EFFLUX-F2: A High Performance Hardware Security Evaluation Board

Arpan Jati[1,2], Naina Gupta[1], Anupam Chattopadhyay[1], and Somitra Kumar Sanadhya[3]

[1] Nanyang Technological University, Singapore
{arpan.jati,anupam}@ntu.edu.sg, naina003@e.ntu.edu.sg
[2] Indraprastha Institute of Information Technology Delhi, India
[3] Indian Institute of Technology Jodhpur, India
somitra@iitj.ac.in

Abstract. Side-channel analysis has become a cornerstone of modern hardware security evaluation for cryptographic accelerators. Recently, these techniques are also being applied in fields such as AI and Machine Learning to investigate possible threats. Security evaluations are reliant on standard test setups including commercial and open-source evaluation boards such as, SASEBO/SAKURA and ChipWhisperer. However, with shrinking design footprints and overlapping tasks on the same platforms, the quality of the side channel information as well as the speed of data capture can significantly influence security assessment.
In this work, we designed EFFLUX-F2, a hardware security evaluation board to improve the quality and speed of side-channel information capture. We also designed a measurement setup to benchmark the signal differences between target boards. Multiple experimental evaluations like noise analysis, CPA and TVLA performed on EFFLUX-F2 and competing evaluation boards showcase the significant superiority of our design in all aspects.

Keywords: SCA · evaluation board · side-channel analysis · CPA · TVLA

1 Introduction

Side channel attacks have been used to attack a wide range of platforms and algorithms. Recently they have also gained importance to attack AI implementations. In order to standardize the side-channel evaluation platforms across different experiments with reproducible results, many works have been reported in literature over the past few years. Katashita et al. started the first project Side-channel Attack Standard Evaluation BOards (SASEBOs) [11,15] in this direction. As part of the project, the authors initially developed four boards SASEBO, SASEBO-G, SASEBO-R and SASEBO-R where SASEBO-R targeted a custom ASIC design and the rest were used to evaluate FPGA based hardware designs. But as the protected cryptographic designs require more resources, the previous platforms did not had sufficient FPGA logic cells. As a result, the authors updated the SASEBO-G design to be more compact, incorporate an FPGA with more logic cells, and support for few other features such as user-controllable configuration. The platform is commonly recognized as SASEBO-GII. This is followed by another platform (SASEBO-W) targeted towards evaluating smartcards [10]. To further extend the evaluation of more complex and integrated designs, Katashita et al. developed another side-channel board SASEBO-GIII [9] which is equipped with a 28-nm Kintex-7 FPGA. Later, the SASEBO project was terminated and SAKURA boards [8] (successors of SASEBO) were made available in the market.

Apart from these SASEBO/SAKURA boards, another effort by [14] led to the development of a modular platform design. The designers targeted to provide a complete setup which includes target device, measurement setup and trace capture and analysis software. The complete platform design is open-source. Another recent effort by [6] in which the authors designed the platform with Kintex Ultrascale FPGA to allow evaluations of larger designs such as Post-quantum cryptographic (PQC) [4,1] implementations. However, this design is dependent on the type of external power supply provided to the board. Hence, the experiments performed will vary with different power supplies.

In any side-channel attack, the signal-to-noise ratio (SNR) [12] of an SCA setup is a very important and crucial metric. A high SNR allows for cleaner signal and better experimental results. This leads to faster secret recovery or assessments and more confidence in the results. Hence, the main focus of our design is to improve the power circuit to enhance the overall signal quality. The next target is to demonstrate how this improved SNR aids in multiple experiments.

To provide security for constrained devices, smaller footprint or low complexity designs (lightweight cryptography [13]) has gained attention over the past few years. NIST has also started standardizing such designs [16]. But due to the smaller footprint, these designs have very low noise profile. Thus, making it quite difficult to perform side-channel leakage analysis on these designs. Hence, in this work we showcase how our low-noise and high performance EFFLUX-F2 board improves the commonly used Test Vector Leakage Analysis (TVLA) [3] experimental results. We also demonstrate Correlation Power Analysis (CPA) [5] attack differences. For this, we targeted GIFT [2] as it is one of the smallest lightweight cipher.

Further, as the side-channel attacks are becoming more prevalent and sophisticated, so are the countermeasures to protect these designs. The countermeasures if not implemented correctly may still leak information. For this, protected designs are also assessed and analyzed for leakage using TVLA. This analysis is typically performed by capturing and analyzing millions of traces as the designs are protected and may not leak initially. Capturing these many traces usually requires multiple hours. Hence, we also demonstrate how enhanced signal quality improves on the number of required traces and henceforth reduce the time to capture. For this, we used a known GIFT-128 protected implementation with known 1^{st} order leakages. The main contributions can be summarized as follows:

- We designed and developed a high performance and low-noise hardware board for side-channel evaluations. A detailed noise evaluation demonstrates the low-noise characteristics of EFFLUX-F2.
- We performed multiple experimental evaluations (SNR, CPA and TVLA) using a leaky unprotected AES implementation and a very lightweight cipher GIFT, having very low leakage levels compared to the former.
- We also compare EFFLUX-F2 with state-of-the-art commonly employed side-channel evaluation board SAKURA-X.
  - Statistical analysis of the raw noise measurements demonstrate that EFFLUX-F2 has at least 4.5× lower noise levels than SAKURA-X.
  - EFFLUX-F2 achieves almost 8.2× higher SNR than SAKURA-X for AES.
  - For CPA attack on GIFT, we achieve a reduction of at least 5× the required number of traces. We also showcase how EFFLUX-F2 clearly distinguishes between the correct key and wrong keys with high probabilities. Whereas, this is not true for SAKURA-X.
  - Furthermore, we show even though the performance improvement (t-value vs number of traces) for TVLA on an unprotected AES is moderate. It is quite drastic in the case of GIFT. This difference becomes more prominent in case of a partially protected design where the leakage detection is ≈20× better for EFFLUX-F2.

2 EFFLUX-F2

EFFLUX-F2 as shown in Fig. 1 is a high performance FPGA board specifically designed to improve measurement accuracy and noise characteristics. While designing the board we used low-noise power supplies and many low-noise design principles to minimize noise and signal interactions. The board is also designed as a general-purpose FPGA board allowing for many use cases in multiple scenarios. While designing the board we had the following requirements in mind:

- **Low EMI and Noise:** In order to improve measurement quality for both power and EM measurements, keeping the noise low is very important. One normally needs a large number of traces to detect leakage from weakly protected implementations, we aim to improve the SNR significantly in order to make the experiments faster and more reliable.
- **Adjustable Voltages and Fault Injection Support:** Possibility to undervolt the VCCINT rails and voltage glitch fault injection support, with high temporal resolution.
- **FPGA and DRAM measurements:** The current offerings do not support DRAM power measurements. We intend to provide support for the same to make it possible to exploit this avenue in future attacks.
- **Single FPGA design:** To simplify the design and reduce system noise and cost. Two FPGAs are always better, but with careful hardware and software design, we can use a single FPGA in most applications, without much side-effects.

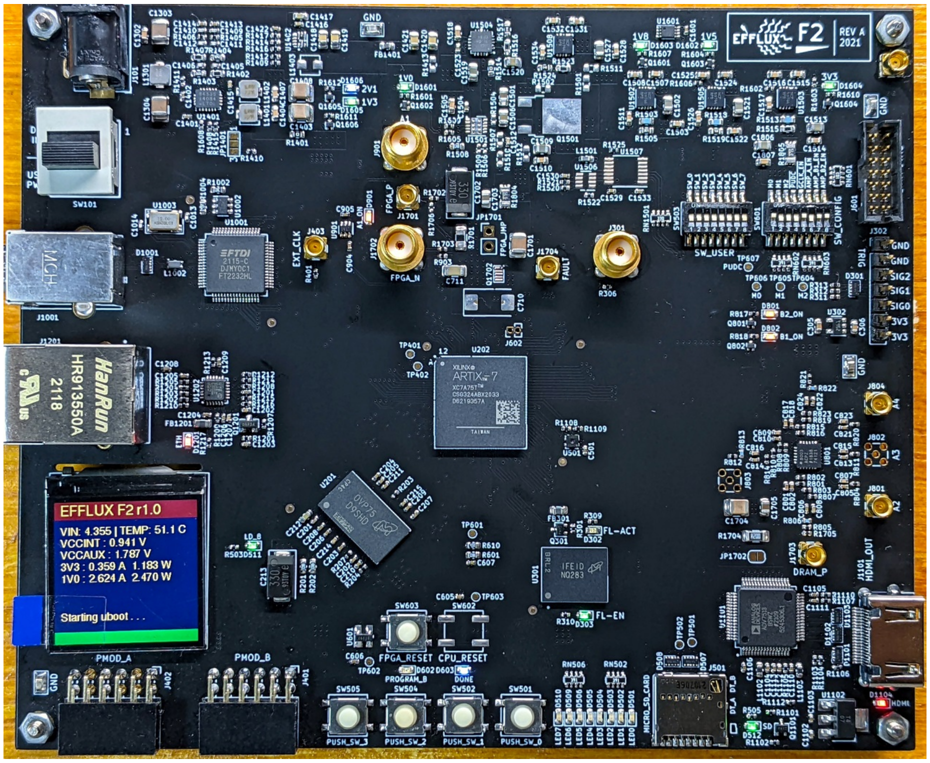

Fig. 1: EFFLUX-F2 Board: LCD showing U-boot starting to load Linux.

- **DRAM, Ethernet, SDIO and HDMI:** To closely replicate real systems and associated noise, DDR3 RAM, Ethernet, SDIO and HDMI output.
- **Built-in amplifiers:** To simplify the trace capture setup, multiple onboard amplifiers are required.
- **General purpose use:** Additional LEDs, LCD display and switches so that the board can be used as a general purpose FPGA development board as well.
- **Compatibility:** By using the same USB-IF chip used in SASEBO/SAKURA boards, the board remains software compatible with existing software setups.

Fig. 2 shows the block diagram for the EFFLUX-F2 board. All the devices like RAM, Flash, USB IF, Display interfaces and Ethernet PHY etc., are connected to the FPGA. A specifically designed power supply targeting low noise and EMI emissions is powering all the devices on the board. The power delivery system contains current-sense resistors and sense amplifiers to measure the current in 1.0V and 3.3V rails. This allows for precise measurement of FPGA VCCINT current measurement, leading to very precise FPGA core power measurement. In addition to voltage measurement, we have also added a very precise environment sensor with accuracy of ±0.2°C and ±2%RH temperature and humidity respectively. To keep noise at a minimum, there are no additional microcontroller or CPLD for housekeeping purposes. Further, all the devices other than the FPGA do not contain any additional processing units to avoid noise generation. In addition to this, all noisy peripherals can be power-gated to reduce noise in the captured power traces, this includes high frequency clock generators.

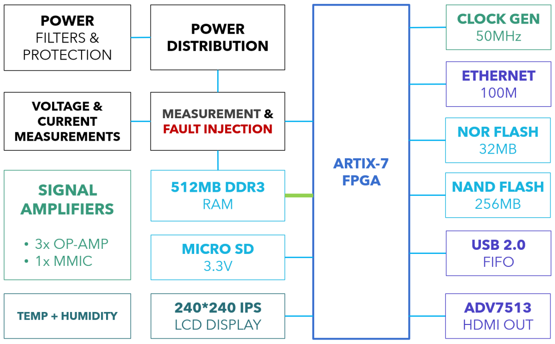

Fig. 2: EFFLUX-F2: Block Diagram

2.1 Choice of FPGA

The board was designed with flexibility and cost in mind. We decided to use Xilinx Artix-7 family of devices. These have wide market adoption and are manufactured using TSMC 28nm HPL process. On a fundamental design level these devices are very similar to the other 7-series FPGA lines like the Kintex and Virtex. The board is designed to support the `CSG324` 0.8mm pitch package. This package supports multiple devices from XC7A15T (16.6K Logic Cells) to XC7A100T (101K Logic Cells). We initially planned on using a 1.0mm pitch `FTG256` package device, but had to migrate to the larger 324-ball package as we quickly ran out of pins during the design phase. Only four pins of the current FPGA are left unused.

Further, having a single FPGA instead of two FPGAs or FPGA + Microcontroller combination has multiple benefits like reduced board complexity and fewer sources of noise. As most of the power leakage comes from switching noise, clock gating the non-cryptographic portion would allow for a design with a single FPGA while having minimal side effects on the noise performance. The clock gating can easily be implemented with small modifications to the hardware.

2.2 Power Supply

Noise sources in any system can vary widely. While designing the board we consider the additional noise generated by the power supply and attempt to

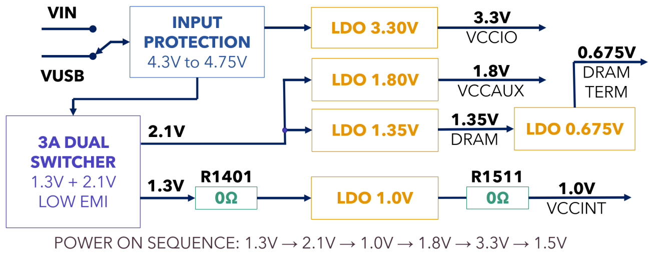

Fig. 3: Voltage Regulation and Power Delivery

keep them at a minimum. The majority of this additional noise comes from the switching noise of power supply components like buck converters. The buck converters are DC-DC power converters which convert a higher voltage to a lower voltage. An inductor is used to temporarily store energy while capacitors are used to reduce ripple. Using such a topology, high efficiencies of around 95% can be achieved, leading to lower power loss through heat and smaller circuits. Unfortunately, these circuits are inherently noisy because the switching activity of the power MOSFETs cause high currents to pass through inductors and capacitors. The inductors store energy in magnetic fields, and rapid switching causes a lot of EM emissions.

In this work, we follow many design techniques, like reducing high current loop areas, appropriately sized and placed capacitors, proper ground planes, protecting sensitive signals from noisy traces and many others during the design optimization process. LDO (Low dropout) linear regulators are known for their low noise, high PSRR (Power Supply Rejection Ratio) and good transient response. We use these devices as post regulators after the initial switching mode power supplies. The switching regulators are also running with spread-spectrum enabled, this distributes the conducted and radiated EM over a wider frequency band. These steps allow us to significantly reduce ripple on the power rails. We also use resistors and components (references, OPAMPs and regulators) with low TCR (Temperature Coefficient of Resistance) of $\pm 10ppm/^{\circ}$C or better and high accuracy 0.1% wherever applicable for better signal drift characteristics. This helps in experiments that run over a long time and face changing DC levels caused by temperature effects.

Fig. 3 shows the voltage regulation and power delivery topology used in the board. The input power first passes through a fuse, a MOSFET based reverse voltage protection circuit and a 5.76V over-voltage protection circuit. It is then filtered using a wide-band high order power filter with an *insertion loss* of more than 60dB for frequencies between 100 KHz and 100 MHz. This filtered power then, passes through a 3rd order π type LC EMI filter built using discrete components, before reaching a low EMI switching regulator.

The idea is to convert the 5V power input (USB or external input) to 1.4V using a switching regulator and then regulate the 1.4V to 1.0V using an LDO for the FPGA VCCINT rail. The LDO input to output voltage difference of

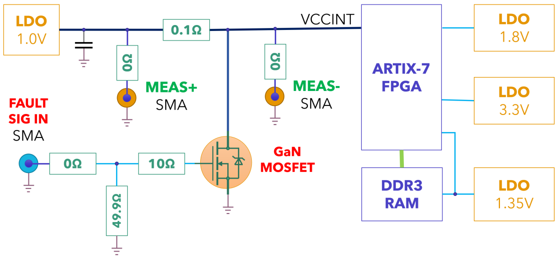

Fig. 4: FPGA Core Power Measurement and Fault Injection

0.4 volt is higher than required, but it is intentional as it helps in improving the PSRR performance of the regulator. The 0.675V, 1.35V and 1.8V rails are generated from a 2.1V switcher, using three LDOs. The 3.3V rail is generated directly from the input 5V supply, this is fine as the current demand on this rail is not very high. We use power-sequencing (as shown in the figure) to ensure safe and reliable FPGA power up. We also use a dedicated power on reset chip (with voltage sense) to ensure reliable reset signal to the FPGA, both on power on and manual reset.

2.3 FPGA power measurements

Fig. 4 shows the FPGA core power measurement setup. The board uses high side current measurement. 0.1Ω and 0.36Ω resistors are used for the FPGA and DRAM measurement respectively. Measurement points are provided on the board to access these power signals. Multiple amplifier circuits are additionally added to amplify these small signals so that an oscilloscope can be directly connected without inline amplifiers. This simplifies the overall setup. In the figure, the voltage drop through the 0.1Ω resistor is amplified by either of the amplifiers and can be used as the leakage signal. An SMA connector also allows for an external amplifier to be used alternatively. Faults can be injected by using an FPGA or pulse generator to generate glitch signals of the required width and then driving it to the gate terminal of the MOSFET. The power supply is designed to handle shorts to the ground for small duration. The MOSFET has a current rating of more than 20A and resistance close to $3\text{m}\Omega$. As we are using a GaN FET the total gate charge is just 6.6 nC, enabling fast switching performance.

Fig. 5a shows the back side of the board with the 5× VCCINT capacitors removed. This is required to improve the captured signal quality.

Onboard Amplifiers We implemented two amplifiers, a 30dB MMIC (Monolithic microwave integrated circuit) amplifier and a 3 channel wide-band Low-noise 20dB op-amp based amplifier as shown in Fig. 5b. Both the amplifiers are 50Ω matched and AC coupled using 0.1μF capacitors. The -3dB bandwidth for

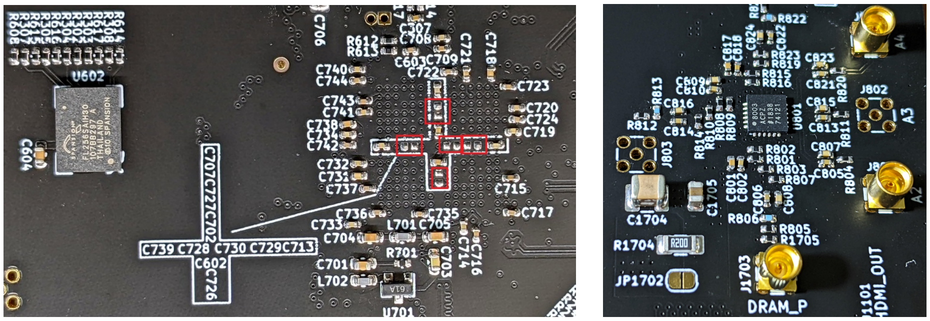

(a) 5× VCCINT rail capacitors removed to improve signal quality.

(b) AD8003 Op-amp Based Amplifier.

Fig. 5: EFFLUX-F2: PCB under FPGA and Amplifer

the amplifiers are 2.2GHz and 446MHz respectively. These separate amplifiers provide flexibility in optimizing the measurement setup. The amplifiers can be disabled on a per channel basis to reduce coupled noise when a channel is not needed.

2.4 Clocks Generation

A precise *low-jitter* clock helps SCA measurements. The 50 MHz clock generator SiT9121AI from SiTime is used. LVDS version of the chip is employed with 10ppm stability and 1.2ps RMS (Root Mean Square) period jitter. External clock input and output is also supported.

2.5 Memories

There are four types of supported memories on the board:

- **32 MB Flash:** To store FPGA bitstream or OS images.
- **512 MB DDR3:** 16-bit High speed DRAM (1600 MB/s).
- **NAND Flash:** Large file or OS storage.
- **SDIO:** SD Card.

These memory features allow for testing of countermeasures in realistic scenarios. Further, hardware/software co-design based designs running on Microblaze/RISCV etc. can be tested and verified. Further, the current offerings do not support direct DRAM chip level power measurements. We enable support for the same to enable new avenues of attacks.

2.6 I/O Interfaces

An FT2232HL USB 2.0 chip from FTDI is employed for the USB interface. It supports 12 Mbaud (UART) and up-to 40MB/s (Sync FIFO) using two independent UART/FIFO interfaces. As the interfaces are independent and come with separate 4KiB TX and RX internal FIFO buffers the chip offers high performance and minimal latency. This allows for fast transfer of data like keys, plain-text, cipher-text etc. The board also has additional protected ports for trigger in and out. To aid in development and debug tasks, multiple devices like LEDs, push switches, slide switches, GPIO pins, and configuration switches are also added. We further implemented 2x `Digilent PMOD` compatible pin-out and physical connectors so that a variety of extension boards can be used. This is a board primarily designed for side channel attacks, adding Ethernet which may be quite noisy is counter-intuitive. We have added Ethernet as a means of easy data transfer especially while using Linux or performing attacks on complex high throughput AXI peripherals. Power gating is supported to fully disable this unit when performing analysis on low noise or low leakage designs. We use `LAN8720A`, which is a 10/100 Mbps `RMII` transceiver. The transceiver supports `IEEE 802.3u` and Auto-negotiation.

2.7 OS + System Support

Additionally, one of the goals behind designing the board was to allow side-channel trace capture in realistic scenarios with significant background noise for certain experiments. To closely replicate real systems, as discussed above we added the support for 2Gb (256 MB) DDR3 RAM. Further, SDIO was added to support large storage devices to enable booting OSes like Linux. HDMI output was added for display support. All these interfaces can be fully disabled when needed especially while preforming noise sensitive experiments.

2.8 ESD Protection

The board is protected against ESD. We used IEC 61000-4-2 Level 4 compliant ESD Protection(±12-kV Contact Discharge protection). TVS Diodes with low capacitance are used resulting in negligible distortion to the protected signal lines. All the I/O interfaces like Trigger, PMOD, USB and HDMI pins are protected. The USB and HDMI interfaces are additionally protected using common mode chokes. The protection devices makes the board much more robust against accidental ESD strikes while handling and touching.

2.9 Power Gating

Power-gating is implemented for multiple modules in the board. The SD-Card, Ethernet, Flash, HDMI, individual signal amplifiers etc. can be disabled to reduce noise. The power gating is implemented using P-Channel MOSFETs controlled by signals from the FPGA and individual slide-switches wherever applicable.

2.10 PCB Design and Routing

We used multiple design techniques, including the ones discussed in section 2.2 while placing components and routing the board to improve the performance of the board. The fully routed PCB is shown in Fig 6. It is evident that a minimum of the traces are routed on the outer layers (red and green), this is done to minimize EMI emissions. Apart from keeping noisy signals away from sensitive ones we also separated critical signals using ground planes. We initially started with an eight layer PCB design, but after routing all the important signals, it was determined that it is possible to have very similar performance and isolation even with six layers; so, the extra layers were removed and the design finished with six layers instead.

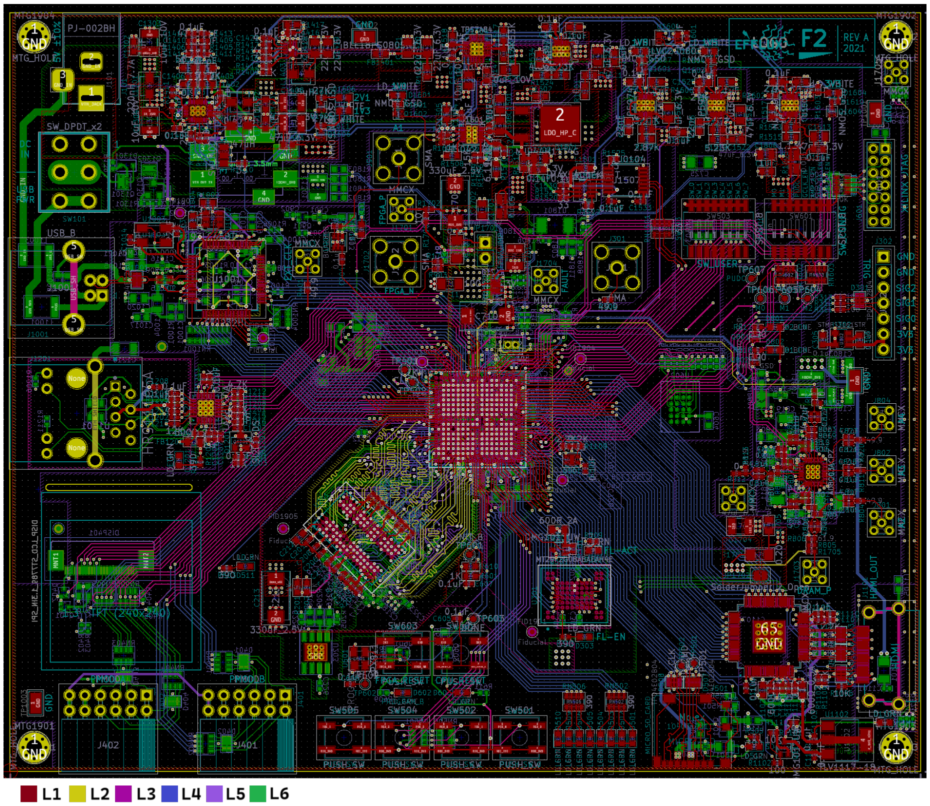

Fig. 6: PCB Routing all internal layers. Ground planes not shown for clarity.

3 Experimental Setup

In this work, we present results for multiple experiments using the following metrics:

- Direct Noise Measurements (FPGA VCCINT noise)
- Signal to Noise Ratio (SNR)
- Correlation Power Analysis (CPA)
- Test Vector Leakage Assessment (TVLA)

In order to obtain consistent repeatable results we perform experiments with fixed settings. The SNR, TVLA and CPA experiments use the same setup, Agilent DSO6034A oscilloscope, PA303 Amplifier from Langer EMV-Technik and a 50 MHz Low pass inline SMA filter from Crystek Corporation. All these experiments use traces captured at 2GS/s.

While measuring noise directly, we were faced with many challenges. The first being the fact that we are interested in measuring noise coming from the power supply and not from the FPGA. As the power supplies are designed to be low noise, the voltage levels are very small in the tens to hundreds of μV range. This level is easily below the noise floor of most oscilloscopes as the lowest gain range is often close to 1mV/division. Second, when we tried to measure noise using off-the-shelf amplifiers (from multiple vendors) we quickly ran into the problem that the noise measurement for EFFLUX-F2 was either below or close to the noise floor. Cascading multiple amplifiers did not help as the input noise density (1.5-2.2 nV/$\sqrt{Hz}$) of the first amplifier in the chain is still quite high. To get around these issues we designed a multi-stage custom amplifier with a differential input section constructed using discrete low-noise transistors (BJTs). The amplifier has an adjustable gain of approx 70dB-80dB, and an extremely low input voltage noise density of 465pV/$\sqrt{Hz}$. Further, two separate 2^{nd} order Sallen-Key LPF filters before the output stage allows us to band-limit the signal to 100 kHz and 1 MHz, the unfiltered signal has a bandwidth of around 10 MHz. The performance is enough to measure power supply noise from very clean LDO devices with good accuracy. For noise measurements, the traces are captured at 200 MS/s given the reduced bandwidth.

4 Evaluation Results

In this section, we present and discuss detailed comparison of EFFLUX-F2 with SAKURA-X. Apart from SAKURA-X, CW305 and CW310 boards from NewAE also allows power measurement. Unfortunately, the boards do not use low noise power supplies; the FPGA in the boards are powered directly using noisy switching regulators. This is good enough for many applications, but more that an order of magnitude higher noise levels is measured from CW305 ($\approx 72\mu$V RMS) compared to EFFLUX-F2. Additionally, as the noise level is too high for our high gain noise measurements setup (without added attenuation) and the measured SNR is lower compared to SAKURA-X, we do not include it in our detailed evaluation.

4.1 Noise Measurements

FPGA core voltage noise measurement allows us to compare multiple boards in a direct manner. For this, we power the board from a fairly noise free power supply, USB or battery and then place the target FPGA under reset. A power trace is then captured using an oscilloscope using the 'single' capture mode to provide the longest possible trace length (4 million points). We use the high-gain low-noise amplifier as discussed above to amplify the power signal by around 76dB (6300×). Even though we could use a higher gain, we chose to use 76dB as this limited the peak to peak signal level to around 1V for all experiments, having a higher amplitude would cause other undesirable side effects. For consistency, the same power source is used for both the boards. To transform the time domain results to frequency domain, an FFT is performed on the captured trace points and the results are plotted on a log-log graph. A trace labeled NOISE FLOOR is also added to the graphs. It corresponds to the noise floor of the measurement system, and is obtained by shorting the amplifier input to the ground.

Battery Powered System (board generated noise) This experimental setup captures voltage signal from both the boards powered from two Li-Ion cells in series at a voltage of 8.2V. The setup is designed to show the inherent noise of the power supplies, and as we do not have any external higher frequency noise sources, we are band-limiting the signals to 1 MHz for this experiment.

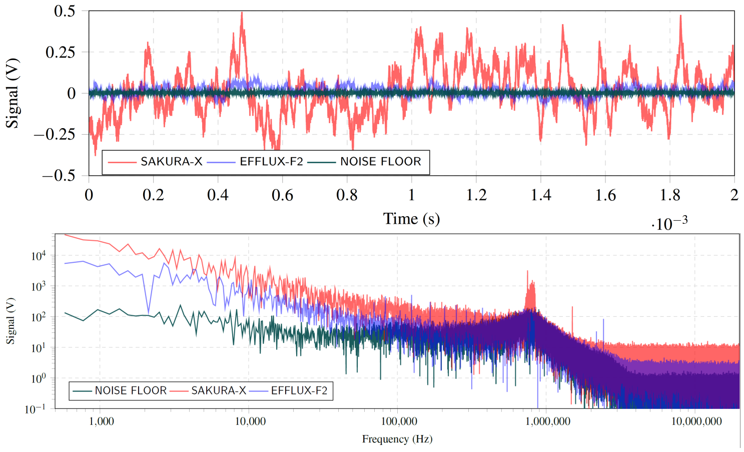

Fig. 7: EFFLUX-F2 vs. SAKURA-X. Battery Powered devices, 1 MHz bandwidth. The graph on top shows amplified voltage signal from the VCCINT rail while the bottom one shows the corresponding FFT (log-log scale).

It can be seen in Fig. 7 that EFFLUX-F2 (shown in blue) has a much lower noise amplitude when compared to the SAKURA-X board. Noise levels of EFFLUX-F2 is very low, but the even lower noise floor of our measurement setup allows us to demonstrate the accuracy of the results with high confidence. One can also notice the small peak in noise around 800 kHz for SAKURA-X, this is the operating frequency of the switching regulator. Such peaks are not visible in EFFLUX-F2.

USB Powered System (typical use case) This experimental setup captured voltage signal from both the boards, powered from USB (connected to PC) at a voltage of 5.1V. To show the input power noise filtering, no filter was used and the signal was band limited by the amplifier's bandwidth which is around 10 MHz (Wideband setup).

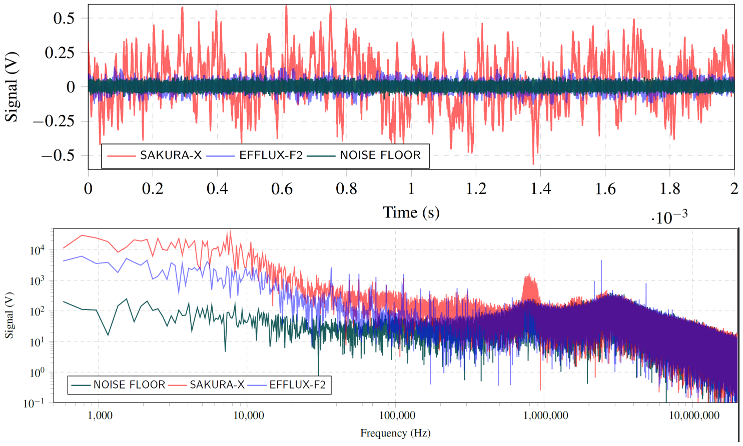

Fig. 8: EFFLUX-F2 vs. SAKURA-X. USB Powered devices, 10 MHz bandwidth. The graph on top shows amplified voltage signal from the VCCINT rail while the bottom one shows the corresponding FFT (log-log scale).

Fig. 8 shows the voltage trace and the corresponding FFT while the boards are operating from a USB power source. It can be seen that EFFLUX-F2 (shown in blue) has a much lower noise when compared to the SAKURA-X.

Statistical analysis of the measured noise Table 1 shows the noise statistics for the boards. To calculate the voltage at input or the amplifier (or the measurement point of the board) we scale the measured signal with the gain of the amplifier which is set to around 76dB for the experiments. Under USB powered condition, when directly comparing the two boards, we can see that power measurement noise levels in EFFLUX-F2 is 4.62× lower than SAKURA-X when comparing the RMS noise value. While operation using batteries leads to 5.14× improvement for the same metric.

Table 1: Noise measurement statistics, battery-powered @ 1 MHz B/W, USB-powered @ 10 MHz B/W. The amplifier input is connected to the boards.

	Parameter	Measured at amplifier output			Calculated at amplifier input		
Battery-powered		NOISE FLOOR	SAKURA-X	EFFLUX-F2	NOISE FLOOR	SAKURA-X	EFFLUX-F2
	MEAN	2.504 mV	10.907 mV	12.437 mV	396.838 nV	1.729 μV	1.971 μV
	RMS	5.626 mV	149.032 mV	28.977 mV	891.819 nV	23.622 μV	4.593 μV
	Vpp 6σ	33.759 mV	894.193 mV	173.862 mV	5.351 μV	141.733 μV	27.558 μV
	STDEV	5.039 mV	148.633 mV	26.172 mV	798.662 nV	23.559 μV	4.148 μV
	VARIANCE	25.389 μV	22.092 mV	684.992 μV	0.638 pV	555.018 pV	17.209 pV
USB-powered	Parameter	Measured at amplifier output			Calculated at amplifier input		
		NOISE FLOOR	SAKURA-X	EFFLUX-F2	NOISE FLOOR	SAKURA-X	EFFLUX-F2
	MEAN	3.419 mV	32.492 mV	4.387 mV	541.936 nV	5.150 μV	695.386 nV
	RMS	20.584 mV	176.850 mV	38.207 mV	3.263 μV	28.031 μV	6.056 μV
	Vpp 6σ	123.502 mV	1.061 V	229.241 mV	19.576 μV	168.188 μV	36.335 μV
	STDEV	20.298 mV	173.839 mV	37.954 mV	3.217 μV	27.554 μV	6.016 μV
	VARIANCE	411.998 μV	30.220 mV	1.441 mV	10.351 pV	759.234 pV	36.191 pV

4.2 Signal-to-Noise Ratio (SNR)

The SNR of a side-channel trace is a very important metric as it helps determine the overall quality of the measurement setup. The SNR is the ratio of $\frac{Var(\mathrm{V_{signal}})}{Var(\mathrm{V_{noise}})}$ [12], where $\mathrm{V_{signal}}$ is the data-dependent signal component, $\mathrm{V_{noise}}$ is the random noise component and Var denotes the variance. The SNR of a side-channel setup is inversely related to the number of traces ($\mathrm{N}_{traces} \propto \frac{1}{SNR}$). This means a high SNR setup requires fewer number of traces for attack or analysis compared to a low SNR setup, where the noise dominates. There are many ways of measuring SNR [12,7]. In our experiments we calculate two sets of traces. The $Var(\mathrm{V_{noise}})$ is captured by using random plaintext, whereas for calculating $Var(\mathrm{V_{signal}})$ we captured averaged traces with $N = 100$, in other words, 100 traces with the same plaintext were captured and pointwise averaged. For our evaluation, we captured 10K traces for noise and $10K * 100 = 1M$ traces for the signal. Thus, for final comparison between the two boards, we utilized 10K raw traces for noise and 10K averaged traces for signal.

Fig. 9 shows the trace captured using both the boards for AES and also the signal and noise traces for all the sample points. One interesting thing to note from Fig. 9b is that in case of SAKURA-X, the noise trace dominates the signal trace. Whereas in case of EFFLUX-F2, the signal dominates than noise Fig. 9d, which should be the ideal case for any side-channel measurement setup. To further highlight this fact, we also show the comparison of Signal-to-Noise ratio in Fig. 10. One can see that there is a huge difference in the SNR of the captured traces between the two boards. The SNR for EFFLUX-F2 is almost 8.2× better than SAKURA-X for the same settings. This clearly demonstrates that even for a highly leaky design such as unprotected AES, EFFLUX-F2 captured signal quality is much better than currently used SAKURA-X.

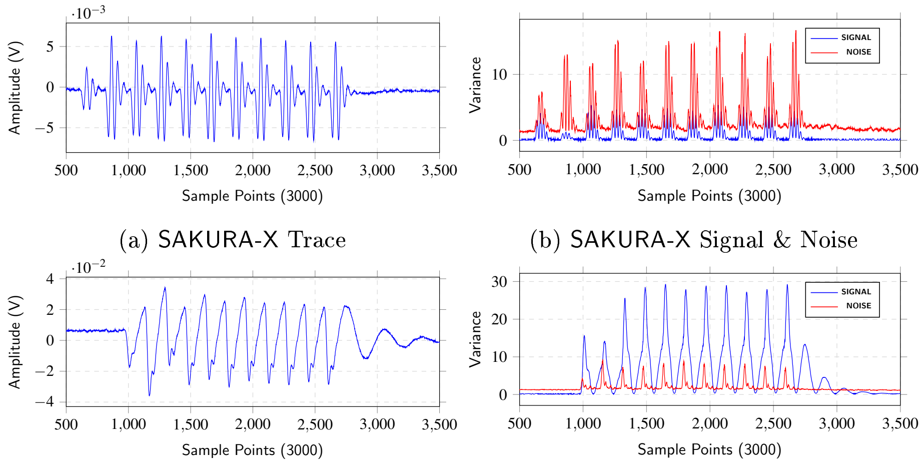

(a) SAKURA-X Trace

(b) SAKURA-X Signal & Noise

(c) EFFLUX-F2 Trace

(d) EFFLUX-F2 Signal & Noise

Fig. 9: Signal and noise traces for all the sample points.

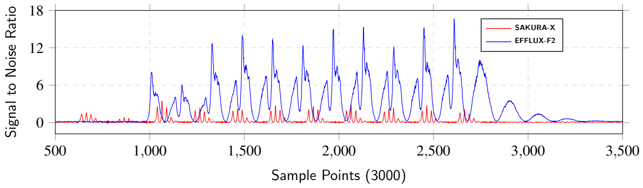

Fig. 10: Signal-to-Noise Ratio Comparison. The maximum SNR in the points of interest (last round of AES) for SAKURA-X is measured at sample point 2664 and is 2.020. Whereas, for EFFLUX-F2, the maximum SNR is measured at sample point 2609 and is 16.562.

4.3 Correlation Power Analysis (CPA)

CPA is a well-known side-channel attack which exploits the correlation of the power with the data to extract secret key. We performed CPA attack on an unprotected implementation of lightweight cipher GIFT. We intentionally chose this as the target design for demonstrating the comparison between the two boards as the power consumption for GIFT is close to the noise floor. We targeted the last round of the cipher and considered hamming distance model for our attack.

Fig. 11 shows results for correlation values for all possible keys corresponding to number of traces. Due to space constraint, the results are shown corresponding to key bytes 0, 1 and 2 for both the boards. Similar trends are observed for other keys as well. One can see that for key byte 0, almost 60K traces are required using SAKURA-X to distinguish it from other possible key values. Whereas, the same key can be recovered using only 12K traces from EFFLUX-F2 achieving a

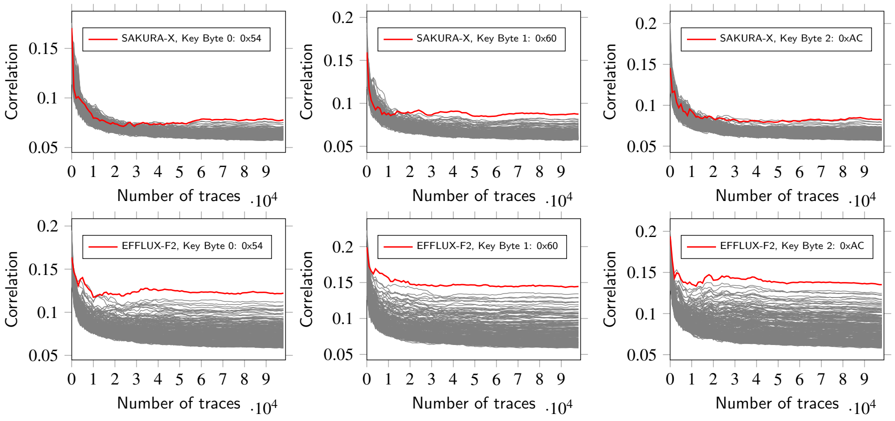

Fig. 11: CPA results for unprotected GIFT

reduction of almost 5×. Similarly, EFFLUX-F2 requires 5× and 17× less traces than SAKURA-X to recover key byte 1 and 2 respectively. One should also note that in case of SAKURA-X, the correlation values between top two-three values are quite close. Whereas, in case of EFFLUX-F2, it is quite consistent and clearly distinguishable for all the three key bytes. Thus, providing high confidence towards key recovery. We would also like to highlight the fact that this huge reduction in required number of traces is quite significant in terms of time required to collect these many traces.

4.4 Test Vector Leakage Assessment (TVLA)

TVLA is a commonly used technique to detect any type of leakage (source may be unknown) rather than exploit the leakage in any system. If the t-value crosses a certain threshold (commonly used threshold ±4.5), then it is considered that the leakage is detected. We used incremental formulae for our calculations.

To present comparison between the two boards, we performed our evaluation around two cryptographic ciphers; an unprotected AES implementation (composite implementation) and a lightweight low noise implementation of GIFT-128. We first performed non-specific TVLA analysis on unprotected implementations of both AES and GIFT. The results are shown in Fig. 12. We show pointwise t-values as well as incremental t-values in the graph. As the AES implementation leaks significantly, the threshold value easily crosses 50 just after 2000 traces for both the boards. EFFLUX-F2 shows more leakage (higher t-value) than SAKURA-X. This fact is more evident when comparing the results for GIFT. GIFT is a lightweight cipher and consumes power within noise floor level. Hence, it becomes difficult to analyse leakage of such a design. As can be seen from Fig. 12c, the t-values obtained from traces captured by SAKURA-X is barely crossing the threshold ±4.5. Whereas, using EFFLUX-F2, it is quite evident the design is highly leaky as an unprotected design will be. This is also visible from the incremental t-values for GIFT in Fig. 12d, where EFFLUX-F2 t-value is almost five times the SAKURA-X t-value after analyzing 20000 traces.

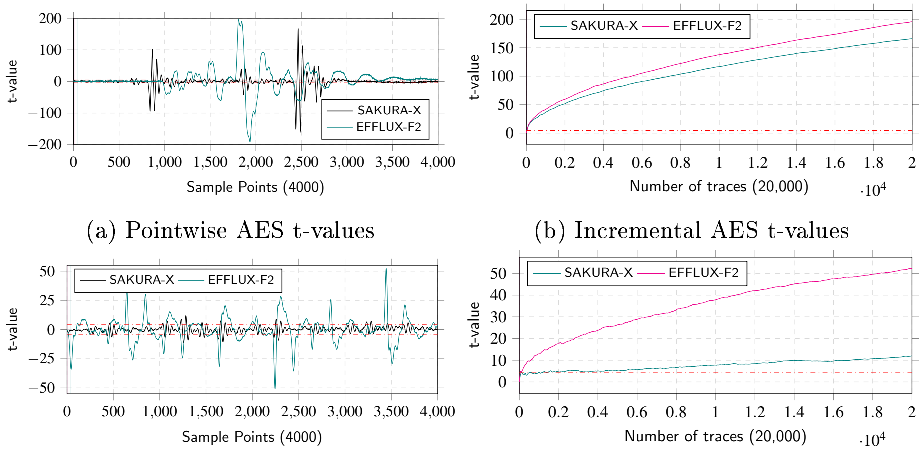

(a) Pointwise AES t-values

(b) Incremental AES t-values

(c) Pointwise GIFT t-values

(d) Incremental GIFT t-values

Fig. 12: TVLA results for unprotected implementations

In order to protect designs from side-channel attack, it has become common to integrate side-channel countermeasure such as masking, threshold implementation, etc. There has also been an effort towards development of side-channel resistant cryptographic ciphers. As TVLA analysis can detect leakage from any source, it is also used to evaluate whether a design is indeed side-channel resistant. If the countermeasures are not properly implemented then the design may still leak but at a much later point in time. Hence, it is a common practice to capture and analyze millions of traces for a protected design which typically requires few hours.

We also performed TVLA analysis of GIFT protected using a threshold countermeasure. To show the significance and performance characteristics of our low-noise board design, we utilized a known partially protected design of GIFT. This is done by intentionally removing a register layer between the decomposed S-boxes. The results are presented in Fig. 13. The t-values in the case of SAKURA-X at different sample points is mostly within the threshold as is expected from a fully protected design, but not from a partially protected design. It shows that the threshold has crossed only at a few sample points (somewhere around 250). Whereas, using EFFLUX-F2, it is quite prominent from multiple sample points that the design still leaks. One should also note the comparison results of incremental TVLA values from Fig. 13b. In case of SAKURA-X, the threshold is crossed only after 100,000 traces are captured and analyzed. Whereas, EFFLUX-F2 shows leakage even before 5000 traces have been analyzed, and the t-value increasing steadily. Thus, significantly reducing the leakage analysis time of a

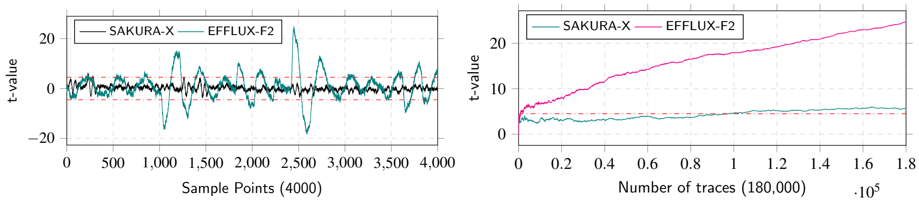

(a) Pointwise t-values

(b) Incremental TVLA

Fig. 13: TVLA results for partially protected GIFT

partially protected design from a few hours to a few minutes or less. The comparison results clearly shows that EFFLUX-F2 outshines SAKURA-X.

5 Conclusion

In this work, we present EFFLUX-F2 a SCA evaluation board designed with the targets of low noise and high SNR. These features leads to the reduction in the number of power traces required for experiments. With detailed experiments we show that the board provides significantly improved performance compared to the current platforms. Apart from the PCB design details, we also delve into the factors involved in achieving high SNR and discuss the reasoning behind the design choices. We show that EFFLUX-F2 has 4.5× lower noise and 8.2× higher SNR compared to SAKURA-X. We also show that EFFLUX-F2 required ≈20× fewer traces compared to the latter when analysing a protected leaky lightweight cipher implementation using TVLA.

Acknowledgement

This research is partially supported by the National Research Foundation, Singapore, and Cyber Security Agency of Singapore under its National Cybersecurity Research & Development Programme (Cyber-Hardware Forensic & Assurance Evaluation R&D Programme NRF2018NCRNCR009-0001). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the view of National Research Foundation, Singapore and Cyber Security Agency of Singapore.

References

1. Alagic, G., Apon, D., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Miller, C., Moody, D., Peralta, R., et al.: Status report on the third round of the nist post-quantum cryptography standardization process. US Department of Commerce, NIST (2022)

2. Banik, S., Pandey, S.K., Peyrin, T., Sasaki, Y., Sim, S.M., Todo, Y.: Gift: A small present: Towards reaching the limit of lightweight encryption. In: Cryptographic Hardware and Embedded Systems–CHES 2017: 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings. pp. 321–345. Springer (2017)
3. Becker, G., Cooper, J., DeMulder, E., Goodwill, G., Jaffe, J., Kenworthy, G., Kouzminov, T., Leiserson, A., Marson, M., Rohatgi, P., et al.: Test vector leakage assessment (tvla) methodology in practice. In: International Cryptographic Module Conference. vol. 1001, p. 13. sn (2013)
4. Bernstein, D.J., Lange, T.: Post-quantum cryptography. Nature **549**(7671), 188–194 (2017)
5. Brier, E., Clavier, C., Olivier, F.: Correlation power analysis with a leakage model. In: Cryptographic Hardware and Embedded Systems-CHES 2004: 6th International Workshop Cambridge, MA, USA, August 11-13, 2004. Proceedings 6. pp. 16–29. Springer (2004)
6. Fujimoto, D., Kim, Y., Hayashi, Y., Homma, N., Hashimoto, M., Sato, T., Danger, J.L.: Sasimi: Evaluation board for em information leakage from large scale cryptographic circuits. In: 2022 IEEE International Symposium on Electromagnetic Compatibility & Signal/Power Integrity (EMCSI). pp. 299–302. IEEE (2022)
7. Guilley, S., Maghrebi, H., Souissi, Y., Sauvage, L., Danger, J.L.: Quantifying the quality of side channel acquisitions. COSADE, February (2011)
8. 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 (2014)
9. Hori, Y., Katashita, T., Sasaki, A., Satoh, A.: Sasebo-giii: A hardware security evaluation board equipped with a 28-nm fpga. In: The 1st IEEE Global Conference on Consumer Electronics 2012. pp. 657–660. IEEE (2012)
10. Katashita, T., Hori, Y., Sakane, H., Satoh, A.: Side-channel attack standard evaluation board sasebo-w for smartcard testing. Power **3**(2012), 400 (2012)
11. Katashita, T., Satoh, A., Sugawara, T., Homma, N., Aoki, T.: Development of side-channel attack standard evaluation environment. In: 2009 European Conference on Circuit Theory and Design. pp. 403–408. IEEE (2009)
12. Mangard, S., Oswald, E., Popp, T.: Power analysis attacks: Revealing the secrets of smart cards, vol. 31. Springer Science & Business Media (2008)
13. McKay, K., Bassham, L., Sönmez Turan, M., Mouha, N.: Report on lightweight cryptography. Tech. rep., National Institute of Standards and Technology (2016)
14. O'flynn, C., Chen, Z.: Chipwhisperer: An open-source platform for hardware embedded security research. In: Constructive Side-Channel Analysis and Secure Design: 5th International Workshop, COSADE 2014, Paris, France, April 13-15, 2014. Revised Selected Papers 5. pp. 243–260. Springer (2014)
15. Satoh, A., Katashita, T., Sakane, H.: Secure implementation of cryptographic modules-development of a standard evaluation environment for side channel attacks. Synthesiology English edition **3**(1), 86–95 (2010)
16. Turan, M.S., McKay, K., Chang, D., Bassham, L.E., Kang, J., Waller, N.D., Kelsey, J.M., Hong, D.: Status report on the final round of the nist lightweight cryptography standardization process (2023)