Physical-layer Identification of RFID Devices

Boris Danev
Dept. of Computer Science
ETH Zürich, Switzerland
boris.danev@inf.ethz.ch

Thomas S. Heydt-Benjamin
IBM Zürich Research
Laboratory, Switzerland
hey@zurich.ibm.com

Srdjan Čapkun
Dept. of Computer Science
ETH Zürich, Switzerland
capkuns@inf.ethz.ch

Abstract

In this work we perform the first comprehensive study of physical-layer identification of RFID transponders. We propose several techniques for the extraction of RFID physical-layer fingerprints. We show that RFID transponders can be accurately identified in a controlled environment based on stable fingerprints corresponding to their physical-layer properties. We tested our techniques on a set of 50 RFID smart cards of the same manufacturer and type, and we show that these techniques enable the identification of individual transponders with an Equal Error Rate of 2.43% (single run) and 4.38% (two runs). We further applied our techniques to a smaller set of electronic passports, where we obtained a similar identification accuracy. Our results indicate that physical-layer identification of RFID transponders can be practical and thus has a potential to be used in a number of applications including product and document counterfeiting detection.

1 Introduction

Passively powered Radio Frequency Identification Devices (RFID) are becoming increasingly important components of a number of security systems such as electronic passports [3], contactless identity cards [4], and supply chain systems [16]. Due to their importance, a number of security protocols have been proposed for RFID authentication [46, 25, 17], key management [31, 28] and privacy-preserving deployment [6, 29, 26, 37, 19, 14, 13]. International standards have been accepted that specify the use of RFID tags in electronic travel documents [3]. Although the literature contains a number of investigations of RFID security and privacy protocols [27, 5] on the logical level, little attention has been dedicated to the security implications of the RFID physical communication layer.

In this work, we focus on the RFID physical communication layer and perform the first study of RFID transponder physical-layer identification. We present a hardware set-up and a set of techniques that enable us to perform the identification of individual RFID transponders of the same manufacturer and model. We show that RFID transponders can be accurately identified in a controlled measurement environment based on stable fingerprints corresponding to their physical-layer properties. The measurement environment requires close proximity and fixed positioning of the transponder with respect to the acquisition antennas.

Our techniques are based on the extraction of the modulation shape and spectral features of the signals emitted by transponders when subjected to both well formed reader signals, and to out of specification reader signals. We tested our techniques on a set of 50 RFID smart cards of the same manufacturer and type and show that these techniques enable the identification of individual cards with an Equal Error Rate of 2.43% (single run) and 4.38% (two runs). We further applied our techniques to a smaller set of electronic passports, where we obtained a similar identification accuracy. We also tested the classification accuracy of our techniques, and show that they achieve an average classification error of 0% for a set of classes corresponding to the countries of issuance. We further show that our techniques produce features that form compact and computationally efficient fingerprints. Given the low frequencies of operation of the transponders in our study, the extraction of the fingerprints is inexpensive, and could be performed using a low-cost purpose-built reader.

Although the implications of physical-layer identification of RFID transponders are broad, we believe that the techniques we present can potentially find their use in the detection of cloned products and identity documents, where the (stored) fingerprints of legitimate documents are compared with those of the presented documents. Our experimental setup corresponds to this application in which the transponders are fingerprinted from close proximity and in a controlled environment.

It has been recently shown that despite numerous protections, RFIDs in current electronic documents can be successfully cloned [18, 34, 33, 47], even if they apply the full range of protective measures specified by the standard [3], including active authentication. We see our techniques as an additional, efficient and inexpensive mechanism that can be used to detect RFID cloning. More precisely, to avoid detection of a cloned document, an adversary has to produce a clone using a transponder with the same fingerprint as the original document. Although, it may be hard to perform such task, the amount of effort required is an open research problem. We discuss two methods of applying RFID physical-layer identification to cloning detection and compare it to other anti-cloning solutions, like those based on physically-unclonable functions (PUFs) [12].

Our results show the feasibility of RFID transponder fingerprinting in a controlled environment. Using the proposed methods is not enough to extract the same or similar fingerprints from a larger distance (e.g., 1 meter). In our experiments, such remote feature extraction process resulted in incorrect identification. Therefore, we cannot assert that chip holder privacy can be compromised remotely using our techniques. This result further motivates an investigation of physical-layer features of RFID transponders that would allow their remote identification, irrespective of (e.g., random) protocol-level identifiers that the devices use on the logical communication level. Our current results do not allow us to conclude that such distinguishable features can be extracted remotely.

The remainder of this paper is organized as follows. In Section 2, we present our system model and investigation parameters. In Section 3, we detail our fingerprinting setup (i.e., a purpose-built reader), signal capturing process and summarize the data acquisition procedure and collected data. The proposed features for transponder classification and identification are explained in Section 4 and their performance is analyzed in Section 5. We discuss an application of our techniques to document counterfeiting detection in Section 6, make an overview of background and related work in Section 7 and conclude the paper in Section 8.

2 Problem and System Overview

In this work, we explore physical-layer techniques for detection of cloned and/or counterfeit devices. We focus on building physical-layer fingerprints of RFID transponders for the following two objectives:

1. RFID transponder classification: the ability to associate RFID transponders to previously defined transponder classes. In the case of identity documents classes might, for example, be defined based on the country that issued the document and the year of issuance.
2. RFID transponder identification: the ability to identify same model and manufacturer RFID transponders. In the case of identity documents, this could mean identifying documents from the same country, year and place of issuance.

A classification system must associate unknown RFID transponder fingerprints to previously defined classes C. It performs "1-to-C" comparisons and assigns the RFID fingerprint to the class with the highest similarity according to a chosen similarity measure (Section 5.1). This corresponds to a scenario in which an authority verifies whether an identity document belongs to a claimed class (e.g., country of issuance).

An identification system typically works in one of two modes: either identification of one device among many, or verification that a device's fingerprint matches its claimed identity [8]. In this work, we consider verification of a device's claimed or assumed identity. This corresponds to a scenario in which the fingerprint of an identity document (e.g., passport), stored in a back-end database or in the document chip, is compared to the measured fingerprint of the presented document. The verification system provides an Accept/Reject decision based on a threshold value T (Section 5.1). Identity verification requires only "1-to-1" fingerprint comparison and is therefore scalable in the number of transponders.

In this study we use a single experimental setup for examination of both classification and identification. Our setup consists of two main components: a signal acquisition setup (i.e., a purpose-built RFID reader) (Section 3) and a feature selection and matching component (Section 4). In our signal acquisition setup we use a purpose-built reader to transmit crafted signals which then stimulate a response from the target RFID transponders. We then capture and analyze such responses. In particular, we consider transponder responses when subjected to the following signals from the reader: standard [4] transponder wake-up message, transponder wake-up message at intentionally out-of-specification carrier frequencies, a high-energy burst of sinusoidal carrier at an out-of-specification frequency, and a high-energy linear frequency sweep.

To evaluate the system accuracy, we make use of two different device populations (Table 1). The first population consists of 50 "identical" JCOP NXP 4.1 smart cards [2] which contain NXP RFID transponders (ISO 14443, HF 13.56 MHz). We chose these transponders since they are popular for use in identity documents and access cards, and because they have also been used by hackers to demonstrate cloning attacks against e-passports [47]. The second population contains 8 electronic passports from 3 different countries[1]. These two populations allow us to define different transponder classes (e.g., 3 issuing countries, and a separate class for JCOP cards) for classification and include a sufficient set of identical transponders to quantify the identification accuracy of the transponders of the same model and manufacturer.

In summary, in this work, we answer the following interrelated questions:

1. What is the classification accuracy for different classes of transponders, given the extracted features?
2. What is the identification accuracy for transponders of the same model and manufacturer, given the extracted features?
3. How is the classification and identification accuracy affected by the number of signals used to build the transponder fingerprint?
4. How stable are the extracted features, across different acquisition runs and across different transponder placements (relative to the reader)?

3 Experimental Setup and Data

In this section, we first describe our signal acquisition setup. We then detail the different types of experiments we performed and present the collected datasets from our population of transponders.

3.1 Hardware Setup

Figure 1 displays the hardware setup that we use to collect RF signals from the RFID devices. Our setup is essentially a purpose-built RFID reader that can operate within the standardized RFID communication specifications [4], but can also operate out of specifications, thus enabling a broader range of experiments. The setup consists of two signal generators, used for envelope generation (envelope generator) and for signal modulation (modulation generator), and of transmitting and acquisition antennas. The envelope generator is fed with a waveform that represents the communication protocol wake-up command[2] required for initiating communication with RFID transponders. The envelope waveform is then sent to the modulation generator and is modulated according to the ISO/IEC 14443 protocol Type A or B, depending on the transponders being contacted. The modulated signal is then sent over a PCB transmitting antenna. Finally, the wake-up signal and the response from the transponder are received at the acquisition antenna and captured at the oscilloscope. The separation of the envelope generation and modulation steps allowed us to independently vary envelope and modulation characteristics in our experiments.

In order to collect the RF signal response, we built a ”sandwich” style antenna arrangement (Figure 2b) where an acquisition antenna is positioned between the transmission antenna and the target RFID transponder. An wooden platform holds the transmission and acquisition antennas in a fixed position to avoid changes in antenna polarization[3]. The platform is separated from the desk by a non-metallic wooden cage. The transmission and acquisition antennas are both connected to an oscilloscope. We used the RF signal on the transmission antenna to trigger the acquisition and then record the transponder's response at the acquisition antenna. It should be noted that we can also observe the transponder's response at the transmission antenna, however as the acquisition antenna had a higher gain than the transmission antenna, we used the described setup to obtain better signal-to-noise ratio.

3.2 Performed Experiments

Using the proposed setup, we performed four major experiments:

Experiment 1 (Standard): In this experiment we initiate communication with the transponders as defined by Type A and B protocols in the ISO/IEC 14443 standard. The envelope generator generates the Type A and B envelopes and the modulation generator modulates the signal at a carrier frequency F_c= 13.56 MHz, using 100% ASK for Type A and 10% ASK for Type B at the nominal bit rate of $F_b \sim$ 106kbit/s.[4] The experiment consists of the following steps: a period of unmodulated carrier is transmitted to power the transponder at which time the oscilloscope begins recording the data. The carrier is then modulated according to the envelope such that it corresponds to a WUQA (Type A) or WUQB (Type B) wake-up command. When the commands are no longer transmitted, an unmodulated period of carrier is then sustained while the oscilloscope records the response from the transponder. The carrier is turned off between each

[1] The small quantity of the electronic passports used in the experiments is due to the difficulty of finding people who are in possession of such passports and at the same time willing to allow experimentation on them.

[2] ISO/IEC 14443 for RFID communication defines two different communication protocols, Type A and B, which use different wake-up commands: WUQA and WUQB, respectively.

[3] It has been observed that such changes can reduce the identification accuracy [11].

[4] For 100% ASK modulation we used pulse modulation as standard built-in amplitude modulation (AM) in our generators could not reach the required precision.

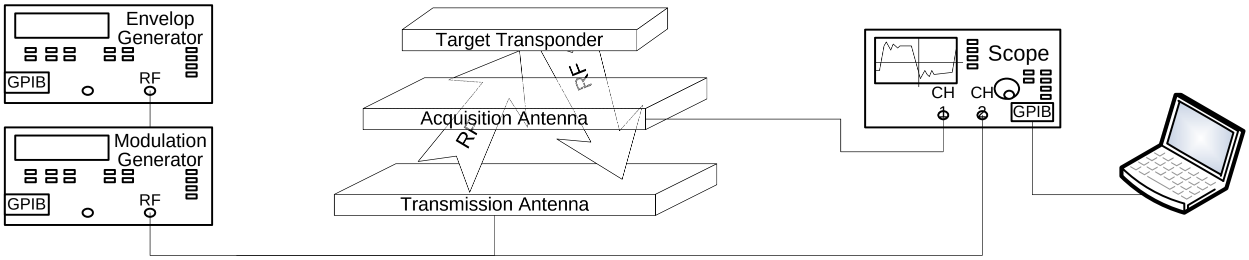

Figure 1: Signal acquisition setup. Envelope and modulation generators generate wake-up signals that initiate the response from the RFID transponder. This wake-up signal is transmitted by the transmitting antenna. The acquisition antenna captures both the wake-up signal and the response from the transponder. The signal from the acquisition antenna is then captured and recorded by the oscilloscope.

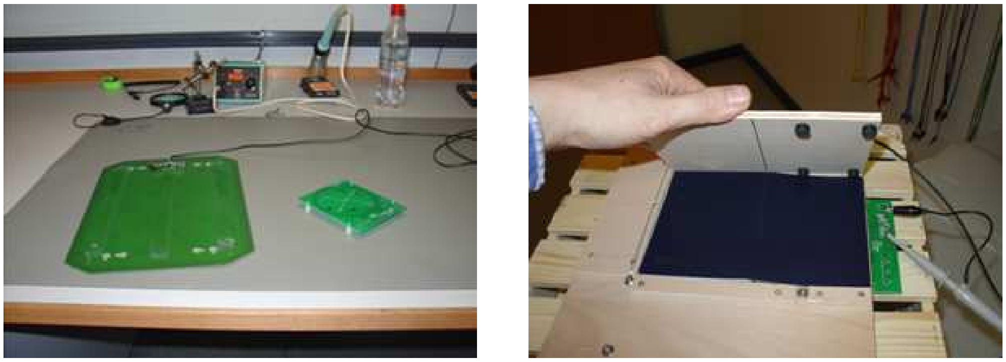

(a)

(b)

Figure 2: a) Transmission and acquisition antennas. b) An electronic identity document being placed in the fingerprinting setup.

observation to ensure the transponder reboots each time. Figures 3a and 3b show the collected samples from Type A and Type B RFID transponders, respectively. This experiment enables us to test if the transponder's responses can be distinguished when they are subjected to standard signals from the reader.

Experiment 2 (Varied F_c): In this experiment, we test transponder responses to the same signals as in Experiment 1, but on out of (ISO/IEC 14443) specification carrier frequencies. Instead of on F_c=13.56 MHz, our setup transmits the signals on carrier frequencies between F_c=12.96 MHz and 14.36 MHz. Figures 3c and 3d display sample transponder responses to signals on F_c=13.06 MHz. We expect the variation in the transponder responses to be higher when they are subjected to out of specification signals, since the manufacturers mainly focus on transponder responses within the specified frequency range.

Experiment 3 (Burst): In this experiment, we tested transponder responses to bursts of RF energy. We subjected the transponders to 10 cycles (2 μs) of non-modulated 5 MHz carrier at an amplitude of V_{pp}=10 V (the maximum frequency and amplitude supported by our generators while in burst mode). Figure 4a shows a sample transponder response to such an RF burst. This experiment tests the response of transponders to an additional out-of-specification signal. We expect to see variation in different transponders' responses for a variety of reasons. For example since each transponder's antenna and charge pump is unique, we believe that during power-up it will present a unique modulation of an activating field.

Experiment 4 (Frequency Sweep): This experiment consists of observing transponder responses to a non-modulated carrier linear sweep from 100 Hz to 15 MHz at an amplitude of V_{pp}=10 V (as measured at transmitting antenna). The duration of the sweep is fixed to the maximum allowed by our generator, 10 ms. In this test we examine how the transponders react to many different frequencies. Among other things, such an experiment provides information about the resonances of the RF circuitry in each transponder. Figure 4b shows a sample transponder response to a frequency sweep. Note the different shape artifacts.

We found that samples collected from Experiment 2 were well suited for transponder classification, whereas

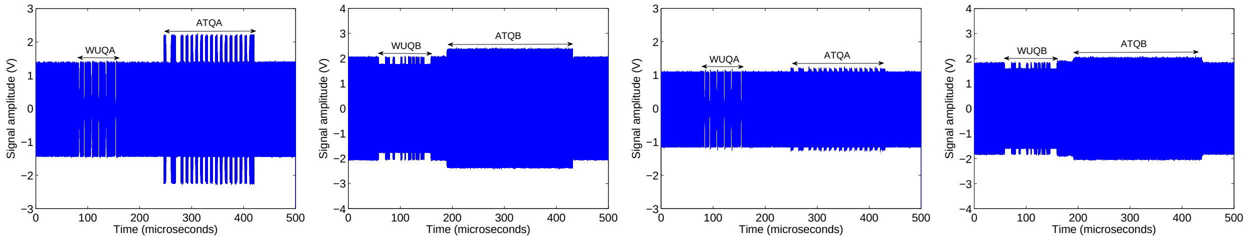

(a) (b) (c) (d)

Figure 3: Experiment 1: Type A (a) and Type B (b) RFID transponder responses to WUQA and WUQB commands sent on the ISO/IEC 14443 specified carrier frequency (F_c=13.56 MHz). Experiment 2: Type A (c) and Type B (d) RFID transponder responses ATQA and ATQB to WUQA and WUQB commands respectively sent on an out of ISO/IEC 14443 specification carrier frequency (F_c=13.06 MHz)

.

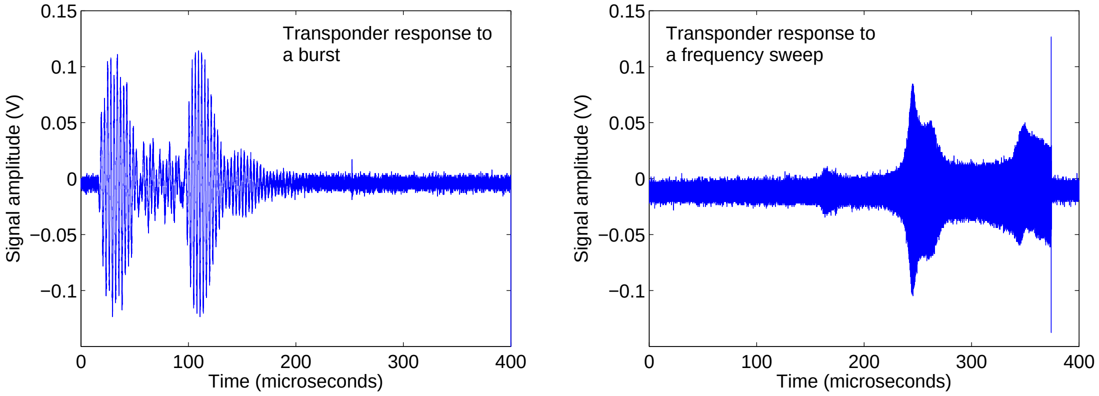

(a) (b)

Figure 4: a) Experiment 3: transponder response sample to a non-modulated 5 MHz carrier in duration of 10 cycles. b) Experiment 4: transponder response sample to a non-modulated carrier linear sweep from 100 Hz to 15 MHz. The duration of the sweep is 10 ms.

those collected from Experiments 3 and 4 were better suited for identification of individual RFID transponders. We discuss this result at greater length in Section 4.

3.3 Collected Data

Using the proposed setup, we performed the experiments described in Section 3.2 and collected samples from 8 passports and 50 JCOP NXP 4.1 smart cards (same model and manufacturer). The types of devices used in the experiments are shown in Table 1. For the privacy of our research subjects we arbitrarily labeled the passports as ID1 to 8. To further protect their privacy we give the country and place of issuance under the pseudonyms C1 to C3 and P1 to P6 respectively.

Our data collection procedure for a single experiment ”run” was as follows: We positioned the target RFID device on the experimental platform with all other transponders being at an out-of-range distance from the activating field. We then placed a heavy non-metallic weight on top of the transponder to position it firmly and horizontally on the platform. For each device we then performed Experiments 1-4 at fixed acquisition timing offset and sampling rate and saved the samples to a disk for later analysis. For each transponder we performed two runs, completely removing and replacing the target transponder on the experimental platform between runs. This ensures that extracted features are stable across repositioning.

In each iteration of Experiment 2 we collected 4 samples per run for 14 different carrier frequencies starting from F_c=12.96 up to 14.36 MHz with a step of 100 KHz. This resulted in 64 samples per transponder per run. In Experiments 3 and 4 we collected 50 samples per device per run.

Table 1: RFID device populations (passports and JCOP NXP smart cards).

Type	Number	Label	Country	Year	Place of Issue
Passport	2	ID1, ID2	C1	2006	P1
	1	ID3	C1	2006	P2
	1	ID4	C1	2006	P3
	1	ID5	C1	2007	P4
	1	ID6	C2	2008	P5
	1	ID7	C3	2008	P6
	1	ID8	C1	2008	P1
JCOP	50	J1..J50	JCOP NXP 4.1 cards (same model and manufacturer)		

4 Feature Extraction and Selection

The goal of the feature extraction and selection is to obtain distinctive fingerprints from raw data samples collected in the proposed experiments, which most effectively support the two objectives in our work, namely classification and identification. In this section, we detail the extraction and matching procedures of two types of features effective for that purpose: modulation-shape features (Section 4.1) and spectral PCA features (Section 4.2). We also investigated the use of some timing features, such as the time interval within which the transponder responds to an WUQ command and the duration of that response (Figure 5a). These timing features performed poorly in both tasks, hence in this work we focus on the modulation-shape and spectral features.

4.1 Modulation-shape Features

In this section, we describe the extraction and matching procedures for the features extracted from the shape of the modulated signal of the transponder responses at a given carrier frequency F_c (Experiment 1&2). Figure 5 b) shows the shape of the On-Off keying modulation for the JCOP NXP 4.1 card for the first packet in a transponder's response to a wake-up command. All Type A transponders in our study had a logically identical first packet.

For a given transponder, the features of the modulated signal are extracted from the captured transponder response (see Figure 3) denoted as $f(t,l)$, using Hilbert transformation. Here, $f(t,l)$ is the amplitude of the signal l at time t. The Hilbert transformation is a common transformation in signal processing used to obtain the signal envelope [38].

In Step (i), we apply Hilbert transformation on $f(t,l)$ to obtain $H(t,l)$:

$$H(t,l) \quad = \quad \text{Hil}(f(t,l)) \tag{1}$$

where Hil is a function implementing the Hilbert transform [36].

In Step (ii), the starting point of the modulation in $H(t,l)$ is determined using the variance-based threshold detection algorithm described in [40]. The end point is fixed to a pre-defined value (see Section 5) and then the modulation-shape fingerprint is extracted.

Feature matching between a reference and a test fingerprints is performed using standardized Euclidean distance, where each coordinate in the sum of squares is inverse weighted by the sample variance of that coordinate [35].

4.2 Spectral Features

In this section, we describe the extraction and matching of spectral features from data collected from Experiments 3 (Burst) and 4 (Sweep) (Section 3.2).

Both frequency sweep and burst data samples are extremely high-dimensional: each sweep data sample contains 960000 points (dimensions) and each burst data sample contains 40000. Such high-dimensional data typically contain many noisy dimensions which are detrimental to finding distinctive features. Therefore, it is critical to remove the noise as much as possible from the raw data samples.

We explored two basic approaches to solve the dimensionality problem. In the first approach, we considered transforming the data in the frequency domain by means of the Fast Fourier Transform (FFT) and remove the high frequencies (usually considered noisy) by filtering. However, matching experiments using direct vector similarity measures such as Euclidean and Cosine distance failed to produce distinctive enough features. This may be because in removing the high frequencies we are also removing frequencies that contain discriminative capabilities. Such behavior is commonly noticed in biometrics research [10]. In the second approach we down-sampled the signal at different rates in order to reduce the dimensionality. We then transformed the data in the frequency domain by FFT and applied standard vector similarity measures. Again reducing the dimensionality in this way did not prove to be effective in extracting sufficiently dis-

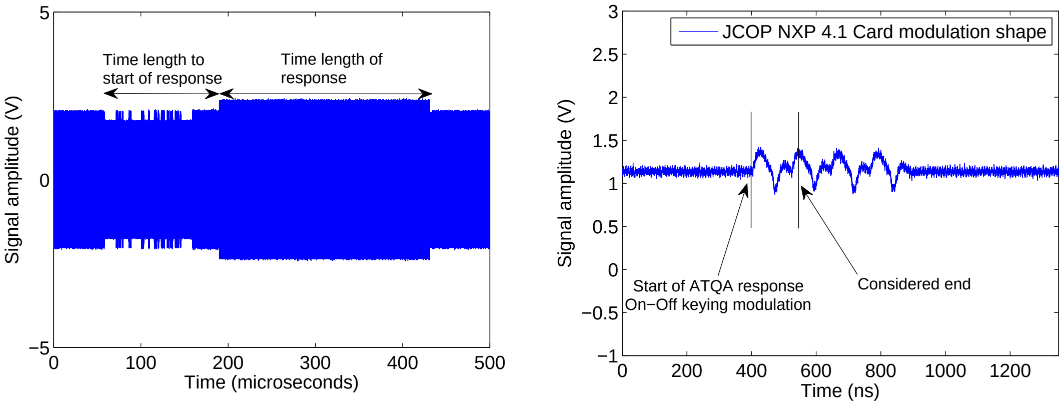

(a)

(b)

Figure 5: a) Timing features extracted from Type B transponder responses. b) Modulation-shape features.

criminative features.

To overcome the above problems, we use a modification of Principal Component Analysis (PCA) for high-dimensional data [7], that reduces data dimensionality by discarding dimensions that do not contribute to the total covariance. Given that the number of dimensions is very high, orders of magnitude higher than the number of data samples we can process, a standard PCA procedure cannot be applied. In the following subsection, we briefly describe the used PCA modification.

4.2.1 Feature Extraction and Matching

For a given RFID device, spectral PCA features are extracted from N captured samples using a linear transformation derived from PCA for high-dimensional data. We denote a signal by $f(t,l)$, where $f(t,l)$ is the amplitude of the signal l at time t. The features are extracted in the following three steps:
In Step (i), we apply a one-dimensional Fourier transformation on $f(t,l)$ to obtain $F(\omega,l)$:

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

where M is the length of signal considered and $0 \leq t \leq M-1$ is time. We then remove from $F(\omega,l)$ its DC component and the redundant part of the spectrum; we denote the remaining part of the spectrum by $\vec{s_l}$.
In Step (ii), a projected vector $\vec{g_l}$, also called a spectral feature, is extracted from the Fourier spectrum using a PCA matrix W_{PCA}:

$$\vec{g_l} = W_{PCA}^t \vec{s_l} \qquad (3)$$

The feature extraction from N captured samples for a given transponder is then given by $G = W_{PCA}^t S$ where G is an array of $\vec{g_l}$ and S is a matrix $S = [\ \vec{s_0}\ ..\ \vec{s_l}\ ..\ \vec{s_N}\ ]$.
Finally, in Step (iii), the feature template (fingerprint) h used for matching is computed:

$$h = \{\hat{G}; \Sigma_G\} \qquad (4)$$

where $\hat{G}$ denotes the mean vector of G and Σ_G denotes the covariance matrix of G. The number of captured samples N used to build the feature template and the number of projected vectors in W_{PCA} (i.e., the subspace dimension) are experimentally determined.

Mahalanobis distance is used to find the similarities between fingerprints[5]. The result of matching a reference h^R and a test h^T feature templates is a matching score, calculated as follows.

$$scr(h^R,h^T) = min(\sqrt{(\hat{G}^T - \hat{G}^R)^t \Sigma_{G^R}^{-1} (\hat{G}^T - \hat{G}^R)}, \sqrt{(\hat{G}^T - \hat{G}^R)^t \Sigma_{G^T}^{-1} (\hat{G}^T - \hat{G}^R)}) \qquad (5)$$

Values of the matching score closer to 0 indicate a better match between the feature templates. The proposed matching uses the mean and covariance of both test and reference templates. It also ensures the symmetric property, that is $scr(h^R,h^T) = scr(h^T,h^R)$.

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

4.2.2 PCA Training

In order to compute the eigenvalues and corresponding eigenvectors of the high-dimensional data (the number

[5] We discovered that the feature templates are distributed in ellipsoidal manner and therefore use Mahalanobis distance that weights each projected sample according to the obtained eigenvalues.

of samples $\ll$ the number of dimensions), we used the following lemma:

Lemma: For any $K \times D$ matrix W, mapping $x \rightarrow Wx$ is a one-to-one mapping that maps eigenvectors of W^TW onto those of WW^T.

Here W denotes a matrix containing K samples of dimensionality D. Using this lemma, we can first evaluate the covariance matrix in a lower space, find its eigenvectors and eigenvalues and then compute the high-dimensional eigenvectors in the original data space by normalized projection [7]. Based on this description, we compute the PCA matrix W_{PCA}=$[\vec{u_1}\vec{u_2}\ldots\vec{u_i}]$ by solving the eigenvector equation:

$$(\frac{1}{K}X^TX)(X^T\vec{v_i}) = \lambda_i(X^T\vec{v_i}) \qquad (6)$$

where X is the training data matrix $K \times D$ and $\vec{v_i}$ are the eigenvectors of XX^T. We then compute the eigenvectors of our matrix $\vec{u_i}$ by normalizing:

$$\vec{u_i} = \frac{1}{\sqrt{K\lambda_i}}(X^T\vec{v_i}) \qquad (7)$$

It should be noted that other algorithms like probabilistic PCA (e.g., EM for PCA) can potentially be also used given the fact that we discovered that only 5-10 eigenvectors are predominant. We intend to investigate these as a part of our future work.

5 Performance Results

In this section, we present the performance results of our fingerprinting system. First, we review the metrics that we use to evaluate the classification and identification accuracy.

5.1 Evaluation Metrics

As a metric for classification, we adopt the average classification error rate, defined as the percentage of incorrectly classified signatures to a predefined set of classes of signatures (e.g., countries). We used the 1-Nearest Neighbor rule [7] for estimating the similarity between testing and reference signatures from a given class; that is, a testing signature is matched to all reference signatures from all classes and assigned to the class with nearest distance similarity. It should be noted that more sophisticated classifiers can be devised such as Support Vector Machines (SVM), Probabilistic Neural Networks (PNN) [7]. However these classifiers require more training which we do not consider in this work.

As metrics for identification, we adopt the Equal Error Rate (EER) and the Receiver Operating Characteristic (ROC) since these are the most agreed metrics for evaluating identification systems [8]. The False Accept Rate (FAR) and the False Reject Rate (FRR) are the frequencies at which the false accept and the false reject events occur. The FAR and FRR are closely related to each other in the Receiver Operating Characteristic (ROC). ROC is a curve which allows to automatically compute FRR when the FAR is fixed at a desired level and vice versa [8]. 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 identification systems [1]. The operating threshold value at which the EER occurs is our threshold T for an Accept/Reject decision.

To increase the 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 = 1%).

5.2 Classification Results

In this section, we present the results of the classification using modulation-shape and spectral features. In this evaluation, we consider all our passport samples and 5 of the JCOP NXP 4.1 cards. Here, the identity documents ID1, ID2, ID3, ID4, ID7, ID8 (see Table 1) and the JCOP cards implement Type A communication protocol, whereas ID5 and ID6 use Type B protocol. It is interesting to notice that within the same country class (C1) we have documents with two different communication protocols (ID1-ID4 and ID8 implement Type A, whereas ID5 implements Type B protocol).

5.2.1 Classification using Modulation-shape Features

The modulation-shape features described in Section 4 show the discriminant artifacts in the transponder's response. In particular, we discovered that these artifacts (shapes) vary from one transponder to another on out-of-specification carrier frequencies.

Figure 6 shows the modulation envelope shapes of the initial sequence of the RFID transponder's response after Hilbert transformation for 4 different classes of Type A protocol devices. These were recorded at an out of specification carrier frequency F_c=13.16MHz. Visual inspection shows that the modulation shapes not only differ from class to class but also are stable within different runs.

In order to quantify these observations more precisely, we considered classification with 3 classes (2 countries + JCOP cards) with all fingerprints from two different runs. The classification process was repeated 8 times with 8 different reference fingerprints per class for validation.

Table 2: Classification using modulation-shape features (Experiment 2)

Number of Classes	Class structure	Average Classification Error Rate
3	(C1),(C2),(JCOP)	0%
4	(ID1,ID3,ID4,ID8), (ID2), (ID7), (JCOP)	0%
2	(ID5-C1),(ID6-C3)	0%

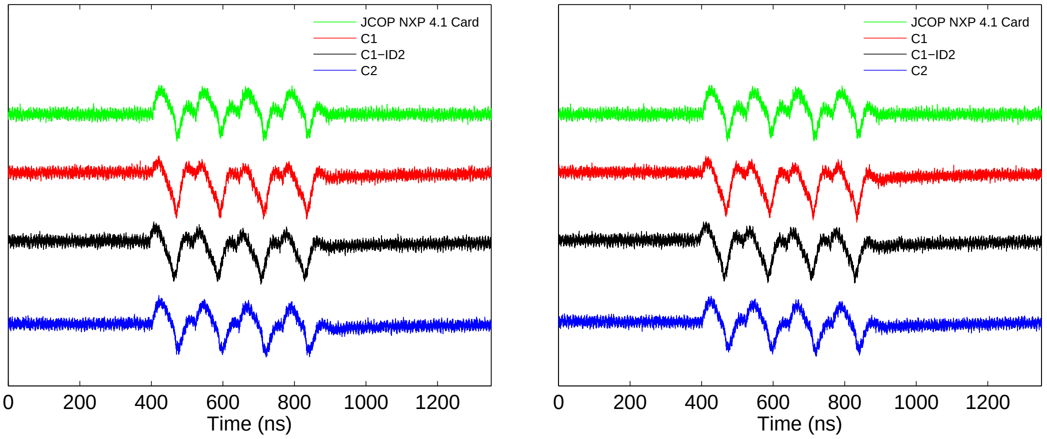

(a)

(b)

Figure 6: Modulation shape of the responses of 4 different classes (C1),(C1-ID2),(C2),(JCOP): a) first run b) second run. In each run, the sample transponders were freshly placed in the fingerprinting setup. These plots show the stability of the collected modulation-shape features across different runs.

The results show perfect separability of the classes with average classification error rate of 0%. In addition, after detailed inspection of the modulation-shape features we discovered that ID2 from C1 differs significantly from the representatives of that class. We therefore formed a new classification scenario with 5 classes and obtained again a classification error rate of 0%. It is an interesting result given that ID1 and ID2 are issued by the same country, in the same year and place of issue. However, their transponders are apparently different. The modulation-shapes of ID1,ID3 and ID4 from C1 could not be further distinguished using the combination of modulation-shape features and Euclidean matching. Table 2 shows the results.

Similar to Type A, the 2 Type B transponders from two different countries (C1,C3) available in our population showed complete separability with classification error rate of 0%. We acknowledge that our data set is insufficient due to the difficulty of obtaining e-passports. We believe however that our results are promising to stimulate future work with a larger set of e-passports.

In summary, the modulation shapes at an out-of-specification carrier frequency are successful in categorizing different classes of transponders (e.g., countries). They are quickly extractable and stable across different runs. For the classification task, there is no need of statistical analysis in contrast with the proposed spectral features analyzed in the next sections. An additional advantage is that specialized hardware is not required as current RFID readers can be easily adapted.

5.2.2 Classification using Burst and Sweep Spectral Features

We also performed classification using burst and sweep spectral features (Experiment 3 & 4) on the same set of classes as with modulation-shape features (Table 2). Similar to the modulation-shape features, this classification achieved a 0% classification error rate on the proposed classes. Moreover, using the spectral features we were also able to distinguish individually each of our 9 identity documents with an EER=0%, i.e. we were able to verify the identify of each individual document with an accuracy of 100% with FRR=FAR=0%. This result motivated us to estimate the identification accuracy of spectral features on a larger set of identical (of the same make and model) transponders.

5.3 Identification results

In this section we present the results of the identification capabilities of the (burst and sweep) spectral features for

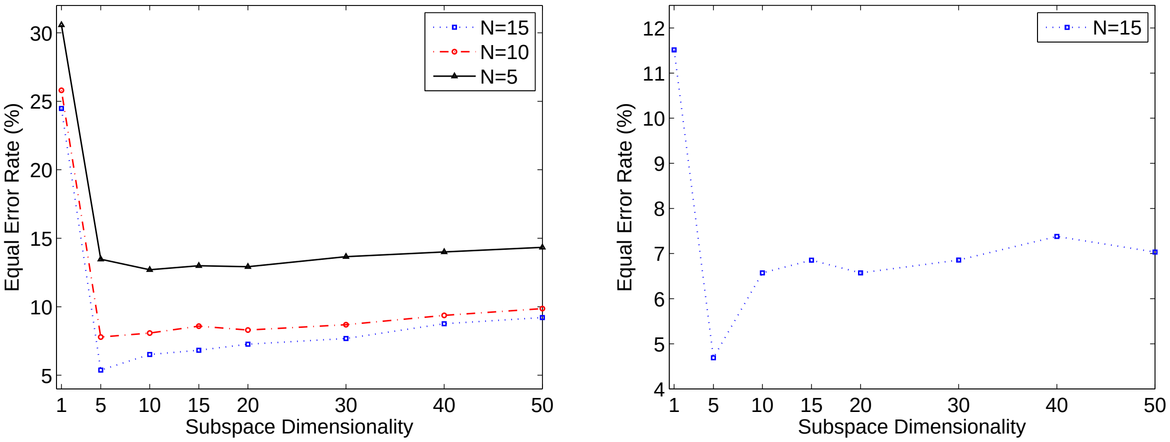

(a)

(b)

Figure 7: Spectral features identification accuracy for different number of samples N used to built the fingerprint and for different subspace dimensions: a) burst spectral features, b) sweep spectral features. 50 identical (same manufacturer and model) transponders are used in the computation.

our data population (50 identical JCOP NXP 4.1 cards). We adopt the following approach. We first evaluate the accuracy over the data collected in a single run of the experiment (Section 5.3.1 and 5.3.2). We then quantify the feature stability of the spectral features by considering samples from two independent runs together (Section 5.3.3).

We validate our results using cross-validation [7]. We measured 50 samples per transponder per run of which we use 5-10 samples for training and the remaining 40-45 samples for testing depending on the number of samples N used to build the fingerprint. The training and testing data are thus separated and allow validation of the identification accuracy.

5.3.1 Identification using Burst Spectral Features

In this evaluation, we consider the samples from the burst dataset, from a single experiment run (Experiment 3) in order to obtain a benchmark accuracy. We varied two parameters: the number of samples N used to build the feature templates (fingerprints) and the dimension of the PCA subspace used to project the original features into. The dimension of the PCA subspace is also related to the feature template size which we discuss below.

The results of this analysis are presented in Figure 7a for different N and subspace dimensionality. The dimension of the features before the projection is 19998. The results show the EER of the system reaching 0.0537 (5.37%) for N=15. This means that our system correctly identifies individual identical transponders with an accuracy of approximately 95% (GAR at the EER operating point) using the features extracted from the burst samples. We later show that this accuracy is preserved in cross-matchings between different runs. Table 3 summarizes the underlying data, namely the number of samples N, total genuine and imposter matchings performed for EER computation[6], Accept/Reject threshold, EER and confidence interval (CI).

The results in Figure 7a also confirm that using the first 5 eigenvectors to project and store the feature template provides the highest accuracy. Our proposed features therefore form compact and computationally efficient fingerprints (see Section 5.4).

5.3.2 Identification using Sweep Spectral Features

Similarly to the above analysis, we considered the first run of samples from the sweep experiment (Experiment 4) dataset. For computational reasons, we did not consider the entire sample. Instead, we extracted the spectral features from the part of the sample between 220 to 270 microseconds. As it can be seen in Figure 4, this part contains the biggest shape changes in the frequency sweep. This decision reduced the considered space to 100000 points which allowed reasonably fast feature extraction (26 s per sample). This clearly excludes some discriminant information from our analysis, and future work should include other sections of the sample signals.

The results are presented in Figure 7b for N=15 and

[6] The number of genuine and imposter matchings depends on the number of available fingerprints per transponder. For N=10, we are able to built 4 different fingerprints with the testing data within a run. This results in 6 different matchings of fingerprints from the same device (i.e., genuine matchings) and 392 different matchings of fingerprints from different transponders (i.e., imposter matchings). For 50 transponders, this makes 300 genuine and 19600 imposter matchings.

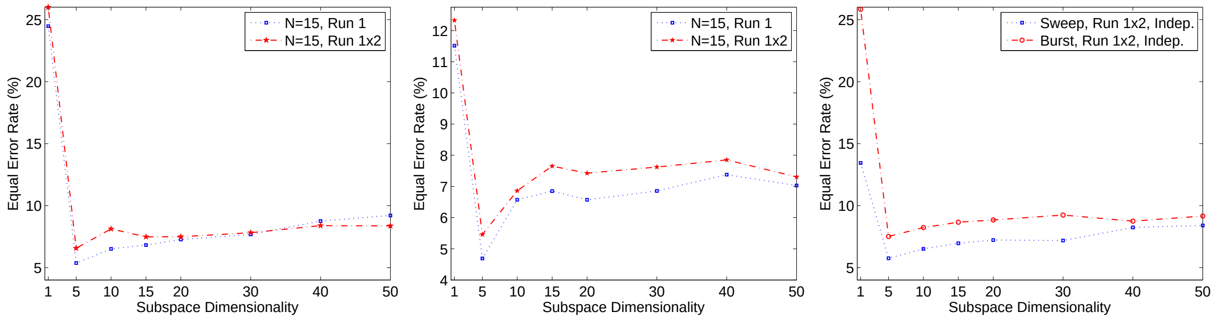

(a)

(b)

(c)

Figure 8: Feature stability in identification: a) burst spectral features b) sweep spectral features. 50 identical (same manufacturer and model) transponders are used in the experiments. c) burst and sweep spectral features on independent transponder sets for training and testing; 20 transponders are used for training and 30 transponders - for testing; N=15.

different subspace dimensions. The dimension of the original features before projection is 49998. We computed the EER for N=15 (see Burst analysis in Section 5.3.1). The obtained EER is 0.0469 (4.69%), when using the first 5 eigenvectors to project and store the feature template. The obtained accuracy is therefore similar to the one obtained with the burst features, i.e. our system correctly identifies the individual identical transponders with an accuracy of approximately 95% (GAR at the EER point). Table 3 shows the confidence intervals.

5.3.3 Feature Stability

In the previous sections we have analyzed the identification accuracy using burst and sweep spectral features within a single experiment run. This allows us to have a benchmark for estimating the stability of the features. In particular, we performed the following stability analysis:

1. Using the linear transformations W_{PCA} obtained in the first run, we selected 4 feature templates (2 from each run) and computed again the EER by considering only the cross matching scores of fingerprints from different runs[7]. The process was repeated 3 times with different feature templates from the two runs to validate the feature stability.

2. We trained the system over the first 20 transponders and then used the obtained linear transformation to estimate the accuracy over the remaining 30 transponders. This analysis tests the stability of the obtained linear transformations to discriminate independent transponder populations[8].

Figure 8 compares the EER accuracy obtained with the first run (Run 1) and the accuracy obtained by mixing fingerprints of both runs (Run 1×2) for a fixed N=15. Table 3 displays the confidence interval for subspace dimension of 5 eigenvectors. The obtained EERs do not show a statistically significant difference between the two experiments for both the burst and sweep features using 4-fold validation.

Figure 9 displays the EER accuracy obtained using independent transponder sets for training and testing for a fixed N=15. Here, the fingerprints from both runs are mixed as in the previous analysis. Table 4 summarizes the numeric results together with confidence intervals of the EER. Even though the testing population (30 transponders) is smaller, we observe that the sweep features do not show any significant accuracy deviation from the benchmark accuracy on Run 1×2 (Table 3). On the other hand, the burst features slightly decreased the accuracy on average (Table 3). The reason for this might be that 20 different transponders are not sufficient to train the system; however, we cannot assert this with certainty.

5.3.4 Combining Sweep and Burst Features

Given that the identification accuracies of both burst and sweep spectral features are similar; in order to fully characterize the identity verification we computed the ROC curves for the burst and sweep features as shown in Figure 9b. We notice that while the EERs are similar, the curves exhibit different accuracies at different FARs. In particular, for low FAR≤1% the sweep features show lower GAR.

The burst and sweep features discriminate the fingerprints in a different way, and therefore these features can be combined in order to further increase the accuracy. Such combinations are being researched in multi-modal

[7] This procedure is required in order to remove any possible bias from cross matching scores of fingerprints from the same run. We point out that this results in a reduced number of genuine and imposter matchings for the EER computation, 200 and 9800 respectively (see Table 3).

[8] The motivation behind this division (20 vs. 30) is that it gives reasonable number of transponders for both training and testing.

Table 3: Summary of accuracy for the 5-dimensional spectral features (50 transponders).

Type	Run	N	Test matchings		Threshold T	EER (%)	EER CI (%)		Validation
			Genuine	Imposter			lower	upper	
Burst	1	15	150	11025	1.88	5.37	4.38	6.36	4-fold
	1	10	300	19600	2.91	7.79	5.29	10.28	4-fold
	1	5	300	19600	7.56	13.47	13.22	13.72	4-fold
	1x2	15	200	9800	2.64	6.57	6.25	6.89	4-fold
Sweep	1	15	150	11025	1.68	4.69	3.65	5.74	4-fold
	1x2	15	200	9800	1.93	5.46	5.08	5.84	4-fold

Table 4: Accuracy when independent sets are used for training (20) and testing (30) transponders.

Type	Run	N	Test matchings		Threshold T	EER (%)	EER CI (%)		Validation
			Genuine	Imposter			lower	upper	
Burst	1x2	15	120	3480	2.78	7.33	6.01	8.65	3-fold
Sweep	1x2	15	120	3480	2.03	5.75	5.45	6.05	3-fold

biometrics [42] where different ”modalities” (e.g., fingerprint and vein) are combined to increase the identification accuracy and bring more robustness to the identification process [42].

A number of integration strategies have been proposed based on decision rules [32], logistic functions to map output scores into a single overall score [24], etc. Figure 9 shows the EERs and ROC curves of feature combination by using the sum as an integration function. The overall matching score between a test and a reference template is the sum of the matching scores obtained separately for the burst and sweep features. Table 5 summarizes the results.

For the benchmark datasets (Run 1), we observe significant improvement of the accuracy reaching an EER=2.43%. The improvement is also significant for all target FARs (e.g., 0.1%, 1%) as shown in Figure 9b. We also observe a statistically significant improvement on using fingerprints from both Run 1 and 2. The accuracy is slightly lower (EER=4.38%). These results motivate further research on feature modalities and novel integration strategies.

5.4 Summary and Discussion

In this section, we have experimentally analyzed the classification and identification capabilities of three different physical-layer features with related signal acquisition, feature extraction and matching procedures.

The results show that classification can successfully be achieved using the modulation shape of the transponder's response to a wake-up command at an out-of-specification frequency (e.g., F_c=13.06 MHz). This technique is fast, does not require special hardware and can be applied without statistically training the classification process.

For identification, we proposed using spectral features extracted from the transponder's reaction to purpose-built burst and linear frequency sweep signals. Our proposed signal acquisition and feature extraction/matching techniques achieved separately an identification accuracy of approximately EER=5% over 50 identical RFID transponders. The proposed features are stable across acquisition runs. In addition, our spectral features showed that they can be combined in order to further improve the accuracy to EER=2.43%.

The results also confirm that using the first 5 eigenvectors is sufficient to represent the proposed features while keeping the identification accuracy high. Therefore, our proposed features also form very compact and computationally efficient fingerprints. Typically, if each dimension is represented by a 4-byte floating-point number, the size of the corresponding feature template $h = \{\hat{G}; \Sigma_G\}$ is 20 (5×4) bytes for $\hat{G}$ and 100 (5x5x4) bytes for the square covariance matrix Σ_G resulting in a total of 120 bytes.

In terms of feature extraction performance, given the much lower dimensionality of the burst samples (40000 vs. 960000 for the sweep), they are much faster to digitally acquire and extract with approximately 2 sec. compared to 26 sec. for the sweep data samples. The times are measured on a machine with 2.00 GHz CPU, 2 GB RAM running Linux Ubuntu. It should be noted that all the components of the feature extraction can be implemented efficiently in hardware which would significantly improve the performance.

6 Application to Cloning Detection

The classification and identification results presented in Section 5 indicate that physical-layer fingerprinting can be practical in a controlled environment. In this section,

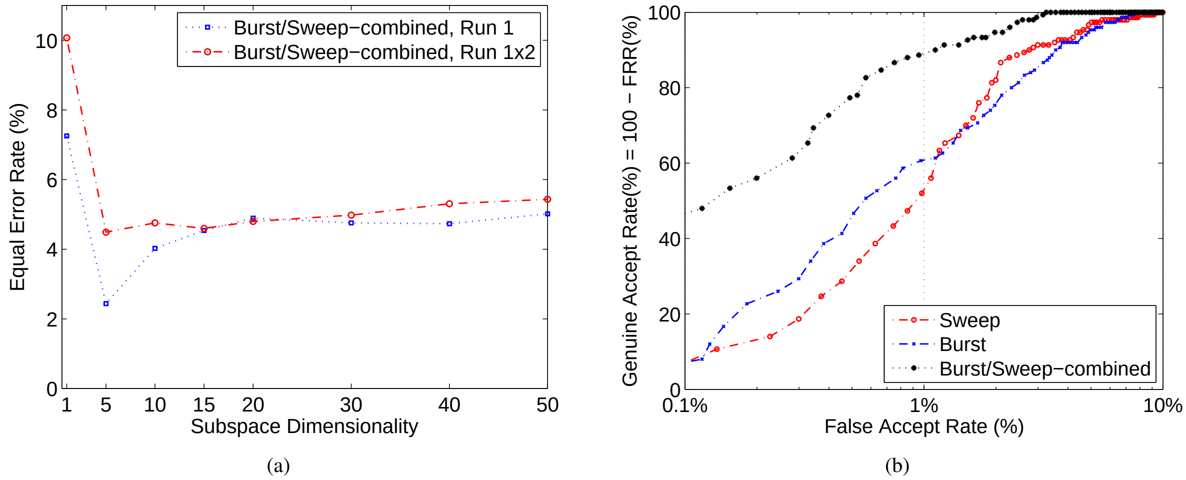

Figure 9: a) The identification accuracy combining the sweep and burst features b) Receiver Operating Characteristic (ROC) for N=15 for burst and sweep spectral features and their combination. 50 identical transponders are used. The subspace dimension is fixed to 5. See Table 5 for the underlying data.

Table 5: Summary of accuracy when a combination of burst and sweep features used (50 transponders).

Type	Run	N	Test matchings		Threshold T	EER (%)	EER CI (%)		Validation
			Genuine	Imposter			lower	upper	
Burst/Sweep	1	15	150	11025	1.56	2.43	1.54	3.33	4-fold
Burst/Sweep	1x2	15	200	9800	2.18	4.38	3.9	4.9	4-fold

we discuss how it could be used in the context of product or document cloning detection. We point out however that the cloning detection will obey to the achieved error rates. Despite a number of protective measures, it has been recently shown [18, 34, 33, 47] that even RFID transponders in electronic identity documents can be successfully cloned, even if the full range of protective measures specified by the standard [3], including active authentication, is used. We consider the physical-layer fingerprinting described in this work as an additional efficient mechanism that can be used to detect document counterfeiting.

We foresee two use cases in which fingerprints can be applied for anti-counterfeiting. In the first use case, the fingerprints are measured before RFID deployment and are stored in a back-end database, indexed with the unique transponder (document) identifier. When the authenticity of the document with identifier ID is verified, the fingerprint of the document transponder is measured, and then compared with the corresponding transponder fingerprint of document ID stored in the database. In order to successfully clone the document, the attacker needs to perform two tasks:

1. Obtain the fingerprint template of the transponder in the original document and
2. Produce or find a document (transponder) with the same fingerprint.

In order to extract a fingerprint template the attacker needs to fully control the target document (hold it in possession) for long enough to complete the extraction. Using the methods from our study, it would be hard, if not infeasible, for the attacker to extract the same fingerprints remotely (e.g., from few meters away). In our experiments, such remote feature extraction process resulted in an EER of approximately 50%. We assume that this is due to the change of acquisition antenna orientation and lower signal-to-noise ratio. We do not exclude the possibility that other discriminant features could be found that could be extracted remotely. However, this does not appear to be the case for our features. After obtaining the original fingerprint, the attacker now needs to produce or find an RFID transponder with that fingerprint (i.e., such that it corresponds to the one of the original document), which is hard given that the extracted fingerprints are due to manufacturing process variation. Although manufacturing process variation effects the RFID micro-controller itself, it is likely that the main source of detectable variation lies in the RFID radio circuitry. However, we cannot conclude with certainty which component of the entire transponder circuit contributes most to the fingerprints. We leave this determination to future work. Because of the complexity of these circuits this is a challenging task in the lab, let alone in ”the wild” environment of the attacker.

In the second use case, transponder fingerprints are measured before their deployment as in the first case, but are stored on the transponders instead of in a back-end database. Here, we assume that the fingerprints stored on the transponders are digitally signed by the document-issuing authority and that they are protected from unauthorized remote access; the digital signature binds the fingerprint to the document unique identifier, and both are stored on the transponder. When the document authenticity is validated, the binding between the document ID and the fingerprint stored on the transponder is ensured through cryptographic verification of the authority's signature. If the signature is valid, the stored fingerprint is compared to the measured fingerprint of the document transponder. The main advantage in this use case is that the document authenticity can be verified ”off-line”. The main drawback is that the fingerprint is now stored on the transponder and without appropriate access protection, it can be remotely obtained by the attacker. Here, minimal access protection can be ensured by means of e.g., Basic Access Authentication [3] although, that mechanism has been shown to have some weaknesses due to predictable document numbers [33]. As we mentioned in Section 5.4, our technique generates compact fingerprints, which can be stored in approximately 120 bytes. This means that they can easily be stored in today's e-passports. The ICAO standard [3] provides space for such storage in files EF.DG[3-14], which are left for additional biometric and future use; transponder fingerprints can be stored in those files. Our proposal does not require the storage of a new public key or maintenance of a separate public-key infrastructure, since the integrity of the fingerprints, stored in EF.DG[3-14] will be protected by the existing passive authentication mechanisms implemented in current e-passports.

The closest work to ours in terms of transponder cloning protection is the work of Devadas et al. [12], where the authors propose and implement Physically Unclonable Function(PUF)-Based RFID transponders. Processors in these transponders are specially designed and contain special circuits, PUFs, that are hard to clone and thus prevent transponder cloning. The main difference between PUF-based solutions and our techniques is that our techniques can be used with existing RFID transponders, whereas PUF-based solutions can detect cloning only of PUF-based transponders. However, PUF-based solutions do have an advantage that they rely on ”controlled” randomness, unlike our techniques, that relies on randomness that is unintentionally introduced in the manufacturing of the RFID tags.

7 Related Work

Besides PUF-based RFIDs [12], that we discuss in the previous section, the following works relate to ours.

In [41], Richter et al., report on the possibility of detecting the country that issued a given passport by looking at the bytes that an e-passport sends as a reply in response to some carefully chosen commands from the reader. This technique therefore enables classification of RFID transponders used in e-passports. Our technique differs from that proposal as it enables not only classification, but also identification of individual passports. Equally, the technique proposed in [41] cannot be used for cloning detection since the attacker can modify the responses of a tag on a logical level.

The proliferation of radio technologies has triggered a number of research initiatives to detect illegally operated radio transmitters [44, 45, 23], mobile phone cloning [30], defective transmission devices [48] and identify wireless devices [20, 22, 43, 40, 39, 9] by using physical characteristics of the transmitted signals [15]. Below, we present the most relevant work to ours in terms of signal similarities, features and objectives.

Hall et al. [20, 21] 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% [22]. Ureten et al. [39] extracted the envelope 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%. Rasmussen et al. [40] 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%. Brik et al. [9] proposed a device identification technique 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. In [11] the authors demonstrate the feasibility of transient-based Tmote Sky (CC2420) sensor device identification with an EER of 0.24%. The same work considered the stability of the proposed fingerprint features with respect to capturing distance, antenna polarization and voltage, and related attacks on the identification system.

8 Conclusion

In this work we performed the first comprehensive study of physical-layer identification of RFID transponders. We showed that RFID transponders have stable fingerprints related to physical-layer properties which enable their accurate identification. Our techniques are based on the extraction of the modulation shape and spectral features of the response signals of the transponders to the in- and out- of specification reader signals. We tested our techniques on a set of 50 RFID smart cards of the same manufacturer and type and we showed that these techniques enable the identification of individual transponders with an Equal Error Rate of 2.43% (single run) and 4.38% (two runs). We further applied our techniques to a smaller set of electronic passports, where we obtained a similar identification accuracy. We tested the classification accuracy of our techniques, and showed that they achieve 0% average classification error for a set of classes corresponding to manufacturers and countries of issuance. Finally, we analyzed possible applications of the proposed techniques to the detection of cloned products and documents.

Acknowledgements

This work was partially supported by the Zurich Information Security Center. It represents the views of the authors.

References

[1] Fingeprint verification competitions (FVC). http://bias.csr.unibo.it/fvc2006/.

[2] IBM JCOP family. ftp://ftp.software.ibm.com/software/ pervasive/info/JCOP_Family.pdf.

[3] ICAO. http://www.icao.int/.

[4] ISO/IEC 14443 standard. http://www.iso.org/.

[5] RFID security and privacy lounge. http://www.avoine.net/rfid /index.html.

[6] Avoine, G., and Oechslin, P. RFID traceability: A multilayer problem. In *Financial Cryptography* (2005), A. Patrick and M. Yung, Eds., vol. 3570 of *LNCS*, pp. 125–140.

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

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

[9] Brik, V., Banerjee, S., Gruteser, M., and Oh, S. Wireless device identification with radiometric signatures. In *Proc. ACM MobiCom* (2008).

[10] Costen, N., Parker, D., and Craw, I. Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics 58*, 4 (1996), 602–612.

[11] Danev, B., and Čapkun, S. Transient-based identification of wireless sensor nodes. In *Proc. ACM/IEEE IPSN* (2009).

[12] Devadas, S., Suh, E., Paral, S., Sowell, R., Ziola, T., and Khandelwal, V. Design and implementation of PUF-based "unclonable" RFID ICs for anti-counterfeiting and security applications. *Proc. IEEE Intl. Conf. on RFID* (2008), 58–64.

[13] Dimitriou, T. A lightweight RFID protocol to protect against traceability and cloning attacks. In *Proc. ICST SecureComm* (2005).

[14] Duc, D. N., Park, J., Lee, H., and Kim, K. Enhancing security of EPCglobal Gen-2 RFID tag against traceability and cloning. In *Proc. Symposium on Cryptography and Information Security* (2006).

[15] Ellis, K., and Serinken, N. Characteristics of radio transmitter fingerprints. *Radio Science 36* (2001), 585–597.

[16] EPCglobal. Architecture framework v. 1.2. standard, 2007. http://www.epcglobalinc.org/standards/ architecture/architecture_1_2-framework-20070910.pdf.

[17] Feldhofer, M., Dominikus, S., and Wolkerstorfer, J. Strong authentication for RFID systems using the AES algorithm. In *Workshop on Cryptographic Hardware and Embedded Systems* (2004), M. Joye and J.-J. Quisquater, Eds., vol. 3156 of *LNCS*, pp. 357–370.

[18] Grunwald, L. Cloning ePassports without active authentication. In *BlackHat* (2006).

[19] Halamka, J., Juels, A., Stubblefield, A., and Westhues, J. The security implications of VeriChip^{TM}cloning. Manuscript in submission, 2006.

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

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

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

[23] Hippenstiel, R., and Payal, Y. Wavelet based transmitter identification. In *Proc. ISSPA* (1996).

[24] Jain, A., Prabhakar, S., and Chen, S. Combining multiple matchers for a high security fingerprint verification system. In *Pattern Recognition Letters* (1999).

[25] Juels, A. Minimalist cryptography for low-cost RFID tags. In *Intl. Conf. on Security in Communication Networks* (2004), C. Blundo and S. Cimato, Eds., vol. 3352 of *LNCS*, pp. 149–164.

[26] Juels, A. Strengthening EPC tags against cloning. Manuscript, 2005.

[27] Juels, A. Rfid security and privacy: A research survey. *IEEE Journal on Selected Areas in Communications 24*, 2 (2006).

[28] Juels, A., Pappu, R., and Parno, B. Unidirectional key distribution across time and space with applications to RFID security. In *Proc. 17th USENIX Security Symposium* (2008), pp. 75–90.

[29] Juels, A., Rivest, R., and Szydlo, M. The blocker tag: Selective blocking of RFID tags for consumer privacy. In *Proc. ACM CCS* (2003), pp. 103–111.

[30] Kaplan, D., and Stanhope, D. Waveform collection for use in wireless telephone identification, 1999.

[31] Kerschbaum, F., and Sorniotti, A. RFID-based supply chain partner authentication and key agreement. In *Proc. ACM WiSec* (2009).

[32] KITTLER, J., HATEF, M., DUIN, R., AND MATAS, J. On combining classifiers. *IEEE Trans. on Pattern Analysis and Machine Intelligence 20*, 3 (1998).

[33] LAURIE, A. Reading ePassports with predictable document numbers. In *news report* (2006).

[34] M, W. Cloning ePassports with active authentication enabled. In *What The Hack* (2005).

[35] MANLY, B. *Multivariate Statistical Methods: A Primer*, 3rd ed. Chapman & Hall, 2004.

[36] MARPLE, S. Computing the discrete-time analytic signal via FFT. *IEEE Trans. on Signal Processing 47*, 9 (1999).

[37] MITRA, M. Privacy for RFID systems to prevent tracking and cloning. *Intl. Journal of Computer Science and Network Security 8*, 1 (2008), 1–5.

[38] OPPENHEIM, A., SCHAFER, R., AND BUCK, J. *Discrete-Time Signal Processing*, 2nd ed. Prentice-Hall Signal Processing Series, 1998.

[39] O.URETEN, AND N.SERINKEN. Wireless security through RF fingerprinting. *Canadian J. Elect. Comput. Eng. 32*, 1 (Winter 2007).

[40] RASSMUSSEN, K., AND CAPKUN, S. Implications of radio fingerprinting on the security of sensor networks. In *Proc. SecureComm* (2007).

[41] RICHTER, H., MOSTOWSKI, W., AND POLL, E. Fingerprinting passports. In *NLUUG Spring Conference on Security* (2008).

[42] ROSS, A., AND JAIN, A. Multimodal biometrics: An overview. In *Proc. EUSIPCO* (2004).

[43] TEKBAS, O., URETEN, O., AND SERINKEN, N. Improvement of transmitter identification system for low SNR transients. In *Electronic Letters* (2004).

[44] TOONSTRA, J., AND KISNER, W. Transient analysis and genetic algorithms for classification. In *Proc. IEEE WESCANEX* (1995).

[45] TOONSTRA, J., AND KISNER, W. A radio transmitter fingerprinting system ODO-1. In *Canadian Conf. on Elect. and Comp. Engineering* (1996).

[46] VAJDA, I., AND BUTTYÁN, L. Lightweight authentication protocols for low-cost RFID tags. In *Proc. 2nd Workshop on Security in Ubiquitous Computing – Ubicomp* (2003).

[47] VANBEEK, J. ePassports reloaded. In *BlackHat* (2008).

[48] WANG, B., OMATU, S., AND ABE, T. Identification of the defective transmission devices using the wavelet transform. *IEEE PAMI 27*, 6 (2005), 696–710.