Power Analysis Attacks on the Customizable MK-3 Authenticated Encryption Algorithm

Peter Fabinski
Department of Computer Engineering
Rochester Institute of Technology
pnf9945@rit.edu

Steve Farris
L3Harris Technologies
Steve.Farris@L3Harris.com

Michael Kurdziel
L3Harris Technologies
Mike.Kurdziel@L3Harris.com

Marcin Łukowiak
Department of Computer Engineering
Rochester Institute of Technology
mxleec@rit.edu

Stanisław Radziszowski
Department of Computer Science
Rochester Institute of Technology
spr@cs.rit.edu

***Abstract*—MK-3 is an authenticated encryption scheme based on the duplex sponge construction, suitable for both hardware and software. It provides broad factory and field customization features. The same security claims are valid for the original and all recommended customizations. Extensive security analyses of MK-3 were performed in our previous work: differential, linear, cube, and brute force attacks, as well as statistical analysis. In this work we report on new experiments involving Correlation Power Analysis (CPA), which is considered one of the most powerful side-channel attack (SCA) techniques.**

Two CPA attacks on MK-3 were developed: the first directly after the key absorption, and the second after the S-boxes in the first round of IV absorption. In the first attack, under strong assumptions about an attacker's capability to collect traces, we can recover 128 of the 512 state bits in a physical test on an FPGA. The second attack builds on top of the first one, but it assumes that special registers have been embedded after the S-boxes. Even under such ideal conditions, this attack can potentially reduce the brute-forcing difficulty only by an additional 88 to 194 bits. Overall, this gives the CPA attack no advantage over brute-forcing for the original 128-bit key. The previous and current results ensure that MK-3 and its customized versions effectively conceal its plaintext input.

***Index Terms*—FPGA, side-channel attacks, correlation power analysis, customizable encryption, sponge construction**

I. Introduction

One of the most common side-channel attacks, power analysis attacks, can pose a real threat to any system implementing cryptographic algorithms or other operations involving sensitive information. Power analysis is based on the fact that performing operations in an electronic device requires energy, and that the amount of energy is related to the operations and data in use. By observing power consumption externally, an attacker can potentially deduce information about the internal state of the device, often strongly correlated to secret keys. This internal information is often considered inaccessible and left as an implementation detail, meaning that an attacker able to gain access to this information could exploit a design which is secure when considered as a whole.

MK-3 is an authenticated encryption scheme based on the duplex sponge construction. It is suitable for both hardware and software, but its design features are targeted specially for hardware implementations. The MK-3 scheme is a proprietary algorithm of L3Harris Technologies, formerly also proprietary of its original developer, Harris Corporation. MK-3 provides broad customization features, in the form of both factory customization and further field customization. This allows many different organizations to use non-interoperable variations of the same design [1]–[3]. The same security claims are valid for the original and all factory customizations, and for further algorithm customizations which can be easily adopted by the users. Extensive security analyses of the MK-3 scheme were performed in our previous work: classic cryptographic analysis including differential and linear attacks, cube attacks, and brute force attacks, as well as statistical analysis of bit positions and ciphertext statistical analysis [1]–[3]. However, until now there have been no detailed studies of MK-3's resistance to Side-Channel Attacks (SCA).

The MK-3 scheme is based on the sponge construction, a cryptographic structure introduced in Keccak, the winner of the NIST SHA-3 competition [4], [5]. Given that this construction is still somewhat new, it has a relatively small body of work on power analysis when compared to other schemes such as the Advanced Encryption Standard (AES). The goal of this work is to explore the application of power analysis attacks to the MK-3 algorithm, and to give a set of security recommendations for its implementation [6]. Section II provides an overview of the MK-3 scheme and its customizations, followed by a description of Correlation Power Analysis (CPA), the attack method used in this work. Section III describes the design and implementation of the experimental setup, as well as the analysis process and the reasoning behind the two-step approach taken for the attacks. The results of these attacks are presented and discussed in Section IV, with Section V providing a summary of the design recommendations and concluding remarks.

II. Background

A. MK-3 Algorithm

The MK-3 scheme uses the duplex sponge construction, which allows a single pass to provide both authenticity, the ability to verify that information has not been changed in transit, and confidentiality, the ability to prevent such information from being observed by anyone not in possession of the key. Figure 1 shows the duplex sponge construction used in MK-3 [1], [2].

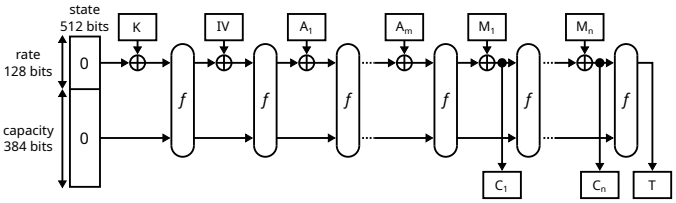

Fig. 1: Duplex sponge construction for MK-3 (128-bit key).

In each iteration of the sponge, one of three operations is performed. Absorption is when data is XORed into the first 128 bits of the 512-bit state, called the rate. This is used for the key K, IV (initialization vector), and additional authenticated data A_i. Squeezing is when the rate is taken as output with no modification, as with the authentication tag T. Duplexing combines the first two operations, XORing the input message block M_i with the rate and taking the result as the output ciphertext C_i. In the duplexing mode, the sponge acts similarly to a stream cipher, while also incorporating the inputs M_i into the state to create the final authenticating tag T.

The main component specific to MK-3 is the bijective function f applied in-between steps of the sponge. It is composed of 10 (for a 128-bit key) or 16 (for a 256-bit key) consecutive applications of the round function g. The round function itself has four stages; substitution, permutation, mixing, and addition of a round constant, as shown in Figure 2.

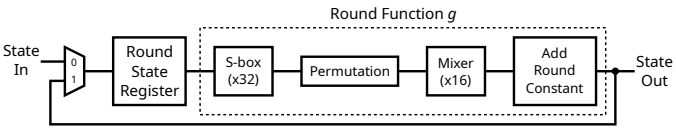

Fig. 2: Block diagram of the MK-3 bijective function f.

The substitution stage is composed of 32 identical 16-bit S-boxes. The permutation stage remaps S-box output bits to mixer input bits, with the requirement that each output of a given S-box serves as an input to a different mixer. The mixing stage is made up of 16 mixers, each taking a 32-bit input and producing a 32-bit output through several XOR operations. Some of these operations are defined by the irreducible polynomial chosen for that mixer as part of customization. Finally, the add-round-constant stage performs an XOR with a 512-bit value, which is predetermined based on the round number. A state register is used to store each output and pass it to the next round until the required round count is reached.

B. MK-3 Customization

Customization of an encryption scheme allows several different organizations to each use their own independent, non-interoperable versions of the scheme without redefining the scheme as a whole. MK-3 is highly customizable, providing opportunities to modify its functionality in all stages of the round function, with some customizations able to be applied dynamically in the field [3].

In this work, only the customization of the irreducible polynomials used in the mixing stage is considered. As will be seen in Section III, the specific dependencies of a mixer's output bits on its input bits are essential for the second power analysis attack, and these dependencies are directly influenced by the choice of mixer polynomial. The other customizations do not have any significant effect on the analysis, only changing constants or modifying which particular bits are targeted rather than influencing the feasibility of the attacks as a whole.

C. Correlation Power Analysis

Correlation Power Analysis (CPA) and Differential Power Analysis are the two most known methods for extracting internal information about the operation of a device based on fluctuations in its power consumption [7]–[10].

In CPA, a model is developed for the circuit's expected power consumption at a certain point in its execution, given a known input and a guess for part of the unknown target value. After modeling, each of the known inputs is applied to the target circuit, and its power consumption with each is sampled. Finally, the Pearson correlation is computed between the modeled and measured power consumption, for each guess and at each sampling time. Out of the resulting matrix, the guess having the highest-magnitude individual correlation coefficient is taken to be the correct value for the target. Figure 3 shows a visual representation of this process.

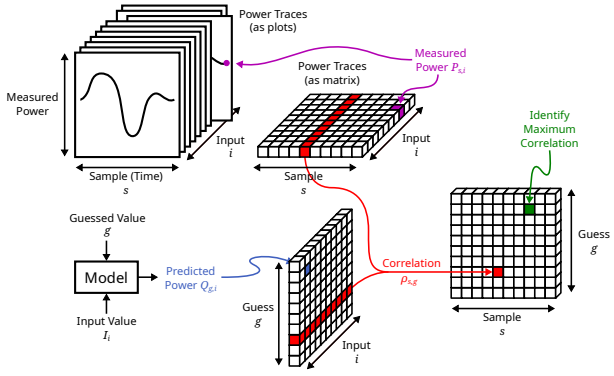

Fig. 3: Diagram of CPA computations.

III. Methodology

A. MK-3 Implementation and Measurements

Our attack scenario targets the start of IV absorption. The absorption of the key itself is not a viable target, as CPA requires a known, changing input in addition to the constant value being guessed. All steps after the IV do not influence

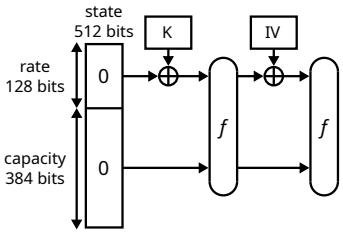

Fig. 4: Abbreviated sponge construction for power analysis.

the attack, so for the sake of measurement speed, they are removed. This leaves the final configuration shown in Figure 4.

In order to execute the algorithm and take measurements, the ChipWhisperer platform [11] was used. It provides both a Field-Programmable Gate Array (FPGA)-based target board and a specialized data acquisition system with software tooling for measurement automation. Based on previous work implementing MK-3 [12]–[15], the reduced configuration was integrated into the ChipWhisperer system. Several variants of the round function were also implemented in order to test their effects on the attack. Figure 5 shows a register transfer logic (RTL) diagram for one such variant including an additional register after the substitution stage. The S-box design is logic-heavy, so this register may be desirable to improve timing. Other variations included resetting this register in between each round and removing it outright. Additional details on variants are presented in [6].

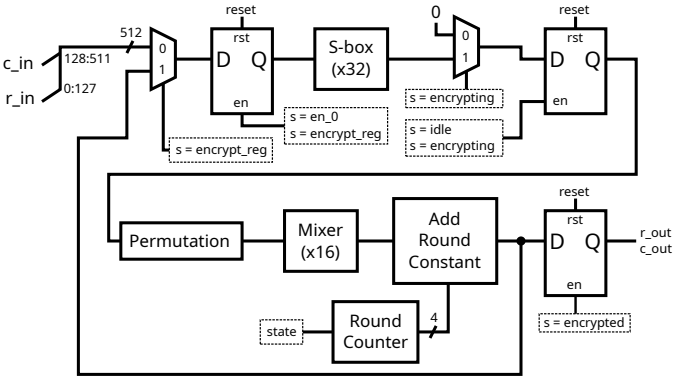

Fig. 5: 10-round MK-3 RTL with register after S-box.

When collecting power measurements, the key is fixed, while the IV is randomized. Each measurement records the power consumption of the FPGA during computations with one IV, producing a set of samples called a power trace. As an example, Figure 6 shows a plot of 500 power traces collected from the design shown in Figure 5.

In this plot, the X-axis is the sample number, representing time. Each sample covers approximately 11 ns. The Y-axis is the relative supply voltage at the FPGA; because the measurements are AC-coupled, there is no calibrated scale for this axis, but such information is not needed for CPA. These measurements show the first three rounds of IV absorption. The S-boxes are the most logic-heavy and power-consuming

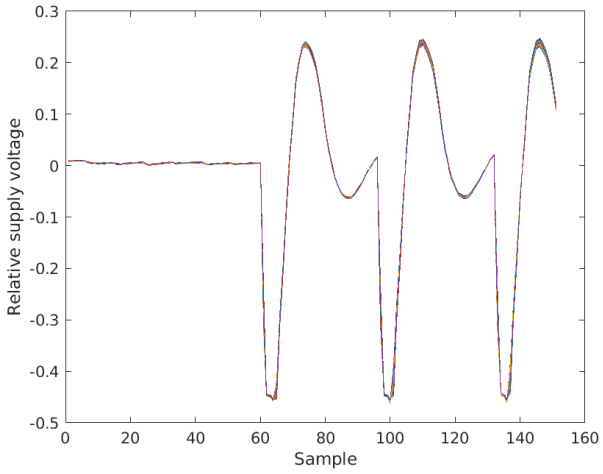

Fig. 6: Plot of 500 power traces collected from 10-round registered design.

parts of the design, and each downward spike indicates a point where their inputs changed at the beginning of each round.

B. Rate Attack

The first attack uses the rate after key absorption as the guessed location, with the IV taken directly as the known input. Based on the rate and IV, a model was created predicting the power consumption corresponding to the Hamming weight of the computed S-box outputs.

For this attack, the rate is guessed in 16-bit segments called words. Each word corresponds to the inputs and outputs of a single S-box. In this case, the division into 8 words reduces the total number of guesses to find the rate from 2^{128} with brute-forcing to $(8 \cdot 2^{16}) = 2^{19}$ with CPA.

C. Capacity Attack

For the remaining 384 bits of the state, called the capacity, this first attack method does not apply. In the sponge construction, the IV is only XORed with the rate portion of the state, and it does not interact with the capacity during the first round of IV absorption. Therefore, the first-round S-box outputs corresponding to the capacity are unaffected by the IV. This prevents CPA from working because there will only be one modeled power consumption for each guess, while CPA relies on having many different cases for each guess to compute its correlations. In order to resolve this problem, we need to move the power modeling point further into the round function, where our desired capacity bits combine with the changing bits of the rate in the mixers. Specifically, the modeling point is moved to the register at the beginning of the next round; this is chosen because the Hamming weight power model is most effective on registers.

Unfortunately, in addition to making the capacity attack dependent on the success of the rate attack, moving the modeling point past the mixers makes power modeling significantly more complicated. Instead of working with one S-box in isolation, the permutation and mixing stages force

computations to use outputs from multiple S-boxes. As defined in [1], the permutation stage ensures that each of a mixer's 32 input bits comes from the output of a different S-box. Additionally, due to the mixer structure, each individual mixer output bit is the XOR of 2, 3, or 5 of the mixer's inputs. Every output of an S-box depends on all 16 of its inputs, so performing CPA on even a single mixer output bit will require at least 2^{16} guesses at the S-box input. A single bit is not an effective target for CPA, but in order to use even two mixer output bits, two S-boxes would need to be guessed at once. This would mean applying the model 2^{32} times for each IV, as well as computing 2^{32} correlations. These computational requirements scale very quickly with the multiple bits needed to obtain a result in a reasonable number of traces, and are considered infeasible in the scope of this work.

In order to circumvent these computations, the developed attack targets S-box outputs rather than their inputs. This means that it no longer directly provides bits of the state; however, it still allows cases to be eliminated, reducing the search space for a brute-force attack on the remaining bits.

Even with this modification, it is not possible to fully determine the capacity. As before, the requirement for an output bit to be modelable with CPA is that it depends on both a bit being guessed and a changing known input, and that its value can be calculated given a guess and a known input. In this situation, a modelable bit must depend on at least one bit in the rate, at least one bit in the guessed bits of the capacity, and none of the unknown capacity bits not being guessed. Based on the design of the mixing and permutation stages, only some output bits will meet this requirement.

In this paper, only a proof-of-concept attack was performed in hardware. Upper bounds were also calculated for the number of capacity bits which can potentially be recovered for each mixer customization. These upper bounds are determined based on the total number of bits which could be guessed if all capacity bits could be used at once as valid modeling inputs. Details are given in [6].

IV. Results

A. Rate Attack

For the rate attack, the results are presented in correlation plots as shown in Figure 7, where each subplot corresponds to the attempts to obtain one word of the rate. Within a single plot, the X-axis is the sample number and the Y-axis is the absolute correlation between the real and modeled power. Each line represents a single guess, so for the rate attack, each plot contains 2^{16} lines, though only 200 are rendered here. The line for the correct value is green, while the red line is the guess with the highest peak correlation. Cases with no red line indicate that the attack worked correctly, i.e., that the guess with the peak correlation was indeed the correct value.

In this attempt, performed on the registered MK-3 variant using 5,000 power traces, 3 out of 8 words were correctly identified. Additionally, when compared with Figure 6, we can see that as expected, the peaks of the correct guesses all lie at approximately the point where the first round of IV absorption occurs, which was the point selected for the power model.

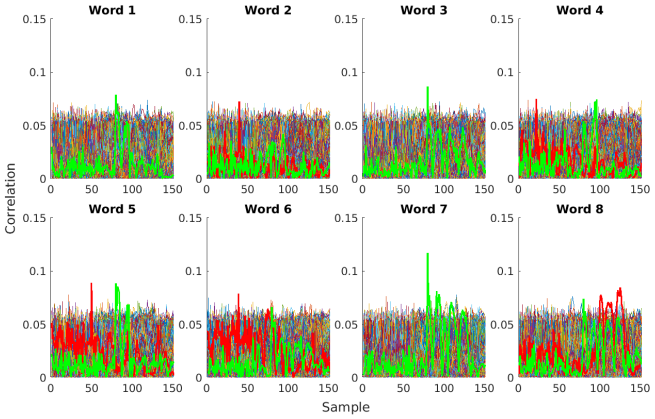

Fig. 7: Analysis results for 10-round registered variant using 5,000 traces.

By increasing the number of power traces to 20,000, the correct value was obtained for all eight words, as shown in Figure 8. Note that between these two figures, the Y-axis scale is fixed. In most cases, the correlation of the correct guess did not increase significantly when compared with the earlier attempt. Rather, the baseline correlation of the incorrect guesses decreased due to additional points in the correlation, revealing the previously-indistinguishable correct guess.

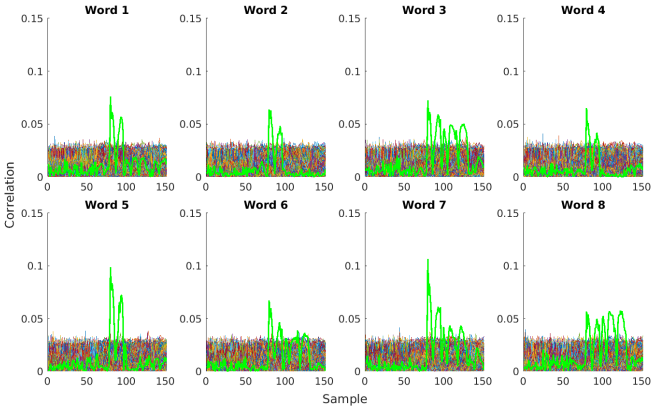

Fig. 8: Analysis results for 10-round registered variant using 20,000 traces.

In order to evaluate the effects of trace count on the performance of the rate attack, it was performed repeatedly, changing the number of traces included in each analysis. For each data point, ten random selections of N traces were made from a large pool. The analysis was performed on each of these trace sets, recording the number of correctly-recovered words and the number of words whose correct guess was within the top 10 ranked by correlation. The overall experiment was also performed with two different keys. Figure 9 shows the results of this testing for the registered MK-3 implementation shown

in Figure 5. From this plot, we can observe that at approxi-

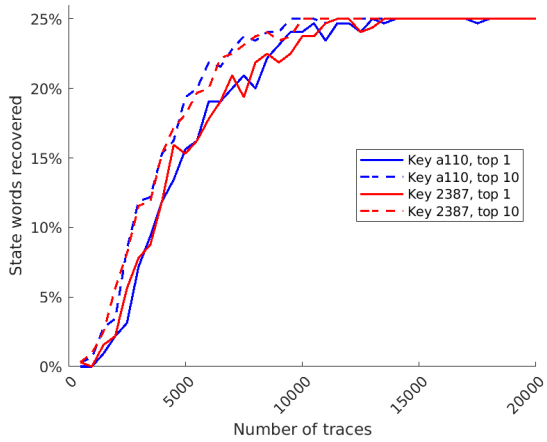

Fig. 9: Success rate results for 10-round registered variant at various trace counts.

mately 15,000 traces, all 8 words were correctly identified for both keys in all 10 iterations. This only corresponds to a value of 25% because the rate is only 1/4 of the total state bits, and this attack does not target the capacity.

The same analysis was then performed on an implementation of MK-3 with no register in between the S-boxes and the permutation stage. From the results in Figure 10, we can see

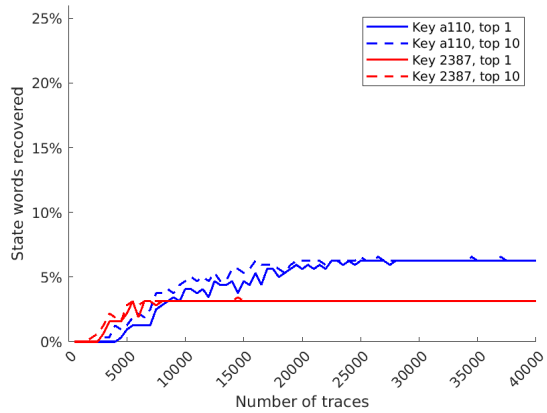

Fig. 10: Success rate results for 10-round unregistered variant at various trace counts.

that even with twice as many traces, the success rate plateaus at one or two words depending on the key. This is because, with the register removed, there is no longer a single point in time where the model correctly predicts the power usage, spreading out the results over several samples. The register removal also shifts the focus onto the mixing stage, where the key has a stronger influence on power consumption.

B. Capacity Attack

The results for the capacity are composed of both the proof-of-concept attack in hardware and the theoretical upper-bound calculations based on customization. Figure 11 shows the results of the proof-of-concept attack using a plot with a similar format to those in Figures 7 and 8. As before, the correct result

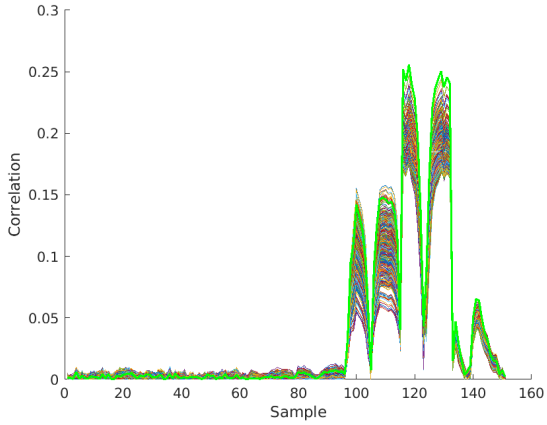

Fig. 11: Capacity proof-of-concept results for 10-round registered variant with reset using 150,000 traces.

is highlighted in green. Due to software limitations, only those lines with the top 200 peaks are shown. The correct value does indeed have the highest peak correlation; however, due to a bitwise inversion case described in [6], one other guess exists with an identical correlation trace. There is still a visible gap before the next trace, so an attacker would be able to narrow down the possibilities to only these top two guesses.

For the upper-bound calculations, each result depends both on the polynomial selected for a mixer and the position of that mixer in the design. Due to the structure of the permutation stage, mixer locations are not interchangeable. For each of the 16 mixers, there are 4080 independently-selectable irreducible polynomials which could be used for customization. If we assume all mixers are configured identically, Figure 12 shows a histogram of the upper bound on the number of recoverable capacity bits. These values range from 102 to 194 bits. In the worst case, the 194 potentially recoverable capacity bits add to the 128 bits from the rate attack, giving a total of 322.

Because the mixers are independent of one another, we

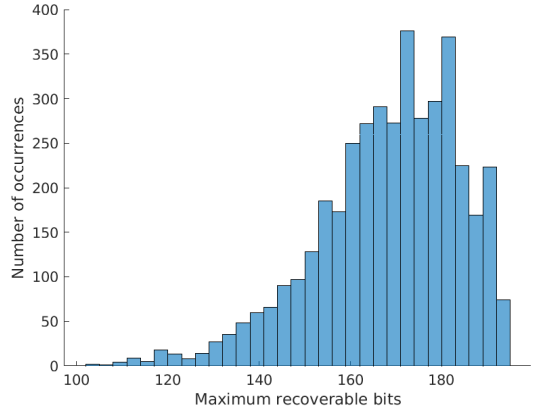

Fig. 12: Histogram of maximum recoverable capacity bits across all irreducible mixer polynomials.

TABLE I: Mixer configuration producing the lowest possible number of recoverable capacity bits.

Mixer Number	Polynomial (hex)	Recoverable Bits
0	0x04c1	5
1	0x0a81	5
2	0x0a03	5
3	0x0807	5
4	0x040b	5
5	0x0807	5
6	0x002b	5
7	0x002b	5
8	0x002b	5
9	0x002b	5
10	0x002b	5
11	0x3801	6
12	0x5003	6
13	0x200d	8
14	0x8013	7
15	0x002b	6

can also evaluate each one separately to find an optimal configuration. Table I shows the mixer configuration with the smallest possible number of recoverable capacity bits. In this case, configuring mixers separately reduces the minimum to 88. However, this change does not have any effect on the maximum number of recoverable bits, which remains at 194.

V. Conclusions

Our prior work included extensive security analyses of the MK-3 scheme and its customized variants, which implied MK-3's resistance to classical cryptographic attacks. In this work we reported on new experiments with attacks on the MK-3 implementation using Correlation Power Analysis, which is considered one of the most powerful general side-channel attack techniques.

Two different CPA attacks on MK-3 were developed targeting different locations in the encryption: the first right after the key absorption, and the second during the absorption of the IV, after the S-boxes of the first round. In the first attack, under strong assumptions about an attacker's capability to collect traces, we can recover 128 of the 512 state bits in a physical test on an FPGA. This is far from sufficient to attempt the inversion of the bijection f and thus recover the key. The second attack builds on top of the first one, but it assumes that additional registers have been embedded after the S-boxes in order to make the attack feasible. Even under such ideal conditions, this attack can potentially reduce the brute-forcing difficulty only by an additional 88 to 194 bits, giving the CPA attack no advantage over brute-forcing for the original 128-bit key. Overall, the power analysis attacks used in this work, despite being successful in that they confirm the technique of modeling with statistical analysis of power traces, do not present any significant threat to the security of the MK-3 scheme.

Side-channel attacks typically require physical access in order to take useful measurements. If these can be prevented by means of physical security, such as wiping keys if a device is opened, then methods such as power analysis simply cannot be applied. Another possibility is to limit the number of encryptions which can be performed with a given key, via a technique such as key rolling, which could prevent collection of a sufficient number of traces for SCA.

The previous and current results together ensure that the MK-3 encryption algorithm and its customized versions effectively conceal its plaintext input.

References

[1] M. Kelly, "Design and cryptanalysis of a customizable authenticated encryption algorithm," Master's thesis, Rochester Institute of Technology, August 2014.

[2] M. Kelly, A. Kaminsky, M. Kurdziel, M. Łukowiak, and S. Radziszowski, "Customizable sponge-based authenticated encryption using 16-bit S-boxes," in *MILCOM 2015 - 2015 IEEE Military Communications Conference*, 2015, pp. 43–48.

[3] P. Bajorski, A. Kaminsky, M. Kurdziel, M. Łukowiak, and S. Radziszowski, "Customization modes for the Harris MK-3 authenticated encryption algorithm," in *MILCOM 2018 - 2018 IEEE Military Communications Conference (MILCOM)*, 2018, pp. 1–5.

[4] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche, "Cryptographic sponge functions," January 2011.

[5] ——, "The Keccak SHA-3 submission, version 3," January 2011.

[6] P. Fabinski, "Side-channel attacks and countermeasures for the MK-3 authenticated encryption scheme," Master's thesis, Rochester Institute of Technology, December 2022.

[7] E. Brier, C. Clavier, and F. Olivier, "Correlation power analysis with a leakage model," in *Cryptographic Hardware and Embedded Systems - CHES 2004*, M. Joye and J.-J. Quisquater, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 16–29.

[8] M. Tehranipoor and C. Wang, *Introduction to Hardware Security and Trust*. New York, NY: Springer New York, 2012.

[9] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Advances in Cryptology — CRYPTO' 99*, M. Wiener, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 388–397.

[10] P. Kocher, J. Jaffe, B. Jun, and P. Rohatgi, "Introduction to differential power analysis," *Journal of Cryptographic Engineering*, vol. 1, no. 1, pp. 5–27, April 2011.

[11] C. O'Flynn and Z. D. Chen, "ChipWhisperer: An open-source platform for hardware embedded security research," in *Constructive Side-Channel Analysis and Secure Design*, E. Prouff, Ed. Cham: Springer International Publishing, 2014, pp. 243–260.

[12] C. A. Wood, "Large substitution boxes with efficient combinational implementations," Master's thesis, Rochester Institute of Technology, Aug 2013.

[13] G. Werner, S. Farris, A. Kaminsky, M. Kurdziel, M. Łukowiak, and S. Radziszowski, "Implementing authenticated encryption algorithm MK-3 on FPGA," in *MILCOM 2016 - 2016 IEEE Military Communications Conference*, 2016, pp. 1225–1230.

[14] D. F. Stafford, "Evaluating performance and efficiency of a 16-bit substitution box on an FPGA," Master's thesis, Rochester Institute of Technology, June 2021.

[15] ——, "Correlation power analysis of MK-3 and countermeasures," Project report, Rochester Institute of Technology, July 2017.