Using Wireless Physical Layer Information to Construct Implicit Identifiers

Kevin Bauer[1], Damon McCoy[1], Ben Greenstein[2],
Dirk Grunwald[1], and Douglas Sicker[1]

[1] University of Colorado
{mccoyd, bauerk, grunwald, sicker}@colorado.edu
[2] Intel Research Seattle
benjamin.m.greenstein@intel.com

Abstract. Recent work has focused on removing explicit network identifiers (such as MAC addresses) from the wireless link layer to protect users' privacy. However, despite comprehensive proposals to conceal all information encoded in the bits of the headers and payloads of network packets, we find that a straightforward attack on a physical layer property yields information that aids in the profiling of users. In this paper, a statistical technique is developed to associate wireless packets with their respective transmitters solely using the signal strengths of overheard packets. Through experiments conducted in a real indoor office building environment, we demonstrate that packets with no explicit identifiers can be grouped together by their respective transmitters with high accuracy. We next show that this technique is sufficiently accurate to allow an adversary to conduct a variety of complex traffic analysis attacks. As an example, we demonstrate that one type of traffic analysis—a website fingerprinting attack—can be successfully implemented after packets have been associated with their transmitters. Finally, we propose and evaluate techniques that can introduce noise into the measurements of such physical layer phenomena to obfuscate the identifiers derived from them.

1 Introduction

The inherent broadcast nature of wireless communication coupled with widespread availability of commodity receivers and little regulation in the industrial, scientific, and medical (ISM) radio bands poses a significant privacy concern for users of 802.11 technology. The threat is that third parties who eavesdrop on communications may profile users by their actions and track their movements. For example, even when message confidentiality is provided by mechanisms such as WEP or WPA, traditional 802.11 packets reveal every user's identifying MAC address, which enables any third party to monitor or track other users in the network. Given an explicit MAC address identifier, it is trivial to associate packets with their transmitting device.

However, even without explicit network identifiers, it has been demonstrated that other *implicitly identifying* information can be used to profile and possibly track users [1]. For example, suppose that a user predictably visits a set of webpages on a regular basis. Such browsing habits can be used to construct an identifying profile for that user.

This type of identifying information is defined as an *implicit identifier*. In this paper, we present a technique that constructs implicit identifiers using information obtained from the physical layer in combination with more traditional traffic analysis methods. We show that such identifiers can be constructed even when state of the art privacy protections are applied.

Motivation: An Identifier-Free Link Layer. To eliminate the transmission of implicitly and explicitly identifying information at the link layer, recent work has focused on providing identifier-free link layer protocols that obfuscate all transmitted bits to increase privacy with respect to third party eavesdroppers [2–4]. By concealing fields such as addresses, these protocols greatly increase the difficulty for unintended recipients to associate sequences of packets to their sources or destinations. Thus, to a large extent, they defend against attacks such as location tracking and traffic analysis that require correlating sequences of packets. While previous work has been limited in scope to addressing the removal of identifiers from the link layer, we demonstrate that there exists a significant amount of identifying information preserved within the properties of the wireless physical layer.

Reconstructing Packet Sequences Using Physical Layer Information. Our goal is to develop a method to classify packets by their transmitter from information revealed by the physical layer. When only a single device is active and transmitting in a wireless network, it is trivial to associate the transmitted packets with that device. However, when several devices are active, their transmissions can be interleaved or mixed, thus making it difficult to separate sequences by their sources. In order to implicitly identify devices using traffic analysis techniques, it is necessary to isolate and group the sequences of packets transmitted by each device. We present, implement, and evaluate a statistical technique to associate packets with their respective sources when multiple devices are transmitting simultaneously. This method relies only on information provided by the physical layer and, thus, it is effective even when implicit and explicit identifiers are removed from the link layer. Our approach is based on recording the strength of received signals from several locations and applying a clustering algorithm to classify the packets by transmitting device. The method is practical, since it utilizes commodity hardware and requires no training or cooperation from the wireless devices in the network.

While our approach will determine which packets originated at the same source, it won't identify sources by name. However, we demonstrate that our packet grouping is accurate enough to enable complex traffic analysis attacks, which use features such as packet timings and sizes, that reveal more about who the user is and what he is doing. Examples of the types of information that can be inferred through traffic analysis attacks include videos watched [5], passwords typed [6], web pages viewed [7, 8], languages and phrases spoken [9], and applications run [10]. While any of these traffic analysis methods can be applied, we show that the reconstructed packet sequences that we derive can be used to perform a practical website fingerprinting attack with a high enough level of success for wireless users to be concerned.

Experimental Validation. In order to demonstrate the efficacy of our method, we evaluate our technique by conducting experiments in a real indoor office building environment. A set of passive listening sensors are deployed and signal strength read-

ings are collected for wireless transmitters scattered throughout the building. We apply our packet clustering technique, which uses straight-forward statistical methods, and the results show that packets are correctly correlated to their transmitting devices with 71%–85% accuracy, depending on the number of transmitters in the network. We evaluate how these reconstructed sequences of packets can be used to perform a website fingerprinting traffic analysis attack, as in [7, 8]. Our results demonstrate that we can fingerprint a website with 41%–55% accuracy using the reconstructed sequences of packets. As more sophisticated clustering approaches are possible, we consider these results as an establishment of a lower bound on attainable accuracy.

Obscuring Physical-Layer Information. Finally, we explore two unique methods to mitigate the effectiveness of using the physical layer to associate packets to transmitters. Both of these methods introduce additional noise to obscure the properties of the physical layer. The first method is for the transmitter to randomly vary the power level at which it sends packets. The second is to use a directional antenna and change the antenna's orientation while transmitting. Both of these techniques show promise; we demonstrate that both the packet clustering accuracy and traffic analysis accuracy are degraded in the presence of devices that either vary their transmit power level or utilize a directional antenna.

Outline. The remainder of this paper is organized as follows: In Section 2, we discuss our threat model and in Section 3, we provide a brief description of wireless physical-layer properties and present our method at a high level. Section 4 describes the experimental validation with which we demonstrate the effectiveness of our technique. Section 5 presents our approach to traffic analysis and the results of the website fingerprinting attacks on packets that have been clustered by source device. Preliminary solutions are evaluated in Section 6. Finally, avenues of future work and concluding remarks are provided in Section 7.

2 Threat Model

The primary goal of this attack is to passively and accurately correlate packets to their transmitting device using only information revealed by the wireless physical layer. We assume that messages are sent using an identifier-free link layer, eliminating an eavesdropper's ability to associate messages to devices using explicit identifiers such as a MAC address or other implicit information leaked at the link layer. It is trivial to correlate packets when it is known that only a single device is active at any particular time. However, we assume a more common situation in which several devices may transmit at arbitrary times, possibly with interspersed transmissions. Analyses of wireless traces have shown that there are often many simultaneously active devices at tight time scales [3].

Attack. Even when identifiers are removed from the link layer, it is possible to label packets by their sources using only the implicit characteristics of the physical layer. An eavesdropper can use the sequence of packets associated with a particular device to perform more complex traffic analysis attacks, even if the packets are encrypted. The attack is completely passive and users can be subjected to this attack without their knowledge.

In addition, this technique requires only commodity hardware and no resource-intensive training.

Having accurately grouped packets by their transmitting device, an adversary can perform traffic analysis attacks. The effectiveness of these attacks is dependent on an adversary's ability to associate packets sent to their respective transmitting devices. We assume that the adversary performs a traffic analysis attack on packets observed during a short time period, as wireless users are more likely to leave the network as time progresses.

Adversary. The adversary places several commodity 802.11 wireless sensors in chosen positions around a target location (such as a building). Also, the attacker has the ability to estimate how many devices are present in the area. This can be achieved through visual inspection, an automated machine vision method, or a reliable estimate of the expected number of devices.

Victims. The victims use a standard 802.11 wireless device to communicate using an identifier-free link layer protocol, and transmit at a constant power level. Also, the user is using an application, such as a web browser, that is vulnerable to traffic analysis attacks. Users remain stationary while they transmit, but are free to move when their transmitters are silent.

3 Packet Source Classification

Background and Intuition. Our packet association technique uses information leaked solely at the physical layer. When a commodity 802.11 wireless card receives a packet, it records the signal strength of the received packet as a received signal strength indication (RSSI) value. In a simplified signal propagation model, wireless signals fade with distance as they propagate over physical space. Thus, the RSSI values roughly correlate with the distance between the transmitter and receiver. This means that the same transmission will be received at different RSSI levels depending on the distance between the transmitter and receiver. Using these RSSI values, we show that it is possible to passively identify the device that transmitted a particular set of packets.

However, several factors affect a packet's RSSI value in real world environments, which makes accurately associating packets to their transmitting devices using physical layer information a very challenging task. From the perspective of the receiver, the RSSI values of different packets from the same transmitter often vary over time due to unobservable factors such as multipathing, interference from other devices, and unpredictable fading [11]. A robust packet association method cannot assume a static environment. Hence, our technique uses statistical methods to cluster packets to their true transmitting device. In practice, we find it necessary to obtain multiple RSSI readings from several spatially disparate sensors for accurate classification. This is due both to inherent ambiguity when using one sensor (transmissions from two different locations might result in similar RSSI values when the transmissions propagate over roughly the same distance) as well as to the high level of unpredictable temporal variability in RSSI readings. Figure 1(a), for instance, shows the RSSI values recorded from multiple packets sent over time from five distinct transmitting devices, whose corresponding physical locations are given in Figure 1(b). While the values are similar for each device, there is

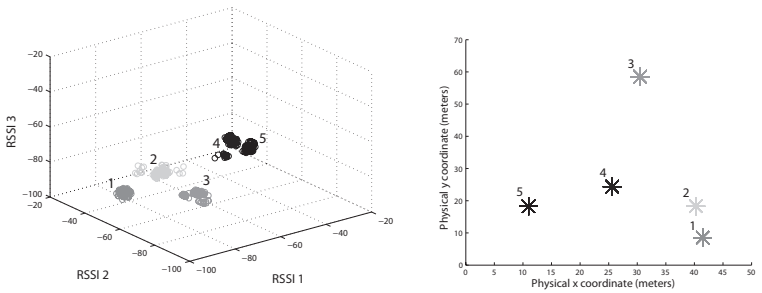

(a) RSSI values from three sensors for five different devices.

(b) The corresponding physical locations of the five devices.

Fig. 1—A visualization of the RSSI values from five transmitters at different locations. The packets are color coded by physical location.

some unpredictable fluctuation due to the inherent noise in the physical environment. When transmitters are close to one another, it is helpful to apply statistical techniques and multiple RSSI readings to overcome this noise.

Packet Association Technique. As the first step in the method to associate packets to their respective transmitters using physical layer information, we assume that an adversary can deploy n passive sensors to record the RSSI values of every packet received. For each received packet p_i, the RSSI values are combined into a feature vector of the form $(RSSI_{i1}, RSSI_{i2}, ..., RSSI_{in})$. Since the RSSI values are inherently noisy, we use the k-means clustering algorithm to group packets by their respective transmitting devices. In order to perform source classification, k-means requires the RSSI feature vectors and the number of devices (k), which we assume is known (or closely estimated) by the attacker. While k-means is computationally inexpensive, it is probabilistic in nature and, therefore, is not guaranteed to provide a globally optimal solution. For this reason, it is common to execute k-means repeatedly to arrive at a stable clustering result. More details on k-means clustering can be found in [12].

4 Experimental Validation

To demonstrate the efficacy of our physical layer clustering technique, we present a series of experiments conducted with 802.11 devices in a real indoor office building environment. In this section, we describe the methodology used to collect real RSSI values. To understand how the packet clustering technique performs in practice, we present a set of metrics with which to evaluate its ability to accurately associate packets to wireless devices. We characterize the clustering technique's performance with respect to how the number of devices effects clustering accuracy, the influence of device proximity (or the physical distance between devices) on clustering accuracy, and how the number of listening sensors effects clustering accuracy.

4.1 Clustering Evaluation Metrics

It is necessary to define the metrics that will be used to evaluate the correctness of the packet clustering technique. In general, it is assumed that each packet transmitted from device *i* is initially labeled in *class i*. *k*-means (or any clustering algorithm) attempts to provide a mapping between the initial class labels and an arbitrarily chosen *cluster* label *j*. It is trivial to determine when the clustering algorithm has returned a perfect clustering solution; this occurs when every initial class label *i* is mapped uniquely to *precisely one* cluster label *j*. However, evaluating the accuracy of clustering results when the class to cluster label mapping is not perfect is, indeed, a challenging task.

In order to describe the accuracy of the clustering solution that adequately captures the variety of clustering errors that occur, we apply the *F-Measure* metric from information retrieval [13]. Suppose that C is the set of class labels, K is the set of cluster labels, and n_{ij} is the number of data points from class i that are placed in cluster j. The F-Measure is defined in terms of *precision* and *recall*, which are defined as follows.

Definition 1 *Recall is the fraction of data points from initial class i that are in cluster j. Recall is calculated as* $R(c_i,k_j) = \frac{n_{ij}}{|c_i|}$*, for each class* $c_i \in C$ *and cluster* $k_j \in K$.

Definition 2 *Precision is the fraction of data points within cluster j that are members of class i. Precision is calculated as* $P(c_i,k_j) = \frac{n_{ij}}{|k_j|}$*, for each class* $c_i \in C$ *and cluster* $k_j \in K$.

Recall and precision can be combined into the F-Measure, which provides an adequate measure of clustering accuracy. The F-Measure is a weighted harmonic mean of recall and precision in which both are weighted equally; it is defined below in Equation 1.

$$F(c_i,k_j) = \frac{2R(c_i,k_j)P(c_i,k_j)}{R(c_i,k_j)+P(c_i,k_j)} \tag{1}$$

A single F-Measure value can be derived for a particular clustering result by performing the class to cluster mappings, and scaling by the fraction of the total data points N that are in that class, as illustrated in Equation 2.

$$F(C,K) = \sum_{c_i \in C} \frac{|c_i|}{N} \max_{k_j \in K} \{F(c_i,k_j)\} \tag{2}$$

We use F-Measure as our primary metric for expressing clustering accuracy in the following evaluation of our physical layer packet clustering technique.

4.2 Experimental Setup

In order to understand how our physical layer packet clustering technique works in practice, we deployed five 802.11 wireless devices to act as sensors in a $75\,m \times 50\,m$

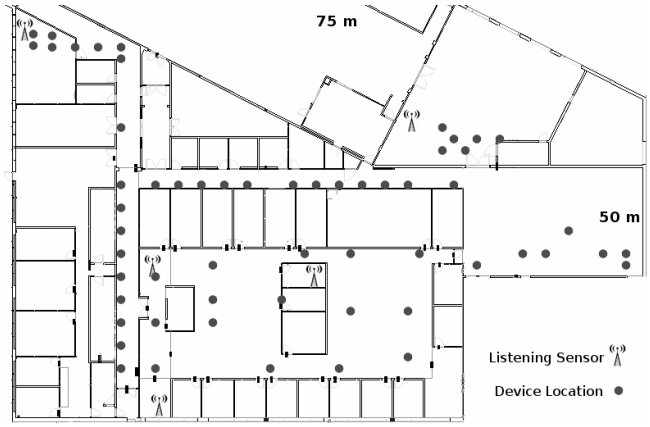

Fig. 2—Wireless devices are placed at 58 distinct physical locations in an office building.

indoor office building. Each sensor, a commodity Linux desktop machine, passively listens for packets on a fixed 802.11 channel. This allows the sensors to record RSSI values from all audible packets on that particular channel. To collect RSSI measurements, we moved around with a laptop computer and transmitted packets at a constant power level of 16 dBm at 58 distinct physical locations throughout the office space. The layout of the office space, marked with the positions of the passive sensors and the locations where the laptop transmitted packets, is provided in Figure 2.

Since we only used a single wireless device to transmit packets at multiple locations, in order to construct scenarios with multiple devices, we combined packets transmitted at multiple locations. For example, combining packets from five randomly chosen transmitter locations into a single network trace effectively represents a scenario in which five distinct devices are present. Finally, in order to evaluate the clustering technique, we assume that each device transmits precisely 100 packets.

Setup for Evaluating Variable Number of Devices and Sensors. To evaluate how the number of devices effects the clustering algorithm's accuracy, we vary the network size from 5, 10, 15, 20, and 25 devices. In order to ensure that there is no bias in the selection of the device's locations that may influence the performance of the clustering algorithm, we generate 100 randomly chosen device location configurations for each network size[3]. Next, we perform clustering on these device location configurations. Recall that since *k*-means is not guaranteed to provide a globally optimal solution, it is necessary to perform the clustering several times to arrive at a stable clustering solution; we observed that the algorithm stabilized after approximately 100 runs, therefore we perform *k*-means clustering 100 times on each device location configuration. First, we keep the number of sensors fixed at five to measure the effect of varying the number of devices. Next, to evaluate the number sensors on clustering accuracy, we also vary the number of sensors from one to five.

[3] Although we collected RSSI measurements at 58 distinct positions, we chose to limit the number of devices to 25 in any experiment to allow for variety in the randomly chosen locations of the devices included in the experiments.

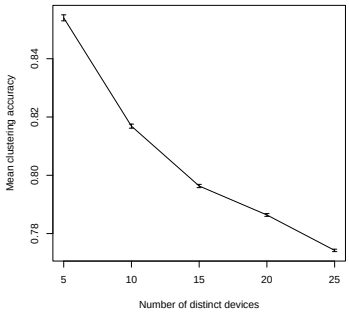

Fig. 3—Mean clustering accuracy for 5, 10, 15, 20, and 25 devices with 95% confidence intervals. Each device transmits an equal number of packets.

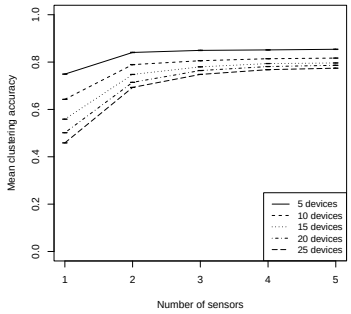

Fig. 4—Mean clustering accuracies (with 95% confidence intervals) for each device configuration as the number of sensors varies.

Setup for Evaluating Device Proximity. To examine the effect of device proximity on clustering accuracy, we conducted experiments in which the network size is fixed at two devices. Device location configurations are constructed by choosing each unique pair of device locations from the 58 total points. We use Euclidean distance between the two devices as a proximity metric.

4.3 Packet Clustering Results

We next present the results of our physical layer packet clustering technique in terms of its ability to accurately associate packets with their respective transmitting device. In particular, we examine three factors that we believe to be significant with respect to clustering accuracy: (1) the number of devices in the observation space, (2) the number of sensors in the observation space, and (3) the device proximity, as defined by the physical Euclidean distance between devices.

Effect of Number of Devices on Accuracy. The average clustering accuracy across the different number of devices is provided in Figure 3 (with 95% confidence intervals). In general, the accuracies decrease as the number of devices increase; thus, it can be concluded that the clustering algorithm performs better on a smaller number of devices and produces additional clustering errors as more devices are introduced. However, the 20 and 25 device experiments produced similar clustering accuracies, so there is evidence that the clustering accuracy may, in fact, level off as the number of devices reaches a critical threshold. In Section 5 we show that the clustering accuracy is marginally worse when the number of packets transmitted by each device is not constant for all devices.

Effect of Number of Sensors on Accuracy. As shown in Figure 4, the clustering accuracy is surprisingly high when just one sensor is used for each device configuration. However, as more sensors are added, the accuracy for each configuration increases gradually, with diminishing returns, as the number of sensors reaches three or higher. This indicates that the resources required—in terms of number of sensors to deploy—are very minimal, making the packet clustering technique very practical for a low resource adversary.

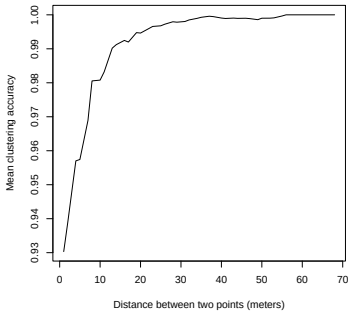

Fig. 5—Clustering accuracy steadily increases as the distance between devices increases.

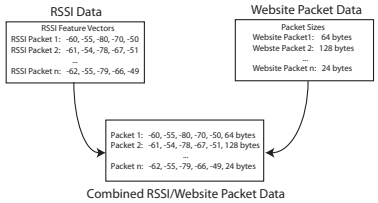

Fig. 6—Example of combining RSSI and website datasets.

Effect of Device Proximity on Accuracy. Since there is a correlation between RSSI values and physical locality, it is the case that packets transmitted by wireless devices that are close to one another will be received with similar RSSI values at the listening sensors. There is, in fact, a relationship between proximity and clustering accuracy (as depicted in Figure 5). In the case where the two devices are closest, the clustering technique has an average accuracy of approximately 93%, while the average approaches perfect accuracy when the two devices are at least 25–30 *m* apart. This demonstrates that there is a relationship between device proximity to one another and clustering accuracy.

5 Application: Website Fingerprinting

Having evaluated the packet clustering technique in isolation, we now explore how this technique can be used to perform complex traffic analysis attacks. In particular, we demonstrate that the ability to associate packets with their transmitting devices with relatively high accuracy provides sufficient information to perform a sophisticated website fingerprinting traffic analysis attack in which the source of an encrypted HTTP session is discovered using only packet size information [7]. While we could have chosen to demonstrate the utility of our packet clustering technique with a variety of other classes of traffic analysis attacks, website fingerprinting is a sufficiently complex problem which can be practically implemented by an attacker.

In this section, we first present the traffic analysis methodology. Next, using our real RSSI data in combination with encrypted HTTP traces, we demonstrate the efficacy of a website fingerprinting attack using packets that have been associated to their respective transmitters using our packet clustering method. The results of this traffic analysis attack are presented in terms of website identification accuracy. Recall that our packet clustering technique assumes that the number of devices is known a priori; we finally explore the extent to which the number of devices can be specified imprecisely, and show the website fingerprinting accuracy that results in those cases.

5.1 Traffic Analysis Methodology

In order to apply our real RSSI data to the problem of website fingerprinting, it is necessary to combine the RSSI data with an encrypted HTTP dataset. Liberatore and Levine [7] provide a dataset consisting of several instances of encrypted connections to many distinct websites over the course of several months. To perform a simplified website fingerprinting attack in combination with our packet clustering, we extract multiple traces of 25 different websites from this dataset. In general, to perform a website fingerprinting attack it is necessary to partition the website trace data into two disjoint sets, a training set, and a validation (or testing) set, and consider the task of website identification as a classification problem: given an unlabeled website instance, it is necessary to identify the website using the training set. We construct our website training set by collecting precisely 20 instances of each of the 25 websites that we wish to identify.

Next, the validation set is constructed. Since the encrypted HTTP dataset was collected in a wired setting, it is necessary to affix wireless physical layer RSSI values to the encrypted HTTP packets. This is accomplished by appending an RSSI feature vector onto each website packet as follows. Suppose that the set W consists of 25 websites and that $w_i \in W$ consists of a set of j packets, with their respective packet sizes s_j.[4] Next, suppose that the set of physical locations L consists of the 58 positions where RSSI values were collected. $r(l_k)$ consists of all RSSI feature vectors that are recorded when the transmitter is at location $l_k \in L$. The following procedure is used to combine the RSSI dataset with the encrypted HTTP dataset.

1. A website $w_i \in W$ is chosen at random and w_i is removed from W.
2. A physical location $l_k \in L$ is chosen at random and l_k is removed from L.
3. For each packet $p_j \in w_i$, an RSSI feature vector $r_j(l_k)$ is appended onto p_j.

This process is repeated for each of the wireless devices. After the procedure is executed, each packet consists of the tuple $p_j = (r_j(l_k), s_j)$, and thus each website packet in our validation set has been marked with a feature vector of RSSI values from our real RSSI data. This process is illustrated in Figure 6. It is now possible (1) to execute our physical layer packet clustering technique to associate each individual packet with its respective transmitter, and next (2) to launch a website fingerprinting attack to identify the website that is contained within the encrypted packets. For the website classification phase, we apply the naïve Bayes classifier provided by `Weka` [14], as in [7].

5.2 Traffic Analysis Results

We now present experimental results of our physical layer clustering technique applied to the task of performing a website fingerprinting traffic analysis attack. As in the experiments presented in Section 4, we construct different scenarios by varying the number of wireless devices from 5, 10, 15, 20, to 25. However, instead of including an equal number of generic packets, we make the assumption that every device downloads a single randomly selected webpage w_i and include all $p_j \in w_i$ packets with affixed RSSI vectors from a randomly selected position $l_k \in L$ using the process described above.

[4] Each packet's size is a multiple of the encryption algorithm's block size, which is 8 bytes, since 3DES encryption is used.

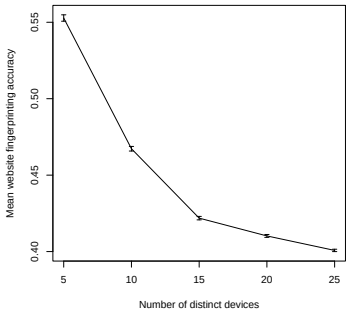

Fig. 7—Mean website classification accuracy (with 95% confidence intervals) as the number of devices varies.

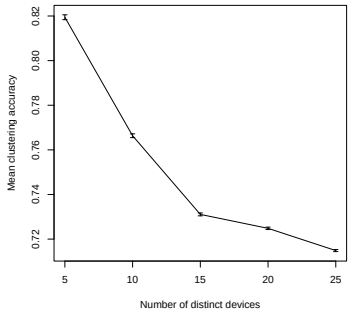

Fig. 8—Mean clustering accuracy for 5, 10, 15, 20, and 25 devices with 95% confidence intervals. Devices transmit different numbers of packets.

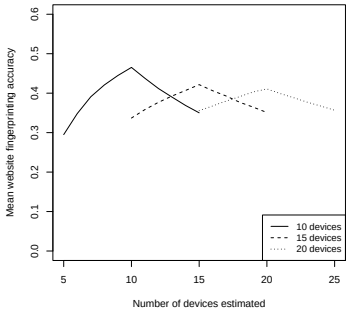

Fig. 9—Traffic analysis accuracy is best when the estimated number of devices is accurate.

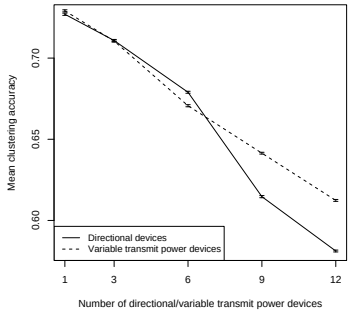

Fig. 10—The average clustering accuracy (with 95% confidence intervals) as the number of devices with directional antennas or the number of devices with variable transmit power levels increases.

We first explore the performance of the clustering algorithm on the website data. A key distinguishing feature of the website data is that each website has a different and arbitrary number of packets. For some websites, the device transmits several hundred packets, while for others, the device transmits less than ten packets. The clustering performance for our different device configurations is shown in Figure 8. Clustering devices that transmit a non-uniform number of packets does not appear to be a significant factor in clustering accuracy; the clustering accuracy for the website data is only marginally lower (72%–82% accuracy) than for the equal packet data (given in Figure 3).

Given the clustering algorithm's ability to accurately cluster encrypted website data, we next perform a website fingerprinting attack on the packets that are clustered by wireless device. Using the naïve Bayes classifier, the attack is able to correctly identify the encrypted web page between 40%–55% of the time. This accuracy is significantly greater than random chance, in which an adversary simply guesses the website. In this

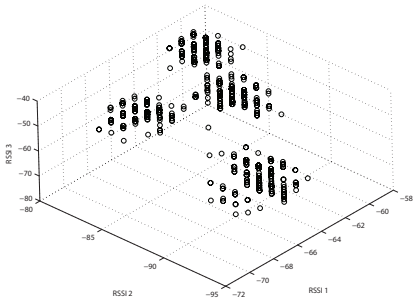

Fig. 11—Visualization of noise created from a directional antenna transmitting at the same location rotated 0, 90, 180, and 270 degrees.

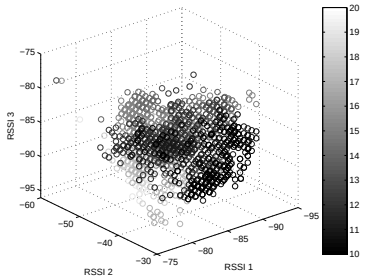

Fig. 12—Visualization of the high RSSI variance introduced when a single device varies its transmit power.

case, the expected accuracy is 1/25 = 4%. For comparison, if packets are perfectly clustered, the website fingerprinting attack achieves 92% accuracy (for each device configuration). The accuracy of the website identification is strongly linked to the accuracy of the clustering result; for example, in the 5 device network, both the clustering and website identification accuracies are the highest, and each respective accuracy degrades as the number of devices increases. The website fingerprinting accuracies for each number of devices are provided in Figure 7.

Traffic Analysis When *k* is Not Known Precisely. We finally explore the likely scenario in which the adversary cannot obtain the precise number of devices (k) in the wireless network. To explore the feasibility of performing traffic analysis in this case, we conduct experiments with 10, 15, and 20 devices where the k value is not correct. In particular, for a 10 device network, we vary k from 5 to 15, for a 15 device network, k varies from 10 to 20, and for a 20 device network, k varies from 15 to 25. The website fingerprinting accuracies for each network as k varies are given in Figure 9.[5] Even in the cases when the adversary's k estimate is off by as much as 5 (either higher or lower than the true value), the website fingerprinting accuracy decreases by at most 17%.

6 Solution: Obfuscating Physical Layer Information

We demonstrated a unique method to correlate packets with their transmitting devices using only information that is provided by the physical layer. Given the ease and relative accuracy with which this method can be applied to perform more sophisticated traffic analysis tasks, we explore techniques to mitigate the amount of information present at the physical layer that an adversary can use to associate packets with wireless transmitters. This serves as a basis to protect users' privacy from identifying information that is leaked by the physical layer. In particular, the solutions focus upon manipulating the physical layer properties that influence the RSSI values that are recorded by (1) varying the transmission power at which packets are sent, and (2) exploiting directionality (with a directional antenna) to focus the signal in specific directions.

[5] We do not provide clustering accuracies for the experiments where k is imprecise, since the F-Measure is ill-defined in this situation.

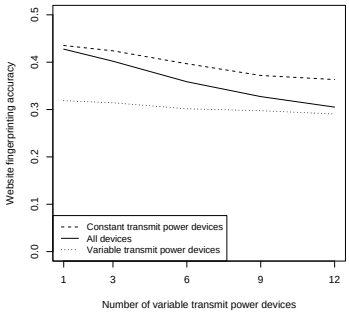

Fig. 13—The average website classification accuracy for devices that transmit at variable power levels, constant transmit power devices, and all devices in a network with 15 total devices (95% confidence intervals are very small so they are omitted).

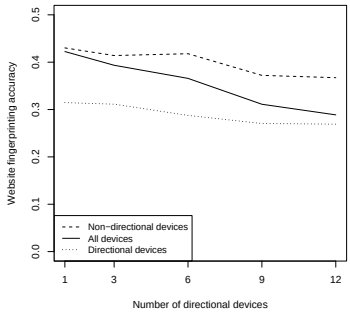

Fig. 14—The average website classification accuracy for devices with directional antennas, normal devices, and all devices in a network with 15 total devices (95% confidence intervals are very small so they are omitted).

Intuition. To provide an intuition behind these techniques, first consider an RSSI plot for a single wireless device with a directional antenna oriented in four different directions provided in Figure 11. The packets sent in each of the four directions appear to form their own distinct clusters—this phenomenon, as we will demonstrate, has an adverse effect on clustering accuracy, and, therefore, reduces an adversary's ability to perform traffic analysis attacks on the correlated packets. Next, the additional variance that is introduced in the signal space as a result of a single transmitter varying its power levels from 10–20 dBm, as depicted in Figure 12, results in a cluster that encompasses a significantly large portion of the signal space. As a result, we show that clustering accuracy degrades, which reduces its usefulness for performing traffic analysis.

6.1 Transmit Power Fluctuation

To understand the extent to which variable transmission power levels can be used to protect devices from leaking identifying information at the physical layer, we conduct experiments in which devices may wish to mitigate the effectiveness of our attack by sending each packet at a randomly chosen transmit power level between 10–20 dBm. All other devices in the network transmit their packets at a fixed 16 dBm. Experiments are conducted with 15 total devices in which 1, 3, 6, 9, and 12 devices transmit their packets at a randomly chosen power level. The impact of this technique on clustering accuracy is shown in Figure 10. As the number of devices with variable transmit power levels increases, the clustering algorithm's accuracy decreases as a result of the additional variance in the received signal strengths that is introduced.

This reduction in clustering accuracy has a negative impact on the accuracy of the website fingerprinting traffic analysis attack. As depicted in Figure 13, the traffic analysis accuracy remains relatively high for the devices that transmit at a constant power level in comparison to the case in which all devices transmit at a variable power level. However, for those devices that vary their transmit power levels, the website fingerprinting accuracy is reduced to approximately 30%. This demonstrates the effective-

ness of varying transmit power levels on mitigating traffic analysis for those devices that actively transmit packets at variable power levels. The traffic analysis accuracy for normal devices is marginally reduced, and thus, if a device desires protection, then it should actively fluctuate its transmit power levels.

6.2 Directional Antennas

We next explore how directional antennas can be used to obscure identifying information leaked by the physical layer. To understand how directionality effects an adversary's ability to correlate packets with their transmitters, experiments are conducted with 15 total devices in which the number of directional transmitters varies from 1, 3, 6, 9, to 12. Clustering accuracy, provided in Figure 10, decreases as the number of directional transmitters in the network increases. In general, the clustering accuracy decreases in a similar fashion as in our experiments with the variable transmit power levels, but the experiments with 9 and 12 directional antenna transmitters have lower clustering accuracies than their counterparts in the transmit power fluctuation experiments.

The degradation of clustering accuracy helps to mitigate traffic analysis. As shown in Figure 14, the website fingerprinting attack yields an accuracy of 30% for the directional antenna transmitters. However, the non-directional transmitting devices show a similar vulnerability to traffic analysis as in the experiments of Section 5.

6.3 Discussion

Note that transmitting at a variable power level provides a similar degree of protection from traffic analysis as the directional antenna technique—and thus either technique provides a viable solution. Providing variable transmit power capabilities is straightforward, by modifying an existing commodity 802.11 driver to effectively randomize the card's transmit power levels. In addition, low-cost directional antennas, such as sectored or MIMO antennas, are becoming widely available. However, it is likely that both of these techniques may reduce network performance. Therefore, it is necessary to understand the trade-offs between privacy and performance when considering these solutions.

7 Future Directions and Conclusions

In this paper, we experimentally demonstrate that even when explicit identifiers are removed from wireless packets at the link layer, a significant amount of identifying information remains preserved within the wireless physical layer. As future work, we plan to address several open issues with regard to the methods presented. First, the packet clustering technique assumes that the adversary can closely estimate the correct number of devices in the network—we plan to investigate methods to derive the number of devices automatically, perhaps from the properties of the physical signals themselves. In addition, while *k*-means clustering provides relatively high accuracy, it is not optimal. We plan to explore additional clustering algorithms that may be better suited to

our problem. As a final avenue of future work, we will study the potential benefits of combining timing information with the information contained at the physical layer to improve our technique's ability to accurately correlate packets with their transmitting devices.

Acknowledgments. We thank Jeff Pang for his insightful comments on early drafts of this work, James Martin for granting access to our office building testbed, and Eric Anderson for assisting with the data collection. This research was partially funded by NSF Awards ITR-0430593 and CRI-0454404.

References

1. Pang, J., Greenstein, B., Gummadi, R., Seshan, S., Wetherall, D.: 802.11 user fingerprinting. In: MobiCom. (2007)
2. Armknecht, F., Girão, J., Matos, A., Aguiar, R.L.: Who said that? Privacy at link layer. In: INFOCOM, IEEE (2007)
3. Greenstein, B., McCoy, D., Pang, J., Seshan, T.K.S., Wetherall, D.: Improving wireless privacy with an identifier-free link layer protocol. In: Mobisys. (2008)
4. Singelée, D., Preneel, B.: Location privacy in wireless personal area networks. In: WiSe. (2006)
5. Saponas, T.S., Lester, J., Hartung, C., Agarwal, S., Kohno, T.: Devices that tell on you: Privacy trends in consumer ubiquitous computing. In: Proc. 16th USENIX Security Symposium. (2007)
6. Song, D.X., Wagner, D., Tian, X.: Timing analysis of keystrokes and timing attacks on ssh. In: 10th USENIX Security Symposium. (2001)
7. Liberatore, M., Levine, B.N.: Inferring the source of encrypted HTTP connections. In: CCS '06: Proceedings of the 13th ACM conference on Computer and communications security, New York, NY, USA, ACM (2006)
8. Sun, Q., Simon, D.R., Wang, Y.M., Russell, W., Padmanabhan, V.N., Qiu, L.: Statistical identification of encrypted web browsing traffic. In: IEEE Symposium on Security and Privacy. (2002)
9. Wright, C., Ballard, L., Monrose, F., Masson, G.: Language identification of encrypted VoIP traffic: Alejandra y roberto or Alice and Bob? In: Proceedings of the 16th USENIX Security Symposium. (2007)
10. Wright, C., Monrose, F., Masson, G.: On inferring application protocol behaviors in encrypted network traffic. Journal of Machine Learning Research (2006)
11. Reis, C., Mahajan, R., Rodrig, M., Wetherall, D., Zahorjan, J.: Measurement-based models of delivery and interference in static wireless networks. SIGCOMM Comput. Commun. Rev. **36**(4) (2006)
12. Hastie, T., Tibshirani, R., Friedman, J.H.: The Elements of Statistical Learning. Springer (2001)
13. Van Rijsbergen, C.J.: Information Retrieval, 2nd edition. Dept. of Computer Science, University of Glasgow (1979)
14. Witten, I.H., Frank, E.: Data mining: Practical machine learning tools and techniques. Morgan Kaufmann, San Francisco, CA, USA (2005)