Side Channel Authenticity Discriminant Analysis for Device Class Identification

Eric Koziel, Kate Thurmer,
Lauren Milechin, Peter Grossmann,
Michael Vai, Roger Khazan
MIT Lincoln Laboratory
Lexington, MA, USA

Keith Bergevin, Philip Comer
Defense Microelectronics Activity (DMEA)
McClellan, CA, USA

Contact Author: eric.koziel@ll.mit.edu

Abstract: *Counterfeit microelectronics present a significant challenge to commercial and defense supply chains. Many modern anti-counterfeit strategies rely on manufacturer cooperation to include additional identification components. We instead propose Side Channel Authenticity Discriminant Analysis (SICADA) to leverage physical phenomena manifesting from device operation to match suspect parts to a class of authentic parts. This paper examines the extent that power dissipation information can be used to separate unique classes of devices. A methodology for distinguishing device types is presented and tested on both simulation data of a custom circuit and empirical measurements of Microchip dsPIC33F microcontrollers. Experimental results show that power side channels contain significant distinguishing information to identify parts as authentic or suspect counterfeit.*

Keywords: counterfeit detection; device authentication; supply chain security; hardware identity

Introduction and Background

The prevalence of counterfeit microelectronics in government and commercial supply chains poses a significant threat to the reliability of government systems. The Department of Defense now mandates that organizations and contractors proactively detect and avoid counterfeit parts [1]. This is complicated by the fact that many current counterfeit detection methods are too expensive, time-consuming, or destructive to scale to a full acquisitions supply chain [2].

Counterfeit microelectronics are any electronic components that are misrepresented in sale. Counterfeit types include remarked parts, reproduced or “cloned” parts, and recycled parts that have been used in prior systems. The sophistication of counterfeits can vary from simple re-etching and blacktopping to full netlist-level reproduction of parts [3].

Several current and developing technologies aim to prevent counterfeiting. In particular, DARPA’s SHIELD program [4] aims to add dielets to microelectronic components that assert the part’s identity and authenticity. Other technologies add additional circuitry [5] to give devices a unique, uncloneable signature. These anti-counterfeiting approaches require manufacturer participation and incur some additional overhead cost to execute. These approaches also cannot be easily applied to parts produced in the past.

Instead of leveraging other components to assert identity, it may be possible to utilize a device’s intrinsic operational characteristics. Such a technique could be used to detect misrepresented parts regardless of manufacturer participation. In [6], Cobb et al. use electromagnetic side channels to uniquely identify small microcontrollers. Their work was successful in identifying individual parts based on generated templates, but also noted that parts of a given type demonstrated many similar characteristics.

In this paper, we propose a methodology for utilizing power side channel information to compare devices. This approach expands [6] by matching entire classes of parts instead of individual devices, and also uses power information instead of electromagnetic emanations. This initial work investigates differences in similar device types in side channel space by analyzing both simulations and commercial devices. Our experiments test several counterfeit-relevant use cases to determine the effectiveness of this approach at distinguishing authentic classes of devices from misrepresented types. We also examine which waveform characteristics contribute most strongly to classification decisions. Due to sharing restrictions, please contact the authors for related work references.

Methodology

Our main approach was to identify and compare power signatures of similar devices with a focus on applicability to cases of misrepresented parts. This includes devices produced with different manufacturing processes, or functionally-equivalent devices with different internal structures. SPICE-level simulations are ideal for such cases since the manufacturing process and circuit can be altered relatively easily. We also collected empirical data from commercial microcontrollers. The experimental setup aimed to compare several variants of the same core architecture, for example one with more memory, and another with additional functional units. Different environmental grades and date codes of the same device type were also compared.

Experimental Setup: Simulation data were collected from a custom low-power 32-bit Multiply-Accumulate (MAC) circuit designed in IBM's 45nm process technology. The MAC was simulated with Cadence's UltraSim accelerated SPICE simulator. The UltraSim speed and accuracy settings were chosen to balance simulation runtime with accuracy of the power dissipation estimate. A functionally equivalent variant of the MAC circuit requiring 3% fewer logic gates was also simulated. The larger circuit was also translated to IBM's 65nm process to isolate the effects of process technology on power analysis results. All variants were measured at the average, slow-slow, and fast-fast process technology corners (Table 1). We collected 1000 iterations of 40-cycle operational loops for each simulation.

Table 1: MAC SPICE Simulation Type List

Label	Process	Gate Count
45nm	IBM 45nm	15177
loose	IBM 45nm	14667
65nm	IBM 65nm	15177

Empirical measurements were taken from Microchip's dsPIC33F family of microcontrollers (Table 2). We used j12 parts as the base device, while j32 and j128 parts had larger memory sizes. j128 devices have more comparators, DACs, and timers compared to either j12 or j32 devices. Each type is further broken down by grade, where I-type parts are standard industrial grade and E-type parts are extended temperature grade parts.

Table 2: Microchip dsPIC33F Type List

Label	Device	Date Code	Quantity
j12i	J12GP-202 I	074047Y	9
j12i	J12GP-202 I	13490M1	6
j12e	J12GP-202 E	13226M0	6
j32i	J32GP-202 I	1421134	18
j32e	J32GP-202 E	1332P0Q	6
j128i	J128GP-802 I	1431RK9	15
j128i	J128GP-802 I	1320J7W	6
j128i	J128GP-802 I	1526RBE	3
j128e	J128GP-802 E	1229MST	10
j128e	J128GP-802 E	1421YRM	6

Power information was collected via custom sensing circuits and captured on an oscilloscope. Each microcontroller was loaded with a program to execute 1000 iterations of a consistent operational loop utilizing a mixture of arithmetic, register, and memory operations. This operational loop comprised 53 individual clock cycles. Each device was clocked via an external pulse generator at 10 MHz to avoid timing inconsistencies in factory calibration settings.

Analysis Process: We processed both simulation and empirical data using a feature generation approach adapted from [5]. Power traces were broken down by clock cycle and transformed into Hilbert analytical signals, from which the instantaneous amplitude, instantaneous phase, and instantaneous frequency were derived (Figure 1). The first four statistical moments (mean, variance, kurtosis, and skewness) were gathered from each of these waveforms, in addition to the standard deviation. We also applied this analysis to the full signal. This generated a total of 615 features for simulations and 810 features for empirical measurements. Each individual feature is considered a dimension of information, however many features were highly correlated.

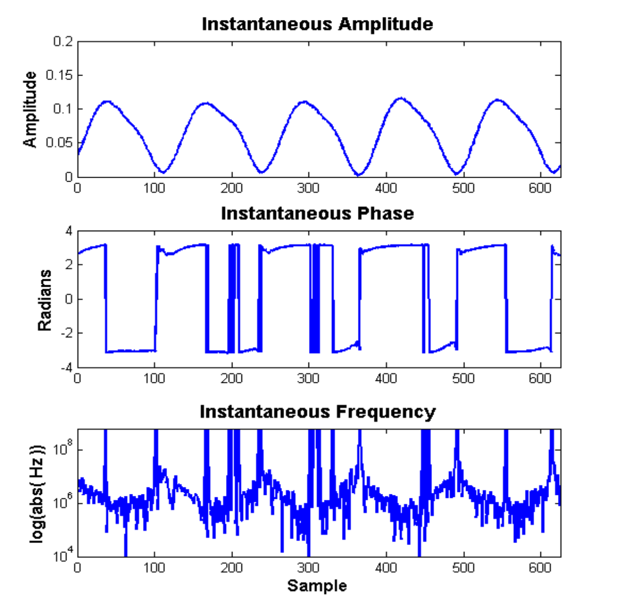

Figure 1: dsPIC33F Instantaneous Data Example

We hypothesize that these features contain significant distinguishing information so that unlike parts can be differentiated from one another. We tested this hypothesis by using a support vector machine (SVM) classifier to attempt to separate a "golden" set of observations from other groups of observations. The null hypothesis is that both sets are different, and thus would be easy to separate. If cross-validation of the resulting model shows significant error, it means that the two sets could not be easily separated and are likely from the same type of device.

For each comparison, we centered and normalized the data before projecting to a lower dimensionality using Principal Components Analysis (PCA). We settled on projecting down to 3 principal components to avoid overfitting, which typically explained more than 60% of variance. All projected observations were then used to create a 10-fold cross-validation SVM model.

Another goal of this research was to identify which source features contribute the most towards correct classification decisions. Narrowing the number of features used in analysis can lead to more accurate classification decisions and a simplified methodology. Additionally, individual features could potentially be tied back to manufacturing process variation differences or other physical phenomena for a better understanding of what contributes to a device's side intrinsic side channel characteristics. We used the Multiple Cluster Feature Selection (MCFS) [7] algorithm to generate ordered lists of features according to their contribution to a given comparison. A Naïve Bayes classifier was used for MCFS comparisons to allow for training with more than 2 classes.

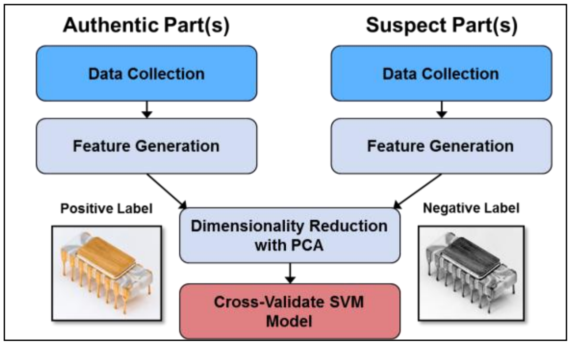

Figure 2: SICADA Analysis Workflow

Results

Simulation Results: We used the simulation data to test two primary counterfeit cases: slight circuit variations and "cloned" parts produced with a different manufacturing process. The former case deals with misrepresented parts, especially for authentic devices within the same architectural family but of a different model or grade. Cloned parts comprise any part reproduced with the exact netlist as an authentic part, but using a different manufacturing process.

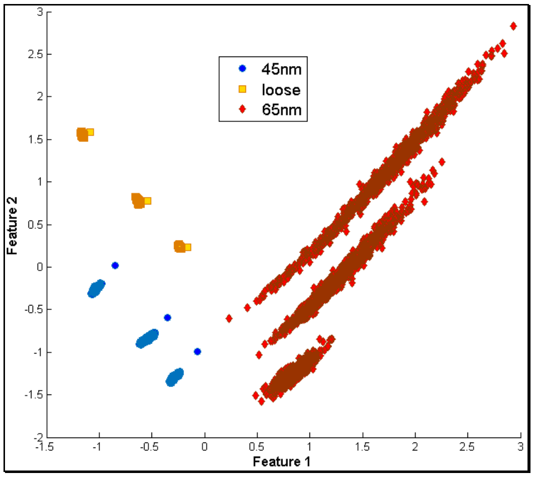

Figure 3: Simulation Data PCA Projection – 2 Principle Components

Figure 3 is a PCA projection of the data down to 2 principal components and then re-centered and normalized for ease of presentation. Even without classification, each major device type is visibly distinct. The three unique clusters for each type represent the individual corners, so typical device performance would be expected to fall between the clusters. Note that the measurements are much more widely distributed with the 65nm part compared to either 45nm part. This makes sense that the larger processing node has a larger magnitude of variations in operating characteristics, but this is a significant observation in that detecting a part produced with a different node can be as simple as testing the variance of measurements.

We achieved perfect classification for every comparison tested, meaning that there was no cross-validation loss. Instead, we report the cross-validation edge in Table 3, which is the average distance of observations from misclassification.

Table 3: Simulation Data Edges and Best Features

Simulation	Cross-Validation Edge	Best Feature	1-Feature Error Rate
45nm vs loose	2.408	Inst. Frequency Skew of clock 16	0
45nm vs 65nm	4.642	Inst. Amplitude Skew of clock 15	0
loose vs 65nm	4.794	Inst. Amplitude Skew of clock 28	0

The best features for individual comparisons are also listed. For those specific comparisons, the listed feature alone was enough to achieve perfect classification. Using only the instantaneous frequency and amplitude skewness for clocks 16 and 15, respectively, resulted in perfect classification of all 3 classes with a Naïve Bayes classifier.

Empirical Results: We designed the experimental tests around determining how different individual authentic devices are from one another. Specifically, we tested whether devices exhibit substantially different operating characteristics from one another based on the presence of additional memory or functional units. We also tested whether different temperature grades of parts could be distinguished. Results for a subset of the comparisons are tabulated in Table 4. We consider error rates greater than 0.05 to indicate parts as similar, and rates greater than 0.10 to indicate likely equivalence. These values are not intended to be universal and would need to be tuned for other experiments.

The first thing to note is that through most separability comparisons, j12i parts were not distinguishable from j32i parts. The difference between these parts is only memory size. However, j128i were easily distinguishable. j128i parts have additional functional units compared to j12i and j32i parts. Each E-type part is easily separated from other E-Types. However, different temperature grades of the same part are not easily separated. This makes sense as they likely are identical parts separated only by the binning process of the manufacturer.

We applied MCFS to each comparison as well. We found that relatively few features were required to achieve perfect classification in most comparisons. Instantaneous frequency mean of the full signal was particularly strong at distinguishing E-type parts from each other as well as from their I-type counterparts. Figure 4 shows a few of the major

Table 4: dsPIC33F Comparison Error Rates and Best Features

Comparison	Cross-Validation Error Rate	Best Feature	One-Feature Error Rate
j12i vs j32i	0.3256	Inst. Amplitude StdDev of full signal	0.1718
j12i vs j128i	0.0001	Inst. Amplitude Mean of clock 32	0
j32i vs j128i	0	Inst. Amplitude Mean of clock 32	0
j12e vs j32e	0.0001	Inst. Frequency Mean of full signal	0
j12e vs j128e	0.0108	Inst. Frequency Mean of full signal	0
j32e vs j128e	0	Inst. Frequency Mean of full signal	0
j12i vs j12e	0.0533	Inst. Frequency Mean of full signal	0
j32i vs j32e	0.1714	Inst. Frequency Mean of full signal	0
j128i vs j128e	0.3754	Inst. Amplitude Kurtosis of clock 0	0.2833

distinguishing features for I-type parts; using these three features gives a 0.09 error rate. One observation from this set is that the most remote groupings of j12i measurements belong almost entirely to the parts manufactured in 2007. Other parts in the comparison were produced in 2013 or later, indicating that physical age may affect some features.

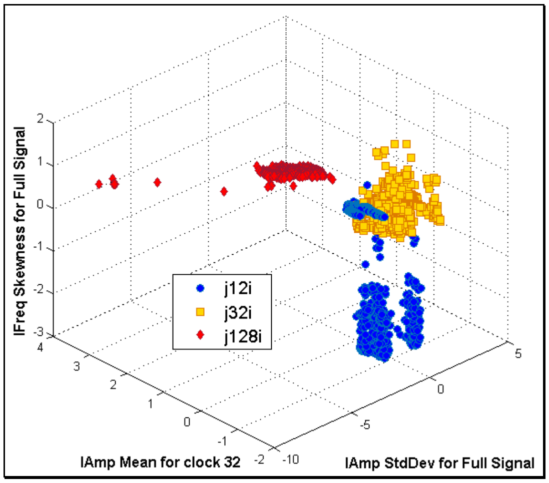

Figure 4: dsPIC33F Data - 3 Features

We also compared different date codes for the j128i part. The SVM classifier was unable to clearly separate the parts, and most cross-validation error rates were greater than 0.30. However, using instantaneous phase skewness for the full signal as the only feature gives perfect classification for each of the 3 datecodes. Distinguishing by date code may be useful in specific counterfeit detection use cases.

Conclusions and Future Work

These initial results indicate that power side channels of devices contain significant distinguishing information. The presented experiments provide a basic example of using the SICADA methodology to investigate device differences and identify classes of devices.

The simulation results provide evidence that SICADA could be used as a foundation for detection of certain types of clone parts. Simulations also support the ability to distinguish functionally-equivalent devices with differing implementations. The empirical results indicate an ability to differentiate parts with different functional units. Several features were very useful in differentiating device types.

Currently, we are refining the SICADA approach as we expand it to other types of counterfeits, such as recycled parts. Other possible improvements include tying specific features to physical phenomena, including additional features, applying more advanced machine learning techniques, and examining other types of side channels.

Acknowledgements

This material is based upon work supported by DMEA under Air Force Contract No. FA8721-05-C-0002. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DMEA.

References

1. Federal Register (2014). *Defense Federal Acquisition Regulation Supplement: Detection and Avoidance of Counterfeit Electronic Parts*. DFARS Case 2012-D055.

2. Government Accountability Office (2012). *Suspect Counterfeit Electronic Parts Can Be Found on Internet Purchasing Platforms*. GAO -12-375.

3. Guin, U; DiMase, D; Tehranipoor, M. *Counterfeit Integrated Circuits: Detection, Avoidance, and the Challenges Ahead*. Journal of Electronic Testing, vol. 30 no. 1, Feb 2014.

4. Defense Advanced Research Agency. (2014). *Supply Chain Hardware Integrity for Electronics Defense (SHIELD)*. DARPA-BAA-14-16.

5. Suh, G.E.; Devadas, S. *Physical Unclonable Functions for Device Authentication and Secret Key Generation*. 44th ACM/IEEE Design Automation Conference, June 2007.

6. Cobb, W; Laspe, E; Baldwin, R; Temple, M; Kim, Y. *Intrinsic Physical-Layer Authentication of Integrated Circuits*. IEEE Transactions on Information Forensics and Security, vol.7, no.1, Feb. 2012.

7. Cai D; Zhang C; He X. *Unsupervised Feature Selection for Multi-cluster Data*. 16th ACM Conference on Knowledge Discovery and Data Mining, July 2010.