Prototype IC with WDDL and Differential Routing – DPA Resistance Assessment

Kris Tiri[1], David Hwang[1], Alireza Hodjat[1], Bo-Cheng Lai[1], Shenglin Yang[1], Patrick Schaumont[1], and Ingrid Verbauwhede[1,2]

[1]Electrical Engineering Dept.
UC Los Angeles, USA
{tiri, ingrid}@ee.ucla.edu
[2]Dept. ESAT/SCD-COSIC
K.U.Leuven, Belgium

Abstract. Wave dynamic differential logic combined with differential routing is a working, practical technique to thwart side-channel power attacks. Measurement-based experimental results show that a differential power analysis attack on a prototype IC, fabricated in 0.18µm CMOS, does not disclose the entire secret key of the AES algorithm at 1,500,000 measurement acquisitions. This makes the attack de facto infeasible. The required number of measurements is larger than the lifetime of the secret key in most practical systems.

Keywords: side-channel attack (SCA), differential power analysis (DPA), countermeasure, dual rail with precharge, wave dynamic differential logic (WDDL), differential routing, parasitic capacitance matching

1. Introduction

A prototype IC has been fabricated in 0.18µm CMOS to demonstrate the secure digital design flow [11]. This design flow creates correct-by-construction side-channel power attack resistant integrated circuits. It starts from any HDL design and does not need custom layout, iterative design processes, or complex algorithm-specific countermeasures. It is based on employing logic cells with a single switching event per clock cycle and a place and route approach that balances the interconnect capacitance of the output wires.

Side-channel power attacks can be mounted on ASICs, FPGAs, DSPs and microprocessors because in standard CMOS technology, power is only drawn from the power supply when a 0 to 1 output transition occurs. Therefore, by measuring the power supply current during the encryption, and then performing statistical analysis of the measured power traces, the secret key can readily be determined. The secure digital design flow pursues a constant power dissipation by balancing the power consumption of the logic gate. When the power dissipation of the smallest building block is constant and independent of the signal activity, no information is leaked through the power supply. As a result, it protects against all power attacks including simple power analyses, differential power analyses and higher order power analyses.

Two functionally identical coprocessors have been fabricated on the same die. The first 'secure' coprocessor is implemented using WDDL and differential routing. The second 'insecure' coprocessor is implemented using regular standard cells and regular routing techniques. We fabricated two functionally identical coprocessors to allow us to compare the side-channel attack resistance of a typical IC versus one with special circuit techniques. Measurement-based experimental results show that a DPA attack on the insecure coprocessor requires only 8,000 measurements to disclose the entire 128-bit secret key. The same attack on the secure coprocessor still does not disclose the entire secret key at 1,500,000 measurements.

The remainder of the paper is organized as follows. The next section describes the prototype IC. It also discusses in brief the secure digital design flow and the architecture of the AES cryptographic engine. In section 3, our measurement setup is presented and an attack is mounted on the fabricated IC to assess the increase in DPA resistance of the secure coprocessor. This section also presents area, timing and power numbers. Section 4 presents related state-of-the-art. Finally, a conclusion will be formulated.

2. Prototype IC

The prototype IC, depicted in figure 1, consists of two functionally-identical coprocessors and is fabricated on the same die using a TSMC 6M 0.18µm process. An insecure coprocessor serving as benchmark is implemented using standard cells and regular routing techniques. A secure coprocessor is implemented through the secure

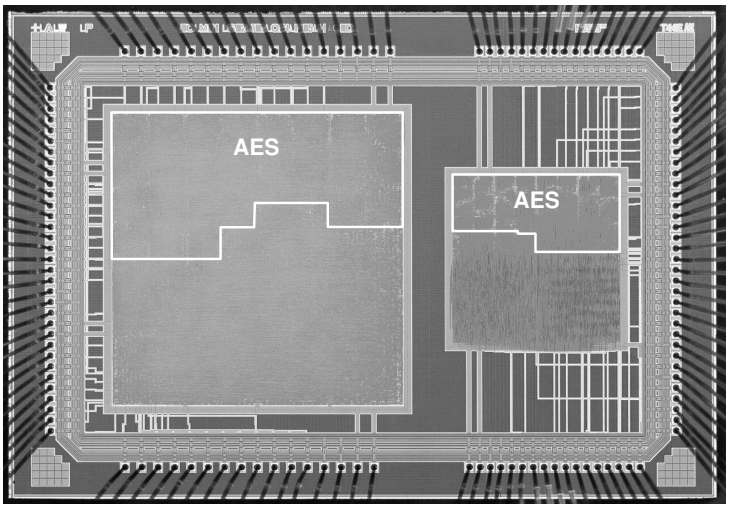

Fig. 1. IC micrograph: secure coprocessor using WDDL and differential routing (left); and insecure coprocessor using standard cells and regular routing (right)

digital design flow using WDDL and differential routing. Both coprocessors have been implemented starting from the same synthesized gate level netlist. The WDDL gates have been derived from the commercial static CMOS standard cell library used in the regular insecure design.

The IC, which is used for embedded cryptographic and biometric processing, consists of four components: an AES based cryptographic engine, a fingerprint-matching oracle, a template storage, and an interface unit. The coprocessor is part of a portable biometric and cryptographic authentication device that is called ThumbPod [5]. Architectural partitioning has been performed to divide the system into insecure (LEON SPARC V8 processor) and secure (coprocessor) modules, such that the processing and storage of all sensitive information is done on the secure module [6]. This ensures that the entire system does not need to be protected. Only the secure module must be protected for the system to remain secure.

2.1 Secure digital design flow

The secure digital design flow is completely supported by mainstream EDA tools and uses a commercially available static CMOS standard cell library. The differences with a regular synchronous CMOS standard cell design flow are minor. The secure digital design flow starts from a normal design in a hardware description language (HDL) and only a few key modifications are incorporated in the backend of the design flow. A cell substitution phase and an interconnect decomposition phase parse intermediate design files. The former procedure modifies the gate level description, the latter duplicates and translates the interconnect wires. The additional steps only required six minutes of CPU time for the prototype IC.

The design flow is based on a constant power dissipating logic: in one clock cycle the power consumption of each individual logic gate is constant and independent of its input signals. In other words, 0 to 0, 0 to 1, 1 to 0, and 1 to 1 output transitions all draw the same power from the supply. Two conditions must be satisfied to have constant power dissipating logic: a logic gate must have exactly one charging event per clock cycle; and the logic gate must charge a constant capacitance in that event.

Dynamic differential logic, also known as dual rail with precharge logic, has a single charging event per cycle. The design flow uses wave dynamic differential logic to implement dynamic differential behavior using static CMOS standard cells [13]. A WDDL gate consists of a parallel combination of two positive complementary gates. In the precharge phase, both true and false inputs are set to 0. This puts the output of the gate at 0. This 0 precharge value travels as the input to the next gate, creating a precharge ‘wave’. In the evaluation phase, each input signal is differential and the WDDL gate calculates a differential output. Special registers and input converters launch the precharge value. They produce an all-zero output in the precharge phase but let the differential signal through during the evaluation phase.

Besides a 100% switching factor, it is essential that a fixed amount of capacitance is charged during the transition. Thus, the total load at the true output of the differential gate should match the total load at the false output. With shrinking channel-length of the transistors, the interconnect capacitances have become the dominant capacitance. Hence, the issue of matching the interconnect capacitances of

the signal wires is crucial. The best strategy to achieve matched interconnect capacitances is differential routing [12]. The true and false output signals are routed at all times with parallel routes in adjacent tracks of the routing grid, on the same layers, and of the same length. Independent of the placement, the two routes have the same first order parasitic effects.

2.2 AES-based cryptographic engine

The cryptographic engine consists of an AES core with multiple modes of operation. The datapath is based on a single round of the AES-128 algorithm which consists of byte substitution, shift row, mix column, and key addition phases along with on-the-fly key scheduling in (see figure 2). Byte substitution is implemented using look-up tables. A full encryption of 128-bit data using a 128-bit key takes precisely eleven cycles. For a detailed discussion on the architecture, the reader is referred to [4].

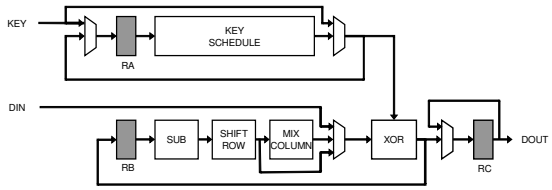

Fig. 2. Architecture of AES core

3. DPA resistance assessment

3.1 Measurement setup

The measurement and analysis setup is depicted in figure 3. The core supply current is measured between the PCB decoupling capacitances and the IC. A CT1 current probe from Tektronix [10] with a 25KHz to 1GHz bandwidth measures the supply current variations. For every mA, it provides 5mV output to the HP54542C oscilloscope [1]. The oscilloscope filters the waveform transients at 500MHz and digitizes with a 2GHz sampling frequency. With a standard GPIB interface, we have made up to 400 measurements a second, including data transfer. Such a setup only requires four minutes to make 100,000 power measurements. A ‘measurement’ refers here to multiple data points which are used as one acquisition in a side-channel attack.

To facilitate the synchronization of the measurements, we have access to the encryption start signal. A clock of 50MHz is provided to the coprocessor under attack. During the attack, only the AES core processes data. This means that for the attack on the insecure processor, the other circuits and modules are quiet, while for the attack on the secure processor, the other circuits and modules are constantly charging and precharging in the same manner.

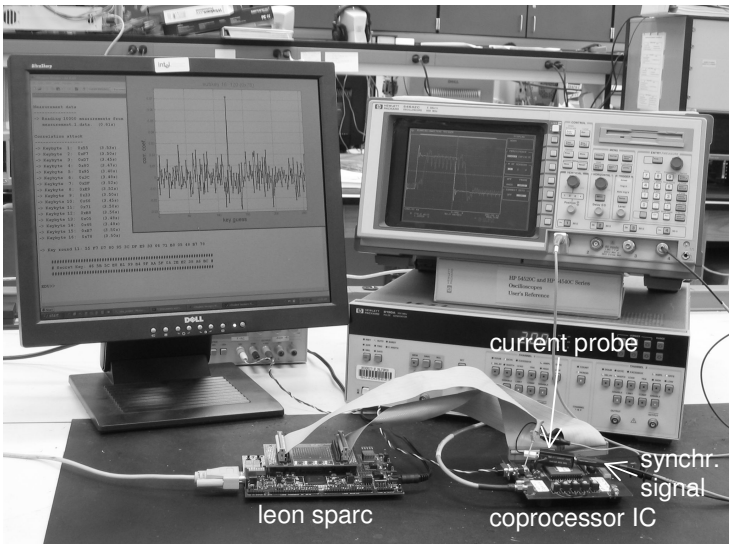

Fig. 3. DPA measurement and attack setup

Figure 4 shows the encryption start signal and the supply current of the AES coprocessors in output feedback (OFB) mode. The supply current of the insecure coprocessor exhibits large variations. It broadcasts the eleven encryption rounds and a high power peak exposes the starting point of each new encryption. The power consumption profile of the secure implementation on the other hand is invariant and does not reveal any information in a simple power analysis. In each clock cycle, nominally the same total load capacitance is charged and thus the same power is consumed, regardless of the operation been performed.

3.2 Differential power analysis

The DPA attack is performed as the AES core encrypts a plaintext P, using a key K, to produce a ciphertext C_{11} after eleven rounds. Note that the original K is broken up into different round keys (K_1 through K_{11}), where K_{11} is the round key for round eleven. Once K_{11} is deduced, it is easy to trace it back to find the original key K.

The influence of the datapath on the power consumption of the AES core is estimated through the Hamming distance of two successive values of register RB, (see figure 2) or in other words, through the number of changing state bits in a clock cycle. Most AES operations work with bytes and eight state bits can be calculated using a guess on one key byte. Using the same measured data, each of the sixteen bytes of the 128-bit key is cracked separately in the following manner.

We compare the estimations and the measurements with the correlation test [2]. The correct key guess is the one that results in the highest correlation coefficient

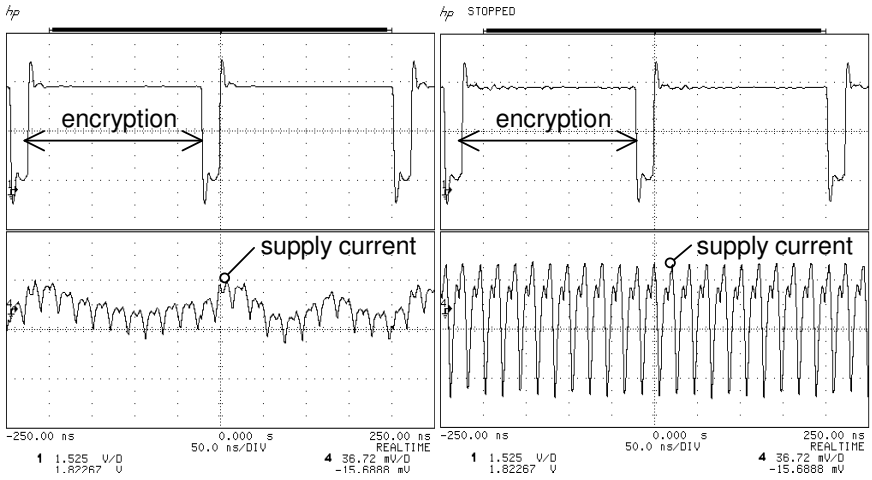

Fig. 4. Transient measurement (2 encryptions, 22 clock cycles) of encryption start signal (top) and core supply current (bottom): insecure coprocessor (left); and secure coprocessor (right)

between the vector of Hamming distances and the vector of representative measurements, for which we use the maximum supply current in a clock cycle.

We choose to attack register RB as it transitions from round eleven to the following round. As shown in figure 5, RB in round eleven (D_{11}) can be found by tracing back the signal obtained after xor-ing the final ciphertext (C_{11}) and a key guess (K_{11}) through both the shift row operation and the substitution box. RB in the next round, during which we perform the supply current measurement, is the known final ciphertext (C_{11}).

Each key byte (0 to 15) of K_{11} can be a value between 0x00 and 0xFF, for a total of 256 possibilities. Thus, for each key byte, there are 256 power estimations, one of which is the correct estimation. Of course the correlation may be inaccurate for only a few measurements (i.e., sets of $P_{measurement}$ and C_{11}). Hence thousands of different ($P_{measurement}$,C_{11}) pairs were measured using the same key (and hence the same K_{11}) in order to filter out the noise and provide a truthful correlation.

The correct key is found by evaluating:

$$\max_{K_{11}} f_{cost}(K_{11}) = \mathrm{corr}(P_{measurement}, P_{estimation}) \quad (1)$$

$$\textit{where}\ \ P_{measurement} = \max(I_{supply,11+1})$$
$$P_{estimation} = \mathrm{HamDist}(D_{11}, C_{11})$$
$$D_{11} = \mathrm{sub}^{-1}(\mathrm{shiftrow}^{-1}(K_{11} \otimes C_{11}))$$

For the secure design, we only need to look at one round, as all signals are at 0 at the start of the evaluation phase. The number of changing bits of RB in round eleven, during which we also carry out the measurements, is the Hamming weight of RB.

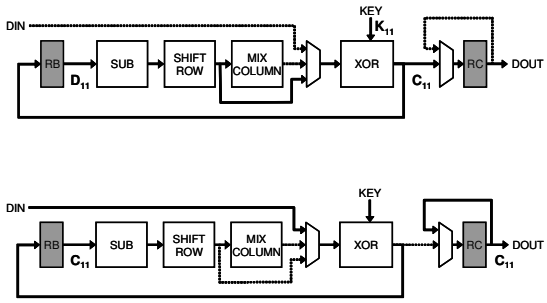

Fig. 5. AES core: round 11 (top); and round 11 + 1 (bottom)

Figure 6 shows the encryption start signal and the core supply current during the attack. The supply current of the insecure coprocessor reveals the encryption operation by showing exactly eleven peaks. The secure coprocessor has a continuous current whether or not data is being processed, either cryptographic or other. It has an identical power consumption profile in figures 4 and 6. If an attacker does not have access to the encryption start signal, it is almost impossible to know when the IC is encrypting.

For the actual attack, we only measure the round of interest. The dynamic range is set to cover the variation of the maximum current. The other irrelevant samples may be clipped. For the remainder of this manuscript, we will refer to the maximum value of one acquisition as the measurement.

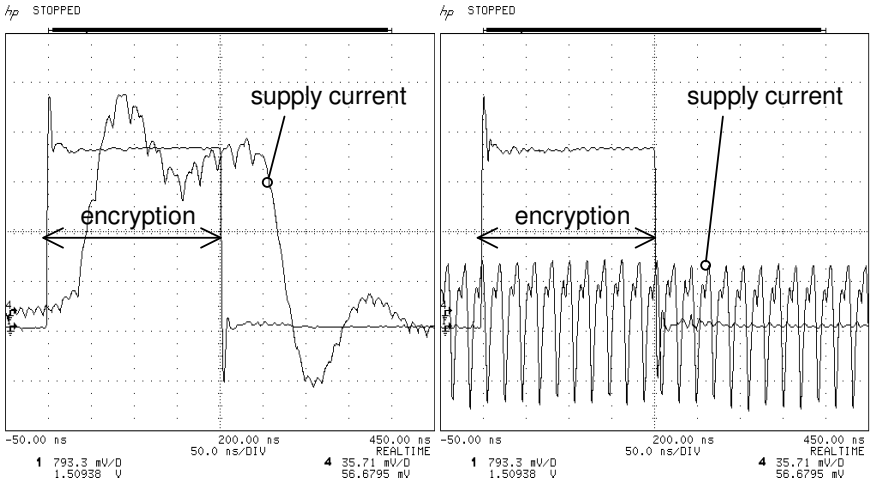

Fig. 6. Transient measurement of encryption start signal and core supply current for single encryption: insecure coprocessor (left); and secure coprocessor (right)

3.3 DPA resistance

The resistance against DPA is quantified with the number of measurements to disclosure (MTD). This number expresses how many measurements are necessary to correctly distinguish the correct secret key from all the other wrong key guesses.

We define MTD as the cross-over point between the correlation coefficient of the correct key and the maximum correlation coefficient of all the wrong keys guesses. For both coprocessors, an example of an attack on one key byte is shown in figure 7. MTD is depicted in the 'Correlation vs. Number of Measurements' graphs as the point where the black line (correct key) crosses the grey envelope (wrong keys). The results for the other fifteen key bytes are similar. The maximum number of measurements we took is 15,000 for the insecure coprocessor and 1,500,000 for the secure coprocessor.

For the insecure implementation, the correct key bytes are found very easily. On average, 2,000 measurements are required to disclose a key byte. In one case, a mere 320 samples were sufficient to mount a successful attack. There is also a large

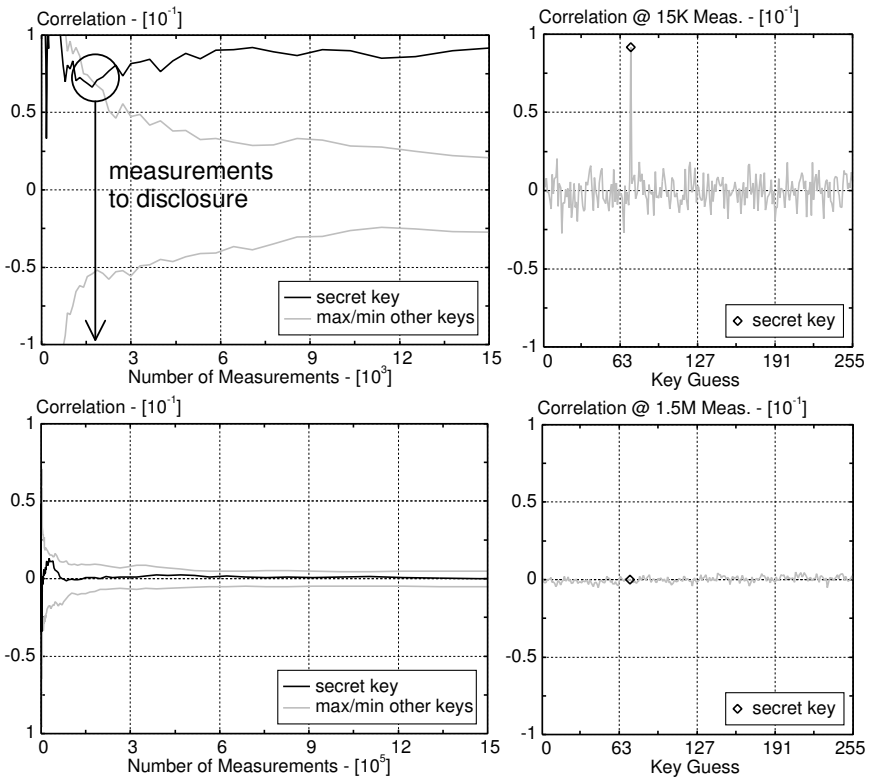

Fig. 7. Cracking the secret key: insecure coprocessor 15K measurements (top); and secure coprocessor using 1.5M measurements (bottom)

resolution. In the top right plot of figure 7, there is no doubt about the correct key guess.

The secure coprocessor on the other hand substantially reduces this resolution of correlation, as shown by the small correlation peaks in the 'Correlation vs. Key Guess' graph in figure 7. Our measurements show that out of sixteen keys bytes, WDDL effectively protects five key bytes. One and a half million measurements are not sufficient to disclose the correct key bytes. One example is shown on the bottom of figure 7. The eleven key bytes that are found require on average 255,000 measurements, an increase of more than two orders of magnitude when compared with the insecure coprocessor. One example is shown in figure 8.

The analysis also revealed that for a dual rail design, the correlation coefficient of the correct key guess can be negative, as shown in figure 8. This means that less power is consumed as more bits change. This implies that the 0 to 1 switching of the false net uses more power than the 0 to 1 switching of the true net. The parasitic capacitances affected by the false signals are larger than the ones affected by the true signals. On the other hand, for the five bytes that have not been found, the capacitances have an almost perfect matching between the differential nets. Hence it is crucial to guarantee matched capacitances consistently for all the logic.

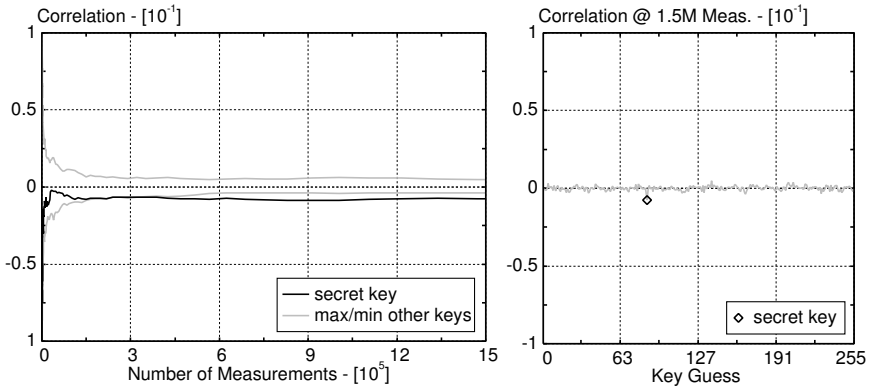

Fig. 8. Negative correlation coefficient

Shielding the differential routes on either side with a power line improves the matching. This eliminates the cross-talk to adjacent wires in the same metal layer. Alternatively, increasing the distance between different differential pairs would reduce the effect, or an improved iterative design flow could be used to identify and correct mismatches.

Table 1 summarizes the results. WDDL and differential routing is a functional technique to thwart power attacks. The trade-off is a three times increase in area, and a four times increase in power consumption and minimum clock period.

Table 1. IC results summary

Parameter	Unprotected AES	Protected AES
Gate Count (eq. gates) [K]	79	245
Area [mm^2]	0.79	2.45
Maximum Frequency (@1.8V) [MHz]	330.0	85.5*
Maximum Throughput (@1.8V) [Gb/s]	3.84	0.99
Power Consumption (@1.8V, 50 MHz) [mW]	54	200†
Measurements to Disclosure‡		
min	320	21,185
mean	2,133	255,391
max	8,168	1,276,186
Key bytes not found (@1.5M Meas.)	n/a	5

Duty factor of clock > 50% to guarantee precharge of all gates
†*Estimation based on area ratio AES vs. Entire System*
‡*Based on correctly guessed key bytes*

Recall from section 2 that security partitioning, the careful division of the architecture into two parts (a secure and a non-secure part) [6], minimizes the cost for complex systems. Only the relatively small part that processes sensitive information requires realization in a coprocessor with specialized logic and routing. This minimizes the area and reduces the power and time penalty. Even with these penalties, the secure coprocessor still runs orders of magnitude faster and expends less energy than a software implementation on the main processor.

The protected AES achieves a figure of merit, obtained by normalizing the throughput with the power consumption, of 2.9Gb/s/W. On the other hand, an unprotected AES implemented with C code on an embedded Sparc processor attains a mere 0.0011Gb/s/W (gcc, 1mW/MHz @120Mhz Sparc, 0.25µm CMOS). Research papers on algorithmic countermeasures unfortunately do not document the overhead in cycle count (and in byte code). Given a likely penalty of a factor 2 to 3, the figure of merit of the protected coprocessor is 4 orders of magnitude better than a software implementation of an algorithmic countermeasure on a microprocessor.

As future work, we foresee the need to explore the EMA resistance and the impact of 'noisy' regular components. Electromagnetic Analysis (EMA) is the equivalent of a power attack but instead uses the electromagnetic fields. The electromagnetic fields are generated by the (dis)charging gates. Since ideally each gate uses always the same amount of charge, a significant increase in EMA resistance is also expected. The impact of regular components, which process the insensitive data, causes a substantial increase in power variations. In figures 4 and 6, the maximum current has only minimal variations for the secure implementation while it has large variations for the regular implementation. Consequently, if both a secure and a regular component are present on an IC, the dynamic range of the measurements must be set to cover the maximum variation of the maximum current. As a result, the measurements of the side-channel information leaked by the secure module will be much less accurate and possibly within the quantization error. This analysis is possible with the prototype IC, as both the secure and the insecure processor can be operated at the same time.

4. Related work

To our knowledge, this work is the first to deliver and demonstrate a working, practical DPA countermeasure implemented and tested in actual silicon. All other published techniques have never been implemented in silicon, or have never been measured and attacked, or did not offer any significant DPA resistance.

A dual rail asynchronous chip has been presented previously [3]. The implementation, however, did not provide a significant increase in DPA resistance. This failure has been attributed to unbalanced signal paths caused by routing differences. Note that if asynchronous logic is used to increase the DPA resistance, dual rail encoded asynchronous logic must be used. Because of the dual rail logic, there is also a factor three area increase compared with a single ended synchronous benchmark [8].

Algorithmic countermeasures are mathematically DPA resistant. In practice, however, proposed solutions actually have been insecure [7]. To the best of our knowledge, published results of algorithmic countermeasures have never been successfully demonstrated on any platform with an actual measurement-based DPA resistance assessment. We are aware of one silicon implementation of an algorithmic countermeasure [9]. Measurements and assessment of the DPA resistance, however, have not yet been performed.

5. Conclusions

We have presented a secure coprocessor that does not leak information through the power supply, which is a major and easy to access side-channel leakage source. Built in a 0.18μm CMOS technology, we believe that this is the first IC that is practically immune to DPA attacks. Its immunity has been experimentally verified and compared to a second IC that is built with a regular standard cell approach. The design approach relies on WDDL, a logic style that has a single switching event per cycle, and differential routing, a place and route technique that controls the load capacitance. An actual power attack has been mounted on the IC to experimentally assess the increase in DPA resistance. Experimental results showed that 1,500,000 acquisitions are not sufficient to fully disclose the 128-bit secret key. This makes the attack de facto infeasible. The required number of measurements is larger than the lifetime of the secret key in most practical systems.

Acknowledgements. This work was supported in part by the National Science Foundation (CCR-0098361), UC-Micro 02-079, Panasonic Foundation, SUN Microsystems, Atmel corporation and the Fannie and John Hertz Foundation.

References

1. Agilent technologies, 54542C 4 Channel 2 GSa/s Color Digitizing Oscilloscope, http://www.home.agilent.com/USeng/nav/-536894779.536881118/pd.html.
2. J. Coron, P. Kocher, and D. Naccache, “Statistics and Secret Leakage”, Financial Cryptography (FC 2000), Lecture Notes in Computer Science, vol. 1962, pp. 157–173, February 2000.
3. J. Fournier, S. Moore, H. Li, R. Mullins and G. Taylor, “Security Evaluation of Asynchronous Circuits,” Cryptographic Hardware and Embedded Systems (CHES 2003), Lecture Notes in Computer Science, vol. 2779, pp. 137–151, September 2003.
4. A. Hodjat, D. Hwang, B. Lai, K. Tiri, and I. Verbauwhede “A 3.84 Gbits/s AES Crypto Coprocessor with Modes of Operation in a 0.18-µm CMOS Technology”, accepted at Great Lakes Symposium on VLSI (GLSVLSI 2005), April 2005.
5. D. Hwang, P. Schaumont, Y. Fan, A. Hodjat, B. Lai, K. Sakiyama, S. Yang, and I. Verbauwhede, “Design flow for HW/SW acceleration transparency in the ThumbPod secure embedded system”, 40th Design Automation Conference (DAC 2003), pp. 60-65, June 2003.
6. D. Hwang, P. Schaumont, K. Tiri, and I. Verbauwhede, “Making embedded systems secure”, accepted at IEEE Security & Privacy Magazine.
7. S. Mangard, T. Popp, and B. Gammel, “Side-Channel Leakage of Masked CMOS Gates”, Cryptographers’ Track - RSA Conference (CT-RSA 2005), pp. 351-365, February 2005.
8. S. Moore, R. Anderson, R. Mullins, and G. Taylor, “Balanced Self-Checking Asynchronous Logic for Smart Card Applications,” Journal of Microprocessors and Microsystems, vol. 27, pp. 421–430, 2003.
9. N. Pramstaller, F. Gürkaynak, S. Häne, H. Kaeslin, N. Felber, and W. Fichtner, “Towards an AES Crypto-chip Resistant to Differential Power Analysis”, 30th European Solid-State Circuits Conference (ESSCIRC 2004), pp. 307-310, September 2004.
10. Tektronix, CT1 current probe, http://www.tek.com/site/ps/60-12572/pdfs/60W_12572.pdf.
11. K. Tiri, and I. Verbauwhede, "A VLSI Design Flow for Secure Side-Channel Attack Resistant ICs", accepted at Design, Automation and Test in Europe Conference (DATE 2005), March 2005.
12. K. Tiri, and I. Verbauwhede, "Place and Route for Secure Standard Cell Design", 6th International Conference on Smart Card Research and Advanced Applications (CARDIS 2004), pp. 143-158, August 2004.
13. K. Tiri, and I. Verbauwhede, "A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation", Design, Automation and Test in Europe Conference (DATE 2004), pp. 246-251, February 2004.