Electromagnetic Analysis: Concrete Results

[Published in Ç.K. Koç, D. Naccache, and C. Paar, Eds., *Cryptographic Hardware and Embedded Systems – CHES 2001*, vol. 2162 of *Lecture Notes in Computer Science*, pp. 251–261, Springer-Verlag, 2001.]

Karine Gandolfi, Christophe Mourtel, and Francis Olivier

Gemplus Card International, Card Security Group
Parc d'Activités de Gémenos, B.P. 100, 13881 Gémenos France
{Karine.Gandolfi, Christophe.Mourtel, Francis.Olivier}@gemplus.com
http://www.gemplus.com/smart

Abstract. Although the possibility of attacking smart-cards by analyzing their electromagnetic power radiation repeatedly appears in research papers, all accessible references evade the essence of reporting conclusive experiments where actual cryptographic algorithms such as DES or RSA were successfully attacked.

This work describes electromagnetic experiments conducted on three different CMOS chips, featuring different hardware protections and executing a DES, an alleged COMP128 and an RSA. In all cases the complete key material was successfully retrieved.

Keywords. Smart cards, side channel leakage, electromagnetic analysis, SEMA, DEMA, DPA, SPA.

1 Introduction

In addition to its usual complexity postulates, cryptography silently assumes that secrets can be physically protected in tamper-proof locations.

All cryptographic operations are physical processes where data elements must be represented by physical quantities in physical structures. These physical quantities must be stored, sensed and combined by the elementary devices (*gates*) of any technology out of which we build tamper-resistant machinery. At any given point in the evolution of a technology, the smallest logic devices must have a definite *physical extent*, require a certain *minimum time* to perform their function and dissipate a minimal *switching energy* when transiting from one state to another.

This paper analyzes an area of recent interest – electromagnetic side-channel attacks – which exploits correlations between secret data and variations in power radiations emitted by tamper-resistant devices.

Since any electrical current flowing through a conductor induces electromagnetic (EM) emanations, it seems natural to look for the same phenomenon in the

vicinity of a semiconductor. As the power consumption of a tamper-resistant device varies while data are being processed, so does the EM field and one may legitimately expect to extract secret information from a relevant EM analysis.

In some cases, power curves appear to convey no information: this happens when power does not vary or does vary but in a way seemingly uncorrelated to the secret data. Very much simplified, the chip's global current consumption can be looked upon as a big river concentrating the sum of the small tributaries flowing into it. If the subcomponents' contributions could be determined, then the small streams would be isolated. This is impossible by direct electrical measurement but should become possible by eavesdropping local EM radiations. By opposition to power analysis, this requires the design of special probes and the development of advanced measurement methods that focus very accurately selected points of the chip.

For the sake of scientific accuracy, we would like to precise that this paper does not claim the discovery of EM information leakage (which is attested by numerous accessible sources [1–3, 8–10, 12]); we rather report *complete* and conclusive experiments where secrets used by *specific* cryptographic algorithms running on eight-bit CMOS microcontrollers were *thoroughly* disclosed.

Intentionally, none of the tested programs featured software counter-measures against power or EM attacks and in each case the EM information leakage was compared to the result of power attacks performed under identical experimental conditions.

The rest of this work is organized as follows: in section 2 we describe the experimental conditions under which our results were obtained. The results themselves are presented, commented and compared to power leakage in section 3.

2 Electromagnetic Analysis

2.1 Probe design

Chip-scale electromagnetic analysis requires very small probes, similar in dimension to the chip areas to be isolated. The standard layout of a smart card chip shows functional blocks of a few hundred microns (CPU, cryptoprocessor). This defines an upper bound for the probe size.

Although this experimental study was carried out by successively trying different kinds of sensors such as hard disk heads, integrated inductors and magnetic loops [5, 7], the best EM signals were collected using simple hand-made probes. These are solenoids made of a coiled copper wire of outer diameters varying between 150 and 500 microns. An example is shown in Figure 1.

2.2 Electrical behavior

An important advantage of such inductive sensors is their broadband. In other words, a resonance frequency which is much higher than the highest frequency that the analyzed chip is able to generate. The characterization of such sensors is a rather difficult task requiring the generation of a constant-magnitude magnetic-field over a very broad spectral band (several tens of MHz).

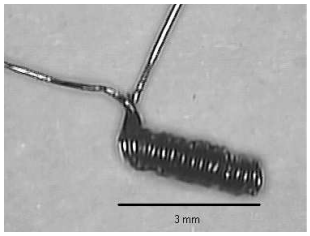

Fig. 1. Electromagnetic Probe.

The main drawback of such probes is their very low output signal (typically 2 to 4 mV peak to peak). Sensitivity can be enhanced at the expense of bigger frequency selectivity, thereby resulting in some bandwidth loss. The trade-off was finally settled for the benefit of bandwidth as radiation spectra were unknown. The amplitude's weakness was compensated by the use of an advanced acquisition chain featuring a very efficient amplification stage.

Most chips are designed in CMOS technology. Figure 2 shows a CMOS logic inverter. The inverter can be looked upon as a push-pull switch: `in` grounded cuts off the top transistor, pulling `out` high. A high `in` does the inverse, pulling `out` to ground. CMOS inverters are the basic building-block of all digital CMOS logic, the logic family that has become dominant in very large scale integrated circuits (VLSI).

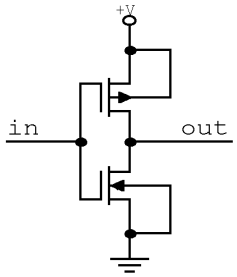

Fig. 2. Elementary CMOS gate.

During a transition from 0 to 1 or *vice-versa*, the device's n and p transistors are on for a short period of time. This results in a short current pulse from V_{dd} to V_{ss}. This (very partially) explains why information leaks when data flips and why the power curve is correlated to the transition's Hamming distance.

This sudden current pulse causes a sudden variation of the EM field surrounding the chip which can be monitored by inductive probes which are particularly

sensitive to the related impulsion [5, 7]. The electromotive force across the sensor (Lentz' law) relates to the variation of magnetic flux as follows:

$$V = -\frac{d\phi}{dt}$$

where V, ϕ and t denote the probe's output voltage, the magnetic flux sensed by probe and the time. In practice, parasitic resistors and inescapable measurement imprecisions require a slight correction of the probe's output.

Whenever a bit flips, the resulting time signal exhibits a high frequency damped oscillation. Acquisition was optimized to better reflect these variations and data dependencies. Frequency-tuned signal processing can be applied. This may require, approximately, a 1 GHz sampling frequency.

Figure 3 shows a power consumption example while Figure 4 shows the corresponding EM signal. The monitored signal is caused by the execution of a *transfer into accumulator* instruction (TIA) applied to `00h` and `FFh` (the specific instruction name was deliberately changed to keep the chip's identity secret).

EM curves appear to be more noisy than power curves, but feature sharper data signatures. Moreover, EM signals can be phase-reversed given the minus sign in Lentz' law and the probe's spatial position: the magnetic flux is inverted by changing the side of the source where the sensor is present.

To reduce parasite signals, an attempt was made to host the chip and the probe in a Faraday cage. This had little effect and finally proved to be unnecessary. Isolating an experiment from external high frequency radiations proves to be a nontrivial engineering exercise for even if the probe can be hosted in a pollution-free cage, most elements in the acquisition chain remain sensitive to ambient EM noise and prone to mutual (cross-talk) perturbations.

2.3 Spatial positioning

To increase the chances of capturing data-dependent signals, the probe was positioned in the neighborhood of a region that radiates while the program runs. Areas radiate with different intensities and various code dependencies but, experimentally, the most active points appear to be located near the CPU, data buses and power supply lines. Amongst these three, the CPU seemed to be the most data-dependent.

Each curve in Figure 5 is the difference between two traces: that of $\texttt{00h} \oplus \texttt{00h}$ and $\texttt{FFh} \oplus \texttt{00h}$. This simple experiment illustrates the information leakage of the exclusive-or instruction *via* the power consumption and EM radiation as measured at five different locations : ROM, EEPROM, RAM, the supply line and the CPU. Each area features a distinct signature either through the signal's shape or magnitude. The CPU clearly stands out by radiating the most informative signal.

Approximating the source as a long linear wire (or the probe as negligibly small), the field's magnitude B decreases (Biot and Savart's law) as the inverse of the distance r between the wire and the probe:

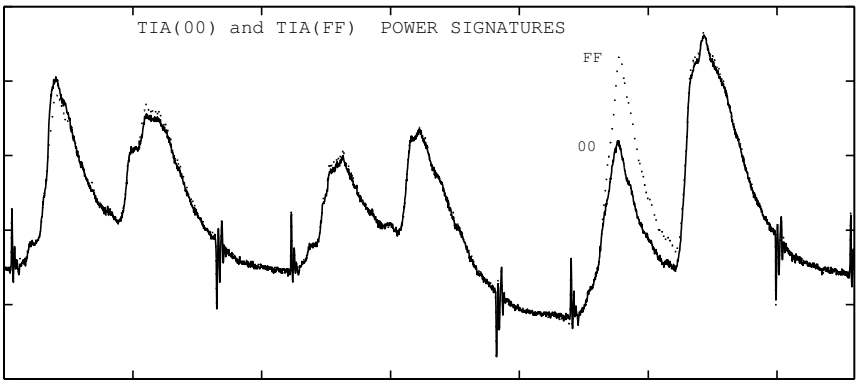

Fig. 3. Current consumption during TIA.

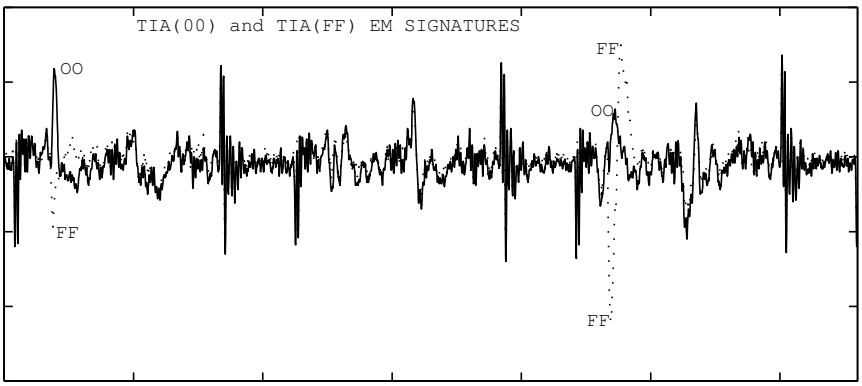

Fig. 4. EM radiation during TIA.

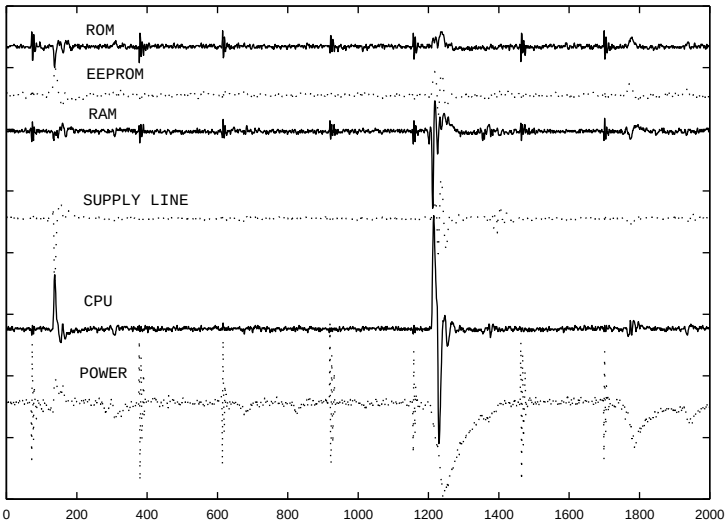

Fig. 5. Differentials between $B(\texttt{00h} \oplus \texttt{00h})$ and $B(\texttt{FFh} \oplus \texttt{00h})$: significative spikes are located near $t = 100$ and $t = 1200$; other regular spikes are just clock residues.

$$B = \frac{\mu_0 \, I}{2\pi \, r}$$

where I denotes the current flowing through the wire. It is thus important to perform measurements as closely as possible to the chip. Since the standard thickness of a card is 800 microns, landing the sensor on its back sets the observation point at 400 or 500 microns away from the target.

This distance may sometimes appear to be prohibitive given the weakness of the EM power radiant and its low signal to noise ratio. However, in some cases the chip's surface can be eroded by mechanical or chemical means [9, 3]. This operation (called decapsulation) offers two important advantages: once the chip is bare (if still functional), the probe's coil can be lowered so as to touch the passivation layer and thereby capture the highest possible field. As a side effect, the chip becomes optically visible and its specific blocks can be pinpointed more accurately. Recapsulating the chip after the attack remains possible for industrially-equipped attackers.

3 Practical Results

We conducted practical experiments on various devices and algorithms. In this section three of the most significant results are presented. Interestingly, the three were conducted on different chips made by different manufacturers.

One of these chips is protected by a shield and the other two feature randomly synthesized logic (RSL). This means that the CPU is scrambled with other functional and useless blocks to make specific functions difficult to identify. Such designs thwart physical intrusions using Focused Ion Beam test equipment (FIB).

The attacked algorithms were respectively the alleged COMP128 (described in [6] and hereafter denoted ACOMP128), DES and RSA. In all cases software counter-measures were deliberately turned off. For a comparative study, test cards were calibrated with known keys. Power and EM signals were systematically acquired simultaneously. Once conditioned, the EM signals were digitized and processed exactly the same way as classical power signals, using the same sampling frequency, digitizer and software tools. Only their physical nature differed.

J.-J. Quisquater and D. Samyde suggested [12] the following acronyms: DEMA for **D**ifferential **E**lectro**M**agnetic **A**nalysis, by analogy to P. Kocher's **D**ifferential **P**ower **A**nalysis (DPA) [8]. **S**imple **E**lectro**M**agnetic **A**nalysis (SEMA) relates in a similar way to **S**imple **P**ower **A**nalysis (SPA). DA and SA will be used for **D**ifferential and **S**imple **A**nalysis, when the leakage's physical nature happens to be irrelevant. D. Naccache coined the Greek term *cryptophthora* to generically address the phenomenon of side channel leakage.

3.1 Alleged COMP128

In this experiment, the smart card was not decapsulated and therefore the probe was positioned rather approximately. No software DPA/DEMA counter-measures were activated.

A DPA and a DEMA were performed simultaneously on the same batch of 256 chosen messages and related sets of curves. Results are shown in Figure 6: the two attacks generated differential spikes for the same right guess. Despite more noisy measurements, the DEMA provided better peaks than the DPA both in terms of contrast and signal to noise ratio. Equivalently, the DEMA required less acquisitions than the DPA. The sign opposition in the raw signatures remained visible throughout the whole differential analysis process. Moreover, since wrong guesses provided no peaks, the experimental evidence was brought that DEMA could work successfully.

3.2 DES

Having obtained these first results, a new DEMA was attempted on another component. Again, no decapsulation was performed thereby preventing a very accurate positioning. The attacked algorithm was a DES featuring no software counter-measures against DPA. For DPA and DEMA, 500 acquisitions and messages were necessary to infer the secret key.

While performing a DPA, it is expected that the right guess would yield the maximum peaks but, experimentally, strong differential peaks are often observed for wrong guesses. Such false alerts may even rise higher than the right spikes and confuse an attacker trying to make a final decision. This phenomenon stems

from the consumption model that underlies classical DPA. Indeed, the signature is usually supposed to be correlated to the data's Hamming weight. In reality this may not match each and every VLSI behavior as other subparts of the chip may also consume power in a correlated manner.

In the present case, a DPA spotted the right guess successfully but with many difficulties. As shown in Figure 7, there were even examples of wrong guesses (39) whose peaks were higher than the right one (15) in absolute value. The corresponding DEMA yielded correctly ordered spikes, smearing the wrong guess peaks and enhancing the right one.

Compared to DPA and for a relevant pinpointed area, experiments generally showed that DEMA tended to reduce the dispersion of peaks to the benefit of the right guess. In other words, the number of wrong guesses was reduced and the final decision made easier. DEMA can therefore be potentially considered experimentally at least as efficient as DPA, in absence of specific software counter-measure.

3.3 Modular exponentiation

The third experiment concerned an RSA exponentiation performed in a decapsulated smart card. The chip's visibility allowed a very close positioning of the sensor and the monitoring of the most energetic part of the EM field.

No software EM/PA counter-measures were implemented to protect the exponentiation (except a constant time implementation). As shown in Figure 8 (lower traces) the power traces did not suggest any apparent pattern that could have exposed the chip to a potential SPA.

Having observed this, the target of our study became the isolation a data-correlated location which is why the chip's surface had to be scanned. The success of this tedious operation was not guaranteed but a suitable point was finally found after several manual positioning attempts.

Two EM signals monitored at this point are shown in Figure 8 (upper traces). They look less noisy than the power curves and happen to contain patterns that leak the key. This illustrates how complementary SEMA and SPA can be.

4 Conclusion and Work in Progress

The purpose of this work was to find out if EM attacks can be implemented in practice; the answer is clearly positive.

Our experiments suggest that although more noisy, EM measurements finally yield better differentials than power signals. DEMA's SNR was higher than DPA's SNR and the correct guess identification was easier, as there were no false alerts due to erroneous peaks. The third experiment is particularly instructive as it shows that SEMA $\not\Rightarrow$ SPA. As is obvious, this shouldn't lead to the fallacious conclusion that SEMA is in some manner "more powerful" than SPA: we haven't encountered yet the opposite case (SEMA-proof, SPA-vulnerable) but nothing rules

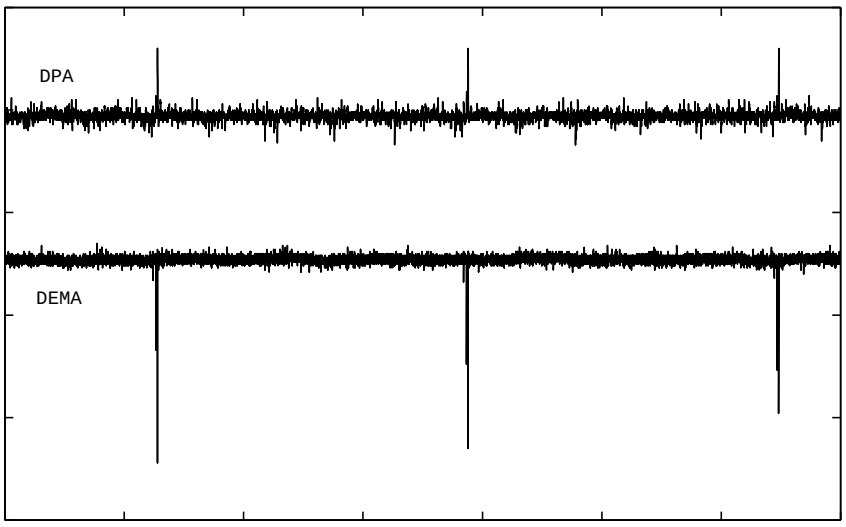

Fig. 6. DPA and DEMA right guess curves for ACOMP128.

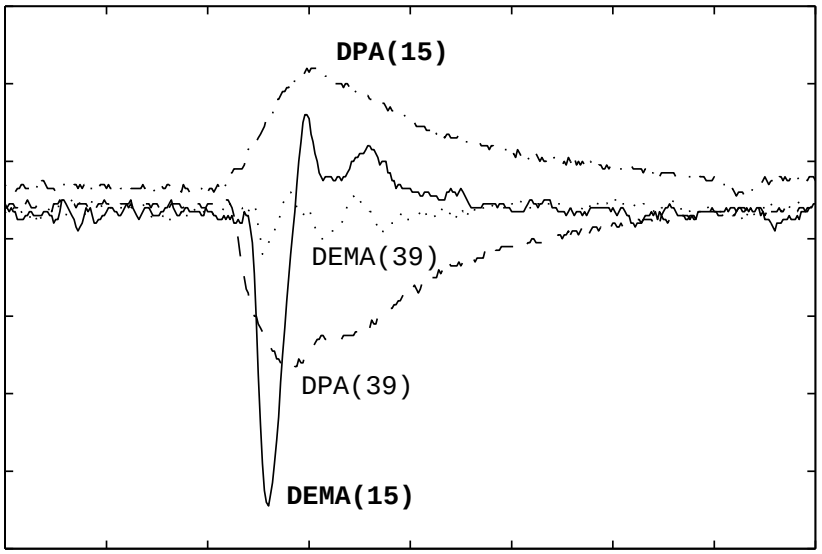

Fig. 7. DPA and DEMA DES curves for a right (15) and a wrong (39) guess.

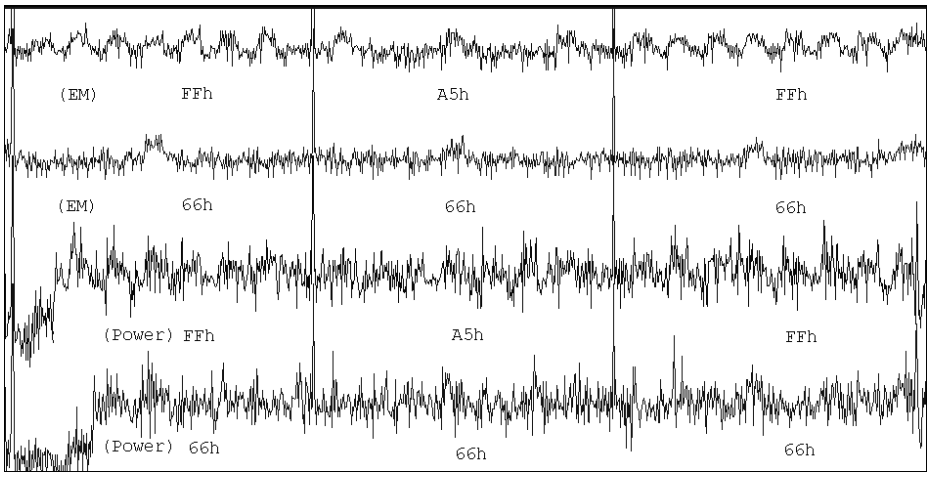

Fig. 8. EM and power traces for two different exponentiations involving three bytes of the private key : `FFA5FFh` and `666666h` (Same message and modulus). Artificial spikes delimitate the three-byte windows where patterns clearly appear.

out *a priori* that it might exist. In other words, when PA or EMA does not suffice alone, both can be attempted simultaneously.

EMA's advantage is definitely its capability of exploiting local information. This geometrical degree of freedom is useful as it allows to pinpoint the problematic spots that leak information. PA's major advantage is undoubtedly the relative simplicity of electric measurements as opposed to EM ones.

The manual scanning of the chip's surface performed during this work are, of course, non-exhaustive. The next step in our investigation is the implementation of automatic cartography tools. Note that chip-spots characterized by intensive power radiations (*e.g.* clock lines) do not necessarily leak data-correlated EM signals. Procedures for evaluating the likelihood of data-correlated leakage are described in [4]. By running such tests on EM signals collected at various locations on a given chip, a cartography of leakage probabilities can be performed. This would give an immediate bird-eye view of the potentially problematic spots in each chip and allow cross-platform comparisons.

Natural EM hardware counter-measures typically include an upper metal layer (*contain* the radiation), variable random currents, flowing through an active grid and generating noisy fields (*blur* the radiation[1]) and successive technology shrinks that regularly reduce the elementary transistors' size and make the functional areas more compact (*reduce* the radiation). Particular synthesis of problematic functions (coding ones as $\{1, 0\}$ and zeros as $\{0, 1\}$) tries to partially *cancel* the radiation.

[1] Such a counter-measure, called *surface oscillators*, is reportedly present in Clipper [3]

It is our opinion that the combination of such hardware counter-measures with particular software coding techniques that inherently prevent specific forms of leakage, provides an acceptable security-level for most commercial applications.

Acknowledgments

We are very grateful to David Naccache, Pascal Moitrel, Christophe Clavier and Marc Joye for their contribution and help which greatly improved the development of this study.

References

1. SEPI'*88, Primo simposio nazionale su sicurezza elettromagnetica nella protezione dell'informazione*, Rome (Italy), 1988.
2. SEPI'*91, Symposium on electromagnetic security for information protection*, Rome (Italy), 1991.
3. R. Anderson, M. Kuhn, *Tamper Resistance - a Cautionary Note*, Proc. of the Second USENIX Workshop on Electronic Commerce, USENIX Association, 1996.
4. J-S. Coron, P. Kocher, and D. Naccache, *Statistics and Secret Leakage*, Financial Cryptography 2000 (FC'00), Lecture Notes in Computer Science, Springer-Verlag, To appear.
5. Y. Gao and I. Wolff, *A new miniature magnetic field probe for measuring three-dimensional fields in planar high frequency circuits*, IEEE Trans. on Microwave Theory and Techniques, vol. 44 no. 6, pp. 911–918, 1996.
6. H. Handschuh and P. Paillier, *Reducing the collision probability of alleged* COMP128, In J.-J. Quisquater and B. Schneier, editors, Smart Card Research and Applications (CARDIS'98), vol. 1820 of Lecture Notes in Computer Science, pp. 380–385, Springer-Verlag, 2000.
7. T. Harada, H. Sasaki and Y. Kami, *Investigation on radiated emission characteristics of multilayer printed circuits boards*, IEICE Trans. Commun, E80-B, no. 11, pp. 1645–1651, 1997.
8. P. Kocher, J. Jaffe and B. Jun, *Differential power analysis*, In M. Wiener, editor, Advances in Cryptology – CRYPTO'99, vol. 1666 of Lecture Notes in Computer Science, pp. 388–397, Springer-Verlag, 1999. Also available at: `http://www.cryptography.com/dpa/Dpa.pdf`.
9. O. Kömmerling and M. Kuhn, *Design principles for tamper-resistant smartcard processors*, In Proc. of the USENIX Workshop on Smartcard Technology (Smartcard'99), pp. 9–20. USENIX Association, 1999.
10. M. Kuhn and R. Anderson, *Soft tempest: Hidden data transmission using electromagnetic emanations*, In D. Aucsmith, editor, Information Hiding, vol. 1525 of Lecture Notes in Computer Science, pp. 124–142. Springer-Verlag, 1998.
11. T. Messerges and E. Dabbish, *Investigations of power analysis attacks on smartcards*, In Proc. of the USENIX Workshop on Smartcard Technology (Smartcard'99). USENIX Association, 1999.
12. J-J. Quisquater and D. Samyde, *A new tool for non-intrusive analysis of smart cards based on electro-magnetic emissions, the* SEMA *and* DEMA *methods*, Presented at the rump session of EUROCRYPT'2000.