HAL
open science

Silicon-level Solutions to Counteract Passive and Active Attacks

Sylvain Guilley, Laurent Sauvage, Jean-Luc Danger, Nidhal Selmane, Renaud Pacalet

To cite this version:

Sylvain Guilley, Laurent Sauvage, Jean-Luc Danger, Nidhal Selmane, Renaud Pacalet. Silicon-level Solutions to Counteract Passive and Active Attacks. FDTC, Aug 2008, Washington, DC, United States. pp.3-17, 10.1109/FDTC.2008.18 . hal-00311431

HAL Id: hal-00311431

https://hal.science/hal-00311431v1

Submitted on 17 Aug 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Silicon-level Solutions to Counteract Passive and Active Attacks

Sylvain GUILLEY, Laurent SAUVAGE, Jean-Luc DANGER, Nidhal SELMANE and Renaud PACALET

Institut TELECOM, TELECOM ParisTech
CNRS LTCI (UMR 5141)
Département COMELEC, 46 rue Barrault
75 634 PARIS Cedex 13, FRANCE
Email: < sylvain.guilley@telecom-paristech.fr >

Abstract

This article presents a family of cryptographic ASICs, called SecMat, designed in CMOS 130 nanometer technology by the authors with the help of STMicroelectronics. The purpose of these prototype circuits is to experience with the published "implementation-level" attacks (SPA, DPA, EMA, templates, DFA). We report our conclusions about the practicability of these attacks: which ones are the most simple to mount, and which ones require more skill, time, equipments, etc. The potential of FPGAs as security evaluation commodities at design time is also detailed. Then, we discuss about "dual counter-measures", that are meant to resist both passive and active attacks. This study started four years ago with TIMA (Grenoble), in the framework of the project MARS [31]. We highlight some research directions towards dependable and cost-effective dual counter-measures.

Keywords: side-channel attacks (SCA), differential power analysis (DPA), SecMat ASIC family, dual-rail with precharge logic (DPL), SecLib DPL style, differential fault attack (DFA), FPGA as evaluation platforms, attacks mitigation techniques, dual DPA-DFA counter-measures.

1 Introduction

Since the seminal publication of Paul Kocher on the differential power analysis (DPA [36]), a research community committed to study attacks carried out directly on hardware devices has built up. International conferences devoted to hardware security have been created: IACR **CHES** [49] in 1999, IEEE **FDTC** [50] in 2004 and IEEE **HOST** [51] in 2008. The major threat tackled by the academia has been the retrieval of secret information from a cryptographic device. The canonical attack scenario consists in extracting a key from a block cipher (for instance DES [34] or its successor AES [35]) while it is encrypting. Other attacks have been devised, that target specific algorithms or devices. However, for the sake of generality, we restrict our scope to the basic abovementioned scenario.

Attacks on DES or AES can be divided into two categories, depending whether the hardware device under analysis is only **passively** observed or whether its functional behavior is **actively** altered. It is customary to refer to them respectively as side-channel attacks (SCA) and fault injection attacks. They both are **key recovery** attacks.

Common side-channel attacks record physical emanations of a device, such as its execution time, instant power consumption or emitted radiations. The differential power analysis (DPA [36]) and its avatars, such as the correlation power analysis (CPA [6]) or the partitioning power analysis (PPA [30]) attempt to correlate these collected "leakages" with an assumed internal "power model". As in unprotected or in poorly protected devices the power model depends on the cryptographic key, the very power model that has the best correlation with the leakages observation betrays the correct key concealed in the device. Similar attacks using radiated (magnetic and/or electric) fields instead of conducted emanations have also been described; they are referred to as electromagnetic analyses (EMA [14]). Apart from some subtleties related to the spectral nature of the EM emanations [2], the rationale for the EMA attacks resemble that of the DPA. The simple power attack (SPA [36]) and the "template attacks" [7, 1] work without any physical model. Instead, they consist in matching the observations with either a known (as for SPA) or pre-characterized (as for templates) data base. In the later case, they are also qualified of "blind attacks" [17]. Again, a correct matching discloses all or some partial information on the key.

The principle of fault attacks is to cause the device to malfunction, so as to gain some information from a faulty cryptogram. At first glance, it might seem contradictory to retrieve correct information from incorrect encryption re-

sults. The working factor of most fault attacks is to mobilize differential analysis [4] on only the few last rounds: the simultaneous knowledge of correct and incorrect cryptograms, associated with a fair intuition of the fault location (in value, time or space), can indeed be used to discard some hypotheses made on the cryptographic key candidates. It turns out that some differential fault attacks (DFA) are extremely powerful. As an unprecedented example, Gilles Piret shows that only two well-behaved faults can suffice, under some circumstances, to disclose the integrality of an AES-128 key [37].

These attacks can be carried experimentally on devices without even touching the cryptographic chip. According to the terminology introduced by Sergei Skorobogatov [42], we qualify them as **non-invasive** attacks. Some SCAs can be improved if the naked chip surface can be illuminated accurately, say by a focused continuous or pulsed laser beam [43]. Similarly, some DFAs can be surgically efficient under the same conditions, referred to as semi-invasive. During the collaboration with STMicroelectronics, we have focused on non-invasive attacks. Indeed, they are obviously the first ones that a prospective attacker will use to break a device, and it is thus a necessary condition to first understand and resist them.

The rest of the paper is organized as follows. Section 2 imparts the feedback gained during SCA campaigns led on the prototype ASICs called SecMat, and discusses the relative harmfulness of these "observation" attacks. Non-invasive DFAs on SecMat are reported in Sec. 3. Section 4 provides perspectives in factoring counter-measures against various threats. Finally, conclusions are drawn in the Sec. 5.

2 Side-Channel Attacks

2.1 SecMat Family ASIC as Security Evaluation Platforms

The SecMat ASICs have been manufactured to compare the relative strength of implementation-level attacks. They use HCMOS9GP, the general purpose CMOS 130 nanometer process from STMicroelectronics (fab located at Crolles, France), with the low leakage option for the transistors. The actual fabrication has been managed by the "CMP" multi-wafer projects broker. With respect to the state-of-the-art in the smartcard industry, the HCMOS9GP technology is advanced, *i.e.* very deep-submicronic. This choice is deliberate; it makes sure that the security evaluations carried out on SecMat chips to remain representative for a couple of years, after some technological nodes have elapsed.

The complete SecMat family is shown in Fig. 1. The four circuits are respectively nicknamed SecMat v1, v$^3/_2$, v2 and v3. SecMat v$\{1,2,3\}$ have a fully-fledged ISO 7816 smartcard architecture, whereas SecMat v$^3/_2$ is a proof-of-concept ASIC designed to evaluate the relative leakage-freedom of several logic styles and backend-level balance strategies. They all embed facilities to ease power measurement and analysis. SecMat v1 has one supply (couple of `vdd`/`vss` ports) per cryptoprocessor. The advantage is that the signal can be recorded with high fidelity by an oscilloscope. The disadvantage is that, when comparing various co-processors, the acquisition conditions are not exactly the same: the measurement probe must be manually displaced from one supply to another on the test board. Now, the SecMat v$^3/_2$, v2 and v3 are designed to enable comparisons between various implementations. Therefore, we opted for a different access to the side-channel: there is only one couple of supply pads `vdd`/`vss`, mutualized for all the modules to test. The chip's driver either does not use modules others than the one under test (as in SecMat v$^3/_2$) or programs a power management module to clock gate the irrelevant modules (as in SecMat v2 and v3). With this solution, the traces loose fidelity, because the power network is all the same larger since distributed, and thus parasitically cross-coupled to other active parts of the circuit. The execution of a typical DES encryption using both side-channel measurement strategies (SecMat v1 *versus* v3) is depicted in Fig. 2. Let apart the fact the signal is lower for SecMat v3 (for multiple reasons: novel architecture of DES, different PCB and probes, *etc.*), it is in addition clearly more noisy. This is the cost we accept to pay to guarantee fair comparisons. However, as discussed below, the loss is not that detrimental to the analyses, that continue to succeed, albeit with a need for slightly more traces to disclose the full key. Anyway, the power measurements on SecMat ASICs are of a great quality, much superior to that obtainable in practice on real tamper-proof circuits. This apparent superiority merely models a purportedly empowered attacker, and additionally enables sound interpretations of DPA-like attacks. To further facilitate the attacker's analyses, all the SecMat circuits provide an unambiguous synchronization signal (to properly trigger the acquisition apparatus).

The SecMat v1 experimental circuit is designed to validate counter-measures against the DPA. The overall architecture is a bus-centric system-on-chip (SoC). Standardized modules, implementing the Advanced-VCI interface, are plugged together onto a fixed priority bus mastered by an 8-bit 6502 CISC micro-processor. The processor boots a "monitor" from an embedded 2kb ROM and loads its program from the outside through an UART (connected to a serial RS232 or a USB cable) into a embedded 32kb RAM. The main feature of the chip is the activation of its four cryptographic accelerators – one AES and three DES – to lead SCA campaigns. The rationale of classical attacks has been studied on this ASIC, and the lessons we learnt are summarized in the subsections 2.1.1, 2.1.2, 2.1.3, 2.1.4 and 2.1.5.

SecMat v1	**SecMat v3/2**, alias SubBytes (13 instances of the AES [35] direct byte-wise substitution box)
⇒ CMP run `S12C5_1` (13/01/2005) ⇒ 4.0 mm^2, 2.0 million transistors	⇒ CMP run `S12C6_4` (10/11/2006) ⇒ 1.0 mm^2, 1.0 million transistors
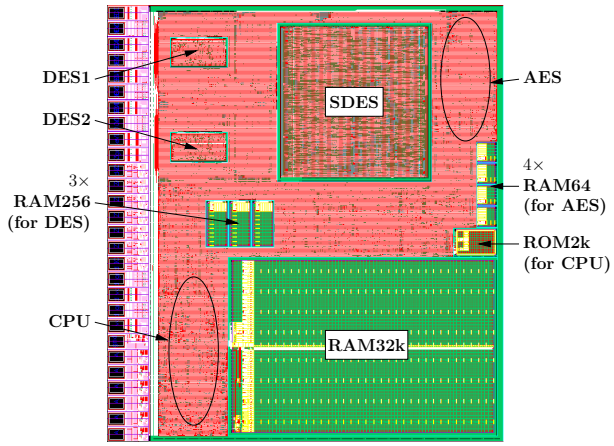	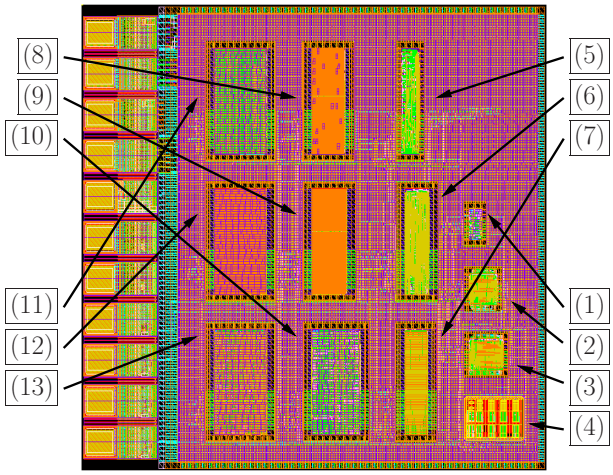
SecMat v2	**SecMat v3** [8, 18]
⇒ CMP run `S12C6_1` (17/02/2006) ⇒ 4.3 mm^2, 2.3 million transistors	⇒ CMP run `S12C7_1` (03/01/2007) ⇒ 4.4 mm^2, 2.4 million transistors
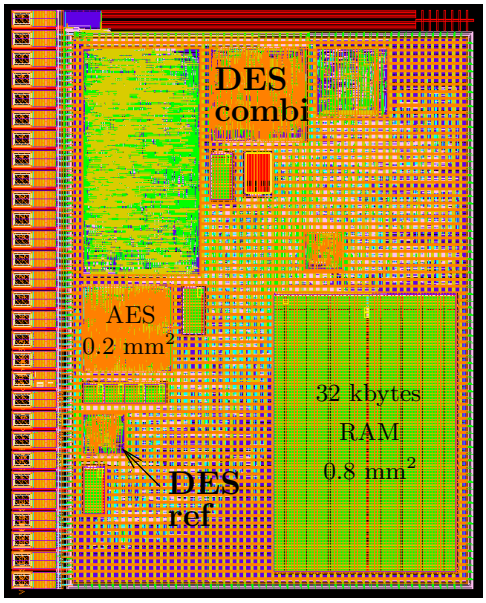	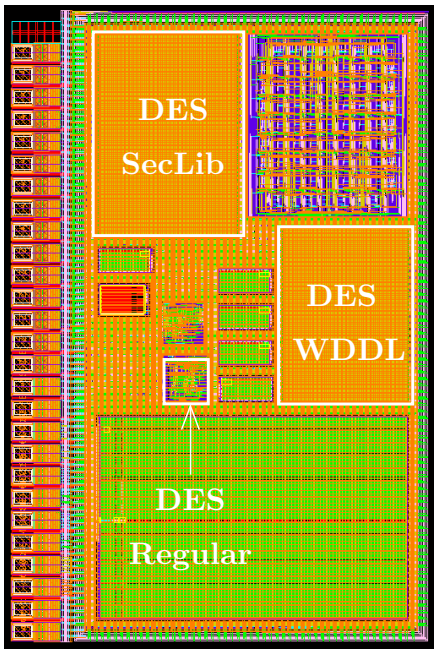

Figure 1. Run information, dimensions (in terms of die area & transistor count) and layout of the four 130 nanometer ASICs forming the "SecMat" family.

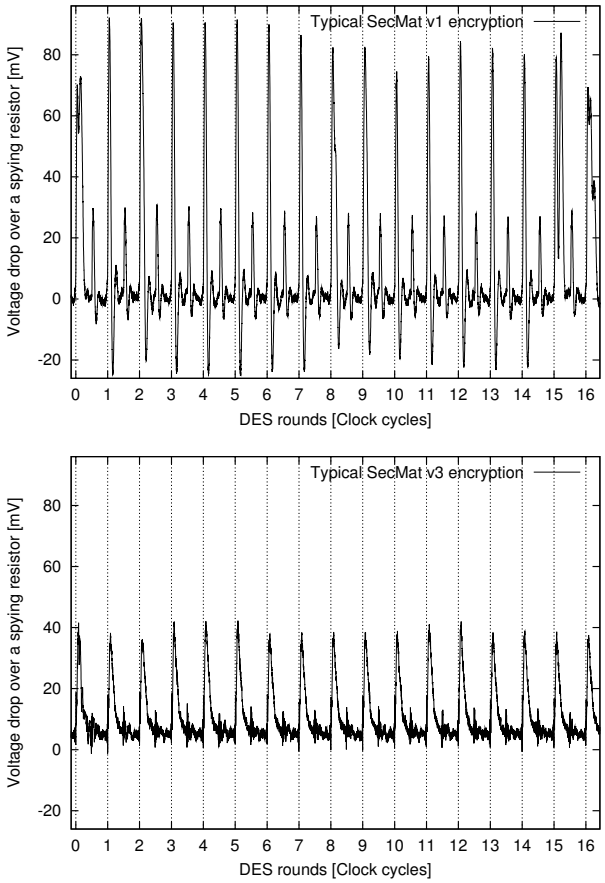

Figure 2. Typical DES encryption power curve for SecMat v1 (*top*) and v3 (*bottom*).

2.1.1 DPA as Boolean Oracle

SecMat v1 is a sandbox for plain attacks, typically the DPA. In an acquisition campaign, m power curves (also casually called "traces") $\mathbf{T}_i$ are garnered. Each trace corresponds to a given leakage. For the sake of simplicity, we focus on a part of the leakage, related to a very restricted set of gates in the netlist. We assume this leakage can be summarized into two typical behaviors:

- high dissipation to charge a net *versus* low dissipation to discharge it (referred to as the Hamming weight model),
- activity *versus* no-activity (referred to as the Hamming distance model), *etc*.

These two behaviors are related to an inner Boolean variable, called a decision function D_i. If the secret key is unknown, this variable is also unknown; however, it depends in practice in only a couple of bits from the key, which can be tested exhaustively. The decision function $D_i \in \{0, 1\}$ is an oracle for an attacker: if it is the correct decision function, it can allow for an exhibition of the two behaviors, otherwise it is simply irrelevant.

Amongst the many oracles that have been proposed, we focus on three of them, noted DPA_{diff}, DPA_{cov} and CPA, defined in equations (1), (2) and (3).

The idea behind the DPA_{diff} is to exhibit an asymptotic difference between the behaviors. The *ad hoc* criterion introduced by Paul Kocher [36] is:

$$\text{DPA}_{\text{diff}} \doteq \frac{1}{m_0} \sum_{i/D_i=0} \mathbf{T}_i - \frac{1}{m_1} \sum_{i/D_i=1} \mathbf{T}_i \,, \qquad (1)$$

where m_0 and m_1 denote the number of traces for each decision. More specifically, $m_0 \doteq \#\{i \in [0, m[/D_i = 0\}$ and, symmetrically, $m_1 \doteq \sum_{i=0}^{m-1} D_i$, with the following complementation property $m_0 + m_1 = m$.

A seemingly different approach consists in computing a covariance between the m traces and their associated decision functions. The DPA covariance estimator is:

$$\text{DPA}_{\text{cov}} \doteq \frac{1}{m} \sum_i \mathbf{T}_i \times D_i - \frac{1}{m} \sum_i \mathbf{T}_i \times \frac{1}{m} \sum_i D_i \,. \quad (2)$$

This estimator extracts the leakage linked to the Boolean variable D_i from the overall traces, seen as the superimposition of myriads of leaking elements (all the nodes of the netlist, in fact) [23].

This definition of the DPA actually coincides with the previous one, as far as the decision function is balanced, which is the case when random messages are processed or when the attack is carried out on the last round (cryptograms are exemplarily equidistributed).

Proof. Assuming that $m_0 = m_1 = m/2$,

$$\begin{aligned}
\text{DPA}_{\text{cov}} &= \frac{1}{m}\sum_i \mathbf{T}_i \times \left(D_i - \frac{1}{2}\right) \\
&= \frac{1}{2m}\sum_i \mathbf{T}_i \times (-1)^{D_i} \quad \left\{ \begin{array}{l} \text{Covariance with} \\ \text{the character} \\ \text{function of } D. \end{array} \right. \\
&= \frac{1}{4}\,\text{DPA}_{\text{diff}}\,.
\end{aligned}$$

□

Similarly, we can prove that even if $m_0 \neq m_1$, the two DPA concepts are linked, by the following relationship:

$$\text{DPA}_{\text{cov}} = \frac{m_0 \cdot m_1}{m^2}\,\text{DPA}_{\text{diff}}\,.$$

The multiplicative factor $m_0 \cdot m_1/m^2$ depends on the key hypothesis. Therefore, the two approaches do not lead to identical attack results. However, we notice that asymptotically $m_0 \cdot m_1/m^2 \to 1/4$, which reunifies them.

As the decision function D is *a priori* unknown, a wise tactic to limit the number of possible candidates is to select it either close to the first or close to the last encryption round. This way, D depends only on a few bits of the key. In block ciphers, be them Feistel schemes (*e.g.* DES [34]) or substitution-permutation networks (*e.g.* AES [35]), the decision function involves only a couple of secret key bits, equal to the substitution boxes (sboxes) entrance. Therefore, DES (resp. AES) implies 2^6 (resp. 2^8) key guesses per sbox. Now, when guessing a key for a given sbox, the predicted bits are in fact the sortance of the targeted sbox. If the sbox is multi-bit (4 bits as for DES or 8 bits as for AES), multi-bit DPAs can be devised. As mono-bit DPA insulates the dissipation from one net, we can deduce that DPA is **extensive**. This means that the contribution of a whole is equal to the arithmetic sum of its components. The Boolean decision function D can therefore be advantageously replaced by a vectorial Boolean decision function. This option expresses naturally for DPA_{cov} as the correlation between the power curves and the Hamming weight of D. An illustration of this additivity property is shown in Fig. 4 on the example of the 32-bit left datapath register, named L in the standard, of DES in SecMat v1. The multi-bit selection functions (in the Hamming distance model) are respectively:

$$\sum_{n=1}^{N} \text{L}_{\text{r-1}}[n] \oplus \text{L}_{\text{r}}[n] = |(\text{L}_{\text{r-1}} \oplus \text{L}_{\text{r}})[1:N]|\ ,$$

at round $\text{r} = 3$ for a number of bits $N = 8, 16, 24$ or 32. The graph clearly shows the extensivity of the multi-bit DPA, which confirms its ability to simultaneously extract unrelated bitwise covariances. This proves that DPA_{cov} is more likely to handle cases where the sbox is implemented in hardware, where all the output bits are computed concomitantly.

By definition [6], CPA is a normalization of the DPA. Although both DPA_{diff} and DPA_{cov} could used, only DPA_{cov} is mentioned in the literature. It is defined as a correlation coefficient, estimated by:

$$\text{CPA} \doteq \frac{\text{DPA}_{\text{cov}}}{\sigma_{\mathbf{T}} \cdot \sigma_D} \in [-1, +1] \subset \mathbb{R}\,, \tag{3}$$

where σ_X is the standard deviation of the random variable X, for which an unbiased empirical estimator is $\sqrt{\frac{1}{m-1}\sum_{i=0}^{m-1}\left(X_i - \frac{1}{m}\sum_{j=0}^{m-1} X_j\right)^2}$. Notice that because of the term σ_D, the CPA is not linear w.r.t. the decision function D. Although for the DPA_{cov}, if $D = D_1 + D_2$ then $\text{DPA}_{\text{cov}}(D) = \text{DPA}_{\text{cov}}(D_1) + \text{DPA}_{\text{cov}}(D_2)$, for the CPA this property does not hold any longer: $\text{CPA}(D_1 + D_2) \neq \text{CPA}(D_1) + \text{CPA}(D_2)$. Thus, if DPA_{cov} is unequivocally a leakage extraction tool, the interpretation to give to CPA is less intuitive. This remark does not presume of their efficiency, discussed in the next subsection 2.1.2.

In the sequel, the term DPA is employed in lieu of the sole multi-bit variant, namely DPA_{cov}. Also, DPA and CPA are tacitly understood as multi-bit.

2.1.2 DPA *versus* CPA

We have investigated why CPA [6] concretely performs better than DPA [22]. For this purpose, we compare the two differential curves obtained after $m = 1,000$ traces estimation of covariance (leading to DPA) and correlation (leading to CPA). The result is shown for the DES first substitution box of SecMat v1 in Fig. 3. We find that CPA outreaches DPA not because peaks appear with fewer traces, because it reduces the noisy peaks better than DPA. In the example of Fig. 3, the encryption starts at the clock period labeled 16. We therefore expect covariance and correlation peaks to build up around clock cycles 16 or 17. For instance, in the DPA differential trace on top of Fig. 3, the peak at clock period 38 is irrelevant because definitely outside of the correlation area. After accumulation with more traces, it vanishes, whereas those at clock cycles 16 and 17 remain. Notice that if the analysis window is very well focused on the samples corresponding to the maximum correlation, the DPA is always better, because it is not handicapped by spurious noisy peaks at the early stages of covariance estimation (*i.e.* when m is small).

2.1.3 DPA *versus* EMA

Given that the spatial extension of the SecMat chips is tiny (millimetric), a precise cartography [39] using electromagnetic analysis (EMA [2]) is completely illusory. For deep

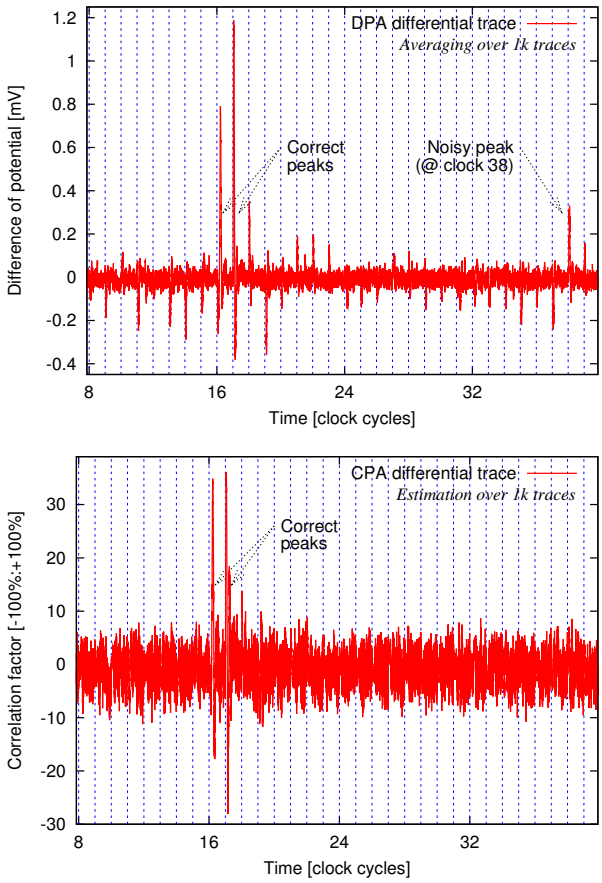

Figure 3. Comparison between a DPA (*top*) and a CPA (*bottom*) differential curve on the same 1k (*i.e.* very few) traces collected from SecMat v1.

Table 1. Ad hoc sensors fabricated for EMA analysis exploration on SecMat v1.

Stranded copper core denuded coaxial cable (1)	PCB coil (2)	Loop coil (3)
Field: $\vec{E}$	Field: $\vec{H}$	Field: $\vec{H}$

submicron technologies, EMA (and its actual signal processing variants: SEMA and CEMA) are to be considered as global leakage interception from the targeted device. Therefore, we carried out acquisitions on the decoupling capacitor located on the PCB, aside SecMat v1 itself, using a denuded coaxial cable and two home-made magnetic sensors as fortune antennas. The dimensions of our antennas are shown in Tab. 1. The antenna (1) is detuned, as its central passing frequency is 2 GHz. On the contrary, the antennas (2) and (3) are adapted at 30 MHz, which corresponds to the clock frequency of SecMat v1. Antennas (2) and (3) cutoff frequencies are respectively at 70 and 100 MHz.

This lightweight and totally non-invasive setup allowed us to reproduce the DPA/CPA experiments, with a key retrieval time comparable to power-line attacks. The number of measurements to disclose the last round key (subkey of the first sbox and all the eight sboxes) of DES with CPA is given in Tab. 2.

From these results, we draw some conclusions. On the one hand, power-line attacks are less noisy than EM attacks. This is due to the fact that we first pass the EM signal through an *imperfect* large band (9 kHz – 1 GHz) 32 dB amplifier and that we do not confine the experiment within a Faraday cage. This is typically shown by the fact that $1\times$ or $64\times$ averaged power curves yield roughly to the same key retrieval speed. With EM signals, the direct averaging by the oscilloscope greatly improves the signal quality, as shown by the need for a $256\times$ averaging for all EM mea-

Table 2. Number of measurements to disclose the key for various power and EM sensors.

Sensor	Averaging	Sbox #1	All sboxes
50 Ω **resistor**	1×	176	940
50 Ω **resistor**	64×	231	518
Antenna (1)	256×	1,695	1,883
Antenna (2)	256×	1,066	1,786
Antenna (3)	256×	707	1,008

surement campaigns. On the other hand, EM acquisitions are totally non invasive, whereas power-line measurements require the insertion of a 50 Ω SMC resistor in series with the positive supply.

Regarding EM sensors, measuring the $\vec{H}$ field with antenna (3) is more relevant than measuring the $\vec{E}$ field with antenna (1). The more loops in magnetic captor, the better the acquisition: antenna (2) with 3 loops is less efficient than antenna (3) which has 4 loops.

2.1.4 DPA as Abstract Variable Tracer

In the covariance graph of Fig. 4, there are two distinct peaks at two successive clock periods. In DES as in any Feistel schemes implemented iteratively, transitions actually occur in the right-hand register (R) one clock cycle before they manifest by a power signature in the left-hand register (L). More formally, $\forall r \in [1, 16], \mathrm{R}_{r-1} = \mathrm{L}_r$, which implies that the decision function $|\mathrm{L}_{r-1} \oplus \mathrm{L}_r|$ is equal to $|\mathrm{R}_{r-2} \oplus \mathrm{R}_{r-1}|$ by chance. This accounts for the presence of twin peaks, and also explains why one is anticipated by one clock period.

This teaches us that DPA insulates abstract variables, *wherever* and *whenever* they are manipulated while percolating through the circuit. Put differently, DPA is not a tool to extract the activity of a precise gate at a desired moment. Instead, the decision function is to be apprehended as a delocalized and intemporal wave-function over the circuit.

Nonetheless, if the decision function is defined in terms of **abstract** Boolean variables, the extractions themselves translate a **concrete** leakage. The leakage is indeed a valued quantity with a physical unit. It can be expressed in millivolts for instance, when the measurement consists in observing a voltage drop across a spying resistor artificially inserted in series with the power supply. The voltage is the transduction of the device inner mechanism to create and destroy the abstract variable(s) the decision function wishes to capture. What really happens within the circuit with the abstract variable(s) neither needs to be understood nor even known. The SCA relies only on the access to the side-channel, an accessible conversion of the mysterious internal information processing. To illustrate this point, we notice that the flip-flops (DFFs) that make up R and L register banks are the same, namely 32 instances of the standard cell `HCMOS9GP:CORE9GPLL:FD2QLL`. Nonetheless, the two peaks identified previously as activity in respectively R and L are of different height. The reason behind this noting can be at least twofold: despite R and L being the same gates, their environment might differ, making R more dissipative than L. Additionally, the abstract variable $\mathrm{R}_{r-2} \oplus \mathrm{R}_{r-1}$ can be used by more instances than only the register R: buffers, downstream registers, *etc.* The two hypotheses are plausible, given the structure of DES: at one and the same time, R is more loaded than L, and R drives a larger combinatorial logic cone than L.

2.1.5 DPA as Underlying RTL Architecture Recovery

We notably discovered how to use DPA for the reverse-engineering of the algorithm scheduling. When the key is known (chosen or not), it becomes possible to get rid of the hypothesis testing step: correct correlations can be computed directly. In the two graphs making up the Fig. 5, the DPA correlations with $|L_3 \oplus L_4|$ (resp. $|L_4 \oplus L_5|$) are computed. On the two resulting differential curves, one can notice that a similar extraction of that displayed in Fig. 4, except that it is offset by one (resp. two) clock periods. This proves that the DES architecture is iterative and that one round is executed every clock period. Incidentally, this technique can be generalized to enable arbitrary register transfer level (RTL) structure recovery. Some pioneering works from Christophe Clavier [11] illustrate this tremendous potential for SCA techniques. Incidentally, Eli Biham and Adi Shamir explain in [4, §4] how to take advantage of fault attacks (see Sec. 3) to reconstruct unknown ciphers.

2.2 SecMat Counter-Measures against SCAs

We have studied two types of counter-measures against SCAs while specifying the SecMat cryptographic accelerators. The first fully combinatorial encryption algorithm reported publicly is implemented in SecMat v2, as explained in Sec. 2.2.1. Power-constant logic CMOS styles are experimented in SecMat v3, as discussed in Sec. 2.2.2. Globally, our conclusion is that "*it is easy to make SCAs difficult*". Moreover, the counter-measure we propose are not necessarily costly in terms of power consumption. However, for low-end and low-profit applications, the cost in terms of area of the counter-measures is often prohibitive. In most efficient counter-measures, the area is indeed at least doubled w.r.t. an unprotected implementation; this is the most

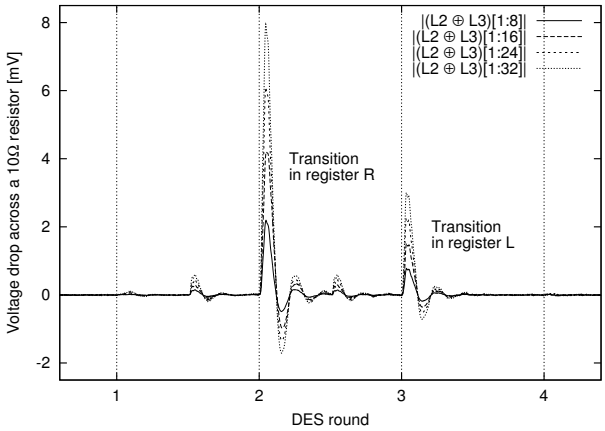

Figure 4. Known key multi-bit DPA extraction of the power dissipation correlated to one quarter, one half, three quarters and the full L register of DES at round $r = 3$.

serious bottleneck that refrains the industrial actors from massively adopting SCA-proof solutions.

2.2.1 SecMat v2 Combinatorial DES

SecMat v2 embeds a fully unrolled DES co-processor. Initially, the LR register (concatenation of the left and right datapath halves) contains public information, as it is, the previous cryptogram. It is loaded with the new plaintext block, which discloses no information about the key. Indeed, in the Hamming weight model, the plaintext is leaked, and in the Hamming distance model, the exclusive or between the plaintext and the previous encryption result is leaked. These messages does not convey a single bit of information from the key. Then, the rest of the encryption unfolds without any register transfer. The fundamental DPA hypothesis that traces remain synchronized is thus violated. Surprisingly enough, the unrolled (hence combinatorial from end to end) DES datapath is only about 5 times larger than the reference iterative DES, as shown in the bottom left picture of SecMat v2 in Fig. 1. This means that logic factorization can be achieved at each of the sixteen rounds borders.

2.2.2 SecMat v3 WDDL and SecLib DES

SecMat v3 contains three hardwired DES accelerators. They have been obtained from the same VHDL description, but are implemented into different micro-architectures. The most secure of those architectures uses a full-custom cell library, called SecLib [21], assembled into a synchronous

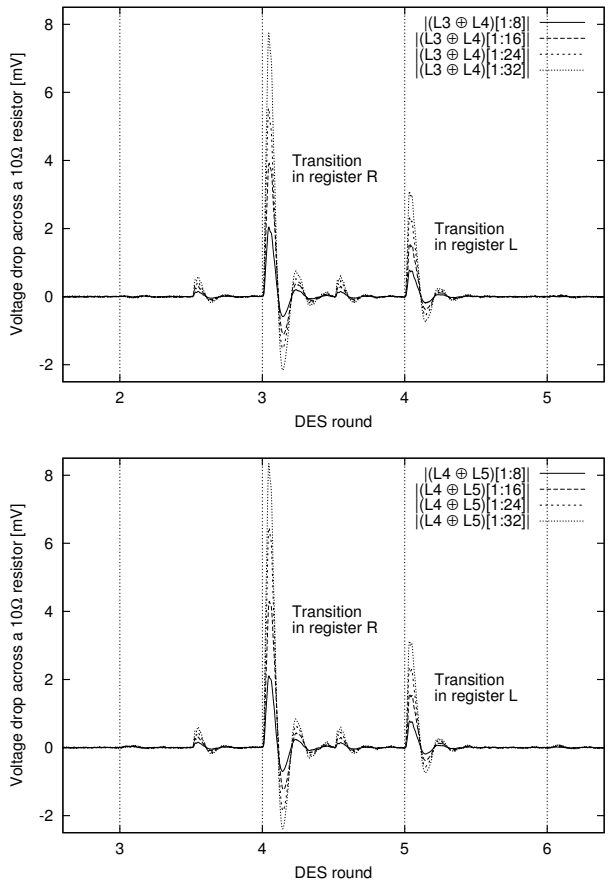

Figure 5. Extraction similar to that of Fig. 4, however with a selection function delayed by one (*top*) or two (*bottom*) encryption rounds.

dual-rail return-to-zero netlist. The second module is implemented in wave dynamic differential logic (WDDL [47]). Those two DES modules feature a high level of immunity against side-channel attacks that exploit information leakage through the power consumption [18]. Their security relies on a careful backend design that balance every possible dissymmetry [20]. The third module is an unprotected reference, that can be easily broken with any SCA.

The SecLib and WDDL modules resist that furiously that a breakthrough is required to successfully exploit their leakage. Indeed, they are expected to leak second order information. For instance:

- WDDL encryptions have a bias due to early precharge and evaluation [45], whereas
- SecLib is immune to signal races attacks. However, it might be vulnerable to subtle disbalance effect caused by technological dispersion.

2.3 FPGAs as Security Prototyping Platforms

The implementation of attacks on an actual ASIC is definitely the best method to evaluate the real security level. However, this process requires some time (for both the design and the manufacturing) and a consequent budget (the creation of the photo-lithographic masks set is expensive).

Therefore, using FPGAs as security prototyping platforms to assess the principle of counter-measures is a relevant alternative. The rationale is that if a counter-measure is strong in FPGA, it is likely to be even more secure once tweaked in an ASIC. Moreover FPGAs perfectly suits low to medium volume markets where security and flexibility are major constraints. At a first glance FPGAs leak much more (in intensity), because (*i*) the computing logic is generic and (*ii*) the routing is active, thus dissipating.

A study from [29] shows that FPGAs are 35 times bigger and 12 more consuming than ASICs on average. Hence the attacks can take advantage of this increase in side channel leaks. The attacks described in the section 2.1 have been successfully ported from the ASIC version of SecMat to the FPGA version. The floorplan under QUARTUS of SecMat v3 and the experimentation board are shown in Fig. 6.

The FPGA can also be an indicated platform to perfect both attacks and countermeasures. Therefore, to gain confidence in the attack methodology, a tunable countermeasure (from weak to strong) is welcome. This cannot be done in ASICs, as it would leave the door open to attacks on the tuning mechanism. FPGA prototypes can play this role, as they can be reprogrammed at will. We illustrate this usage of FPGA in counter-measures devising and evaluation on a concrete case-study. The SecMat v3 ASIC embeds two protected DES co-processors that we fail to

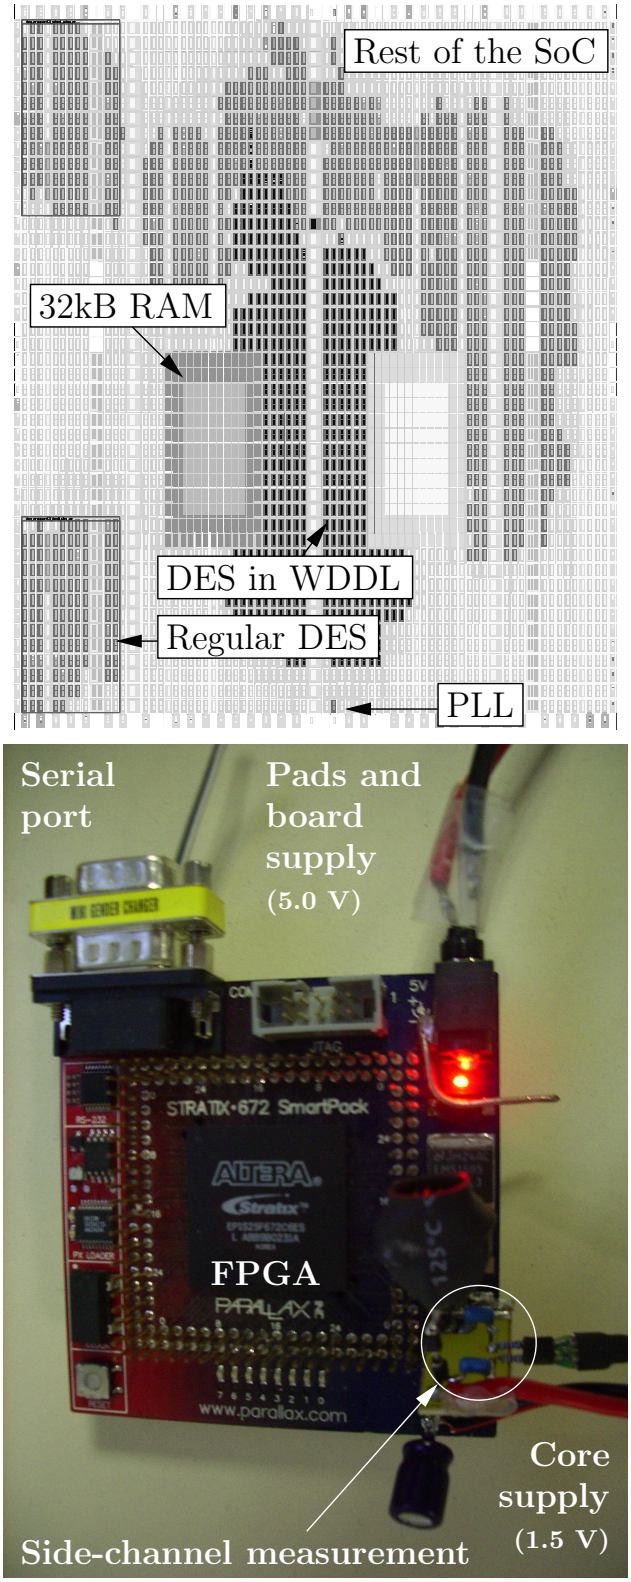

Figure 6. The SoC SecMat v3 in a Stratix (*top*) and the DPA board (*bottom*).

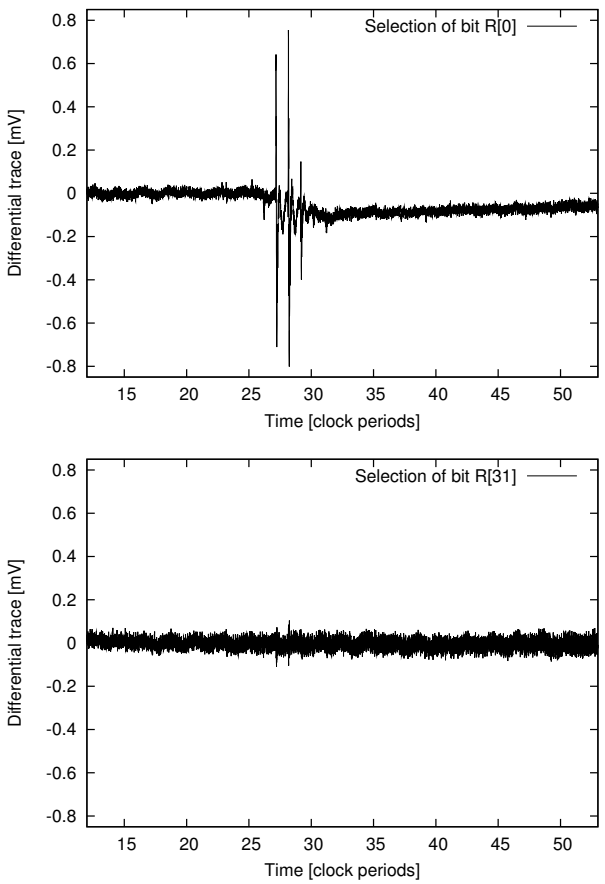

Figure 7. Intentionally unbalanced (*top*) and normally balanced (*bottom*) dual-rail register.

attack. The WDDL [47, 24] version has been ported in FPGA, as shown in the right part of Fig. 6. In order to adequately setup the oscilloscope acquisition window and to motivate the optimal choice for the selection function, it has been interesting to intentionally disbalanced some dual nets. We have namely disbalanced the WDDL registers R[1,10,20,26] (output bits of the 4th sbox). The differential traces are shown in Fig. 7.

Incidentally, it is also amazing to note the FPGAs as such are emerging as a viable trusted computing platform. The security weaknesses, mainly caused by a costly programmable structure, can become a strength when used judiciously. For instance the reprogrammability can provide dynamic countermeasures to confuse side channel attacks or can be done on purpose when a fault attack is detected. However programmability implies configuration memories which can be hacked into, as every software system. FPGA manufacturers already provide bitstream protection based on triple DES or AES but other forms of piracy exist as described in [12]. The community of FCCM publishes a top 10 prediction list `http://www.fccm.org/top10.php` that reflects the prospective vision of the experts in field-programmable custom computing machines. Although focused in the 1990's on the increase of the FPGA market, it is now concerned with the issue of secure FPGA design. The two themes "viruses/malware" and "on-field bitstream patches" have emerged for the 2007–2012 prediction.

3 Fault Attacks

As explained by David Naccache in [33], fault attacks can be conducted in many forms and do not restrict to digital circuits. An astonishing variety of specialized attacks on cryptographic devices have been invented and experimented on smartcards, as reported for instance by these surveys realized by Gemplus [3] and by the UCL [26].

3.1 Experimental Realization of Non-Invasive Faults: Failure Analysis

We have experienced with non-invasive fault attacks on the unprotected DES and AES accelerators of the SecMat v1 circuit. This section presents and extends the experimental results we obtained in [40] on AES-128. We considered soft stresses:

1. **Supply deprivation**, based on the power voltage decrease at nominal clock frequency (32 MHz),

2. **Over-clocking**, based on clock frequency increase at nominal power voltage (1.2 volt).

We gradually raise the stress level, and show that an attacker can accurately choose the quantity of faults induced within the device. The figure 8 shows that there is comfortable range of vulnerable voltage and frequency where the cryptographic device outputs faulty results while not crashing.

This is a notable difference with other non-invasive fault attacks, such as clock or power glitches. Indeed, these brutal (in speed or in amplitude) stresses lead to poorly reproducible failures, making their analysis somehow more tricky.

The same type of faults observed on SecMat v1 has been reproduced in a Stratix (platform depicted in Fig. 6), by reducing the core voltage supply from 1.5 V to 1.0 V. Given the resemblance of the fault occurrence profile, we conclude that the two studied stress modalities generate faults of the same type. The errors are caused by a setup time violation: the output of a combinatorial logic block is sampled in its downstream synchronous register prematurely. In the case of supply deprivation, the combinatorial logic becomes slower at constant clock frequency, whereas in the

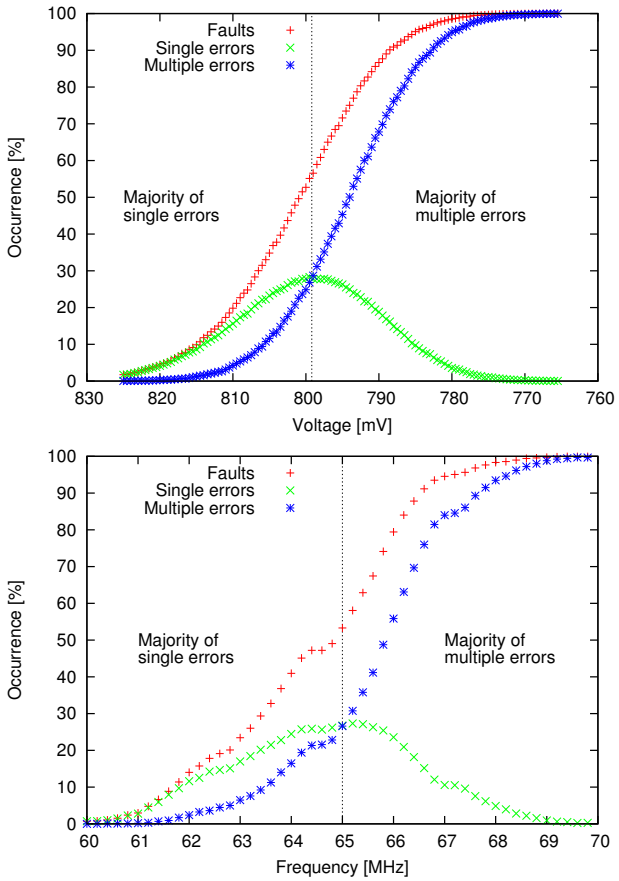

Figure 8. Occurrence of faults when the stress, namely the power deprivation (*top*) or the frequency (*bottom*), increases globally on SecMat v1.

case of over-clocking, the combinatorial logic keeps its natural speed but the registers latches more often. With both fault injection pathes, a critical path is violated, resulting in single errors for low level of stresses: in Fig. 8, single errors are dominant above $\approx$ 800 mV and below $\approx$ 65 MHz.

Differential fault attacks are devastating insofar as very few couples of (correct, incorrect) encryptions suffice to disclose the key. A state-of-the-art overview of DFAs that retrieve the key the more easily is given in Tab. 3. Given the small number of faults to effectively crack the cryptoprocessor, a very stringent security level is mandatory.

3.2 Analysis of Observed Faults

We show that non-invasive practical faults are positively at the advantage of the attacker.

A1 The probability of faults occurrence is controllable very finely by the attacker.

A2 The faults are random, *i.e.* we do not face the case where the same fault always manifests. Instead, a large variety of faults show up.

Table 3. State-of-the-art of DFA.

DFA targeting the AES datapath

Attack	Model	Faults required
Blömer and Seifert [5]	Byte	90,112
Giraud [15]	Byte	250
Dusart and al. [13]	Byte	50
Piret and Quisquater [37]	Byte	2/1

DFA targeting the AES key schedule

Attack	Model	Faults required
Blömer and Seifert [5]	Bit	128
Chen and Yen [10]	Byte	44
Takahashi and al. [46]	Byte	7/4/2
Kim and Quisquater [27]	Byte	2

Those two characteristics are required by most attacks, since all fault models expect that:

1. Single faults (bit-wise or byte-wise) are injected, which implies that the stress can be chosen to be mild: not to weak, because faults must happen, but not to strong, because faults with too high a multiplicity cannot be exploited.

2. Fault models rely on the occurrence of several faults, which means that the faults entropy must be high. If an attack requires two distinct faults to work but that the device always malfunctions in producing the same fault, then the attack fails to be applied. This happens not to be the case experimentally.

The explanation for A1 arises from the fact that DFFs are designed to be highly reliable. The timing window in which DFF risks to become metastable is extremely narrow, making its decision vary continuously with the skew between its input and the clock.

The reason for A2 comes from two factors. First of all, the algorithms such as AES have an highly regular datapath structure: any bit of the state roughly passes through the same amount of logic of similar complexity. Second, the physically-aware and timing-oriented synthesizers tend to still flatten any residual discrepancy that might remain from the netlist unbalance or from the layout. Therefore, the datapath ends up with all bits being potentially on the critical path.

We have not investigated so far whether or not faults are likely to fall in the key-path. However, in algorithms where

Table 4. DPL protocol with a single spacer.

VALID_0		NULL		VALID_1
(1, 0)	⇆	(0, 0)	⇆	(0, 1)

the key schedule is as complex as the datapath, there is no reason why the critical path would not be in the key-path.

3.3 Software Attacks Exploiting Faults

Fault attacks can target directly the cryptographic algorithms, whatever their implementation. However, an alternative attack strategy is to strike at a higher level. The possibility to read/write the full contents of an embedded memory is evoked in [16, 48], by running in background a malicious applet that becomes active once a fault (almost of any type) happens. Without the assistance of an ancillary applet, an attacker can locate weak zones in the code and launch an attack (*e.g.* by glitches) at the very instants when the processor enters them [9]. Therefore, the detection of faults in any organ of an embedded system (not limited to crypto-processor or crypto-code) is very mandatory.

4 Combined Counter-Measures against Side-Channel and Fault Attacks

4.1 SecLib: one DPL Counter-Measure against Side-Channel Attacks

When it is impossible to prevent an attacker from measuring a side-channel, a secure designer can consider two options: either information **hiding** or **masking**. By leaking always the same quantity, dual-rail with precharge logic (DPL) is one paradigm to achieve power-constant circuits, hence hiding internal activity. In order to enable an indiscernible process, every variable is actually ported in the circuit by two wires, said "true" and "false". Every computation is split into two phases: **precharge**, where all the nets are set to a constant value (for instance the ground), and **evaluation**, where exactly one amongst each true/false pair is set to one. This way, every transition from precharge to evaluation consists in a positive edge of half the nets, while the return to precharge consists in a negative edge of the same ones that had previously been set.

A combinatorial gate can implement the protocols of Tab. 4 and 5. However, if we consider an input skew, a stateless gate is not adequate anymore, because it will necessarily evaluate early or late, thus leaking transient information on which input has arrived.

SecLib is a DPL answer to those two combined requirements: the computation must not allow an attack to distinguish between the true and false networks, neither in time

Table 5. DPL protocol with two alternating spacers NULL $\in \{(0,0),(1,1)\}$ **[44].**

VALID_0		NULL		VALID_1
(1, 0)	⇆	(0, 0)	⇄	(0, 1)
(1, 0)	⇄	(1, 1)	⇆	(0, 1)

(X component of traces, discussed in Sec. 4.1.1) nor in intensity (Y component of traces, discussed in Sec. 4.1.2).

4.1.1 SecLib Balance in Time

Although DPL gates are presented pairwise, their internal timing does not guarantee that the evaluation occurs in constant time, Without special caution, there can be a dissymmetry between the true and the false output evaluation dates. Actually, the rationale of some publications can be misleading: pairwise placement can help balance the load of the dual gate, but certainly does not balance the timing.

The CMOS gates power or EM leakage is logical (dependent on data): if a gate evaluates twice the same value, then the second time, there is no (or little) emanation. Otherwise, if the new evaluation value differs from the previous one, then an emanation is emitted. In a similar way, it is important to notice that the skew problem is logical. It is related to the fact the OR gate always evaluates early whereas the AND gate always evaluates late when the spacer NULL is $(0,0)$. The role of OR and AND is swapped when the spacer is NULL=$(1,1)$.

The solution is to delay the evaluation until all inputs have arrived, which cannot be done by a global signal, as otherwise combinatorial logic cannot be implemented. For the return to precharge phase, a global signal can be used. Alternatively, the precharge can propagate. There are thus two options: either the precharge is always early (it is always the faster input to become NULL that forces the output to NULL), or it is always late (*i.e.* it waits for the all the inputs to be NULL).

In all the cases, thwarting early precharge/evaluation cannot be done with combinatorial cells. In SecLib, we use C-element gates [41], layouted by hand. To avoid redrawing cells, the only solution to build level-sensitive sequential cells is use cyclic combinatorial logic. For instance, in the iMDPL gates [38] presented at CHES'07 by Thomas Popp, the synchronization is realized by RS (Reset–Set) flip-flops, built out of standard NOR cells.

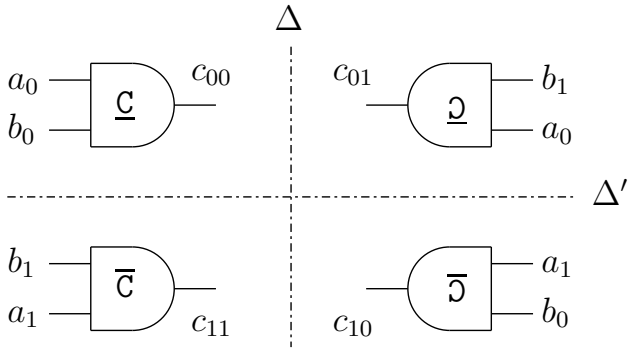

Figure 9. Geometrical symmetries in the placement of the four `C-element`s making up a two-input SecLib gate.

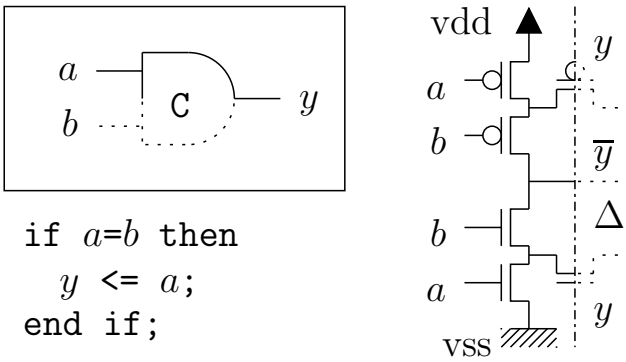

Figure 10. Symbol, VHDL description, and schematic of the half `C-element` gate.

4.1.2 SecLib Balance in Space

The two-input SecLib gate [21] first decodes the input value: either c_{00}, c_{01}, c_{10} or c_{11}. This part is realized by four `C-element` gates, that are placed indiscernibly on the floorplan, as shown in Fig. 9. Now, in turn, each `C-element` is laid out symmetrically w.r.t. an axis Δ: they are built by the assembly of two half-gates, that differ only by the order of their inputs, as depicted in Fig. 10.

The ultimate limitation of SecLib is the technological dispersion. Logic and interconnection dispersion are studied respectively in the papers [19] and [25].

4.2 Dual DPA-DFA Counter-Measures

This subsection reports some outcomes of the project MARS [31]. The goal of this project is to attempt to merge counter-measures against a variety of attacks, so as to combine resistance and efficiency.

Usual counter-measures against fault attacks rely on redundancy insertion and invariant checking. The redundancy can be temporal or spacial:

- a computation is realized several times, and the results are compared, or
- several computations are realized in parallel, the results of which are checked for consistency.

This technique requires modifications either in the control (algorithm scheduling) or in the datapath (algorithm computation part). However, only the RTL description of the algorithm is impacted.

For this reason, RTL counter-measures against DFA and backend-level counter-measures against DPA are not exclusive. They can be used together, without facing the risk of one counter-measures weakening the other.

It has been suggested in [32] that the natural redundancy of information encoding in asynchronous circuits can be used to detect faults, that lead to forbidden states. We concur, and add that this philosophy can be enlarged to any logical-level fault detection mechanism combined with a DPL backend resistant to DPA.

Nonetheless, it must be kept in mind that gentle faults, that do not translate into logical faults, can be used subtly to unbalance a SCA-resistant circuit. A cautionary note has notably been published by K. J. Kulikowski *et al.* in [28].

5 Conclusion

Four ASICs have been designed in 130-nanometer CMOS technology to assess the actual threat of known implementation-level attacks. We tested both passive (such as DPA) and active (such as DFA) on these circuits. It appears that those attacks are extremely efficient: DPA proves to actually extract the activity from whatever Boolean variable in an unprotected cryptoprocessor, which empowers the attacker with an impressive introspection capability. By tuning the stress exercised on the circuit's environment, an attacker can have a remarkably fine-grain and reproducible control over the number of faults occurring in its logic. Thus, DFA allows for practical attacks where the attacker can choose to disturb only gently the device so as to keep environmental sensors unalarmed.

The securization of dedicated cryptoprocessor can basically be divided into two levels. On the one hand, backend-level counter-measures such as power-constant libraries are the most appropriate to obtain a leakage-proof design. On the other hand, RTL counter-measures based on error detection are suitable for the malevolent fault-resistance. Those two counter-measures can be safely deployed on top one of each other, without any risk of one compromising the other.

At the opposite, a skilled designer can enhance the fault detection capability by taking advantage of the redundancy in the data encoding of power-constant logics.

Therefore, sound solutions exist to protect hardware against non-invasive attacks. Semi-invasive attacks remain however extremely efficient. The challenge for the next years is to come up with trusted method to counteract them.

Acknowledgements

The work presented in this article has been partly funded by the Conseil Régional de Provence-Alpes-Côte d'Azur (PACA) and the French National Agency for Research (ANR) through the MARS [31] (ACI SI 2004) grant.

The authors thank the AST division of STMicroelectronics Rousset (France) and Agrate (Italy) for its support in the SecMat (Sécurité du Matériel) project. In addition, the authors are grateful to Dr. Ronan Keryell (GET / ENSTBr, "Trusted Computing Platform" project, now with HPC Project) for his valuable advices and encouragements.

We also sincerely thank Florent Flament, currently with Trust-IC (`http://www.Trust-IC.com/`), for his precious technical and managerial help.

References

[1] M. A. E. Aabid, S. Guilley, and P. Hoogvorst. Template Attacks with a Power Model. Cryptology ePrint Archive, Report 2007/443, December 2007. `http://eprint.iacr.org/2007/443/`.

[2] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi. The EM Side-Channel(s). In *CHES*, volume 2523 of *LNCS*, pages 29–45. Springer, 2002.

[3] H. Bar-El, H. Choukri, D. Naccache, M. Tunstal, and C. Whelan. The Sorcerer's Apprentice Guide to Fault Attacks. *Proceedings of the IEEE*, 94(2):370–382, 2006. DOI: 10.1109/JPROC.2005.862424.

[4] E. Biham and A. Shamir. Differential Fault Analysis of Secret Key Cryptosystems. In *CRYPTO*, volume 1294 of *LNCS*, pages 513–525. Springer, 1997.

[5] J. Blömer and J.-P. Seifert. Fault based cryptanalysis of the Advanced Encryption Standard. In *Financial Cryptography, LNCS Springer 2003*, volume 2742, pages 162–181, 2003.

[6] É. Brier, C. Clavier, and F. Olivier. Correlation Power Analysis with a Leakage Model. *Proc. of CHES'04*, 3156:16–29, August 11–13 2004. ISSN: 0302-9743; ISBN: 3-540-22666-4; DOI: 10.1007/b99451; Cambridge, MA, USA.

[7] S. Chari, J. Rao, and P. Rohatgi. Template Attacks. In *CHES*, volume 2523 of *LNCS*, pages 13–28. Springer, August 2002.

[8] S. Chaudhuri, J.-L. Danger, and S. Guilley. Efficient Modeling and Floorplanning of Embedded-FPGA Fabric. In *FPL*, pages 665–669. IEEE, Aug 2007. Amsterdam, Netherlands. ISBN: 1-4244-1060-6.

[9] S. Chaumette and D. Sauveron. An Efficient and Simple Way to Test the Security of Java Cards™. In *WOSIS*, pages 331–341, May 2005. Miami, FL, USA.

[10] C.-N. Chen and S.-M. Yen. Differential fault analysis on AES key schedule and some countermeasures. In *Information Security and Privacy, LNCS Springer 2003*, volume 2727, pages 118–129, 2003.

[11] C. Clavier. An Improved SCARE Cryptanalysis Against a Secret A3/A8 GSM Algorithm. In *ICISS*, volume 4812 of *LNCS*, pages 143–155. Springer, 2007. Delhi, India.

[12] S. Drimer. Personal web page at Cambridge University, UK. It is entitled "FPGA design security bibliography". Available at: `http://www.cl.cam.ac.uk/~sd410/fpgasec/`.

[13] P. Dusart, G. Letourneux, and O. Vivolo. Differential Fault Analysis on A.E.S. In *Applied Cryptography and Network Security, LNCS Springer*, volume 2846, pages 293–306, 2003.

[14] K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic Analysis: Concrete Results. In *CHES*, volume 2162 of *LNCS*, pages 251–261. Springer, 2001.

[15] C. Giraud. DFA on AES. In *Advanced Encryption Standard (AES) conference, LNCS springer*, volume 3773, pages 27–41, february 2005.

[16] S. Govindavajhala and A. W. Appel. Using Memory Errors to Attack a Virtual Machine. In *SP'03: Proceedings of the 2003 IEEE Symposium on Security and Privacy*, page 154, Washington, DC, USA, 2003. IEEE Computer Society.

[17] V. Gratzer and D. Naccache. Blind attacks on engineering samples. Cryptology ePrint Archive, Report 2005/468, 2005. `http://eprint.iacr.org/2005/468/`.

[18] S. Guilley, F. Flament, R. Pacalet, P. Hoogvorst, and Y. Mathieu. Secured CAD Back-End Flow for Power-Analysis Resistant Cryptoprocessors. *IEEE Design & Test of Computers, special issue on "Design and Test of ICs for Secure Embedded Computing"*, 24(6):546–555, November-December 2007.

[19] S. Guilley, F. Flament, R. Pacalet, P. Hoogvorst, and Y. Mathieu. Security Evaluation of a Secured Quasi-Delay Insensitive Library. In *DCIS, full text in HAL, `http://hal.archives-ouvertes.fr/hal-00283405/en/`*, pages 1–7, November 2008. DCIS'08, Grenoble, France.

[20] S. Guilley, P. Hoogvorst, Y. Mathieu, and R. Pacalet. The "Backend Duplication" Method. In *CHES*, volume 3659 of *LNCS*, pages 383–397. Springer, 2005. August 29th – September 1st, Edinburgh, Scotland, UK.

[21] S. Guilley, P. Hoogvorst, Y. Mathieu, R. Pacalet, and J. Provost. CMOS Structures Suitable for Secured Hardware. In *Proceedings of DATE'04*, pages 1414–1415, February 2004. Paris, France.

[22] S. Guilley, P. Hoogvorst, and R. Pacalet. Differential Power Analysis Model and some Results. In *Proceedings of WCC/CARDIS*, pages 127–142, Aug 2004. Toulouse, France.

[23] S. Guilley, P. Hoogvorst, R. Pacalet, and J. Schmidt. Improving Side-Channel Attacks by Exploiting Substitution Boxes Properties. In *BFCA – `http://www.liafa.jussieu.fr/bfca/`*, pages 1–25, 2007. May 02–04, Paris, France.

[24] S. Guilley, L. Sauvage, J.-L. Danger, T. Graba, and Y. Mathieu. Evaluation of Power-Constant Dual-Rail Logic as a Protection of Cryptographic Applications in FPGAs. In *SSIRI*, pages 16–23, Yokohama, Japan, jul 2008. IEEE. DOI: 10.1109/SSIRI.2008.31, `http://hal.archives-ouvertes.fr/hal-00259153/en/`.

[25] M. Ikeda, H. Yamauchi, and K. Asada. Tamper Resistivity Analysis for Nano-meter LSI with Process Variations. In *ICECS, 13th IEEE International Conference on Electronics, Circuits and Systems*, December 2006. DOI: 10.1109/ICECS.2006.379806.

[26] C. H. Kim and J.-J. Quisquater. Faults, Injection Methods, and Fault Attacks. *IEEE Design & Test of Computers*, 24(6):544–545, 2007.

[27] C. H. Kim and J.-J. Quisquater. New Differential Fault Analysis on AES Key Schedule: Two Faults are enough. In *CARDIS 2008*. Springer, 2008. Royal Holloway, University of London, UK.

[28] K. J. Kulikowski, M. G. Karpovsky, and A. Taubin. DPA on Faulty Cryptographic Hardware and Countermeasures. In L. B. et al., editor, *FDTC*, volume 4236 of *LNCS*, pages 211–222. Springer, Oct 2006. Yokohama, Japan.

[29] I. Kuon and J. Rose. Measuring the Gap Between FPGAs and ASICs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 26(2):203–215, 2007.

[30] T.-H. Le, J. Clédière, C. Canovas, B. Robisson, C. Servière, and J.-L. Lacoume. A Proposition for Correlation Power Analysis Enhancement. In *CHES*, volume 4249 of *LNCS*, pages 174–186. Springer, 2006. Yokohama, Japan.

[31] "MAtériel Robuste pour Systèmes sûrs", `http://projects.comelec.enst.fr/mars/`, project funded by the French National Research Agency (ANR) in the call "ACI SI", led by Sylvain Guilley (ENST) and Régis Leveugle (TIMA), 2004–2007.

[32] Y. Monnet, M. Renaudin, and R. Leveugle. Designing Resistant Circuits against Malicious Faults Injection Using Asynchronous Logic. *IEEE Trans. Computers*, 55(9):1104–1115, 2006.

[33] D. Naccache. Finding faults. *IEEE Security & Privacy*, 3(5):61–65, 2005.

[34] NIST/ITL/CSD. Data Encryption Standard. FIPS PUB 46-3, Oct 1999. `http://csrc.nist.gov/publications/fips/fips46-3/fips46-3.pdf`.

[35] NIST/ITL/CSD. Advanced Encryption Standard (AES). FIPS PUB 197, Nov 2001. `http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf`.

[36] P. Kocher and J. Jaffe and B. Jun. Differential Power Analysis. In *Proceedings of CRYPTO'99*, volume 1666 of *LNCS*, pages 388–397. Springer-Verlag, 1999. (PDF).

[37] G. Piret and J.-J. Quisquater. A Differential Fault Attack Technique against SPN Structures, with Application to the AES and KHAZAD. In *CHES*, volume 2779 of *LNCS*, pages 77–88. Springer, 2003. (Online PDF version).

[38] T. Popp, M. Kirschbaum, T. Zefferer, and S. Mangard. Evaluation of the Masked Logic Style MDPL on a Prototype Chip. In *CHES*, volume 4727 of *LNCS*, pages 81–94. Springer, Sept 2007.

[39] J.-J. Quisquater and D. Samyde. ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smard Cards. In I. Attali and T. P. Jensen, editors, *E-smart*, volume 1240 of *LNCS*, pages 200–210. Springer-Verlag, 2001. ISSN 0302-9743.

[40] N. Selmane, S. Guilley, and J.-L. Danger. Setup Time Violation Attacks on AES. In *EDCC, The seventh European Dependable Computing Conference*, pages 91–96, Kaunas, Lithuania, may 2008. ISBN: 978-0-7695-3138-0, DOI: 10.1109/EDCC-7.2008.11.

[41] M. Shams, J. Ebergen, and M. Elmasry. Modeling and comparing CMOS implementations of the C-Element. *IEEE Transactions on VLSI Systems*, 6(4):563–567, December 1998.

[42] S. P. Skorobogatov. Semi-Invasive Attacks, Oct 2001. `http://www.cl.cam.ac.uk/~sps32/semi-inv_def.html`.

[43] S. P. Skorobogatov. Optically Enhanced Position-Locked Power Analysis. In *CHES*, volume 4249 of *LNCS*, pages 61–75. Springer, 2006.

[44] D. Sokolov, J. Murphy, and A. Bystrov. Improving the Security of Dual-Rail Circuits. In *CHES*, LNCS, pages 282–297. Springer, Aug 2004.

[45] D. Suzuki and M. Saeki. Security Evaluation of DPA Countermeasures Using Dual-Rail Pre-charge Logic Style. In *CHES*, volume 4249 of *LNCS*, pages 255–269. Springer, 2006. `http://dx.doi.org/10.1007/11894063_21`.

[46] J. Takahashi, T. Fukunaga, and K. Yamakoshi. DFA Mechanism on the AES Key Schedule. In *FDTC 2007 Workshop*, pages 62–74, 2007.

[47] K. Tiri and I. Verbauwhede. A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation. In *DATE'04*, pages 246–251, February 2004. Paris, France.

[48] O. Vertanen. Java Type Confusion and Fault Attacks. In *FTDC*, volume 4236 of *LNCS*, pages 237–251. Springer, 2006. DOI: 10.1007/11889700, ISSN 0302-9743 (Print) 1611-3349 (Online), ISBN 978-3-540-46250-7.

[49] IACR Cryptographic Hardware and Embedded Systems (CHES), http://www.chesworkshop.org/, 1999.

[50] IEEE Workshop on Fault Diagnosis and Tolerance in Cryptography (FDTC), http://conferenze.dei.polimi.it/FDTC08/, 2004.

[51] IEEE International Workshop on Hardware-Oriented Security and Trust (HOST), http://www.engr.uconn.edu/HOST/, 2008.