

Device Identification via Analog Signal Fingerprinting: A Matched Filter Approach

Ryan M. Gerdes, Thomas E. Daniels, Mani Mina, Steve F. Russell
Department of Electrical and Computer Engineering
Iowa State University
Ames, IA 50011
{rgerdes,daniels,mmina,sfr}@iastate.edu

Abstract

As part of the Detecting Intrusions at Layer ONE (DILON) project, we show that Ethernet devices can be uniquely identified and tracked—using as few as 25 Ethernet frames—by analyzing variations in their analog signal caused by hardware and manufacturing inconsistencies. An optimal detector, the matched filter, is utilized to create signal profiles, which aid in identifying the device the signal originated from. Several non-traditional applications of the filter are presented in order to improve its ability to discriminate between signals from seemingly identical devices of the same manufacturing lot. The experimental results of applying these filters to three different models of Ethernet cards, totaling 16 devices, are presented and discussed.

Important applications of this technology include intrusion detection (discovering node impersonation and network tampering), authentication (preventing unauthorized access to the physical network), forensic data collection (tying a physical device to a specific network incident), and assurance monitoring (determining whether a device will or is in the process of failing).

1. Introduction

1.1. Network access control

Current network access control (NAC) mechanisms rely exclusively on the use of digital tokens or identifiers—usernames and passwords, MAC addresses, SSL certificates, WEP/WPA keys, etc—to prevent unauthorized access. Unfortunately, even strong tokens and identifiers, such as SSL certificates, by their purely digital nature, can be discretely copied if improperly secured, and put to use by malicious users. Even worse, popular weak identifiers, such as MAC addresses, may be easily obtained through pas-

sive network monitoring, and spoofed through the use of a programmable network card. In contrast, the analog characteristics of a device are nearly impossible to obtain (a measurement cannot be done without physical access to the medium) and duplicate, which makes them well-suited for NAC purposes.

In the digital age, the physical layer is often regarded as a security impediment, or, at best, overlooked as a source of solutions for today's security needs, because of its non-digital nature. The instinctive reaction to the physical layer has been to focus on securing the layers above it, through the use of encryption, so that some level of authentication is necessary for access to it. These methods often prove intrusive to the end-user; forcing them to remember forever-changing and arcane keys, configure troublesome access clients, or keep track of yet another access token. Clearly, a non-intrusive method, which compliments existing access control methods, is needed to control access to the network infrastructure. We believe that DILON technology can fulfill this need.

1.2. The DILON concept

The DILON project investigates the use of analog and digital characteristics of digital devices for such security purposes as intrusion detection, authentication, forensic data collection, and assurance monitoring. DILON is founded upon the belief that hardware and manufacturing inconsistencies cause minute and unique variations in the signaling behavior of every digital device; furthermore, these variations are manifest by use of the appropriate signal processing technique(s). Central to the security of this concept is the belief that these slight variations are difficult, if not impossible, to control and duplicate. This assumption is founded upon knowledge of the variable tolerances of device components, which are introduced in the design and fabrication processes, used in the construction of digital devices. These tolerances allow for unpredictable variations

in the overall electrical operation of the device. Simply put, because of these variations, no two devices may be made exactly the same, and hence their analog signal characteristics cannot be made the same, without substantial reverse-engineering beyond the reach of all but the most determined attackers.

Figure 1 presents a system-level diagram for an implementation of DILON technology. On the top of the diagram are subject devices that communicate over a physical medium—wired or unwired—to connect with a controlled device, a switch or access point for instance. At the control device an analog tap is used in conjunction with an analog-to-digital converter (ADC) to sample the electrical signals arriving across the medium, at a much higher rate and with greater resolution than is necessary to actually decode the signal. Storage will also be required for past and present fingerprints. A policy engine will make use of a comparison module to determine which devices have access to the network, as well as issue reports concerning the state of the network.

The present approach for DILON focuses on making use of a matched filter to create profiles of signals that are useful in identifying the device the signal originated from. We have found that a traditional matched filter is sensitive enough to easily discriminate between signals produced by different model Ethernet cards. Using advanced techniques, a matched filter, applied in non-traditional ways, can be made to discriminate between Ethernet cards of the same model—even when each component of these cards possesses the same serial numbers, and appear to come from the same manufacturing lot. We have also developed adaptive methods that accurately track fluctuations in signals due to device aging, voltage variations, and temperature changes. These methods provide realistic and consistent false-accept and false-reject rates (FAR and FRR).

1.3. Previous work

Signal detection and identification was one of the major challenges in the research and development of radar and wireless communication systems for a greater part of the 20th century. In particular, identification of radar, radios, and various wireless communications became a very important and popular topic around the time of World War II [13]. Most methods developed for radar identification at this time were based upon transient analysis. As higher frequency and faster responding circuits were introduced, more in-depth transient analysis became necessary for transmitter identification. To this day, many researchers are making use of transient methods for the identification of modern transmitters [3, 25, 1, 9, 10, 4, 17, 5]. However, these methods have only proven successful in situations when the transmitters under consideration were considerably different.

To date, a robust, reliable, and adoptable system for transmitter characterization has yet to be devised to effectively handle multiple transmitters in interconnected systems. While frequency based classification models have been suggested [11, 12, 15, 14, 20], and other general rules for identification have been suggested [6, 24], each is limited to discriminating between different brands and systems. As traditional methods cannot adequately identify similar devices, they will not be able to guarantee the privacy, security, and integrity of sensitive information necessary to medical, legal, governmental, and security management firms.

It should be noted that a similar problem was addressed by cellular phone companies to combat cloning [23, 19, 18]. However, due to propriety nature of their work, there is very little published on their methodology. From what can be determined from the limited literature available, these methods do not have a high success rate in discriminating signals from similar sources.

Recently, work in the development of physical authentication schemes has led to the creation of a physical token that implements a physical one-way function, which is verified using a statistical hashing algorithm [22]. Our work is different from [22] in that we rely on the inherent physical variation introduced as part of the manufacturing process, and do not require extra variation to be explicitly added to the devices for such purposes.

A more closely related physical authentication system was introduced in [8, 7]. Gassend *et. al* investigated the identification of integrated circuits based upon the indirect measurement of their timing characteristics. In contrast, our method focuses on examining the spectral characteristics across the operating bandwidth of the device. Additionally, our work shows that the signaling characteristics of network devices appear to be more amenable to identification than integrated circuits, as we have been able to identify a greater number of devices.

Finally, recent work has investigated the possibility of remotely fingerprinting devices over the Internet by measuring their clock skew [16]. This method shows promise; however, accurate identification seems to require 36 hours of observation, where packets are received from the remote host at a rate of 46 packets per hour. The efficacy of this method is difficult to measure, as the authors do not report their results in terms of false-reject and false-accept rates.

2. Background

The concepts of systems, signals, filtering, and related terminology and tools are discussed. The matched filter operation is defined.

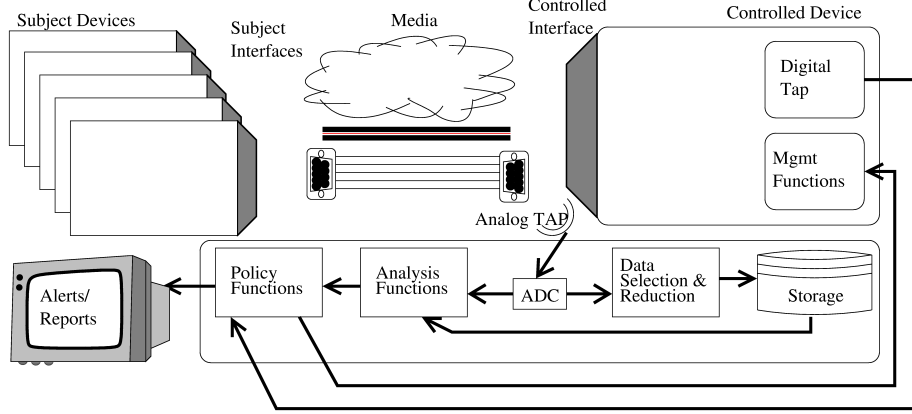


Figure 1. Functional view of a NAC system incorporating DILON technology.

2.1. Signals and systems

A system is a process by which an input signal is transformed to produce an output signal; furthermore, a system is said to be linear time-invariant (LTI) if the system is both additive and multiplicative, and a time shift of the input results in a corresponding time shift in the output. It can be shown that the response, or output, of an LTI system to all inputs can be completely described by determining the systems unit impulse response [21]. For our purposes, the unit impulse response, or transfer function, of a system in the time domain will be denoted by $h(t)$. The response of a system, $y(t)$, to a particular input, $x(t)$, can be found by convolving the transfer function of the system with the input signal. The convolution operation, denoted by \star , between $h(t)$ and $x(t)$ is defined as:

$$y(t) = h(t) \star x(t) = \int_{-\infty}^{+\infty} h(t - \tau)x(\tau)d\tau \quad (1)$$

By taking the Fourier Transform of the input signal, denoted by $\mathcal{F}\{x(t)\} = X(\omega)$, and the transfer function, denoted by $\mathcal{F}\{h(t)\} = H(\omega)$, the convolution operation defined in (1) may be replaced by multiplication:

$$Y(\omega) = H(\omega)X(\omega) \quad (2)$$

It should be noted that (2) gives the frequency response of a system, whereas (1) gives its time-domain response. Of course, these responses are related through the inverse-Fourier and Fourier Transforms, respectively.

2.2. Filters

A filter may be regarded as a special kind of system, where the relative amplitudes and phases of the frequency

components of an input signal are modified, or eliminated. As the filter discussed in this paper is LTI, we may describe its response via a transfer function. In turn, this transfer function may be used in conjunction with either (1 or 2) to determine the response of the filter to an input signal.

2.3. The matched filter

The matched filter is said to be an optimal detector, as it can be shown that the filter maximizes the signal-to-noise ratio of a known input signal in additive white Gaussian noise (AWGN). [2]. The transfer function of the matched filter, in the frequency domain, at sampling time t_0 may be stated as:

$$H(\omega) = \kappa \frac{A^*(\omega)}{P(\omega)} \exp^{-j\omega t_0} \quad (3)$$

where $A^*(\omega)$ is the complex conjugate of the Fourier Transform of a known time-domain signal $\alpha(t)$, $P(\omega)$ is the *power spectral density* (PSD) of the noise associated with an input signal, and κ is an arbitrary constant. By selecting an appropriate value of κ for the operating environment, and assuming AWGN for the PSD, $P(\omega)$ may be eliminated from (3). For a given input signal, $\beta(t)$, the output of the filter, M_{t_0} , at sampling time t_0 , in the Gaussian noise case is then:

$$M_{t_0} = H(\omega)B(\omega) = A^*(\omega) \exp^{-j\omega t_0} B(\omega) \quad (4)$$

where $B(\omega)$ is the Fourier Transform of the time-domain input signal $\beta(t)$.

Taking the inverse Fourier Transform of (3) gives the transfer function of the filter, $h(t)$, in the time-domain, for the AWGN case, as:

$$h(t) = \alpha(t_0 - t) \quad (5)$$

It can be shown that the output of the filter is maximized

when:

$$\mu(t_0) = h(t_0) \star \beta(t_0) = \int_{t_0-T}^{t_0} \alpha(\tau)\beta(\tau)d\tau \quad (6)$$

where T is the period of the known time-domain signal $\alpha(t)$.

As can be seen from (6), the matched filter operation may be interpreted as the inner-product of two signals, or an integrated-correlation.

3. Signal identification

We describe how the matched filter may be used to create a signal profile useful for identifying a signal’s device of origin.

3.1. Signal selection rationale

This work focuses on the profiling of 10Mb wired Ethernet signals. We chose to study 10Mb Ethernet because of the relative simplicity of the electronic devices and signaling involved, and its operation at low speeds. As the electronics and signaling are less complicated than higher-speed systems, we were able to understand the functioning of the devices, and identify common behavior between devices of different makes, which aided us in hypothesis creation and testing while attempting to identify differences and similarities in signals. In addition, capturing accurate samples of 10Mb Ethernet frames may be accomplished using lower resolution, slower, and therefore less expensive ADCs.

Wired Ethernet was chosen due to the low noise environment inherent in wired systems. Environmental noise adds a stochastic and non-stationary component to the signal that must be minimized as much as possible to obtain consistent measurements. On the other hand, noise characteristics of an individual device, or component from a device, may exhibit distinguishing characteristics.

Finally, we believed that if we should fail in discriminating 10Mb Ethernet signals, we would have little chance of succeeding in the high-speed wired and wireless domains. However, we should also note that in some respects profiling 10Mb Ethernet signals may be viewed as a more difficult problem than that of higher-speed systems: there are fewer components per device, and hence less opportunity for signal variability due to perturbation by device components.

3.2. Identifying a common signal

In order to create a profile of the signal characteristics for an Ethernet device, a portion of the frame preamble

common to all devices was identified. At the beginning of each frame a 64-bit sequence of alternating ones and zeros, encoded using differential Manchester encoding with a fundamental frequency of 5MHz, ending with the sequence *10101011* are sent to synchronize the receiver of the destination device to the transmitter of the source device (Figure 2).

This synchronization signal consists of a transient, or turn-on, portion (denoted by ‘- . . .’ in Figure 2), which is the result of the transmitting circuitry of the sending device powering on, as well as a steady state portion (denoted by ‘- - - -’ in Figure 2) that serves as the actual synchronization signal.

As mentioned earlier, most work in signal identification has traditionally focused on the transient portion of a signal. However, as the transient signal in 10Mb Ethernet is so small, in terms of the number of wavelengths of the overall signal, we do not believe that there is physically enough information contained in it for the identification of similar devices. Indeed, it has been shown in the literature that transient analysis is sufficient only for distinguishing between devices of different models, but not devices of the same model. As such, our methodology relies primarily upon the steady-state portion of the signal for profiling purposes.

The final portion of the Ethernet frame shown in Figure 2 (denoted by ‘. . . .’) is the beginning of the MAC address of the receiving device. Preliminary work with this portion of the signal has shown that it may be possible to use the MAC source address for signal profiling.

3.3. Matched filter creation

Having identified a common and repetitive portion of the Ethernet signal suitable for identification purposes, an exact starting position and period of the portion of the signal to be matched to must be chosen. We call this part of the signal the *reference signal*, and choose it to represent the known time-domain signal $\alpha(t)$. As per (5), the reference signal must be reversed in the time-domain, and shifted by t_0 to be used as the filter. In this respect t_0 may be regarded as the final time point of the reference signal.

Initially, the period and position of the reference signal were chosen as an arbitrary number of points spanning the length of the synchronization signal. For 10Mb Ethernet, we have found this acceptable to distinguish between all but the most similar of signals; however, we have also developed algorithms to determine the optimal reference signal for a set of known devices. This type of reference selection would be useful during a training period, where sample data could be taken for a new device introduced on the network, and compared to previously collected data of other devices. For a general study of the matched filter, however, we have selected a reference signal that includes the preamble tran-




Figure 2. The Preamble of an Ethernet frame used for signal profiling.

sient and steady-state portion of the synchronization signal, which is the same, to within five sample points, for each device. While optimally determining a reference signal for a device, in relation to other known devices, may increase performance, our experiments have shown that it is not generally necessary to do so.

3.4. Signal profile

The first step in creating a signal profile is to apply the filter to the signal used to create it; i.e., convolve the filter with the portion of the signal used for the selection of the reference signal. The filter returns a single value from this operation that serves as a baseline. This value represents the filter response when a perfect match is made between the filter and the original signal. If another signal is exactly the same as the original, then we expect that applying the filter to this signal will produce the same value. In general then, applying the filter to a signal produces a measurement of the closeness of the signal to the original, and consequently the alikeness of the devices the signals were acquired from. If a signal from a different device approaches the filter output value for the original signal too closely then we are unable to distinguish it from the device that produced the original.

Due to the noise inherent in any system, we cannot assume that even a properly functioning device will output exactly the same synchronization signal for each frame. Noise from surrounding devices, created by a hard disc or CD-ROM being accessed or variations in system load, and thermal noise assuredly cause slight variations in the signal from frame to frame. In fact, with the aid of temperature recording equipment we have been able to correlate aberrations in the filter output to variations in the ambient temperature of the lab. Furthermore, due to the non-ideal properties of the Ethernet cabling—parasitic resistance, ca-

pacitance, and inductance—even the act of measuring the signal on a different portion of the Ethernet cabling, or using a different cable altogether, may affect the measured signal. This affect, however, gives rise to the interesting possibility of detecting passive taps on the line, which often change the effective material parameters of the medium.

To take into account the inherent variability of every device’s output, as well as external factors such as temperature and system load variations, a signal profile must be created by using a collection of signals taken over a period of time. The filter created by the original signal is applied to this collection of signals and the response to each recorded (Figure 3). We have found that only 25 sequentially sampled signals are necessary to adequately determine the unique signaling behavior of a device.




Figure 3. Filter output for 25 frames of an Ethernet device.

By examining the filter response for a device over a number of hours, we have determined that a device’s synchronization signal is under continuous change. In many cases, we have discovered that slight variations in the amplitude of the signals are the cause of this variation. A subtle change in signal shape, over a period of hours, also changes the

filter response. By using the average of several synchronization signals for the reference signal we have been able to decrease the variation of the filter response; however, this often leads to a corresponding increase in the FAR.

In order to account for these changes in signal characteristics over time, we have introduced a tolerance, δ , for the maximum amount of deviation in filter response acceptable before a signal is labeled as too different from the original. In order to take into account past behavior, we require that the next n -frames resemble the previous n -frames, $\pm\delta$. In this way a device may be adaptively tracked as its signal changes over time. Mathematically, this is stated by defining two thresholds for the maximum amount of positive, th_+ , and negative, th_- , deviation in filter output allowed over the previous n -frames:

$$\begin{aligned} th_+(\mu_i \cdots \mu_{i+n-1}) &= \max(\mu_{i-1} \cdots \mu_{i-n})(1 + \delta) \\ th_-(\mu_i \cdots \mu_{i+n-1}) &= \min(\mu_{i-1} \cdots \mu_{i-n})(1 - \delta) \end{aligned} \quad (7)$$

where μ_i represents the filter output of the i^{th} frame. We have found that setting n equal to the number of samples used to learn the behavior of a device proves adequate for tracking the signal over time.

During our experiments the filter output for the first 25 frames of a device were used as training data, whereby an appropriate value for δ would be determined by stipulating that zero false-rejects would occur for the next 25 filter outputs. A minimum value of .001 for δ was imposed, and incremented by .001 until the aforementioned condition was met. After observing a device's behavior over time the value of δ can be adjusted to better fit the unique behavior of the card. We have also found that large, spurious, deviations do occur for all Ethernet devices, so a perfect acceptance rate cannot be obtained—unless one is willing to allow a certain number of significant deviations every n -frames, or set δ unreasonably high. As with any system with statistical variation, a balance must be found for δ that results in an acceptable number of false-accepts and false-rejects.

3.5. Variations on the matched filter method

To further improve the efficacy of our method, we have devised several variations on the procedure outlined above to improve our ability to discriminate between highly similar devices. Each of these techniques works to amplify slight differences in signal characteristics that are too subtle to be distinguished by the original method. The impetus of this work was based upon the observation that as the matched filter operation is a sum of products, large-scale similarities between signals can often overshadow the small-scale differences useful for signal profiling.

3.5.1. An ensemble of filters. For a given device, multiple matched filters may be created by selecting a refer-

ence signal for each portion of the preamble identified in Section 3.2. Matching filters to the transient, steady-state, and source MAC address sections of the frame gives a full characterization of the broad traits of a signal. An ensemble of filters is utilized, instead of a single large filter, so that strong similarities in certain regions of the signal cannot overshadow smaller differences in others.

Selecting multiple reference signals for each section of the signal may also highlight slight differences; e.g., each transition, or pair of transitions, of the synchronization signal could be matched to different filters. In such a way the granularity of filtering could be arbitrarily increased to take into account the smallest of differences.

3.5.2. Bandpass filtering. By analyzing the spectrum of signals from a multitude of similar devices, we have found that distinguishable differences exist in the frequencies beyond the fundamental frequency of the synchronization signal; however, as the fundamental frequency dominates other frequency components, in terms of relative power, these differences are often obscured. Applying a bandpass filter to the reference signal and signal samples minimizes the influence of the fundamental frequency on the filter response by removing that portion of the signal altogether.

Through experimentation, by use of several bandpass filters with increments of 1MHz in bandwidth, we have determined that, for some devices, the 13-17MHz frequency range exhibits the greatest variation. As the power levels of frequency components beyond 17MHz approach that of the noise level, we have found frequencies higher than 17MHz ill-suited for discriminatory purposes.

3.5.3. Normalization. Normalizing both the reference signal and signal samples, according to the Euclidean norm, desensitizes the matched filter to similarities in shape, and increases its sensitivity to variations in amplitude. This proves advantageous for discriminating between signals where the differences exist primarily in their relative amplitudes. However, if the amplitudes of two signals are closely matched, while their shapes are not, this form of normalization will decrease our ability to distinguish between the two.

3.5.4. Trimming. The concept of time-domain trimming was developed in order to minimize the affect of the signal amplitude on filter response. By eliminating amplitude dominance, variations in the shape of the signal are made apparent. Analogous to the frequency domain trimming used in bandpass filtering, time-domain trimming removes the portions of a signal that tend to overshadow all others by zeroing the signal amplitude for values greater than a predetermined upper bound. By adding a lower bound, and varying the height of each boundary accordingly, a window

is created that allows for any portion of the signal to be scrutinized by its shape alone.

For example, by only setting an upper bound, the zero-crossings—where the signal crosses the horizontal axis—of a signal may be examined in order to ensure that the width of a signal matches that of the filter. We have found that time-domain trimming is most effective when only the signal samples are trimmed.

4. Experimental results

The equipment and methods used to acquire the Ethernet signals for analysis are given. Methods for calculating the FAR and FRR are discussed. Finally, the results of the matched filter approach to signal profiling are given.

4.1. Experimental setup

Our current testbed consists of two PCs running GNU Linux; one to act as the Test PC (TPC), which houses the Ethernet card we wish to fingerprint, while the other, the Data Acquisition PC (DAQPC), makes use of a Tektronix 3052 digital sampling oscilloscope, interfaced via an IEEE 488 card and Labview-6, connected to a passively tapped internal Ethernet card, to capture Ethernet frames sent to it over a crossover cable by the TPC.

In order to generate traffic for the DAQPC to capture, the TPC is instructed to ping the DAQPC. During a typical data acquisition period the TPC will ping the DAQPC 10,000 times. To ensure that only traffic from the TPC is captured, only the receiving pins of the DAQPC’s Ethernet card have been connected to the oscilloscope. In this way we are able to allow the DAQPC to respond to the TPC’s pings, and ensure that the data acquisition process hasn’t caused any packet loss.

Upon detection of an Ethernet frame the oscilloscope begins to sample the signal at a rate of 1Gigasamples/s. The signal is sampled 10,000 times, for a total of 10 microseconds, with 8-bits of resolution. The data collected during sampling is sent to the DAQPC via the IEEE 448 interface, where a custom Labview routine monitoring the interface accepts the data and stores the values in a vector we call a record, which is subsequently written to the disc. Each captured frame is stored in its own record; all of the records collected for a device during a session encompass its dataset.

4.2. Filter application

Having acquired several thousand signal samples from each device over a number of hours, we then create a filter for each of the devices using the procedure outlined in Section 3.3. The reference signal for each device has a period of 4,176 sample points, and is selected from the first

valid record of a device’s dataset. Following this, the reference signal is convolved with each record of its dataset using an FFT-based convolution algorithm. Convolution of the reference signal with each record of its dataset performs the matched filter operation for all possible time-shifts; consequently, an output is created that is equal in length to that of the length of the record. This operation is necessary to determine the time of optimal alignment, t_0 , between the filter and the record, which results in the maximum filter output, as per (6).

Thus, the filter output at the point of maximum alignment corresponds to the maximum of the convolution operation. Letting $\epsilon_i(t)$ represent the reference signal for the i^{th} device, and $\eta_i^j(t)$ the j^{th} record of its dataset, the filter output, $\mu_i^j(t_0)$, is then:

$$\mu_i^j(t_0) = \max(\epsilon_i(t) \star \eta_i^j(t)) \text{ for } j = 1 \dots n \quad (8)$$

where n is the number of records in the device’s dataset (Figure 4). This procedure is followed for each device in order to determine the filter response of each record in its dataset.

Having determined the filter output for each record of its own dataset, we then apply the filter to each record of the other device’s datasets in order to determine the likeness of their respective signals (Figure 5). Letting $\gamma_{i,k}^j(t_0)$ represent the filter output using the i^{th} device’s filter applied to the k^{th} device’s dataset:

$$\gamma_{i,k}^j(t_0) = \max(\epsilon_i(t) \star \eta_k^j(t)) \text{ for } j = 1 \dots n \quad (9)$$

As can be seen from Figure 5, the respective filter outputs of *Device i* and *Device k* do not overlap. Following the explanation set forth in Section 3.4., we are therefore able to discriminate between *Device i* and *Device k*.

4.3. Acceptance testing

Following the procedure set forth in Section 3.4., a value for δ can be determined that is expected to provide an acceptable FRR (less than .009 in our experiments). Using the response of the filter for the i^{th} device to the 26th through 50th records of its own dataset, $\mu_i^{26 \dots 50}(t_0)$, as training data in conjunction with (7), thresholds can be established for the next 25 filter outputs. If the filter response for one of the next 25 records lies outside of the bounds set by these thresholds then its corresponding record is marked as rejected, and is not used in determining the thresholds for the next 25 outputs. This procedure is followed for the remainder of the filter responses in the device’s dataset. The FRR is then calculated using:

$$FRR = \frac{n_r}{n - 25} \quad (10)$$




Figure 4. Filter output for 10,000 records of an Ethernet device.




Figure 5. Filter output for 10,000 records of two different Ethernet devices using the same filter.

where n_r is the number of rejected records and n is the number of total records.

4.4. Intrusion testing

Whereas it is possible to determine the FRR by sequentially applying (7) to each of the next 25 filter outputs, the FAR may not be determined in such a sequential manner, as it cannot be known where to begin comparing the output of the i^{th} device's filter applied to the k^{th} device's dataset. Simply comparing the distributions of the filter output for the two cases would also produce an inaccurate FAR, as the filter output for each device is changing in time, and it would not be unreasonable to assume that at a particular point in time one device will have the same filter response as another device at a different point in time (Figure 6).

Thus, to calculate an accurate FAR, we assume that the filter response for each record of the k^{th} device's dataset us-

ing the i^{th} device's filter, $\gamma_{i,k}^{1 \dots n}(t_0)$, where n is the number of records in a dataset, is equally likely. Based upon this assumption, random numbers between one and n are generated to serve as an index used in deciding the starting value of j , for the filter response $\gamma_{i,k}^j(t_0)$.

Using the first value of the index for j , the next 24 filter responses, $\gamma_{i,k}^{j \dots j+24}(t_0)$, are compared to the threshold values calculated for $\mu_i^{1 \dots 25}(t_0)$ to check whether or not a record from $\gamma_{i,k}^{j \dots j+24}(t_0)$ would be accepted as a record from $\mu_i^{1 \dots 25}(t_0)$. This procedure is followed for each 25 record segment of $\mu_i^{1 \dots n}(t_0)$, where every 25 records a new value of j is chosen by taking the next value in the index. The total number of index values generated should then be n divided by 25. The FAR is then calculated using:

$$FAR = \frac{n_a}{n} \quad (11)$$

where n_a is the number of accepted records and n is the




Figure 6. Filter output for 10,000 records of two different Ethernet devices using the same filter, where at different times the filter output is the same.

number of total records.

This procedure is repeated 1,000 times, with new index values chosen for each iteration. The FAR for each iteration are then averaged to produce the total FAR. Repeated testing using this method has provided consistent values for the FAR.

4.5. Results

The results of the matched filter methodology for signal profiling are shown for 16 devices, consisting of a combination of three different models, via a confusion matrix (Table 1), which indicates the FRR and the FAR. The FRR may be deduced by subtracting the diagonal elements from one, while the FAR is simply the off-diagonal elements. Perfect detection/rejection would result in a matrix where the diagonal is one and off-diagonal elements are zero. The FRR and FAR are reported for 10,000 records per dataset. The naming convention $mXcY$ is utilized to denote card Y of model X .

As can be seen from the table, the FRR is sufficiently low (less than 1%), for different model cards we have near perfect detection, while some cards of the same model are difficult to differentiate. By experimenting with different minimum and incremental values used in determining δ , we have found that minimum and incremental values of .001 allow for too much variation in filter output. In fact, a slightly lower value of δ for each card will result in a negligibly higher false-reject rate; completely eliminate nearly all collisions which occur with frequency less than 20%; decrease collisions which occur with frequency less than 80% by up to 30%; but have no affect on collisions which occur with frequency greater than 80%. In addition, by utilizing the

techniques discussed in Section 3.5., we have been able to substantially reduce or eliminate most collisions. In particular, bandpass filtering proved particularly effective in differentiating $m6c3$ from $m5c3/7$. Through the use of both bandpass filters and an ensemble of filters, we were also able to eliminate almost all of the intra-model collisions of $m5cY$ and $m6cY$, respectively. Time-domain trimming and an ensemble of filters were also employed to dramatically reduce the number of collisions in $m4cY$, although perfect discrimination was not possible.

5. Future work

Several important issues regarding the variability of a device's analog signal require additional consideration. For example, under what conditions does the signal vary, how does device aging affect signaling characteristics, and how can a signal from a system that has lost and re-established a connection with the network be tracked? These questions, amongst others, provide a rich backdrop for future research.

An immediate area of consideration is extending this work to include different networking systems. Initial work has already begun on attempting to profile 100Mb Ethernet signals. Preliminary results indicate that the aforementioned techniques will be adequate for discriminating between different model devices; however, a deeper investigation into the signaling characteristics of 100Mb Ethernet devices may be required in order to provide accurate results for devices of the same model. Work will also continue in the 10Mb realm, as we try to create signal profiles for as many devices as possible. Other work includes analyzing wireless signals from 802.11b, sensor networks, and RFID systems. Currently, we are attempting to optimize

Table 1. Confusion matrix of 16 devices with 10,000 records per dataset

Expected Card	Tested Card															
	m4			m5										m6		
	c1	c2	c3	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c1	c2	c3
m4c1	.9961	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
m4c2	0	.9965	.8470	0	0	0	0	0	0	0	0	0	0	0	0	0
m4c3	0	.8988	.9956	0	0	0	0	0	0	0	0	0	0	0	0	0
m5c1	0	0	0	.9969	.9729	.0003	0	0	.0012	.0002	.0342	0	0	0	0	0
m5c2	0	0	0	.9290	.9970	0	0	0	.0026	0	.0626	0	0	0	0	0
m5c3	0	0	0	.0032	0	1.000	0	0	0	.9982	0	0	0	0	0	0
m5c4	0	0	0	0	0	0	.9999	0	0	0	0	.0020	.0017	0	0	0
m5c5	0	0	0	0	0	0	0	.9928	0	0	0	0	0	0	0	0
m5c6	0	0	0	.0184	.0394	0	0	0	.9999	0	.9584	0	.7792	0	0	0
m5c7	0	0	0	.0003	0	.9751	0	0	0	.9940	0	0	0	0	0	0
m5c8	0	0	0	.0278	.0751	0	0	0	.8873	0	.9957	0	.1821	0	0	0
m5c9	0	0	0	0	0	0	.0001	0	0	0	0	.9932	0	0	0	0
m5c10	0	0	0	0	0	0	.0004	0	.3988	0	.1518	0	.9938	0	0	0
m6c1	0	0	0	0	0	0	0	0	0	0	0	0	0	.9995	.3489	0
m6c2	0	0	0	0	0	0	.0150	0	0	0	.0490	0	0	.3176	.9992	.1037
m6c3	0	0	0	0	0	0	.5769	0	0	0	.7100	0	0	0	.0857	.9994

the matched filter for the profiling of wireless signals. Major challenges include adjusting the sensitivity of the filter to handle fluctuations of amplitude. Possible solutions to this problem include signal normalization and equalization.

6. Conclusion

We have shown that the matched filter can be reliably used to build signal profiles that can be used to discriminate between Ethernet cards of different models. By applying the matched filter in non-traditional ways, we have also demonstrated that it is possible to discriminate between seemingly identical cards, which appear to have originated from the same manufacturing lot. Finally, we have demonstrated that the analog signal characteristics of Ethernet devices can be tracked, and are thus suitable for use in network access control schemes. The techniques used in evaluating the effectiveness of the matched filter method have also been given. Future work will focus on applying our methods to the high-speed (100Mb and 1Gb Ethernet) and wireless domains (802.11b, sensor networks, and RFID systems), as well as exploring how device behavior changes due to environmental factors and aging.

7. Acknowledgments

The authors would like to thank Ed Jackson for his invaluable and insightful contributions to the project, Jason Sytsma for his diligence and exactitude in collecting data, and Adrienne Huffman for her help in proofing the manuscript.

References

[1] M. Barbeau, J. Hall, and E. Kranakis. Intrusion detection and radio frequency fingerprinting in mobile and wireless

networks. Technical report, Carleton University, School of Computer Science, October 2003.

[2] L. Couch. *Digital and Analog Communication Systems*. Macmillan Publishing Company, 1990.

[3] K. J. Ellis and N. Serinken. Characteristics of radio transmitter fingerprints. *Radio Science*, 36:585–598, 2001.

[4] P. J. Ferrell. Method and apparatus for characterizing a radio transmitter. United States Patent 5,005,210, April 1991.

[5] M. B. Frederick. Cellular telephone anti-fraud system. United States Patent 5,448,760, September 1995.

[6] B. Gassend, D. Clarke, M. Dijk, and S. Devadas. Delay-based circuit authentication and applications. In *Proceedings of the 2003 ACM symposium on Applied computing*, Computer security, pages 294–301, Melbourne, FL, 2003. ACM Press.

[7] B. Gassend, D. E. Clarke, M. Dijk, and S. Devadas. Controlled physical random functions. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC 02)*, page 149, Las Vegas, NV, December 2002. IEEE Computer Society.

[8] B. Gassend, D. E. Clarke, M. Dijk, and S. Devadas. Silicon physical random functions. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 148–160, Melbourne, FL, 2002. ACM Press.

[9] J. Hall, M. Barbeau, and E. Kranakis. Detection of transient in radio frequency fingerprinting using phase characteristics of signals. In L. Hesselink, editor, *Proceedings of the 3rd IASTED International Conference on Wireless and Optical Communications (WOC 2003)*, pages 13–18, Banff, Alberta, Canada, July 2003. ACTA Press.

[10] J. Hall, M. Barbeau, and E. Kranakis. Detection of transient in radio frequency fingerprinting using signal phase. ITACS IT Theme Meeting, October 2003. Slide Presentation.

[11] R. D. Hippenstiel and Y. Payal. Wavelet based transmitter identification. In *Proceedings of the Fourth International Symposium on Signal Processing and Its Applications (ISSPA 96)*, Gold Coast, Australia, 1996.

[12] D. N. Hoogerwerf et al. Active waveform collection for use in transmitter identification. United States Patent 6,035,188, March 2000.

[13] R. Jones. *Most Secret War*. Hamilton, 1978.

- [14] J. Kamarainen, V. Kyrki, and T. Lindh. Signal discrimination based on power spectrum of filter response. Technical Report 80, Lappeenranta University of Technology, Department of Information Technology, 2002. Research Report.
- [15] D. Kaplan and D. M. Stanhope. Waveform collection for use in wireless telephone identification. United States Patent 5,999,806, December 1999.
- [16] T. Kohno, A. Broido, and K. C. Claffy. Remote physical device fingerprinting. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 211–225. IEEE, May 2005.
- [17] D. L. Mensa et al. Radar signature evaluation apparatus. United States Patent 6,529,157, March 2003.
- [18] H. Mustafa, M. Doroslovacki, and H. Deng. Automatic radio station detection by clustering power spectrum components. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, (ICASSP 02)*, volume 4, page 4168. IEEE, May 2002.
- [19] T. L. Overman and K. C. Overman. Adaptive radar threat detection and tracker verification system. United States Patent 4,146,892, March 1979.
- [20] Y. Payal. Identification of push-to-talk transmitters using wavelets. Master's thesis, Naval Postgraduate School, Monterey, CA, 1995.
- [21] C. L. Phillips, J. M. Parr, and E. A. Riskin. *Signals, Systems and Transforms*. Prentice Hall, 2003.
- [22] P. S. Ravikanth. *Physical one-way functions*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [23] D. Shaw and W. Kinsner. Multifractal modeling of radio transmitter transients for classification. In *Proceedings of the IEEE Conference on Communications, Power, and Computing (WESCANEX 97)*, pages 306–312, Winnipeg, Manitoba, Canada, May 1997. IEEE.
- [24] J. Toonstra and W. Kinsner. Transient analysis and genetic algorithms for classification. In *Proceedings of the IEEE Conference on Communications, Power, and Computing (WESCANEX 95)*, pages 432–437, Winnipeg, Manitoba, Canada, May 1995. IEEE.
- [25] O. Ureten and N. Serinken. Detection of radio transmitter turn-on transients. *Electronics Letters*, 35(23):1996–1997, November 1999.