

Power Attacks on Secure Hardware Based on Early Propagation of Data

Konrad J. Kulikowski, Mark G. Karpovsky, Alexander Taubin
Reliable Computing Laboratory, Boston University
8 Saint Mary's Street, Boston, MA 02215
{konkul, markkar, taubin}@bu.edu

Abstract

The early propagation effect found in many logic gates is a potential source of data-dependent power consumption. We show that the effect and the corresponding power dependency can be targeted for successful power analysis attacks in cryptographic hardware. Many of the current balanced gate designs did not directly consider the effect and are vulnerable to power analysis attacks.

1. Introduction

Cryptographic algorithms are vulnerable to attacks which exploit the physical characteristics of their hardware implementations. Data dependent power consumption of a circuit is one such characteristic which can be used to infer information about the secret key of a cryptographic algorithm. Differential Power Analysis (DPA) attacks [1] have been shown to recover the complete key from unprotected hardware implementations of symmetric ciphers in minutes [2] by only passively observing and analyzing the power consumption of a cryptographic coprocessor. Power attacks pose serious and realistic security vulnerabilities which have seen a variety of countermeasure strategies.

A promising method for combating these attacks which has seen a considerable amount of design effort is based on the use of balanced gates. The balanced gate approach aims to make the power consumption of each logical gate equal for all valid input and output combinations and transitions ensuring that the power consumption of the gate is completely independent of the data that it is processing. Several such balanced gate designs have been proposed targeted for both synchronous (SABL [3], WDDL [4], DyCML [5, 6]) and asynchronous (BSDT [7], DIMS [8]) implementations. The different designs range in the level of their balance, overhead, and performance they provide. Most of the balanced gate designs have been

evaluated and tested against power analysis attacks and have been shown to provide adequate protection for the attacks and evaluations performed.

However, the balanced gate designs have not been evaluated against all possible sources of data-dependent power consumption. An additional source of the data-dependent power consumption results from an early propagation effect found in most gate designs. In designs with the early propagation effect a gate can evaluate before all gate inputs are valid. This effect creates data-dependent temporal switching behavior. Especially vulnerable are the “low cost” synchronous methods which claim to save hardware overhead or design effort by reusing existing unsecured standard-cell libraries to make larger balanced gates.

We start by analyzing the early propagation effect and show how it can lead to data-dependent power consumption. We then analyze the two most common cryptographic symmetric algorithms, the Advanced Encryption Standard (AES) and the Data Encryption Standard (DES), and describe a feasible attack which can be used to determine the key in these algorithms. We finalize by providing a survey and an analysis of possible methods which can be used to prevent this effect and attack.

2. Balanced-gate considerations

The internal data-dependent power consumption of individual traditional gates has been the primary focus of the initial development of the DPA attacks. In view of this the proposed attacks the sources of the internal imbalances have been thoroughly analyzed and examined. The findings led to several balanced gate designs most of which have been based on a dual-rail return-to-zero (RTZ) type implementations. Due to the dual-rail designs, an additional constraint of balancing the capacitive loads of the dual-rail wires has to be addressed [9]. More recently, in related design considerations for masked gates, it has been shown that glitches within a circuit may also be sources of data-

dependent power consumption [10]. We aim to show that in addition to the balanced internal switching, balanced routing, and glitch free operation requirements there is an additional requirement for balanced gate designs for cryptographic hardware. The additional balanced gate requirement is aimed at preventing data-dependent circuit path differences which create data-dependent power consumption.

3. Data-dependent circuit path differences

Many of the normal Boolean functions implemented in basic logic gates have the property that the values of their logical outputs can be uniquely determined without necessarily having the knowledge of the logical values of all the inputs. This functional property usually directly translates to physical implementations. If some logical values of the Boolean function implemented by a gate can be uniquely determined without having to know all of the logical values than the physical design can be minimized and the gate can physically evaluate without having to wait for all of the inputs. A physical gate can sometimes propagate its logical output early without having to wait for all of the logical inputs. This property is often called the “early propagation effect”. This effect is usually outside the consideration of most gate designs and the gates found in most styles and designs can evaluate before all of the inputs are available.




Figure 1. A simple circuit with data-dependent power consumption due to early propagation.

In traditional digital designs this early propagation effect is usually not seen as detrimental since it allows the design of a minimized gate and can potentially lead to performance improvements. However, this property can result in data dependent power consumption even for circuits implemented with balanced gates and with balanced routing. As an example, consider the circuit in Figure 1 and assume that it is implemented with perfectly balanced, dual-rail, glitch free gates with perfectly balanced routing. The inputs A, B, and C are primary inputs to the circuit and they are synchronized (for example by a register). The data dependence can be observed in the timing differences of the gate transitions.

Using a simple delay and power consumption model, the number of gate transitions with respect to time for all the possible input combinations are shown

in Table 1. Using the delay of a single-gate as the time scale, the number of gate transitions, and hence the power consumption of the circuit is dependent on the logical value of the A input. If the AND1 gate is implemented such that it has the early propagation effect, whenever the input A is a logical zero AND1 can evaluate and does not have to wait for the value of the other input to propagate through the OR1 gate. Thus when A is a logical-zero, both of the gates will evaluate at the same time. On the other hand if the A input is a logical-one, the AND1 gate cannot evaluate before the OR1 gate evaluates and hence the gates will evaluate in series. Thus as Table 1 shows, the number of gate transitions (and hence the power consumption) at time-1 is dependent on the logical value of the A input. If A is a logical-zero there are always two gate transitions at time-1, while if it is a logical one there will be one logical transition at time-1 and another at time-2.

Table 1. Number of gate transitions as a function of data.

A B C	Number of gate transitions at	
	Time 1	Time 2
0 0 0	2	0
0 0 1	2	0
0 1 0	2	0
0 1 1	2	0
1 0 0	1	1
1 0 1	1	1
1 1 0	1	1
1 1 1	1	1

The more precise and realistic behavior of the data-dependent power consumption is shown in the SPICE simulations of the circuit in Figure 2. The circuit from Figure 1 was implemented with the balanced SABL gate designs presented in [3]. The design of the SABL gate from [3] still has the early propagation effect for one of its inputs. The circuit was simulated for all of the possible input combinations. Figure 2A shows the power signatures of the circuit for all of the eight input combinations. As the figure shows, there are essentially two different power signatures. One signature directly corresponds to the case when A is a logical-one and the other to when A is a logical-zero. Despite the balanced gate design, there is still a clear power data-dependence which can be used to easily distinguish the logical value of A if the circuit structure is known.




Figure 2. (A) Simulated power curves of the circuit showing the dependence of the power signature on the logical value of A. Dashed lines are for the case when A is a logical-zero and the solid line shows the power curves when A is a logical-one. (B) Differential curve of the two types of power curves

The data dependent power consumption differences of the simple circuit are on the order of the power consumption of a single gate. An imbalance of such magnitude can be potentially exploited for power analysis. Although the temporal differences in the presented circuit were small (one gate delay), the difference between the delay of the primary and secondary inputs can be much longer in practice. The one OR gate from Figure 1 can be replaced by a larger and deeper combinational cloud creating a longer temporal difference which can distinguishable even under timing measurement uncertainties and errors. Traditional differential power analysis attacks can be adapted to target the logical value of A. The details and experiments which show how these data-dependent power differences can be used in a realistic power attack cryptanalysis of real symmetric ciphers like AES and DES is demonstrated in the next section.

4. Power analysis attacks

The temporal data-dependent power differences can be used in realistic attacks much in the same way as internal gate imbalances were used in the initial DPA attacks. Below is a description of an example of a simple exemplary attack of a specific structure found in AES and DES symmetric ciphers. Other more sophisticated attacks are certainly possible.

The Advanced Encryption Standard and the Data Encryption Standard (as well as other symmetric ciphers) have a structure in which the key addition (XOR) is followed by a nonlinear substitution box (Sbox). Additionally in both of the ciphers, the value of the data prior to the key addition is known. In AES this structure is observed at the input to the algorithm where the plaintext is first XORed with the key and immediately followed by a nonlinear transformation (Figure 3A). In DES, this structure and input can be observed at the output of the algorithm where one of the data outputs is XORed with the key and is fed

through to the output (Figure 3B). In this structure, since the value of the data prior to the key addition is known, if the value of the data after the key addition can be determined than the value of the key easily follows. The early propagation effect combined with the nonlinearity of the Sbox can make the determining the value of data after the key addition feasible.

Assume that the implementation of the Sbox which follows the key-addition is a combinational circuit (not a look-up table) and that for each of the primary inputs a structure similar to that in Figure 3 exists and that structure is known. If the gate-level structure is known then the power signature for each data bit value can be easily determined. The proposed attack is based on checking if the power signature corresponding to each bit entering the key-addition (which is known) corresponds to the predetermined power signature for that bit by using a differential distance-of-mean power analysis. If the power signature matches the input bit that that key-bit had to be a logical zero, since it did not flip the bit. On the other hand if the signature is opposite, one should conclude that the key bit is a logical one, since the input value was flipped prior to the Sbox.

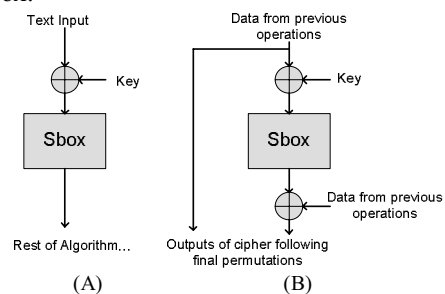


Figure 3. General structures where the data prior to key addition is known from either (A) input as in AES or (B) the output as in DES

Of course it is not possible to directly observe and distinguish the power signatures of the respective bits and circuits by simply looking at the overall power consumption of the whole algorithm. However, as in the traditional DPA attack, it is possible to “filter” the unwanted power consumption by taking the difference-of-mean test on the power consumption measurements by separating them according to the logical value of one of the bits prior to key-addition. Since the value of bit prior to the key addition is known, then power curves can be separated completely accurately. The only additional requirement is that the values of all the other bits are not correlated to the target bit and have a random distribution with respect to the partition of the target bit. Since the values of the target bits are known, the power curves can be actively selected from a large random recorded sample space to meet this

requirement providing a “filtered” power difference of the target bit. The polarity of the differential curve can be used to determine if the value of the data was flipped by the key addition.

The constraint of having a known gate-level structure of the Sbox implementation should not be problematic in practice and is not even strictly necessary. For most of algorithms there are many well established minimized implementations (i.e. AES Sboxes based on tower and sub fields) which can be inferred from the knowledge of the manufacturer and the objective of the design of the chip (i.e. low power, performance, etc.). Alternatively the structure can be inferred by adapting a template attack on a card in which the key is known. Performing the attack with a known key will provide the polarity of the differential trace and the structure of the design which can be used for the final attack.

One possible limitation for the proposed attack can be the number of required samples necessary to obtain a large enough distribution for which all logical values of bits are evenly distributed among the sets of the target bit and are not correlated to the target bit. As a result this attack might be more feasible for very compact implementations in smartcard and other mobile applications where the datapath of the algorithm is not a full width of the size of the cipher block. Smaller datapaths of 32 or even 8-bits are often found in many resource constraint applications such as sensor networks where the low power and not performance are the driving requirements.

5. Simulated attack

To demonstrate the necessity of preventing this imbalance in data propagation simulated power attacks were performed for networks with gates which have and do not have the path dependent power imbalance. The Sbox subcircuit of DES (Figure 4) was implemented and simulated with two versions of the SABL gate. The SABL gate was chosen for the experiment since it is synchronous thus providing a good timing reference from the clock and the early propagation can be eliminated by adding several additional transistors. The additional modifications needed to prevent the early propagation have minimal effects on the power consumption and the balance of the individual gates. Example AND/NAND implementations of the two versions of the SABL gates used are shown in Figure 5. Both implementations are well balanced designs and have virtually the same power consumption characteristics. The only difference is that the pull-down network of the first gate suffers from the early propagation problem. The

second version of the gate uses the enhanced differential pull-down network (DPDN) [11] which prevents the evaluation until all the dual-rail data inputs have valid logical values. The almost identical structure, balance, and behavior of the implementations allow a fair comparison of the results to determine the effect of the path dependent power consumption.




Figure 4. DES sub-circuit used in simulations.

The Sbox circuit was automatically synthesized using the Design Compiler from Synopsys and simulated on the schematic level (pre layout). Due to the nonlinearity of the function of the Sbox the resulting gate-level netlist had a similar arrangement of gates as in Figure 1 for each of the six key dependent inputs. That is, without any special coercion each of the six inputs of the Sbox was a primary input to at least one AND gate for which the other input was not a primary input to the circuit. The power data obtained from the simulations was used to perform a modified power analysis attack based on the presented path imbalances.

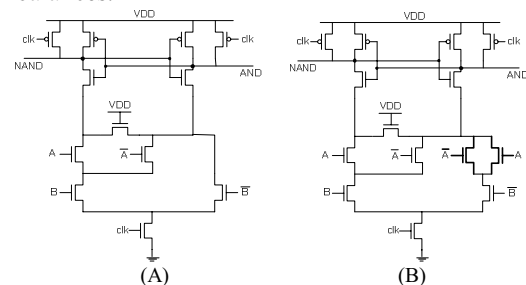


Figure 5. SABL AND/NAND implementations (A) with early propagation effect (B) with the enhanced DPDN without the propagation effect.

The power analysis attack results on the two implementations are shown in Figure 6. Each of the six key inputs of the Sbox test circuit was targeted separately. Since the structure for each input bit of the Sbox was similar to that of Figure 1, taking the difference of the power curves partitioned by the logical value of known bit prior to key addition (logical-one power consumption minus logical-zero power consumption) should result in a differential signature similar to that in Figure 2B. If indeed the input bit to the Sbox was not flipped by the key addition one should observe a differential power trace which has the negative polarity followed by a positive polarity. If the observed polarity of the differential

trace is opposite then it can be concluded that the input bits were flipped by the key addition.

For the first SABL implementation (which has the early propagations problem and is susceptible to the imbalances described) the attacks on all of the six key bits of the test circuit were successful. Since the expected polarity of the differential curve is known, the value of the keys can be easily read off the graphs by examining the polarity of the first spikes of the differential power curves. The first spike of the differential power curves represent the data-dependent power differences of the AND gate structures and the effect described. The differential power curves in Figure 6 (dashed curves) clearly reveal the value of the key as 110100 which is in fact the correct value. That is, based on the polarity of the first spike of the differential curves it is possible to determine if the input value of the data entering the key addition was flipped (key bit was one) prior to the Sbox input.

For the second implementation which had the same exact structure but with gates which had the enhanced DPDN the attack was not successful. The differential power curves for this implementation showed no data dependence with regard to this effect and cannot be used for a successful attack. As Figure 6 shows, the differential traces of this implementation (solid lines) are almost completely flat and featureless demonstrating the data-independent power consumption.

The simulations performed on the SABL gate show the necessity of preventing the early propagation for secure gates. The effect had been considered by Kris Tiri et al. in [11] for the improved version of the SABL gate, and speculations about the possible impact of the effect led to the development of the SABL gate from Figure 5B but the exact analysis was not performed. The SABL gates were easily modified and served as a good demonstration of the effect but other gate styles and methods are much harder to improve. A balanced-gate which does not have a mechanism for preventing this early propagation of data should not be considered secure.

6. Existing balanced-gate designs and countermeasure techniques

Many of the existing balanced gate designs have the early propagation effect and are vulnerable to potential attacks which exploit this data-dependent power consumption. Since these imbalances are manifested

with respect to time, synchronous implementation, where a good time reference is available, are especially susceptible. The synchronous implementations which use existing imbalanced standard cell libraries to make larger balanced gate designs are very vulnerable. As it was mentioned before, virtually all standard cell library designs do not take the early propagation problem into consideration and the resulting larger gates will usually have the same problem. For example, the composite gates proposed in [12] and the WDDL gates [4] have this weakness. The synchronous standard-cell designs can be buffered with adequate delay elements to ensure a constant path depth for all paths. However, such buffering or slack matching is costly and not completely accurate. Our experiments indicate that to ensure a constant depth in the nonlinear Sbox circuits may result in as much as 100% overhead. Increasing by a factor of two the existing 250%-700% overhead of the balanced designs (compared to standard unsecured static CMOS implementations) makes the approaches less appealing.

Balanced asynchronous gates have the advantage that they lack a precise timing reference which is needed to distinguish the demonstrated power differences. Certain fine-grained dual-rail asynchronous micropipeline styles explicitly check and require the validity of all data inputs prior to evaluation by a completion detection circuit. With careful consideration [13] balanced gates based on such asynchronous delay-insensitive design approach are naturally resistant to the power analysis attacks and do not have the early propagation effect. A balanced gate design based on the development of a standard-cell library approach and the fine-grained asynchronous structures has been presented in [7] and the corresponding synthesis and design flow in [14].

7. Conclusions

Balanced gate designs offer an attractive solution for preventing power analysis attacks. It has been previously shown that secure balanced gates require balanced internal switching, balanced routing, and glitch free operation. We add and motivate an additional constraint dealing with the path related data-dependent power consumption of the gates which can be eliminated by preventing the early propagation effect on the gate level. Many of the proposed balanced gate designs have not taken this consideration

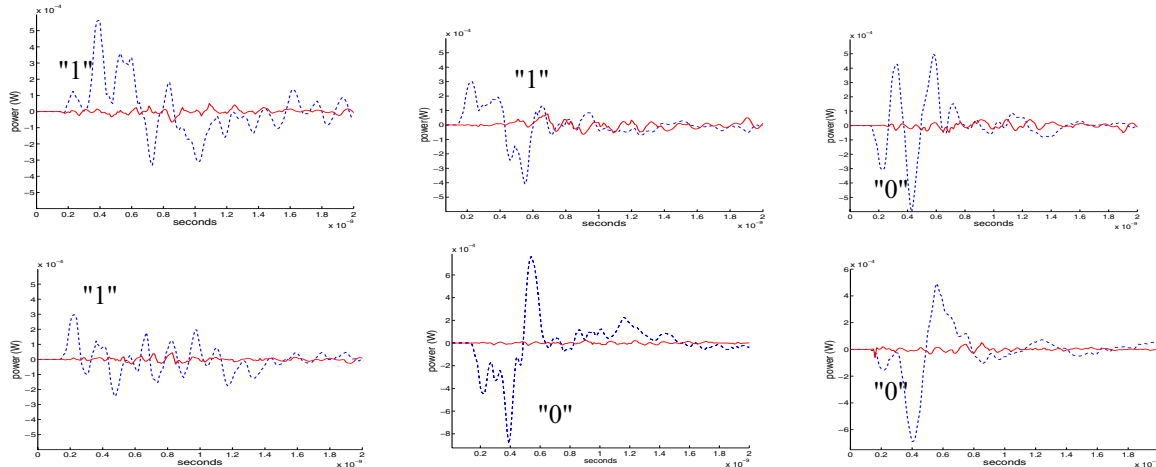


Figure 6. Differential power curves from power analysis attack targeting the path differences for each bit of the input to the Sbox. Dashed line is from an implementation which suffers from early propagation effect. The solid line is from power curves of an implementation which does not have the early propagation effect.

into account and are potentially vulnerable to power analysis attacks and require addition hardware or redesign. Most of the balanced gate designs can be modified or adjusted to eliminate the early propagation effect, but the modifications and adjustments usually require additional hardware overhead and make some “low cost” solutions based on existing standard-cell libraries potentially much more expensive.

8. Acknowledgements

This work was partially funded by Omnibase Logic Inc.

References

- [1] P. Kocher, J. Jaffe, and B. Jun, "Differential Power Analysis", CRYPTO, LNCS 1666, pp. 388-397, 1999
- [2] K. Tiri, D. Hwang, A. Hodjat, B.-C. Lai, S. Yang, P. Schaumont, and I. Verbauwhede, "Prototype IC with WDDL and Differential Routing - DPA Resistance Assessment". CHES, 2005
- [3] K. Tiri, M. Akamal, and I. Verbauwhede, "A Dynamic and Differential CMOS Logic with Signal Independent Power Consumption to Withstand Differential Power Analysis on Smart Cards". ESSCIRC, 2002
- [4] K. Tiri and I. Verbauwhede, "A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation". DATE, 2004.
- [5] M. W. Allam and M. I. Elmasry, "Dynamic Current Mode Logic (DyCML), A New Low-Power High Performance Logic Style," *IEEE Journal of Solid State Circuits*, vol. 36, pp. 550-558, 2001
- [6] F. Mace, F.-X. Standaert, J.-J. Quisquater, and J.-D. Legat, "A Design Methodology for Secured ICs Using Dynamic Current Mode Logic," PATMOS 2005
- [7] D. J. MacDonald, "A Balanced-Power Domino-Style Standard Cell Library for Fine-Grain Asynchronous Pipelined Design to Resist Differential Power Analysis Attacks," Master's Thesis, Boston University, 2005 http://reliable.bu.edu/Projects/MacDonald_thesis.pdf
- [8] F. Bousse, M. Renaudin, and F. Germain, "Asynchronous AES Crypto-Processor including Secured and Optimized Blocks," *Journal of Integrated Circuits and Systems*, vol. 1, pp. 5-13, 2004
- [9] K. Tiri and I. Verbauwhede, "Place and Route for Secure Standard Cell Design," CARDIS, 2004
- [10] S. Mangard, T. Popp, and B. M. Gammel, "Side-Channel Leakage of Masked CMOS Gates," CT-RSA, 2005
- [11] K. Tiri and I. Verbauwhede, "Design Method for Constant Power Consumption of Differential Logic Circuits," DATE, 2005
- [12] J. Jaffe, P. Kocher, and B. Jun, "Balanced Cryptographic Computational Method and Apparatus for Leak Minimization in Smartcards and other Cryptosystems." US Patent # 6510518, Cryptographic Research Inc., 2003
- [13] K. J. Kulikowski, M. Su, A. Smirnov, M. G. Karpovsky, and D. J. MacDonald, "Delay Insensitive Encoding and Power Analysis: A Balance Act". ASYNC, 2005
- [14] A. Smirnov, A. Taubin, M. Su, and M. G. Karpovsky, "Automated Fine-Grain Pipelining Using Domino Style Asynchronous Library," ACSD, 2005