying the appropriate signal band based on a suspected rogue signal.

One assumption that was made in the creation of the profile maps was a fixed signal strength of the generator. When dealing with a rogue access point, the originating signal strength will not be known, and may not lie on the expected contour from the contour map. Thus, this technique will not work if a single sensor picks up the rogue signal. However, if two or

more sensors pick up the signal, then the expected intersection can be displayed as shown in Figure 3. To accommodate a stronger or weaker originating signal, WiVis provides a "signal strength" slider that moves each of the displayed contour lines at the same rate. The rationale for this is that the ratio of signals received by the multiple sensors will stay the same regardless of the originating strength. The actual location is determined by where the multiple contours intersect as the strength slider is adjusted.

5. Results and Future Plans

WiVis is currently in prototype to demonstrate proof of concept and is being used to collect data for accuracy evaluation and analysis. We have used it in a couple of different environments with up to seven profiled sensors. Preliminary results indicate that it does locate previously unknown access points to the room level. Several open questions remain that require additional data collection and analysis. For example, the number of sample points necessary to accurately profile a sensor has not been quantitatively determined. Similarly, additional study is necessary to determine what environmental factors cause the profile map to change and how sensitive rogue detection is to those changes. For the localization problem described earlier, some researchers have looked into ways of dynamically changing profile maps based on measured conditions. These techniques may apply to this approach as well. Additional experiments also need to be conducted with rogue access points at different signal strengths to measure the accuracy and sensitivity of the "signal strength slider" approach. Another area for study is how the accuracy and utility of this approach compares to other rogue access locating methods. Finally, a comprehensive user study needs to be completed of users responsible for network security actually using it in a live environment.

In addition to the expanded data collection and analysis, other refinements are planned for the tool. As stated earlier, the current prototype operates in a standalone mode. That is, it is not directly connected to the remote sensors, but relies on a manual technique for entering signal strengths of potential rogue devices. The desired system configuration is to have WiVis operate on a central server with remote sensor nodes reporting in wireless activity and signal strengths for processing. This is similar to the capability in Kismet. The details of how often sensors sample the wireless environment and report back, and the communication interface is yet to be determined. The capability of distinguishing unauthorized access points from legitimate ones will need to be incorporated in the tool. Once a networked version of the tool is operational, extensive tests of usability and scalability will be conducted. For example, how many sensors are necessary to adequately cover the entire floor of a building, and how well does the tool handle a large number of profile maps.

While results are still preliminary, we are encouraged by the relative simplicity and apparent accuracy of the results. Profile maps can be quickly generated by placing access points at sensor locations and walking around with a receiver that can capture several signals at each sample point with the single click of a button. The visual approach is easy to understand and allows the user to quickly identify the approximate location of a suspected signal.

6. ACKNOWLEDGEMENTS

Several individuals worked on this project through the course of the semester. Special recognition goes to two students in the Information Security class, Cadet Chris Woodward and Cadet Fernando Niclolade, who worked in various capacities including data collection, background research, interpolation testing, and data analysis. Additional thanks go to other department faculty members who offered their wireless and security expertise to this effort.

7. REFERENCES

- [1] Adelstein, F., Alla, P., Joyce, R., Richard, G.G., III, Physically locating wireless intruders, Information Technology: Coding and Computing, ITCC 2004, (April 2004), vol. 1, pp. 482-489.
- [2] Andreas Haeberlen , Eliot Flannery , Andrew M. Ladd , Algis Rudys , Dan S. Wallach , Lydia E. Kavraki, Practical robust localization over large-scale 802.11 wireless networks, Proceedings of the 10th annual international conference on Mobile computing and networking, September 26-October 01, 2004.
- [3] Bahl, P., and Padmanabhan, V. N. Radar: An In-Building RF-based User Location and Tracking System. In IEEE Infocom 2000 (March 2000), vol. 2, pp. 775--784.
- [4] Bahl, P., Padmanabhan, V. N., and Balachandran, A. Enhancements to the RADAR User Location and Tracking System. Tech. Rep. MSR-TR-00-12, Microsoft Research, February 2000.
- [5] Gwon, Y., Jain, R., and Kawahara, T. Robust Indoor Location Estimation of Stationary and Mobile Users. In IEEE Infocom (March 2004).
- [6] Kozup, C. Detecting rogue wireless LAN access points. techupdate.zdnet.com/techupdate/stories/main/0,14179,2914417,00.html. (Aug. 2003).
- [7] Nutter, R. Detecting rogue access points on campus. www.networkworld.com/columnists/2005/062005nutter.html. (Jun. 2005).
- [8] Paul Castro , Patrick Chiu , Ted Kremenek , Richard R. Muntz, A Probabilistic Room Location Service for Wireless Networked Environments, Proceedings of the 3rd international conference on Ubiquitous Computing, p.18-34, September 30-October 02, 2001.
- [9] Teemu Roos , Petri Myllymäki , Henry Tirri, A Statistical Modeling Approach to Location Estimation, IEEE Transactions on Mobile Computing, v.1 n.1, p.59-69, January 2002.