Observability of Stuck-at-Faults with Differential Power Analysis

Giorgio Di Natale, Marie-Lise Flottes, Bruno Rouzeyre

Observability of Stuck-at-Faults with Differential Power Analysis

Giorgio DI NATALE, Marie-Lise FLOTTES, Bruno ROUZEYRE
Laboratoire d'Informatique, de Robotique et de Microélectronique de Montpellier
Université Montpellier II / CNRS UMR 5506
161 rue Ada, 34392 Montpellier Cedex 5, France
{dinatale,flottes,rouzeyre}@lirmm.fr

Abstract

In this paper we propose an innovative method to test integrated circuits based on the use of Differential Power Analysis. We will show that this technique, classically used to perform attacks on cryptographic devices, is very effective in observing single stuck-at faults. Based on the observation of the current consumed by the circuit during net transitions, it does not require observing primary outputs of the circuit and allows the test of hard-to-observe faults. Conversely to Iddq, this technique is not sensible to process variation.

Keywords: Stuck-at-fault, Observability, Current-based testing, DPA

1 Introduction

Classically, stuck-at fault detection lies on two basic concepts: it requires excitation (activation) of the fault by setting up the circuit primary inputs such that fault-free and faulty values at the fault location are different, and propagation of the value of the node under test to a primary output. These two concepts are respectively referred to as control and observation [1].

Achieving both control and observation conditions is a difficult task for large circuits with numerous and possibly divergent/re-convergent propagation paths. Moreover, redundant circuits, which contain undetectable stuck-at faults, prevent either proper activation of the faults or propagation of the fault effects to observable outputs.

Control or observation of hard-to-detect fault can be achieved at the cost of long test pattern generation procedures. Alternatively, controllability and observability can be increased by the exploitation of Design for Testability (DfT) techniques, like the insertion of control/observation points for instance.

However when requirements on test generation time or chip area are too strict, and thus do not allow long ATPG executions or the introduction of DfT structures, circuit testing is affected because of hard-to-detect faults.

In this paper we focus on hard-to-observe faults. We propose a new observation strategy based on the measurement of the power consumed by the circuit during the application of test vectors. Conversely to traditional stuck-at fault test techniques where the logical value of the node under test needs to be propagated to an observable point, i.e. a primary output, we observe the effect of the fault through the power supply line without any requirements in terms of logic propagation. This technique exploits the fact that for CMOS technology there is a correlation between the logic data processed by the circuit and the power consumption of the device. In particular, CMOS device consummation is strictly related to the transitions of transistors. Power supply measures and analysis used for the proposed technique are based on the Differential Power Analysis (DPA) method, originally proposed to tamper crypto devices.

DPA is a form of side channel attack in which the attacker observes the power consumption of the crypto device in order to disclose the processed data, particularly the secret key used for coding [2]. DPA is an extension of the simple power analysis that involves statistical analysis of the power consumed by the device.

The technique we propose is different from classical current based testing like I_{DDQ} or I_{DDT}. I_{DDQ} testing [3] relies on measuring the supply current (I_{DD}) in the quiescent state (when the circuit is not switching). In contrast, I_{DDT} test methods measure transient supply currents ([4] [5] [6] [7]) When these test techniques rely on a golden device measurement (reference), they can be affected by process variations.

Another supply current-based test metric, called Energy Consumption Ratio (ECR), has been presented in [8]. This method is based on the measurement of averages of dynamic power consumptions. The advantage is that average dynamic currents are easier to measure than rapid transient currents and it does not require high-speed measurement circuitry. ECR relies on the fact that a fault alters the number and location of signal transitions that occur due to a change in input. In other words, a fault can alter the energy consumed by the circuit. It uses two pairs of vectors, which are alternated at the input of the circuit. ECR is the ratio of currents (or energies) consumed by the two transitions. ECR is immune to process variations as the effect of process changes affects both the numerator and the denominator and gets canceled.

Similarly to the ECR-based test strategy, the DPA-based test strategy proposed in this paper relies on a test metric related to the average power consumptions of the circuit-under-test computed after execution of an appropriate test set. Both strategies allow the detection of logically redundant faults or hard-to-observe faults that cannot be observed on primary outputs. Compared to ECR-based strategy, diagnostic of the stuck-at fault affecting the circuit is straightforward with our technique even in the case of logically redundant faults. Moreover, our method is based on the measure of dynamic currents that are more difficult to realize but that can better results also in the case of new technologies.

The paper is organized as follows: Section 2 introduces the basic concepts of the Differential Power Analysis method. Section 3 shows how stuck-at fault can be observed using DPA, while Section 4 presents preliminary results. Eventually, Section 6 concludes this paper.

2 Differential Power Analysis

The first goal of the DPA is to identify the secret key of a cryptographic device used for ciphering or deciphering information. The basic idea is to correlate the power consumed by the device and the encryption data including the key.

In practice, the supply current measurements of a large number of encryptions are recorded then divided over two sets by means of a selection function and a guess on the secret key. The difference between the typical supply currents of the two sets will approach zero for a wrong key guess, but has noticeable peaks if the correct secret key has been predicted.

We now introduce some theoretical issues that allow the reader to understand the principle underlying the DPA attack.

The transition on the output of a CMOS gate leads to a charge or a discharge of the parasite capacitance C on inputs of downstream gates. In particular, when the output switches from 0 to 1 (figure 1.a), the supply line sources a current and C is charged. Conversely, when the output switches from 1 to 0 (figure 1.b), the current coming from the input capacitance is discharged to the ground.

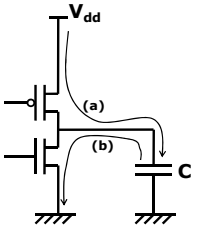

Figure 1: Power consumption in a CMOS circuit.

Power dissipated by the circuit can be monitored:

- by using a small resistor in series between Vdd and the true source. The current measured on this resistor is the sum of all the currents due to gates that switch from 0 to 1. This is the method used for the DPA;
- by measuring the electromagnetic field generated by the circuit. In this case the measure represents the sum of all the currents due to gates that switch either from 0 to 1 or from 1 to 0. This is the method used for the technique called Electromagnetic Analysis (EMA). Except measurements, the principle of the technique is similar to the DPA.

We consider the output of a gate within the circuit and we call it target node. We will detail later the characteristics of the target node and how it is chosen. We consider now a sequence of input patterns P_0, P_1, …, P_n that generate the transitions T_1 ($P_0 \rightarrow P_1$), T_2

($P_1 \rightarrow P_2$), ..., $T_n(P_{n-1} \rightarrow P_n)$ on the circuit inputs. A logic simulation of the circuit while monitoring the target node allows classifying these input transitions in two sets: (PA) transitions that make the target node to commute from 0 to 1 and therefore that make the target gate to consume (commutations from 1 to 0 are also considered in case of EMA), and (B) transitions that do not lead the target gate to participate to the power consumed by the circuit or the electromagnetic field (0 to 0, 1 to 1, and 1 to 0 in case of DPA).

We call *selection function* this partitioning criterion. Different selection functions can be used according to the power consumption models:

- Rising transition model: the set PA contains the input transitions that make the target gate to commute from 0 to 1. The set PB contains all the other input transitions. This function models in an accurate way a CMOS circuit, where power consumption is measured by means of a small resistor.
- Hamming distance models: the set PA contains the input transitions that make the target gate to commute from 0 to 1 and from 1 to 0. The set PB contains all the other input transitions. This function models in an accurate way a CMOS circuit for which the electromagnetic field is measured;
- Hamming weight model: the set PA contains the input transitions that make the target gate to commute from 0 to 1 and from 1 to 1. The set PB contains all the other input transitions. This function is used when it's not possible to calculate the transition on the target gate, but only the final value because the initial state is not computable due to the secret information processed by the circuit.

Figure 2 represents the power consumption of a circuit under attack when stimulated by several input vectors. Each rectangle represents the total power consumed by the circuit when a new vector is applied to the inputs. The set of consumptions has been split in two parts: in the upper part of the figure there are the PA vectors and the related consumptions while in the bottom part there are the PB vectors and corresponding consumptions. Corresponding to the set PA, there is a part of the consumption linked to the power consumed by the target gate. Obviously, the commutation from 0 to 1 of non-target nodes contribute to the power consumption of the circuit but input transitions that leads to such commutations are evenly distributed to set PA and PB. Mean consumptions related to set PA and PB are thus equal except for the contribution of the target gate.

In other words, since we classified the two sets in such a way that the set PA always leads to a component of power consumption that is not present in the set PB, the difference between the two mean powers computed from set PA and set PB must show a noticeable peak.

When the DPA is used for cryptanalysis on a crypto-chip, the target bit is typically chosen as the output of a gate whose state depends on both the plain text under ciphering (primary inputs) and the secret key. The function that links the input text, the key and the target bit must be known by the attacker since she has to perform a logic simulation of the function in order to build the two sets PA and PB. For instance, it is possible to choose the output bit of one Substitution Box of the circuit, this function being defined in the standard of the cryptographic algorithm. In this paper we do not focus on cryptographic devices. We suggest the reader to read the paper [2] for further information.

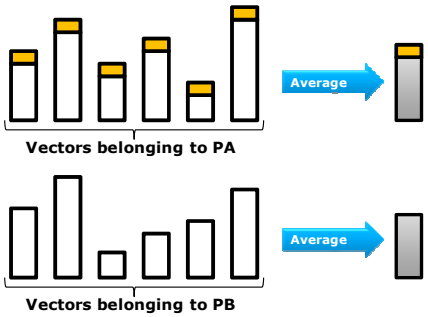

Figure 2: Power consumption after pattern partitioning

Let's now detail the whole DPA experience. Let's consider a generic logic function *f* that depends on two inputs I_1 and I_2 (see Figure 3), and a set of vectors $V = \{v_1, \ldots v_n\}$ used as sequence of input data on I_1.

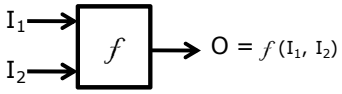

Figure 3: Generic function depending on two multi-bit variables (functional input I1 and key I2)

When DPA is performed during ciphering, I_1 corresponds to the functional inputs of the circuit on which the plain text is applied, while I_2 corresponds to the secret key.

Power consumption traces are recorded for the vectors belonging to the set V. These measurements performed on the real circuit depend on the actual value on I_2 even if this value is not known at this point of the experiment.

The logic value of the output O is computed for all the vectors belonging to V and a guess on I_2; PA and PB are built according to the logical value on O and a selection function.

The average consumption curves are computed from the traces recorded for vectors belonging to PA and to PB. If the hypothesis on I_2 is correct, PA actually includes the patterns that lead to a transition 0 to 1 on the target bit O while PB does not include any of these patterns. The difference between the curves obtained from PA and PB shows a peak in this case. On the contrary, when the curves are classed in PA or PB independently from the actual value on I_2 (wrong assumptions on the target bit value), the two average curves do not present any noticeable difference.

Classification of input vectors into PA and PB sets for every I_2 value allows identifying the guess that leads to the most significant peak on the differential curve.

3 DPA for SAF Observation

We propose to test a possible stuck-at-fault (SAF) by means of differential power analysis. Every stuck at fault in a circuit should bias the statistical results of the differential power analysis since the consumption profile of the faulty circuit should not match with the profile of the fault free circuit or the one of a circuit affected by another fault.

Considering the rising transition model, differential power analysis requires logical simulations for differentiating circuit input transitions that should lead to a transition on the SAF site, from all other input transitions.

Let consider a SAF site S and a downstream node N in the circuit sketched in figure 4. The SAF is logically observable on N, the logic function $N = g(I_A, I_B, I_C)$ (Figure 4.a) is thus affected and becomes $N = g'(I_A, I_B, I_C)$ in the faulty circuit (Figure 4.b).

The basic idea of the approach is to model the set of faults under consideration as a new circuit that performs the function of the original circuit or the function of a faulty circuit according to two input ports I_1 and I_2: $O = f(I1, I2)$; I_1 corresponds to the inputs of the original circuit, while I_2 is a virtual input allowing to set a code corresponding to the SAF under consideration, the "key" we try to discover.

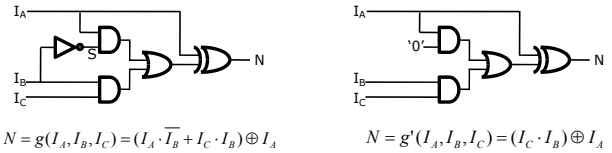

Figure 4: Fault-free (a) and faulty (b) functions

If the DPA is performed according to different "key" guesses, the only key that will leads to an appropriate classification of input transitions into PA and PB sets will correspond to the code of the SAF really presents in the circuit. In this case, PA will effectively include all input transitions that lead to a rising transition on N while PB will contain all other input transitions.

The target bit is a node on which the fault is logically observable: a primary output of the circuit or an internal node when it is not observable on a primary output.

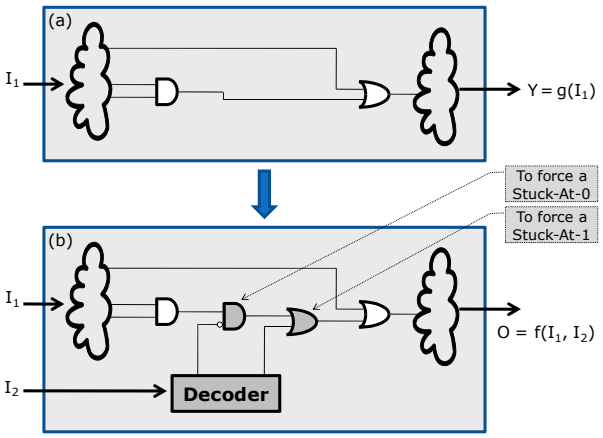

Figure 5: Stuck-at fault codification

Each SAF is modeled by inserting a control-to-0 or control-to-1 point (Figure 5.b) according to the fault, respectively SA0 or SA1. These control points are driven by the output of a decoder, which goal is to force either one single fault in the netlist or no faults at all according to input I_2. Again, the differential analysis of

the power traces must identify the right code (key I_2) that corresponds to the actual faulty circuit.

This modification of the netlist is purely theoretical, in the sense that it will be used to assign a numerical code to each possible fault during the logic simulation of the circuit (Table 1). In other words there is no testability improvement in the real circuit.

Table 1: Meaning of input I_2

I_2	*Meaning*
0	No faults (i.e., $f(x,0)=g(x)$)
1	Stuck-At 0, first gate, first port
2	Stuck-At 1, first gate, first port
3	Stuck-At 0, first gate, second port
4	…

4 Preliminary results

In this version of the paper we illustrate a first result we obtained applying the technique to a small circuit. In particular, we considered a combinational circuit with 8 input bits and 8 output bits that implements a Substitution Box (the main part of a circuit used in cryptographic standard devices [8] [10]). In this first experiment, the fault list contains the SA0 faults on the 8 possible inputs of the circuit.

The experimental environment is composed of three main parts: the electrical measurements, the logical simulation and the differential analysis.

A DPA suite has been developed, details can be found in [11]. This suite takes the VHDL description of the circuit and automatically generates all the scripts required to synthesize it with a given target library. It generates the scripts required to perform simulations of the circuit with a user defined vector set. Starting from the waveforms and the input vectors, the DPA Suite generates the DPA traces for each key supposition (i.e., for each fault). In particular, for each key supposition it classifies the waveforms (PA and PB) based on the logical function implemented by the circuit and according to the chosen power consumption model.

Note that normally, electrical measurements of the circuit under test are used to collect its actual consumption when stimulated with the input vectors. Nevertheless, due to the lack of real faulty circuits, consumption "measures" have been obtained from spice simulations (Synopsys Nanosym [12]) where the SA0 we want to detect has been electrically modeled using a direct connection of the node to ground.

Regarding the input vectors used for both logical simulation and electrical "measurements", the exhaustive test set has been considered (i.e, 256 input vectors).

Figure 5 shows the result of this experiment and the effectiveness of the method. The green curve (a) corresponds to the DPA trace of the circuit when the assumption of the correct stuck-at has been done. On the contrary, the grey traces (2) correspond to all others assumptions: SA0 on wrong inputs or no faults in the circuit.

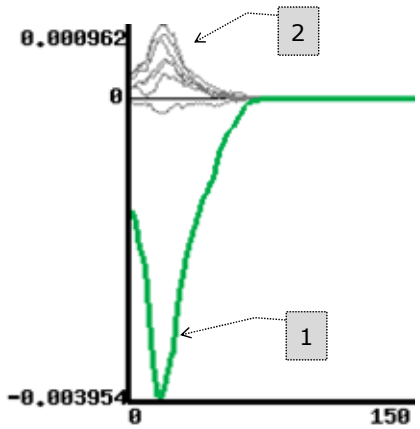

Figure 5: DPA traces under SAF assumption

5 Conclusions

In this paper we proposed an innovative method to test integrated circuits based on the use of Differential Power Analysis. The main advantage of this method is that it allows the observation of stuck-at faults that are not logically observable on primary outputs.

Preliminary experimental results showed the effectiveness of the approach on a small circuit and only one stuck-at fault.

Many issues are still under study. In particular, the selections of input vectors and target nodes have been manually performed in this experiment. In the future we will look for the automatic test pattern generation oriented to DPA.

Another issue we want to investigate is how to addresses the practical implementation of the DPA.

Conducting power analysis in practice (using real measurements) leads to several technical difficulties (e.g., obtaining noise-free measurements) that must be investigated.

6 References

[1] M. Abramovici, M. A. Breuer, A. D. Friedman, “Digital Systems Testing and Testable Design”, Piscataway, New Jersey: IEEE Press, 1994

[2] Paul Kocher, Joshua Jaffe and Benjamin Jun, “Differential Power Analysis", in Proceedings of Advances in Cryptology CRYPTO'99, Springer-Verlag, 1999, pp. 388-397

[3] Manoj Sachdev, “Current-Based Testing for Deep-Submicron VLSIs”, IEEE Design & Test of Computers, Volume 18, Issue 2, Mar/Apr 2001 pp. 76-84

[4] R. Z. Makki, S. T. Su, and T. Nagle, "Transient power supply current testing of digital CMOS circuits", in Proceedings of ITC '95, pp. 892-901, IEEE Computer Society, 1995.

[5] J. F. Plusquellic, D. M. Chiarulli, and S. P. Levitan, "Digital Integrated Circuit Testing using Transient Signal Analysis", in Proceedings of ITC '96, pp. 481-490, IEEE Computer Society, 1996.

[6] M. Sachdev, P. Janssen, V. Zieren, “Defect Detection with Transient Current Testing and its Potential for Deep Sub-micron CMOS ICs”, in Proceedings of ITC '98, pp. 204-213, IEEE Computer Society, 1998

[7] Manoj Sachdev, “Current-Based Testing for Deep-Submicron VLSIs”, IEEE Design & Test of Computers, Volume 18, Issue 2, Mar/Apr 2001 pp. 76-84

[8] “B. Vinnakota, W. Jiang, and D. Sun, "Process-tolerant test with energy consumption ratio", in Proceedings of ITC '98, pp. 1027-1036, IEEE Computer Society, 1998.

[9] Data Encryption Standard”, FIPS-Pub.46. National Bureau of Standards, U.S. Department of Commerce, Washington D.C., January 1977

[10] “Advanced Encryption Standard (AES)”, Federal Information Processing Standards Publication 197, November 26, 2001.

[11] G. Di Natale, M.-L. Flottes, B. Rouzeyre, “An Integrated Validation Environment for Differential Power Analysis”, in Proceedings of the 4th IEEE International Symposium on Electronic Design, Test & Application (DELTA), 2008 (to appear)

[12] http://www.synopsys.com