

An Overlay-Based Correlated Noise Generation Countermeasure against Statistical Power Analysis

Austin Baylis, Greg Stitt, Ann Gordon-Ross
 Department of Electrical and Computer Engineering
 NSF Center for High-Performance Reconfigurable Computing (CHREC)
 University of Florida, Gainesville, FL, USA

Abstract—Statistical power analysis is an increasing hardware-security concern that enables attackers to reverse-engineer encryption keys for field-programmable gate array (FPGA) bitfiles, resulting in intellectual-property (IP) theft and tampering. To address this problem, we present a countermeasure that integrates correlated noise generation with overlays—virtual architectures implemented atop an FPGA—to protect application IP, even on vulnerable FPGAs. By extending existing overlay benefits (e.g., fast compilation, application portability, 1000x smaller bitfiles) with improved hardware security, our approach provides an attractive platform for Internet of Things, defense, and many embedded applications.

Keywords—*overlay, FPGA, differential power analysis, hardware security, virtualization, side-channel attacks*

I. INTRODUCTION

Field-programmable gate arrays (FPGAs) are increasingly used for embedded and data-center applications due to their performance, power, and/or energy advantages over other technologies [3]. However, FPGAs suffer from hardware-security concerns that enable attackers to perform side-channel attacks that can reverse-engineer encryption keys, extract intellectual property (IP), and maliciously tamper with functionality. Such attacks are potentially devastating for Internet of Things (IoT) applications, where billions of deployed units can be maliciously modified. Similarly, attacks on safety-critical applications in defense and automotive applications can result in loss of human life.

One increasing concern is statistical power analysis (e.g., Differential [9] and Correlation Power Analysis [1] attacks), where attackers can reverse-engineer encryption keys via statistical analyses of power consumption during operations that access a key (e.g., [7][8]). Although previous work has introduced numerous countermeasures [10], statistical power analysis is still a common concern, even for FPGAs marketed as secure [11].

One of the main limitations of existing countermeasures is a lack of protection for the FPGA’s encryption key that is used to decrypt application bitfiles during configuration. Such protection is challenging because FPGAs provide no mechanisms to control functionality during configuration, where the device is in an undefined state. Although there are previous countermeasures integrated with AES cores (e.g., [8]), existing FPGAs are not manufactured with these cores, making existing FPGA bitfiles vulnerable to reverse engineering.

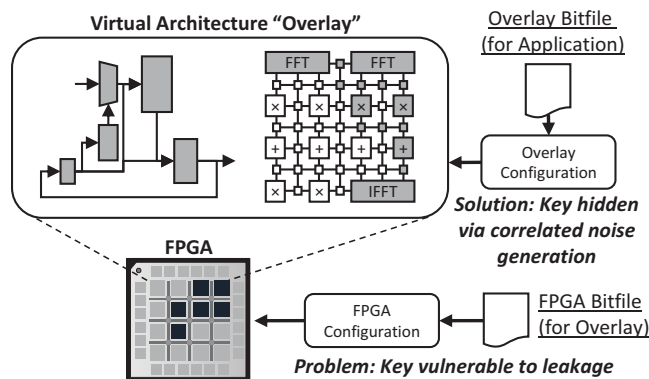


Figure 1: Overview of overlay-based correlated noise generation. In our approach, the FPGA provides an application-specialized virtual architecture (i.e., an *overlay*). Whereas FPGA configuration is vulnerable to power-analysis attacks, overlay configuration protects encryption keys by injecting input-correlated power noise into the FPGA.

To address this concern, we present an approach that overcomes the lack of protection for the FPGA’s key by using a virtual architecture implemented on the FPGA (i.e., an *overlay*) to hide leakage of keys and/or other secret information. Overlays have been introduced for a variety of purposes including fast placement and routing, application portability, simplified debugging, among others [2][4][5]. A recent hardware-security study investigated using unique overlays to limit damage from tampering [12]. In this paper, we complement earlier overlay studies by introducing statistical power analysis countermeasures that can potentially be integrated into any overlay.

Figure 1 provides an overview of our approach, which implements applications on the overlay as opposed to directly on the FPGA. Because of this two-layer approach, the system uses two separate bitfiles: 1) a vulnerable FPGA bitfile that configures the FPGA with the overlay, and 2) the protected overlay bitfile that contains all application IP. Although the FPGA’s encryption key is still vulnerable in this approach, if an attacker were to discover the FPGA key, reverse engineering the bitfile would only provide the architecture of the overlay, which provides little application information [12].

Unlike FPGA configuration, overlay configuration has control over the functionality of the FPGA, which we exploit to inject power into the FPGA via toggling of nets. By correlating this power with an input to be encrypted/decrypted (i.e., *correlated noise generation*), our approach provides the same




Figure 2: Correlation Power Analysis relies on (a) a relationship between Hamming distance and voltage, which we (b) obfuscate using correlated noise generation to prevent leakage.

noise for repeated attacks using the same input, and different noise for different inputs, which prevents Differential and Correlation Power Analysis by yielding correlations with an incorrect key. We evaluate the countermeasure against Correlation Power Analysis attacks on DES using publicly available power traces [6] and show that such attacks yield incorrect keys.

Although we present the countermeasure in the context of overlays, we could similarly apply the approach to protect FPGA bitfiles by using partial reconfiguration, where one region performs correlated noise generation to protect the bitfiles in other regions.

II. STATISTICAL POWER ANALYSIS COUNTERMEASURE

This section presents the details of our statistical power analysis countermeasure. The approach consists of two parts: 1) correlated noise generation to determine the type of noise to be injected to hide leakage (Section II.A), and 2) a power injector core that can increase and decrease power by specific amounts each cycle (Section II.B).

A. Correlated Noise Generator

Statistical power analysis attacks create a statistical relationship between changes that occur during an operation and the measured power to extract leaked information. The classical statistic power analysis is Differential Power Analysis (DPA) [9] where the attacker guesses input values for a model of the operation they wish to attack to generate potential output values. The values are used to sort the power traces into two sets. The attacker then takes the difference of the mean of each set. Incorrect guesses average to be similar, resulting in a small difference, while a correct guess results in a large peak at the location the operation occurs. An improvement on this attack is to approximate power consumption using the Hamming weight of the output value generated by the input guess. The Hamming weight can then be used to build a correlation between the guess and the power traces. The correct guess would be associated with the highest correlation. Both previously

mentioned methods of DPA, however, are limited due to unrealistic modeling of power. Therefore, the attack we used to test the effectiveness of our correlated noise generator is Correlation Power Analysis (CPA), as described in [1]. It is worth noting that while we explore CPA in this paper, CPA countermeasures tend to provide the same defensive efficiency against the other statistical analysis attacks [1].

CPA is a model-based attack, much like DPA using Hamming weight. The fundamental difference, however, is that while DPA uses the model to generate a value based on a guessed state and then takes the Hamming weight of that value, CPA takes the Hamming distance. The reason for using Hamming distance is because data leakage depends on the number of bits that flip from one state to another. The model we use assumes that a bit flipping from 1 to 0 requires the same amount of energy as a bit flipping from 0 to 1. We also assume that there is a linear relationship between current, and in turn power consumption, and the Hamming distance between the reference state and the variable output of our model. This relationship is shown in Figure 2(a), which shows the average peak voltage during round one of DES versus the Hamming distance between the 32-bit right half at the start of the round and the 32-bit right half after the round. It is this concept of Hamming distance versus the amount of power generated that allows CPA to build its correlations and extract the secret information.

The basic idea behind a CPA attack is that an attacker would first identify an operation that would potentially leak secret information. They would then create a model of how that operation would function. The attacker would then guess the secret information, and create Hamming distances between a reference state and the output of their model. A linear correlation factor between Hamming distance and the power traces would then be produced. For this paper, we applied CPA to publicly available DES traces supplied by [6]. The operation we attacked was the Feistel function in round one, specifically the individual substitution boxes (S-boxes). Using the known plain text, we used the right-side values that are fed into the Feistel function as the known state and the right-side values




Figure 3: Power injector architecture that can toggle up to 2^n-1 nets each cycle by varying the power input p .

after the round as the output of our model. The correlation factor is produced using a Pearson Correlation factor as described in [1].

Since the key idea of CPA is for the attacker to build a correlation between Hamming distance and power, we designed the correlated noise generator to obfuscate that relationship by creating additional bit transitions simultaneously with the encryption. This obfuscation is demonstrated in Figure 2(b). While similar to Figure 2(a) in that it relates peak power during round one to the Hamming distance produced by the model, we added correlated noise to the power traces.

Our approach does not use random noise, which would not obscure any correlations because the attacker could simply filter the noise. Instead, the correlated noise generator uses the plain- or cipher-text as a seed to correlate the bit-switching to the encryption/decryption of the same data. This strategy makes the noise repeatable and harder to filter out. That seed is fed into a hash function that performs an exclusive OR with a random number generated at synthesis, referred to as the salt. The salted hash output is then used as the input to the power injector. At the same time, the salted hash output is fed back into the hash function to be used in the next clock cycle. This system may be reset at any time to return to its initial state.

Although we evaluate correlated noise in defending the encryption key used for DES, the technique potentially applies to other statistical power analysis attacks. Such attacks rely on the ability to model and relate potential power consumption to actual measured power. Because our method targets that relationship, it can obscure the leakage of other applications.

B. Power Injector

Whereas the correlated noise generator determines how power should be varied to obscure leakage, the power injector is responsible for generating that power. To accomplish this goal, we created a register-transfer-level power-injector core that can be integrated into any overlay to toggle any number of nets every cycle. In addition to increasing power by specified amounts, the injector can also decrease power by using a baseline that includes some amount of toggling every cycle that can be increased or decreased each cycle.

Figure 3 demonstrates the architecture of the power injector. The injector takes as input an n -bit power setting p that specifies how many nets should be toggled. The value of n is specified pre-synthesis based on the application-specific required range of power, but p can change every cycle using power values from the correlated noise generator. To implement the toggling, the architecture uses 2^n-1 flip-flops

with inverters. The enable for each flip flop is controlled by p , where bit i of p enables 2^i different flip flops. With this architecture, an overlay can change the number of toggled nets from 0 to 2^n-1 every cycle.

III. EXPERIMENTS

In this section, we present a DES case study on the effectiveness of correlated noise against statistical power analysis. Although this case study demonstrates CPA against the first round of DES, our approach could be applied to both other forms of statistical analysis and to other applications requiring protection against such an attack.

To perform CPA, we used publicly available DES power traces [6]. Each trace consists of 20,000 data points that have been pre-processed such that all rounds occur at the same time across traces. The experiments were performed using CPA with 500 traces.

The experiment was done in two stages. The first stage involved performing a CPA attack on the provided traces to act as a baseline, showing the effectiveness of the attack in general at breaking a single S-box. The second stage performed the same attack on power traces that we augmented to simulate the effects of correlated noise as described in Section II.A. The goal was to demonstrate how correlated noise can obfuscate the correlation between Hamming distance and measured voltage. Only one S-box was broken, as the concept is equivalent across all S-boxes.

To perform the attack, we developed a Matlab function that modeled a round of DES and returned the Hamming distance for the 4-bits affected by the specified S-box. We then used a wrapper function that tested all 64 possible partial subkey combinations and collected the respective Hamming distances. After we generated the Pearson coefficients correlating the Hamming distance and power traces, we observed the peak values that occurred during round one for each key guess. The highest coefficient was returned as the best guess. The peak coefficients for each guess can be seen in Figure 4(a). Notice that the correct guess, 56, has a distinctively high correlation compared to the other points.

To evaluate our countermeasure, we generated augmented traces to mimic correlated noise generation by first creating a new trace that averaged 3,000 of the original traces. Then, we divided the traces into three equally sized sets. The first set was used as a base trace value. Next, the average trace was subtracted from the i th trace in both the second and third set, creating traces that represent the additional power contributions of each respective trace. Those were combined, then added to the i th trace from the first set to create an augmented trace that had some amount of power added to it. Although the correlated noise generator generates noise differently than these augmented traces, this experiment demonstrates a proof-of-concept that adding additional power to a trace can break the correlation between power usage and Hamming distance.

Figure 4(b) demonstrates that when performing CPA on the augmented traces, our countermeasure obfuscates the correlation between Hamming distance and power consumption. The figure shows the peak correlation value for




Figure 4: A successful CPA attack on a DES S-box (a) will produce a distinctively high correlation with the correct guess. However, when the correlation between the model used in the attack and the power traces is obfuscated (b), there is no longer a distinctive correlation with any guess, let alone the correct one.

each of the 64 guesses using the augmented traces. As before, the highest correlation should occur at a guess of 56. However, the figures clearly shows two peaks: one occurring at guess 39 and another at guess 30. These peaks render the attack inconclusive, protecting the secret key and power consumption.

Although space constraints prohibit a detailed analysis of the power injector core, we provide a summary of area requirements and power-injection granularity. Because most FPGAs provide lookup tables (LUTs) with two flip-flops, and because the input to the flip-flop is just a one-input inverter, each LUT can implement two toggling flip-flops. For an n -bit power input, most FPGAs will require 2^{n-1} LUTs. The value of n will vary for different applications, but will generally not be larger than \log_2 of the maximum Hamming distance used in an attack because this value of n can often double (or cut in half) the original power. Therefore, the number of LUTs required to implement the countermeasure will generally be similar to the maximum Hamming distance, which is a minor overhead for most envisioned use cases.

IV. CONCLUSIONS

Statistical power analysis is a significant hardware security concern for FPGA applications. Whereas previous work has focused on encryption cores that protect against such analysis, our approach complements those studies by providing an overlay-based countermeasure that can protect widely used vulnerable cores and devices. Our approach integrates correlated noise generation and power injection into FPGA overlays to obscure leakage of secret information by obfuscating the relationship between Hamming distance and voltage. We demonstrated a proof of concept by performing a Correlation Power Analysis attack on publicly available DES power traces, and then demonstrated that our countermeasure prevents the attack by yielding an incorrect encryption keys. Future work includes integrating this countermeasure into existing FPGA overlays and extending the approach to protect

FPGA bitfiles by implementing the correlated noise generator in a partially reconfigurable region.

ACKNOWLEDGMENTS

We gratefully acknowledge help from Draper, Swarup Bhunia, and Kai Yang.

REFERENCES

- [1] E. Brier, C. Clavier, and F. Olivier. *Correlation Power Analysis with a Leakage Model*, pages 16–29. Springer Berlin Heidelberg, 2004.
- [2] D. Capalija and T. S. Abdelrahman. A high-performance overlay architecture for pipelined execution of data flow graphs. In *2013 23rd International Conference on Field programmable Logic and Applications*, pages 1–8. IEEE, 2013.
- [3] P. Cooke, J. Fowers, G. Brown, and G. Stitt. A tradeoff analysis of fpgas, gpus, and multicores for sliding-window applications. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 8(1):2:1–2:24, Mar 2015.
- [4] J. Coole and G. Stitt. Fast, flexible high-level synthesis from opencl using reconfiguration contexts. *IEEE Micro*, 34(1):42–53, Jan 2014.
- [5] J. Coole and G. Stitt. Adjustable-cost overlays for runtime compilation. In *Field-Programmable Custom Computing Machines (FCCM)*, pages 21–24, May 2015.
- [6] DPA Contest. <http://www.dpacontest.org/home/>.
- [7] W. Hnath. Differential power analysis side-channel attacks in cryptography. Worcester Polytechnic Institute, 2010.
- [8] N. Kamoun, L. Bossuet, and A. Ghazel. Correlated power noise generator as a low cost dpa countermeasures to secure hardware aes cipher. In *2009 3rd International Conference on Signals, Circuits and Systems (SCS)*, pages 1–6, Nov 2009.
- [9] P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *Proceedings of the International Cryptology Conference on Advances in Cryptology*, CRYPTO '99, pages 388–397, London, UK, 1999.
- [10] W. Luis, G. R. Newell, and K. Alexander. Differential power analysis countermeasures for the configuration of sram fpgas. In *Military Communications Conference, MILCOM 2015 -2015 IEEE*, pages 1276–1283, Oct 2015.
- [11] S. Skorobogatov and C. Woods. In the blink of an eye: There goes your aes key. *IACR Cryptology ePrint Archive*, 2012:296, 2012.
- [12] G. Stitt, R. Karam, K. Yang, and S. Bhunia. A unqiufied virtualization approach to hardware security. *IEEE Embedded Systems Letters*, PP(99):1–1, 2017.