An Open-Source Platform for Evaluating Side-Channel Countermeasures in Hardware Implementations of Lightweight Authenticated Ciphers

Abubakr Abdulgadir[1], William Diehl[2] and Jens-Peter Kaps[1]

[1] Cryptographic Engineering Research Group
George Mason University, Fairfax, Virginia 22030, USA
email: {aabdulga, jkaps}@gmu.edu
[2] Signatures Analysis Lab
Virginia Tech, Blacksburg, Virginia 24061, USA
email: wdiehl@vt.edu

Abstract. Lightweight implementations of cryptographic algorithms must be evaluated in terms of security, cost, and performance before their deployment in most practical applications. The availability of open-source platforms for such evaluation saves researchers' time and increase reproducibility of results. In this work, we improve upon the previous version of the Flexible Opensource workBench fOr Side-channel analysis (FOBOS) to introduce "FOBOS 2," and utilize it to perform such evaluation tasks for hardware implementations of authenticated ciphers, with special focus on candidates submitted to the NIST Lightweight Cryptography standardization process. We perform power measurements on Artix7 FPGA, and countermeasure evaluation of lightweight hardware implementations of selected NIST Lightweight Cryptography Round-2 candidates and the current NIST standard AES-GCM on the Spartan6 and Artix7 FPGAs. Our results show that Ascon consumes the least power at 50 MHz, and has the lowest change in dynamic power per increase in frequency, while GIFT-COFB consumes the least energy-per-bit. We also show that side-channel countermeasures applied to implementations of Ascon and AES-GCM are effective in both Spartan6 and Artix7 using leakage detection tests.

Keywords: Lightweight cryptography · Side-Channel Analysis · Countermeasure Evaluation · FPGA

1 Introduction

Lightweight cryptography (LWC) is by definition deployed in resource constrained devices like smart-cards, RFID tags and remote-sensor nodes. Among the most critical specifications of LWC applications are power consumption and energy per bit (E/bit) since they determine power supply specifications and battery life.

Also, adversaries can more easily gain physical access to such systems and measure side-channels such as power consumption and Electro-Magnetic emanations (EM). For example, it is easier to obtain physical access to a remote sensor node than to a server that resides in a physically secured data center. This makes side-channel analysis [13] (SCA) especially concerning for lightweight applications.

Power SCA has two major variants. Simple Power Analysis (SPA) uses one or a few traces to recover the key, where amplitude of an observed signal corresponds to a key fragment. On the other hand, Differential Power Analysis (DPA) needs more traces but is very powerful in extracting secret information even if the collected power traces are noisy [13]. Applying countermeasures against DPA becomes a necessity for deploying practical and secure systems. Otherwise, system security can be easily bypassed regardless of the mathematical strength of the cryptography in use.

Many countermeasures have been proposed to protect against SCA at all levels of abstraction. For example, at the protocol level, the number of cryptographic operations performed using a single key can be limited to a predefined maximum depending on the targeted security level. At the algorithmic level, masking [25] and threshold implementations [16] can be used. Examples of applying masking and threshold

implementations are shown in [8] and [6]. Dual rail technology can be used at the logic level among many other countermeasures for ASICs [24] and FPGAs [27, 28].

Evaluating SCA resistance is necessary for test labs and countermeasure designers. A systematic methodology should be used to confirm that leakage is reduced, and that the design meets minimum security requirements. One of the most widely used methodologies is Test Vector Leakage Assessment (TVLA) [10, 23] which applies statistical tests to measure the significance of leakage. Such methodologies and tools will be valuable for efforts like the NIST LWC project that aims to standardize algorithms for resource-constrained devices.

The availability of open-source hardware and software to perform security evaluation saves researchers' time and enables the reproduction of results across research teams. Several solutions to perform SCA are already available for academia and industry. The DPA Workstation from Rambus [20] and Inspector from Riscure [22] are examples of commercial systems, however, they are out of the cost range of many academic and most low-end users. SAKURA boards [11] are also widely used in academia and support FPGAs and smart cards, however, they do not include integrated acquisition and analysis tools. NewAE Chipwhisperer is a platform that has many Design Under Test (DUT) options and synchronous capture (i.e., DUT clock can be synchronized with the sampling clock) which allows using a low sampling frequency, yet getting results comparable to asynchronous capture at higher sampling rates [18].

The Flexible Opensource workBench fOr Side-channel analysis (FOBOS) is a comprehensive SCA platform that uses commercially available low-cost FPGA boards (e.g. Digilent Nexys-A7) whenever possible. Many academic environments already have similar boards used for teaching purposes which further reduces the system's cost.

Since FOBOS is directly compatible with CAESAR (Competition for Authenticated Encryption, Security, Applicability and Robustness) Hardware API [12] and expected to be directly compatible with the upcoming Lightweight Cryptography API, no time is needed to adapt cipher implementations to a new interface. Given the number of candidates in the NIST LWC project, time savings will be a significant factor in evaluating these ciphers. While Chipwhisperer is compatible with state-of-the-art target boards, work is needed to adapt NIST LWC ciphers interface to its interface. Therefore, using FOBOS will save time in the evaluation of NIST LWC candidates.

In this work, we improve the architecture of the FOBOS framework and upgrade FOBOS for compatibility with state-of-the-art Xilinx 7-series FPGAs resulting in the new FOBOS 2. We use FOBOS 2 to measure power and compute E/bit for Round-2 NIST LWC candidates Spoc, Spook, GIFT-COFB and Ascon, and compare them to the current standard, AES-GCM, as a benchmark. We also evaluate SCA countermeasures on protected implementations of Ascon and AES-GCM in the Spartan6 and Artix7. As a result, we claim the following contributions:

- An upgraded test platform capable of power measurement and SCA resistance evaluation that supports state-of-the-art, low-cost, commercially available FPGA boards.
- The first power measurements and energy computations of NIST LWC hardware implementations by 3rd party testers on actual advanced hardware.
- The first verification of SCA countermeasures of NIST LWC candidates in the Artix7 FPGA.

2 Background

2.1 Correlation Power Analysis

Cryptographic implementations leak information through side-channels (e.g. power consumption). This means, the power consumption of the implementation is correlated to the intermediate values processed in the implementation. One SCA variant is Correlation Power Analysis (CPA) [3] where an intermediate value that depends on secret data and known data is chosen. The attacker makes guesses on parts of the key and calculates the intermediate values for all key guesses. A power model is then applied to convert the intermediate values to a value roughly proportional to the power consumption. The most-used power models are based on Hamming weight and Hamming distance. Once an attacker has the hypothetical power for

each key guess, he/she calculates the correlation between data generated by each key guess and the actual measured power. The key with the highest correlation is assumed to be the correct key [14].

2.2 Test Vector Leakage Assessment

Welch's t-test is a moments-based statistical test used in a wide range of scientific research to show if two populations are significantly different. This test is used in the Test Vector Leakage Assessment (TVLA) methodology [10, 23] which has been used in many publications to test if there is significant information leakage from an implementation. To evaluate SCA leakage from a cryptographic implementation, testers can perform a DPA attack on various intermediate values or attack points. However, calculating hypothetical power for each of these attack points is time consuming and requires expert knowledge about the implementation. Alternatively, TVLA may be used to quickly assess the significance of leakage. If an implementation is secure against DPA, its power consumption must be independent of the algorithm's intermediate values. This implies that power traces collected when processing fixed data and traces collected when processing random data should be statistically indistinguishable. We call the two trace sets Q_f and Q_r respectively. A t value is calculated as follows:

$$t = \frac{\mu_f - \mu_r}{\sqrt{\frac{s_f^2}{n_f} + \frac{s_r^2}{n_r}}}$$

Where μ_f and μ_r are the means, s_f and s_r are the standard deviations and n_f and n_r are the number of samples in the sets. The null hypothesis is, that the means of the two trace sets Q_f and Q_r are equal (i.e., the two trace sets are indistinguishable). We use the calculated t value as an indicator to reject the null hypothesis at certain confidence level. If $\mid t \mid > 4.5$ we reject the null hypothesis at a confidence level of 99.999% (i.e., $p < 10^{-5}$). This means the two sets Q_f and Q_r are distinguishable and the cryptographic core is likely leaking information. However, this doesn't prove that the leakage is exploitable, and doesn't recover any secret information [23].

2.3 Frequency-based Leakage Detection

While moments-based leakage detection, e.g., computations on means and variances, can be used, frequency-based leakage detection can also be employed. An example of frequency-based leakage detection is the χ^2-test [15], which is based on frequency of occurrence. Frequencies of occurrences between "classes" are evaluated to χ values, and summed to get χ (normalized expected frequency of occurrence) and ν (degrees of freedom). Classes could be "fixed-vs-random" data D, "random-vs-random" D, etc. A probability p is computed to determine whether "classes are distinguishable."

The first step of the χ^2-test is to construct a contingency table, which is a two-dimensional table with r rows of classes, and c columns of discrete possible test results. In our evaluation of power analysis vulnerabilities in terms of measurements of voltage or current using an oscilloscope with an n-bit analog-to-digital (ADC) converter, there are 2^n possible discrete results, and thus 2^n columns. For analysis of two distribution classes, we can consider a fixed-vs-random test methodology, where Class 1 consists of "fixed D," and Class 2 consists of "random D." The entries $f_{r,c}$ in the respective cells of the table are the frequencies of occurrence of discrete value c occurring in Class r. Note that a separate contingency table is maintained for each "sample," or element in the time-domain. Correspondingly, the frequency of occurrences of discrete analog measurements from each of m trials, or "traces," are summed to their respective cells in each contingency table. The number of degrees of freedom $v = (r-1)(c-1)$. Next, expected values $E_{i,j}$ are computed as

$$E_{i,j} = \frac{(\sum_{k=0}^{c-1} F_{i,k}) \cdot (\sum_{k=0}^{r-1} F_{i,k})}{N},$$

where $F_{i,k}$ is the frequency in cell (i,k), and N is the total number of entries ($2n$ in this figure). In the next step, χ values are computed for each contingency table using the formula

$$\chi = \sum_{i=0}^{r-1} \sum_{j=0}^{c-1} \frac{(F_{i,j} - E_{i,j})^2}{E_{i,j}}.$$

Note that there is a unique degree of freedom ν for each contingency table corresponding to each sample, since the number of observed discrete analog values at every time instant could be different, resulting in a variable number of columns. This requires independent calculation of ν at every sample, which adds to computational complexity.

Finally, given a χ and corresponding ν for every sample, p is calculated as $p = \int_x^\infty f(x, v)dx$. The χ^2-test is interpreted as "passing" for every instance in time where $p > 10^{-5}$, and "failing" when $p < 10^{-5}$. In the t-test, a value of t such that $|t| > 4.5$ corresponds to $p < 10^{-5}$. However, there is no easily derivable equivalent of t in a χ^2-test.

3 Methodology

3.1 FOBOS

FOBOS is a free and open-source tool which provides a single "acquisition to analysis" platform to measure resistance to power analysis side-channel attack. The system was described in [26] and demonstrated at [1]. FOBOS consists of two major components, the data acquisition module and the analysis module. The data acquisition module is used to acquire power traces from the Device Under Test (DUT) and the analysis module is used to process the traces, run attacks and assess SCA leakage.

The ongoing NIST LWC standardization process has 32 Round-2 candidates, and a presumably large number of future later-round candidates. NIST LWC candidates are evaluated partially based on performance (including power) and cost (including energy) [17]. To compare this large number of algorithms, in terms of power, E/bit and SCA resistance, one needs an efficient platform with flexible interfaces that is compatible with the hardware API in use. Academic efforts benefit from low-cost systems that can be assembled using commercially available components, which at the same time promotes result reproducibility. The previous version of FOBOS was limited in speed because of its PC $\leftrightarrow$ control board communication protocol and lack of support for fast oscilloscopes. It also relied on Digilent Nexys-3 boards with embedded Spartan6 FPGA which have been discontinued by the manufacturer.

To address these issues and to have a platform that is suitable for efficient evaluation of a large number of candidates, we have developed an upgraded system with similar architecture, but which runs much more efficiently and uses modern hardware. Our upgraded system is available for download at [4]. Specifically we have performed the following upgrades:

- The test-vector transmission speed is improved by using a UART-USB connection.
- Power trace acquisition speed improved by supporting modern USB3-based oscilloscope (Picoscope).
- Support for NewAE CW-305 Artix7-based DUT.
- New control-board based on Digilent Basys3 has been developed. Using hardware-software codesign based on Microblaze, future upgrades to control firmware can be made primarily in software.
- New analysis scripts have been added such as the χ^2-test script.

Below, we describe the upgraded system in detail.

3.1.1 Data Acquisition module

The data acquisition module uses commercially available boards whenever possible. This lowers cost and makes usage easy for users already familiar with these widely used boards. To reduce trace noise and for modularity, we use separate boards for the controller and the DUT. Fig. 1a shows the major components of the FOBOS capture system. Fig. 1b shows an example setup of the system. In this figure, control board appears on the right, the oscilloscope used is Picoscoe 5000 (top), and the NewAE CW-305 (a low-noise Artix7-based SCA board) was used as DUT (left). The data acquisition module consists of the following components:

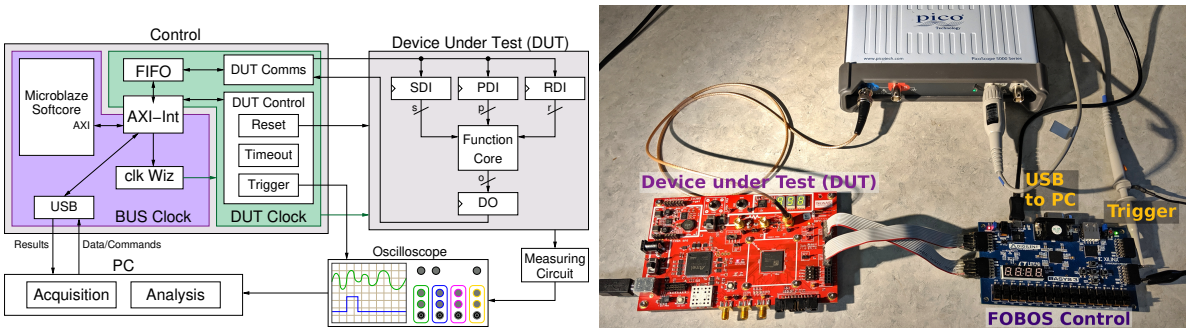

(a) Architecture (b) Typical Setup

Figure 1: FOBOS 2

- **Control PC** The user interacts with the control PC which runs scripts to generate test vectors, communicate with the control board and retrieve traces from the oscilloscope. All scripts are written in Python which provides portability across all major operating systems and good scientific computing libraries (e.g. NumPy for matrix manipulation). Traces are collected from the oscilloscope and stored in the control PC for analysis.

- **Control board** The control board is responsible for communication with the control PC and the DUT, and triggers the oscilloscope to capture power traces. FOBOS 2 supports Digilent Basys3 and Nexys-A7 control boards. Below, we describe the features of the control board.

  Communication The control board handles communication with the control PC. It is connected to the PC using USB-UART. To process a test vector, the PC sends the vector to the control board via USB-UART. The control board stores the vector briefly before sending it to the DUT. A simple protocol is used for PC-control board communication. The protocol provides headers to read/write a specific configuration parameter (e.g. trigger mode) and instructs the control board to execute, i.e., encrypt using the DUT. The control board also handles communication with the DUT. The signals used to interface with the DUT are a subset of the AXI stream protocol [2] where the control board acts as a master when sending the test vector to the DUT and a slave when receiving the result from the DUT.

  Triggering The control board is also responsible for generating the trigger signal which tells the oscilloscope when to start capturing the power waveform. The timing of the trigger signal relative to the beginning of data processing in the DUT as well as the length of the trigger signal is user-configurable.

  DUT Reset In some cases, the control board may need to reset the DUT due to an error condition or because the interesting part of the victim algorithm has already executed. This is specifically valuable for ciphers that take a long time to complete. In this case, the cipher runs for a configurable number of clock cycles and then reset without waiting for it to complete. This helps reduce acquisition time.

  Timeout In some cases, due to communication error or DUT non-responsiveness the control board asserts a timeout error message to the control PC when a configurable time has elapsed.

  DUT Clock Generation The control board is capable of supplying a clock signal to the DUT. This signal is generated using a clock wizard with configurable frequency between 400 KHz and 100 MHz.

- **DUT Board** The DUT board is where the function core (a.k.a victim algorithm) is instantiated. The power consumption of the core FPGA voltage is measured using a power probe or shunt resistor.

Cryptographic hardware interfaces typically use multiple data types as input to cryptographic cores. For example, some algorithms might need plaintext/ciphertext, cryptographic keys, and random data. We provide a simple wrapper to split data provided by the control board to separate streams. This wrapper is directly compatible with CAESAR Hardware API interface and is expected to be directly compatible with a future Hardware API for Lightweight Cryptography (LWC API). We developed a simple, yet versatile protocol to enable the wrapper to split the data types. The wrapper receives data from the control board and distributes it into three FIFOs 1) the Public Data Input (PDI) FIFO (i.e. plaintext) 2) the Secret Data Input (SDI) FIFO (i.e. key) 3) the Random Data Input (RDI) FIFO which stores random data which can be used for protected implementations that use masking schemes. Once the wrapper prepares the data for the function core, it starts the core which consumes the data in the input FIFOs and produces output. The wrapper accumulates the output into a fourth FIFO called the Data Out (DO) FIFO until the expected number of bytes are stored. Then, the wrapper returns the data to the control board which forwards it back to the PC. To date we have validated Digilent Nexys3 boards (Spartan6 FPGA) and NewAE CW-305 SCA DUT (Artix7 FPGA) as DUT in FOBOS 2.

- **Oscilloscope and power measurement** The power consumption of the DUT can be measured using a current probe (e.g. Tektronix CT-1 current probe) if we are interested in relative changes in power consumption rather than absolute values. Alternatively, a shunt resistor can be inserted in the FPGA core voltage rail. The NewAE CW-305 DUT comes with a shunt resistor and amplifier on board which makes power measurements very easy. The current version of the capture module supports two oscilloscopes models, Picoscope 5000 via USB and Agilent DSO6054A via Ethernet.

3.1.2 The Analysis module

The analysis module is used to process the traces acquired by the capture module and run SCA attacks or leakage assessment.

- **Preprocessing** Many preprocessing steps may be taken to prepare traces for attack or evaluation. Traces can be compressed in time domain (i.e. combining multiple contiguous samples into one sample using mean, max or min functions). Also, samples at the beginning or at the end of the traces can be discarded to remove uninteresting parts of the trace.
- **CPA attack capability** The analysis software can run CPA attacks. The user provides the power model. CPA attack takes captured traces and hypothetical power values that are generated for each key guess. CPA is then used to perform a statistical correlation to find which key was likely used in the DUT. The output from the CPA attack is the highest correlation key, and graphs to show the highest correlation compared to the correlation of other key guesses. CPA attacks can also be used by testers to verify countermeasure effectiveness and show if a specific leakage is exploitable. However, as discussed, this requires identification of a specific power model, which can be difficult and time-consuming. CPA attacks are not performed in this work, but have been previously performed using FOBOS in [1, 26].

3.2 Power Measurement

We measure the power consumption of the Vcc_{INT} rail by measuring the amplified voltage drop across a 1 Ohm resistor while the DUT processes test vectors. Specifically, we used the XBP [19, 5] which provides the 1 Ohm resistor and TI-INA225 current sense amplifier. We connect Vcc_{INT} through the resistor in the XBP board to the DUT FPGA. When using the NewAE's CW-305 DUT, we cut the wire bridge between TP2 and TP3 and connected the Vcc wire from XBP to the FPGA through jumper JP7. Then the oscilloscope is used to measure the amplified voltage drop across the XBP's 1 Ohm resistor. A Python script is used to calculate the power using the data collected from the oscilloscope.

3.3 SCA Resistance Evaluation

3.3.1 TVLA Flow

FOBOS includes scripts that can perform fixed-vs-random TVLA. To perform this test, the user generates test vectors with fixed vectors randomly interleaved with random vectors. A meta file that records the type

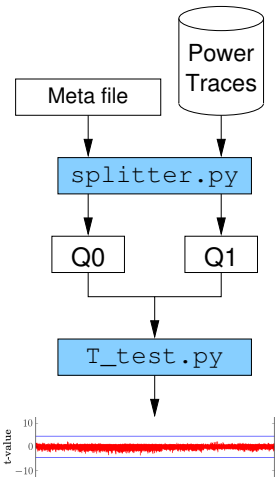

Figure 2: FOBOS TVLA Flow.

of each trace (i.e. fixed vs. random) is also generated. A script takes a fixed test vector as input (e.g., using CAESAR HW API test vectors generated by aeadtvgen in [12]), and generates the fixed-vs-random test vectors and the meta file. The test vectors are then fed to the capture module which processes them and produces power traces measured from the oscilloscope. The power traces are split into the random traces Q_r and fixed traces Q_f and passed to a script that calculates the t-values.

3.3.2 The χ^2-test Flow

The current χ^2 -test flow is based on two frequency classes "fixed" and "random;" as such, test vector generation and trace acquisition are identical to the TVLA. The two-class test differs only in the final analysis script, which calculates p-values for every sample, instead of t-values.

4 Results

4.1 Power Measurements and Benchmarking

We performed power measurements on FPGA implementations of four NIST LWC Round-2 candidates plus AES-GCM with the following breakdown:

- 3 NIST LWC Round-2 implementations using basic-iterative architecture (SpoC, Spook, GIFT-COFB).
- 1 NIST LWC Round-2 implementation using a multi-cycle lightweight approach (Ascon-small).
- 1 existing standard AES-GCM, using a pipelined lightweight approach.

The implementation details for SpoC, Spook, and GIFT-COFB are documented in [21]. The implementation details for Ascon-small and AES-GCM are discussed in [7].

We used the upgraded FOBOS platform with a NewAE CW305 Artix7 target board. Picoscope 5000 oscilloscope with XBP was used to measure power. All five implementations use the same CAESAR and LW Developer's package [29], benchmarked with Minerva hardware optimization tool [9] in Artix7 FPGA. Power is computed using the above methodology on 100 traces of four test vectors each (150 - 450 byte vectors) measured at 10, 25, and 50 MHz. Results of power measurements are highly linear in increasing frequency as shown in Fig. 3, which allows linear interpolation to estimate static power.

The measurements can be found in Table 1 and are shown Figs. 3 and 3b. In Table 1, abbreviations are Opt Freq (optimum frequency), LUT (look-up tables), P (power), E/bit (energy-per-bit). Opt Freq and area of SpoC, Spook, and GIFT-COFB are excerpted from [21]. P_{static} is estimated with linear interpolation.

Table 1: Characteristics of authenticated ciphers and their implementations investigated in this work.

	Opt Freq MHz	Area LUTs	Cycles/ Block	Bits/ Block	P_{static} mW	Freq MHz	P_{mean} mW	P_{max} mW	ΔP %	TP Mbps	E/bit nJ/bit	Gradient dP/dFreq
AES-GCM	240	1532	205	128	26.9	10	28.6	29.7	3.7	6.2	4.59	0.1808
						25	31.4	33.2	5.9	15.6	2.01	
						50	35.9	38.3	6.8	31.2	1.15	
Ascon	232	1808	82	64	26.8	10	28.1	29.0	3.0	7.8	3.60	0.1369
						25	30.3	31.5	4.1	19.5	1.55	
						50	33.6	35.0	4.3	39.0	0.86	
SpoC	265	1344	111	64	27.0	10	28.6	29.4	2.8	5.8	4.96	0.1529
						25	30.8	31.7	2.9	14.4	2.14	
						50	34.7	36.1	4.1	28.8	1.20	
Spook	141	7082	145	256	47.0	10	58.8	71.0	20.8	17.7	3.33	1.642
						25	96.5	116.6	20.9	44.1	2.19	
						50	125.9	161.6	28.4	88.3	1.43	
GIFT-COFB	172	2695	53	128	27.3	10	29.1	30.1	3.5	24.2	1.20	0.1871
						25	32.0	33.5	4.6	60.4	0.53	
						50	36.6	38.5	5.2	120.8	0.30	

ΔP is calculated as $(|P_{max} - P_{min}|/P_{min}) * 100$. E/bit is calculated as $E/bit(nJ/bit) = P(mW)/TP(Mbps)$. Gradient $dP/dFreq$ is $dPwr(mW)/dFreq(MHz)$. Below, we discuss some observations:

- Ascon has the lowest power at 50 MHz, followed by SpoC, however, Ascon, AES-GCM, SpoC, and GIFT-COFB are relatively close. Spook is the outlier in Fig. 3 with power consumption much higher than other ciphers.
- Ascon has the smallest gradient, i.e., slope of increasing power with increasing frequency. There is a 98% correlation between increases in area (LUT) and increases in dynamic power gradient. There is also a 76% correlation between decreasing optimal frequency (or increasing minimum period) and dynamic power gradient. However, Ascon has an area and minimum period larger than that of SpoC or AES-GCM, but has a lower dynamic power gradient.
- SpoC and Ascon have the smallest delta in percentage between maximum and mean power, which is a desirable design and security characteristic. Spook has up to a 28.4% delta between max and mean power at 50 MHz.
- GIFT-COFB has the lowest E/bit, 0.30 nJ/bit, versus the next lowest, Ascon, at 0.86 nJ/bit at 50 MHz. GIFT-COFB uses only slightly more power than Ascon at 36.6 vs 33.6 mW at 50 MHz. Since GIFT-COFB was implemented using basic-iterative architecture and Ascon using a multi-cycle approach, GIFT-COFB can probably be further optimized for power vs. E/bit.
- Static powers of all ciphers (except Spook) are 27.0 mW, ± 1%. The static power of Spook is much higher, likely due to its larger area.

4.2 Countermeasure Assessment

We performed leakage detection-based assessment on select lightweight implementations of the authenticated ciphers. We limit countermeasure assessment to AES-GCM and Ascon, since both unprotected and SCA-protected implementations of these ciphers are documented in [7], but no protected implementations of SpoC, Spook, or GIFT-COFB are documented in [21]. These two ciphers were implemented in RTL-level hardware using VHDL and protected using threshold implementation against first-order DPA. For details the interested reader is referred to [7].

We instantiated the implementations above in a Spartan6 xc6lx16-cs324 FPGA and supplied a 1 MHz clock. We used a Basys3 control board and Picoscope 5000 oscilloscope to collect traces. The oscilloscope sampling frequency was 125 MS/s. We used a 5 mV/1 mA Tektronix CT-1 current probe to measure current

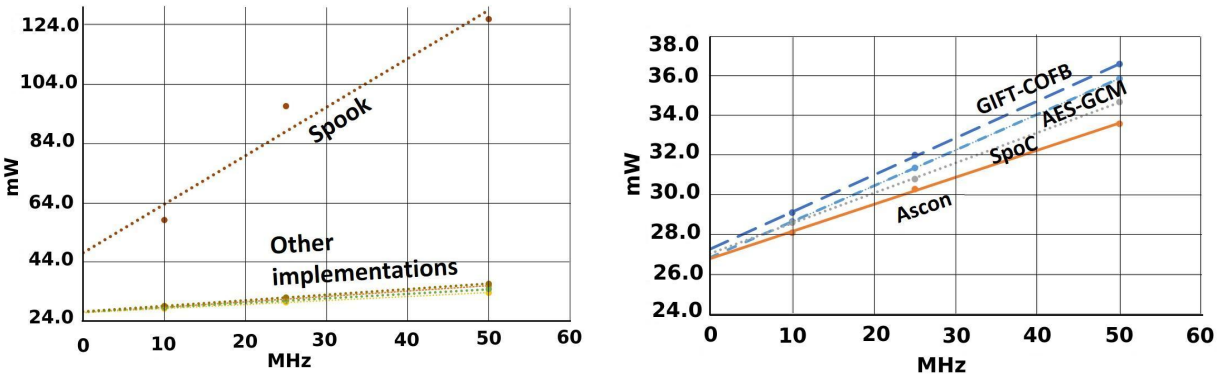

(a) All ciphers

(b) Spook excluded

Figure 3: Measured Power Consumption vs. Frequency

variations which are proportional to power consumption variations of the core voltage by the DUT FPGA. We then collected 2000 traces using fixed-vs-random test vectors. In all cases, trace collection took less than 3 minutes for each cipher implementation.

We then repeated the above experiments in the NewAE CW-305 Artix7 DUT and supplied a 1 MHz clock. This DUT features an Artix7 xc7a100tftg256-3 FPGA. We used a Basys3 control board and Picoscope 5000 oscilloscope to collect traces. The oscilloscope sampling frequency was 125 MS/s. Power measurements were taken from the CW-305 on-board low-noise amplifier (LNA) which amplifies the voltage drop across the on-board 0.1 Ohm shunt resistor inserted between the core Voltage and the FPGA. We then collected 2000 traces using fixed-vs-random test vectors.

After the traces has been collected, they were supplied to the analysis module to run TVLA. The results are shown in Figs 4 and 5. The two horizontal lines at $|t| = 4.5$ mark the threshold. T-values exceeding the threshold indicate detection of leakage with high confidence.

The χ^2 -test has also been performed for the unprotected and protected Ascon implementations on Spartan6 FPGA using the same traces used for t-test. The results are shown in Fig. 6. In this figure, values of $p < 10^{-5}$ are considered a failure as discussed previously.

TVLA results show significant first order leakage in the unprotected versions as expected. On the other hand, the protected versions show t-values within the threshold which implies no significant leakage is detected. The χ^2-test on the unprotected Ason detected leakage while the protected Ascon implementation shows no significant leakage which confirms the result obtained using TVLA.

5 Conclusions

We presented an upgraded FOBOS platform called FOBOS 2 suitable for performing power measurements and SCA resistance evaluation for hardware implementations of lightweight authenticated ciphers on modern Xilinx 7-series FPGA and corresponding target boards. We used the platform above to measure power and compute energy-per-bit for (E/bit) for selected cipher candidates in the NIST LWC standardization process, including Spoc, Spook, GIFT-COFB and Ascon, and included a comparison to a current standard, AES-GCM. Through measurements on the Artix7 FPGA, we found that Ascon has the lowest power consumption at 50 MHz, and lowest incrementally increasing dynamic power with increasing frequency, and that GIFT-COFB has the lowest E/bit. We also reason that GIFT-COFB power can be further reduced through innovative architecture, without large sacrifices in energy efficiency. We additionally validated SCA protection countermeasures on Ascon and AES-GCM on two FPGAs, including Spartan6 and Artix7.

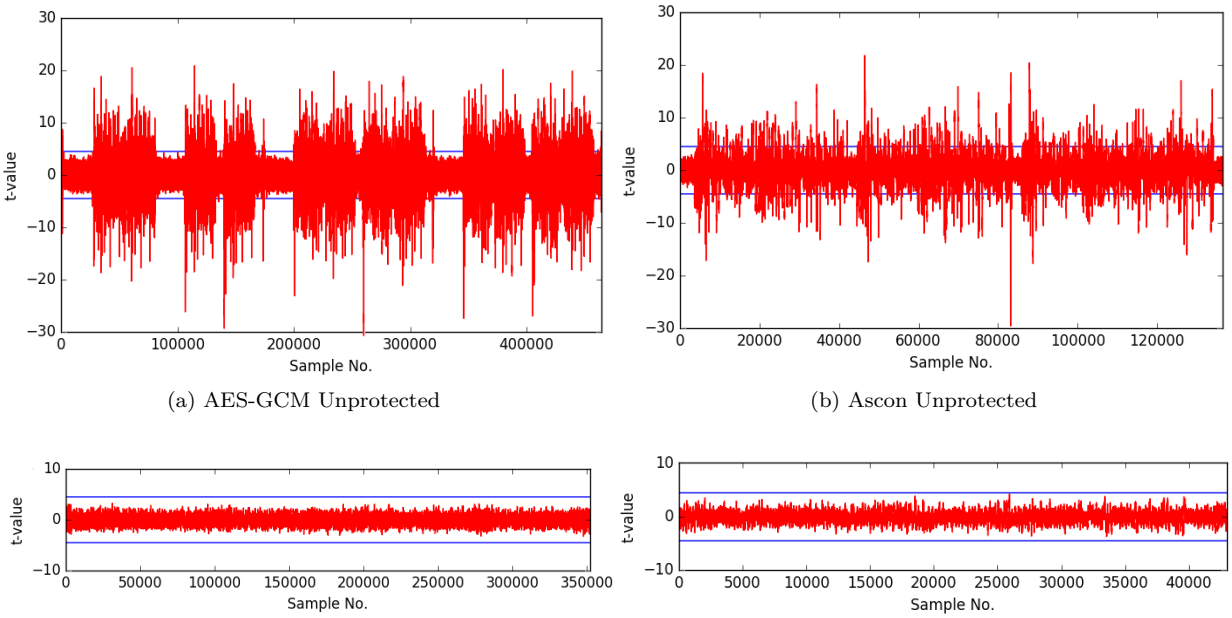

(a) AES-GCM Unprotected

(b) Ascon Unprotected

(c) AES-GCM Protected

(d) Ascon Protected

Figure 4: TVLA results for AES-GCM and Ascon on Spartan6 FPGA.

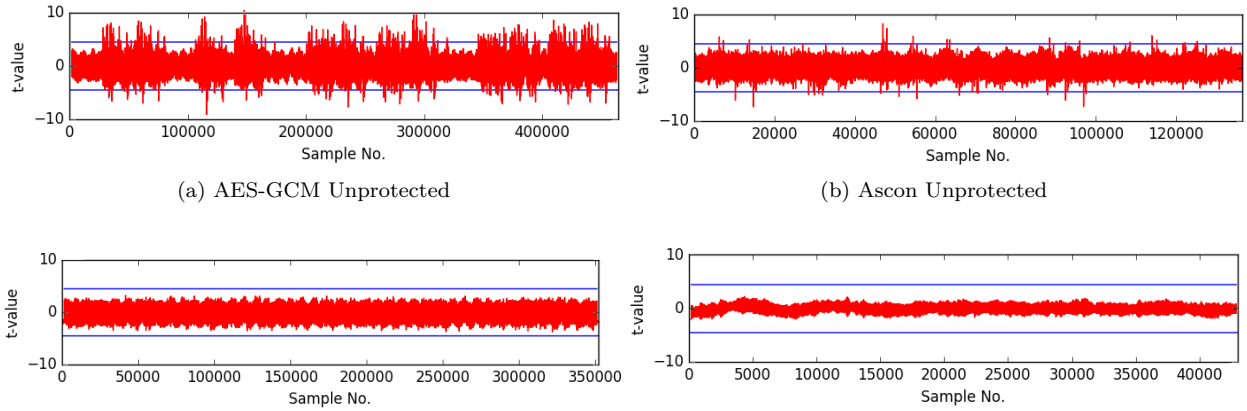

(a) AES-GCM Unprotected

(b) Ascon Unprotected

(c) AES-GCM Protected

(d) Ascon Protected

Figure 5: TVLA results for AES-GCM and Ascon on Artix7 FPGA.

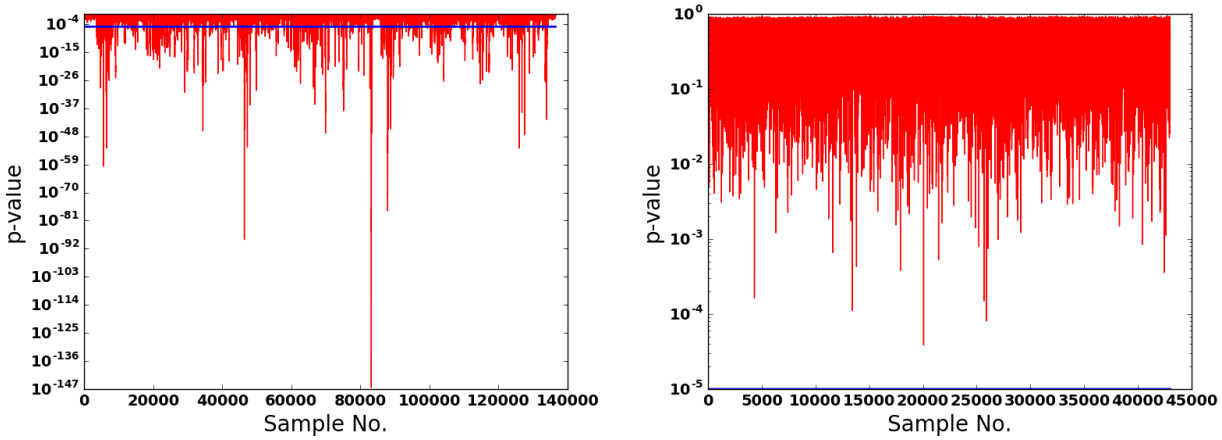

(a) Ascon Unprotected (b) Ascon Protected

Figure 6: χ^2-test on Ascon.

References

[1] Abdulgadir, A., Diehl, W., Velegalati, R., Kaps, J.P.: Flexible, Opensource workBench fOr Side-channel analysis. In: IEEE Hardware Oriented Security and Trust (2018)

[2] ARM: AMBA AXI Protocol Specification. https://developer.arm.com/docs/ihi0022/b/amba-axi-protocol-specification-v10 (2003)

[3] Brier, E., Clavier, C., Olivier, F.: Correlation Power Analysis with a Leakage Model. In: Cryptographic Hardware and Embedded Systems - CHES 2004, pp. 16–29. Springer Berlin Heidelberg (2004)

[4] CERG - GMU: FOBOS Home Page. https://cryptography.gmu.edu/fobos/ (2019)

[5] Cryptographic Engineering Research Group (CERG) at George Mason University: eXtended eXternal Benchmarking eXtension (XXBX). https://cryptography.gmu.edu/xxbx/index.php (2019)

[6] Diehl, W., Abdulgadir, A., Farahmand, F., Kaps, J.P., Gaj, K.: Comparison of cost of protection against differential power analysis of selected authenticated ciphers. In: 2018 IEEE International Symposium on Hardware Oriented Security and Trust, HOST 2018. pp. 147–152. Washington, DC (Apr 2018)

[7] Diehl, W., Farahmand, F., Abdulgadir, A., Kaps, J.P., Gaj, K.: Face-off between the CAESAR Lightweight Finalists: ACORN vs. Ascon. In: 2018 International Conference on Field Programmable Technology, FPT 2018. Naha, Okinawa, Japan (Dec 2018)

[8] Diehl, W., Farahmand, F., Yalla, P., Kaps, J.P., Gaj, K.: Comparison of hardware and software implementations of selected lightweight block ciphers. In: 2017 27th International Conference on Field Programmable Logic and Applications (FPL). pp. 1–4. IEEE, Ghent, Belgium (Sep 2017)

[9] Farahmand, F., Ferozpuri, A., Diehl, W., Gaj, K.: Minerva: Automated hardware optimization tool. In: 2017 International Conference on ReConFigurable Computing and FPGAs, ReConFig 2017. pp. 1–8. IEEE, Cancun (Dec 2017)

[10] Goodwill, G., Jun, B., Jaffe, J., Rohatgi, P.: A testing methodology for sidechannel resistance validation. In: NIST Non-Invasive Attack Testing Workshop. p. 15 (2011)

[11] Guntur, H., Ishii, J., Satoh, A.: Side-channel AttacK User Reference Architecture board SAKURA-G. In: 2014 IEEE 3rd Global Conference on Consumer Electronics (GCCE). pp. 271–274. IEEE, Tokyo, Japan (Oct 2014)

[12] Homsirikamol, E., Diehl, W., Ferozpuri, A., Farahmand, F., Yalla, P., Kaps, J.P., Gaj, K.: CAESAR Hardware API. Cryptology ePrint Archive 2016/626 (2016)

[13] Kocher, P., Jaffe, J., Jun, B.: Diferential Power Analysis. In: CRYPTO '99 - 19th International Conference on Cryptology. p. 10. Santa Barbara, CA (Aug 1999)

[14] Mangard, S., Oswald, E., Pop, T.: Power Analysis Attacks: Revealing the Secrets of Smart Cards. Advances in Information Security, Springer (2008)

[15] Moradi, A., Richter, B., Schneider, T., Standaert, F.X.: Leakage Detection with the $X2$-Test. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 209–237 (2018)

[16] Nikova, S., Rechberger, C., Rijmen, V.: Threshold Implementations Against Side-Channel Attacks and Glitches. In: Information and Communications Security, ICICS 2006. LNCS, vol. 4307, pp. 529–545. Springer Berlin Heidelberg (2006)

[17] NIST: Lightweight Cryptography | CSRC. https://csrc.nist.gov/projects/lightweight-cryptography/ (2019)

[18] O'Flynn, C.: A Framework for Embedded Hardware Security Analysis. Ph.D. thesis, Dalhousie University, Halifax, Nova Scotia (Jun 2017)

[19] Pham, J.: Benchmarking of Cryptographic Implementations on Embedded Platforms. Master's Thesis, GMU (2015)

[20] Rambus: DPA Workstation Analysis Platform - Rambus. https://www.rambus.com/security/dpa-countermeasures/dpa-workstation-platform/ (2019)

[21] Rezvani, B., Diehl, W.: Hardware Implementations of NIST Lightweight Cryptographic Candidates: A First Look. Tech. rep. (Jul 2019)

[22] Riscure: Side Channel Analysis Security Tools. https://www.riscure.com/security-tools/inspector-sca/ (2019)

[23] Schneider, T., Moradi, A.: Leakage Assessment Methodology: A Clear Roadmap for Side-Channel Evaluations. In: Cryptographic Hardware and Embedded Systems – CHES 2015, pp. 495–513. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)

[24] Tiri, K., Verbauwhede, I.: A logic level design methodology for a secure DPA resistant ASIC or FPGA implementation. In: Proceedings Design, Automation and Test in Europe Conference and Exhibition. pp. 246–251. IEEE Comput. Soc, Paris, France (2004)

[25] Trichina, E.: Combinational Logic Design for AES SubByte Transformation on Masked Data. Cryptology ePrint Archive 2003/236 (Nov 2003)

[26] Velegalati, R., Kaps, J.P.: Towards a Flexible, Opensource BOard for Side-channel analysis (FOBOS). In: Cryptographic Architectures Embedded in Reconfigurable Devices, CRYPTARCHI 2013. p. 6 (Jun 2013)

[27] Velegalati, R., Kaps, J.P.: DPA Resistance for Light-Weight Implementations of Cryptographic Algorithms on FPGAs. In: Field Programmable Logic and Applications, FPL 2009. pp. 385–390. IEEE (Aug 2009)

[28] Velegalati, R., Kaps, J.P.: Improving Security of SDDL Designs Through Interleaved Placement on Xilinx FPGAs. In: Field Programmable Logic and Applications, FPL 2011. pp. 506–511. IEEE (Sep 2011)

[29] Yalla, P., Kaps, J.P.: Evaluation of the CAESAR hardware API for lightweight implementations. In: 2017 International Conference on ReConFigurable Computing and FPGAs, ReConFig 2017. Cancun, Mexico (Dec 2017)