Faculteit Industriële Ingenieurswetenschappen

master in de industriële wetenschappen: elektronica-ICT

Masterthesis

Side channel analysis on complementary flexible thin-film transistor technologies

Max-émile Meylaerts

Scriptie ingediend tot het behalen van de graad van master in de industriële wetenschappen: elektronica-ICT

PROMOTOR :
Prof. dr. ing. Kris MYNY

PROMOTOR :
prof. dr. ir. Nele MENTENS

BEGELEIDER :
ing. Jelle BIESMANS

Gezamenlijke opleiding UHasselt en KU Leuven

Universiteit Hasselt | Campus Diepenbeek | Faculteit Industriële Ingenieurswetenschappen | Agoralaan Gebouw H - Gebouw B | BE 3590 Diepenbeek

Universiteit Hasselt | Campus Diepenbeek | Agoralaan Gebouw D | BE 3590 Diepenbeek
Universiteit Hasselt | Campus Hasselt | Martelarenlaan 42 | BE 3500 Hasselt

2024
2025

Faculteit Industriële Ingenieurswetenschappen

master in de industriële wetenschappen: elektronica-ICT

Masterthesis

Side channel analysis on complementary flexible thin-film transistor technologies

Max-émile Meylaerts

Scriptie ingediend tot het behalen van de graad van master in de industriële wetenschappen: elektronica-ICT

PROMOTOR :

Prof. dr. ing. Kris MYNY

PROMOTOR :

prof. dr. ir. Nele MENTENS

BEGELEIDER :

ing. Jelle BIESMANS

UHASSELT KU LEUVEN

Foreword

In the modern world, where the advancement of and reliance on technology continues to accelerate at a relentless pace, the potential harm done by malicious individuals that manage to exploit these devices also increases in magnitude. I therefore think cybersecurity will continue to gain importance and I was eager to learn about how security is implemented on the lowest level of an electronic device.

This thesis delves into the challenges of creating cryptographic circuits on Low-Temperature Polycrystalline Oxide (LTPO), a potential technology for flexible chips, that is resistant to side-channel analysis (SCA). LTPO is currently used in displays to achieve a variable refresh rate and almost all the research available is on this application of the technology. I believe LTPO could be a strong contender for the technology used in flexible chips but only if a logic style is chosen that plays to its strengths. As you read this thesis, you will see that an approach where the different types of thin film transistors (TFTs) present in LTPO are combined to create CMOS-style gates reveals fundamental challenges for the creation of logic gates because of the inherent imbalance between these TFTs.

I spend considerable time focusing on details, like choosing the optimal supply voltage and trying to balance the gates for the SCA resistant design in the chosen approach, instead of trying out alternative logic styles. I believe there is a lot of potential for future research in finding a logic style that is both well suited to LTPO and can be made resistant to SCA.

I am glad that I chose this subject and believe it has deepened my understanding of SCA, its countermeasures and electronics design in general.

I would like to express my sincere gratitude to my promoter, Prof. Dr. Ing. Kris Myny for his continuous support and unwavering faith in me. I am also deeply grateful to my supervisor, Ing. Jelle Biesmans, for his guidance and endless patience.

Finally, I would like to thank my friends and family for their constant support throughout this research, their encouragement was a constant source of motivation.

Contents

List of Tables

List of Figures

Abstract

As new semiconductor technologies like thin-film transistors (TFTs) are developed, their future adoption in security-sensitive applications requires a thorough evaluation of their susceptibility to side channel attacks. These attacks bypass the mathematical security of a cryptographic device by exploiting physical information leakage to extract sensitive data. This master's thesis examines the feasibility of power analysis attacks and its countermeasures on a Low-Temperature Polycrystalline Oxide (LTPO) device. To accomplish this, standard cells were created using device models to create a test circuit of an Ascon s-box. A correlation power analysis (CPA) was performed on this circuit, which was able to extract the secret key. Wave dynamic differential logic (WDDL) was attempted as a countermeasure against the CPA. WDDL standard cells were created and their effectiveness in countering CPA was tested at different transistor sizes.

WDDL was insufficient in protecting against CPA, because of the difference in static and dynamic power, as well as timing behavior between the p-type low-temperature polycrystalline silicon TFTs and n-type indium–gallium–zinc-oxide TFTs used in LTPO. These results confirm the feasibility of power analysis on a cryptographic device made using LTPO technology and showcase the necessity of further research into specific countermeasures.

Abstract in het Nederlands

Naarmate nieuwe halfgeleidertechnologieën zoals dunnefilmtransistoren worden ontwikkeld, vereist hun toekomstige toepassing in veiligheidsgevoelige applicaties een grondige evaluatie van hun kwetsbaarheid voor zijkanaalaanvallen. Deze aanvallen omzeilen de wiskundige beveiliging van een cryptografisch apparaat door misbruik te maken van fysieke informatielekkage om gevoelige gegevens te extraheren. Deze masterproef onderzoekt de haalbaarheid van vermogensanalyse-aanvallen en de tegenmaatregelen daarvoor op een *Low-Temperature Polycrystalline Oxide* (LTPO). Om dit te realiseren, werden standaardcellen gecreëerd met behulp van componentmodellen om een testcircuit van een Ascon s-box te maken. Een *correlation power analysis* (CPA) werd uitgevoerd op dit circuit, waarmee de geheime sleutel kon worden achterhaald. *Wave dynamic differential logic* (WDDL) werd ingezet als tegenmaatregel tegen CPA. Er werden WDDL-standaardcellen gemaakt en hun effectiviteit als beveiliging werd getest bij verschillende transistorgroottes.

WDDL bleek onvoldoende om de CPA tegen te gaan, vanwege het verschil in statisch stroomverbruik tussen de LTPO-componenten. Deze resultaten tonen de haalbaarheid aan van vermogensanalyse op een cryptografisch apparaat dat gebruikmaakt van LTPO-technologie en benadrukken de noodzaak van verder onderzoek naar specifieke tegenmaatregelen.

Chapter 1

Introduction

As new semiconductor technologies, such as thin-film transistors (TFTs) are developed, their future adoption in security-sensitive devices and applications requires a thorough understanding of their vulnerability to side-channel analysis (SCAs). While traditional cryptanalysis focuses on identifying mathematical weaknesses in encryption schemes, SCA uses leaked information to break the encryption. This information includes electromagnetic emissions by the device, timing variations, acoustic attacks, cache attacks, fault injection attacks and much more. This master's thesis focuses on power analysis attacks (PAAs), which exploit the variation in power caused by the cryptographic operations of the device. In particular, a correlation power analysis (CPA) was performed on a test circuit for an Ascon S-box [1], implemented using Low-Temperature Polycrystalline Oxide (LTPO) technology. An S-box is a component of many cryptographic algorithms that obscures the relationship between the input and output data by performing substitutions, this introduces non-linearity and is crucial in delinking the in- and output of the cryptographic device.

The aim of this master's thesis is to examine the viability of performing an SCA on a circuit implemented using LTPO and to examine and implement countermeasures against these attacks. Wave dynamic differential logic, a hiding countermeasure that attempts to eliminate the information leakage at the logic gate level by implementing combining multiple CMOS-style standard cells into one gate, is examined in detail. This is crucial in assessing if LTPO, which is primarily used in displays [2], is a viable competitor for other TFT technologies that are more widely used for flexible integrated circuits today.

This master's thesis will start with a brief dive into TFTs: the different technologies used, their characteristics, advantages and disadvantages. This is followed by an overview of the different types of side-channel attacks, starting broadly before zooming in on the different methods used to perform power analysis attacks. The literary review ends with a broad overview of SCA countermeasures. After this, the choice of the supply is expounded upon. Next, the design of the standard cells is explained followed by the design of the test circuit for an ascon S-box. This is followed by the design of the WDDL-cells and circuit components. Finally, the results of the SCA on the test circuit with and without WDDL protection are shown and discussed.

Chapter 2

Literature Review

2.1 Flexible integrated circuits

Thin-film transistors (TFTs) are commonly used in displays, where they function as in-pixel switches and drivers. The key advantages of TFTs over traditional complementary metal-oxide-semiconductor (CMOS) transistors include the ability to be manufactured on large substrates at a lower cost per unit area and at lower processing temperatures [3]. This low-temperature processing is crucial because it enables direct integration onto various flexible substrates. This capability enables new and exciting product categories, such as foldable smartphones and rollable TVs, and lightweight electronics and Internet of Things (IoT) devices. For example, some modern smartphones achieve a high screen-to-body ratio by internally bending the screen, allowing it to reach the edge of the device.

The primary TFT technologies currently used in consumer electronics are amorphous silicon (a-Si), low-temperature polycrystalline silicon (LTPS), and amorphous metal-oxide semiconductors, most notably indium-gallium-zinc-oxide (IGZO) [3]. Each of these technologies presents a unique set of trade-offs in terms of performance, cost, and complexity. Organic transistors are also widely studied as a potential candidate for flexible circuits, particularly to complement n-type metal-oxide TFTs, as a high-performing p-type metal-oxide equivalent to IGZO has not yet been discovered [3].

Figure 2.1 provides a comprehensive overview of the four primary technologies, their properties and applications [3].

In recent years, Low-Temperature Polycrystalline Oxide (LTPO) has proven to be useful as a sophisticated display backplane technology, enabling features, such as a dynamic refresh rate [2]. Modern smartphones equipped with the technology are commonly able to vary their refresh rate between one and 120 Hz. LTPO is able to accomplish this by combining the benefits of its two underlying technologies: LTPS and IGZO. LTPS can operate at faster speeds, but suffers from a higher power leakage, while IGZO has an extremely low power leakage. IGZO has the additional limitation that no viable p-type TFTs have been discovered, whereas both n-type and p-type LTPS TFTs exist.

For this master's thesis, two device models were used to create standard cells in LTPO: a model for an n-type IGZO TFT and a model for a p-type LTPS TFT. Combining the technologies in

Parameter	a-Si	LTPS	Oxide	Organic
μ ($cm^2 V^{-1} s^{-1}$)	0.5–1	50–100	10–40	0.1–10
Process complexity	Low	High	Low	Low
Manufacturing cost	Low	High	Low	Low
Bias and light stability	Poor	Good	Fair	Poor
Intrinsic properties on mechanical stress, without stack optimization	Poor	Poor	Good	Fair-to-good
Semiconductor	n-type	CMOS	n-type	p-type (n-type possible)
TFT-to-TFT uniformity	Good	Low	Good	Low
Large-area uniformity	Good	Low	Good	Low
L-scaling (lateral device architecture)	–	Limited due to polycrystalline semiconductor (µm-range)	Deep-submicrometre demonstrated	Limited due to contact resistance and polycrystalline semiconductor (µm-range)
Backplane applications	Low-end and large-area display and imagers	High-end display and imagers, with on-panel circuitry	Low-to-high-end display and imager applications and large-area panels	Low-end backplane and circuit applications
Circuit applications	Low-end applications	High-end digital and analogue	Low-to-high-end digital, low-to-medium-end analogue	Low-end applications

Figure 2.1: Comparison of different properties and applications of the four primary TFT technologies [3]

this way is rather unconventional, as LTPS commonly consists of both n-type and p-type TFTs.

2.2 Introduction to Side Channel Analysis

In March 2003, as the United States was preparing its invasion of Iraq, pizza restaurants near the Pentagon noticed an unusually high number of pizza orders. The planning of the invasion caused a lot of people to stay in the office late as they had to meet urgent deadlines. Someone observing this sudden increase in people ordering pizza to the Pentagon could conclude that something significant, like the pending attack, was about to happen [4].

This scenario is an example of a side-channel analysis: the leakage of information, here the sudden spike in pizza deliveries, is exploited to predict a secret. No spies were used, no security was breached, no server had to be hacked. But the people who planned the attack had to work overtime and they also had to eat.

In cybersecurity, attacks can be split into two general categories: cryptanalysis, which exploits weaknesses in the mathematical algorithms used during the encryption, and side-channel analysis, which exploits unintended information leakage. The leaked information can range from changes in power consumption, timing variations, electromagnetic emissions or acoustic information leaked by electronic devices, to the quantity of pizzas ordered to US intelligence facilities. As mentioned above, this master's thesis focuses on power analysis attacks (PAAs). These attacks exploit the variation in power consumption of the integrated circuits used for the cryptographic functions of the device.

2.2.1 Simple Power Analysis (SPA)

The most basic form of PAA is the simple power analysis (SPA). This attack involves visually examining a single power trace measured during the cryptographic operation of the target,

revealing patterns that correlate to specific operations used in the encryption algorithm or to intermediate values in the chip.

Implementations of the Rivest–Shamir–Adleman (RSA) cryptosystem that do not have countermeasures against SCA are vulnerable to SPA attacks. The RSA decryption retrieves the original message by performing a modular exponentiation of the ciphertext with the secret key. In practice this calculation is implemented as a series of square and multiply calculations. When a bit of the private key is zero, only a square operation is performed, while the 'square' operation is followed by a 'multiply' operation when the key bit is one. [5]

Figure 2.2 shows how SPA can differentiate between these two operations on a power trace of an exponentiation, allowing an attacker to retrieve the private key.

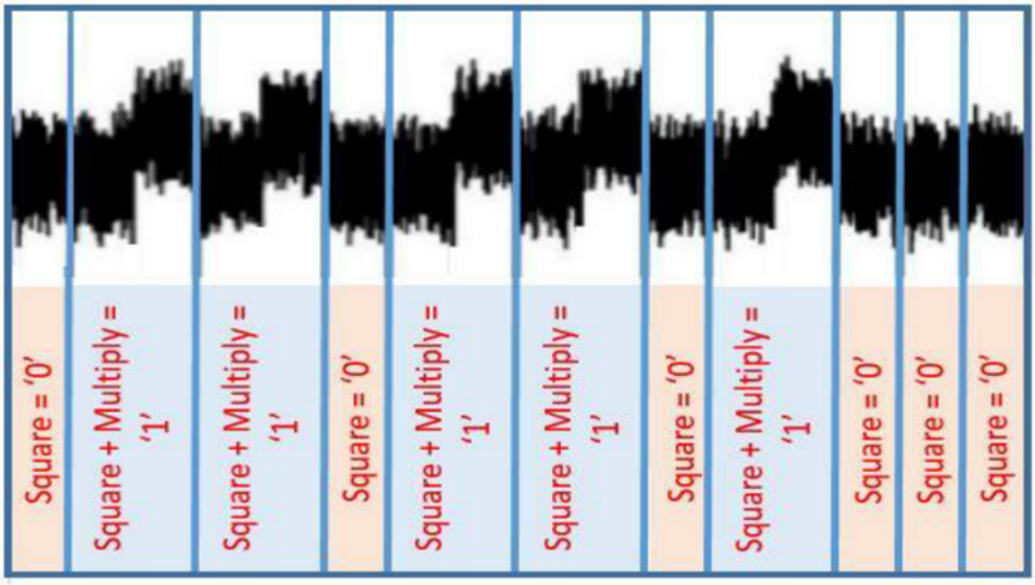

Figure 2.2: Result of an SPA on an RSA implementation that lacks countermeasures [5]

However, by ensuring that the operations performed are independent of the secret key, a designer can protect against this type of SPA. This can be achieved by inserting dummy 'multiply' operations. In general, SPA works best on simple circuits, like smart cards. In more complex systems, where many operations are executed in parallel, all drawing power from the same power rail, too much noise will be introduced. This makes it difficult to visually discern meaningful patterns.

2.2.2 Differential Power Analysis (DPA)

A much more powerful type of PAA is differential power analysis (DPA) [6]. In a classic DPA attack, the attacker collects hundreds to thousands of power traces, while the device encrypts known plaintexts with a secret key. To verify a key guess, the attacker uses it to predict the value of a single intermediate bit within the algorithm. The traces are then partitioned into two sets based on this predicted bit value, and the difference of means of the two sets is calculated. A significant spike in this difference reveals a correct key guess.

This master's thesis uses a more advanced type of DPA, correlation power analysis (CPA). In CPA the attacker creates a power model to hypothesize the expected power consumption for each operation, often the Hamming weight or Hamming distance of the input or output of part of the

cryptographic circuit, such as an S-box, is used. These expected power values are then compared to the actual power traces by calculating a statistical score, such as the Pearson correlation coefficient. The correlation coefficient for the real secret key will be significantly higher than those of the other key guesses. CPA performs better and requires fewer traces than regular DPA [7].

Power analysis can be combined with other types of leaked information such as timing variations or emissions of electromagnetic radiation to perform a high-order DPA attack [8]. In addition, machine learning techniques can be used to help analyze the data or construct better leakage models [9]. These attacks are more resistant to countermeasures.

2.3 SCA countermeasures and mitigation strategies

Countermeasures against SCA can be broadly categorized into masking and hiding techniques. Masking attempts to break the correlation between the processed data and the executed algorithm, while hiding countermeasures aim to break the correlation between the processed data and leaked information, such as the power consumption [10]. Lastly, implementing countermeasures in software offers an inexpensive and flexible solution, but one that is often insufficient by itself. Strategically combining these methods is required for effective protection against SCAs.

2.3.1 Masking techniques

Boolean masking primarily uses the bitwise exclusive-OR (XOR) operation between the mask and the unmasked secret to obtain the masked secret. This makes it well suited for linear logical operations, like XOR. However, for non-linear operations (which include AND and OR) and for arithmetic operations, it requires complex logic or pre-computed tables to prevent security leaks.

Arithmetic masking uses modular addition to obtain the masked secret, making it well suited for algorithms that rely on arithmetic operations, such as addition and multiplication. Algorithms that combine Boolean and arithmetic operations require both masking schemes and must be able to securely convert a variable between the two representations. This is not a trivial problem, [11] shows how early proposals of two conversion algorithms were not side channel secure.

Although first-order masking schemes can defend against first-order DPA and CPA attacks, higher order differential power analysis (HO-DPA) attacks are still able to break their security. This is because of the assumption that an attacker can only evaluate the power consumption caused by one share of the split, either the mask or the masked secret, at a time, while HO-DPA is able to measure the power consumption of both shares at different moments in time, and statistically combine these measurements. The solution is to split the secret into more shares. This exponentially increases the difficulty of the attack.

2.3.2 Hiding techniques

Hiding can be accomplished either by introducing enough noise so that the data dependent variation cannot be differentiated, or by making the power consumption completely uniform and independent of the data [12], [13].

Time randomization can be introduced through clock signal randomization or by random delay insertion [14], [15].

Amplitude randomization can be achieved through various means. A random number generator can be used to create a variable resistance on the power supply circuit, varying the supply voltage [13]. Noise generators can also provide an efficient and simple solution in reducing the information leakage [16].

DPA-resistant logic styles attempt to make the power consumption of the circuit uniform and independent of the data on the level of the logic gates. Most of these logic styles are based upon dual-rail pre-charge (DRP) logic [17]. In a DRP logic style, every one-bit signal 'x' is represented by both the signal itself and its complement, often referred to a 'x_true' and 'x_false'. The logic operates on a two-phase cycle. First all signal rails, including both 'true' and 'false' rails, are brought to a uniform state, typically a logical zero, in a precharge phase. Then, the inputs are applied to the gate, causing exactly one of the output rails to transition from zero to one in an evaluation stage. Because there is always a single transition in each phase, power consumption is theoretically independent of the data being processed.

Sense Amplifier Based Logic (SABL) uses fully custom DRP logic gates based on the principle of sense amplifiers [18]. One major advantage of this style is that unlike CMOS based logic styles, it does not suffer from any "glitches", which cause extra power consumption and could thus lead to leaked information. Because of their differential nature, the sense amplifiers are also resistant to variations in supply voltage and operating temperature.

Wave Dynamic Differential Logic (WDDL) is a DRP logic style that combines complementary CMOS gates as shown in Figure 2.3. To create a WDDL AND-gate, a CMOS AND gate is combined with its complement, a CMOS OR-gate. The desired function is given the 'true' power lines as inputs, while the complement is given the false power lines as inputs. A WDDL circuit can be thought of as two complementary subcircuits as shown in Figure 2.4. One advantage of WDDL is that the precharge signal must only be applied at the beginning of the circuit and is propagated through it in a wave. This significantly simplifies the routing process as there is no need to distribute the precharge signal to every gate independently. This has the additional advantage that there is no large load on the precharge signal and lowers the peak supply current during the precharge stage, which lowers the supply bounce [19].

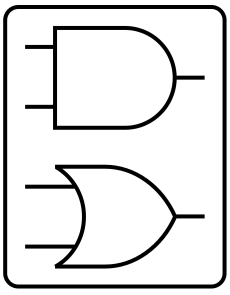

Figure 2.3: Symbolic representation of a WDDL AND-gate

2.3.3 Software-level countermeasures

Earlier in this thesis, it was shown how SPA can differentiate between the different operations required in RSA decryption. This is an example of an attack that can be prevented by using

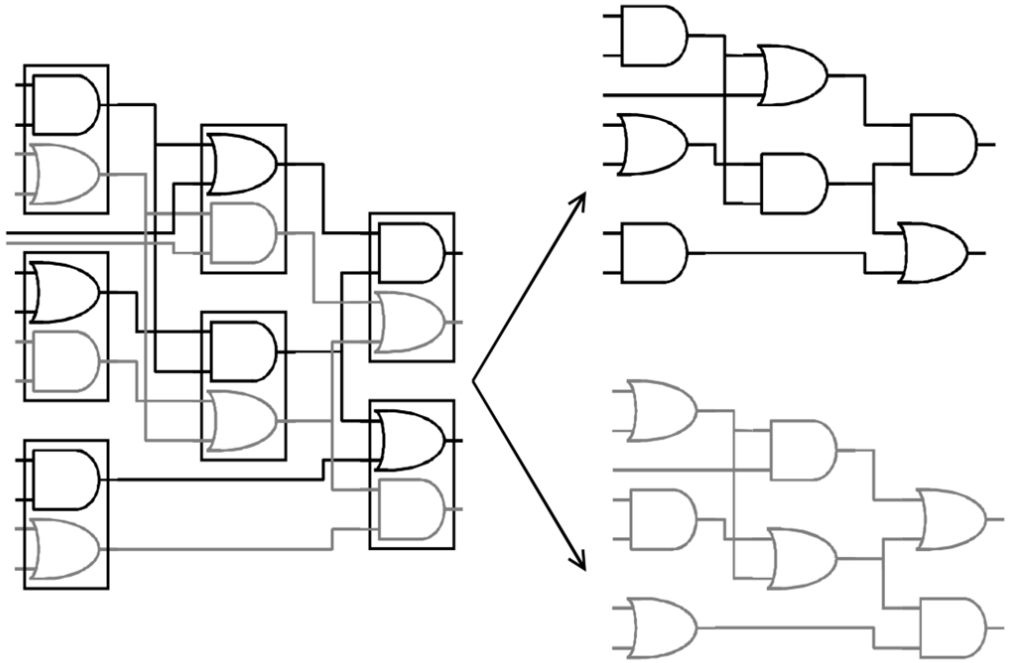

Figure 2.4: WDDL circuit split into subcircuits

software built to protect against SCA.

This can be achieved through constant-time programming, which ensures the program's execution time and the pattern in which it accesses memory are completely independent of the secret data. Another technique is to implement masking through software instead of using specialized hardware. This is often accomplished by splitting the secret value into shares, executing the cryptographic algorithm on each share and recombining them afterwards [20].

Chapter 3

Circuit design

3.1 Supply voltage and inverter balancing

For this thesis, two models based on real characteristics of devices were used: one for an IGZO n-type TFT and one for an LTPS p-type TFT. These TFTs were then combined into complementary TFT logic gates. This leads to several complications, including a different reaction to changes in supply voltage between two types of TFT. Additionally, the n-type TFT has a threshold voltage of 0 V, and the p-type TFT has a threshold voltage of -2.95 V, leading to a skewed switching threshold, asymmetric noise margin and asymmetric switching speeds.

To find the optimal supply voltage that minimizes these effects while achieving the best trade-off between power consumption and performance, different voltages were tested. An inverter was constructed and balanced for every supply voltage that was tested. Balancing was done by achieving equal rise and fall times for the inverter with a fan-out of four. The resulting inverters were then used in several testing circuits to measure their performance.

Rise and fall times were measured for both a fan-out of one and four, a common metric in electronics [21]. Noise margins, switching threshold voltages, and gain were derived from the voltage transfer characteristics of the inverters. Lastly, ring oscillators with a length of 51 inverters were used to measure the inverter delay, as well as the average power consumption.

To compare the performance of these transistors at different supply voltages, the **Energy Delay Product (EDP)** was used as a figure of merit. EDP is calculated as the average power consumption multiplied by the square of the inverter propagation delay:

$$\text{EDP} = P_{\text{avg}} \times t_{\text{p}}^2$$

Table 3.1 shows the supply voltage along with the optimal width for each transistor at that voltage, the propagation delay and the resulting EDP values. Table 3.2 shows the high and low noise margins expressed as a percentage of the supply voltage and the switching threshold voltages for the inverters with the sizes shown in Table 3.1 at each supply voltage.

The complete table with all measurements and calculations can be found in Appendix A. A five-volt supply voltage was chosen because it had the lowest EDP and showed good noise margins and a switching threshold compared to the other voltages. Rephrase for clarity: "Noise margins

Table 3.1: Inverter sizes and properties at different voltages

VDD (V)	$w_{n\text{-}TFT}$ (µm)	$w_{p\text{-}TFT}$ (µm)	t_p (ns)	P_{avg} (µW)	EDP (10^{-21})
3	2	3.7	1819.246	0.009	30.26
4	2	2	320.227	0.067	6.87
5	3	2	120.549	0.460	6.68
6	3.3	2	81.025	1.604	10.53
7	3.2	2	73.871	3.597	19.63
8	3.4	2	75.521	6.541	37.31
9	3.7	2	77.983	10.373	63.08
10	4	2	78.431	15.322	94.25
12	3.9	2	72.665	28.152	148.65
14	3.5	2	62.824	45.207	178.43
16	3.1	2	53.073	66.804	188.17
18	2.9	2	44.822	95.946	192.76
20	2.9	2	38.022	136.673	197.58
25	2.9	2	27.032	292.586	213.80
30	2.9	2	20.727	555.655	238.71

and inverter gain improved slightly with increasing supply voltage up to 7 V, after which they became increasingly asymmetric." The switching threshold voltage improved with increasing supply voltage until it reached 10 V.

3.2 Standard cell design and evaluation

Figure 3.1 illustrates the test circuit designed for performing SCA on an Ascon S-box. The circuit accepts five plaintext bits and five key bits, which are combined using a bitwise XOR operation to generate the S-box input. Registers are placed between each stage to allow for controlled processing, with their operation governed by enable and reset signals.

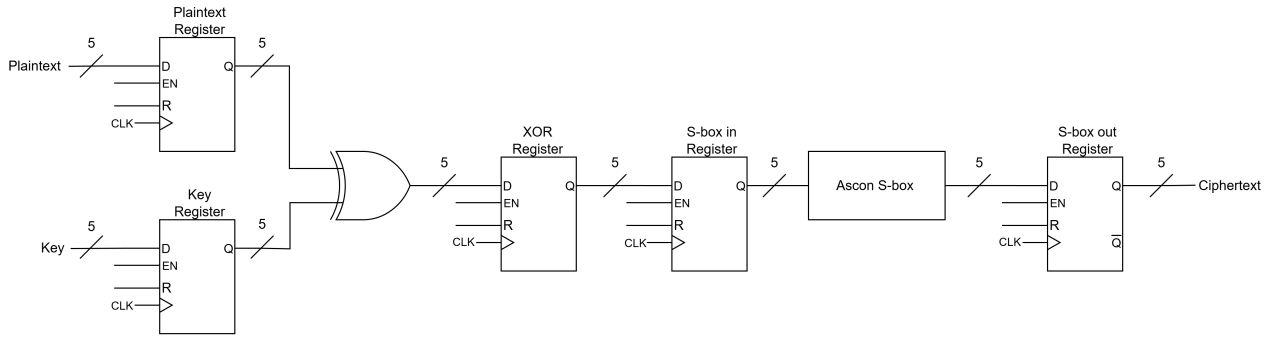

Figure 3.1: Positive edge-triggered D-type flip-flop

The netlist for the test circuit was synthesized from the Verilog description of the circuit using Cadence Genus. This tool uses a library file, which describes the timing, power consumption, logical function, etc. for every standard cell in the library, to create the optimal netlist for the circuit. Due to timing constraints, no custom library file was created for the custom LTPO standard cells. Instead, an existing CMOS technology library was used. This resulted in netlists that might be suboptimal.

Table 3.2: Reliability characteristics at different voltages.

VDD (V)	NMH (%)	NML (%)	V_{inv} (V)	Average Glitch (%)
3	39.44	24.43	1.24	34.81
4	38.76	29.80	1.76	20.83
5	38.47	32.78	2.28	13.48
6	37.35	35.17	2.83	11.26
7	35.89	36.75	3.40	12.17
8	34.63	37.22	3.95	13.90
9	32.88	36.81	4.49	14.55
10	30.99	35.66	5.03	13.90
12	26.34	33.69	6.19	11.22
14	21.48	32.98	7.44	8.54
16	15.82	34.67	8.74	6.40
18	10.68	37.17	9.98	5.42
20	7.85	37.87	11.09	4.75
25	4.11	38.09	13.85	3.63
30	5.11	35.20	16.64	2.91

The standard cells in this thesis were constructed using conventional CMOS topology. P-type LTPS TFTs were used for the pull-up network (PUN), while n-type IGZO TFTs were used for the pull-down network (PDN).

At the five-volt supply voltage the width of the n-type TFT must be 1.5 times larger than that of the p-type TFT to balance the rise and fall times of an inverter. The sizing was kept as small as possible, resulting in a smallest p-type TFT with a gate length of 2 µm and a gate width of 2 µm. The transistors in all logic standard cells were sized to balance the drive strength of the PUN and PDN, equalizing the rise and fall times. By increasing the width of transistors in series, the total resistance is made equivalent to the resistance of a single transistor in the inverter.

These methods were used to construct the required NAND, NOR, XOR, XNOR, two AND into an OR, inverted (AOI22) and one AND into an OR, inverted (AOI21) standard cells and scale their transistors. An implementation using transmission gates was also created for the XOR and XNOR standard cells. Figure 3.2 shows the schematic of the transmission gate implementation of the XOR circuit, which requires four fewer transistors than the implementation based on complementary logic.

Positive edge-triggered D-type flip-flops were implemented using transmission gates, as shown in Figure 3.3. These flip-flops combine a negative level-sensitive latch and a positive level-sensitive latch in a master-slave configuration [22].

3.3 WDDL component design

As previously discussed, WDDL is a dual-rail logic style, in which the precharge wave is propagated throughout the circuit in a wave. Consequently, each logic gate must handle both the true and false signals for all its inputs and outputs, where the true signal is the signal required in an unprotected, non-dual-rail logic style and the false signal is its complement (during the evaluation phase). Figure 3.4 shows the implementation of a WDDL AND-gate. It is made using

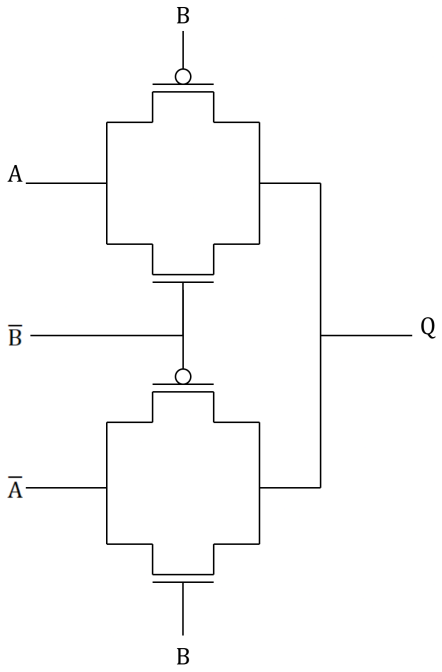

Figure 3.2: Schematic of the XOR standard cell using transmission gates

a NAND standard cell followed by an inverter for the true inputs and a NOR-gate followed by an inverter for the false inputs.

Inverted gates, such as a NAND or NOR, cannot be implemented by combining a NAND and a NOR standard cell. While these standard cells are complementary, they would prevent the propagation of the precharge wave, as their output is high when both inputs are low [23]. Instead, the outputs of the WDDL AND-gate can be switched, as shown in Figure 3.5 to create the inverted output while maintaining the ability to propagate the precharge wave.

Similarly, WDDL XOR and XNOR gates cannot be implemented by combining simple XOR and XNOR standard cells, as the XNOR would stop the propagation of the precharge wave. Instead, these gates can be implemented by combining AOI22 and OAI22 standard cells, essentially expressing the logic functions as a sum of products, as shown for the XOR gate in Equation (3.1).

$$A \oplus B = (A \cdot \overline{B}) + (\overline{A} \cdot B) \tag{3.1}$$

Figure 3.6 shows the implementation of a WDDL XOR gate. Each term in Equation (3.1) requires two TFTs per factor. The entire gate thus requires twenty TFTs in total. For a three-input XOR the sum of products is given by Equation (3.2).

$$A \oplus B \oplus C = (A \cdot \overline{B} \cdot \overline{C}) + (\overline{A} \cdot B \cdot \overline{C}) + (\overline{A} \cdot \overline{B} \cdot C) + (A \cdot B \cdot C) \tag{3.2}$$

Each term of Equation (3.1) requires two TFTs per factor, which results in 52 TFTs in total. A three-input XOR can also be implemented by combining two two-input XNOR gates which only requires 40 TFTs in total. Because of this it was chosen not to implement three-input XOR or XNOR gates. Implementations of WDDL AOI22 and AOI21 gates were created, requiring 20 and 16 TFTs respectively, compared to 36 and 24 TFTs if they were made as a combination of

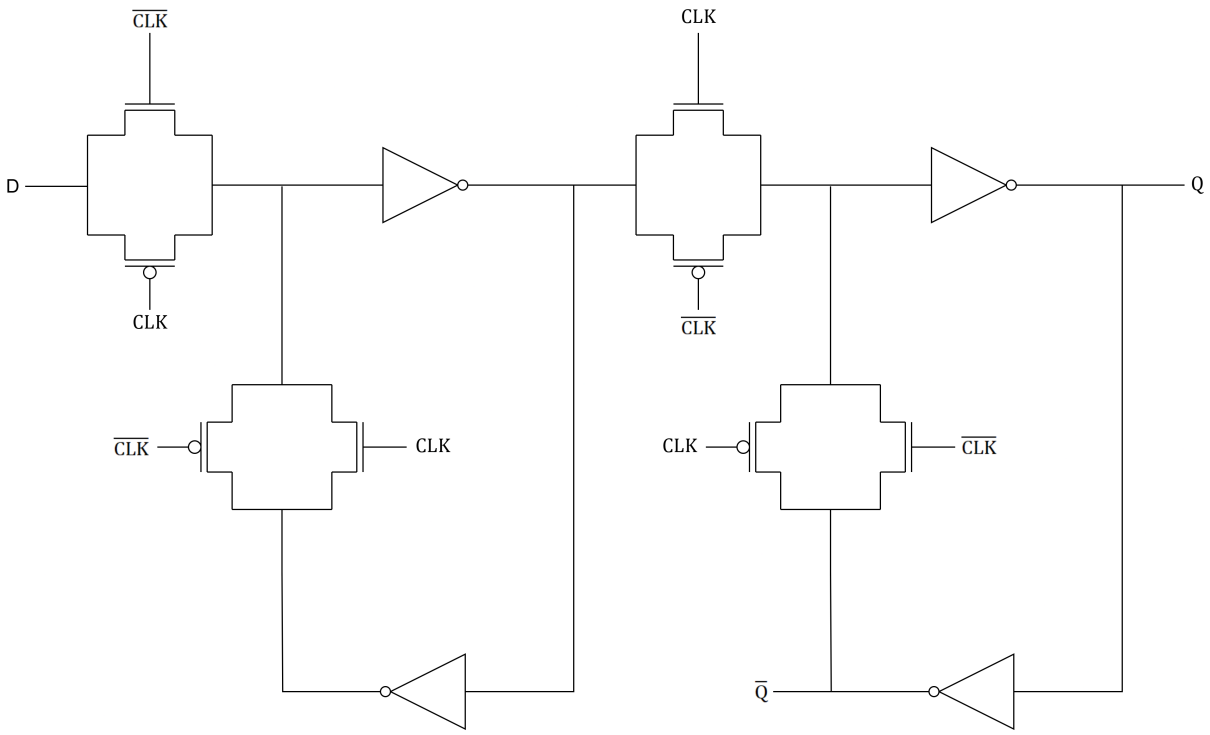

Figure 3.3: Positive edge-triggered D-type flip-flop

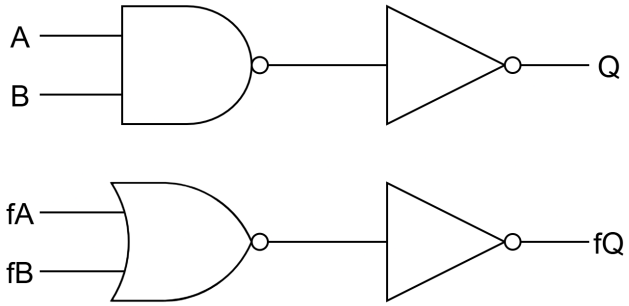

Figure 3.4: Logic gate implementation of a WDDL AND-gate

WDDL AND and OR gates.

The WDDL precharged inputs are created at the input of the WDDL circuit. Figure 3.7 shows the precharge operator in the red, dotted rectangle, consisting of two NOR standard cells, one receiving the input and the other its complement. The other input of each gate is connected to the precharge signal. To propagate the precharge wave beyond a register, precharge operators can be added to the outputs of the register as shown on the left side of Figure 3.7. This method is not preferred as it requires the precharge signal to be applied after every register. Instead, WDDL registers were implemented by combining two flip-flops to create a Master-Slave DDL register, as shown on the right side of Figure 3.7. The precharge and evaluation phases now last an entire clock cycle, requiring the clock frequency to be doubled for the same data rate.

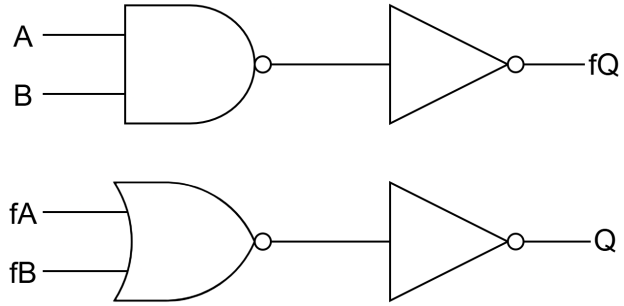

Figure 3.5: Logic gate implementation of a WDDL NAND-gate

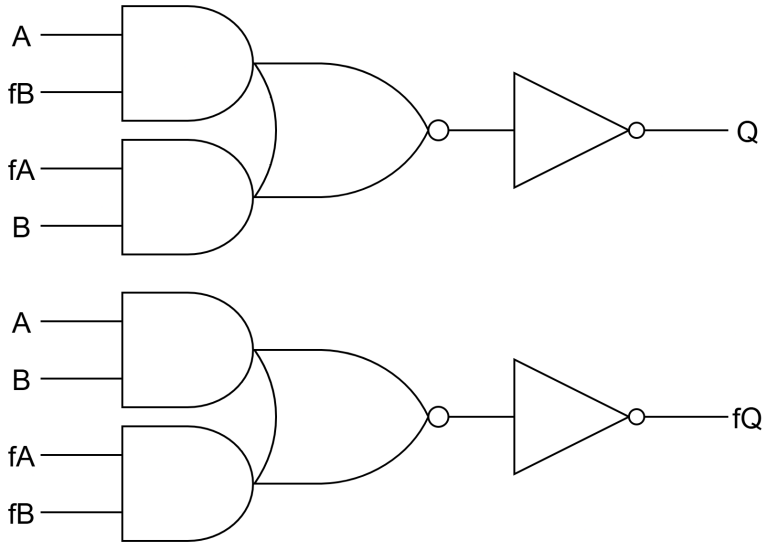

Figure 3.6: Logic gate implementation of a WDDL XOR-gate

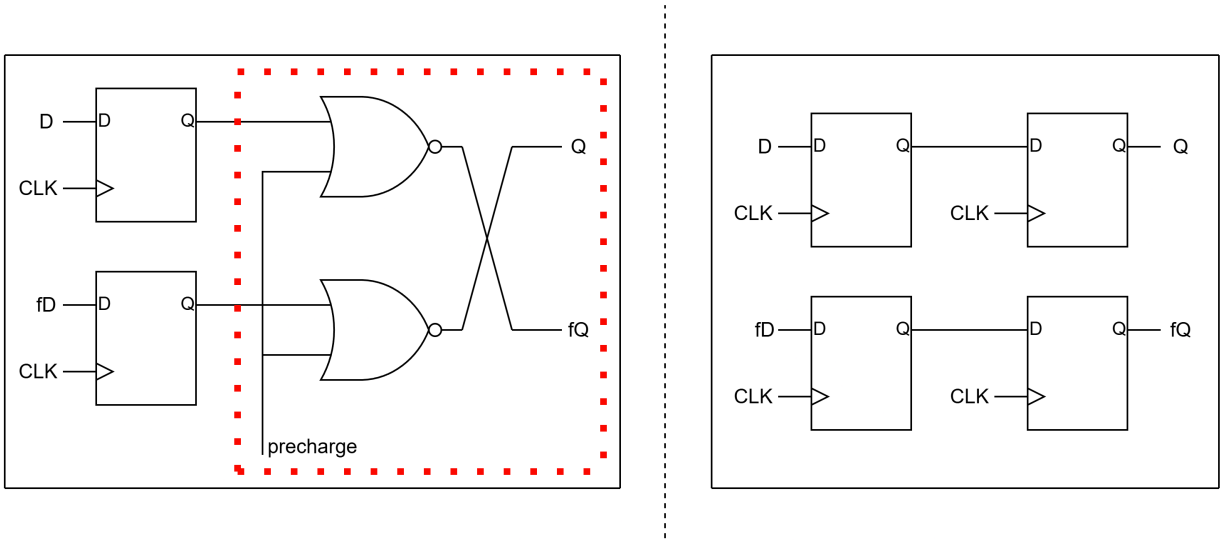

Figure 3.7: The different methods of implementing a WDDL register: with the use of precharge operators (left) or as a master-slave configuration (right)

Chapter 4

SCA on the unprotected S-box: methods and results

4.1 SCA on the unprotected S-Box: method

To perform the side channel analysis power traces were collected from a simulation done using Cadence Virtuoso. Figure sca-test-circ shows the circuit used. Plaintext is created by 'bus-random', a VerilogA logic block that produces five random bits at its output, while the key is chosen beforehand. The simulation was performed for 10,040 µs, with a clock period of 10 µs, resulting in dataset of a thousand power traces, after excluding the first few invalid traces. The results of the simulation are exported to a comma-separated-values file, containing the power usage of the circuit, but also the clock, plaintext, key and resulting ciphertext signal.

A correlation power analysis was performed with the help of a Python script. This script first verifies the correct operation of the S-box circuit by obtaining the key and plaintext values and using them to calculate the expected ciphertext output of the S-box. After this, the script separates the complete power consumption data into individual power traces, one for each clock pulse. The voltage measurements are then transformed into an array of 20,000 interpolation points to have measurements at the same point in time for each power trace.

The predicted power consumption of the S-box is then calculated for each key and each power trace, by taking either the Hamming weight or the Hamming distance of the input or output of the S-box for a particular power trace, or a combination of the in- and outputs. The input is obtained by performing a bitwise XOR operation on the known plaintext and the key guess, and the output is based on this input. The results, which are described in more detail in the following section, were optimal when taking the Hamming weight of the output signal.

The correlation between the predicted power signals for each power trace of a key guess and the power consumption at each interpolation point is then calculated. This correlation calculation is performed for incrementally larger subsets of power traces as well as for the complete dataset, to gauge the increase in accuracy of the side channel analysis as the size of the number of power traces increases.

4.2 SCA on the unprotected S-Box: results

Figure 4.1 shows a sample of ten power traces. Each power trace starts at a rising clock edge, which is also where the flipflops change their output. Spikes in power consumption are noticeable at the beginning and end, as well as the middle of each power trace, corresponding to the rising and falling edges of the clock. After an initial rise, the power consumption drops steadily after the falling edge. This corresponds to the charging of capacitances, which requires more power in the beginning within the circuit, as well as to the shorter subsections of the circuit finishing their operations faster and no longer contributing to the power consumption as significantly. The smaller power spike after the spike from the falling edge of the clock is caused by the switching of the inverters present in the flipflops and the charging of their internal capacitive nodes. A clock period larger than strictly necessary was chosen to include potential correlations that exist at or near the static power consumption of the circuit, as the assumption can be made that an attacker could alter the clock frequency themselves.

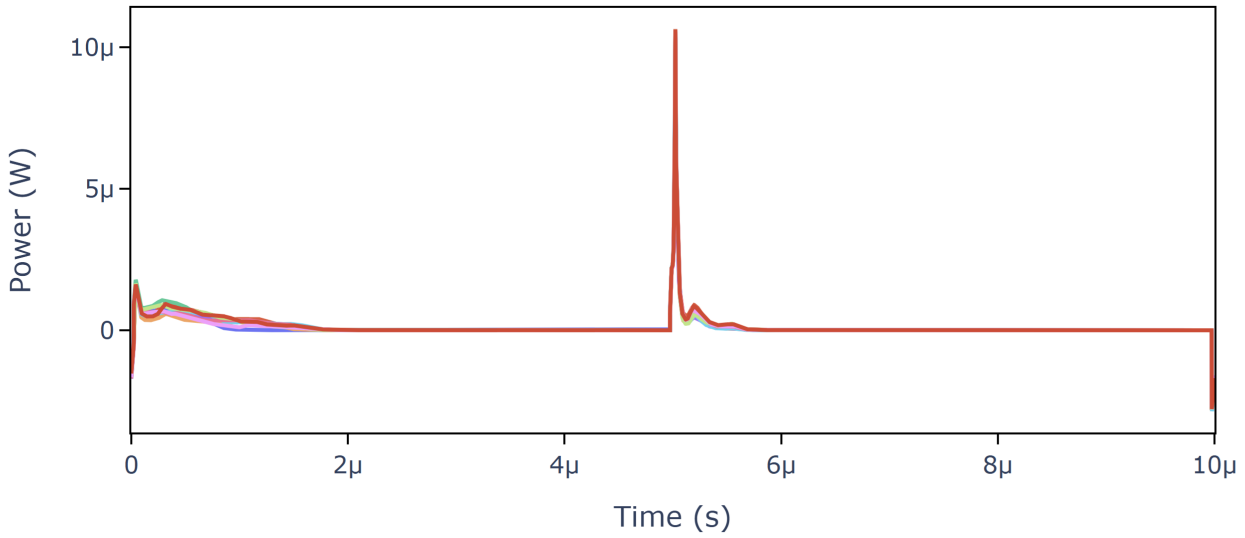

Figure 4.1: Sample of ten power traces

Figure 4.2 shows the correlation coefficient for the different key guesses over time. The x-axis denotes the interpolation point, while the y-axis shows the correlation coefficient, ranging from zero to one. The red line depicts the correlation for the real key.

The graph shows the correlations peaks during the rising and falling edges of the clock signal, with the correlation with the real key being significantly higher. This suggests that the power consumption of the flip-flops is data-dependent, allowing an attacker to reveal the secret key. Another, wider peak can be observed around the 5000th interpolation point, with again the real key having a higher correlation. This illustrates the CPA can differentiate the real key from the power consumed by the switching of the gates as well as by the power consumed by the flip-flops. Near the end of the power trace, before the correlation spikes during the rising edge, the correlation increases again. However, the correlation for the real key is not the highest, which shows the key cannot be differentiated from the passive power consumption of the circuit.

Figure 4.3 shows the result of the CPA. Each bar represents the highest correlation for a particular key guess, with the red bar representing the correlation of actual key. The bar for the real

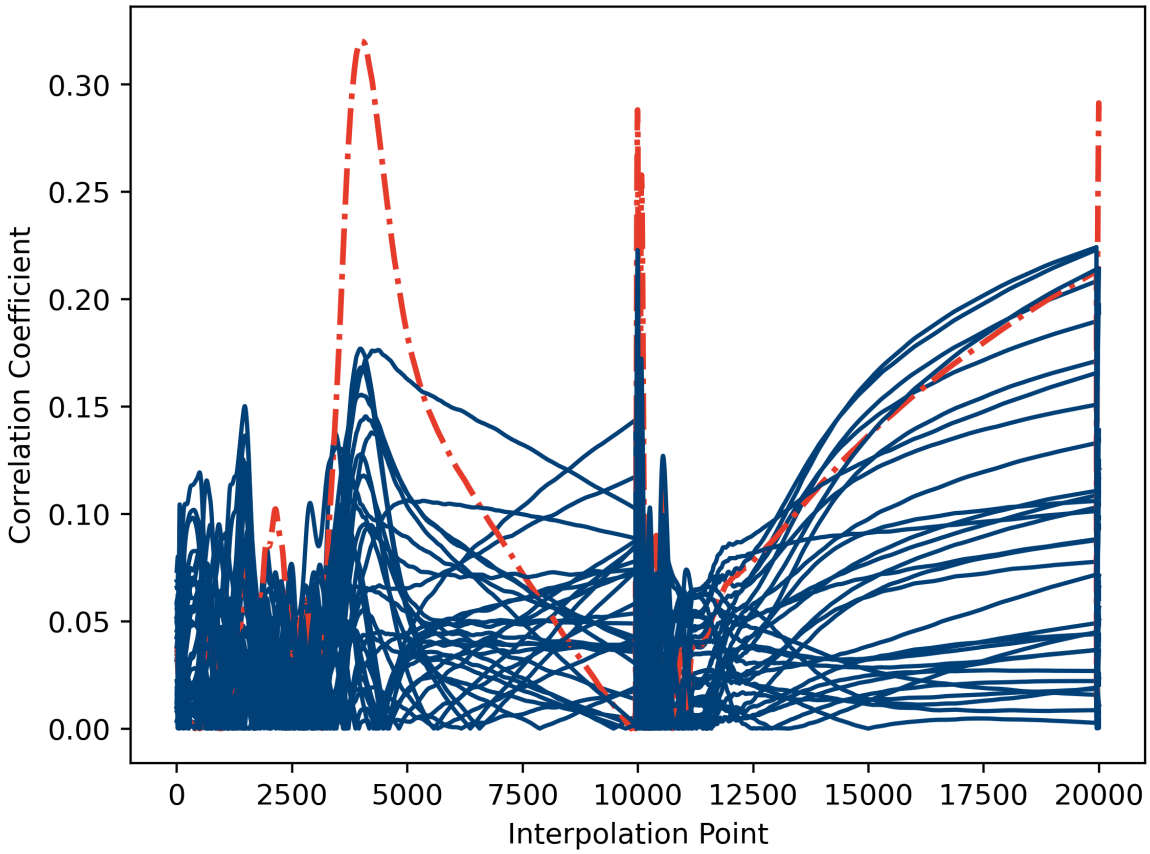

Figure 4.2: Correlation Coefficient for each key guess at different interpolation points

key is significantly higher than those of the other key guesses, indicating that the attacker has successfully obtained the secret key.

Figure 4.4 shows how the correlation coefficient of the different key guesses evolves as the number of power traces used in the analysis increases. All traces start out with a high correlation coefficient, which rapidly decreases as the size of the simulation increases. The correlation coefficient of the real key, shown in red, also decreases, but a lot slower than that of the other key guesses. After around a hundred power traces, the real key became clearly differentiated from the other key guesses.

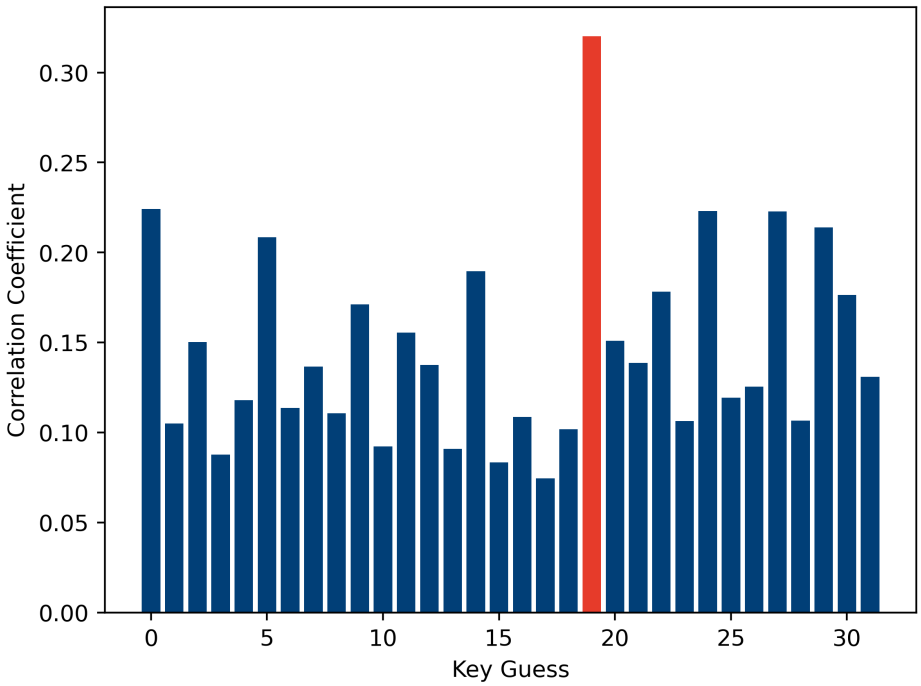

Figure 4.3: Correlation Coefficient for each key guess

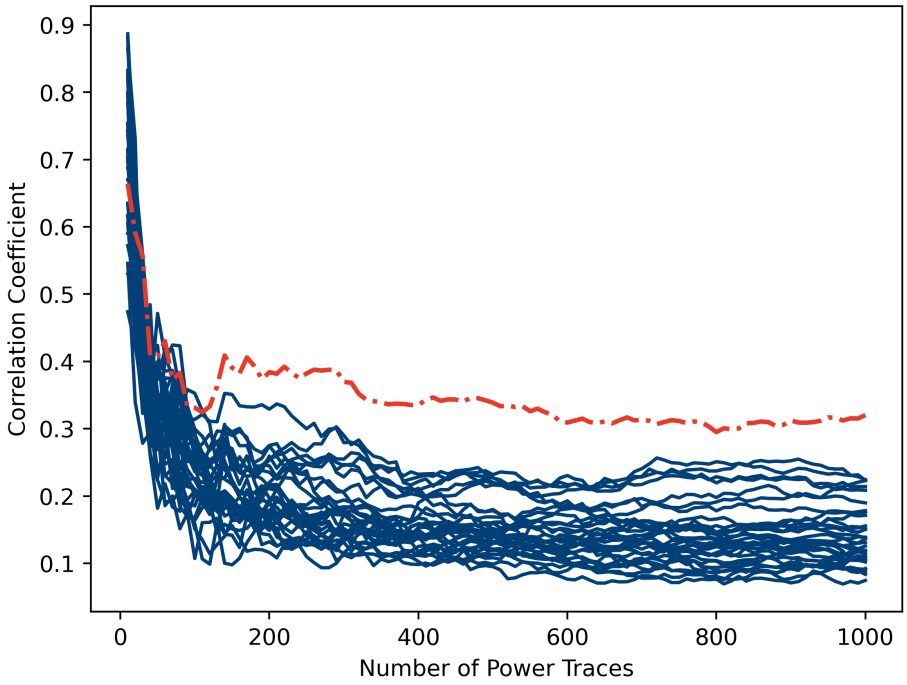

Figure 4.4: Correlation Coefficient for each key guess in function of the number of power traces

Chapter 5

SCA on WDDL S-box

The process of performing a CPA on the S-box that implemented WDDL was similar to the process on the unprotected S-box described in the previous section. The modified test circuit adds precharge operators to all the inputs of the S-box. A precharge signal is added which has half the frequency of the clock. The simulation was executed with a clock period of ten µs and a total runtime of 20,080 µs, twice as long as that of the unprotected design, as every other clock pulse is used for the propagation of the precharge wave.

The python code was modified to accommodate the precharge and evaluation phases. Power traces are created from each phase. Because performing the CPA on a combination of both types of traces was unsuccessful, they were split and the analysis was executed on either the precharge or evaluation power traces.

Figure 5.1 shows the correlation coefficient at different points in evaluation phase of the cycle. The correlation coefficient is highest in the middle of the graph, around the time the gates finish switching. The real key, depicted in red, can be clearly differentiated at this point. The rising and falling edges of the clock are no longer the points with the highest correlation coefficient and the real key cannot be clearly differentiated at these points, indicating that the power consumption of the registers is no longer data-dependent. The real key does however have a significantly higher correlation coefficient near the end of the graph, indicating that the static power consumption strongly correlates with the expected power consumption.

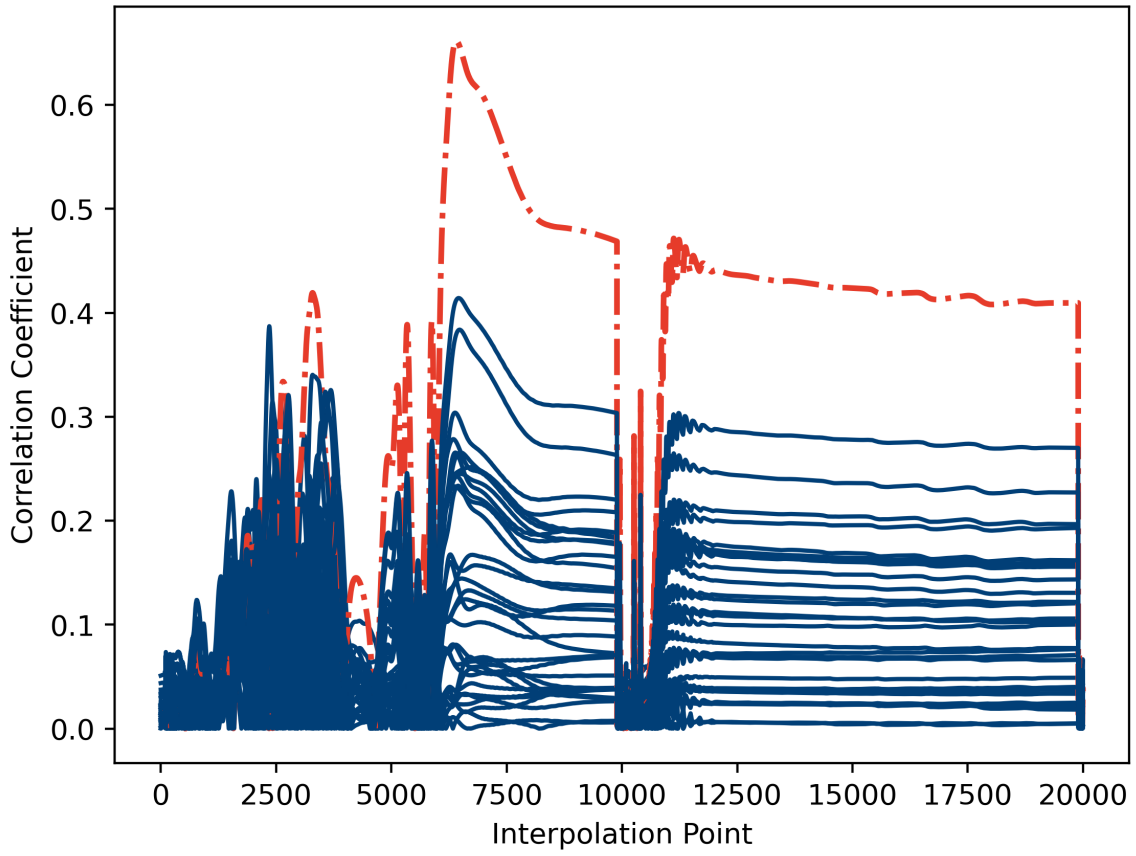

Figure 5.1: Correlation Coefficient for each key guess at different interpolation points for the WDDL circuit

Chapter 6

Discussion

This master's thesis aimed to study countermeasures against side channel analysis on cryptographic circuits on LTPO technology. As the results of the correlation power analysis on the unprotected Ascon S-box show, this attack is viable on LTPO. The five-bit secret key used in the S-box could consistently be retrieved with only a hundred power traces. The best results were obtained when using the Hamming weight of the calculated output of the S-box as the predicted power consumption. The real key could also be differentiated using the Hamming weight of the input of the S-box if the spikes in correlation at the clock edges were ignored. Hamming distance models were similarly successful but showed a smaller difference in correlation between the real key and the other key guesses. The real key could be differentiated both during parts of the clock cycle where the gates in the combinatorial logic were most active, as well as at the clock edges. This shows there is sufficient information leakage from both the dynamic power consumption of the flip-flops and the dynamic power consumption of the gates to differentiate the real key.

The WDDL logic style was ultimately unsuccessful in preventing CPA from revealing the secret key. The real key could be differentiated more easily than in the unprotected circuit. This suggests that the power consumption of uncorrelated parts of the circuit, which essentially function as noise, is more consistent. The WDDL registers were successful in hiding the data-dependent power consumption of the individual flip-flops. However, WDDL was unsuccessful in removing the information leakage caused by the data-dependent power consumption of combinatorial parts of the Ascon S-box. Additionally, the static power consumption by this part of the last few combinatorial parts of the S-box was sufficient to differentiate the real key.

The reduced data-dependent power consumption of the beginning of the S-box can be explained by the fact that the power consumption of the WDDL versions of the gates is less data-dependent than that of the unprotected gates. Especially the WDDL XOR and XNOR gates are data independent, since they consist of the same AOI22 and INV standard cells for both their true and false paths, simply with different inputs. Similarly, the WDDL registers don't leak information as they consist of the exact same master-slave flip-flop configuration in both the true and false paths. The other WDDL gates consist of a combination of different standard cells, such as an AND and an OR standard cell for a WDDL AND gate or an AOI22 and an OAI22 standard cell for the WDDL AOI22 gate.

Because different gates are used, there will always be some data-dependent power consumption. Both the capacitive and resistive properties of each gate would have to be perfectly balanced,

which is not possible. In a real chip, or a more detailed simulation based on a layout, most of the capacitance comes from the interconnects between the logic gates. This effect will become more dominant as the technology continues to improve, and channel-lengths continue to shrink [src-og].

Balancing the WDDL gates adequately to significantly reduce the data-dependent power consumption proved to be a significant challenge. A global optimization in Cadence Virtuoso was used to balance the worst-case rise and fall times of the AND and OR standard cells. Dummy n-type TFTs were added to the NAND standard cell to increase its input capacitance, making it closer to that of the NOR standard cell. A version of the S-box was then created using only WDDL-cells based on these balanced AND and OR standard cells, as well as the balanced WDDL gates mentioned above. However, the results, shown in figure 6.1, indicate that the correlation coefficient of the real key could be differentiated even more clearly. Figure 6.2 shows how the correlation coefficient of the real key approaches one near the end of the evaluate cycle. This is due to the increased imbalance in the static power consumption of the NAND and NOR standard cells paired with the reduced data dependence of the power consumption of gates in earlier parts of the circuit. Furthermore, Figure **??** shows that very few traces are needed to differentiate the key in this circuit.

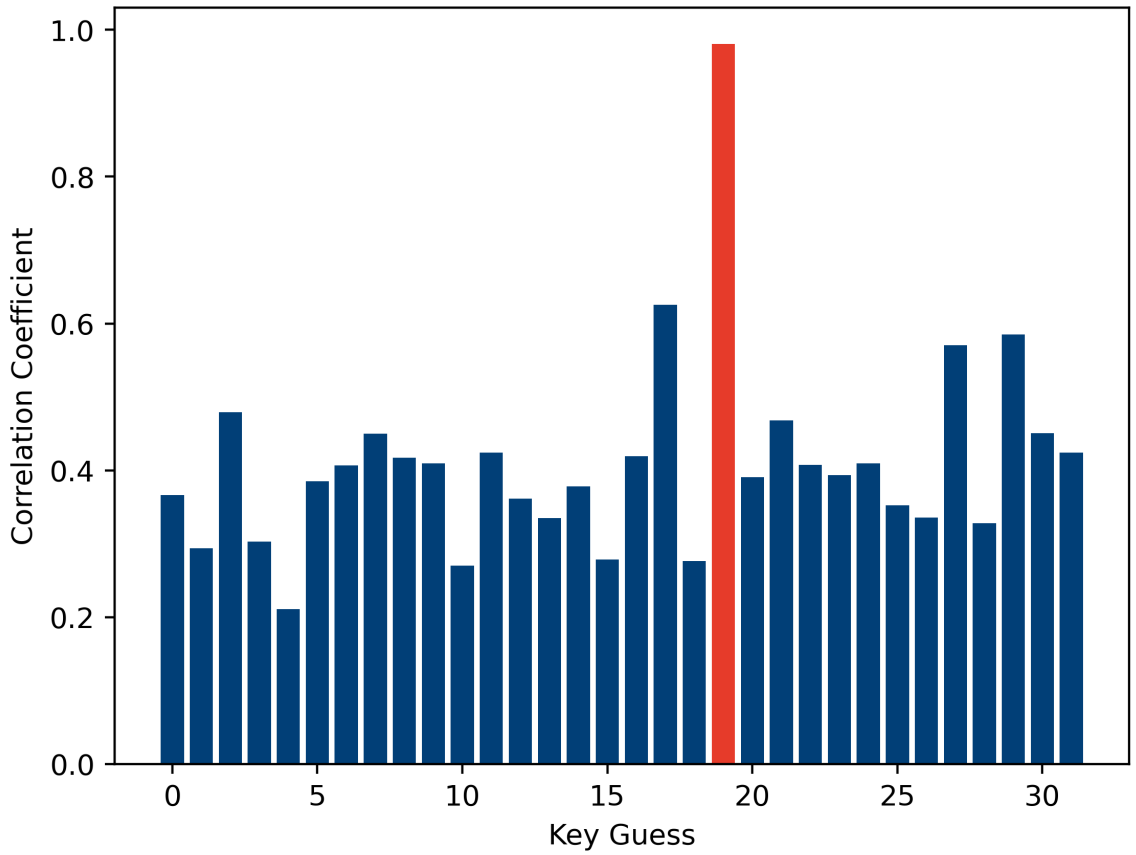

Figure 6.1: Correlation coefficient for each key guess, for the balanced WDDL circuit

The difficulty of balancing the different standard cells that make up the WDDL gates is due to the inherent imbalance between the n-type IGZO TFTs and the p-type LTPS TFTs. The

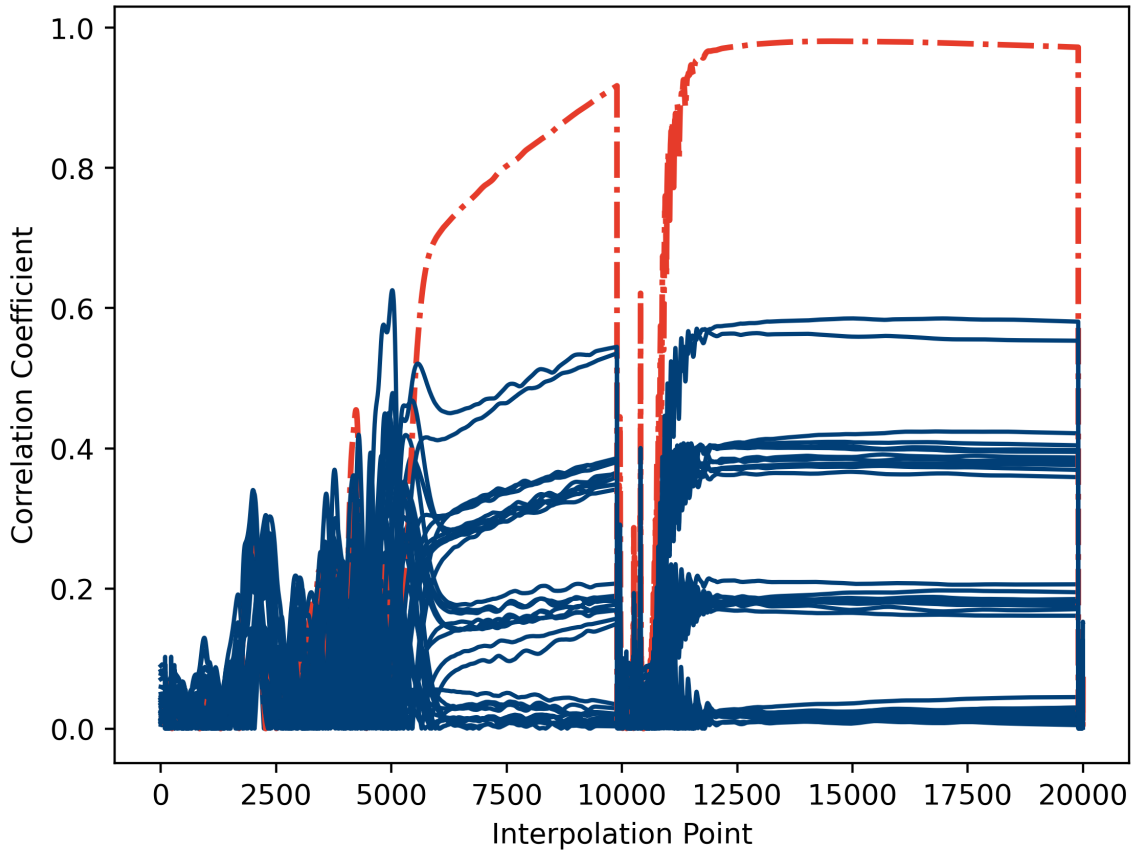

Figure 6.2: Correlation coefficient for each key guess at different interpolation points for the balanced WDDL circuit

LTPS TFTs are better at switching than the IGZO TFTs, but have a much higher static power consumption. The difference between the threshold voltage of both technologies, zero V and -2.95 V for the IGZO and LTPS respectively is also far too large to create gates that have both a switching threshold at half the supply voltage, with symmetrical noise margins to both sides and equal rise and fall times. Even if balanced gates could be created at a given supply voltage, an attacker could simply change the supply voltage of the device, which would bring the gates out of balance again, as the characteristics and necessary size adjustments are strongly dependent on the supply voltage, as was shown in the chapter where the optimal supply voltage was determined. [3] also notes how a substantial difference in charge carrier mobility between n- and p-type TFTs will have a negative impact on other performance characteristics, such as speed and area consumption.

6.1 Limitations

The lack of a layout poses a major limitation of this master's thesis. The capacitive effects caused by interconnects are an important factor in balancing the WDDL gates and preventing information leakage. Additionally, the device models of IGZO and LTPS TFTs utilized in this research have their limitations. [24] shows that their simulation using similar models corresponds closely to their experimental results. However, [**?**] shows an increased inverter gain as the supply

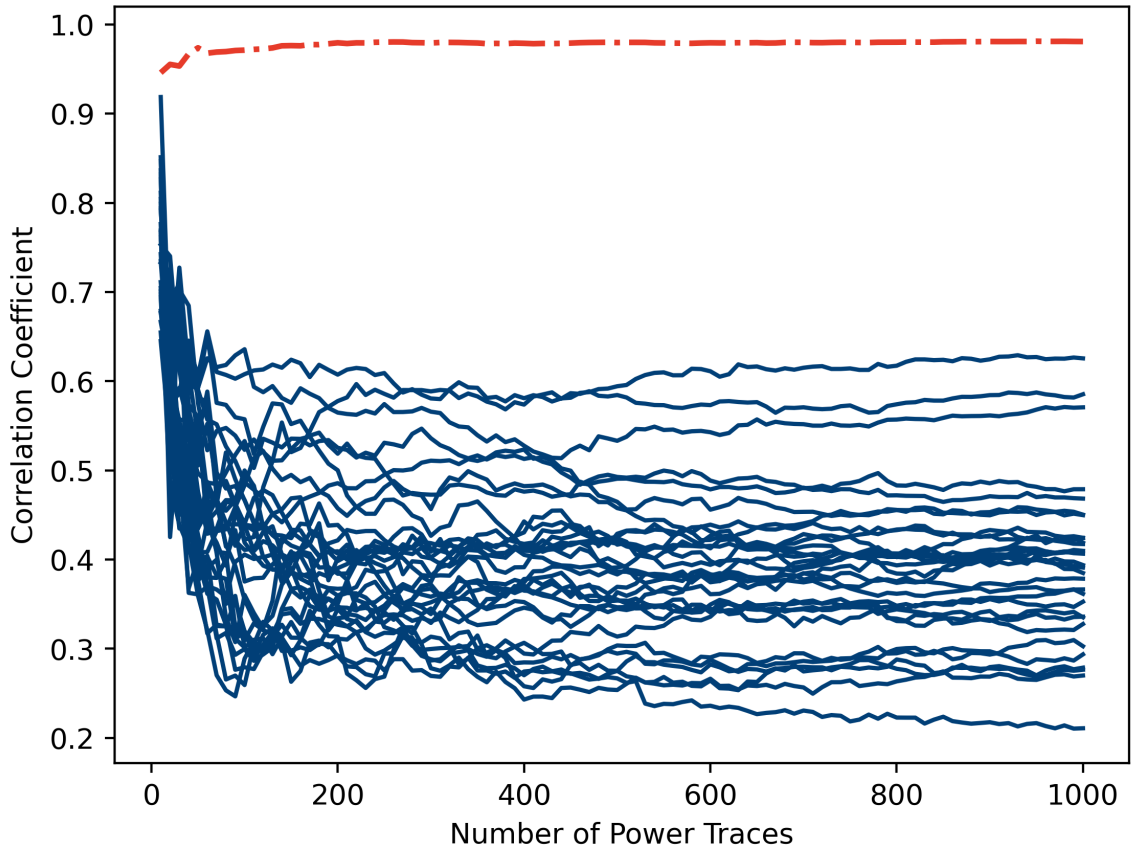

Figure 6.3: Correlation Coefficient for each key guess in function of the number of power traces

voltage increases from six to eight V, while the simulations done in this thesis observed a decreased gain when increasing the supply voltage between these levels. This illustrates that while they accurately model circuits these models are a valuable tool for simulating LTPO circuits, their accuracy may be limited when used at different supply voltages.

6.2 Future Work

The fast switching of LTPS TFTs paired with the low passive power consumption of IGZO TFTs has great potential to create processors that are both fast and power efficient, which would be ideal for some of the new device categories that flexible chips could enable, such as medical devices or Internet of Things applications. However, this will require a logic style that plays to the strengths of the technology. Future research should study which logic styles are most advantageous for LTPO. Using a dynamic logic style where the LTPS functions as a fast switch that rapidly precharges the outputs and the IGZO executes the function of the gate during the evaluate stage, such as in dynamic CMOS-like logic style, would be one way to benefit from the different properties of these TFTs. If such a logic style is also inherently differential, such as in Differential Cascode Voltage Switch (DCVS) logic or Sense Amplifier Based Logic (SABL), such a logic style could also streamline the implementation of WDDL.

Other measures against SCA on LTPO could also be developed. Many of these countermea-

sures, such as time and amplitude randomization or masking techniques, do not suffer from the challenges of combining IGZO and LTPS TFTs.

Chapter 7

Conclusion

This master's thesis examined the feasibility of SCA on LTPO. The implementation of reliable countermeasures against SCA is crucial in establishing LTPO as a viable technology for flexible integrated circuits. With the use of device models for an IGZO and an LTPS TFT, CMOS-style standard cells were created, along with master-slave, rising edge-triggered flip-flops. The optimal supply voltage was examined by obtaining the characteristics of inverters at different voltage levels. A literature review of different SCA countermeasures identified WDDL, a dual-rail logic style that propagates a precharge wave throughout the circuit, as a promising contender.

An Ascon S-box test circuit was implemented on LTPO using the created standard cells, on which a correlation power analysis was performed with the help of a Python script. The analysis proved most successful when using the Hamming weight of the predicted output of the S-box as the power consumption model.

WDDL gates were implemented by combining complementary standard cells and used in an adapted Ascon S-box test circuit. The WDDL proved unsuccessful in preventing the SCA, due to the imbalance between the standard cells used to create the WDDL gates. The large difference in threshold voltages, static and dynamic power consumption and switching speeds between the LTPS and IGZO TFTs makes it a significant challenge to balance CMOS-style LTPO standard cells. Balancing the rise and fall times of the standard cells, along with the input capacitance proved insufficient, as it increased the correlation coefficient during static power consumption. Furthermore, the large variability of the characteristics of these TFTs when the VDD changes makes it impossible to balance them at different supply voltages.

This thesis showed that the proposed solution of using WDDL gates consisting of CMOS-style standard cells as a countermeasure against SCA on LTPO is not a viable option for devices that require rigorous security.

Based on these findings, future research should focus on finding or developing a more advantageous logic style that benefits from the characteristics of LTPO. If standard cells can be created that have similar capacitive characteristics and switching speeds, so that the capacitance of the interconnects becomes the determining factor in balancing these gates, WDDL would be viable. Other SCA countermeasures such as time and amplitude randomization or masking techniques are also promising.

Bibliography

[1] M. S. Turan, K. A. McKay, D. Chang, J. Kang, and J. Kelsey, "Ascon-Based Lightweight Cryptography Standards for Constrained Devices," Tech. Rep. NIST SP 800-232, National Institute of Standards and Technology, Gaithersburg, MD, August 2025.

[2] H. Luo, S. Wang, J. Kang, Y. M. Wang, J. Zhao, T. Tsong, P. Lu, A. Gupta, W. Hu, H. Wu, S. Zhang, J. Kim, C. M. Chiu, B. G. Lee, Z. Yuan, and X. Yu, "Complementary ltpo technology, pixel circuits and integrated gate drivers for amoled displays supporting variable refresh rates," in *SID International Symposium Digest of technical papers*, vol. 51, pp. 351–354, 2020.

[3] K. Myny, "The development of flexible integrated circuits based on thin-film transistors," *Nature electronics*, vol. 1, no. 1, pp. 30–39, 2018.

[4] N. Zaidenberg and A. Resh, "Timing and Side Channel Attacks," in *Cyber Security: Analytics, Technology and Automation* (M. Lehto and P. Neittaanmäki, eds.), vol. 78 of *Intelligent Systems, Control and Automation: Science and Engineering*, Springer, Cham, 2015.

[5] M. M. Sravani and S. A. Durai, "Side-channel attacks on cryptographic devices and their countermeasures - a review," in *Smart Innovations in Communication and Computational Sciences*, (Indore), pp. 209–226, 2018.

[6] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Advances in Cryptology — CRYPTO' 99* (M. Wiener, ed.), (Berlin, Heidelberg), pp. 388–397, Springer Berlin Heidelberg, 1999.

[7] E. Brier, C. Clavier, and O. Francis, "Optimal Statistical Power Analysis," Research Report 2003/152, IACR Cryptology ePrint Archive, July 2003.

[8] B. Gierlichs, L. Batina, B. Preneel, and I. Verbauwhede, "Revisiting Higher-Order DPA Attacks: Multivariate Mutual Information Analysis," in *Cryptographic Hardware and Embedded Systems - CHES 2010*, pp. 233–248, Springer, 2010.

[9] S. Picek, A. Heuser, A. Jovic, S. A. Ludwig, S. Guilley, D. Jakobovic, and N. Mentens, "Side-channel analysis and machine learning: A practical perspective," in *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 4095–4102, 2017.

[10] N. Mentens, "Hiding side-channel leakage through hardware randomization: A comprehensive overview," in *2017 International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS)*, pp. 269–272, 2017.

[11] J.-S. Coron and L. Goubin, "On boolean and arithmetic masking against differential power analysis," in *Cryptographic Hardware and Embedded Systems-CHES 2000* (Ç. K. Koç and C. Paar, eds.), (Berlin, Heidelberg), pp. 231–237, Springer Berlin Heidelberg, 2000.

[12] P. Sasdrich, A. Moradi, T. Güneysu, and H. Handschuh, "Hiding higher-order side-channel leakage: Randomizing cryptographic implementations in reconfigurable hardware," in *Topics in Cryptology – CT-RSA 2017*, Lecture Notes in Computer Science, pp. 131–146, Cham: Springer International Publishing, 2017.

[13] I. Levi, D. Bellizia, D. Bol, and F.-X. Standaert, "Ask less, get more: Side-channel signal hiding, revisited," *IEEE transactions on circuits and systems. I, Regular papers*, vol. 67, no. 12, pp. 4904–4917, 2020.

[14] N. Veyrat-Charvillon, M. Medwed, S. Kerckhof, F.-X. Standaert, X. Wang, and K. Sako, "Shuffling against side-channel attacks: A comprehensive study with cautionary note," in *Advances in Cryptology – ASIACRYPT 2012*, vol. 7658 of *Lecture Notes in Computer Science*, pp. 740–757, Berlin, Heidelberg: Springer Berlin Heidelberg, 2012.

[15] Y. Lu, M. O'Neill, and J. McCanny, "Evaluation of random delay insertion against dpa on fpgas," *ACM transactions on reconfigurable technology and systems*, vol. 4, no. 1, pp. 1–20, 2010.

[16] E. Tena-Sánchez, F. E. Potestad-Ordóñez, V. Zúñiga-González, and A. J. Acosta, "Low-cost full correlated-power-noise generator to counteract side-channel attacks," *Applied Sciences*, vol. 15, no. 6, 2025.

[17] J.-L. Danger, S. Guilley, S. Bhasin, and M. Nassar, "Overview of dual rail with precharge logic styles to thwart implementation-level attacks on hardware cryptoprocessors," in *2009 3rd International Conference on Signals, Circuits and Systems (SCS)*, pp. 1–8, 2009.

[18] K. Tiri and I. Verbauwhede, "Charge recycling sense amplifier based logic: securing low power security ics against dpa [differential power analysis]," in *Proceedings of the 30th European Solid-State Circuits Conference*, pp. 179–182, 2004.

[19] K. Tiri and I. Verbauwhede, "A logic level design methodology for a secure dpa resistant asic or fpga implementation," in *Design, Automation and Test in Europe Conference and Exhibition*, Gielen, G (Editor), IEEE, 2004.

[20] M.-L. Akkar and C. Giraud, "An implementation of des and aes, secure against some attacks," in *Cryptographic Hardware and Embedded Systems — CHES 2001* (Ç. K. Koç, D. Naccache, and C. Paar, eds.), (Berlin, Heidelberg), pp. 309–318, Springer Berlin Heidelberg, 2001.

[21] D. Harris, R. Ho, G.-Y. Wei, and M. Horowitz, "The Fanout-of-4 Inverter Delay Metric," *Stanford University, EE371 Handout*, 1997. Appears to be based on papers presented at ISSCC 1997 and other contemporary conferences.

[22] J. M. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*. Prentice Hall, 2nd ed., 2003.

[23] J. Plusquellic, J. Howard, R. MacKinnon, K. Hoffman, E. E. Tsiropoulou, and C. Chan, "A Physical Layer Analysis of Entropy in Delay-Based PUFs Implemented on FPGAs," *arXiv preprint arXiv:2409.00825*, Sep 2024.

[24] S. Priyadarshi, A. Rahaman, M. Masum Billah, S. Nahar, M. R. M. Arnob, and J. Jang, "High speed level-down shifter using ltpo tfts for low power and interface electronics," *IEEE Journal of the Electron Devices Society*, vol. 12, pp. 587–593, 2024.

Appendix A

Inverter characteristics at varying supply voltages

Table A.1: Inverter characteristics (VDD 3V to 7V).

VDD (V)	**3V**	**4V**	**5V**	**6V**	**7V**
nmos (µm)	2	2	3	3.3	3.2
pmos (µm)	3.7	2	2	2	2
$t_{\text{rise 20 to 80\% (FO4)}}$ (s)	1.82849E-06	3.16739E-07	1.20464E-07	8.10064E-08	7.33443E-08
$t_{\text{fall 80 to 20\% (FO4)}}$ (s)	1.81001E-06	3.23715E-07	1.20634E-07	8.10437E-08	7.43981E-08
$t_{\text{rise 20 to 80\% (FO1)}}$ (s)	4.70502E-07	6.54197E-08	3.74485E-08	3.22532E-08	3.16721E-08
$t_{\text{fall 80 to 20\% (FO1)}}$ (s)	3.39747E-07	6.71539E-08	3.45171E-08	2.90772E-08	2.78216E-08
Inverter delay (s)	8.89946E-07	1.60528E-07	5.95332E-08	3.6908E-08	3.03281E-08
P_{avg} (W)	9.14342E-09	6.70005E-08	4.59792E-07	1.60445E-06	3.59651E-06
gain	6.917376457	8.939282428	10.71808676	12.03886239	12.49380561
Vinv (V)	1.243060078	1.761449335	2.277374983	2.828898096	3.403670143
NML (V)	0.732988321	1.191832868	1.638802827	2.110117853	2.572526272
NML (%)	24.43294402	29.7958217	32.77605653	35.16863088	36.75037531
NMH (V)	1.183131693	1.55053805	1.923560764	2.240851921	2.512165927
NMH (%)	39.43772309	38.76345124	38.47121528	37.34753201	35.88808467
overshoot (FO1) (V)	1.093849007	0.859177901	0.684086935	0.758891418	1.074088728
undershoot (FO1) (V)	-0.99486603	-0.80715771	-0.663956639	-0.592472211	-0.62955358
avNM (%)	44.15180557	49.17754732	52.01166418	53.84239688	54.69441765
overshoot (%)	36.46163356	21.47944753	13.6817387	12.64819029	15.34412469
undershoot (%)	-33.1622011	-20.1789427	-13.27913278	-9.87453685	-8.99362256
avGlitch (%)	34.81191731	20.82919514	13.48043574	11.26136357	12.16887363
prop_delay (s)	1.81925E-06	3.20227E-07	1.20549E-07	8.10251E-08	7.38712E-08
EDP(prop_delay)	3.02616E-20	6.8706E-21	6.68175E-21	1.05333E-20	1.9626E-20
EDP(inv_delay)	7.24163E-21	1.72656E-21	1.62959E-21	2.18558E-21	3.30806E-21

Table A.2: Inverter characteristics (VDD 8V to 14V).

VDD (V)	**8V**	**9V**	**10V**	**12V**	**14V**
nmos (µm)	3.4	3.7	4	3.9	3.5
pmos (µm)	2	2	2	2	2
$t_{\text{rise 20 to 80\% (FO4)}}$ (s)	7.52308E-08	7.75769E-08	7.87703E-08	7.27728E-08	6.3264E-08
$t_{\text{fall 80 to 20\% (FO4)}}$ (s)	7.58122E-08	7.83893E-08	7.80908E-08	7.25568E-08	6.23849E-08
$t_{\text{rise 20 to 80\% (FO1)}}$ (s)	3.08179E-08	2.85432E-08	2.65616E-08	2.20779E-08	1.83657E-08
$t_{\text{fall 80 to 20\% (FO1)}}$ (s)	2.73792E-08	2.66133E-08	2.46692E-08	2.08344E-08	1.76436E-08
Inverter delay (s)	2.75691E-08	2.48986E-08	2.19995E-08	1.73087E-08	1.42747E-08
P_{avg} (W)	6.54143E-06	1.03731E-05	1.53223E-05	2.81521E-05	4.52069E-05
gain	11.76358593	10.50180167	9.757436619	6.85972109	5.415255129
Vinv (V)	3.948511327	4.494401242	5.033148511	6.187887401	7.438484976
NML (V)	2.977416445	3.313147543	3.565558287	4.042373339	4.617085354
NML (%)	37.21770556	36.81275048	35.65558287	33.68644449	32.9791811
NMH (V)	2.77040834	2.959344593	3.099338512	3.161393069	3.006636165
NMH (%)	34.63010425	32.88160659	30.99338512	26.34494224	21.47597261
overshoot (FO1) (V)	1.397463601	1.561946184	1.573234423	1.448683415	1.20735868
undershoot (FO1) (V)	-0.825893136	-1.057920139	-1.20589212	-1.24390678	-1.18416376
avNM (%)	54.53275768	53.25355377	51.15227543	46.85891561	43.71716741
overshoot (%)	17.46829501	17.3549576	15.73234423	12.07236179	8.623990574
undershoot (%)	-10.3236642	-11.75466821	-12.0589212	-10.3658898	-8.458312573
avGlitch (%)	13.8959796	14.5548129	13.89563271	11.21912579	8.541151573
prop_delay (s)	7.55215E-08	7.79831E-08	7.84305E-08	7.26648E-08	6.28245E-08
EDP(prop_delay)	3.7309E-20	6.30828E-20	9.42528E-20	1.48648E-19	1.78428E-19
EDP(inv_delay)	4.97184E-21	6.43075E-21	7.41567E-21	8.43416E-21	9.21169E-21

Table A.3: Inverter characteristics (VDD 16V to 30V).

VDD (V)	**16V**	**18V**	**20V**	**25V**	**30V**
nmos (µm)	3.1	2.9	2.9	2.9	2.9
pmos (µm)	2	2	2	2	2
$t_{\text{rise 20 to 80\% (FO4)}}$ (s)	5.29611E-08	4.4655E-08	3.8289E-08	2.71095E-08	2.07826E-08
$t_{\text{fall 80 to 20\% (FO4)}}$ (s)	5.31844E-08	4.49894E-08	3.77547E-08	2.69546E-08	2.06707E-08
$t_{\text{rise 20 to 80\% (FO1)}}$ (s)	1.5738E-08	1.44154E-08	1.37597E-08	1.14405E-08	9.893E-09
$t_{\text{fall 80 to 20\% (FO1)}}$ (s)	1.53247E-08	1.33588E-08	1.16213E-08	9.10779E-09	8.27692E-09
Inverter delay (s)	1.21143E-08	1.04989E-08	9.27317E-09	6.99291E-09	5.3247E-09
P_{avg} (W)	6.68037E-05	9.59462E-05	0.000136673	0.000292586	0.000555655
gain	4.718616554	4.380924473	4.181404189	4.053643913	4.019748752
Vinv (V)	8.743907712	9.982401965	11.09383221	13.85468807	16.64364699
NML (V)	5.546405499	6.689836115	7.573041626	9.522724357	10.56039329
NML (%)	34.66503437	37.1657562	37.86520813	38.09089743	35.20131097
NMH (V)	2.530959352	1.922409511	1.570651707	1.028646464	1.533582793
NMH (%)	15.81849595	10.68005284	7.853258534	4.114585857	5.111942642
overshoot (FO1) (V)	0.934397044	0.887391615	0.862482466	0.835203922	0.817372299
undershoot (FO1) (V)	-1.114639126	-1.065319933	-1.03757681	-0.97752805	-0.929157554
avNM (%)	42.57428235	42.50578262	41.7918374	40.14819036	37.75728229
overshoot (%)	5.839981527	4.929953418	4.31241233	3.340815686	2.724574331
undershoot (%)	-6.966494537	-5.91844407	-5.18788407	-3.91011219	-3.097191848
avGlitch (%)	6.403238032	5.424198744	4.750148201	3.62546394	2.910883089
prop_delay (s)	5.30727E-08	4.48222E-08	3.80219E-08	2.70321E-08	2.07266E-08
EDP(prop_delay)	1.88167E-19	1.92759E-19	1.97583E-19	2.13802E-19	2.38706E-19
EDP(inv_delay)	9.80381E-21	1.05759E-20	1.17528E-20	1.43077E-20	1.57542E-20