EMFORCED: EM-Based Fingerprinting Framework for Remarked and Cloned Counterfeit IC Detection Using Machine Learning Classification

Andrew Stern, *Student Member, IEEE*, Ulbert Botero, *Student Member, IEEE*, Fahim Rahman, *Member, IEEE*, Domenic Forte, *Senior Member, IEEE*, and Mark Tehranipoor, *Fellow, IEEE*

***Abstract*— Electronics supply chain vulnerabilities have broadened in scope over the past two decades. With nearly all integrated circuit (IC) design companies relinquishing their fabrication, packaging, and test facilities, they are forced to rely upon companies from around the world to produce their ICs. This dependence leaves the electronics supply chain open to counterfeiting activities. In this article, we propose an electromagnetic (EM)-based fingerprinting framework, called EMFORCED, to detect remarked and cloned counterfeit ICs. Here, we demonstrate the benefits of using naturally occurring EM side channels to identify the IC design layout without decapsulating the chip under test. Enabling only the clock, V_{dd}, and ground pins allows us to generate a design-specific fingerprint that is dependent upon the physical parameters of the chip under test. EMFORCED leverages the EM emissions from the clock distribution network to create a holistic, design-level, fingerprint, including both temporal information and spatial information. We utilize the fingerprint information of functionally similar 8051-series microprocessors from three vendors and perform unsupervised (principal component analysis) and supervised (linear discriminant analysis) machine learning methods on all ICs to determine their intravendor and intervendor similarities. We acquired ICs from multiple dates and lot codes along with variants acquired from the gray market and analyzed them for authenticity using physical inspection and X-ray tomography. Statistical analysis and machine learning techniques are used to demonstrate the reference-free and reference-inclusive classification methods based on EMFORCED measurements. We demonstrate the classification accuracies of 99.46% and 100% for unsupervised and supervised approaches, respectively.**

***Index Terms*— Counterfeit detection, electromagnetics (EMs), machine learning, side channels, supply chain security.**

I. Introduction

ELECTRONICS are commonly deployed in critical systems; as such, the need to confirm their authenticity is crucial. Their inclusion in everything from automobiles and airplanes to nation-wide utilities and advanced weapon systems demands reliable and authentic integrated circuits (ICs). As most design houses have outsourced their semiconductor manufacturing to offshore companies, the supply chain can currently be compromised at several points of vulnerability [1]. Electronics counterfeiting has become a multibillion dollar industry as the supply chain has become more convoluted [2]. Detecting these counterfeits is a challenging problem as nondestructively verifying all ICs in a system is a daunting task.

Traditionally, companies have used physical inspection, observing features such as package markings, surface textures, material composition, pin corrosion, and shape to determine the authenticity of a given IC [3]–[5]. These techniques are effective at identifying cosmetic defects although they all share a fundamental limitation, that is, physical inspection techniques do not rely upon the underlying electronic properties of the device under test (DUT). The critical component to verify is the silicon die within the package; hence, assuming the functionality and authenticity of a die based upon the exterior of its packaging may lead to misclassifications. Observing the package of a DUT may provide details of the care taken in packaging and distributing the IC and potentially catch recycled ICs. However, experienced counterfeiters can remark and refurbish an IC to evade detection from physical inspection techniques and sell the IC as new or an entirely different part [2]. Extracting physical parameters from the die itself provides nonrefutable evidence of a device's defining characteristics. Additionally, physical inspection is a traditionally manual process, which requires an experienced workforce, expensive instrumentation, and a golden model to compare against, corresponding to a high-cost solution susceptible to IC misclassification.

The taxonomy of electronic counterfeits categorizes the types of counterfeits into recycled, remarked, overproduced, out-of-spec, cloned, forged documentation, and tampered [4]. Here, we focus on remarked and cloned ICs. Remarked ICs are comprised of new or recycled chips that have their markings modified to falsely present themselves as a different IC. Cloned chips are copies of an existing IC design and are widely used by counterfeiters to reduce the development cost of a component [4]. It is important to distinguish between cloned and overproduced ICs. While both ICs are dependent upon the same IC design, overproduced ICs will be physically and functionally identical to the authentic chip as they originate from the same foundry, while cloned ICs may not

Manuscript received April 27, 2019; revised September 15, 2019; accepted October 3, 2019. Date of publication November 12, 2019; date of current version January 21, 2020. This article was presented at the 2018 International Test Conference (ITC). *(Corresponding author: Andrew Stern.)*

The authors are with the Electrical and Computer Engineering Department, University of Florida, Gainesville, FL 32608 USA (e-mail: andrew.stern@ufl.edu; jbot2016@ufl.edu; fahim034@ufl.edu; dforte@ece.ufl.edu; tehranipoor@ufl.edu).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TVLSI.2019.2949733

necessarily be physically identical but share the same function since they may be fabricated in separate foundries. Ideally, integrating a fingerprint into a device, which could provide a trusted digital ID, would provide confidence to distributors, system integrators, and end users. Electronic chip ID (ECID) [6] and physical unclonable functions (PUFs) [7], [8] are the most common methods of providing unique identification to modern ICs. ECID values are not inherently tied to the physical properties of the silicon die and are not available on many legacy components. Additionally, PUFs incur additional overhead, are susceptible to reliability issues, and similarly are not applicable to legacy components.

Counterfeit detection techniques that rely upon internal device parameters have previously been proposed [5], [9]–[11]. To confidently classify an IC as authentic, side channels, such as power, infrared, and electromagnetic (EM) emission, can be used to extract information that is directly correlated with the physical properties of the DUT. Power-based side-channel information is easily obtained from any IC by monitoring the transient current on the ICs' ground pin [12]. Power consumption can provide useful information about a given IC although it is limited by its single point of reference and inherently lacks spatial information. Infrared and photon-based side channels provide both temporal and spatial side-channel information [13] although they require extended acquisition times and a direct line of sight to the silicon die. Certain ICs are distributed with the silicon exposed although the majority of ICs would require decapsulation prior to measurement, inevitably rendering a subset of DUTs unusable in the process. Finally, EM-based side-channel techniques provide temporal and spatial information while remaining completely noninvasive although challenges, such as environmental noise, spectrum resolution and bandwidth, and probe stability, need to be overcome for successful detection [14].

A. Contribution

In this article, we propose EMFORCED, an EM fingerprinting framework for remarked and cloned counterfeit IC detection. EMFORCED leverages EM side-channel information as a design-specific fingerprint to determine the authenticity of a given DUT. Its unique features can be summarized in the following key points.

1) EMFORCED enables nonintrusive, fast, and accurate counterfeit detection allowing all ICs to be authenticated. In contrast to invasive or otherwise damaging techniques, EMFORCED will not assume that all ICs within a larger group are the same as a single IC.
2) The generated fingerprints are independent of circuit functionality, thus requiring no knowledge of the DUT.
3) EMFORCED can be applied to any IC with a clock and does not require any prior modification to the circuit or its original programming.
4) EM fingerprint measurements enable both rotational and spatial analyses of a given DUT.
5) Finally, EMFORCED uses a low-cost experimental setup that can be used at any time in the IC life cycle after packaging.

Our main contributions are summarized as follows.

1) We propose a novel noninvasive EM-based framework for counterfeit IC detection that extracts information from the physical die characteristics to generate a design-specific signature.
2) We utilize two different measurement environments, discussing the potential ubiquity of the approach.
3) We demonstrate high-confidence device classification on more than 60 ICs from three vendors using both unsupervised and supervised machine learning methods to encompass multiple utilization scenarios.
4) We provide a resilience analysis of our proposed framework, discussing the theoretical difficulty of cloning a device signature.

The remainder of this article is organized as follows. Relevant background information regarding counterfeit detection using EM emissions is provided in Section II. The EMFORCED framework is detailed in Section III. An overview of our experiments and results is provided in Section IV, where we also explore the utilization of multiple variables for increased classification confidence. We provide a resilience analysis of our classification methods in Section VI. Finally, we conclude this article in Section VII.

II. Preliminary

A. Near-Field EM Emissions

EM signals are prominent in our everyday lives. This mostly invisible radiation is utilized in applications ranging from microwave cooking to cellular communications. EM waves radiate from any conductor containing a moving current, making nearly every object emit its own EM field. Notably, modern ICs are constructed on billions of current-carrying metal lines. Positioned above the active region of the device, these lines act as antennas that propagate EM waves from numerous sources within the IC. Defined by its close measurement proximity, near-field radiative effects can be observed at distances r less than $r = 0.1 * \lambda/(2 * \pi)$, where λ is the wavelength of the desired measurement signal [15]. Within our EM measurement environment (see Section IV for details) with an observation distance of 1 mm, near-field measurements apply for frequencies below 10 GHz ($\lambda < 3$ cm), which corresponds to our desired spectrum. Measured near-field emanations can be directly correlated with the activity of the IC itself, mapping core functionality and localized power consumption to a measurable physical characteristic [16]. As each metal line corresponds to a unique EM emission source, the layout of a given design determines the emission profile emanating from a given DUT. The two comprehensive metal structures within any digital or mixed-signal IC are the power and the clock distribution networks making them primary candidates for a design-based EM fingerprint. Here, we fluctuate the clock input signal to generate EM fingerprints from the clock distribution network. This was selected to maintain operating

characteristics of the ICs while propagating repeating oscillations throughout the complex clock distribution network.

B. Counterfeit Detection

Counterfeit IC detection techniques have been developed for over a decade. However, the expansive taxonomy of known counterfeit types presents numerous obstacles to overcome [4]. Modern counterfeit detection solutions are categorized by: 1) use of design-for-anticounterfeit (DfAC) technologies; 2) physical inspection; and 3) electrical parameter testing using functional test, structural test, or side-channel measurement. Among them, DfAC methodologies can be used exclusively on new IC designs, as they require modifications to presilicon designs. Alternatively, physical inspection can be performed on all ICs (new and legacy), but it is traditionally a manual process, making it expensive and time consuming [4]. Functional and structural (scan-based) testing are also expensive techniques as they require a costly tester and proper test pattern development, which require sophisticated knowledge of the device [17]. Solutions utilizing side-channel measurements are the most viable and low-cost methods. Among the various side channels, EM is the most attractive due to its low cost, applicability to legacy components, speed/ease of acquisition, and ability to extract temporal and spatial information correlated with the die's physical characteristics, with or without test vectors.

C. Distinguishing Characteristics

To determine whether an IC is a remarked or cloned counterfeit, we must examine the characteristics of the DUT that can be directly correlated with the counterfeiting method. Physical inspection methods dependent upon the external characteristics of the DUT, such as packaging size, texture, markings, and pin appearances, often fail to successfully distinguish between an authentic and a counterfeit part. Examining the physical properties of a suspect IC provides insights into the design itself. In the case of EM fingerprint measurements, the billions of metal traces within the IC's die can be approximated by an array of radiating point sources. These points are accumulated by the EM probe and convolve into the extracted fingerprint. The sources can be approximated by (1), which expresses the intensity of EM emissions (E) accumulated over the entire surface of the IC. The collected EM-based fingerprints represent a complex summation of the currents propagating throughout the design. If the die's area is represented as a $\langle x, y\rangle$ grid of point sources, the measured result can be approximated by the following equation [18]:

$$E \propto \int_y \int_x \frac{\vec{S}}{4\pi \vec{r}^2} dx\, dy = \int_y \int_x \frac{(I_{\langle x,y\rangle})^2 Z_{\langle x,y\rangle}}{4\pi (r_{\langle x,y\rangle})^2}\, dx\, dy. \quad (1)$$

Each of these positions requires a current (I) and complex load (Z) to determine the apparent power (S), in addition to the observation distance (r). Differences between the circuit designs will materialize as differences in the observed fingerprints. Determining physical differences between the suspect ICs will expose remarked and cloned counterfeit ICs as their designs are not physically identical to the authentic ICs.

The use of side-channel signatures to distinguish between similar circuits can also be found within the domain of hardware Trojan detection [19]–[22]. These works require input patterns to be provided to the circuit similar to prior counterfeit detection work. This requires additional time from the engineers to develop custom test environments per device type. However, our method does not require prior knowledge of the DUT and can directly be translated to new devices that share the same form factor. Postfabrication Trojan detection methods, whether using power or EM, require custom test environments, which would not be convertible to other devices without new software/firmware development. Additionally, the techniques utilized in [20] represent similar processing techniques to those proposed here, namely the use of PCA. However, Trojan detection techniques cannot be directly transferred to remarked and cloned counterfeit detection. For example, Muehlberghuber *et al.* [20] detected a Trojan with the prior knowledge of having exactly two unique classes of ICs. In contrast, counterfeit detection techniques, such as EMFORCED, should not assume that the ICs are bound to a given number of classes. They should be able to encounter new types of counterfeits whose responses may vary greatly and eliminate extraneous fingerprints prior to classification.

D. Threat Model

Remarked and cloned counterfeit ICs pose a substantial threat to various entities within the electronics supply chain. Remarked ICs can be introduced by adversaries that manipulate package identifiers to present used or alternative components as new and higher grade versions of a desired IC. Cloned ICs are functionally similar or identical ICs although they deviate from their authentic counterparts in implementation. In certain applications, die revisions of an IC can be considered insufficient for direct substitution; as such, we treat these as cloned ICs as well. To determine the primary entities that remarked and cloned ICs have infiltrated, we present the following scenarios.

1) The entity possesses knowledge of the ICs' hardware design, such as the GDSII file or fabrication node, and requires verified ICs (example entities: fabless design house or government).
2) The entity replaces a legacy component that they possess, but it does not have a reference signature or design from a trusted party (example entities: government or critical system integrator).
3) The entity acquires ICs, which have verified EM reference signatures available from a trusted party, for integration into their systems or further distribution, and is required to verify their authenticity (example entities: original equipment manufacturer (OEM) or component retailer).

Determining the authenticity of ICs is a critical part of each entity's value proposition and could be scrutinized if they were identified by an end user. For example, fabless design houses currently place substantial trust in the foundries which fabricate their designs. Identifying authentic ICs is challenging for them as their "golden" IC is sourced directly from the foundry.

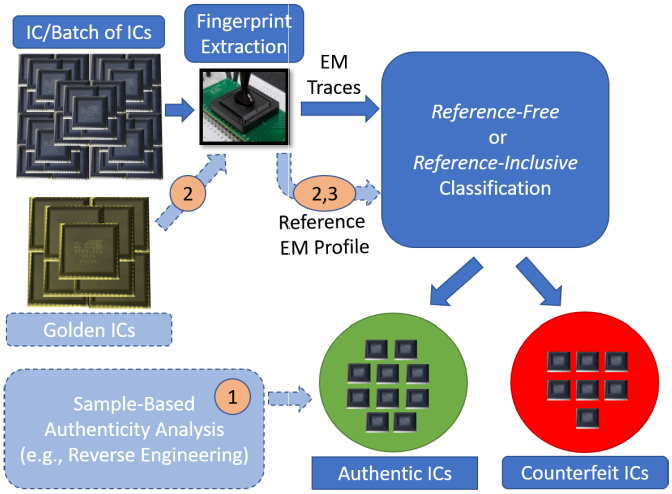

Fig. 1. EMFORCED framework overview with scenario-specific elements labeled in orange.

To verify the design, the design house may reverse engineer the fabricated IC and compare it with their knowledge of the technology node or original GDSII file. This process is very time consuming and destroys the device in the process. However, if the EM fingerprints of all ICs are extracted prior to reverse engineering, the remaining ICs can be confidently categorized depending upon the correlation between their fingerprints and the verified golden IC. Component retailers with access to a verified reference signature may verify an IC's authenticity easily, as they possess verified fingerprints that can be directly compared against. Entities that often replace legacy components, such as government, are able to extract golden fingerprints from the legacy components that they are replacing. Once the EM-based fingerprints are measured, newly acquired suspect devices can be categorized accordingly.

III. EMFORCED Framework

The EMFORCED framework can be applied to a number of entities. Fig. 1 shows the three different scenarios used to classify a set of DUTs into authentic and counterfeit groupings. After acquiring the ICs, device-specific EM fingerprints are extracted and then processed using machine learning methods. Finally, after the ICs are separated into groups, single samples from each group can be analyzed using alternative techniques (e.g., reverse engineering) to properly identify groups of authentic and counterfeit ICs.

A. Scenario Identification

In this article, we demonstrate the effectiveness of EM-based measurements and subsequent fingerprinting for remarked and cloned counterfeit IC detection. To confidently distinguish between the counterfeit and authentic ICs, we defined three scenarios in Section II-D. Scenarios 1–3 encompass all available combinations of reference-free and reference-inclusive classification methods. The scenario corresponding to a given entity impacts the process by which EMFORCED must be executed, as shown in Fig. 1. For example, scenario 1 requires reference-free classification and sample-based authenticity analysis, while scenarios 2 and 3 would not.

B. EM Fingerprint Extraction

The EMFORCED framework relies upon near-field EM traces obtained from ICs operating without test vector application or prior knowledge of the IC. To generate a device-specific fingerprint, the clock distribution network is stimulated by an input clock signal within the typical device operating range. The clock network is chosen as it permeates all sections of the die and naturally supports an oscillating or pulsed signal, enabling high signal-to-noise EM measurements. EM traces are generated most applying a recurring pulsed signal present from a function generator or crystal oscillator. To extract these traces, each DUT is sequentially placed into a socket for consistent measurements. In this article, we first evaluate the validity of our framework on a custom-breadboard setup and then evaluate the potential of our framework with a standard 8051-development board. The first case represents the optimal extraction environment reminiscent of a laboratory setup, while the second provides a more realistic testing environment with commercial-off-the-shelf (COTS) components. Once situated into the test platform, a near-field EM probe is centered over the die and lowered until it contacts the surface of the package. The signal from the near-field probe is amplified and then sampled to provide a digital representation of the EM fingerprint. The extracted fingerprints can then be processed and classified into their respective groups.

C. Suspect Classification Methods

To properly categorize suspect ICs within a specific group, various techniques can be used to confidently identify similar fingerprints. Here, we separate the three scenarios detailed in Section II-D into two groups. Scenario 1 requires additional postprocessing (e.g., reverse engineering) to determine if an IC is authentic although can be conducted upon a sampling basis thereafter. Scenarios 2 and 3 are able to leverage information from an existing database of known authentic EM fingerprints or extract their own reference from their golden IC. Here, we detail our reference-free (scenario 1) and reference-inclusive (scenarios 2 and 3) classification methods.

1) Reference-Free Classification: Assuming that the entity does not have access to a reference profile, it is essential to first separate the parts into similarly performing groups. To quickly and efficiently separate the collected fingerprints into their respective groups, unsupervised classification methods enable grouping regardless of the information an entity possesses regarding the DUTs. For such, we first use principal component analysis (PCA) that is a very powerful and popular technique used for dimensionality reduction as well as feature extraction [23]. The basis behind PCA is to convert potentially correlated variables into a set that is linearly uncorrelated (referred to as principal components). These principal components represent the data in a manner which highlights the most expressive features of each signal by projecting the data in

orthogonal directions that contain the most variance. The first principal component represents the projection of the data with the most variance, the second principal component represents the second most, and so forth. Furthermore, the number of distinct principal components in a data set will be defined as the smaller of the number of observations or original variables. By determining the principal components that contain the most variance, one is able to quantify the importance of each dimension of the data and provide a reasonable characterization within a reduced dimensional space. In our analysis, PCA is used to reduce the 250 000 sample-per-fingerprint feature vector down to a dimensionality of $N-1$, where N is the number of measurements used for training (either 45 or 340 in our experimentation; see Section IV).

Alternatively, cross-correlation analysis can provide a statistical method for separating the ICs into their relative groups. Cross correlation determines the similarity between two time-series signals. This method can provide a simple percentage of correlation for each comparison and serve as a similarity metric between the collected fingerprints. In EMFORCED, cross correlation is used to verify that the fingerprint under test is a member of one of the groups defined by PCA or reference profiles.

2) Reference-Inclusive Classification: Assuming that the entity has access to an authentic IC or a reference profile provided by a trusted entity, reference-inclusive classification can confidently determine the authenticity of the DUT without any additional postprocessing. To confidently address all possible scenarios of remarked and cloned IC insertion, a layered approach must be used. If an entity possesses a single reference fingerprint or IC, then cross correlation will be optimal for determining the device's authenticity. Acquiring a number of suspect parts to classify may benefit from either cross correlation or PCA as there is enough data to extract variance. However, clear fingerprint separation using PCA does not guarantee that the devices are a part of different groups. Should all of the fingerprints used for classification belong to a single device category, PCA may still separate them regardless of their likeness. Reference-inclusive classification should prioritize a cross-correlation analysis to determine whether the collected signatures align well with the reference before moving on to further techniques.

The reference-inclusive processing technique introduced here utilizes several established techniques to properly differentiate remarked and cloned ICs from authentic ones. Assuming access to reference profiles from several authentic device types, a layered approach must be used to ensure proper classification. For dimensionality reduction, PCA is used on all genuine reference fingerprints. Note that from this point forward, all data will be preprocessed through the constructed PCA projection model. This will project all fingerprint data in the $(N-1)$-dimensional space defined by the reference fingerprints. Once the dimensionality of the data has been reduced, observing the known good groups and determining to which group each suspect DUT belongs, similar to our reference-free approach, will not be solved by simply calculating the distance from the suspect to the group. This method only works if one projects all classes through PCA or can definitively state that all part belongs to 1 of the N existing groups. If the suspect belongs to group $N+1$ and the group is not represented in the PCA training set; then the suspect could be misclassified. To address this issue, outlier detection can be utilized. Our outlier method defines an outlier as three standard deviations away from the median. When calculating outliers, not all principal components are equally weighted. For example, our first three principal components account for $\approx$96% of the variance from the projected value. Each principal component value is individually compared as an outlier 1 or inlier 0 against the corresponding authentic references. This binary value is then multiplied by the principal component weight to define the likelihood that a measurement is an outlier. For our measurements, we empirically determined that a likelihood of 1% provided optimal results for suspect classification. After performing outlier detection against all known good groups, the supervised machine learning technique, linear discriminant analysis (LDA), can be used [24]. LDA utilizes group label information and feature vectors to form boundaries between the groups and can classify samples into the most closely related group. This machine learning classification method works by maximizing the distance between the mean values of the class and minimizing the scatter within each class. Using the linear transformation and dimensionality reduction of PCA with the nonlinear multiclass classification of LDA provides increased classification accuracy. LDA can provide the final classification for suspects that classify as inliers in two or more groups, as well as verify the classification of your single-class inliers. True outliers (outliers in all of the reference groups) should be assumed to be part of a different $M+1$ class. Properly incorporating this layered approach is critical for proper counterfeit detection. For example, when assessing a component as a potential counterfeit and do not have a reference from the counterfeit group available, using LDA without outlier detection will misclassify all counterfeits as authentic.

D. Sample-Based Authenticity Analysis

After successfully separating the DUTs into similar groups, reference-free classification (predominantly used in scenario 1) requires a sample from each group to be analyzed for authenticity. Depending upon the entity's available knowledge of the authentic IC and capabilities, this analysis can take different forms. The most common due to its relatively reduced overhead would likely be physical inspection. This process would include comparing package markings against datasheet or other publicly available information along with observing characteristics that may result from the counterfeiting process. These characteristics include uneven or discolored packaging surfaces, black-top coating, painted or corroded pins, and surface material content analysis. For entities that have access to more valuable equipment, IC decapsulation and potentially reverse engineering would provide vastly increased confidence at the expense of time and cost. Once the sample DUTs have been determined to be either authentic or counterfeit, the remainder of the DUTs can be confidently classified within their designation. As EMFORCED relies upon the parameters

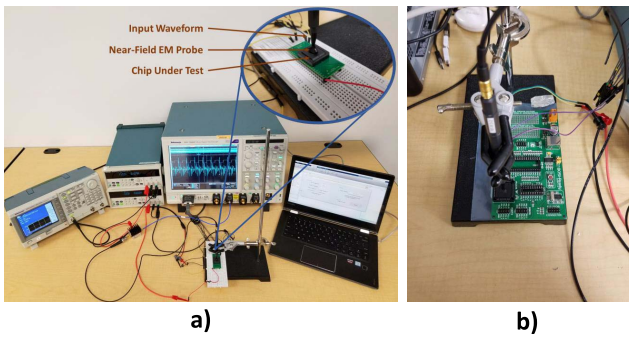

Fig. 2. EMFORCED measurement environments. (a) Breadboard implementation. (b) Development board implementation.

of the die, sacrificing a single IC allows the remainder of its group to be properly identified.

IV. Experimental Procedure and Results

To emulate remarked and cloned ICs, we acquired authentic chips with the same IP core (8051) from three vendors and distinguish between them. Within these three vendors, each variation has at least one set of DUTs (7–10 ICs) acquired from a "trusted" supplier and one set of DUTs (ten ICs) from a gray market seller. For our fingerprint classification experiments, 8051-series microcontroller ICs were selected from three different vendors: Atmel, Maxim, and NXP. These ICs were specifically selected because they had very similar functional characteristics (e.g., instruction set architecture, operating voltage, and package type). Since we originally acquired these ICs through a single retailer in one transaction, all the original devices had the same characteristic markings and there was no variation in fabrication lots, dates, or packaging location within a given vendor set. As obtaining chips with various date and lot codes is preferred, we ordered a second lot of Atmel devices six months after our original order. Additionally, gray market variants were acquired for each of the IC types for more realistic suspect classification. All nongray market DUTs are assumed to contain the same 8051 IP core and are housed within a 44-PLCC package [25]. Specific device fabrication facility and technology node information are not known, as they are not openly provided by the devices' supporting documentation. Estimation of these parameters is beyond the scope of this article although we believe that this information could potentially be recovered from similar EM-based measurement techniques. When attempting to detect remarked and cloned ICs, it is important to utilize all the existing information about the DUT. Various scenarios arise when attempting to overcome this problem, although they can be summarized as either reference-free or reference-inclusive.

A. Reference-Free Experimental Overview

Experimental measurements were taken on two different testing platforms for reference-free classification and only on the COTS 8051 development board for the reference-inclusive classification, as shown in Fig. 2(b). Fingerprints were extracted using a near-field EM probe amplified by 40 dB with wideband amplifiers. The breadboard setup relied on an external clock signal from a function generator, while the DUTs tested on the development board are supplied by a crystal oscillator on the PCB. The probes used in this experiment were factory-tuned to collect near-field EM radiation from a single direction and suppress perpendicular fields [26]. This assists in suppressing noise from external sources, such as PCB wiring, board-level jumper cables, and nearby equipment, alleviating the need for additional EM shielding around the setup. A mixed-signal oscilloscope was used with a sampling rate of 25 GS/s for all data acquisition. It should be noted that bandwidth requirement for proper EM fingerprint acquisition scales proportionally with the target clock frequency and inversely to the variance between device types. When comparing ICs with identical clock frequency, sampling above the Nyquist frequency is required. This enables transient effects to be extracted and used for more accurate classification. Given our unshielded, manually operated test environment, additional variation will be present in the results. Through experimentation, it appears as though these nonidealities contribute more variance than sourcing numerous date and lot codes would. This nonideal test environment alleviates the need for large, diverse, sample sets in our classification methods.

1) Breadboard Implementation: To collect EM emissions from the breadboard implementation, a 5-V peak-to-peak square wave with 50% duty cycle was applied at the clock input to enable external observation of transient currents from forced high-speed transitions. The recurring square wave input was chosen for this implementation as the low transition delay enables larger EM emissions. To ensure that the collected signatures are solely dependent upon the device characteristics, the voltage, frequency, and pulse shape remained the same. Altering these parameters will modify the collected signature and ease the classification of different IC types, although they should remain constant for a given device type. The clock stimulus was provided by a dedicated function generator at 16 MHz to comply with the 8051 devices' typical operation frequencies. The devices were powered by a 5-V dc input from a programmable power supply. This was to ensure that all clock buffers were activated and that the supplied signals could fully propagate through the circuit. The breadboard test environment required a total of three connections to the device (V_{dd}, ground, and the clock input) and was mirrored within the development board setup. A Langer EMV RF-K 7-4 near-field probe was used for fingerprint extraction with the breadboard implementation. This relatively large (6 mm $\times$ 10 mm) probe provides approximately 5 mm of placement accuracy for manual probe alignment.

It is important to note that our experimental setup only requires a clock pulse and does not demand any programs or specific input vectors to be loaded, which allows for black-box analysis of COTS components whose functions are usually not fully known. Furthermore, physical modification of the die or packaging is not required. EMFORCED could also prove useful for one-time-programmable (OTP) chips, on which specific test programs for fingerprinting cannot be loaded. Additionally, the tests conducted on the DUTs do not damage the original functionality or reliability of the devices.

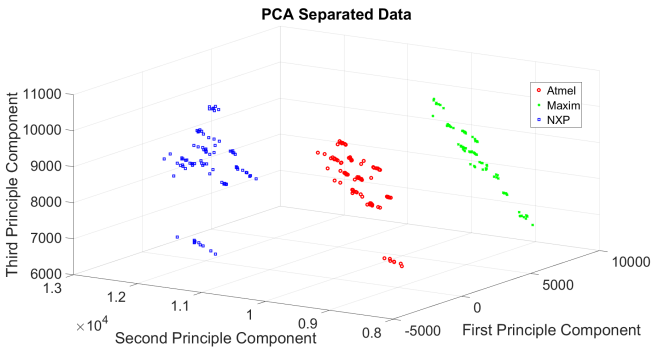

Fig. 3. PCA comparison of 8051 s from three vendors.

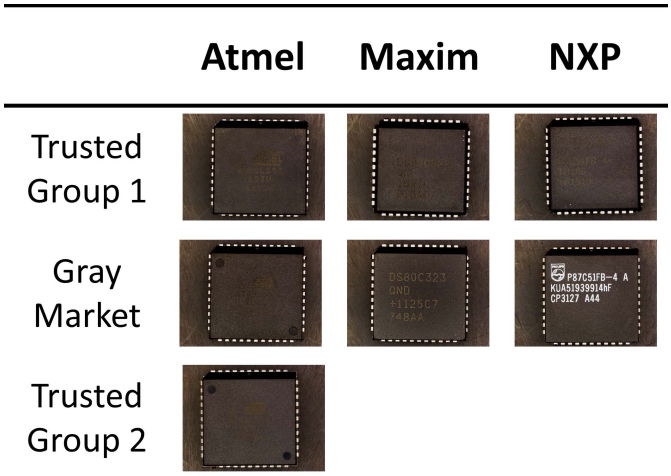

Fig. 4. Visual inspection comparison of authentic and gray market ICs.

This is because the clock signal and power we applied is within the specified operating range, and the time to collect and preprocess EM emissions (manually for ten acquisitions per DUT) was less than 2 min causing no impact compared to the device lifetime. The EM collection window per acquisition was set as 10 ms, in accordance with the mixed-signal oscilloscope's maximum sampling frequency of 25 GS/s and to observe traces over several clock cycles. It is important to note that this method is not restricted to the specific equipment used in this experimentation. However, for optimal results and scalability, real-world solutions should attempt to limit the differences between the measurement environments for the most consistent results.

Our first experiment consisted of collecting 300 EM fingerprints from 30 authentic ICs across Atmel, Maxim, and NXP. This breadboard setup shown in Fig. 2(a) ensured that the only signals measured were derived from the DUT's EM characteristics resultant from the oscillating clock input. The setup ensured that the IC was not actively looping to look for a program to execute, but rather remained inactive. In this setup, it was unnecessary to use averaging techniques for noise and anomaly reduction as the intravendor variations were very small despite different triggering points. The intervendor differences were quite pronounced, despite testing all DUTs with an identical, externally supplied, 16-MHz clock. The physical design differences in the power distribution network and clock tree accounted for the complex loading effects seen in the device fingerprints. To ensure our "golden" ICs were authentic, we processed them using PCA training on 45 of the 300 total fingerprints, randomly selecting traces from three of the ten DUTs per vendor. The first three principal components of one probe set are shown in Fig. 3. Fig. 3 clearly shows three well-defined groups as we should expect from our authentic IC classification. It is important to note that PCA is not a clustering algorithm but rather an unsupervised machine learning method. Therefore, the groups that have formed are a result of the PCA transformation increasing the variance between the points and did not require labeled data. A classification accuracy of 99.46% was determined by averaging the Euclidean, Minkowski, and City Block distance measurements for time-domain fingerprints from 100 randomly generated models [27]. The measurement bandwidth and emission spectrum provided by EM measurements typically yield improved results when performing analysis on frequency-domain data rather than time-domain data. Hence, when transformed to the frequency domain prior to classification, we observed 99.99% accuracy from 100 randomly generated models. Thus, we used frequency-domain data for the remainder of our experiments, excluding time-series cross correlation.

2) Development Board Implementation: Transitioning to a more standard environment with which any entity could easily test an IC, a COTS development board was used for additional experimentation. By substituting the custom-breadboard solution for a COTS board, the entire EMFORCED setup could be easily duplicated using any type of IC. Standard development boards provide the IC (DUT) with all required pull-up resistors to boot into an operating state and are readily available to entities as a standard COTS component. To eliminate variations from executed instructions and reduce overall noise margins, 50× averaging was utilized for each fingerprint collected. The development board supplies a clock signal from its crystal oscillator, requiring only power to be applied externally. The development board is used alongside all acquired ICs from both trusted and gray market suppliers from this point forward. Additionally, a Langer EMV RF-R 3-2 near-field probe was used for fingerprint extraction with the development board implementation and multivariate analysis. The probe was chosen to increase sensitivity in multiple directions and decrease the footprint for increased spatial variation response.

Gray market variations of the authentic ICs were acquired to create a more realistic scenario for sourcing legacy components. To provide a proper baseline for our technique, the gray market suspects were physically inspected. Upon initial inspection, the gray market components appeared fairly different from the authentic ICs. Fig. 4 shows a top-down comparison of all tested IC groups. The first and third rows (including the two authentic Atmel lots) show the authentic ICs acquired from a trusted supplier, and the second row shows the gray market suspect ICs. Given simple marking deviation analysis, a couple of Atmel suspects would be flagged due to varying marking types (e.g., laser engravings compared to ink). Additionally, the markings on the Maxim and NXP suspects strongly deviated from their authentic counterparts.

TABLE I
CROSS-CORRELATION ANALYSIS COMPARING INTRAGROUP, INTRAVENDOR, AND INTERVENDOR SEGMENTS OF 8051 FINGERPRINT DATA

Device Type	Average Intra-Group X-Corr	Exclusive Average Intra-Vendor X-Corr	Average Inter-Vendor X-Corr
Atmel Trusted Group 1	92.3%	93.3%	49.4%
Atmel Trusted Group 2	96.9%	94.6%	50.5%
Atmel Gray Market	96.3%	90.7%	48.7%
Maxim Trusted	97.9%	94.9%	55.0%
Maxim Gray Market	96.4%	94.6%	54.8%
NXP Trusted	90.7%	61.4%	53.8%
NXP Gray Market	88.4%	52.5%	44.1%

Upon further investigation, we identified variations within the reflection patterns of the package surface depending upon which DUT we tested. These variations could be seen even between the original and second Atmel authentic groups. The observations made from physical inspection should not be considered conclusive in determining the authenticity of a given DUT. In this article, authentic ICs are identified as identical designs in silicon (i.e., the exact placed and routed design). For example, if the Atmel 8051-series ICs are considered authentic, the Maxim and NXP ICs can represent remarked or cloned ICs, as they are functionally similar. Here, our EMFORCED analysis will demonstrate the importance of investigating the physical parameters of the die rather than the visible differences between DUTs.

As previously alluded to, cross-correlation analysis is the fastest implementation of comparing a small number of device fingerprints and determining whether the DUTs are the same type of IC. The primary difference between the output analysis of PCA and cross correlation is that PCA will classify all DUTs within one of N groups, based upon the distance to the nearest training data, while cross correlation will not group all DUTs unless specified above a given cutoff correlation. Additionally, cross correlation will work in all of the scenarios outlined in Section II-D, as it is inherently a reference-free technique as it does not require a golden reference. Table I summarizes our comparisons of all 640 fingerprints against one another. Grouping ICs based upon their vendor and date of acquisition, the seven categories that we tested are provided. It should be noted that in our experimentation, we disabled cross-correlation calculations at nonzero lag conditions as our data were prealigned with the triggering mechanism on our oscilloscope, expediting the processing time. The average intragroup cross correlation determines how alike ICs of a given group are to one another. Empirically, we determined that time-domain cross correlation at or above 90% is typically a good indication of similarity, as two fingerprints from the same device can vary due to measurement noise. In Table I, we note that all but the gray market NXP chips appear well correlated within their respective groups. The exclusively averaged intravendor cross correlation describes how well the tested group aligns with the other group(s) from the same vendor. Here, we begin to identify that the NXP trusted and gray market ICs behave very differently. While both Atmel and Maxim groups are very well correlated with intravendor cross correlations of over 90%, NXP appears to be limited to ≈57%. To ensure that the NXP gray market ICs were not remarked from another IC, we also provided the average intervendor cross correlation that appears consistently low across all device groups.

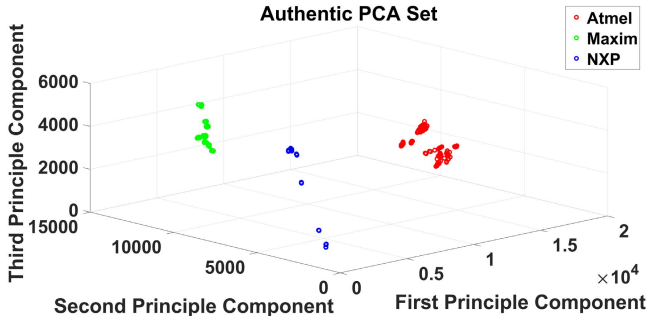

Fig. 5. Projected (3-D) PCA training data for all authentic ICDs using the first three principal components.

B. Reference-Inclusive Experimental Overview

To properly classify any suspect DUT, we have developed a layered approach, as previously described in Section III-C2, which accounts for known reference information during classification. First, a PCA model was trained using our reference fingerprints, which reduced our fingerprint feature dimensionality of 250 000 down to 339. Fig. 5 shows the projections of all authentic fingerprints plotted using their first three principal components. Our authentic reference set consists of all ICs from Atmel trusted groups 1 and 2, Maxim trusted, and NXP trusted. The groupings are easily distinguishable, similar to those seen in Fig. 3, although this model was trained on 340 samples rather than 45, thus strengthening the model. Next, the suspect fingerprint measurements from the gray market ICs are projected into our existing PCA model. The projected suspect data (denoted by an S) has been overlaid onto the reference data in Fig. 6. It can be seen that the Atmel and Maxim suspect DUTs mostly remain within the existing authentic region, while the NXP fingerprints deviate significantly from their designated group, spreading across all different groups. This suspicious behavior reiterates the results acquired in our reference-free cross-correlation analysis. To ensure classification accuracy, we processed the suspect data through outlier detection, which compares the suspect's PCA projections against each of the known good groups. From our analysis, we observed 96% and 97% inlier classification for Atmel and Maxim suspect ICs, respectively. However, NXP suspects classified as outliers 80% of the time. Given this level of deviation from the trusted NXP parts, we can

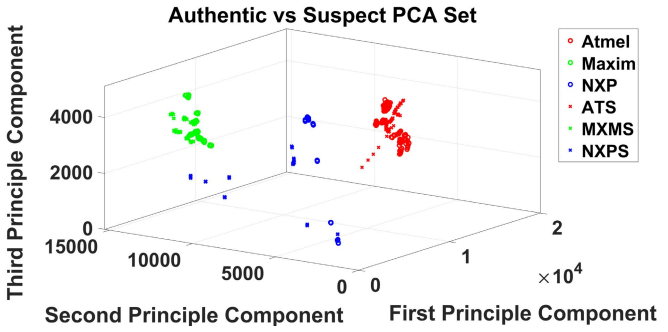

Fig. 6. Comparison between authentic and suspect fingerprints using the first three components of their PCA projection. `ATS`, `MXMS`, and `NXPS` refer to suspect ICs of the respective manufacturers.

comfortably assume that these parts are counterfeit. Alternatively, they may have been recycled and preprogrammed or there was a die revision between the Phillips to NXP transition that physically changed the IC's design [28]. At this point, we had not had any cases of multiple inliers being detected, but to confirm that the 4% outliers identified within Atmel and Maxim were, in fact, authentic, we classified them with a supervised machine learning technique named LDA. We entered all of the PCA projected trusted fingerprint data to LDA and observed 0% resubstitution classification loss. This means that the PCA projected data were perfectly separated as the groups could be clearly defined. Once the suspect Atmel and Maxim parts were predicted using our trained LDA model, we saw 100% accurate classification. Likewise, the NXP chips appeared to deviate into all the class options, leaving some to be identified as Atmel, Maxim, and NXP, furthering our hypothesis. As alluded to earlier during the physical inspection process, the Atmel and Maxim parts would have likely been misclassified as counterfeit parts when they are, in fact, authentic designs. The suspect NXP parts, however, do not exhibit the same characteristics that we should expect from a new NXP 8051 and are identified as highly suspect.

C. Multivariate Analysis

So far, our experimentation has focused exclusively on identifying a given device by using a single parameter. Here, we explore the potential of introducing multiparameter testing and provide insights into how this could maintain high classification accuracy while expanding the number of IC types. Utilizing multiple parameters for counterfeit detection also makes EMFORCED classification more robust and difficult to attack (see Section VI). The tests discussed next focus on EM-based approaches, but combining EM with another measurement, such as power analysis or optical inspection, is also possible.

1) Exploiting Spatial Probing Parameters: In the approach described in Section IV-B, all measurements were taken from the center of the DUT. This was done to maximize the EM response collected from the die itself. In order to study the impact of probe location on the DUT, we measured EM emission from various regions on the IC package. A sample snapshot, with package locations, is shown in Fig. 7, which clearly shows that EM measurements vary between different locations. This approach will provide another degree of difficulty for an attacker attempting to fool the EMFORCED detection framework. When attempting to simulate/clone the EM fingerprint, they would now be required to incorporate the spatial location of the measurement probe in addition to the on-chip EM radiative grid itself.

2) Measurement Distance Accuracy: Having explored the $\langle X, Y \rangle$ plane on the surface of the DUT, we sought to determine whether modifying the proximity of the probe to the surface, or Z-axis, provided any additional information. Changing the probe's vertical location relative to the chip surface is not ideal for extracting circuit switching noise for applications, such as crypto-key extraction (since we desire the highest signal-to-noise ratio). However, since our technique is targeted at extracting a design-specific signature, we are able to see added physical effects from slightly farther-field measurements. Fig. 8 shows the example result of how normalized cross correlation of the acquired signal changes when measuring EM with a near-field probe while introducing a separation distance. Modification of the probe height allows for the radiation to travel a bit further. This would allow the wave to interfere with itself to create new information similar to what was seen in the $\langle X, Y \rangle$ plane, as shown in Fig. 7.

To determine the potential effectiveness of varying the probe proximity, the normalized maximum cross correlation between the signals obtained at different heights was calculated for an Atmel device. In Fig. 8, the blue Atmel line shows the cross correlation between an Atmel measurement while contacting the device (i.e., a separation distance of 0 mm) and at varying distances. For the low separation values with high cross correlation, this implies confidence in the measurement environment, i.e., even with a small height variation, the classification results should remain the same. Upon moving further away from the device (below $\approx$90% cross correlation or $\approx$2 mm), the collected waveforms should contain an adequate amount of new information to be introduced into the analysis. The green and blue lines provide insights into the similarity between the Atmel measurements at a given distance when compared to the Maxim and NXP responses on contact. Notice that at every separation distance, the cross correlation between the Atmel device and the other vendors remains fairly consistent. This shows that the waveforms collected from measuring at a distance provide additional information that not only deviates from the intravendor measurements but also maintains a considerable uniqueness in intervendor comparisons. Finally, it should be noted that the dimensions of the near-field probe control the spatial sensitivity of collected fingerprints. In our experimentation "authentic" ICs are defined as identical silicon layouts. However, in future applications, finer probe dimensions may enable differentiation between device date codes and packaging locations.

3) Rotational Variations: In addition to modifying the relative position in the 3-D space of the near-field probe to the DUT, changing the angle at which the probe is positioned can dramatically impact the observed results. Fig. 9 shows the cross correlation between the rotational measurements on a single DUT. Depending upon the probe type and rotation,

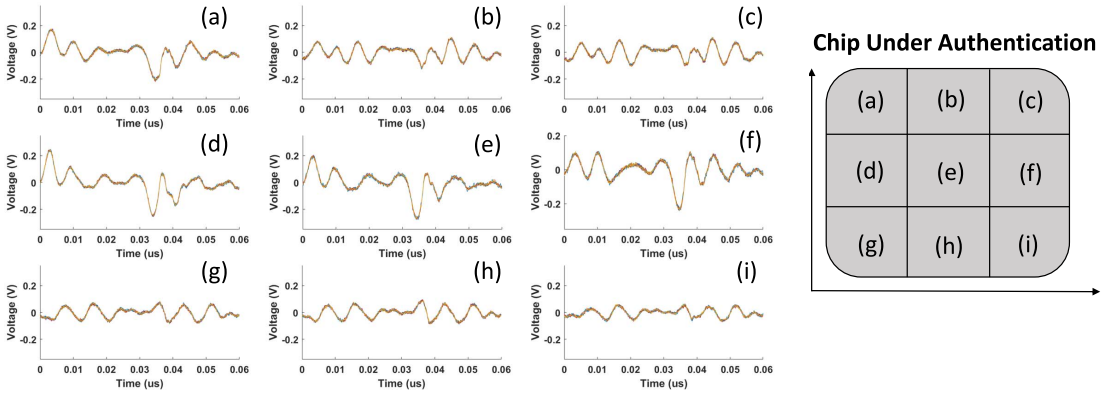

Fig. 7. Coarse sample of the numerous EM fingerprints available through changing the location of the probe relative to the die. (a)–(i) Probe positioning.

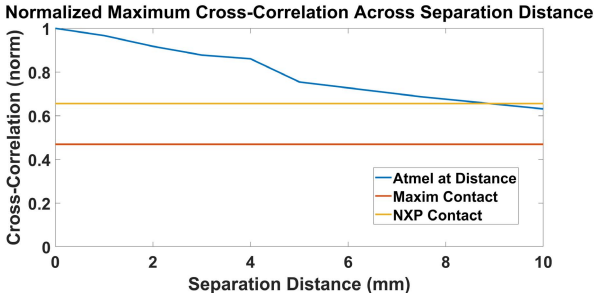

Fig. 8. Normalized maximum cross correlation between measurements from various distances from the DUT.

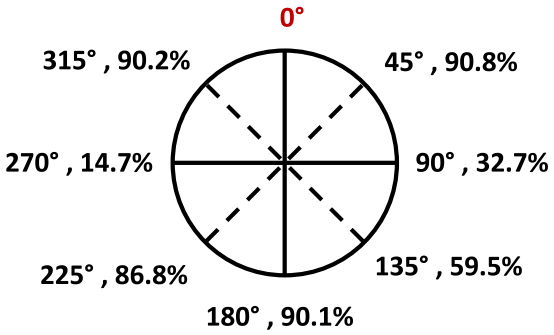

Fig. 9. Cross correlation between rotational measurements on a single chip under test.

the extracted fingerprint can provide significantly different results. The angle measurements (in degrees) in Fig. 9 show that the minimum temporal cross correlation between the original and the new rotation can be seen at the 270° mark with a mere 14.7% likeness. A cross-correlation value such low signifies that there is less in common with this signal than an entirely different IC design. It is important to note that the cross correlation described in Fig. 9 is actually the absolute value of the cross correlation; as such, the 180° correlation value appears as 90.1% rather than the ideal −100% correlation.

4) Input Voltage Variation: Until now, EM signatures were extracted from the ICs while maintaining typical operating conditions, that is, no burn-in tests were performed that might damage the DUT. This is because the EMFORCED framework

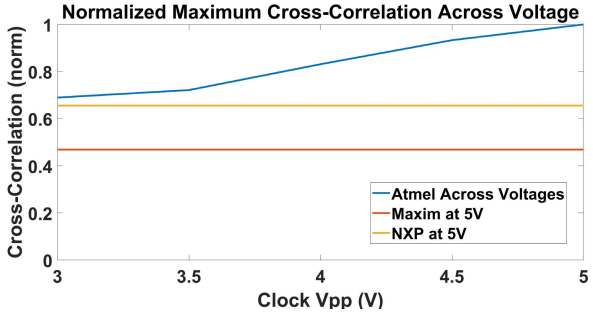

Fig. 10. Normalized maximum cross correlation between the near-field measurements taken with various input voltages.

allows for the DUT, once proven authentic, to be used in the system without losing any projected lifetime. However, to observe the impact of the voltage variation, we applied various input voltages ranging from 3- to 5-V peak-to-peak for the clock signal input. Variation of signal magnitude could provide another dimension of interest when discussing multiparameter analysis. However, it should be noted that utilizing out-of-spec voltages on parts that are intended for use in a system after testing is not recommended because this may create additional faults and cause a reduction in projected device lifetime. Such a method, which puts additional/out-of-spec stress on the device, may be carried out in a postseparation device sampling test to increase confidence in characterization while maintaining low cost. Fig. 10 shows a similar comparison as Fig. 8, with the cross correlation between Atmel devices across a range of voltages and the specific vendor at 5 V. The cross correlation within this voltage range seems to be slightly higher than that seen from the distance analysis although the same trend of uniqueness among both intravendor and intervendor is apparent. As we do not have access to information regarding the die-level operating voltage or the surrounding voltage regulators and level shifters, it is difficult to identify the physical parameters that could be the contributing factors to introduce this deviation. The DUT used in the experiment was not designed to accept an input voltage less than 4.5 V. Therefore, we could fairly assume that the decrease in cross correlation, shown in the blue Atmel line of Fig. 10 below 4.5 V, might be attributed to such factors.

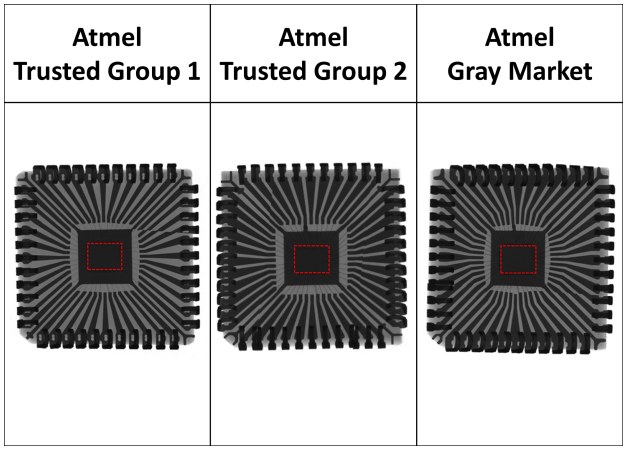

Fig. 11. X-ray comparison between Atmel ICs from each group.

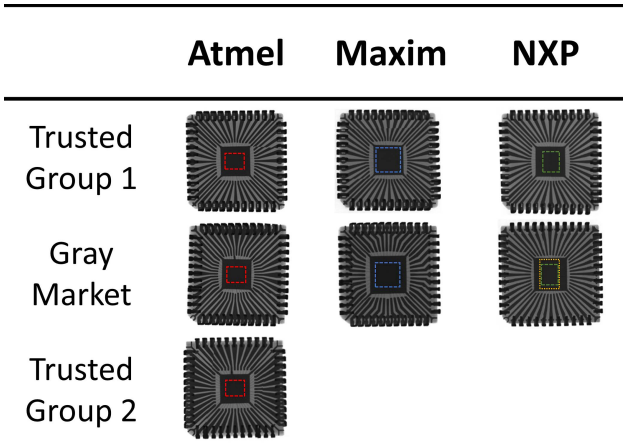

Fig. 12. X-ray comparison of all IC groups with dies outlined in vendor-specific colors, Atmel, Maxim, and NXP in red, blue, and green, respectively. Note that the yellow dotted outline shows the actual die outline of the gray market NXP device.

D. Device Authenticity Analysis via X-Ray

To verify the results shown in section IV-B, an X-ray machine was utilized to see inside of the package and presents further insight. One should note that this process was conducted after fingerprint extraction was completed as this technique can damage older technology devices [29]. X-ray images of the three Atmel groups are shown in Fig. 11. The die area is outlined in red and the dimensions of the die are consistent across all components, providing some confidence that the ICs share the same physical design. However, a closer inspection of the IC frames and bond wires showed slight variations among each group, including double bond wires to a single pin, substrate connections to the ground pin, and pin width variations. These devices would require further inspection as the frame may vary by packaging facility, but given our observations, it is likely that all suspect DUTs would have failed visual inspection or required additional invasive inspection [30].

Expanding our X-ray analysis to the remaining groups shown in Fig. 12 shows that the pad frame varies within each of the vendors. Atmel and Maxim devices maintain the same die dimensions regardless of the pad frame and bond wire characteristics. However, NXP's trusted sample (shown in green) appears to be significantly smaller than the gray market variant (outlined in yellow). Fig. 12 shows the NXP gray market IC with the actual die outline overlaid with the green, expected, die dimensions. Here, we can confidently state that the silicon design is different between the NXP authentic and suspect devices. This difference in die size may be a result of a die revision, likely at a more advanced technology node, since the gray market part was sourced from a batch labeled as Phillips (which was acquired by NXP). Functionally testing these NXP ICs functionally may yield equivalent results as they are listed as the same product, although various applications could take advantage of the benefits from a larger or smaller technology node for reasons such as power efficiency. These X-ray results confirm our EMFORCED results, along with providing evidence against standard physical inspection techniques which would likely misclassify our suspect parts, while showing a significant speed advantage while maintaining classification accuracy.

V. Setup Variability Case Study

The experimental setup used throughout this article remained consistent unless otherwise noted. To transition novel detection methods from laboratory environments to real-world solutions, tolerance to product variations must be accounted for. To demonstrate the variability introduced by the setup itself, three additional Mikroelectronika 8051-Ready development boards were sourced. Upon receipt, the PCBs had noticeable differences among the surface-mounted components, including a different brand or tolerance of power regulation IC, communication IC, and electrolytic capacitor. Despite maintaining the same board revision number of 1.10, these components had been modified. Fingerprints were collected on each of the four development boards using a single Atmel IC for all measurements. Fig. 13 shows the cross correlation among each of the 25 samples per board with representative fingerprints from each PCB. The inter-PCB cross correlation averaged across all PCBs tested was 97.51%. The new PCB group showed an average 97.04% cross correlation across all three PCBs. This shows that the variations introduced by switching identical PCBs were negligible for classification using the EMFORCED framework. As previously stated, a cross-correlation value of below 90% should provide enough differentiating information to be considered a different fingerprint. The average cross correlation between the original and new groups amounted to 89.03%. The differences between the original and newly sourced PCBs may introduce ample variation to be identified as counterfeit. To mitigate this, sourcing identical measurement components and applying calibration methodologies to a collected fingerprint may be desirable. Within the constraints of this article, variations of 89% can be included in the authentic group if desired as the Maxim and NXP components separate below 62% cross correlation.

VI. Attack Resilience Analysis

The primary goal of an attacker would be to replicate an authentic EM fingerprint on a design of his or her choosing. As our fingerprinting technique enables the verifying party to gain insights into the inner workings of the DUT, the attacker would be required to emulate the complex RLC network responsible for modulating the input waveform and emitting it with spatial accuracy. This would likely require substantial

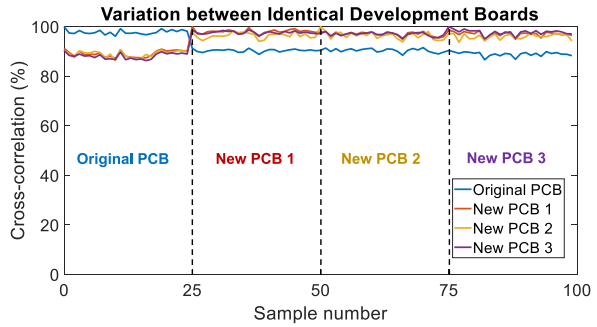

Fig. 13. Cross-correlation comparisons between two iterations of the Mikroelectronika 8051 development board.

assets to not only fabricate the malicious design but also attempt to closely match the EM emissions with temporal and spatial resolution. Modern EM-based simulation tools lack the ability to estimate complex model structures with spatial accuracy while accounting for the effects of the surrounding network. Most EM emission approximation methods rely upon either utilizing the physical parameters of a single fundamental device (e.g., a single transistor) for nano-scale experiments or, through circuit simulation, calculating the total circuit switching activity. These methods will not be successful in generating an accurate EM profile of a given design, as simulations based on physical parameter testing are not scalable at this time. Should an attacker gain access to a simulation tool with this capability, they would be required to modify the placement and routing constraints of their circuitry to comply with the desired fingerprint, which, for more complex circuits, could be nearly impossible. If the attacker was able to create an effective model and an accurate cloned signature when probed from the center of the die, the fabricated version of the device may still not reproduce the same fingerprint, depending upon foundry-specific variations and packaging constraints.

It should be noted that there are certain limitations to the robustness of our approach. For example, the noise introduced by the near-field EM probe and measurement equipment provides an inherent tolerance within the measurement environment. This could create cases where an authentic and counterfeit IC would exist within a certain threshold. Manufacturing variations and spatial resolution of the probe further expand this tolerance. An attacker could potentially use this tolerance to create a cloned signature that conforms to these metrics although modification of a desired design to replicate an authentic fingerprint would require a trial and error methodology. We note that the use of a programmable stage that allows accurate positioning of the probe and reduction of measurement-to-measurement variation should reduce this vulnerability by providing tighter bounds on the classification. Additionally, the introduction of multivariate authentication would further limit the potential for misclassification.

VII. Conclusion

We have demonstrated that our EM-based fingerprinted framework EMFORCED can effectively detect remarked and cloned ICs. Multiple testing setups, including custom breadboard, and standard development board implementations were tested for effective fingerprint extraction with several groups of 8051 ICs. EMFORCED reliably separated the DUTs using reference-free classifications methods that consisted of unsupervised machine learning technique, PCA, and cross correlation. Reference-inclusive scenarios were also addressed with a layered approach using PCA, outlier detection, and, supervised machine learning technique, LDA. This allowed for a confidence interval to be determined when classifying an IC within a known group. To provide a glimpse of the potential of EMFORCED to be scaled with multiple device types and tunable confidence intervals, multiple variables enabled by EM were analyzed within this article, showing the numerous fingerprints that can be extracted from a single IC. The low cost, ease of implementation, and potential scalability of EMFORCED were provided a competent base for remarked and cloned counterfeit IC detection within today's increasingly vulnerable electronics supply chain.

References

[1] *Inquiry into Counterfeit Electronic Parts in the Department of Defense Supply Chain*, United States Senate Committee Armed Services, Washington, DC, USA, 2012.
[2] J. Villasenor and M. Tehranipoor, "Chop shop electronics," *IEEE Spectr.*, vol. 50, no. 10, pp. 41–45, Oct. 2013.
[3] M. Rostami, F. Koushanfar, J. Rajendran, and R. Karri, "Hardware security: Threat models and metrics," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2013, pp. 819–823.
[4] U. Guin, D. DiMase, and M. Tehranipoor, "Counterfeit integrated circuits: Detection, avoidance, and the challenges ahead," *J. Electron. Test.*, vol. 30, no. 1, pp. 9–23, 2014.
[5] K. Huang, Y. Liu, N. Korolija, J. M. Carulli, and Y. Makris, "Recycled IC detection based on statistical methods," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 34, no. 6, pp. 947–960, Jun. 2015.
[6] U. Guin, D. Forte, and M. Tehranipoor, "Anti-counterfeit techniques: From design to resign," in *Proc. 14th Int. Workshop Microprocessor Test Verification*, Dec. 2013, pp. 89–94.
[7] G. E. Suh and S. Devadas, "Physical unclonable functions for device authentication and secret key generation," in *Proc. 44th ACM/IEEE Design Autom. Conf.*, New York, NY, USA, Jun. 2007, pp. 9–14.
[8] B. Gassend, D. Clarke, M. Van Dijk, and S. Devadas, "Silicon physical random functions," in *Proc. 9th ACM Conf. Comput. Commun. Secur.*, Nov. 2002, pp. 148–160.
[9] K. Ahi, N. Asadizanjani, S. Shahbazmohamadi, M. Tehranipoor, and M. Anwar, "Terahertz characterization of electronic components and comparison of terahertz imaging with X-ray imaging techniques," *Proc. SPIE*, vol. 9483, May 2015, Art. no. 94830K.
[10] K. Huang, J. M. Carulli, and Y. Makris, "Parametric counterfeit IC detection via support vector machines," in *Proc. IEEE Int. Symp. Defect Fault Tolerance VLSI Nanotechnol. Syst. (DFT)*, Oct. 2012, pp. 7–12.
[11] T. D. Bergman, C. P. Manager, and K. T. Liszewski, "Battelle barricade: A nondestructive electronic component authentication and counterfeit detection technology," in *Proc. IEEE Symp. Technol. Homeland Secur. (HST)*, May 2016, pp. 1–6.
[12] T. S. Messerges, E. A. Dabbish, and R. H. Sloan, "Examining smart-card security under the threat of power analysis attacks," *IEEE Trans. Comput.*, vol. 51, no. 5, pp. 541–552, May 2002.
[13] P. Song, F. Stellari, and A. Weger, "Counterfeit IC detection using light emission," in *Proc. Int. Test Conf.*, Oct. 2014, pp. 1–8.
[14] H. P. Romero, K. A. Remley, D. F. Williams, and C.-M. Wang, "Electromagnetic measurements for counterfeit detection of radio frequency identification cards," *IEEE Trans. Microw. Theory Techn.*, vol. 57, no. 5, pp. 1383–1387, May 2009.
[15] C. Capps, "Near field or far field?" *EDN*, vol. 46, no. 18, pp. 95–99, 2001.
[16] M. Lacruche and P. Maurine, "Electromagnetic activity vs. logical activity: Near field scans for reverse engineering," in *Smart Card Research and Advanced Applications*. Cham, Switzerland: Springer, 2018, pp. 140–155.
[17] O. Sinanoglu *et al.*, "Reconciling the IC test and security dichotomy," in *Proc. 18th IEEE Eur. Test Symp. (ETS)*, May 2013, pp. 1–6.
[18] A. Stern, U. Botero, B. Shakya, H. Shen, D. Forte, and M. Tehranipoor, "EMFORCED: EM-based fingerprinting framework for counterfeit detection with demonstration on remarked and cloned ics," in *Proc. IEEE Int. Test Conf. (ITC)*, Oct./Nov. 2018, pp. 1–9.

[19] A. Lakshminarasimhan, "Electromagnetic side-channel analysis for hardware and software watermarking," Univ. Massachusetts Amherst, Amherst, MA, USA, Tech. Rep. 693, 2011.
[20] M. Muehlberghuber, F. K. Gürkaynak, T. Korak, P. Dunst, and M. Hutter, "Red team vs. blue team hardware trojan analysis: Detection of a hardware trojan on an actual ASIC," in *Proc. 2nd Int. Workshop Hardw. Architectural Support Secur. Privacy*, Jun. 2013, Art. no. 1.
[21] O. Söll, T. Korak, M. Muehlberghuber, and M. Hutter, "EM-based detection of hardware trojans on FPGAs," in *Proc. IEEE Int. Symp. Hardw.-Oriented Secur. Trust (HOST)*, May 2014, pp. 84–87.
[22] J. Balasch, B. Gierlichs, and I. Verbauwhede, "Electromagnetic circuit fingerprints for hardware Trojan detection," in *Proc. IEEE Int. Symp. Electromagn. Compat. (EMC)*, Aug. 2015, pp. 246–251.
[23] J. Shlens, "A tutorial on principal component analysis," *CoRR*, vol. 1404, pp. 1–12, Apr. 2014.
[24] G. J. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*, vol. 544. Hoboken, NJ, USA: Wiley, 2004.
[25] *8-bit Low-Voltage Microcontroller with 8K Bytes In-System Programmable Flash*, Atmel, San Jose, CA, USA, 2008.
[26] *RF1 set, Near-Field Probes 30 MHz up to 3 GHz*, Langer EMV-Technik GmbH, Bannewitz, Germany, 2019.
[27] S.-H. Cha, "Comprehensive survey on distance/similarity measures between probability density functions," *City*, vol. 1, no. 2, p. 1, 2007.
[28] J. Yoshida, "Philips semiconductors to become NXP," EE Times (now ASPENCORE), Cambridge, MA, USA, Tech. Rep. 1163319, Aug. 2006.
[29] M. Alam, H. Shen, N. Asadizanjani, M. Tehranipoor, and D. Forte, "Impact of X-ray tomography on the reliability of integrated circuits," *IEEE Trans. Device Mater. Rel.*, vol. 17, no. 1, pp. 59–68, Mar. 2017.
[30] K. Mahmood, P. L. Carmona, S. Shahbazmohamadi, F. Pla, and B. Javidi, "Real-time automated counterfeit integrated circuit detection using X-ray microscopy," *Appl. Opt.*, vol. 54, no. 13, pp. D25–D32, 2015.

Andrew Stern (S'12) received the B.S. degree in electrical and computer engineering from the University of Rochester, Rochester, NY, USA, in 2016, and the M.S. degree in electrical and computer engineering from the University of Florida, Gainesville, FL, USA, in 2018, where he is currently working toward the Ph.D. degree.

He is currently a GAANN Fellow with the University of Florida. His current research interests include the domain of hardware security and trust and include electronics supply chain security, counterfeit IC detection and prevention, IP piracy protection, and side-channel analysis.

Ulbert Botero (S'16) received the B.S. degree in electrical engineering from the University of Central Florida, Orlando, FL, USA, in 2016, and the M.S. degree in electrical and computer engineering from the University of Florida, Gainesville, FL, USA, in 2019, where he is currently working toward the Ph.D. degree at the Electrical and Computer Engineering Department.

His current research interests include reverse engineering and hardware security and trust.

Mr. Botero is an NSF Fellow.

Fahim Rahman (S'13–M'19) received the B.S. degree in electrical and electronic engineering from the Bangladesh University of Engineering and Technology, Dhaka, Bangladesh, the M.S. degree in electrical and computer engineering from the University of Connecticut, Storrs, CT, USA, in 2015, and the Ph.D. degree in electrical and computer engineering from the University of Florida, Gainesville, FL, USA, in 2018.

He is currently a Research Assistant Professor with the Electrical and Computer Engineering Department, University of Florida. His current research interests include the domain of hardware and cybersecurity and trust, including investigation of hardware security primitives, CAD for security and automatic assessment, electronic supply chain security, and hardware-assisted cybersecurity. His research has been sponsored by Semiconductor Research Corporation (SRC), Air Force Office of Scientific Research (AFOSR), Air Force Research Laboratory (AFRL), Defense Advanced Research Projects Agency (DARPA), Cisco, Texas Instruments (TI), and National Institute of Standards and Technology (NIST).

Dr. Rahman is a member of the ACM.

Domenic Forte (S'09–M'13–SM'18) received the B.S. degree from the Manhattan College, Riverdale, NY, USA, in 2006, and the M.S. and Ph.D. degrees from the University of Maryland, College Park, MD, USA, in 2010 and 2013, respectively, all in electrical engineering.

From 2013 to 2015, he was an Assistant Professor with the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT, USA. Since July 2015, he has been with the Electrical and Computer Engineering Department, University of Florida, Gainesville, FL, USA, where he is currently an Associate Professor. His current research interests include the entire domain of hardware security from nanodevices to printed circuit boards (PCBs), including hardware security primitives, hardware Trojan detection and prevention, security of the electronics supply chain, security-aware design automation tools, reverse engineering, and antireverse engineering.

Dr. Forte serves on the organizing committees of top conferences in hardware security, such as the IEEE Symposium on Hardware Oriented Security and Trust (HOST) and the AsianHOST. He has been serving on the technical program committees in the areas of electronic design automation, VLSI design and test, and cybersecurity. He was a recipient of the Presidential Early Career Award for Scientists and Engineers (PECASE), the NSF Faculty Early Career Development Program (CAREER) Award, and the Army Research Office (ARO) Young Investigator Award. His research has been also recognized through nine best paper awards and nominations.

Mark Tehranipoor (F'18) is currently the Intel Charles E. Young Preeminence Endowed Chair Professor of Cybersecurity with the University of Florida (UF), Gainesville, FL, USA. Prior to joining UF, he has served as the Founding Director for CHASE and CSI Centers, University of Connecticut, Storrs, CT, USA. He has been serving as the Founding Director for the Florida Institute for Cybersecurity Research (FICS). He has published over 400 journal articles and refereed conference articles and has delivered about 200 invited talks and keynote addresses. He has published 11 books and more than 20 book chapters. His current research interests include hardware security and trust, supply chain security, IoT security, VLSI design, test, and reliability.

Dr. Tehranipoor is a Golden Core Member of IEEE CS and a member of ACM and ACM SIGDA. He was a recipient of a dozen best paper awards and nominations, as well as the 2008 IEEE Computer Society (CS) Meritorious Service Award, the 2009 NSF CAREER Award, the 2012 IEEE CS Outstanding Contribution, and the 2014 Air Force Office of Scientific Research (AFOSR) MURI Award. He serves on the program committee of more than a dozen leading conferences and workshops. He has served as the Program Chair of a number of IEEE and ACM sponsored conferences and workshops, such as the IEEE International Symposium on Hardware-Oriented Security and Trust (HOST), International Test Conference (ITC), International Symposium on Defect and Fault Tolerance in VLSI Systems (DFTS), Defect and Data Driven Testing (D3T) Workshop, Defect-Based Testing (DBT) Workshop, and North Atlantic Test Workshop (NATW). He has served as the Program Chair of the 2007 IEEE DBT and the 2008 IEEE D3T, the Co-Program Chair of the 2008 DFTS, the General Chair of the D3T-2009 and DFTS-2009, and the Vice-General Chair of the NATW-2011. He co-founded the IEEE HOST and has served as the HOST-2008 and HOST-2009 General Chair. He has been serving as the Founding EIC for the *Journal of Hardware and Systems Security* (HaSS) and an Associate Editor for the *Journal of Electronic Testing: Theory and Applications*, the *Journal of Low Power Electronics*, the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS, and the *Transactions on Design Automation of Electronic Systems* (ACM).