



HAL
open science

Security Evaluation of WDDL and SecLib Countermeasures against Power Attacks

Sylvain Guilley, Laurent Sauvage, Philippe Hoogvorst, Renaud Pacalet, Guido
Marco Bertoni, Sumanta Chaudhuri

► **To cite this version:**

Sylvain Guilley, Laurent Sauvage, Philippe Hoogvorst, Renaud Pacalet, Guido Marco Bertoni, et al.. Security Evaluation of WDDL and SecLib Countermeasures against Power Attacks. IEEE Transactions on Computers, 2008, 57 (11), pp.1482-1497. 10.1109/TC.2008.109 . hal-02893103

HAL Id: hal-02893103

<https://telecom-paris.hal.science/hal-02893103>

Submitted on 20 Mar 2023

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.

Security Evaluation of WDDL and SecLib Countermeasures against Power Attacks

Sylvain Guilley, Laurent Sauvage, Philippe Hoogvorst, Renaud Pacalet, Guido Marco Bertoni, and Sumanta Chaudhuri

Abstract—Logic styles with constant power consumption are promising solutions to counteract side-channel attacks on sensitive cryptographic devices. Recently, one vulnerability has been identified in a standard-cell based power-constant logic called WDDL. Another logic, nicknamed SecLib, is considered and does not present the flaw of WDDL. In this paper, we evaluate the security level of WDDL and SecLib. The methodology consists in embedding in a dedicated circuit one unprotected DES co-processor along with two others, implemented in WDDL and in SecLib. One essential part of this article is to describe the conception of the cryptographic ASIC, devised to foster side-channel cryptanalyses, in a view to model the strongest possible attacker. The same analyses are carried out successively on the three DES modules. We conclude that, provided that the backend of the WDDL module is carefully designed, its vulnerability cannot be exploited by the state-of-the-art attacks. Similarly, the SecLib DES module resists all assaults. However, using a principal component analysis, we show that WDDL is more vulnerable than SecLib. The statistical dispersion of WDDL, that reflects the correlation between the secrets and the power dissipation, is proved to be an order of magnitude higher than that of SecLib.

Index Terms—side-channel attacks, differential power analysis, secured logic style, WDDL, SecLib, backend-level countermeasures.



1 INTRODUCTION

Much equipments must conceal secret information, such as personal data, credentials or intellectual properties. Now, these devices can be stolen or simply bought by any attacker who wishes to retrieve the secrets. Indeed, attackers can eavesdrop the information directly within the equipment. In this context, the digital information can no longer be protected by sole cryptographic means. For this reason, many applications delegate the low-level security to a specialized circuit. It usually takes the form of a smartcard, a trusted platform module (TPM) or an embedded crypto-processor. For instance, in some countries, the access to operated mobile telecommunication networks is protected by a subscriber identity module (SIM) card. The authentication at automated teller machines (ATMs) is often realized by a smart card. Worldwide, personal computers are equipped with TPMs. Some FPGA manufacturers now implement on-chip configuration bitstream decryption.

To avoid on-board bus probing, the secured system consists most of the time of a monolithic ASIC. Securing those chips is of major importance. Two threats have been identified in the last decade: side-channel attacks and fault injection attacks. The principle of fault attacks is to force the circuit to malfunction so as to gain illegitimate information [14]. These attacks are very powerful and some circuits have been successfully broken with this technique.

However, given that this attack is active, the circuit can embed fault detection logic. If an error is detected, the circuit can for instance erase its secrets, which implies that an attack might require to sacrifice many circuits. Side-channel attacks consist in observing whatever physical emanation that leaks from the circuit, in a view to derive some secret information about the secrets it handles. They are more sneaky because they are passive: if they are carried out carefully, the circuit is not aware that it is being attacked. Usual side-channels are the timing, power consumption [26] or electromagnetic emanations.

Many successful attacks on unprotected circuits have been reported publicly since 1996. Standard side-channel attacks (SCAs) are SPA [23], DPA [23], [29], inferential power analysis (IPA) [11], CPA [6], [24], EMA [12], [32] and template attacks [4], [8], [37]. To mitigate side-channel attacks, several types of countermeasures have been proposed and implemented. It is possible to balance or randomize the sensitive design at the algorithmic, logical or physical levels: the overall strength of the design will be that of its weakest countermeasure. The security evaluation of the protected circuits usually proves that the efforts to spend to break the circuit is higher than without protections. Unfortunately, many protected implementations were actually partially broken, albeit with more expansive means. Two reasons are mentioned to explain the attack success. Either the attacker exploits a leakage that is not covered by the countermeasure. Or the hypothesis about a countermeasure is made at one level, say logical, but is not ported at a lower level, say physical.

It is now widely admitted by the side-channel community that the SCAs have the potential to extract information about any net of the design. It is thus very often advised to protect the circuit down to the logic gate. In the field of gate-level countermeasures, two options are generally

- S. Guilley, L. Sauvage, Ph. Hoogvorst, R. Pacalet, and S. Chaudhuri are with Institut TELECOM / TELECOM ParisTech, CNRS LTCI (UMR 5141), Département COMELEC, 46 rue Barrault, 75634 PARIS Cedex 13, PARIS, France.
E-mail: sylvain.guilley@TELECOM-ParisTech.fr
- G. M. Bertoni is with STMicroelectronics, Advanced Systems Technologies, Centro Direzionale Colleoni, palazzo la Dialectica, via cardano 2, 20041 AGRATE – MILANO, Italy.
E-mail: guido.bertoni@st.com




Figure 1. Floorplan and acquisition board of the SecMat v3 ASIC.

considered: static or dynamic countermeasures. The goal of the former is to ensure a power-constant execution, whereas the second consists in ensuring a power-constant execution in average, with the help of an ancillary TRNG.

In this article, we specify an attacker that is able to perform DPA, CPA and template attacks. We investigate experimentally her potential to break implementations protected against the specified attacks, with both logical and physical countermeasures. More specifically, the WDDL logic [44] with wire shielding is assessed. We observe that a reported flaw against WDDL [40] cannot be exploited. Additionally, we investigate another logic, called SecLib [18], immune from the WDDL flaw. The second goal of the paper is to quantify the security gain when switching from WDDL to SecLib. This information is very valuable to adapt the security level to the cost of the assets to protect.

The rest of this article is structured as follows. The ASIC designed for the security evaluation is described in Sec. 2. The three DES modules implementation is detailed in Sec. 3. In Sec. 4, the attack methodology and results are given. Finally, section 5 concludes the paper and opens further research perspectives.

2 PROTOTYPE ASIC DEDICATED TO SIDE-CHANNEL INFORMATION LEAKAGE EVALUATION

A dedicated ASIC has been designed to evaluate the security level reached by the two competing logic styles. In the following, we refer to this chip as “SecMat v3”. SecMat v3 has been taped-out on 2007 January 3rd (STM 0.13 μm technology HCMOS9GP with 6 layers of metallization) through the CMP (Circuits Multi-Projets) silicon broker [9]. The ASIC’s die area is 4.4 mm^2 and contains 2.4 million transistors. The circuit is DRC & LVS clean and has been tested fully functional. A picture of the floorplan and of the acquisition printed circuit board (PCB) is given in Fig. 1. The knowledge of the accurate RTL description of the system is an important feature: it enables us to relate side-channel analyses to the circuit’s operations.

The architectural choices made during the design of SecMat v3 are detailed in this section.

2.1 Security Evaluation Target: ASIC versus FPGA

It makes sense to attack both targets. However, in our context, we endeavor to:

- 1) implement sound and robust countermeasures and
- 2) foster the access to the side-channel. Indeed, to increase our level of confidence in an evaluation, the usual methodology consists in choosing the experimental setup that maximizes the attack’s strength.

The ASICs are thus compared to the FPGAs in these two respects.

The implementation of some countermeasures is either impossible or more difficult in FPGAs. Full-custom logic styles cannot be implemented in FPGAs, since the finest reconfiguration grain is the look-up table (LuT), and not the transistor as in ASICs. The placement can be constrained in both targets. Kris Tiri showed how to place WDDL in F and G LuTs in Xilinx FPGAs [46], [47]. Altera proposes the `logiclock` feature to achieve a similar result, albeit at the logic array block (LAB) level. Native FPGA CAD tools do not implement pair-wise dual-rail routing. Concerning ASICs, some tools start to feature this functionality. For instance, Cadence “chip optimizer” provides a space-based router. Although this post-processing functionality is intended to balance only “special wire” couples, it is conceivable to declare all the nets to be special. Nevertheless, other strategies to achieve this functionality have emerged: fat-wire routing [45] and backend duplication [17] operate on top of the CAD tool. In FPGAs, these methods would require the knowledge of the interconnect resources and the ability to forge a bitstream. Additionally, either the routing graph description must be changeable (in the fat-wire method, the channel width must be halved), or it must be possible for the user to set constraints (in the backend duplication, every other routing track must be blocked). The shielding of signals seems difficult in FPGAs: there are no publicly available papers dealing with this aspect.

Finally, the accurate power measurement in FPGAs is a challenge: spying a part of the FPGA consumption is possible under some product families. However, this constrains the module under test to be placed in a partition of the floorplan close to the power pads. Nonetheless, no application note guarantees that the power will not be modulated by the neighbor logic. As too many parameters remain unknown in FPGA designs, we opted for an evaluation in an ASIC, where every aspect is under control.

2.2 System-Level Architecture

The DES modules must be as indiscernible as possible. Hence the choice to place them on a same silicon die.

Besides, SecMat v3 is a system on chip (SoC), where the modules are slaves of a CPU, playing the role of the master. The interconnect is based on the VCI (Virtual Component Interface [3]) standard. Seen from the CPU, the DES modules share the same interface, and differ only from their addressing space. This organization greatly facilitates their control: the same program is typically used for all DES modules. This program repeatedly installs the cryptographic data (key and message) in each module’s memory, asserts a line to trigger an oscilloscope and launches the encryption.

As the main goal of the ASIC is to realize accurate and fair side-channel measurements, a couple of power pads, called (`gnd_des`, `vdd_des`) is devoted specifically to the DES modules energy supply.




Figure 2. SecMat v3 system-level power management.

Another power requirement is to avoid coupling between the DES modules and other parts of the ASIC (CPU, pads, *etc.*) The solution to lower the “substrate noise” is to insulate the ground of the DES modules from the wafer bulk. The HCMOS9GP technology is triple-well: the NISO CAD layer allows to vertically insulate the P-well of a region. The addition of the NISO mask cannot be done by automatic placers and routers, such as Cadence SOC/Encounter. Therefore, we wrote a SKILL script that post-processes the layout by adding a surrounding NISO rectangle around every DES module. The same script also computes the equivalent diode created between `gnd_des` and the bulk; this information is indeed required by the LVS tool.

Anyway, it remains essential to avoid I/O pad activity during the encryption, especially if the pads carry sensitive data (such as a key). In all the experiments presented in the remainder of this paper, there are no I/O operations during the cryptographic operations.

As already stated, the goal of the ASIC is to be able to measure as accurately as possible the power dissipation of the DES modules, with the additional constraint that the power measurement be the same for all the modules. We opted for a shared power supply for the three blocks, but distinct from that of the rest of the core. The modules can be disabled by clock gating, so that only the attacked module absorbs energy. The clock gating suppresses the dynamic power consumption but not the static leakage current. However, given that the deactivated modules are left in a random state, no relevant information is expected to be leaked this way. For the sake of completeness, we mention that a constant leakage of $180 \mu\text{A}$ is measured on SecMat v3. In Fig. 11 at page 7, a 9 mV offset is observed through a 50Ω “spy” surface-mounted component (SMC) resistor.

A module, called “power management”, decides whether the clock delivered to the DES modules is active or zeroed. The architecture of the “SecMat v3” SoC with the clock gating controller is depicted in Fig. 2.

3 REFERENCE, WDDL & SECLIB DES MODULES

The data encryption standard (DES [30]) was chosen as the algorithm to evaluate the security of WDDL and SecLib countermeasures. This algorithm is the preferred one in ASIC implementations, because it is very small, and because of the confidence people have on its cryptographic strength (when used as triple-DES with three distinct keys). For example, DES is used in the electronic passport, in Europay-Mastercard-Visa (EMV) banking applications and in the bitstream encryption for Virtex 2 Xilinx FPGAs.

The architecture of the DES co-processors of SecMat v3 is detailed in [19]. It is an iterative implementation that processes 64-bits of data and that schedules the round key in parallel with the data encryption; one round of DES is thus computed each clock period. In our setup, the DES is made to operate on one single message block, according to the following schedule:

- clock period 0–7: byte-wise key loading from RAM,
- clock period 8–15: byte-wise message loading from RAM,
- clock period 16–31: encryption (16 rounds), in dedicated registers,
- clock period 32–39: byte-wise ciphertext saving into RAM.

For a fair comparison, the modules were designed to be as similar as possible. The VHDL source code is shared. The reference module has been realized using unprotected gates and straightforward automatic CAD tools. More precisely, we have used the Cadence toolchain for the design (`bgx_shell` for the logic synthesis, `SOC/Encounter` for the place/route step and `icfb` for the layout finishing) and Mentor Graphics `calibre` for the verifications (DRC and LVS). The WDDL and SecLib modules resort to advanced physical design techniques, that differ only regarding their logic style.

A description of the three DES modules embedded in SecMat v3 is already provided in [15]. We summarize the main security attributes of these modules in this section.

3.1 Logic Styles

Constant-power computations often use a *dual-rail with precharge logic* (DPL). This logic is also known as *dynamic differential logic*. The protocol of this logic consists of two phases: precharge and evaluation. The precharge phase allows to start new computations from a known electrical state. It thus prevents unexpected transitions between two computation steps. The dual-rail signalization of the data is conveyed by two wires for each Boolean variable: `NULL = 00` while in precharge and `VALID $\in \{01, 10\}$` while in evaluation. Therefore, every evaluation consists in the transition of exactly one wire (`00 \rightarrow 01` or `00 \rightarrow 10`). If the design is adequately balanced, which transition occurred is indiscernible by an attacker.

3.1.1 State-of-the-art about DPL.

In 2002, Kris Tiri introduces the “Sense Amplifier Based Logic” (SABL) logic style [42], which aim is to make power consumption independent of both the logic values and the sequence of the data. It is therefore the first DPL proposal. Its principle consists in combining Differential and Dynamic Logic (DDL) like in the “Dynamic Cascode Voltage Switch Logic” (DCVSL) style, while fixing second order asymmetry in the gate (especially for complex logic functions), due to parasitic capacitances [36]. This allows to decorrelate the power consumption from the inputs. In 2006, Marco Bucci *et al.* [7] show that the balance of DPL gates can be improved by adding a systematic discharge after the evaluation. The resulting computations are thus based on a ternary pace:

(1) pre-charge, (2) evaluation and (3) post-discharge. When applied to SABL, simulations reveal that a gain of two-order of magnitude is obtained in terms of balance.

As these techniques require the full-custom design of new standard cells, Tiri proposes two years later the “Wave Dynamic Differential Logic” (WDDL) style [44]. WDDL uses a standard cell flow, where an original single-ended gate netlist is duplicated to obtain a differential netlist. In addition, the precharge is not global; instead precharge values are imposed only at the inputs, and propagate as a “wave” through the combinatorial netlist. Finally, the total load capacitance is assumed to be dominated by the interconnect capacitance, so the constant load capacitance is obtained by careful routing.

SecLib is introduced in 2004 by Sylvain Guilley *et al.* [18]. This logic is based on an quasi-delay insensitive asynchronous primitives, that are balanced to provide constant evaluation and precharge time and dissipation. Specially crafted transistor-level symmetry grants SecLib a higher resistance level to attacks than WDDL, albeit at a high cost in terms of silicon area [15], [16].

In 2005, SABL and “Dynamic Current Mode Logic” (DyCML) [2] are compared by François Macé *et al.* [25]. In DyCML, only one of the output nodes is discharged during the precharge phase. This leads to better performances, such as a reduction by 80 % of the power delay product and by 50 % of the power consumption. In addition, DyCML is assessed to be more resistant to DPA than SABL.

Recently, Francesco Regazzoni *et al.* explore the resistance of “MOS Current Mode Logic” (MCML) against DPA [10] up to simulated attacks. Preliminary results show that MCML has a strong potential for protecting circuits.

3.1.2 Early Evaluation Flaw

All the DPL styles presented previously feature a problem mentioned in [40] linked to the intrinsic evaluation in CMOS logic. This logic is memoryless, and thus evaluates as soon as an input changes. Now, in dual-rail logic, the levels of the wires act both as signalization (two wires equal to ‘0’ implies a *precharge* stage), and data (two wires with opposite values mean *evaluation*).

For the sake of illustration, we continue the flaw analysis with the example of WDDL with a 00 spacer to precharge the circuit. This choice makes OR gates evaluate faster than AND, because OR gates simply need one input to have a rising transition to change output values, whereas AND gates must wait for two rising transitions to update their output. A scenario that illustrates the data-dependency of the computation flow (and of the power dissipation) is given in Fig. 3. This testbench shows an OR3 gate, receiving its three inputs A , B and C from synchronized registers. The circuit is synthesized in two two-input OR gates in cascade. As depicted in Fig. 4, we assume that the attacker is able to place the circuit in the state $A = B = 0$ (*i.e.* $A_0 = B_0 = 1$ and $A_1 = B_1 = 0$) and tries to guess the value of C by power analysis. The circuit is in precharge state for the negative values of the time t , and the evaluation starts synchronously for all signals at $t = 0$. We observe that, depending on the value of C , the structure of the dissipation differs:




Figure 3. WDDL testbench in which a data-dependency in the power usage is observed.

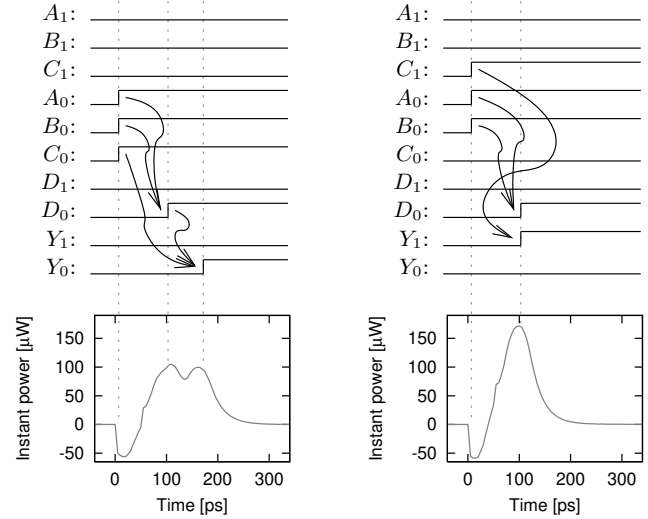


Figure 4. Power signature that betrays the value of the Boolean variable C , in the setup of Fig. 3.

- When $C = 0$, the AND gate called &1 evaluates to true (independently of input C , b.t.w.) and, about 50 ps after that, the second AND gate called &2 evaluates to one, resulting in two distinct power consumption peaks.
- When $C = 1$, the AND gate &1 and the OR gate denoted |2, evaluate simultaneously (at first order), which results in a single power peak.

The power signature thus depends on the value of the variable C . Notice that the problem happens because two paths with different delays converge on the same gate, namely &2 in the false network half and |2 in the other. Incidentally, following the *early evaluation*, WDDL also suffers from an *early precharge* symptom in the next clock cycle.

However, it must be underlined that for this bias to be exploited, the attacker must have an acquisition apparatus that is able to detect 50 ps timing variations. In addition, if the acquisition is somehow low-passed filtered, then the difference vanishes. In the SPICE simulations shown in Fig. 4, the energy consumed by the total transitions is 10.8 fJ for the late evaluation case ($C = 0$) and 11.0 fJ for the early evaluation case ($C = 1$). As these values are very close one from each other, the detection of the difference seems chancy. Nonetheless, the skews add up when descending into the combinatorial logic netlist. A successful attack on a masked DLP (MDPL) circuit exploits a skew of 1 nanosecond at the end of a combinatorial path [34].

In this article, we study SecLib (see Sec. 3.1.4), a DPL

style that does not evaluate early. We compare it with WDDL, because, to the authors' knowledge, it is the only DPL style actually implemented in real cryptographic chips (namely ThumbPod [43] and SCARD [38]). In addition, WDDL does not draw a large current peak at precharge, which simplifies the power planning.

3.1.3 Wave Dynamic Differential Logic (WDDL)

WDDL is a DPL implementable with standard cells. Its principle is that, when a Boolean function $f(x)$ is to be computed, its dual $g \doteq \overline{f(\overline{x})}$ is computed in parallel, so as to mask its activity. Provided that the gate is precharged to zero before every evaluation, either f or g has a transition (exclusively), which ensures a power-constant computation. In SecMat v3, the WDDL synthesis was realized based only on AND and OR instances. Standard cells of several "drive force" from a design kit are armored, so as to:

- ease the pins accessibility and to make pins symmetrical, as required by the backend duplication method, and
- to wrap the standard cells into an electromagnetic cage.

The standard cell is made up of transistors, polarization well-taps and interconnect wires up to the first metal layer (M1). The added coating consists in the superimposition of stripes of the second metal layer (M2). The steps involved in the construction of the armored AND and OR gates are detailed in Fig. 5: the standard cell (1) is added M2 coating (2) to end up with the armored cell (3) = (1) + (2).

3.1.4 Secure Library (SecLib)

SecLib is a balanced quasi-delay insensitive (QDI) cells library that enables power-constant and timing-constant computations. The design of each cell involves two stages:

- 1) a front one in charge of inputs synchronization and
- 2) a back one in charge of the output computations.

Muller C-elements [39] realize the synchronization task. At this stage, the input is decoded. The second stage consists in the redirection of the value to the adequate output, thanks to OR or XOR gates. Redundant logic is added to balance the paths to the *direct* (Y_1) and *dual* (Y_0) output couple, resulting in the schematic given in Fig. 6. In SecLib, The computation is realized for both the direct and its dual output with the same logic, namely a three-input OR gate, which provides a protection against an attacker that would be capable of distinguishing the two halves side-channel signature. The use of C-elements increases the cost in terms of area, delay and power consumption of SecLib cells. However, they do fix the "input skew" issue.

Another advantage of SecLib over WDDL is the large range of logic functions that are affordable – security-wise. For instance, as opposed to WDDL, the SecLib gates can be "logically" inverting and non-positive. Indeed, the C-elements of SecLib handle the precharge state; the evaluation is thus unrestricted. The SecLib library includes the following combinatorial cells: $(A, B) \mapsto \{A \cdot B, A \cdot \overline{B}, \overline{A} \cdot B, A \oplus B, A + B, \overline{A} \cdot \overline{B}, \overline{A \oplus B}, A + \overline{B}, \overline{A + B}\}$. This variety of gates helps to reduce the silicon area overhead of SecLib over WDDL [16].

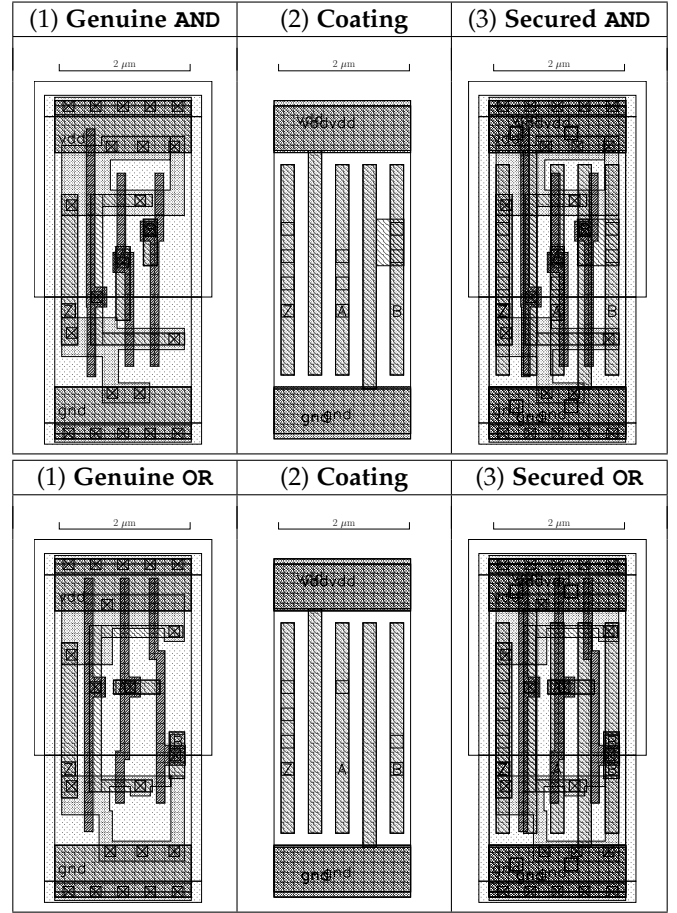


Figure 5. Two-input logic AND and OR gates armoring, suitable for WDDL.

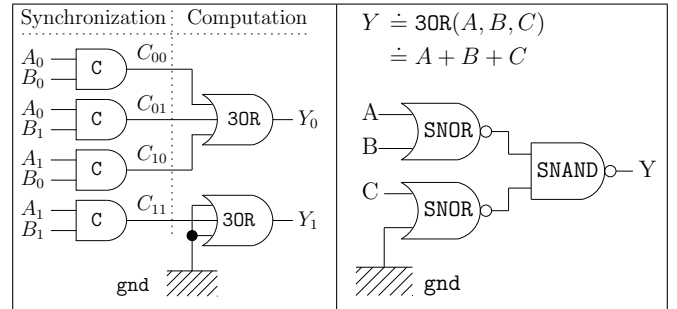


Figure 6. Schematic of the SecLib QDI secured AND gate (left) and its internal 3OR architecture (right).

3.1.5 Common WDDL and SecLib Cells

The DFFs and the buffers are reused directly from the design kit libraries.

For both WDDL and SecLib, the inverter is implemented as a hard-wired cell, depicted in Fig. 7. As those two logic styles expect the netlist to be reset to zero during precharge, the inverter cannot be implemented by the application: $(a_{\text{true}}, a_{\text{false}}) \mapsto (\overline{a_{\text{true}}}, \overline{a_{\text{false}}})$. Instead, the wire crossing $(a_{\text{true}}, a_{\text{false}}) \mapsto (a_{\text{false}}, a_{\text{true}})$ is adequate. Consequently, the inverter of Fig. 7 does not contain any transistor.

The special cells added for WDDL and SecLib synthesis are compatible with standard cells. The height is equal to 12 pitches, divided into a 5-pitch P-well and a 7-pitch N-well.

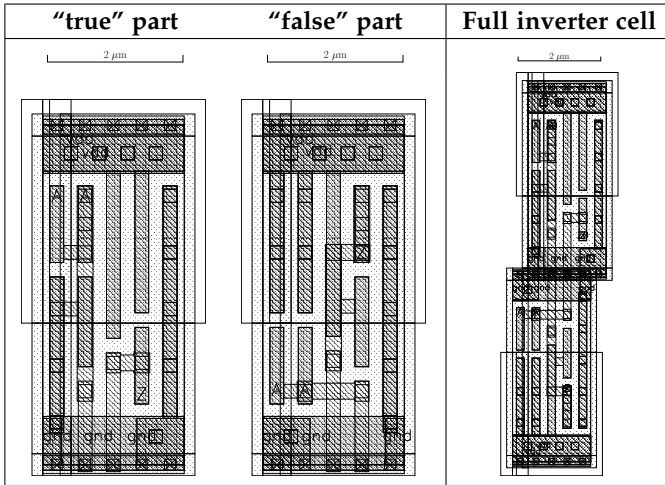


Figure 7. Logical inverter in dual-rail logic, suitable for both “WDDL” and “SecLib” DPL styles.

3.2 Placement and Routing

Two standard methods exist to achieve a balanced dual-rail routing. With the **fat wire** [45] technique, the router tool is tricked into seeing one large wire instead of a couple. The conversion from the resulting single-ended to the dual-rail design is done afterwards by a script. The **“backend duplication”** [17] technique consists in a copy-and-paste of half of the design, placed-and-routed (P&R) with half of the resources obstructed, so as to leave room for a subsequent duplication. The true part of the design is first placed every other row. The false part can therefore fit in the free (because firstly obstructed) placement rows. The same strategy is applied to the interconnect: for every level of metallization, half of the routing tracks is blocked. This precaution makes it possible to route the dual nets in the tracks that have been reserved for them, without creating any short circuit with the regular nets. Compared to the fat wire technique, the backend duplication does not require to tamper with design rules used by the P&R tool, because it relies solely on constraints. Although defining routing constraints are sometimes described as “practically too complex”, we report here that no more than about two hundred lines of TCL scripts (generated automatically from the floorplan description file) can actually suffice to implement the “backend duplication” technique.

The principles of the two placement and routing methods are illustrated in Fig. 8. As the access to the pins of the dual-rail gate instances is difficult with the first method, we have opted for the second one.

Both methods can be enhanced by a systematic shielding of the pairs. This option improves drastically the balance of the pairs in each wire couple, albeit at the expense of routability. In our quest to design a DES co-processor as secure as possible, we decided to apply a systematic shield, which resulted in the design being constrained by the wires. As discussed in [15], in SecMat v3, the placement density of WDDL (resp. SecLib) is 35 % (resp. 95 %).

The shielding method used for both WDDL and SecLib is based on a periodic routing track allocation depicted in the left part of Fig. 9. The corresponding layout is illustrated




Figure 8. Fat wire (*upper*) and backend duplication (*lower*) paths balancing illustration.




Figure 9. Horizontal and vertical routing tracks allocation.

in Fig. 10 for a typical area of the DES WDDL module. We clearly see that the minimally-sized metal layers are mostly crowded, which is characteristic of a **routing congestion** problem.

The shielding method used in SecMat v3 can be optimized. The number of shielding signals can be divided by two, in both directions, without reducing the insulation between the pairs. The corresponding power planning layout is described in the right part of Fig. 9. Instead of 4 tracks to route a dual-rail signal, only 3 are now necessary, both vertically and horizontally; this new shielding scheme enables a $100 \times \left(1 - \left(\frac{3}{4}\right)^2\right)$ % silicon area saving. The density of WDDL can thus be increased from 35 % to 62 %. SecLib density is already 95 %: possible silicon savings




Figure 10. Typical congested routing zone in the SecMat v3 WDDL DES module.

Table 1
Performance of SecMat v3 DES modules.

	Reference	WDDL	SecLib
Area [μm^2]	25 368	299 824	382 871
Energy [nJ]/encryption]	97.2	2×106	2×197
DES-CBC speed [Mbit/s]	266.7	$266.7 / 2$	$266.7 / 2$
3DES-OBC speed [Mbit/s]	88.9	$88.9 / 2$	$88.9 / 2$

are not significantly impacted by a new shielding method. Therefore, the overhead for WDDL can be reduced from 11.8 to $6.6 = 11.8 \times (3/4)^2$. This ratio is still twice larger than in the implementation reported in [41].

3.3 Performances

Table 1 reports the performance of the DES modules. The area of the WDDL module is larger than the factor 3 of overhead claimed in [41] because in SecMat v3 every pair of wire is shielded individually. As the dual-rail modules are limited by the routing, it is not surprising that WDDL and SecLib modules have roughly the same area. The power dissipation has been measured experimentally at 8 MHz under the nominal voltage (1.2 volt). It is expressed as the energy per ECB encryption of one 64-bit block. The DES modules were synthesized to run at 66.7 MHz. At this frequency, the regular DES is able to encrypt or decrypt:

- at 266.7 Mbit/s in DES-CBC mode with a 56-bit key, or
- at 88.9 Mbit/s in 3DES-CBC mode with a 112-bit key.

The dual-rail modules operate twice slower, because every computation step is interleaved with a precharge step.

The performance table shows that securing a chip with WDDL or SecLib has definitely a non-negligible impact both on the cost and on the power budget of the cryptoprocessors. However, these co-processors have been designed with the primary goal to resist power attacks. Actually, as proved in the next section 4, this goal has been reached. Improving the performances while remaining SCA-proof is a challenge we need to address in future research. Second, it must be kept in mind that if the area bloat is undebatably impressive, it can remain acceptable in absolute value. For instance, in the same technology, the 0.3 or 0.4 mm^2 of the secured DES module can be contrasted to an unprotected AES module encrypting an 128-bit block in 44 cycles (0.2 mm^2 [21]) or a 32-kbyte RAM (0.8 mm^2 [21]). Regarding the dissipation, WDDL does not consume much more than twice the power the reference module does (the factor two accounts for the necessary precharge/evaluation dynamic). Roughly speaking, WDDL is built with twice more gates than a single-end logic, but only half of it is activated. SecLib consumes more because each gate is actually made up of several CMOS gates (C-elements followed by OR). As compared to WDDL, one can argue this weakens SecLib. However, as explained in Sec. 4.3, the power consumption is higher but the information leakage it conveys is lower.

4 ATTACKS

We assume that the attacker is able to collect power traces from a circuit. We give the attacker the maximum strength

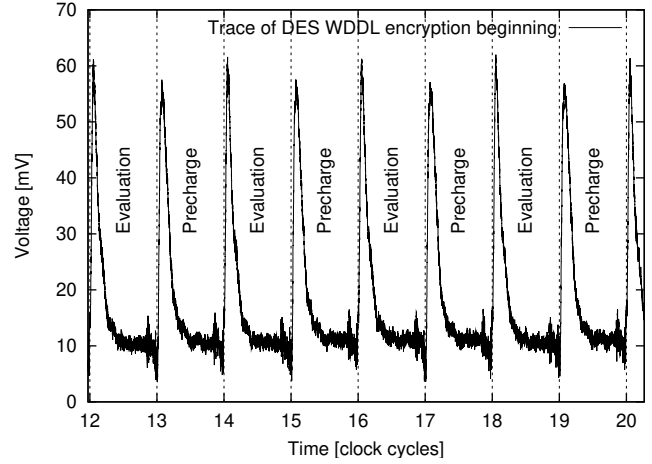


Figure 11. Under-clocked DES WDDL module current trace.

by easing the access to the side-channel and to the synchronization with the encryption. The attacker is fair – it has the same strength irrespectively of the attacked DES module. The exact strength of the attacker is described in the following sections.

4.1 Experimental Traces Collection

Given the small spatial extension (a few tenths of square millimeters) of the cryptoprocessors, a local electromagnetic attack (EMA) is not realistic. With standard antennas, the signal collected would be that emitted globally by the DES cryptoprocessor. This brings down the EMA to a powerline analysis. Thus, we decided to focus on power measurements instead.

We measure the differential voltage across a spying resistor, when SecMat v3, running at 33 MHz, performs an ECB encryption of an all-zero message with the key $0x6b65796b65796b65$. The power traces are averaged 64 times by the oscilloscope, in order to remove the ambient noise and to increase the vertical resolution from 8 to 12 bits.

A typical waveform is shown in Fig. 11. The trace shows that a static leakage current exists.

The Fourier transform of typical traces for each module is given in Fig. 12. The clock harmonics (33 MHz) are visible on all spectra. A peak at half the clock frequency is observable for the WDDL version of DES. This frequency is characteristic of the (precharge, evaluation) dynamic, illustrated in Fig. 11. The reason why the SecLib module does not feature this peak is not intrinsic; it is rather an acquisition artifact, documented in Appendix B. In this Appendix, it is shown that this peculiarity does not affect the fairness of the security evaluation of SecLib. In the WDDL spectrum, some additional peaks are visible for multiples of half the clock period (e.g. 50, 100 MHz). Beyond 100 MHz, all the three spectra feature the same high-frequency components. Therefore we do not expect to exhibit any special side-channel in the $[100 \text{ MHz}, +\infty[$ bandwidth. Consequently, the traces are used plain, without any initial signal processing.

In order to assess the security level of each DES module, we collected 6,400,000 traces for each of them. Gilles Piret

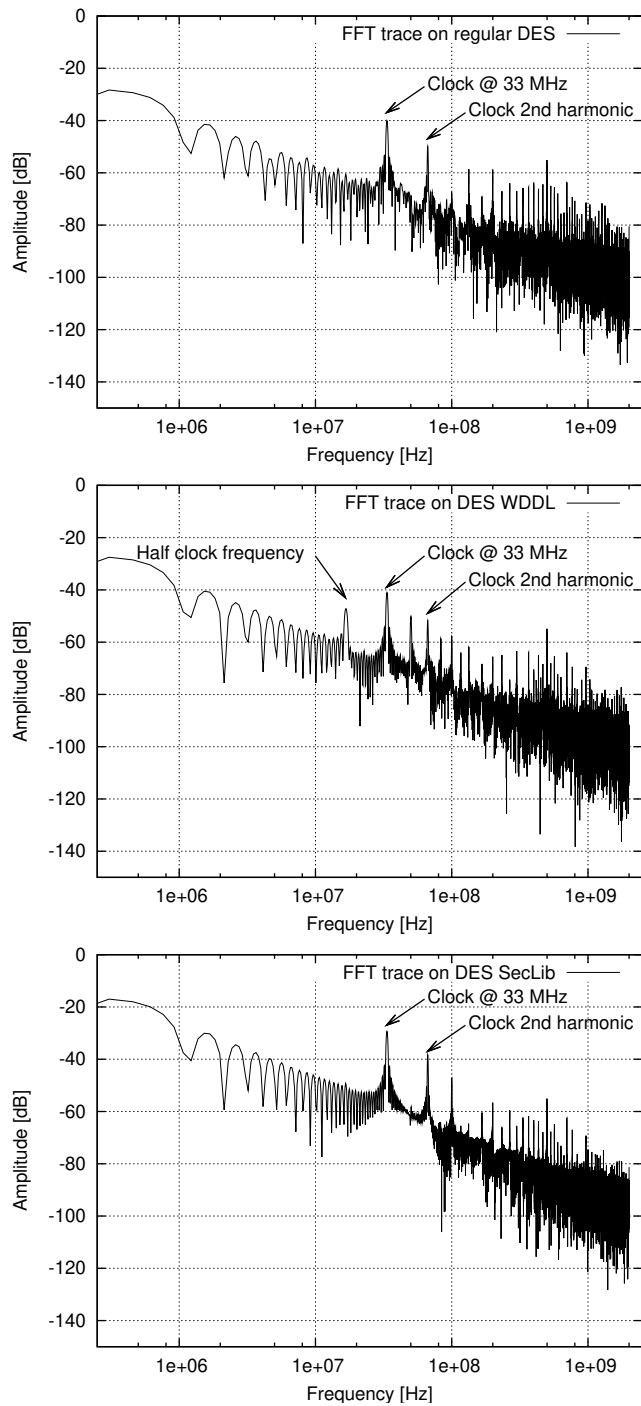


Figure 12. FFT of three power traces from regular, WDDL and SecLib DES modules.

suggests in [33] a method to optimize the number of measurements to disclose the key. He basically proposes two complementary ways to accelerate an attack:

- 1) If the plaintexts are chosen uniformly in front of the attacked substitution box, the selection function bias in the early stages of the correlation attack is minimized.
- 2) If the plaintexts bits not involved in the sub-key attack are chosen constant, the algorithmic noise is minimized.

These two ideas help accelerate an attack, but do not impact its success or the failure with an unlimited amount of side-channel information. As our goal is to test whether the circuits can be asymptotically broken, we simply chose the plaintext randomly with UNIX `rand(3)`.

From a pure cryptographical standpoint, the number of measurements is not large: $6,400,000 \approx 2^{22.6}$, to be contrasted to the $2^{168} = 2^{3 \times 56}$ number of keys in triple-DES with three independent keys.

However, it can give some insights about how much security is available in hardware: it lets the security strategy be partitioned into a hardware/software mixture. For instance, in the context of stream encryption with DES in CFB, OFB or GCM modes of operation, it can give an indication on the frequency of keys renewal: diversified keys regenerated at the rate of one per 6,400,000 encrypted blocks is enough.

4.2 Off-line Attack on the Reference DES Module

In this subsection, efforts are devoted to identify the strongest attacks against the reference DES module. The incentive is to define the best analyses suitable for the protected instances, discussed in the forthcoming subsection 4.3.

4.2.1 Description of the Power Attacks

It is customary to divide power attack into two classes:

- i)* mono-variate analyses, such as IPA, DPA or CPA, and
- ii)* multi-variate analyses, such as template attacks.

4.2.1.1 Correlation Attacks.: We discard IPA because it is too unfavorable from the attacker viewpoint and too specific (it targets software implementations). Instead, we wish to describe the most powerful attacker against a hardware parallel implementation.

Other mono-variate attacks can be nicely unified by the enhanced CPA [24], a heuristic technique that bridges the gap between CPA and DPA. For each side-channel instant, it consists in computing a biased correlation coefficient between the acquired trace (denoted W , as in waveform) and the expected dissipation (denoted H , as in Hamming weight or distance).

If W and H are considered random variables, we note $\mathbb{E}W$ the expectation of W and $\sigma_W \doteq \sqrt{\mathbb{E}(W - \mathbb{E}W)^2}$ its standard deviation (idem for H). The covariance between W and H is defined by: $\text{cov}(W, H) \doteq \mathbb{E}((W - \mathbb{E}W) \cdot (H - \mathbb{E}H)) = \mathbb{E}(W \cdot H) - \mathbb{E}W \cdot \mathbb{E}H$. The

correlation factor between W and H is the normalized quantity, constructed as: $\rho_{W,H} \doteq \frac{\text{cov}(W,H)}{\sigma_W \cdot \sigma_H}$. The Cauchy-Schwarz theorem implies that the correlation factor is normalized:

$$-100 \% \leq \rho_{W,H} \leq +100 \%$$

The H random variable is actually parametrized by a subkey to guess. In DES, the dissipation can be split into eight contributions, each of which corresponding to the substitution boxes (sbox) layer. In each sbox, 6 bits of the key are mixed with the datapath, both at the first and at the last rounds. We thus end up, for every sbox (there are 8 of them in DES), with 2^6 H functions.

The DPA consists in guessing the key according to the greatest value of $\text{cov}(W, H)$, when H explore all the possible key guesses weighting functions. The resulting waveforms are called differential traces, and consist in the extraction of a selected dissipating phenomenon from the overall cryptoprocessor power consumption.

The CPA [6] simply differs from the DPA in that it uses the correlation factor $\rho_{W,H}$ instead of the plain correlation $\text{cov}(W, H)$ to choose which key candidate is the best. It is customary to designate CPA by the term DPA, and to distinguish them as “correlation-based” or “distance of mean” for the classical one.

The enhanced CPA introduces an empirical parameter $\varepsilon \in [0, +\infty[$. The correct key decision is made based on the biased parameter comparison for the 64 key guesses:

$$\frac{\text{cov}(W, H)}{(\sigma_W + \varepsilon) \cdot \sigma_H}.$$

For $\varepsilon = 0$, the enhanced CPA is equal to the regular CPA. When $\varepsilon \rightarrow +\infty$, and provided σ_H is not noisy (for instance using the chosen plaintext methodology described in [33]), the contribution of σ_W is cancelled and the enhanced CPA tends towards the DPA.

The empirical ε offset makes up for a possible statistical artifact: the uninteresting instants in the power curves also correspond to the minimal variance σ_W . However, if this value is too low, $\rho_{W,H} \propto 1/\sigma_W$ becomes artificially large; there is thus the risk that an automatic peak detection software be fooled by such a spurious peak. As on our measurements $\sigma_W > 2.5$ mV, the protection offered by ε is useless.

4.2.1.2 Template Attacks.: Template attacks [8] consist of a two-phase strategy. First, a probabilistic model of the dissipation is built based on the training on a clone device. Second, an intercepted trace is matched against the pre-characterized templates. The practical problem raised by template attacks is the high dimensionality of the data used in the training phase. To alleviate the memory and computational requirements, Archambeau *et al.* [4] proposed to use the principal components analysis (PCA [22]). In many concrete cases, PCA is appropriate. The basic assumption made in PCA is that all templates share a common diagonalization basis; it has been shown to be realistic in many cases.

Unlike correlation attacks (DPA or CPA), that target a single sample in the traces, templates with PCA collect a distributed leakage. Indeed, PCA constructs a linear combination of samples that maximizes the variance (dependency

in the key). This analysis is thus able to capture the skews induced by the early evaluation problem of un-synchronized DPL styles, such as WDDL.

4.2.1.3 Vulnerability Metrics.: The two attack classes just presented allow to qualitatively compare two implementations. If one implementation is broken by an analysis and not the other, then the former is weaker than the later.

However, in the case where two implementations resist an attack¹, correlation and template analyses can produce quantitative metrics that reflect the intrinsic degree of vulnerability of an implementation. For such a vulnerability estimator to enable security comparisons, it must be homogeneous for the various implementations to compare.

We propose three homogeneous metrics that are proportional to the vulnerability criticality.

The first metric is the amplitude of the DPA peak. In [20], it is shown that the differential traces are the extraction of a relevant part from the chip’s overall activity. The targeted logic gates are identified by the DPA selection function. This quantity is thus expressed in the units of the side-channel measurement. As we use a differential voltage probe, the side-channel unit is the volt. This metric might not be appropriate for two unrelated experiments, with different acquisitions apparatuses and conditions. However, the SecMat v3 architecture has been devised to enable comparisons: the side-channel is measured from the same power pads, with the same probe and the same oscilloscope setup.

The second metric is the best correlation factor obtained by CPA. This metric does not have any unit, because it is a ratio. The correlation factor also allows to compare two different setups, since it is relative to the acquisition noise (σ_W).

Finally, the third metric is the largest eigenvalues obtained by template attacks in PCA. Its interpretation is the maximal variance (dependency in the secret) that can be extracted from the side-channel. The units of the eigenvalues are the square of the side-channel, because they represent the square of a standard variation. Thus, as already discussed for the first metric, they are applicable only to setups designed specifically to enable comparisons. It is thus relevant for the comparison of the three SecMat v3 DES modules.

4.2.2 Attack Results of the Reference DES Module

The reference DES module is easily broken with both DPA and CPA. The number of measurements to disclose (MTD) the key is given in Tab. 2. The CPA appears to be the best attack on average. We provide in Fig. 13 the correlation factors obtained after 80k traces accumulations.

We tried the enhanced CPA. This technique is supposed to improve the speed of the CPA; however, apart from sbox #2, the gain is marginal or null, and sbox-dependent. As the protected DES modules have different sboxes (synthesis and P&R differ), the improvement is not expected to be portable. The results are given in Fig. 14.

The thorough analyses made on the reference DES module led to conclusions stated in Tab 3. Based on these results,

1. This happens to be the case for WDDL & SecLib modules (see Sec. 4.3).

Table 3
Analysis of the attack strategies relevant for SecMat v3.

Table 2
Number of traces required to attack the reference DES co-processor with DPA and CPA.

Sbox Index	First round		Last round	
	DPA	CPA	DPA	CPA
#1	146,368	163,008	92,480	65,024
#2	183,040	206,080	201,920	146,816
#3	263,296	227,456	109,440	96,640
#4	191,360	149,376	84,608	72,192
#5	160,384	136,256	79,680	81,984
#6	92,992	89,856	32,000	18,304
#7	241,152	247,552	47,744	47,808
#8	41,280	37,888	227,840	191,744
Worst	263,296	227,456	227,840	191,744
Best	41,280	37,888	32,000	18,304

Attack	Relevance	Description
SPA	no	The control of DES is data-independent
IPA	no	Less powerful than CPA
DPA	no	Less powerful than CPA
CPA	yes	Appropriate
Enhanced CPA	yes	But the improvements are not statistically representative
Templates with PCA	yes	Eigenvalues describe the optimal dependency on the key

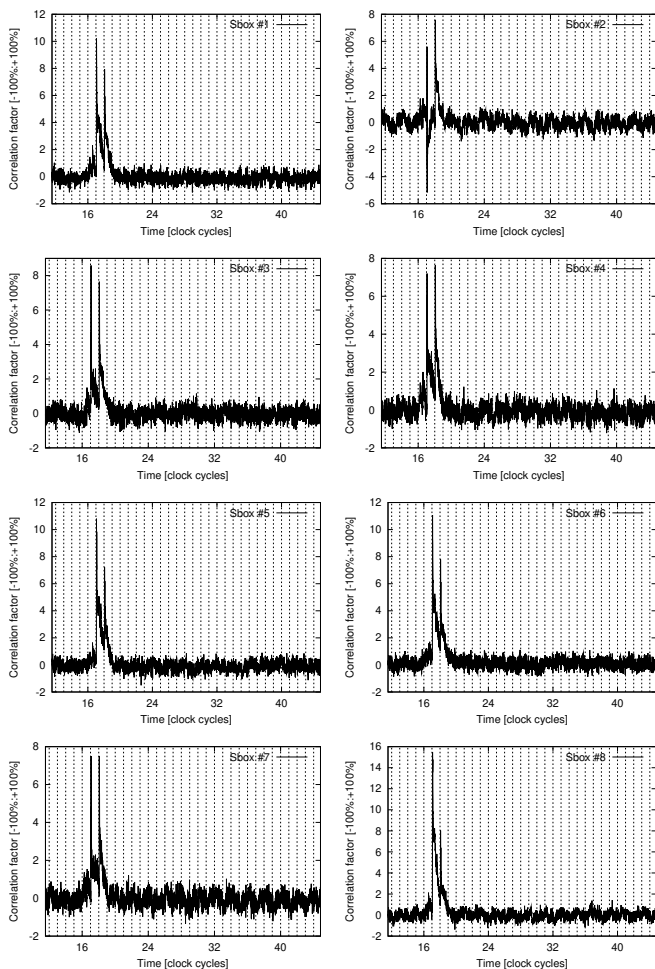


Figure 13. Correlation factor for the correct key guesses obtained when attacking the first round of the reference DES module's eight sboxes.

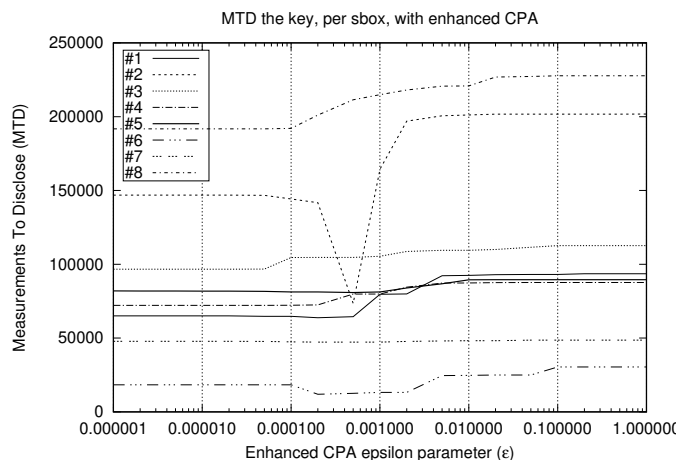


Figure 14. MTD on the reference DES module, attacked with enhanced CPA on the last round (for the eight sboxes).

we can motivate a trustworthy model of an empowered attacker against the two protected instances. To summarize the information gained by an adversary from the preliminary tests, we can say that:

- current traces are preferred over electromagnetic traces,
- traces are used without preprocessing,
- regular CPA or templates with PCA are definitely the best attacks,
- the correlation attacks are slightly better on the last round than on the first one. However, for reasons disclosed in Appendix A, the attack on the last round is more subtle. Therefore, in order to present unambiguous results, the attacks are performed on the first round.

4.3 Off-line Attack on the Protected DES Modules

The CPA has been realized on the first round of the WDDL and SecLib DES modules. The only difference between this CPA and the one used for the regular DES is the switch from the Hamming distance to the Hamming weight selection function. Indeed, because of the precharge, the reference state is plain zero. The Hamming distance, as a tool to count transitions, thus degenerates into a Hamming weight.

Table 4
Extremal correlation factors of CPA on the first round of WDDL and SecLib DES.

Sbox index	DES WDDL		DES SecLib	
	Min.	Max.	Min.	Max.
#1	-1.10 %	+1.10 %	-5.3 %	+4.2 %
#2	-0.82 %	+0.84 %	-5.2 %	+6.6 %
#3	-0.87 %	+1.00 %	-5.2 %	+6.5 %
#4	-0.90 %	+1.10 %	-5.0 %	+6.7 %
#5	-0.93 %	+1.20 %	-6.5 %	+3.9 %
#6	-1.00 %	+1.00 %	-4.7 %	+5.4 %
#7	-1.00 %	+0.95 %	-5.3 %	+5.3 %
#8	-1.20 %	+1.30 %	-7.2 %	+7.8 %

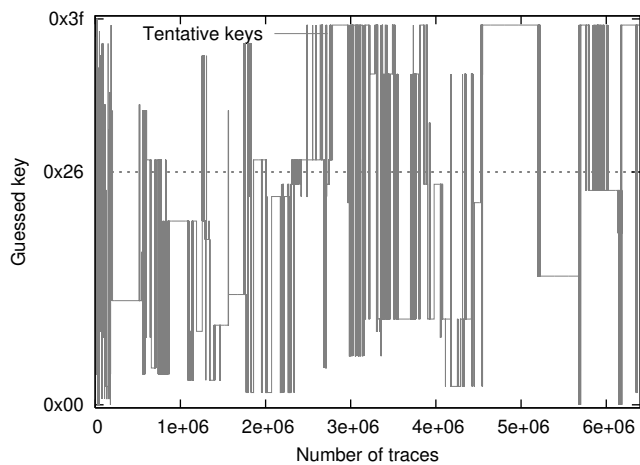


Figure 15. Key automatically selected by CPA on SecLib DES (correct key: $0x26 \in [0x00, 0x3f]$).

The correct key fails to be found by the CPA with 6,400,000 traces. The extremal (minimal and maximal) correlation factors over the whole trace (5,000 points) found for the two protected instances are reported in Tab. 4. It must be emphasized that none of these extremal values correspond to the correct key guess. To illustrate this fact, we show in Fig. 15 how the correlation power analysis on WDDL and SecLib is erring.

The whole correlation traces are shown in Fig. 16 for the first sbox. The highlighted trace corresponds to the correct key guess; the others, superimposed in the background, are those obtained by an erroneous key hypothesis. The correlation traces for the other sboxes are similar: no significant peak appears at the encryption beginning (clock period 16).

The template construction results are shown in Tab. 5. Principal component analysis [4] is used to quantify the amplitude of the variances. The WDDL implementation has two significant eigenvalues, whereas SecLib does not have any overwhelming eigenvalue. The dispersion of WDDL, compared with that of SecLib, after 6,400,000 traces is about $15 = \sqrt{181.2 \text{ mV}^2 / 0.8 \text{ mV}^2}$. This figure means that the WDDL traces depend on the key about one order of magnitude more than the SecLib traces.

Despite the high values taken by WDDL eigenvalues,

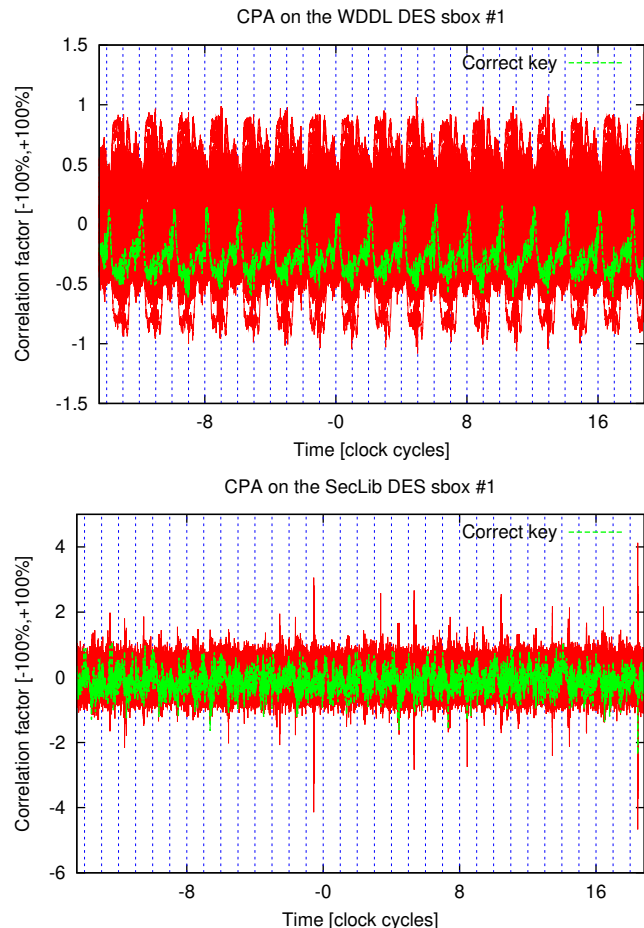


Figure 16. Trace of the correlation factor for WDDL (top) and SecLib (bottom).

Table 5
Three principal eigenvalues, expressed in μV^2 , for the template on the sboxes inputs.

Sbox index	WDDL			SecLib		
	λ_0	λ_1	λ_2	λ_0	λ_1	λ_2
#1	178.3	22.5	0.3	1.0	0.5	0.3
#2	171.5	20.8	0.3	0.9	0.5	0.3
#3	153.6	17.5	0.2	0.8	0.4	0.2
#4	201.5	21.0	0.4	0.8	0.4	0.2
#5	196.7	17.0	0.3	0.7	0.3	0.2
#6	194.8	14.3	0.3	0.7	0.4	0.2
#7	171.4	18.9	0.3	0.8	0.5	0.2
#8	182.3	20.0	0.3	0.9	0.5	0.2
Average	181.2	19.0	0.3	0.8	0.4	0.2

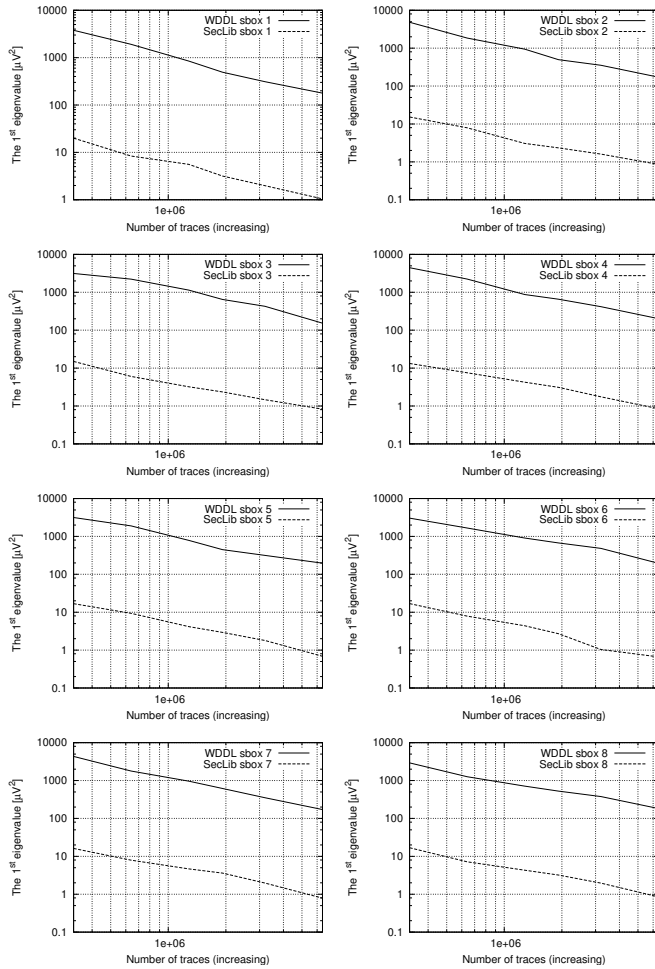


Figure 17. Decay of the largest eigenvalue for WDDL and SecLib modules when characterized by PCA.

the matching of an unseen trace does not work. This can be understood by the fact that the templates quality is not sufficient after their estimation with 6,400,000 traces. To give an idea on the speed of the dispersion convergence, the evolution of the largest eigenvalue with the number of traces used to build the templates is given in Fig. 17. Indeed, the templates are in practice empirical estimators, whose variance decreases with the number of samples used to build them.

4.4 Comparison with the State-of-the-Art

The results obtained on SecMat v3 are compared with the state-of-the-art attacks on circuits protected against SCAs in Tab. 6. The resistance is evaluated with the number of measurements to disclose (MTD) one bit of the key. It can be misleading to compare the MTD the key between different circuits, because setups, acquisition conditions, target algorithms, attacks, *etc* may all differ. We quantify the **security gain** as the ratio between the MTD of a protected and unprotected modules. The selected results have all been validated in silicon. They are listed chronologically.

In 2004, the ADIDES family of asynchronous QDI circuits in 0.18 μm technology (ASIC1) has been successfully attacked [5] because the backend was unbalanced. In 2005,

Table 6
Resistance assessment of protected ASICs, based on real attacks.

Circuit id.	Algorithm	MTD		Security gain
		Unprotected	Protected	
ASIC1	DES	10,000	200,000	20.0
ASIC2	AES	320	21,185	66.2
ASIC3.1	AES	25,000	30,000	1.20
ASIC3.2	AES	25,000	130,000	5.20
ASIC4	CPU	279	471	1.69
ASIC5.1	DES	18,304	6,400,000 is not enough	> 350
ASIC5.2	DES	18,304	6,400,000 is not enough	> 350

the ThumbPod synchronous power-constant WDDL circuit with parallel routing (ASIC2), implemented in 0.18 μm technology, leaks some key bytes [43]. Possible reasons could be the early evaluation problem or an insufficient wires shield against cross-talk. In 2005, a SoC, realized in 0.25 μm technology, embedding various AES processors protected with algorithmic masking (ASIC3) is broken by correlation analysis [28]. The selection function targets glitches in the sboxes [27]. The two masking schemes are that of M.-L. Akkar [1] (ASIC3.1) and of E. Oswald [31] (ASIC3.2). In 2007, the 0.13 μm SCARD [38] evaluation circuit (ASIC4), containing, amongst others, one reference 8051 CPU and seven protected versions, plus some AES hardwired co-processors, is evaluated. The MDPL [35] version of the 8051 is broken because of the early evaluation issue [34]: the MOV instruction leaks the transferred data. Also in 2007, an attack on the SCARD circuit suggests that the MDPL version of AES has a serious breach, due to flaws in the assumptions made on the randomness source [13]. However, the practicability of this attack is still uncertain, notably because no indication about the number of power measurements to break the implementation is mentioned. Finally, the WDDL (ASIC5.1) and SecLib (ASIC5.2) DES co-processors of the SecMat v3 system-on-chip, the 0.13 μm circuit described in this article, remain unbroken.

5 CONCLUSION

A prototype ASIC, called SecMat v3, has been designed and fabricated in 0.13 μm technology. Its purpose is to evaluate the security level of DES co-processors implemented in two power-constant logic styles: WDDL and SecLib. WDDL is subject to a security flaw: under some circumstances, for instance when a skew exists between two signals, the computation duration does depend on some intermediate data. The SecLib logic features a synchronization stage that prevents early evaluation: in addition to being power-constant, SecLib is also timing-constant. The maximal level of efforts has been spent to obtain an accurate idea of the resistance of the protected DES instances. The circuit's architecture, thanks to a power management IP that controls modules clock-gating, allows for fair comparisons of side-channel measurements. The protected modules are carefully designed, especially at the backend level: dual-placement, parallel routing and systematic wire shielding techniques have been used for both WDDL and SecLib modules.

We have found that both secured DES modules feature biases, but that they fail to be exploited by an attack. This does not mean that the DES protected modules are invulnerable. It merely implies that some yet-to-discover attack might defeat them, but that with nowadays attacks, they resisted all our assaults. As of today, the “SecMat v3” ASIC is the most robust power-constant cryptographic implementation because its security gain is the largest published so far (> 350).

ACKNOWLEDGEMENTS

We are grateful to the anonymous reviewers for their help in improving the presentation of the results and in suggesting a way to reduce the area overhead of the WDDL version of DES. We also wish to thank Florent Flament for designing the SecMat v3 ASIC, Karim Benkalaia for producing the PCB, Jean-Luc Danger and Yves Mathieu for assistance with CAD tools and Ronan Keryell for valuable comments on the project in general. Sumanta Chaudhuri has brought an inestimable help in the specification and the validation of the SecMat v3 ASIC; his encouragements were very beneficial to the success of this “trusted computing” project. This work has been partly financed by the french conseil régional “Provence Alpes Côte d’Azur” (Région PACA).

REFERENCES

- [1] Mehdi-Laurent Akkar and Christophe Giraud. An Implementation of DES and AES Secure against Some Attacks. In *LNCS*, editor, *CHES*, volume 2162 of *LNCS*, pages 309–318. Springer, May 2001. Paris.
- [2] M.W. Allam and M.I. Elmasry. Dynamic current mode logic (DyCML), a new low-power/high-performance logic family. In *CICC*, pages 421–424, 2000. DOI: 10.1109/CICC.2000.852699.
- [3] VSI Alliance. On-Chip Bus Development Working Group. Virtual Component Interface (VCI) Standard Version 2 (OCB 2.2.0), April 2001. <http://www.vsia.org/>.
- [4] Cédric Archambeau, Éric Peeters, François-Xavier Standaert, and Jean-Jacques Quisquater. Template Attacks in Principal Subspaces. In *CHES*, volume 4249 of *LNCS*, pages 1–14. Springer, October 10-13 2006. Yokohama, Japan.
- [5] G.F. Bouesse, M. Renaudin, B. Robisson, E. Beigné, P.-Y. Liardet, S. Prevosto, and J. Sonzogni. DPA on Quasi Delay Insensitive Asynchronous Circuits: Concrete Results. In *DCIS*, 24–26 Nov 2004.
- [6] Éric Brier, Christophe Clavier, and Francis Olivier. Correlation Power Analysis with a Leakage Model. *CHES*, 3156:16–29, August 11–13 2004. DOI: 10.1007/b99451; Cambridge, MA, USA.
- [7] Marco Bucci, Luca Giancane, Raimondo Luzzi, and Alessandro Trifiletti. Three-Phase Dual-Rail Pre-charge Logic. In *CHES*, volume 4249 of *LNCS*, pages 232–241. Springer, October 10-13 2006. Yokohama, Japan. DOI: 10.1007/11894063.
- [8] Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template Attacks. In *CHES*, volume 2523 of *LNCS*, pages 13–28. Springer, August 2002. San Francisco Bay (Redwood City), USA.
- [9] “Circuits Multi-Projets (CMP)” website, <http://cmp.imag.fr/>.
- [10] Francesco Regazzoni et al. A Simulation-Based Methodology for Evaluating DPA-Resistance of Cryptographic Functional Units with Application to CMOS and MCML Technologies. In *SAMOS IC*, July 2007. Samos, Greece.
- [11] Paul N. Fahn and Peter K. Pearson. IPA: A New Class of Power Attacks. In *CHES*, volume 1717 of *LNCS*, page 173. Springer Berlin / Heidelberg, August 1999. Worcester, MA, USA. ISSN 0302-9743.
- [12] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic Analysis: Concrete Results. In *CHES*, volume 2162 of *LNCS*, pages 251–261. Springer, May 14-16 2001. Paris, France.
- [13] Benedikt Gierlichs. DPA-Resistance Without Routing Constraints? – A Cautionary Note About MDPL Security –. In *CHES*, volume 4727 of *LNCS*, pages 107–120. Springer, September 2007. Vienna, Austria.
- [14] Christophe Giraud and Hugues Thiebauld. A Survey on Fault Attacks. In Kluwer, editor, *Smart Card Research and Advanced Applications VI, IFIP 18th, World Computer Congress, TC8/WG8.8 & TC11/WG11.2 Sixth International Conference on Smart Card Research and Advanced Applications (CARDIS)*, pages 159–176, 22-27 August 2004. Toulouse, France.
- [15] Sylvain Guilley, Florent Flament, Renaud Pacalet, Philippe Hoogvorst, and Yves Mathieu. Secured CAD Back-End Flow for Power-Analysis Resistant Cryptoprocessors. *IEEE Design & Test of Computers*, 24(6):546–555, November-December 2007.
- [16] Sylvain Guilley, Florent Flament, Renaud Pacalet, Philippe Hoogvorst, and Yves Mathieu. Security Evaluation of a Balanced Quasi-Delay Insensitive Library. In *DCIS*, Grenoble, France, nov 2008. IEEE. 6 pages, Session 5D – Reliable and Secure Architectures, ISBN: 978-2-84813-124-5, full text in HAL: <http://hal.archives-ouvertes.fr/hal-00283405/en/>.
- [17] Sylvain Guilley, Philippe Hoogvorst, Yves Mathieu, and Renaud Pacalet. The “Backend Duplication” Method. In *CHES*, volume LNCS 3659, pages 383–397. Springer, Aug 2005.
- [18] Sylvain Guilley, Philippe Hoogvorst, Yves Mathieu, Renaud Pacalet, and Jean Provost. CMOS Structures Suitable for Secured Hardware. In *DATE’04 – Volume 2*, pages 1414–1415. IEEE Computer Society, February 2004. Paris, France. DOI: 10.1109/DATE.2004.1269113 (Online version).
- [19] Sylvain Guilley, Philippe Hoogvorst, and Renaud Pacalet. A Fast Pipelined Multi-Mode DES Architecture Operating in IP Representation. *Integration, The VLSI Journal*, 40:479–489, July 2007.
- [20] Sylvain Guilley, Philippe Hoogvorst, Renaud Pacalet, and Johannes Schmidt. Improving Side-Channel Attacks by Exploiting Substitution Boxes Properties. In *BFCA*, pages 1–25, May 2007. Paris, France.
- [21] Sylvain Guilley, Laurent Sauvage, Jean-Luc Danger, Nidhal Selmane, and Renaud Pacalet. Silicon-level solutions to counteract passive and active attacks. In *FDTC, 5th Workshop on Fault Detection and Tolerance in Cryptography, IEEE-CS*, pages 3–17, Washington DC, USA, aug 2008.
- [22] Ian T. Jolliffe. *Principal Component Analysis*. Springer Series in Statistics, 2002. ISBN: 0387954422.
- [23] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. In *CRYPTO*, volume 1666 of *LNCS*, pages pp 388–397. Springer, 1999.
- [24] Thanh-Ha Le, Jessy Clédière, Cécile Canovas, Bruno Robisson, Christine Servière, and Jean-Louis Lacoume. A Proposition for Correlation Power Analysis Enhancement. In *CHES*, volume 4249 of *LNCS*, pages 174–186. Springer, 2006. Yokohama, Japan.
- [25] François Macé, François-Xavier Standaert, Jean-Jacques Quisquater, and Jean-Didier Legat. A Design Methodology for Secured ICs Using Dynamic Current Mode Logic. In *PATMOS*, volume 3728 of *Lecture Notes in Computer Science*, pages 550–560. Springer, September 21–23 2005. Leuven, Belgium.
- [26] Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer, December 2006. ISBN 0-387-30857-1, <http://www.dpabook.org/>.
- [27] Stefan Mangard, Thomas Popp, and Berndt M. Gammel. Side-Channel Leakage of Masked CMOS Gates. In *CT-RSA*, volume 3376 of *LNCS*, pages 351–365. Springer, 2005. San Francisco, CA, USA.
- [28] Stefan Mangard, Norbert Pramstaller, and Elisabeth Oswald. Successfully Attacking Masked AES Hardware Implementations. In *LNCS*, editor, *Proceedings of CHES’05*, volume 3659 of *LNCS*, pages 157–171. Springer, August 29 – September 1 2005. Edinburgh, Scotland, UK.
- [29] Thomas S. Messerges, Ezzy A. Dabbish, and Robert H. Sloan. Investigations of Power Analysis Attacks on Smartcards. In *USENIX — Smartcard’99*, pages 151–162, May 10–11 1999. Chicago, Illinois, USA.
- [30] NIST/ITL/CSD. Data Encryption Standard (DES). FIPS PUB 46-3, Oct 1999.
- [31] Elisabeth Oswald, Stefan Mangard, Norbert Pramstaller, and Vincent Rijmen. A Side-Channel Analysis Resistant Description of the AES S-box. In *LNCS*, editor, *Proceedings of FSE’05*, volume 3557 of *LNCS*, pages 413–423. Springer, February 2005. Paris, France.
- [32] Éric Peeters, François-Xavier Standaert, and Jean-Jacques Quisquater. Power and electromagnetic analysis: Improved model, consequences and comparisons. *Integration, The VLSI Journal*, 40:52–60, January 2007.

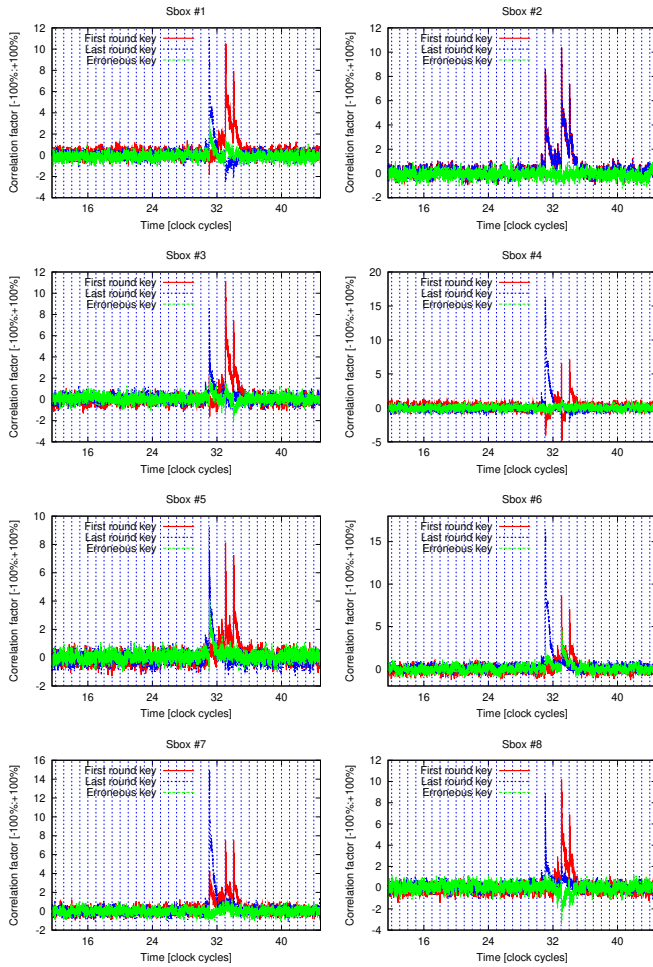


Figure 19. Correlation factors obtained when attacking the last round of the reference DES module’s eight sboxes.

- $R_{14} \oplus R_{15} = L_{15} \oplus L_{16}$. This correlation matches (1) when the key guess is correct, *i.e.* when $K = K_{16}$.
- 2) **At clock period 33**, the transition $L_{16} \rightarrow R_{16} \oplus f(L_{16} \oplus K_1)$ happens in R. The dissipation is correlated with $L_{16} \oplus R_{16} \oplus f(L_{16} \oplus K_1)$, that matches (1) when $K = K_1$.
 - 3) **At clock period 34**, the same transition takes place in register L, hence an echo of the previous strong correlation with $K = K_1$.

Consequently, the DES module embedded in SecMat v3 leaks two non-overlapping 6-bit sets of the key when analyzed by a correlation attack on the last round. From an attacker viewpoint, it is thus profitable to restrict side-channel acquisitions to the clock periods [31-34], because the signal is more intense here than during the encryption beginning.

APPENDIX B DETAILS ABOUT SYNCHRONIZATION

In our setup, the encryption is announced by a trigger signal. The CPU of SecMat v3 executes the snippet of code given in Fig. 20.

```
for(*ever*(;)) // Go!
{
    // This block must be executable at least once,
    // otherwise the trigger is skipped and the
    // message is never encrypted:
    do
    {
        memcpy( msg_addr, msg_backup, msgSize );
        // The synchronization signal for the 54622D
        // oscilloscope is PO[0]. The rising edge of
        // PO[0] announces the next encryption:
        PO_write( 0x01 );
        launch_cipher();
        PO_write( 0x00 );
    }
    while( !UART_is_char_in() );
    // Ciphertext message is in memory at the plain
    // message's address when exiting.
    switch( UART_get_char() )
    {
        case EXIT: return 0;
    }
}
```

Figure 20. Code in C programming language executed by the on-chip CPU (VCI master) to realize side-channel acquisitions.

Due to the system-level VCI [3] management, the encryption is starting few cycles (deterministic value) after the rising edge of a PO signal.

However, the dual-rail modules have their own dynamic. During the precharge stage, they cannot accept data to start a new computation. If they receive all the same a request, it is delayed by one clock period for it to arrive in the evaluation stage.

The behavior of the circuit executing the abovementioned code has been simulated under Mentor Graphics MODELSIM. It happens that:

- The WDDL module always starts after 26 cycles,
- The SecLib module starts one encryption over two after 25 cycles or after 26 cycles.

Thus, when averaging the signals 64 times, the SecLib DES is accumulating 32 traces starting on time with 32 traces starting one clock earlier. This explains why the FFT spectrum of SecLib (Fig. 12) does not show a peak at half the clock frequency. We have captured an unaveraged trace of SecLib, and for this signal, the FFT does show the peak that vanished because of the averaging. This “on-chip communication” problem can be safely ignored for our analyses. To bring an experimental evidence to this assertion, we simulated the timing offset on WDDL. The WDDL traces were split into two groups, the second being additionally offset by one clock period. Of course, this helps neither CPA nor DPA. In the template attacks, that are multi-variate, the temporal position of the leak does not affect the results. Indeed, the results shown in Fig. 21 for sbox #1 confirm this assumption; the other sboxes exhibit the same independence w.r.t. the probabilistic timing offset.

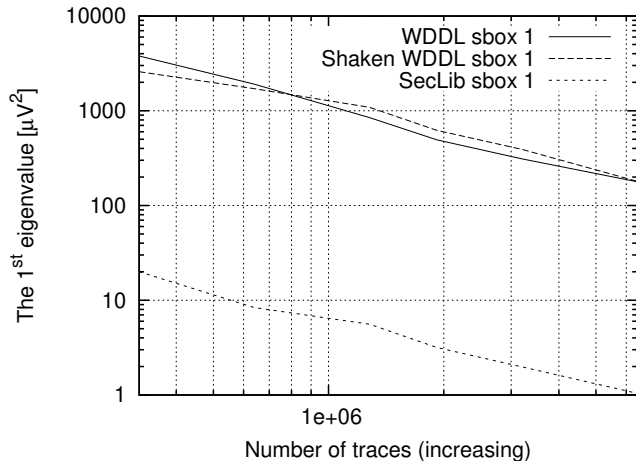



Figure 21. Decay of the largest eigenvalue for “consistent” and “artificially shaken” WDDL, compared to SecLib, while characterizing protected DES modules by templates in PCA.



Renaud Pacalet received his M.S. from the ENST in 1988. From 1993 to 1995 he worked on various industrial projects as a research engineer at TELECOM ParisTech. From 1996 to 2003 he was responsible for the Integrated Systems group at TELECOM ParisTech. From 2003 on, he created and now leads the Systems-on-Chip laboratory of TELECOM ParisTech at Sophia-Antipolis. His research interests are the flexible architectures for the software defined radio; the methods and tools for the specification, design and validation of integrated systems; the security of embedded systems (shielding against side-channel attacks, privacy and integrity of memory buses, formal proof of critical embedded software).



Sylvain Guilley belongs to the french inter-ministerial body of telecommunication engineers. He graduated from the École Polytechnique (X1997) and from the École Nationale Supérieure des Télécommunications (ENST) in 2002. In 2002, he also received the M.S. of quantum physics from the École Normale Supérieure (ENS). He got a PhD *summa cum laude* from TELECOM ParisTech (new brand name of the ENST) in 2007 on the topic of backend countermeasures against side-channel


attacks. Since 2002, he is associate professor with the VLSI group at TELECOM ParisTech. His research interests are the security of cryptographic hardware (ASIC and/or FPGA) and the specification of provable trusted computing platforms.



Guido Marco Bertoni received the Dr. Eng degree in computer engineering and the Ph.D degree from Politecnico di Milano in 1999 and 2004 respectively. He joins ST in fall 2003 as researcher in the field of cryptography. His research interests include the cryptographic algorithms, hardware and software implementations, and problems related to side channels attack. He teaches cryptography at Politecnico di Milano as contract professor.



Laurent Sauvage received his M.S. in electronics, electrotechnique and cybernetics in 1998 and the “agrégation” (french national competitive exam for high school teachers) of electrical engineering in 2002. He is currently pursuing a PhD in practical side-channel attacks (mainly DPA & EMA). He is responsible for the experimental aspects linked to the physical cryptoanalysis platform of TELECOM ParisTech.



Philippe Hoogvorst graduated from École Normale Supérieure, Paris. He got his PhD in 1974 for a study on languages without assignments. He was one of the creators of the “Laboratoire d’Informatique Expérimentale of the École Normale Supérieure”. He defended a “thèse d’État” in 1983 on the same subject as the PhD. Philippe Hoogvorst is currently researcher at the CNRS and detached to the [LTCI/UMR 5141](#). He is working on innovative ways to attacks on electronic circuits; more specifically, he devises signal

and information processing techniques ranging from correlation to template attacks.



Sumanta Chaudhuri received the BTech degree in electronics and communication engineering from the National Institute of Technology, Warangal, India, in 2000 and the MSc degree in electrical engineering from the Ecole Nationale Supérieure des Télécommunications, Paris, in 2005. He worked as a research engineer at the Center for Development of Telematics, Bangalore, India, from 2000 to 2004 and is currently working toward his PhD thesis with CNRS and ENST, Paris. His research interests include reconfigurable computing, asynchronous circuits, cryptography, and, in a broader sense, physical implementation of computing and communication networks.