Transient-based Identification of Wireless Sensor Nodes

Boris Danev
System Security Group
ETH Zurich, Switzerland
bdanev@inf.ethz.ch

Srdjan Capkun
System Security Group
ETH Zurich, Switzerland
capkuns@inf.ethz.ch

ABSTRACT

Identification of wireless sensor nodes based on the characteristics of their radio transmissions can provide an additional layer of security in all-wireless multi-hop sensor networks. Reliable identification can be means for the detection and/or prevention of wormhole, Sybil and replication attacks, and can complement cryptographic message authentication protocols. In this paper, we investigate the feasibility of transient-based identification of CC2420 wireless sensor nodes. We propose a new technique for transient-based identification and show that it enables reliable and accurate sensor node recognition with an Equal Error Rate as low as 0.0024 (0.24%). We investigate the performance of our technique in terms of parameters such as distance, antenna polarization and voltage and analyze how these parameters affect the recognition accuracy. Finally, we study the feasibility of certain types of impersonation attacks on the proposed technique.

1. INTRODUCTION

Identification of components in a networked environment(e.g., operating systems, drivers, physical device) can benefit a number of applications such as authorized access, forensics, device cloning and malfunctioning detection, inventory management, tracking. This identification is commonly referred to as fingerprinting since it relies on distinctive characteristics (fingerprints) of network components, obtained with or without their cooperation. In a typical scenario, the fingerprinter observes traffic to and from a targeted device (fingerprintee) in order to find characteristics that (uniquely) distinguish the device or its components. Fingerprinting spans physical [1, 2, 3], link [4, 5] and application [6] layers for a variety of purposes such as identifying the type of a device [4], operating system [7, 8], particular drivers [5] or the physical device itself [2, 6, 9, 10, 11].

In wireless sensor networks, reliable sensor node identification can be means for detection and/or prevention of wormhole [12, 13], Sybil [14] and replication attacks [15], and can complement cryptographic message authentication protocols [13]. We focus on fingerprinting of wireless sensor nodes by distinguishing characteristics of their radio signals. This approach is commonly referred to as Radio Frequency Fingerprinting (RFF). More specifically, we investigate the feasibility of transient-based RFF [1, 2] of wireless sensor nodes. This fingerprinting technique consists of observing unique features in the radio turn-on transients, that appear at the beginning of each transmission. Device fingerprinting based on turn-on transients has been investigated in the past and has been shown to be useful in identifying radars, 802.11 devices [3, 16], Bluetooth mobile phones [11] and Mica2 CC1000 (433MHz) sensor nodes [13]. The majority of those works focused on the identifications of the device manufacturer or model.

In this work, we propose a new transient-based fingerprinting technique and show that this technique can be successfully used to identify individual 802.15.4 CC2420 radio transceivers of the same manufacturer and model. For this purpose, we propose an improved signal acquisition setup and related spectral FFT-based Fisher-features for sensor node identification. Our system enables highly accurate device identification both from short (<1m) and large (>40m) distances with an Equal Error Rate (EER) as low as 0.0024 (0.24%). We analyze the recognition accuracy of our system in terms of the number of signals used to build the device fingerprint, distance, antenna polarization, voltage and temperature. We show how changing these parameters affects the recognition accuracy. The obtained results expose the limitations of using transient-based techniques in dynamic environments. To validate the applicability of the proposed system to other radio transceivers,

IPSN'09, April 15–18, 2009, San Francisco, California, USA.

we also use it to identify CC1000 radio transceivers and show that it achieves similar performance to the CC2420 radios. This result indicates that our technique might be applicable to a wider range of transceivers. We further test the resiliency of our scheme to impersonation by hill-climbing antenna polarization attacks. We show that the system becomes highly vulnerable to such attacks if the number of signals used to build the device fingerprints is small. Finally, we demonstrate that malicious interference (jamming) can easily prevent accurate device identification.

To the best of our knowledge, this is the first work that analyzes the feasibility of fingerprinting of 802.15.4 CC2420 devices, evaluates the robustness of the transient-based identification in dynamic environments and its resiliency to certain types of impersonation attacks.

The remainder of this paper is organized as follows. In Section 2, we present our investigation parameters and system model. In Section 3, we detail our signal capturing process and summarize the data acquisition procedure and collected data. The proposed features for sensor node identification are explained in Section 4. Their performance is analyzed in Section 5. In Section 6, we develop a number of attacks and evaluate the resiliency of our fingerprinting approach. We describe possible application scenarios in Section 7, make an overview of background and related work in Section 8 and conclude the paper in Section 9.

2. PROBLEM STATEMENT AND SYSTEM OVERVIEW

In the paper, we will address the following questions:

1. What recognition accuracy can be achieved for identical wireless sensor nodes?
2. How is the recognition accuracy affected by the number of radio signals used to build the device fingerprint?
3. What are the effects of distance, antenna polarization and voltage on the recognition accuracy?
4. How susceptible is the recognition system to impersonation and denial-of-service (DoS) attacks?

Answers to the above questions will help identify the types of applications that the described transient-based identification methods are suitable for.

Device recognition systems typically work in one of the two modes: either identification of one device among many, or verification that a device's fingerprint matches its claimed identity [17]. Positive identification determines that a given device is in a (member) database. Functionally it is the same as verification. Negative identification determines if a device is not on a negative list of devices. In this work, we consider positive identification and more precisely verification of a device's claimed or assumed identity. The verification procedure matches a collected fingerprint of a device to the fingerprint that corresponds to its claimed or assumed identity. The verification system then provides an Accept/Reject decision based on a threshold value T (Section 5.1). Verification requires only "1-to-1" fingerprint comparison (compared to "1-to-N" in the case of positive identification) and is therefore scalable.

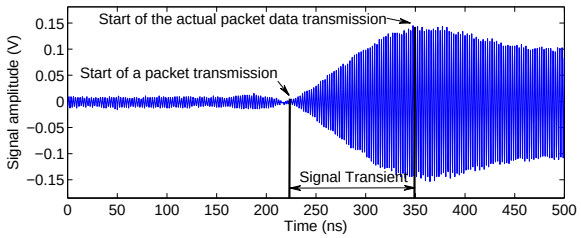

Figure 1: CC2420 radio signal transient shape at the start of each new packet transmission. Before the packet data transmission starts, the amplitude rises from channel noise to full power.

Our fingerprinting system is based on the extraction of the radio signal transient and distinctive features. Figure 1 shows the radio signal at the start of a new transmission (for CC2420, this effect occurs at the start of each packet). The transient is the part of the signal where the amplitude rises from channel noise to full power. The exact beginning and end of the transient is discussed in Section 3.3. The unique properties of the transient are generated by the analog part of the radio transmitter which includes an amplifier, band-pass filter, frequency mixer as well as the physical properties of the transmitting antenna. Each of these entities contains a number of passive (e.g., resistance) and active (e.g., capacitance) components which contribute to the unique behavior of the transient signal. We explore the features that make the transient distinguishable to each sensor node (the same manufacturer and model).

Our system consists of two primary components: a signal acquisition setup (Section 3) and a feature selection component (Section 4).

3. SIGNAL ACQUISITION

In this section, we describe the hardware setup for signal capture and present the collected datasets.

3.1 Hardware Setup

Figure 2 displays the hardware setup used to capture radio signals. The signals are acquired by a Standard Horn directional antenna and subsequently amplified by an ultra low-noise and low-power amplifier (NF=0.15 dB). Due to the low power of the sensor devices, it is critical to amplify the signal without losing its unique characteristics, as the signal-to-noise ratio degrades drastically within a couple of meters. An ultra

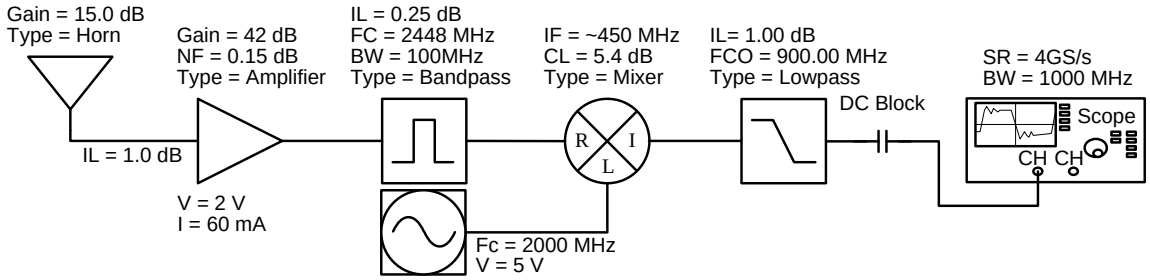

Figure 2: Radio signal-capturing hardware setup.

low-noise and low-power amplifier proved to be the best choice among a number of amplifiers we tested.

We used an ultra low insertion loss bandpass filter to eliminate radio frequencies outside the IEEE 802.15.4 band [18]. We then down-mixed this amplified and filtered signal to an intermediate frequency of 450 MHz using a standard frequency mixer and a voltage controlled oscillator. We down-mixed the signal to capture it with sufficient precision on the 1 GHz oscilloscope we had at our disposal. If the sensor (2.4 GHz) signals are not down-mixed, the oscilloscope significantly attenuates (-25 dB) their high frequency components, which in result significantly degrades the recognition accuracy.

Due to the frequency artifacts in the down-mixing process, we passed the intermediate frequency signal through a lowpass filter and a DC blocking capacitor before it was recorded by our oscilloscope (1 GHz bandwidth, 4 GS/s sampling rate). In all experiments we used high quality SMA cables with low insertion loss (approximately 0.5 dB depending on the cable length used). Our first experiments with standard BNC cables showed that these cables attenuated the signals such that they could not be used for accurate recognition.

The fact that our acquisition setup supports accurate recognition even when the signal is down-mixed to 450 MHz shows that a compact setup can be built for transient-based identification with off-the-shelf components. The primary component of such a setup would be an acquisition board (FPGA with an 1-2 GS/s ADC). It would even be possible to build this setup in a printed-circuited board (PCB) by using surface mount components instead of the currently used coaxial ones. We acknowledge that the price of such boards is currently high (10-15 K) which is a limiting factor in civilian compared to military applications. Therefore further investigation is needed to see if lower intermediate frequencies (<450 MHz) also preserve sufficient discriminant information in the transient part of the signal. This could significantly reduce the price of building the device.

3.2 Collected Data

Using the above described signal-capturing setup we collected sample signals from the sensor nodes. Our population of devices (P) consisted of 50 COTS Tmote Sky sensor nodes with manufacturer signature "4M 94V-0 H014-4787" (i.e., the same manufacturer and model). Given that they were purchased in 2 separate sets, we cannot fully assert that they were all produced at the same production line, even though such an assumption is highly plausible. The recorded datasets and main measurement parameters are summarized in Table 1.

Table 1: Data acquisition sets.

	Goal	Dist.	# Signals	# Nodes P	Total
1	Accur.	10m	600	50	30000
2	Accur.	40m	600	10	6000
3	Volt.	10m	200	10	2000
4	Polar.	-	600	10	6000
5	Attack	10m	350	3	1050

During data acquisition, each node was positioned on the same tripod, previously fixed at a given distance from the fingerprinter's antenna. Polarizations of the sensor devices' antennas (all sensors were equipped with standard on-board integrated antenna) and of the fingerprinter's antenna were aligned and perpendicular to the ground. The devices were run on 2 x 1.5V AA batteries (Dataset 1,2,4,5) and 2 x 1.2V AA batteries (Dataset 3). The experiments were made indoors (Dataset 1,3,4,5) and in a covered parking space (Dataset 2) for about 20 minutes with equally spaced packet transmissions in order to acquire a large number of signal samples for performance evaluation. The data acquisition phase might last shorter or longer depending on the sensor network application. The ambient temperature of the environment was between 18 and 23°C.

3.3 Transient Extraction

From each acquired signal (one signal corresponds to one packet), we extracted its transient. It should be noted that in a regular transmission from the nodes, the transient is present before each transmitted packet. Each acquired signal trace lasted 500 ns, of which the transient consistently lasted approximately 125 ns for all the nodes in our population set (Figure 1). Given the 4GS/s sampling rate of our oscilloscope, this cor-

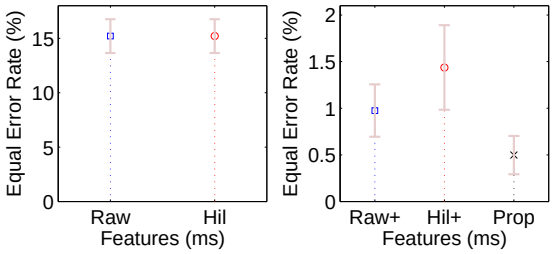

Figure 3: Recognition accuracy of the initial transformations (P=50, D=10m).

responded to approximately 500 data points. We defined the transient data sample as the 512 data points from its detected starting point. The starting point was determined by the variance-based threshold detection algorithm described in [13].

4. FEATURE SELECTION

The goal of feature selection is to obtain distinctive feature templates (fingerprints) from raw transient signals. Our feature selection procedure consists of two stages: (1) initial transformation and (2) feature extraction using statistical analysis. The initial transformation is selected from a set of known transformations and is an input into a Linear Discriminant Analysis (LDA) feature extraction. The feature extraction is done using a linear transformation derived from Fisher LDA [19].

In the initial transformation stage, we experimentally test a number of signal transformations to find initial features that capture most discriminant information in a device's transient. In the statistical analysis stage, we statistically determine linear boundaries between the initial features in order to efficiently reduce the dimensionality and increase the system accuracy. The used Fisher LDA has been effectively applied to discriminate human biometrics [20, 21] and outperforms related methods when the training data is sufficiently large [22].

4.1 Initial Transformations

We considered the following initial transient transformations: *Raw* - the original transient data sample, no transformation; *Hil* - the envelope of the transient data samples obtained by the Hilbert transformation, proposed in [23]; *Raw+* - the FFT spectra of the transient data samples; *Hil+* - the FFT spectra of the envelope of the transient data samples; *Prop* - differences between adjacent FFT spectra of the transient data samples.

We tested the use of these initial transformations in our recognition system. The results of the test over Dataset 1 are summarized in Figure 3. The figure shows the Equal Error Rate (EER) defined in Section 5.1. The obtained results show that when using the original transient data samples (*Raw*) or their envelopes (*Hil*), our recognition system scores a high EER (15%) which translates into a low recognition accuracy. This makes these two transformations unsuitable for further analysis. Using FFT spectra significantly improves the recognition accuracy (*Raw+, Hil+, Prop*), with (*Prop*) scoring the highest. We therefore chose the proposed relative differences between adjacent FFT spectra (*Prop*) as the transformation for further feature extraction.

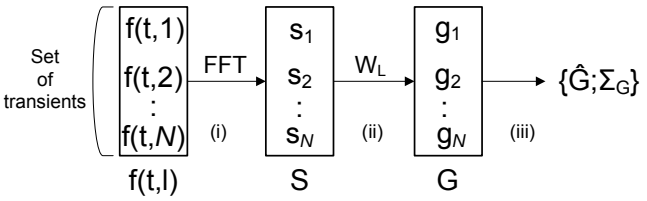

Figure 4: Feature extraction process.

The above results were validated with 4-fold cross validation [19]. Three folds of Dataset 1 were used for training and the remaining one fold for testing. Each fold contained 150 transient data samples per sensor node. This resulted in a total of 300 genuine and 22050 imposter matchings per fold[1] to compute the EER.

4.2 Feature Extraction

In this section, we describe our feature extraction process. It assumes the relative differences between adjacent FFT spectra as the initial transformation.

For a given sensor device, spectral Fisher-features are extracted from N captured signals using a linear transformation derived from LDA. Figure 4 illustrates the process. First, we extract the transient part of the recorded signal l. We denote this part by $f(t,l)$, where $f(t,l)$ is the amplitude of the signal l at time t.

In Step (i), we apply a one-dimensional Fourier transformation on $f(t,l)$ to obtain $F(\omega,l)$:

$$F(\omega,l) = \frac{1}{\sqrt{M}}\sum_{m=0}^{M-1} f(t,l)\exp(-2\pi i\frac{t\omega}{M}) \quad (1)$$

where M is the length of transient and $0 \le t \le M-1$. We then compute the relative difference between the adjacent spectra of the $|F(\omega,l)|$ denoted in a vector form as: $\vec{s_l} = [\,|F(2,l)|-|F(1,l)|\;|F(3,l)|-|F(2,l)|\;\cdots\;|F(M/2-1,l)|-|F(M/2-2,l)|\,]^t$ where the DC component and redundant half of the spectrum are removed.

In Step (ii), a projected vector $\vec{g_l}$, also called a Fisher-feature, is extracted from the Fourier spectrum using an LDA matrix W_L:

$$\vec{g_l} = W_L^t\vec{s_l} \quad (2)$$

[1]Each fold contains 3 feature templates (fingerprints) per sensor node. This results in 6 different matchings of fingerprints of the same sensor node (i.e., genuine matchings) and 441 different matchings of fingerprints from different sensor nodes matching(i.e., imposter matchings). This makes 300 genuine and 22050 imposter matchings for 50 sensor nodes.

Based on the above description, the Fisher-feature extraction from N captured signals for a given sensor device is written as $G = W_L^t S$ where G is an array of g_l and S is a matrix $S = [\ s_0\ ..\ s_l\ ..\ s_N\ ]$.

Finally in Step (iii), the feature template $\boldsymbol{h}$ used for matching (recognition) is computed:

$$\boldsymbol{h} = \{\hat{G}; \Sigma_G\} \quad (3)$$

where $\hat{G}$ denotes the mean vector of G and Σ_G denotes the covariance matrix of G.

The number of captured signals N used to build the feature template and the number of projected vectors in W_L (i.e., the Fisher subspace dimension) are experimentally determined.

4.3 Training and Mahalanobis Matching

The LDA matrix W_L is derived by a standard LDA procedure based on scatter matrices [19]. Here, W_L is the optimal Fisher discriminant projection given as the set of κ eigenvectors in matrix W that correspond to the κ-highest eigenvalues in the generalized eigenvalue problem: $S_b W = \Lambda S_w W$, where Λ is the eigenvalue matrix, S_w is the within-class scatter matrix showing the average scatter of sample features $\boldsymbol{h}$ from the same sensor device and S_b is the between-class scatter representing the average scatter of sample features $\boldsymbol{h}$ from different sensor devices.

Mahalanobis distance is used to find the similarity between feature templates (fingerprints). The result of matching a reference $\boldsymbol{h}^R$ and a test $\boldsymbol{h}^T$ feature templates is a matching score, calculated as follows.

$$Matching\ score = \sqrt{(\boldsymbol{h}^T - \boldsymbol{h}^R)^t \Sigma_G^{-1} (\boldsymbol{h}^T - \boldsymbol{h}^R)} \quad (4)$$

Values of the matching score closer to 0 indicate a better match between the feature templates.

It should be noted that the proposed feature extraction and matching method can be efficiently implemented in hardware as it uses only linear transformations for feature extraction and inter-vector distance matching to compute similarity. These operations have a low memory footprint and are computationally efficient.

5. PERFORMANCE EVALUATION

In this section, we present the performance results of our fingerprinting system. First, we review the metrics used to evaluate the recognition accuracy of the system.

5.1 Evaluation Metrics

We adopt Equal Error Rate (EER) and Receiver Operating Characteristic (ROC) as the metrics for evaluating the accuracy of the proposed system since these are the most agreed way for evaluating identification systems [17]. The metrics are briefly discussed below.

Hypothesis testing is a common approach to statistically establish matching between two samples. The *null* hypothesis H_o states that the two samples match and the *alternative* hypothesis H_a - that the two samples do not match. In such a setting, there are two possible errors: False Accept and False Reject. False Accept means that the system decides H_o when H_a is true. In our system this is equivalent to a decision that a device's (claimed) identity is legitimate while in reality it is an imposter device. False Reject means that the system decides H_a when H_o is true. In our system, this is equivalent to a decision that a device's identity is not legitimate while in reality it is.

The False Accept Rate (FAR) and False Reject Rate (FRR) represent the frequencies at which the above errors occur. The FAR and FRR are closely related to each other in the Receiver Operating Characteristic (ROC). The ROC is a curve which allows to automatically compute FRR when the FAR is fixed at a desired level and vice versa [17]. The operating point in ROC, where FAR and FRR are equal, is called the Equal Error Rate (EER). The EER represents the most common measure of the accuracy of a recognition system [24]. The operating threshold value at which the EER occurs is our threshold T for an Accept/Reject decision.

To increase clarity of presentation, we use the Genuine Accept Rate (GAR = 1 - FRR) in the ROC because it shows the rate of Accepts of legitimate identities. In addition, we also compute FRR for common target values of FAR (e.g., FAR = 0.01%, 0.1%).

5.2 Results

In our evaluation, we first consider the recognition results obtained using Dataset 1 (Table 1) that contains signals from all sensor nodes (P=50) taken at distance D=10m. The number of captured signals N used to build feature templates was fixed to N=50. The results are validated with the 4-fold cross validation procedure described in Section 4.1.

The results are presented in Figure 5(a) (Fisher-features) and show the dependency of the recognition accuracy (EER) of our system on the fingerprint size (i.e., the dimension of the Fisher subspace used to project the initial features into). The dimension of the features after the initial transformation (Section 4.1) is 254.

The results show very small EER of our system, which is, for fingerprint sizes $\geq$3 between 0.0024 (0.24%) and 0.005 (0.5%). This means that our system correctly identifies sensor nodes with an accuracy higher than 99.5% (GAR at the EER operating point). We later show that the accuracy achieved in this set is equally preserved for other datasets.

The results in Figure 5(a) confirm that using the first 5 eigenvectors of Fisher-features for projection scores

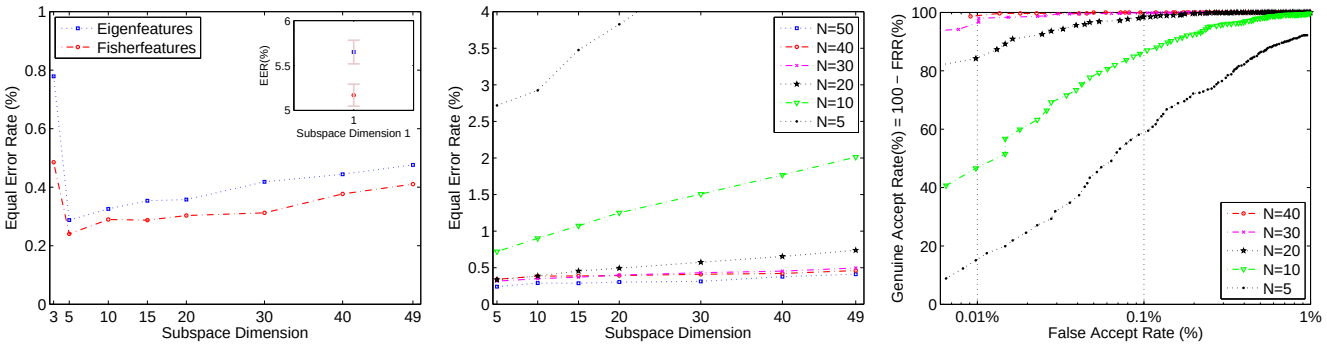

Figure 5: (a) Eigen- and Fisher-features accuracy for different subspace dimension. Dimension 1 is in the inner plot (P=50, D=10m, N=50). (b) Fisher-features accuracy for different subspace dimension and nbr. of signals N used to build the feature templates (P=50, D=10m). (c) Receiver Operating Characteristic (ROC) for different number of signals N used to build the feature template (P=50, D=10m). The Fisher-feature subspace dimension is fixed at 5. See Table 2 for the underlying data.

the highest recognition accuracy. EER degrades progressively in higher dimensional subspaces. This phenomenon is even more pronounced when the number of signals N used to build the feature template decreases, in particular for feature templates built with N<30 signals as shown in Figure 5(b).

The results also demonstrate that our proposed features keep the EER low even when fewer signals N<50 are used to build the feature template. This is exhibited in Figure 5(b) which gives the EER for different dimensions and N. Reducing N allowed us to perform 5-fold cross validation (5 folds x 120 signals) which increased the genuine and imposter matchings per fold (Table 2).

Figure 5(a) also presents the comparison between Eigen- and Fisher-feature extraction. Eigen-feature extraction is based on Principal Component Analysis (PCA). The validated EERs show that the Fisher-subspace is more efficient for lower dimensional subspaces (1-3 eigenvectors) compared to the eigenspace. However, we cannot assert with statistical confidence such behavior for higher dimensional subspaces. This is probably due to the 4-fold cross validation (the maximum for N=50) which produces large (overlapping) confidence intervals.

In summary, the above results demonstrate the recognition efficiency of our proposed acquisition setup and related spectral FFT-based Fisher-features. They also show that a 5-dimensional linear subspace is enough to represent a device feature template (fingerprint). Therefore, our proposed features also form very compact and computationally efficient fingerprints. If each dimension is represented by a 4-byte floating-point number, the size of the corresponding feature template $\boldsymbol{h} = \{\hat{G}; \Sigma_G\}$ is 20 (5x4) bytes for $\hat{G}$ and 100 (5x5x4) bytes for the square covariance matrix Σ_G resulting in a total of 120 bytes. It should be noted that optimization techniques exist to reduce the bit size per dimension to 1-2 bytes.

In order to fully characterize the accuracy trade-offs, we draw the ROC curves for the selected 5-dimensional features and different number of signals N as shown in Figure 5(c). Table 2 summarizes the underlying data, namely the number of signals N, total genuine and imposter matchings performed, Accept/Reject threshold T (at EER point), EER and its confidence interval (CI) and FRR for two common FAR=0.01%, 0.1% targets.

The ROC curve allows us to conclude that reducing the number of signals N used to built the feature templates, degrades the Genuine Accept Rate for lower targets of FAR (e.g., 0.01%). This is not readily visible from Figure 5(b) where the differences in EER for N>10 are statistically insignificant in the range between 0.24% and 0.34% (Table 2). The ROC analysis shows that if an application is required to operate at low FARs (<0.1%), it must use more signals to build the feature template for a reliable recognition with a high GAR.

5.3 Feature Stability

In the following analysis, we investigate the stability of our proposed technique in terms of distance, antenna polarization, voltage and temperature. We also show that our scheme can be used for identification of sensor nodes that use CC1000 radios.

5.3.1 Distance

For any practical use of physical-layer recognition, we must consider the effect of channel attenuation. For this purpose, we performed measurements in the university parking, which allowed us to collect signals from 40m line-of-sight (LoS). We used the first 10 sensor devices from our population set (Dataset 3, Table 1).

Table 3 compares the validated EERs for different N

Table 2: Summary of recognition accuracy for Dataset 1 (P=50, D=10m).

N	Test matchings		Threshold T	EER (%)	EER CI (%)		FRR (%)		Validation
	Genuine	Imposter			lower	upper	FAR=0.01%	FAR=0.1%	
50	300	22050	3.01	**0.24**	0	0.49	0.72	0.65	4-fold
40	300	22500	3.95	**0.34**	0.02	0.66	1.10	0.46	5-fold
30	600	39200	3.87	**0.32**	0.07	0.56	2.92	0.61	5-fold
20	1000	61250	4.10	**0.34**	0.21	0.47	12.94	1.24	5-fold
10	1000	61250	6.74	**0.72**	0.62	0.82	52.00	9.60	5-fold
5	1000	61250	16.04	**2.72**	2.38	3.06	82.12	40.10	5-fold

and a distance of 10m and 40m respectively. The system is trained separately for each distance. We do not observe statistically significant differences in the recognition accuracy. This shows that our capturing setup (Section 3.1) is successful in preserving the discriminant power of the transient signal.

It should be pointed out that for N=30,40,50 the algorithm achieves EER=0%. This confirms that the EER must be computed for a larger set of devices in order to have a more accurate estimation of the recognition capabilities. In biometric recognition systems hundreds and even thousands of different biometric identifiers (e.g., fingerprints, faces) are usually used for evaluation (e.g., NIST, FERET databases). In our experiments, however, due to limited resources, we could not evaluate on such large sets of devices.

Even though all signal capturing was performed in a university parking place with numerous possibilities for reflection (e.g., cars, concrete columns), we did not observe multipath propagation problems. We acknowledge that superposition of signal transients might prevent accurate recognition. In such scenarios, excess signals need to be detected and eliminated from the extraction of the matching features.

In order to complete the analysis on the effect of distance on the recognition accuracy, we performed cross-matching between feature templates extracted at 10m and 40m distance from the capturing antenna. We registered an average recognition accuracy of EER=0.38 (38.01%) for N=50. This result shows that while the frequency information in the transient signal is unique within a given distance, it changes across different distances for the same antenna polarization. The impact of antenna polarization is discussed in Section 5.3.3.

5.3.2 Voltage and Temperature

Given that sensor nodes are generally run on battery supply, we evaluated the effect of voltage. For this purpose, we used transient data samples captured with 2x1.5V alkaline and 2x1.2V NiMH batteries.

Figure 6 shows the matching scores between transient data samples taken at the same voltage level (blue triangles) and between transient data samples taken at different voltage levels (2.4V and 3V respectively) (red circles) for 10 sensor nodes. We do not observe a significant difference between genuine matching scores coming from the same and cross voltage levels. The scores are close to 0 and within the boundary of the genuine matching score distribution (i.e., below T=3.01) for N=50. The EER for this set of 10 sensor nodes (same set of nodes as in the previous section) is 0%.

Table 3: EER at D=10m and 40m (P=10).

N	Test matchings		EER (%)		Valid.
	Genuine	Imposter	10m	40m	
50	60	810	**0**	**0**	4-fold
40	60	810	**0**	**0**	5-fold
30	120	1440	**0**	**0**	5-fold
20	200	2250	**0.57**	**0.36**	5-fold
10	200	2250	**1.35**	**3.41**	5-fold

This is an expected result given that the sensor nodes are equipped with a low-power micro-controller. It requires 2.1-3.6V for its normal operation. It should be noted that such a result is not necessary true for high-power transmitters (e.g., VHF FM) as observed in [25].

Our experiments did not suggest any effect on the recognition accuracy from the surrounding temperature changes (indoor air-conditioned environment or non air-conditioned parking place). We point out however that the ambient temperature during signal acquisition did not vary substantially, the variance being approximately 5°C between the two environments used. We did not investigate extreme changes of temperature (e.g., intentional heating) and higher variance of the ambient temperature which usually occurs in outdoor environments. We intend to consider the latter in future work to quantify the effect.

5.3.3 Polarization

The polarization of an antenna is defined as the polarization of the wave radiated by the antenna. At a given position, the polarization describes the orientation of the electric field. This orientation will change in sensor network applications when the nodes change their position with respect to the receiving antenna. A direct consequence of changing polarization is the change in

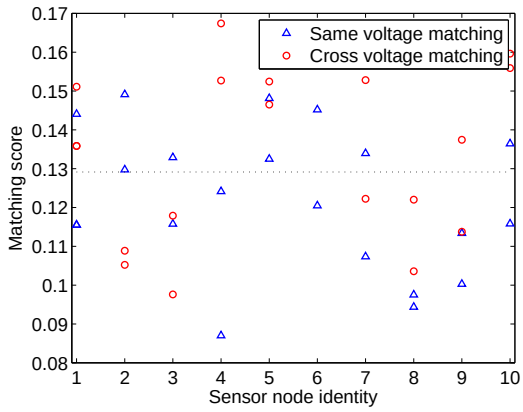

Figure 6: Matching score with variable voltage: the (blue) triangles represent matching scores of fingerprints from the same sensor node and same voltage level; the (red) circles represent matching scores of fingerprints from the same sensor node at different voltage levels (2.4V and 3V). All matching scores are below the threshold T=3.01, thus within the genuine score distribution (P=10, D=10m, N=50).

the shape of the transient signal as shown in Figure 7.

In order to quantify the effect of polarization, we collected transient data samples under the same conditions as in Dataset 1 (Table 1), but with a changed polarization of the antenna on the sensor node by 45° with respect to the fingerprinter antenna. We then matched the extracted feature templates to the reference feature templates in Dataset 1. This resulted in a degraded recognition accuracy (EER = 0.39 (39%)).

As this result could have been influenced by the training procedure where only training data from one type of polarization was used, we collected transient data samples from 10 sensor nodes at 3 different antenna polarizations (Dataset 4, Table 1). The recognition accuracy did not improve. This finding show that varying the polarization changes the frequency information in the transient signal. These changes cannot be well separated by a linear discriminant. The low accuracy is due to incorrect identification of 4 out of the 10 nodes, the other 6 being correctly identified. We acknowledge that further work is needed to quantify how much change in polarization can be tolerated (e.g., small perturbations) as the above results are for a 45° change. We also intend to consider non-linear feature boundaries which may overcome this limitation in future work.

5.3.4 *Results for CC1000 radios*

We applied our proposed features to the dataset collected by [13]. That dataset consisted of 2000 transient data samples captured from 10 identical Mica2

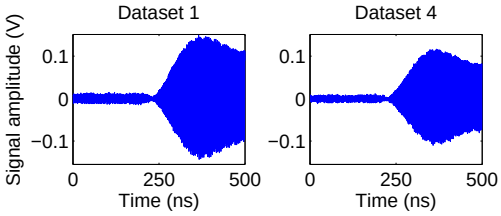

Figure 7: Transient signal shapes from a sensor node at two different antenna polarizations.

sensor nodes equipped with CC1000 (433Mhz) radios from 15 cm distance. The transient part occupied approximately 100 ns (200 data points). Our proposed features scored an EER=0.0167 (1.67%) on that data, showing that CC1000 radios can also be recognized with high accuracy. It should be noted that this result can possibly be improved if the linear transformation W_L was trained specifically for CC1000 radios. This was not possible due to the small size of the considered dataset.

In order to directly compare our features to the ones used in [13], we computed the performance metric used, namely the classification error rate [2]. In our case, it is 3.2% which is a significant improvement compared to the 30% classification error rate reported in [13].

5.4 Summary of Results

Our results show that sensor nodes can be recognized with high accuracy by analyzing the transient part of the transmitted signals. Such recognition proves to be robust to distance, multipath propagation and voltage changes. As such, it can be effectively used in applications where the sensor nodes do not often move.

Transient shape changes due to antenna polarization (mobility) introduce variability that degrades the recognition accuracy. This finding limits the usability of only transient-based features in applications where sensor nodes frequently move. Nevertheless, our features can be combined with other techniques (e.g, directionality, RSSI) to further reduce the set of probable sensor nodes from which the signals came. We acknowledge however that other statistical methods in particular non-linear (kernel) analysis[19] may be more effective in overcoming this issue. More investigation and experimentation is needed to assert this finding.

In application scenarios where the number of sensor devices is known, the classification error rate [19] can be used to evaluate the ability of the fingerprinting approach to classify (map) the transmitted signals to their corresponding devices. Table 4 displays the average classification error rates using our proposed technique on the full set of 50 nodes (Dataset 1) for typical 1-NN and 2-NN classifiers. The results show that the classifi-

[2]The classification error rate is the percentage of incorrectly classified samples to a predefined set of classes of samples.

Table 4: Average classification error rate (%).

N	# Samples	1-NN (%)	2-NN (%)	Valid.
50	300	0.07	0	4-fold
40	300	0.07	0	5-fold
30	600	0.25	0.07	5-fold
20	1000	0.97	0.45	5-fold
10	1000	3.71	2.43	5-fold

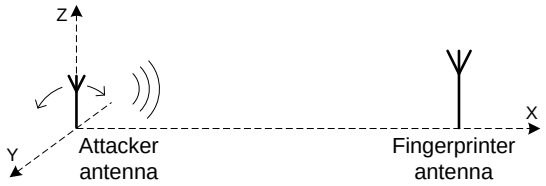

Figure 8: Hill-climbing attack setup. An attacker sensor node with external rotational antenna is positioned at the same X-axis as the fingerprinter antenna. The attacker changes the radio waves by rotating its antenna in the Y-Z axis to find a polarization that impersonates a sensor node from the targeted network.

cation error rate reduces when N increases. It reaches 0.0007 (0.07%) for 1-NN and 0 for 2-NN classifier.

Comparison of the classification error rates in Table 4 with related work (Section 8) can be misleading given the difference in the device population (same vs. different manufacturers), device hardware and radio specification, capturing distance. Nevertheless, our approach outperforms previous work on transient-based identification of identical CC1000 wireless sensor nodes [13] as demonstrated in Section 5.3.4. An advantage of our approach to a recent modulation-based identification technique [26] is that the classification error rate reduces significantly when the number of signals N increases.

It should be noted that the classification error rate is by definition not a suitable metric for recognition (verification) as outlined in Section 8. Furthermore, the obtained results show that the classification error rate significantly differs from the EER (Table 4 vs. Table 2).

We also point out that the results in Table 4 may be improved by using more sophisticated classifiers (e.g., SVM, PNN). However, these classifiers need to be augmented with doubt and outlier classes to fit the application requirements. They are also memory expensive and require more computational resources.

6. ATTACKING FINGERPRINTING

In this section, we analyze the robustness of our identification approach to impersonation and denial-of-service (DoS) attacks. In particular, we demonstrate a hill-climbing attack for impersonating a sensor device through variable antenna polarization and show that impersonation would be possible if a small number of signals is used for feature extraction. We also show that DoS attacks can prevent accurate identification. Finally, we discuss the implications of other attacks.

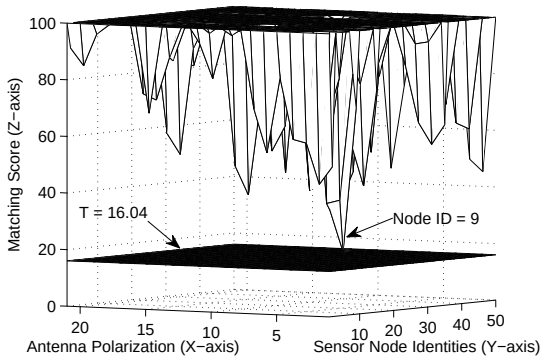

Figure 9: Hill-climbing attack scores. The X-axis contains the 21 (3 sensor nodes x 7 antenna polarizations) attacking features; the Y-axis shows the reference features of the 50 sensor nodes targeted for impersonation; the Z-axis is the matching score obtained between each attacking and reference features. The thick surface is the Accept/Reject threshold (T=16.04).

Table 5: Hill-climbing attack on sensor ID=9.

N	50	20	10	5
Hill-attack distance	42.74	38.12	35.89	21.61
Threshold	3.01	4.10	6.74	16.04

6.1 Hill-climbing Attack

A hill-climbing attack is a well-known attack on biometric recognition systems [17]. This attack consists of repeatedly submitting data to an algorithm with slight modifications. Only modifications that preserve or improve the matching score are kept in the process. Eventually, a score that exceeds the operating threshold (Table 2) might be achieved. This results in successful impersonation without providing the genuine biometric.

To perform the attack, we would ideally need a specialized device that is able to create transient signals (similar to the ones generated by the sensor nodes) and at the same time allow for introducing variations in it.

We decided to use 3 additional sensor nodes that are not part of the population of 50 sensor nodes used so far. In order to create variations in the shapes, we mounted external antennas on the 3 sensor nodes and change their antenna polarization as shown in Figure 8.

We collected 50 transient data samples from 7 different polarization positions of the antennas of the 3 sensor nodes. We then supplied these transient data samples to our proposed matching algorithm. Figure 9 displays the matching scores obtained during the attack in a 3D

representation for N=5. For clarity reasons, all scores that exceed 100 are not displayed.

The identification procedure becomes more vulnerable to the impersonation attack when N decreases. In particular, the matching scores against sensor node ID=9 for N=5 were consistently very close to the Accept/Reject threshold T=16.04 (Table 5). Device impersonation is possible for $N \leq 5$. A real system needs to consider acquiring $N > 5$ signals to build the fingerprint to ensure protection against this type of impersonation.

6.2 Denial-of-service Attacks

Due to the low output power and limited spectral diversity of sensor node transceivers, wireless sensor networks are particularly vulnerable to jamming-based DoS attacks [27]. We therefore decided to quantify the effect of jamming on the recognition in our system.

We collected transient data samples in the presence of a jammer. For jamming purposes we used an USRP device with GNU radio software [28]. Figure 10 displays 2 different transient data samples acquired in the presence of a Gaussian noise jamming signal.

The matching experiments showed that it is impossible to recognize the device due to the superposition effect of the jamming and the original sensor node signal. Furthermore, even jamming a small amount of the sensor node signals (5-10 out of 50 that formed the template features) was sufficient to prevent accurate recognition. These findings show that an identification procedure based on physical signal characteristics must be complemented by a jamming detection mechanism.

It should be noted that a sophisticated jammer can jam only the signal transient, which will result in successful data transmission, but inaccurate identification. As a result, there is a need for devising a jamming detection procedure not only at the data layer [27, 29], but also for the transient part of the transmission.

This attack also shows that if the network authority wants to prevent fingerprinting by an attacker, it could do so by appropriately jamming the communication between the sensor nodes (i.e., jamming only the transient and not affecting the transmitted data).

We did not investigate intentional heating of the circuit of the sensor node as a possible DoS attack. We point out that even if such an attack succeeds, it might be easily detected by appropriate temperature sensors or tamper-responsive shielding [30].

6.3 Other Attacks

The possibility of an attack which records the transient part of the signal and subsequently concatenates it to some data needs to be investigated. There is a number of points which make this attack hard to achieve. First, the replaying device needs to have a zero-length

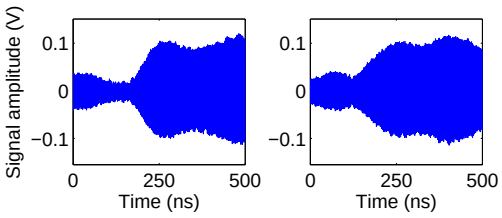

Figure 10: Jammed transient signals (to be contrasted with the not jammed signals in Fig. 7).

transient in order to successfully transmit the originally recorded transient. Second, the concatenation needs to be also very precise to allow accurate demodulation of the signal for data extraction. Third, the replayed transient part features will score exactly the same matching score when matched to the reference template features of the attacked device. As a result, the attack is easily detectable unless some variability is introduced to prevent same matching score. In addition, the introduced variability needs to stay within the genuine distribution scores of the attacked device. This is not trivial to achieve as demonstrated in our hill-climbing attack.

Hardware circuit replication (cloning) is another attack that can be performed to compromise the system. The instrumentation of such an attack needs physical sensor node capturing and subsequently very accurate replication of the circuitry (i.e., matching as much as possible the characteristics of all integrated circuit components). In addition, if the devices are equipped with special shielding or a node capture detection mechanism is in place, such a task becomes even harder.

These attacks require further investigation.

7. APPLICATION SCENARIOS

In this section, we describe applications of physical-layer identification in all-wireless multi-hop sensor networks. We focus on protection against wormhole, Sybil and node replication attacks as well as enhancement of cryptography-based protocols for authentication.

In a wormhole attack [12], an attacker forwards packets received at one point of the network to another point that is usually multiple hops away. This is achieved by tunneling between two attackers' devices positioned at the respective points. This attack is particularly harmful to routing protocols [31] and is very challenging to detect because it can be executed by external attackers and the packet information does not need to be changed. Physical-layer identification helps identifying the attacker's device (intruder) when trying to forward packets, as the physical characteristics of the transmitted signal differ. Such detection can be achieved by a centralized or distributed approach detailed in [13].

Physical-layer identification can be used to prevent Sybil [14] and node replication (cloning) [15] attacks.

In the Sybil attack, the attacker gives several identities to the same sensor node with the purpose to fool the routing and data aggregation in the network. The replication attack consists of assigning the same (legitimate) identity to several nodes. With a physical-layer identification mechanism in place, and given the difficulty of compromising the identification, these attacks can successfully be prevented.

Finally, physical-layer identification can also be used to complement cryptography-based protocols for authenticating the communication between sensor nodes. It provides a second layer of security that cannot be easily subverted even if the attacker has compromised or is in the possession of the cryptographic keys for communication (internal attacker)[13]. An (internal) attacker who holds the cryptographic keys will not be able to authenticate to the network with her own device unless she is able to replicate the sensor node radio circuit to impersonate a legitimate device from the target network. In addition, in some scenarios, our technique can be used alone for device authentication which saves power compared to cryptography-based authentication[32, 33].

8. RELATED WORK

The proliferation of radio technologies triggered a number of research initiatives to detect illegally operated radio transmitters [1, 9, 10], device cloning [34], defective transmission devices [35] and identify wireless devices [3, 11, 36, 13, 23] by using physical characteristics of the transmitted signals [2]. Below, we present the most relevant work to ours in terms of signal similarities, features and objectives.

Hall et al. [3, 16] explored a combination of features such as amplitude, phase, in-phase, quadrature, power and DWT of the transient signal. The authors tested on 30 IEEE 802.11b transceivers from 6 different manufacturers and scored a classification error rate of 5.5%. Further work on 10 Bluetooth transceivers from 3 manufacturers recorded a classification error rate of 7% [11]. One weakness of the approach is that the classification error rate highly depended on the device's manufacturer.

Ureten et al. [23] extracted the envelop of the instantaneous amplitude by using the Hilbert transformation and classified the signals using a Probabilistic Neural Network (PNN). The method was tested on 8 IEEE 802.11b transceivers from 8 different manufacturers and registered a classification error rate of 2%-4%.

Both works differ from ours in terms of the features and type of wireless devices used. Devices from different manufacturers ease the recognition task due to significant differences in the signals. An attacker could easily compromise such a system by using a device from the same manufacturer.

Rasmussen et al. [13] explored transient length, amplitude variance, number of peaks of the carrier signal and the difference between mean and maximum value of the transient. The features were tested on 10 identical Mica2 (CC1000) sensor devices (approx. 15cm from the capturing antenna) and achieved a classification error rate of 30%. This work is the closest to ours as it considered wireless sensor devices from the same model and manufacturer. We tested our approach on the data they have used and scored a much improved classification error rate of 3.2%.

None of the above works considered the stability of their proposed features with respect to capturing distance, antenna polarization and voltage, or attacks.

Very recently, Brik et al. [26] proposed a device identification based on the variance of modulation errors. The method was tested on 100 identical 802.11b NICs (3-15 m from the capturing antenna) and achieved a classification error rate of 3% and 0.34% for k-NN and SVM classifiers respectively. No evidence about feature stability or attacks have been presented in that work. Given that only classification error rate is used to evaluate that system, we cannot compare our achieved recognition accuracy to that work. We therefore show the trade-offs of our technique with respect to that metric as well. We point out that even if our classification error rate is comparable and even lower, a direct comparison can be misleading given the different radio type and signal physical properties considered.

Our work also differs from previous work in the use of Equal Error Rate (EER) and Receiver Operating Characteristic (ROC) for performance evaluation. Prior work [3, 16, 11, 13, 23, 26] considered standard classifier (e.g., k-NN, PNN, SVM) and classification error rate as performance metric. While such a metric is appropriate for applications with well-known type and number of classes (e.g., [35]), it is not suitable for applications such as intrusion detection, device authentication, wormhole detection, etc. due to: 1) In intrusion-related applications, the number of classes (i.e., devices) is unlimited. 2) A standard classifier will classify test signals coming from a device that does not belong to the considered classes of devices to one of these classes.

We therefore use EER and ROC to quantify the accuracy of our system. It should be noted that a standard classifier can be adapted for security applications by considering doubt and outlier classes. This additional overhead however unnecessary complicates the design, and it is not scalable for large number of devices.

9. CONCLUSION

In this paper, we investigated the feasibility of transient-based identification of 802.15.4 CC2420 Tmote Sky wireless sensor nodes. We proposed a new technique for transient-based identification and we showed that it en-

ables reliable and accurate sensor node recognition, with an Equal Error Rate as low as 0.0024 (0.24%). Our system works equally well on CC1000 sensor nodes and improves previously reported results. We also investigated the performance of our technique in terms of parameters such as distance, antenna polarization and voltage. We showed that large fixed distances and variable voltage preserve fingerprint properties, whereas varying distance and antenna polarization distort the fingerprints and cause significantly lower recognition accuracy. This result limits the usability of the proposed technique in dynamic networks, however other statistical methods may be more appropriate in such scenarios. We also investigated the feasibility of impersonation and denial-of-service attacks on the recognition. We showed that if the parameters of the system are not well chosen, sensor nodes can be impersonated using a hill-climbing attack with antenna polarization. Finally, we showed that transient-based sensor node identification can be disabled by carefully structured denial-of-service attacks.

Acknowledgments

The authors thank Hansruedi Benedicker, Kasper Bonne Rasmussen and Thomas Schmid for their valuable suggestions and assistance. This work was partially supported by the Zurich Information Security Center. It represents the views of the authors.

10. REFERENCES

[1] J. Toonstra and W. Kisner, "Transient analysis and genetic algorithms for classification," in *Proc. IEEE WESCANEX*, 1995.

[2] K. Ellis and N. Serinken, "Characteristics of radio transmitter fingerprints," *Radio Science*, vol. 36, pp. 585–597, 2001.

[3] J. Hall, M. Barbeau, and E. Kranakis, "Enhancing intrusion detection in wireless networks using radio frequency fingerprinting," in *Proc. CIIT*, 2004.

[4] S. Bratus, C. Cornelius, D. Peebles, and D. Kotz, "Active behavioral fingerprinting of wireless devices," in *Proc. ACM WiSec*, 2008.

[5] J. Franklin, D. McCoy, P. Tabriz, V. Neagoe, J. Randwyk, and D. Sicker, "Passive data link layer 802.11 wireless device driver fingerprinting," in *Proc. USENIX*, 2006.

[6] T. Kohno, A. Broido, and K. Claffy, "Remote physical device fingerprinting," *IEEE TDSC*, vol. 2, no. 2, 2005.

[7] "Nmap security scanner." http://www.insecure.org/nmap/2004

[8] "Xprobe." http://www.sys-security.com

[9] J. Toonstra and W. Kisner, "A radio transmitter fingerprinting system odo-1," in *Canadian Conf. on Elect. and Comp. Engineering*, 1996.

[10] R. Hippenstiel and Y. Payal, "Wavelet based transmitter identification," in *Proc. ISSPA*, 1996.

[11] J. Hall, M. Barbeau, and E. Kranakis, "Detecting rogue devices in bluetooth networks using radio frequency fingerprinting," in *Proc. CCN*, 2006.

[12] Y. Hu, A. Perrig, and D. Johnson, "Packet leashes: A defense against wormhole attacks in wireless networks," in *Proc. IEEE INFOCOM*, 2003.

[13] K. Rassmussen and S. Capkun, "Implications of radio fingerprinting on the security of sensor networks," in *Proc. SecureComm*, 2007.

[14] J. Newsome, E. Shi, D. Song, and A. Perrig, "The sybil attack in sensor networks: analysis and defenses," in *Proc. IEEE IPSN*, 2004.

[15] B. Parno, A. Perrig, and V. Gligor, "Distributed detection of node replication attacks in sensor networks," in *Proc. IEEE S&P*, 2005.

[16] J. Hall, M. Barbeau, and E. Kranakis, "Radio frequency fingerprinting for intrusion detection in wireless networks," *Submission to IEEE TDSC (Electronic Manuscript)*, 2005.

[17] R. Bolle, J. Connell, S. Pankanti, N. Ratha, and A. Senior, *Guide to Biometrics*. Springer, 2003.

[18] "IEEE 802.15.4 standard," 2006.

[19] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[20] B. Moghaddam and A. Pentland, "Probabilistic visual learning for object representation," *IEEE PAMI*, vol. 19, no. 7, pp. 696–710, 1996.

[21] W. Zhao, R. Chellappa, and A. Krishnaswamy, "Discriminant analysis of principal components for face recognition," in *Proc. Conf. on Automatic Face and Gesture Recognition*, 1998, pp. 336–341.

[22] A. Martinez and A. Kak, "Pca versus lda," *IEEE PAMI*, vol. 23, no. 2, pp. 228–233, 2001.

[23] O.Ureten and N.Serinken, "Wireless security through rf fingerprinting," *Canadian J. Elect. Comput. Eng.*, vol. 32, no. 1, Winter 2007.

[24] "Fingeprint verification competitions (fvc)." http://bias.csr.unibo.it/fvc2006/

[25] O. Tekbas, N. Serinken, and O. Ureten, "An experimental performance evaluation of a novel radio-transmitter identification system under diverse environmental conditions," *Canadian J. Elect. Comput. Eng.*, vol. 29, no. 3, 2004.

[26] V. Brik, S. Banerjee, M. Gruteser, and S. Oh, "Wireless device identification with radiometric signatures," in *Proc. ACM MobiCom*, 2008.

[27] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proc. ACM MobiHoc*, 2005.

[28] "Gnu software radio." http://www.gnu.org/software/gnuradio/

[29] M. Strasser, C. Poepper, S. Capkun, and M. Cagalj, "Jamming-resistant key establishment using uncoordinated frequency hopping," in *Proc. IEEE S&P*, 2008.

[30] S. Weingart, "Physical security devices for computer subsystems: A survey of attacks and defenses," in *Proc. CHES*, 2000.

[31] C. Karlof and D. Wagner, "Secure routing in wireless sensor networks: attacks and countermeasures," in *Proc. IEEE Workshop on Sensor Network Protocols and Applications*, 2003.

[32] A. Wander, N. Gura, H. Eberle, V. Gupta, and S. Shantz, "Energy analysis of public-key cryptography for wireless sensor networks," in *Proc. PerComm*, 2005.

[33] G. Guimaraes, E. Souto, D. Sadok, and J. Kelner, "Evaluation of security mechanisms in wireless sensor networks," in *Proc. Systems Communications*, 2005.

[34] D. Kaplan and D. Stanhope, "Waveform collection for use in wireless telephone identification," U.S. Patent 5,999,806, 1999.

[35] B. Wang, S. Omatu, and T. Abe, "Identification of the defective transmission devices using the wavelet transform," *IEEE PAMI*, vol. 27, no. 6, pp. 696–710, 2005.

[36] O. Tekbas, O. Ureten, and N. Serinken, "Improvement of transmitter identification system for low snr transients," in *Electronic Letters*, 2004.